Enhancing Process Safety Through a Holistic Approach to Safe Operating Limits

Mr. Jeffrey Miller, P.E.
Cognascents Consulting Group, Inc.
14701 St. Marys Lane, Suite 300
Houston, TX 77079
jeffrey.miller@cognascents.com

Mr. Dalton Carey
Cognascents Consulting Group, Inc.
14701 St. Marys Lane, Suite 300
Houston, TX 77079
dalton.carey@cognascents.com

Keywords: AIChE Global Congress on Process Safety (GCPS), Safe Operating Limits (SOLs), Hazard and Operability Study (HAZOP), Layers of Protection Analysis (LOPA), Occupational Safety and Health Administration (OSHA), Process Hazard Analysis (PHA), Process Safety Management (PSM).

Abstract

In today's competitive industrial landscape, organizations strive to maximize efficiency and throughput while maintaining an unwavering commitment to safety. This tension creates competing priorities for operations personnel, who must make time-sensitive decisions – continue troubleshooting or shut down to preserve safety. These situations highlight the essential role of engineering teams in equipping operators with practical, accurate Safe Operating Limits (SOLs).

Despite regulatory mandates and industry best practices, many companies struggle to establish clear and effective SOLs. Without robust limits, facilities are exposed to variability in human judgment. Even experienced operators may respond differently: some act conservatively and initiate costly shutdowns prematurely, while others push processes too far, increasing the risk of severe incidents.

This paper examines common pitfalls in SOL development, particularly the disconnect between Process Hazard Analyses (PHAs) and the SOL framework. In many facilities, PHAs and SOLs are developed independently – or PHAs serve as the sole input. Both approaches frequently produce limits that are misaligned with operational reality or insufficiently protective.

A real-world case study is presented in which a holistic SOL strategy was adopted – one that supplements and strengthens the PHA process rather than relying on it exclusively. The case includes corporate guidance spanning unit- and equipment-level specificity, promoting

consistency even where PHA quality varies. The outcome demonstrates a refined, actionable set of limits that reduce risk and improve field decision-making.

Readers will gain insight into:

- Challenges facilities face in defining effective SOLs
- Practical, repeatable solutions for addressing PHAs that fall short of expectations
- Common vulnerabilities that persist in PHAs and LOPAs
- Enhancements for SOL processes that already utilize strong PHA inputs
- Key benefits of integrating a holistic approach to SOL and PHA development

Whether a facility is beginning its SOL journey or optimizing an existing program, this paper offers actionable strategies and valuable lessons learned. The target audience includes those responsible for (1) leading within organizations required to comply with OSHA 1910.119, (2) establishing effective SOL and PHA guidance documents, and (3) implementing SOL initiatives that provide requisite process safety information to operations personnel.

Introduction

Over the past five years, Safe Operating Limits (SOLs) have received increased attention across the process industries. This increased focus reflects the maturation of Process Safety Management (PSM) since OSHA 1910.119's inception in 1992. Having strengthened core elements such as PHA, Management of Change (MoC), Mechanical Integrity (MI), Operating Procedures, and Pressure Relief / Flare System Design, industry focus has expanded toward areas that bridge these disciplines – SOLs being a prominent example[1].

This growing focus also reflects broader industry pressures. U.S. refining and chemical throughput has increased significantly without corresponding growth in new facility construction. Existing assets are expected to run harder, longer, and with greater reliability. Meanwhile, OSHA and insurance auditors have intensified scrutiny on damage-prevention elements of PSM, emphasizing the need for clearly defined, enforceable SOLs. This trend is highlighted in Figure 1 below.

Figure 1 – Refinery Count vs US Oil Capacity [1]

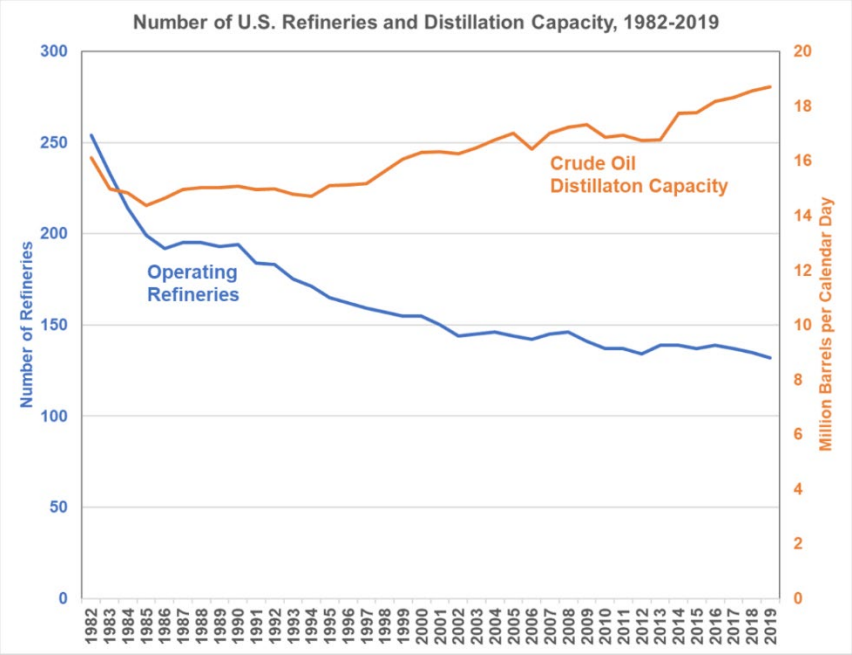

Effective SOL programs enable organizations to identify and manage the critical process parameters and protective instruments that prevent catastrophic equipment damage and high-severity consequences. When clearly defined and integrated into operations, they strengthen decision-making, reduce incident likelihood, and improve operational reliability and regulatory performance. This paper introduces a structured workflow for integrating PHA outputs, engineering validation, and corporate guidance into a repeatable Safe Operating Limit development process which is not consistently implemented in existing industry practice.

[1] For a list of abbreviations, refer to the clarifying table at the end of the paper.

The Role of SOLs in Operational Safety

Imagine if every roadway only had two posted speeds: 0 mph (parked safely) or the maximum speed your car is physically capable of reaching – say 160 mph. Technically, those are the limits of the machine. But as a driver, those numbers do not help in deciding how to safely navigate a school zone, a curvy mountain road, or a congested city street.

With only "park" and "redline" as guidance, some drivers would slam on the brakes at the first sign of traffic (costly overreaction), while others would push dangerously close to the maximum speed (risking catastrophic failure). The real value of speed limits comes from practical, context-specific numbers – 30 mph in a neighborhood, 65 on a highway, 20 in a school zone – that enable consistent, safe decisions in real-world conditions.

In the same way, setting Safe Operating Limits (SOLs) only at MAWP (Maximum Allowable Working Pressure) or MAWT (Maximum Allowable Working Temperature) provides little operational guidance. These values define absolute design boundaries but do not help operators determine when to act to prevent harmful excursions. Effective SOLs, like realistic speed limits, are contextual and actionable, designed to keep the process running safely with sufficient lead time on emerging events to prevent equipment damage that can lead to catastrophic outcomes.

Picture an operator on an isomerization unit observing reactor outlet temperature creeping upward during a feed upset. The posted "limit" is the MAWT of the reactor shell. Yet long before that value is reached, catalyst deactivation accelerates, hot spots form, side reactions increase, the exotherm intensifies beyond control, and metallurgical damage begins. Compounding this is the inherent lag between actual temperature rise and transmitter indication.

By the time the severity is fully recognized, no safe maneuver may remain – cooling cannot be applied quickly enough to prevent damage. The operator saw the trend developing, but the design limit offered no practical window to act. Effective SOLs must account for process dynamics and provide intervention time before escalation becomes irreversible and people are placed at risk.

When implemented correctly, SOLs serve as practical tools for operators, not administrative burdens. They define the transition between normal and emergency operations – the "line in the sand" where continued troubleshooting must stop and predefined shutdown actions must begin. In the absence of SOLs, operator response to abnormal events relies on experience and judgment alone, introducing variability and potential error. SOLs standardize this response by providing validated, pre-approved boundaries informed by engineering analysis, operating knowledge, design intent, and risk assessments. Whether identified by alarms or enforced by automated trips, each SOL breach represents a defined transition requiring immediate, consistent shutdown-style actions. Ask yourself: would you prefer operators determine where emergency operation begins during the emergency itself, or have a multidisciplinary team pre-establish those limits through a structured SOL effort?

In essence, SOLs operationalize process safety knowledge. They ensure protective limits are not abstract engineering values buried in datasheets but real, actionable thresholds guiding decision-making during emergency conditions. Figure 2 below contextualizes SOLs within the full operating window and illustrates how operations progress from normal conditions to emergency operations and, if unmitigated, to emergency shutdown and formal emergency response. This concept is supported by CCPS's book on effective operating procedures which states emergency operating procedures are required and intended to respond to system upsets where any operating parameter falls outside established safe operating limits, resulting, or likely to result in unstable operation, operation outside of design limits, or a potential release [2].

Figure 2 – Operating Window Relative to Operating Procedures

The figure above highlights how SOLs bridge process design and operational execution. Within this framework, operators are empowered to act decisively and consistently, knowing precisely where controlled operation ends and protective action must begin. By clearly delineating these limits, SOLs transform uncertainty into structured response, enabling both safety and operational discipline.

Identifying and Setting SOLs

Developing robust SOLs begins by identifying parameters tied to high-severity, fast-developing events – those that can reasonably transition a process from normal to emergency operation. Low-severity and/or slow-developing events (e.g., water pump seal damage or gradual corrosion/erosion) are better managed through alarm management, Integrity Operating Windows (IOWs)[2], and reliability programs. This distinction is underscored by the fact that neither of the above examples would elevate to an "emergency operation."

PHA files are a logical starting point for SOL identification – but not the endpoint. Each high-severity scenario identified in a PHA should be evaluated to determine the final process variable that precedes equipment damage and loss of primary containment (LOPC). In a distillation column experiencing loss of reflux, for example, the SOL should correspond to the pressure rise signaling imminent overpressure – not the initial low-flow condition, which may be recoverable.

This guidance is not absolute. SOLs must remain actionable. Waiting for high pressure in an overpressure scenario caused by liquid overfill may be inappropriate, as pressure rise could occur rapidly, if not instantaneously, after becoming hydrostatically filled. In that case, high level is the more appropriate SOL. PHAs serve as high-severity screening tools; engineering and operational judgment finalize the parameter to place the SOL.

PHAs provide an essential starting point for SOL identification – but knowing they are a risk assessment tool and not intended to define operating limits, it's important to know their limitations. Even high-quality PHAs can overlook critical limits. Anyone who has participated in an extended PHA has likely observed inconsistent results – where a scenario agreed upon early in the study is later revisited under similar conditions and assigned a different risk ranking. If a single team can contradict itself within the same study, variability across different teams, facilities, and years is even greater. The inherently qualitative nature of PHAs underscores the need for additional inputs when defining SOLs – parameters that may be lifesaving.

[2] Per API 584, IOWs are established limits for process variables (parameters) that can affect the integrity of the equipment if the process operation deviates from the established limits for a predetermined length of time (includes Critical, Standard and Informational IOWs). Critical IOWs are typically the only classification considered for inclusion within a SOL program as they are the ones determined by the mechanical / corrosion teams as requiring immediate predetermined actions by operations as exceedance of those limits can lead to rapid deterioration of the equipment [3]. Critical IOWs typically make up a small percentage of identified IOWs.

A recent poll conducted by Cognascents Consulting showed that industry experts believe only 25–50% of the trips and interlocks documented in their units' Cause and Effect (C&E) Matrices end up in the associated PHAs as credited safeguards or IPLs. Admittedly, some instrumented protections exist primarily for reliability or equipment protection and are therefore not safety critical. However, that same polling group estimated that 50–75% or even >75% of the trips and interlocks in C&E Matrices should have been credited as safeguards and IPLs. This 25–50% gap between execution and expectation is significant for SOL program effectiveness if identification relies only on instruments recognized by the PHA team as protection layers.

One effective strategy to address potential PHA oversight is corporate-level guidance identifying high-severity scenarios or protections for inclusion in the SOL program at the unit and equipment level. When applying such guidance, balance between standardization and site ownership is essential. Making potential SOLs recommended rather than mandatory encourages adoption without rigid uniformity. Requiring each site to evaluate relevance – and justify exclusions with corporate SME approval – ensures accountability and flexibility. In this two-pronged approach, PHA-derived inputs drive accuracy, while corporate guidance drives consistency. When SOL programs achieve both, they better align corporate standards with operational reality, improving process safety outcomes.

Once parameters are identified, setting the SOL requires balancing response time and design margins. Reading the following guidance relative to Figure 2 may be beneficial. Data trending and variability analysis help define realistic Target Operating Limits (TOLs), equivalent to PVHI or PVLO alarms, placed outside normal operating ranges to avoid nuisance alarms yet far enough from associated PVHH/PVLL values to allow troubleshooting time before SOL breaches and shutdown requirements. SOLs, equivalent to PVHH/PVLL alarms or trips, should be established beyond TOLs but within the equipment's design envelope (SDLs – safe design limits), ensuring adequate time for operator or instrumented response before damage occurs. Where the window between SOLs and SDLs is too narrow for adequate operator response, automatic trips or interlocks may be warranted to avoid equipment damage and LOPC.

The rigor applied to SOL validation should be proportional to risk. High-severity scenarios relying on instrumented protection to satisfy LOPA criteria warrant IPL validation considerations. The preferred validation method is process safety time (PST) analysis to confirm protective layers respond within the necessary timeframe[3]. However, there are cases where manufacturer recommendations, engineering judgment, or operating experience are more valuable. For example, a PST calculation to determine how long a pump seal may last in a cavitated or deadheaded state is unlikely to produce better insight than deferring to a rotating equipment specialist familiar with

[3] Per the Center for Chemical Process Safety (CCPS), *Guidelines for Initiating Events and Independent Protection Layers in Layers of Protection Analysis*, Process Safety Time is defined as the "time period between a failure occurring in the process or its control system and the occurrence of the hazardous event," and further notes that "an evaluation of an IPL is important to confirm that the IPL can successfully complete its action and that the process can return to a safe operating condition within the PST." [4]

the pump, seal design, process conditions, and fluid composition. In either case, if credited to save a life in LOPA, an instrument deemed suitable for SOL purposes should undergo validation to confirm effectiveness when called upon.

Not every parameter or SOL requires that level of IPL validation. Conducting PST calculations for scenarios that do not drive LOPA risk reduction may place unnecessary time burden on engineers and reveal insufficient protection capability on instruments that were not even being called upon in the PHA to begin with. Instead of modifying or upgrading an instrument that was not needed as an IPL, certain situations simply need to rely upon non-instrumented IPLs to provide necessary risk reduction. The SOL can still be beneficial; it is simply not required for the worst case, fastest developing, consequence analysis identified in the LOPA. For situations where IPL demand is not being placed on the SOL-based instrument, but an SOL is still determined to be prudent, engineering judgment may appropriately replace detailed calculations, maintaining scalability without sacrificing safety. A common example involves high-pressure trips on compressor discharge. If a maintenance-based manual valve were closed on an online compressor, only a small volume of outlet piping would be available to absorb the pressure rise. It is not reasonable to expect a high-pressure trip to shut down a compressor rotating at thousands of rpm before the design pressure is reached. The fact that such a calculation would fail PST criteria does not mean the high-pressure trip should be removed because it is ineffective in the single worst case. Instead, the LOPA would acknowledge the extremely low likelihood of that initiating event and rely on the discharge PSV for the remaining risk reduction. The SOL team can then focus on more credible initiating events (e.g., downstream pressure control failure), in which the high-pressure trip would satisfy PST criteria and remain acceptable.

The “SUITable” Principle

To maintain practical and effective limits, each SOL should be **SUITable** for its intended use, meaning that it meets the following criteria:

- **Severity:** The SOL protects against scenarios with high potential severities.
- **Utility:** The limit must provide actionable information and clear corrective steps that seek to prevent the associated high severity.
- **Independent:** Where possible, the SOL’s instrument should not be defeated by the failure which caused the deviation.
- **Timely:** Response – automatic or operator-initiated – must occur quickly enough to prevent exceeding design limits.

Ultimately, operators must understand what the limit represents, why it matters, and what action is required when it is reached. SOL programs that align technical accuracy with operational practicality are far more likely to sustain compliance, maintain engagement, and meaningfully improve safety performance.

Example SOL Workflow

A proven workflow combines PHA outputs, engineering validation, and corporate guidance:

Figure 3 – Example SOL Workflow

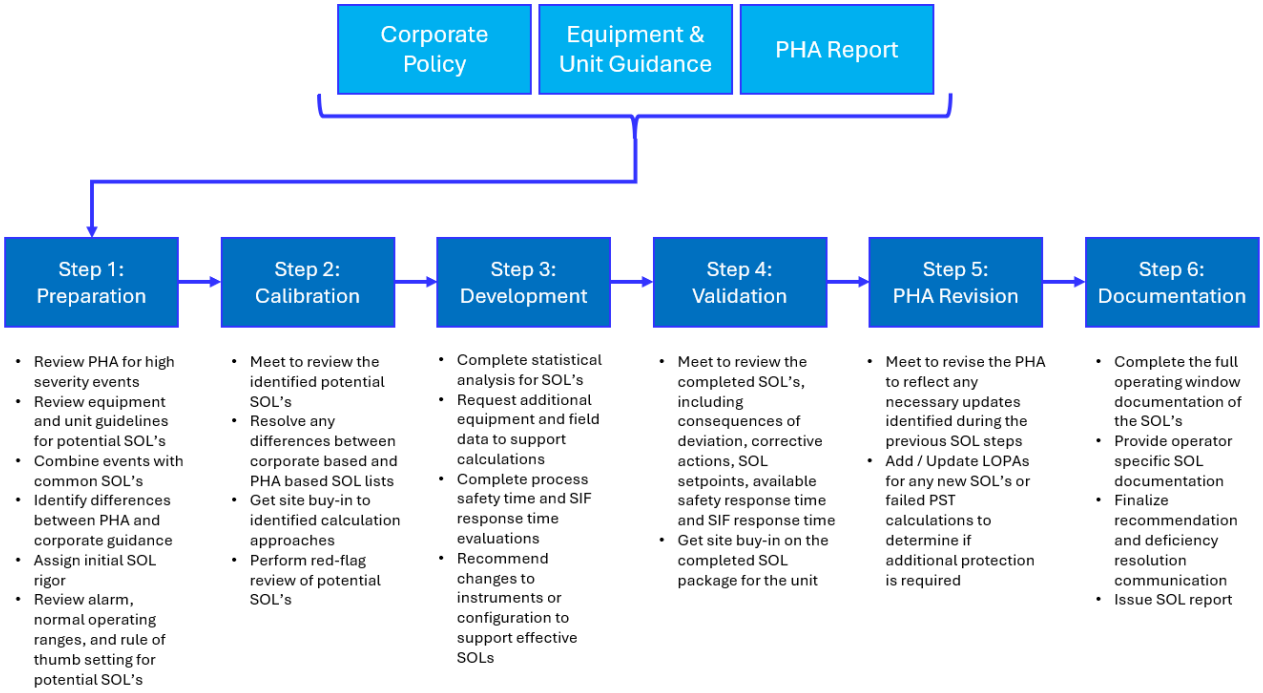

When applied effectively, this workflow balances the site-specific insights captured in PHAs with the standardization and technical depth provided by corporate guidance. It focuses attention on high-severity scenarios identified in the PHA while enabling corporate SMEs to recognize recurring hazards at the unit and equipment level. As resources become increasingly limited, efficiency in team time is critical. Documenting SME expertise within corporate guidance allows facilities to make technically sound decisions even when those experts are unavailable in real time.

Corporate guidance should remain advisory, not prescriptive. Sites must retain flexibility to deviate from standardized recommendations when they do not align with actual operating configurations. Requiring justification and SME approval for deviations preserves both flexibility and technical integrity.

This framework has proven effective across sites with varying PHA quality. Where PHAs are robust, they accelerate SOL development by pre-identifying high-severity scenarios and protections. Where PHAs are weaker, corporate guidance fills the gaps, providing a consistent baseline. Even at sites with strong PHAs, this workflow adds value by recognizing that PHAs identify risk–they do not define operating limits. Translating risk results into actionable SOLs requires an additional layer of analysis and standardization. Figures 4 and 5 illustrate results from multiple units applying this workflow.

Figure 4 – Breakdown Based on SOL Identification Source

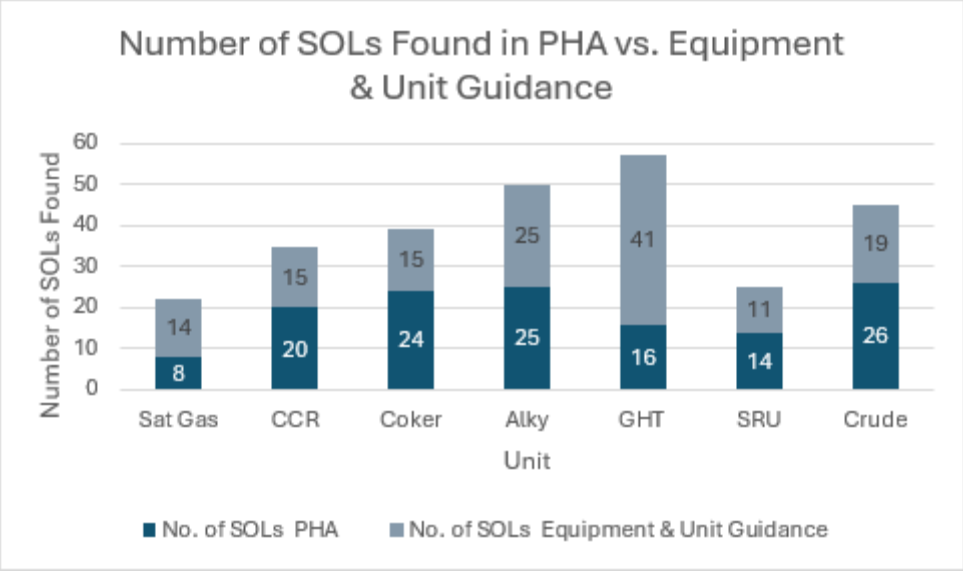

Figure 5 – Instrument Passing Rate for PST Considerations

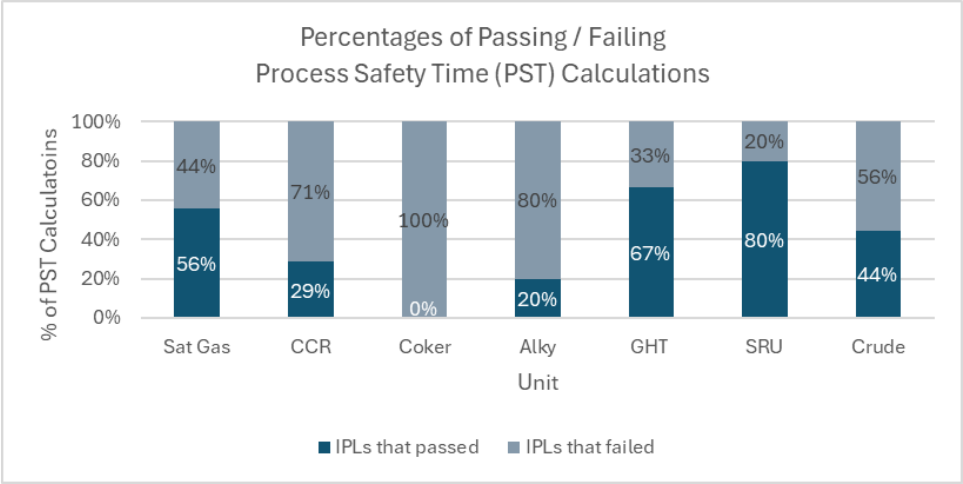

The data used to generate Figures 4 & 5 come from three different refineries across two years of SOL implementation[4]. Across these facilities and operating companies, PHAs alone often missed high-severity scenarios or credited protections incapable of meeting process safety time (PST) requirements. Looking at Figure 4, a company that only generates their SOLs by referencing PHA results could have missed 140 of the 270 SOLs identified by these studies (52% of the overall SOLs). Looking at Figure 5, a company could have been relying on instrumented protections that were not capable of responding in time to prevent the hazards in question (instances where the required response time was more than the scenario's PST).

[4] Although derived from refinery operations, the workflow is applicable to any process industry where acute process excursions can challenge operator and control system response time.

By following the workflow identified above, the SOL implementation efforts corrected these deficiencies by aligning protective capability with actual process dynamics – reinforcing that facilities cannot protect against scenarios they have not identified, nor rely on safeguards that cannot respond in time.

Common Pitfalls in SOL Development

Despite their criticality, many organizations struggle to define effective SOLs. The below examples highlight common pitfalls that reinforce the need for holistic programs as described above.

1. Overreliance on PHAs

Overly conservative outcomes:

Many facilities derive SOLs directly from HAZOP safeguards or LOPA independent protection layers (IPLs), assuming equivalency between PHA protection layers and enforceable operating limits. However, HAZOPs are qualitative, LOPAs semi-quantitative, and both are designed to assess risk – not define precise operational thresholds. PHA teams typically identify safeguards that maximize response time; SOLs must define the final, non-negotiable boundary where safe operation ceases. Without this distinction, limits may be set too early in an upset, causing nuisance shutdowns and unnecessary disruption.

Expanding on the example provided previously, an SOL team may be reviewing a distillation column trending toward overpressure following loss of reflux. A PHA may credit a low reflux flow alarm as an IPL because the alarm activates immediately upon loss of reflux, providing the operator with the maximum available response time to prevent the pressure escalation. From a risk perspective, this may be appropriate. But if directly converted into an SOL, it becomes overly conservative. If operators must open flare valves and cut feed and reboiler heat every time low reflux flow alarms – even during brief, recoverable upsets – operations suffer unnecessary interruptions, off-spec product, and downstream instability.

Allowing operators to troubleshoot loss of reflux while monitoring column pressure – and only initiating SOL-driven action as pressure approaches a defined not-to-exceed limit – preserves both safety and continuity. Early-warning safeguards should not automatically become enforceable SOLs without considering their role within the broader protection hierarchy.

Non-direct outcomes:

PHA protection credits can also misalign with SOL strategy when teams rely on indirect indicators for safeguard or IPL credit.

Consider a blocked bottoms pump outlet on a distillation column. If the hazard is pump deadhead leading to seal failure, LOPC, ignition, and injury, a PHA team might credit a high column level alarm. While logical – blocked outlet leads to rising level and operator response – this approach

creates ambiguity during emergency operations. A high level alarm could just as easily result from a pump trip. In that case, restarting the pump may be appropriate; in a true deadhead event, shutting down the pump is required. Conflicting responses tied to the same alarm increase the risk of error during critical moments.

The solution here is to tie the SOL to the parameter that more directly reflects the potential hazard. If the pump run status is on and a low flow alarm is annunciated, it is more directly correlated with the hazardous deadhead event. The added benefit is that the low flow SOL also protects against low-level cavitation scenarios, potentially replacing two level-based SOLs with a single, hazard-focused SOL. Direct monitoring reduces confusion and improves clarity in emergency response.

Incomplete outcomes:

Even high-quality PHAs can produce incomplete SOL documentation because their objectives differ. PHAs determine risk acceptability; SOLs define enforceable operating boundaries.

For example, a PHA may conclude that overpressure risk is tolerable due to a low initiating event frequency, a properly sized PSV, and ignition probability or occupancy factor – requiring no instrumented IPL. Even if acceptable from a risk-ranking standpoint, pressure still warrants monitoring if overpressure can lead to a high severity event. Critical equipment operating near or above MAWP – even if protected by an adequately sized relief valve – warrants a high-pressure SOL. Without it, a scenario deemed "acceptable" in the PHA could result in extended operation near design limits without operator awareness.

Another example involves a vessel equipped with a single level transmitter. In a PHA, the transmitter may be rejected as an IPL for a gas blowby scenario because it does not meet the independence requirements if the transmitter fails high creating the low level as the initiating event. However, the purpose of a SOL is not necessarily to serve as an IPL, but to define an operating boundary that helps protect equipment and prevent escalation of abnormal conditions. Many credible failures such as an incorrect setpoint, bypass valves left open, loss of feed with continued outflow, or procedural errors could still result in a low-level condition while the transmitter itself remains functional. Establishing a low-level SOL therefore remains beneficial for managing real-world operating risks and protecting equipment integrity, even when the instrumentation does not qualify as a fully independent protection layer in the PHA.

PHAs apply conservative independence rules and may not evaluate all alternative failure modes without specialized studies (e.g., CHAZOP or RCA). This creates a gap between risk-based reasoning and the operator's need for actionable guidance. Effective SOL development must bridge that gap – ensuring PHAs establish risk acceptability while SOLs define measurable, enforceable boundaries that preserve response opportunity for critical parameters.

2. PHA Quality Variability

PHA quality varies widely across sites. Poorly facilitated studies may overlook high-severity scenarios or credit safeguards incapable of responding within required PST. Relying exclusively on such analyses yields SOLs that are incomplete, inconsistent, or technically weak.

Underdeveloped consequences create real vulnerabilities. Conversely, applying risk reduction credit to instruments that cannot respond in time creates a false sense of security. As stated in CCPS's book on Guidelines for Safe and Reliable Instrumented Protective Systems, "The successful operation of manually-initiated instrumented protective functions is heavily dependent on having sufficient process safety time. Manual systems require the operator to detect the process condition, recognize the need to take action, determine the correct action, and complete the specified action." The same is true for automated instrumented protective functions; however, the response time required is often greatly reduced as compared to the human initiated responses. [5] The workflow presented earlier directly addresses both issues, as demonstrated in Figures 4 and 5.

3. Excessive SOL Identification

Given that OSHA 1910.119 references operating limits only briefly, it is unsurprising that implementation strategies vary widely:

1910.119(d)(2)(i)(D) – Safe upper and lower limits

1910.119(d)(2)(i)(E) – Evaluation of consequences of deviations

1910.119(f)(1)(ii) – Operating limits: consequences and corrective steps

In the absence of structure, some sites default to creating limits for nearly every transmitter to avoid noncompliance. The result can overwhelm operators. Imagine searching for guidance on a high-high reactor temperature while navigating pages of supplemental instrumentation documentation.

Excessive limit generation not only buries critical information but often results in poorly refined setpoints. When hundreds of limits compete for attention instead of focusing on the handful that truly matter, design-limit-based setpoints become common. While appearing compliant, this approach drives frequent trips, operator frustration, alarm fatigue, and erosion of SOL credibility.

Ultimately, SOL programs that rely exclusively on PHA outputs – or default to over-implementation and conservatism – fail to deliver full value. Effective programs integrate engineering validation, operational practicality, and corporate consistency. The workflow presented in this paper addresses these challenges and improves the likelihood of successful SOL implementation.

Implementation Challenges with Potential Solutions

While the previous section addressed technical and structural weaknesses in SOL development, even well-designed programs encounter practical challenges during implementation. Acknowledging these obstacles allows sites to capture program benefits while managing more complex SOLs transparently and pragmatically.

1. Cultural Resistance

Cultural resistance to shutdowns is persistent. Operators may hesitate to take protective action – particularly shutdowns – due to production pressure or perceived blame. This is amplified for operator-based SOLs.

Proposed Solution:

Facilities must reinforce that operator-initiated responses and instrumented trips serve the same safety function. An interlock is not reprimanded for acting at its setpoint – and neither should an operator. Consistent messaging from leadership and supervisors is essential.

Early operator engagement builds trust. A "Red Flag Review" at the end of the Calibration Meeting (Figure 3) overlays historical data with alarm and proposed SOL setpoints. This visual review helps identify impractical or overly sensitive limits before implementation.[5]

For example, one site identified a pump seal failure as high severity and proposed a low-flow SOL – but the pump lacked a flow transmitter and operated only intermittently. Installing a transmitter could have created nuisance alarms when the pump was intentionally offline. Instead, a pump run timer was implemented. If the vessel should drain in five minutes but the pump runs for fifteen, this deviation flags potential starvation, blockage, or valve misalignment – without creating unnecessary alarms.

Operator involvement ensures both technical soundness and ownership, improving reliability during real emergencies.

[5]The Red Flag Review is an informal validation step used to identify limits that are impractical, overly sensitive, or inconsistent with historical operating behavior. During the Red Flag Review, all SOLs and their associated instrumentation that were agreed upon during the Calibration Meeting have their historical data trends reviewed to evaluate whether the proposed SOL setpoints will be operationally sustainable and protective of equipment integrity. If a high exceedance frequency is observed in the historical data, the team evaluates whether historical operation has occurred with insufficient margin to failure, whether the proposed SOL appropriately represents an emergency condition, whether past exceedances reflect the absence of a defined target rather than an impractical limit, or whether alternative protection strategies are more appropriate for the scenario.

2. Tight Operating Windows

Processes operating near design or relief limits present unique challenges. Narrow margins between TOLs, SOLs, and design limits can lead to frequent trips or unsafe workarounds.

Proposed Solution:

Interim measures can provide risk reduction while long-term engineering solutions are developed. Many older facilities operate near or above 90% of design pressure due to debottlenecking, while good practice maintains operating pressure below 90% of PSV set pressure to prevent valve simmering and premature lifting.

If pressures fluctuate between 90–95% of PSV set pressure, an immediate trip at 90% may be impractical. Interim steps – enhanced monitoring, increased inspection, targeted training – can provide partial protection while longer-term solutions are engineered, such as:

- Increasing the design pressure rating of critical equipment
- Installing additional or larger relief devices
- Replacing conventional relief valves with pilot-operated relief valves (which permit tighter operating windows)
- Reducing operating pressure or throughput

Though interim SOLs may be less robust than final solutions, they provide measurable risk reduction. An interim SOL partially meeting desired corporate outcomes within a structured program is safer than deferring all SOL implementation until the long term solution is fully developed.

3. SOLs with Imperfect Monitoring

In some cases, the process variable selected for an SOL is appropriate in principle, but the existing instrumentation lacks the precision or span needed for effective action. As a result, the SOL may trigger before a true hazardous condition develops, leading to nuisance shutdowns or alarm fatigue.

Proposed Solution:

Two practical strategies can improve performance in such cases: (1) implementing time delays to filter out transient or non-emergency conditions, and (2) using state-based alarm settings that adjust limits according to the operating mode (startup, shutdown, regeneration, etc.).

For instance, a distillation column high-level SOL intended to prevent hydrostatic overfill/overpressure may be associated with a transmitter whose full span only includes the column sump. A limit set at 90% of transmitter span might not represent an actual overfill risk at that set point. By calculating the time required to reach an unsafe level at design feed rates, a time delay can be applied to prevent spurious trips while preserving the safety function. Ultimately, regardless of technical complexity, every SOL must clearly define the point at which normal operation ends and emergency response begins.

Conclusions

Effective SOL programs deliver value far beyond regulatory compliance. They validate IPL assumptions, strengthen PHA credibility, and reinforce value-based risk management. A consistent SOL framework also functions as an auditing lens – exposing gaps between design intent, operating practice, and documented safeguards. Over time, this discipline improves hazard recognition, sharpens alarm management, and elevates focus on truly safety-critical elements.

Safe Operating Limits are where complex engineering analysis meets frontline decision-making. They translate HAZOPs, LOPAs, and design envelopes into clear, enforceable boundaries for the people closest to the process – the operators. When properly developed and integrated, SOLs unify risk analysis, operations, and protection systems into a single, actionable framework.

A holistic approach – grounded in engineering rigor, operational realism, and corporate consistency – produces SOLs that are both protective and practical. Organizations that embrace this balance do more than satisfy auditors; they give operators clarity under pressure, strengthen system resilience, and prevent the slow normalization of risk.

In high-demand industrial environments, clarity at the boundary between safe and unsafe operation is not optional. It is the difference between controlling risk and being controlled by it.

Abbreviations

API	American Petroleum Institute	**NOL**	Normal Operating Limit
CCR	Continuous Catalyst Regeneration	**NOP**	Normal Operating Procedure
C&E	Cause and Effect	**OSHA**	Occupational Safety and Health Administration
CHAZOP	Controls Hazard and Operability Study	**OSP**	Optimal Set Point
EOP	Emergency Operating Procedure	**PHA**	Process Hazard Analysis
ERP	Emergency Response Plan	**PSM**	Process Safety Management
ESD	Emergency Shutdown	**PST**	Process Safety Time
ESL	Equipment Safe Limits	**PSV**	Pressure Safety Valve
GHT	Gasoline Hydrotreater	**PVHH**	Process Variable High-High
HAZOP	Hazard and Operability Study	**PVLL**	Process Variable Low-Low
IEC	International Electrotechnical Commission	**RAGAGEP**	Recognized & Generally Accepted Good Engineering Practice
IOW	Integrity Operating Window	**RCA**	Root Cause Analysis
IPL	Independent Protection Layer	**SDL**	Safe Design Limit
ISA	International Society of Automation	**SIF**	Safety Instrumented Function
LOPA	Layers of Protection Analysis	**SIS**	Safety Instrumented System
LOPC	Loss of Primary Containment	**SME**	Subject Matter Expert
MAWP	Maximum Allowable Working Pressure	**SOL**	Safe Operating Limit
MAWT	Maximum Allowable Working Temperature	**SRS**	Safety Requirement Specification
MI	Mechanical Integrity	**SRU**	Sulfur Recovery Unit
MoC	Management of Change	**TOL**	Target Operating Limit

References

[1]. U.S. Department of Energy (DOE), *Fact of the Week #1117: The Number of U.S. Crude Oil Refineries Has Declined but Total Distillation Capacity Has Risen From 1982 to 2019*, January 20, 2020, data from U.S. Energy Information Administration (EIA).

[2]. Center for Chemical Process Safety (CCPS). *Guidelines for Writing Effective Operating and Maintenance Procedures*. Hoboken, NJ: John Wiley & Sons, 1996, Section 6.2.

[3]. American Petroleum Institute (API). *API Recommended Practice 584: Integrity Operating Windows*. 1st ed. Washington, DC: American Petroleum Institute, 2014.

[4]. Center for Chemical Process Safety (CCPS), *Guidelines for Initiating Events and Independent Protection Layers in Layers of Protection Analysis* (Hoboken, NJ: John Wiley & Sons, 2015), Section 3.3.1, p. 41.

[5]. Center for Chemical Process Safety (CCPS). *Guidelines for Safe and Reliable Instrumented Protective Systems*. New York, NY: American Institute of Chemical Engineers, 2001, Section 5.5.5.