UNSTATED ASSUMPTIONS IN ALARM MANAGEMENT

David A. Strobhar, PE
Beville Engineering, Inc.
7087 Corporate Way
Dayton, Ohio 45459
dstrobhar@beville.com

ABSTRACT

Alarm management in the process industries has been an increasing area of activity since the first publications of EEMUA 191 in 1999 and ISA18.2 in 2008. Most major oil and chemical companies have adopted some or all of the metrics for alarm management presented in these documents. However, simply targeting for these metrics might not be the best course of action for a company. Many of the metrics involve assumptions or logic not explicitly stated in the guidelines, assumptions that might not be valid for a given console or plant. In other cases, metrics are the result of implied, rather than explicit, relationships between alarm guidelines. Again, the implied connection might not be true under certain circumstances. Understanding these unstated underpinnings of alarm management can facilitate the implementation of a useful and efficient alarm management program. This paper explores some of the key unstated assumptions in alarm management and how changes in those assumptions can alter both the interpretation of alarm metrics as well as what metrics are utilized.

Key Words: alarm metrics, alarm management

1 BACKGROUND

An unintended consequence of the introduction of distributed control systems in the process industries is the proliferation of alarms. Since alarms can now be added at relatively low cost, alarms that would not have been justifiable for a hardwired panel have quickly become the majority. The result has been incidences of alarm floods, with 50 or more alarms actuating every minute during upsets. These alarm floods have been identified as a contributing factor in numerous plant incidents and prompted an industry response.

The first official guidance on alarm management came from the Engineering Equipment and Materials Users Association's (EEMUA) "Alarm Systems: A Guide to Design, Management and Procurement" published in 1999. This was followed almost a decade later with the International Society of Automation's SP18.2, "Alarm Management for the Process Industries". Both documents provide guidance on multiple aspects of alarm management, such as selection, presentation, and rationalization. Both also establish metrics or benchmarks for the alarm system, metrics with unstated assumptions.

Many companies are striving to achieve the metrics outlined in these documents. However, different assumptions can lead to different interpretations of the value used in the metric as well as the metric itself. This paper will examine the metrics and some of the assumptions upon which they are based.

2 ALARM METRICS

It is worth understanding the basis for the metrics themselves. The ISA18.2 metrics largely reflect those in the earlier EEMUA document. The EEMUA document clearly states that the metrics are neither based upon experimental research nor theoretical constructs. The entire guide reflects this in its statement -

> *"The guide has been written by users in industry. It is based on what some leading companies are doing, but it is also intended to be challenging and to push standards forward."*

As to the metrics themselves, the guide quickly points out that they are not intended to be definitive nor static, only

> *"indicative, but may give some pointer to whether designers are likely to have installed too many or too few alarms on a plant"*

In fact, the guide explicitly warns against adhering too closely to the metrics,

> *"It is emphasized that the numbers in Figure 24 & Figure 25 (target alarm rates and priority distribution) should be taken as approximate indicators of effective discrimination between priorities rather than exact targets."*

Despite the obvious qualifier to the metrics, they have largely been adopted by most oil and chemical companies. The process industries currently utilize several key alarm metrics derived from either EEMUA or ISA: (1) steady state alarm actuation rate, (2) upset alarm actuation rate, and (3) priority distribution (Table 1). What should be readily apparent is that these metrics have a prima fascia issue. If a plant were to have eight alarms due to an upset in ten minutes, and then none for the rest of the hour, does it meet the metric? Is the alarm system acceptable for upsets but not steady state (intuitively seems wrong)? Should the upset period be excluded from the steady state period, thereby meeting both metrics? How this inconsistency should be handled and why it even exists is not detailed.

Table I. Current alarm system metrics

Metric	Likely acceptable threshold
Steady state alarm rate	Less than six per hour
Upset alarm actuation rate	Less than 10 in 10 minutes
Priority Distribution	5% (highest priority), 15%, and 80% (lowest priority)

3 UNSTATED ASSUMPTIONS

There are four critical assumptions that underlie the metrics being used. The first is the volume of non-alarm workload. The second is the span of control for the user. The third is the impact of priority on alarm rate. The fourth is the impact of presentation.

Anyone familiar with process plant operation understands that operators do more than respond to alarms. So establishment of any limits for operator workload must be multidimensional, e.g., total workload = alarm work + non-alarm work. While it could be argued that the upset alarm rate is related to information processing limits (not stated by documents), the steady state rate clearly is geared to some aspect of workload. However, to establish an upper limit for steady state alarm rate would require that the amount of non-alarm work be specified. Perhaps my operators have more or less non-alarm work than those upon which the alarm actuation rate is based. In either case, a different limit or target should be utilized. Without knowing how much non-alarm work is assumed in the steady state alarm rate, the metric can at best provide some general target region rather than any specific limit.

Research sponsored by the Center for Operator Performance evaluated the impact of alarm rate on performance. Alarm rates of up to 10 in 10 minute's result in no performance decrement [1]. Only at 20 alarms in ten minutes did performance degrade for novice operators (engineering students trained on a simulator). When replicated with actual plant operators, the decrement was not seen until 30 alarms in ten minutes. Is the 10 alarms in 10 minute target based upon ensuring that novice operators are not overwhelmed? If that is the basis, then it should be stated as such.

Unstated in the metrics is the span of control of the operator, that is, the amount of equipment for which the operator is responsible. This is critical in evaluating the quality of the alarm system. Under the current guidelines, a single gas fired boiler with 50 control valves that had five alarms per hour would pass, while the hot side of an ethylene plant with 12 furnaces and 350+ control valves that had seven alarms per hour would fail. Clearly the ethylene plant has the better alarm system when normalized for the span of control. The metrics penalize plants that have large spans of control for no reason.

While the metrics include guidance on the priority distribution, they do not account for priority in the actuation rate. Under the current guidelines, six emergency priority alarms per hour are the same as six low priority alarms. Some companies argue that the hourly rate should be weighted by the target distribution, but this is not stated in the standard. Research by the Center for Operator Performance has shown that operators will utilize the priorities to handle higher priority alarms over lower priority alarms[2]. If the alarm rate exceeds what they can manage, then low priority alarms start to queue up. Given that low priority alarms have a longer time to respond, say 30 minutes, is counting their actuation in a ten minute period valid?

Finally, the metrics do not account for the presentation method. This is the most forgivable of the unstated assumptions as information on presentation impact has only recently come to light. In high alarm situations (over 10 per hour), presenting alarms by priority rather than time of actuation results in statistically significant better performance[i]. So, if metrics are to be created, then how the alarms are presented should be either a stated assumption or different metrics for different modes of presentation provided.

4 NEXT STEPS

The metrics used to monitor the performance of process plant alarm systems need to evolve. Any metric that is based upon certain assumptions needs to have those assumptions made explicit, so that a user can determine their applicability. In addition, the metrics need to reflect differences in span of control, priority distribution, and presentation methods. Finally, the entire field of operator performance in general, and the impact of alarms in particular, needs to be based upon solid research.

5 REFERENCES

1. Strobhar, D., Uhack, G., & Harvey, C., "Operator Performance as a function of alarm rate and interface design", World Batch Forum, Proceedings of WBF Make2Profit Conference, Austin, TX, 2010

2. Ibid

[i] Ibid.