Technical Paper

IIAR Natural Refrigeration Conference
& Heavy Equipment Expo

March 2 – 5, 2025
Phoenix, Arizona

ACKNOWLEDGEMENT

The success of the IIAR Natural Refrigeration Conference is due to the quality of the technical papers in this volume and the labor of its authors. IIAR expresses its deep appreciation to the authors, reviewers, and editors for their contributions to the ammonia refrigeration industry.

ABOUT THIS VOLUME

The views expressed in this technical paper are those of the author(s), not the International Institute of All-Natural Refrigeration. They are not official positions of the Institute and are not officially endorsed.

International Institute of Ammonia Refrigeration
1001 North Fairfax Street, Suite 503
Alexandria, VA 22314

703-312-4200 • info@iiar.org • www.iiar.org

Technical Paper #4

Livin' la Vida LOPA: A Case Study in the Application of HAZOP and LOPA for Ammonia Refrigeration Systems

Michael Lacher, Process Safety Lead,
Food Solutions North America, Cargill

James Topp, PHA Manager,
Provenance Consulting, A Trinity Consultants Company

Abstract

Since the advent of process safety, the process hazard analysis (PHA) method of choice across the refrigeration industry has been the what-if/checklist technique. Over the last 20 years, many companies' and industries' hazard assessment methods have matured. While the hazard and operability (HAZOP) study and layer of protection analysis (LOPA) have become the norm outside of ammonia refrigeration, the methods are less commonly applied in the refrigeration industry. In early 2022, Cargill's Protein North America business began employing these PHA methods to evaluate process safety risk in ammonia refrigeration systems. Since that time, we have completed nearly 30 studies. This case study examines the practical applications of HAZOP and LOPA for ammonia refrigeration systems based on experience and discusses one company's journey to improve PHA practices

The case study begins with a brief overview of the HAZOP/LOPA method, then delves into the practical application of those methods. It discusses system noding, brainstorming causes, and evaluating consequences and risk. Then, we investigate safeguards and their effectiveness while sharing key lessons and challenges along the way. Finally, we discuss the various benefits of the method as well as its limitations.

Prologue

"Hello…" I answered my cell phone groggily. It was around 3 o'clock in the morning. As an on-call supervisor, 3 AM wake-up calls were not out of the ordinary, but this was not my week to be on-call. On the other end of the line was Jane*, a trusted night shift operator on my team. She was calm but clearly a bit shaken.

"The surge tank collapsed," she said matter-of-factly with a hint of disbelief.

"What?" I asked, still half asleep and obviously very confused.

"The surge tank, it collapsed," she reiterated. "John* is already here…"

"Is everyone ok?" I interrupted her.

A few days before the incident, I was working with several operators and maintenance technicians on a piece of equipment directly underneath the collapsed tank. In wet weather, we frequently had problems with equipment alignment and plugging—nothing that could not be resolved by a couple of experienced plant folks. A walkway a few feet from the bottom of the surge tank allowed operators and management to make frequent trips between the site's offices and the plant control rooms. Another piece of equipment, prone to leaking grain, sat directly beneath the surge tank, and the area required frequent cleaning. It seems as if someone was always standing underneath the surge tank cleaning spilled grain.

I arrived to work shortly before dawn that day. When the sun rose, and we could finally survey the damage from the night's incident, its severity became clear. Because of a previously unrecognized hazardous flow characteristic, the structural integrity of the tank had been compromised over time. Thousands of pounds of grain suddenly fell several floors, bringing with it several pieces of equipment and twisting steel beams in the process. The energy was enough to twist a nearby 24” I-beam. Anyone standing under or near the area would certainly have been crushed.

“Everyone’s ok,” Jane* replied. “We’re going to be down for a while, but everyone is accounted for.”

My plant was lucky that day. I was lucky; my team was lucky. By chance, the failure happened early in the morning, and no one happened to be standing there at that time. In the aftermath of this incident, what became clear to me was that we, as an industry, must do better when it comes to understanding our process hazards. We must improve our ability to identify and control process hazards, and we need to understand our technologies and processes. We cannot accept being lucky because, someday, our luck will run out.

In an effort to continually improve our process safety programs, to provide the critical safeguards our employees, families, neighbors, and communities count on, and to push ourselves beyond luck, we present a key area for potential improvement. It is one of our most critical and effective tools for identifying, controlling, and ultimately reducing risk: the PHA.

Introduction

Since the promulgation of OSHA’s (Occupational Safety and Health Administration’s) and the EPA’s (Environmental Protection Agency’s) rules regarding process safety, the process hazard analysis (PHA) method of choice across the refrigeration industry has been the what-if/checklist technique. Over the last 20 years, many companies’ and industries’ hazard assessment methods have matured. While the hazard and operability (HAZOP) study and layer of protection analysis (LOPA) have become the norm outside of ammonia refrigeration, the methods are less commonly applied in the refrigeration industry. In early 2022, Cargill’s Protein North America business began employing these PHA methods to evaluate process safety risk in ammonia refrigeration systems. Since that time, we have completed more than 30 studies. This case study examines the practical applications of the HAZOP study and LOPA for ammonia refrigeration systems based on experience and discusses one company's journey to improve PHA practices.

This case study begins with a brief overview of the HAZOP/LOPA method, then delves into its practical application. We discuss system noding, brainstorming causes, and the evaluation of consequences and risk. Then, we investigate safeguards and their effectiveness. In each section, after the presentation of how HAZOP and LOPA have been applied to refrigeration systems, we discuss Cargill's learnings, challenges, and efficiencies gained along the way. In addition, we also share, where appropriate, areas for potential industry involvement, discussion, and improvement in various areas, such as incident sharing, automation and controls practices, and system documentation. Finally, we discuss the various benefits of the method and its limitations when applied to typical ammonia refrigeration systems.

Cargill's application of HAZOP and LOPA is unique, blending the approaches together at times as well as incorporating some aspects of an FMEA (failure mode and effects analysis), and we attempt to dig beyond the surface and probe our systems further. We tend to apply many principles of LOPA in the HAZOP studies we complete, assessing and assigning values to cause frequencies, safeguarding effectiveness, and using semi-quantitative values for event frequency and severity. Our studies also use a cause-by-cause approach, which provides more precise results compared to others (Center for Chemical Process Safety, 2008).

As a final piece of introduction, we must state the obvious: we at Cargill are not perfect. We have not found any magic formula for perfectly consistent, perfectly acceptable, all-encompassing PHAs, where no risk goes unmitigated or unmanaged. We do not have all the answers, and sometimes, we do not even know what question to ask. We still have much to learn about the HAZOP/LOPA process, how to best apply it to our refrigeration systems, and how it might fit more broadly in such a large, diverse industry like ammonia refrigeration. However, we have put forth a concentrated and concerted effort to improve our hazard analyses. In addition to challenging many assumptions, we have spent countless hours working to properly identify and evaluate process safety risks. While we work tirelessly to eliminate luck as a factor, our focus is on reducing risk to eliminate catastrophic releases of ammonia in the future. With still so much to learn and improve upon, we have valuable insights to share with you.

Basics of HAZOP

A HAZOP study is a qualitative, systematic process that identifies and evaluates the safety and operability hazards of a process. Its purpose is to systematically review a process in its entirety to determine if deviations from design conditions can result in unwanted or dangerous consequences, such as a fire or release of ammonia. In a HAZOP study, every system component is assessed and analyzed to identify potential causes of these deviations. Once a consequence is identified, the team can then assess whether appropriate safeguards and mitigation are in place or if additional safety measures may be required to meet the company's risk tolerance (Center for Chemical Process Safety, 2008).

During the HAZOP study, an experienced facilitator and a team with detailed knowledge of the process being evaluated use a series of guidewords and process parameters to identify potential deviations from safe operating conditions. For example, the guideword "no" when combined with the parameter "flow" creates the deviation "no flow." The team can then brainstorm causes that would lead to a loss of flow in the node being studied. Tables 1 and 2 demonstrate how guidewords and process parameters can be paired to create a set of deviations for a HAZOP study.

Table 1. HAZOP study guidewords

Guideword	Meaning
NO or NOT	Negation of the design intent
LESS or LESS OF	Quantitative decrease
MORE or MORE OF	Quantitative increase
PART OF	Qualitative decrease
AS WELL AS or MORE THAN	Qualitative increase
REVERSE	Logical opposite of the intent
OTHER THAN	Complete substitution

Table 2. Common HAZOP study process parameters

Flow	Time	Frequency	Mixing
Pressure	Composition	Viscosity	Addition
Temperature	pH	Voltage	Separation
Level	Speed	Information	Reaction

As is apparent, many deviations can be created by combining a guideword with other process variables. Other examples of deviations include “REVERSE FLOW,” “MORE TEMPERATURE,” and “LESS PRESSURE.”

Once the team settles on a set of deviations for a node and brainstorms a possible cause, it identifies the potential consequences of the cause. When describing a consequence, the team must complete its evaluation in the absence of any safeguards to identify the worst credible outcome. The team members should consider the design of the system, the location of equipment, and accessibility to various areas of the facility, along with the known hazards of the process being studied. For example, high pressure can be caused by the inadvertent operation of a compressor discharge valve. It can lead to the overpressurization of a system component and release ammonia, which can fatally injure an operator who is in the area.

After evaluating the likelihood of the cause and severity of the worst credible consequence (details on which are provided later in this paper), the team is left with a qualitative measure of unmitigated risk for the given scenario. Then, the team members list available safeguards that may prevent or mitigate the severity of the ultimate consequence (more details on safeguards are provided later in the paper as well). For our example, almost all compressors are equipped with high-pressure and high-temperature shutdown devices that prevent an overpressure significant enough to result in a compressor failure. Each machine is equipped with pressure safety valves, and machinery rooms have ventilation systems that mitigate this consequence. The team would list and document these safeguards. At this point, the team can assess the risk and make a judgment for this cause–consequence pair. Is the risk effectively mitigated? If the team feels so, they move to the next cause; if not, they document the discrepancy with a recommendation.

This process is repeated until the team has studied all causes and deviations for a node and then further applied to other nodes in the process until the entire system has been analyzed.

Causes that result in costly outcomes, such as lost production, increased maintenance expenditures, and product spoilage, can be studied by the HAZOP team if tasked. These same procedures can be used to assess many types of business risk (in addition to safety and environmental). This includes product spoilage risk, production volume losses, cost exceedance risk, and recall/reputational risk. However, focusing on process safety risk is common, that is, the risk associated with the release of hazardous material or energy that can impact the health and safety of internal stakeholders, the community, or the environment. Going forward, we focus exclusively on process safety risk.

Basics of LOPA

LOPA is a form of risk assessment that employs order-of-magnitude category estimations of event frequency, consequence severity, and likelihood of failure of protection layers to estimate risk. The LOPA method offers a semi-quantitative approach to risk assessment and provides a middle ground between strictly qualitative PHA methods, such as what-if/checklists and HAZOP, and more complex quantitative risk analyses. LOPA requires another analysis to identify the causes and consequences. Typically, this more qualitative risk analysis is a HAZOP PHA. LOPA takes the results of the HAZOP a step further, which is particularly useful for potential high-severity events.

Once a scenario is selected for LOPA evaluation, the team first confirms the scenario has a specific initiating event based on a single failure related to human error, equipment failure, or an external cause. Order-of-magnitude frequencies, modifiers, and independent protection layer (IPL) failure probabilities are applied to estimate unmitigated and mitigated risks. The team can then compare the mitigated risk to the facility's accepted risk tolerance to determine if a certain scenario is thoroughly mitigated or if additional risk reduction is required. For example, does the initiating event have a 1% or 0.1% chance of occurring per year? Are operators present 100% of the time or only 10% of the time? Is the probability of failure of an IPL 10%, 1%, or 0.1%? These probabilities, which can be based on actual site failure rates, if enough data are available, or industry guidance, are applied to a scenario and combined to obtain a semi-quantitative risk level that can inform the LOPA team if the scenario is properly mitigated.

HAZOP and LOPA for Ammonia Refrigeration Systems

In the following sections, we dig into the practical application of HAZOP and LOPA for an ammonia refrigeration system. Starting with the selection of deviation guidewords and noding, through the risk assessment, each step of the process includes a description of how it was applied to our refrigeration systems, using examples when possible, and a discussion around its application.

Selecting Deviations

Explanation

Typically done well in advance of the study, often by a central process safety group or perhaps pre-populated by a contractor facilitating the PHA, deviation selection is a straightforward process that is not discussed at length in this paper. If not already provided for by a central corporate entity or HAZOP facilitator, the team can brainstorm deviations using Tables 1 and 2 or refer to Table 3 for deviations that have worked well in our experience.

Application and Discussion

Table 3 represents a commonly applied set of deviations along with some typically identified causes and immediate consequences. These deviations are usually applied to all nodes of the PHA.

Table 3. Common deviations and causes

Deviation	Causes
Low/No Flow	Inadvertent closure of manual valves. Malfunctioning of automated valves and/or regulators closed.
More/High Flow	Inadvertent opening of manual metering or expansion valves. Malfunctioning of automated valves and/or regulators open.
Reverse/Misdirected Flow	Failure of three-way mixing valves in either position.
Less/Low Temperature	Failure of three-way mixing valve on a compressor. Load changes on a vessel resulting in a decrease in temperature.
More/High Temperature	Failure of three-way mixing valve on a compressor. Blocked compressor discharge/suction. Failure of an oil heating element.
Less/Low Pressure	Malfunction of an automated compressor load control system resulting in low pressure in the suction piping and vessel.
More/High Pressure	Inadvertent closure of manual compressor discharge valve. Loss of condensing water or fans.
Less/Low Level	Malfunction of a vessel level controller resulting in vessel emptying.
More/High Level	Malfunction of a vessel level controller resulting in vessel overfill.
Leak/Rupture	Possible evaporator coil leak. Oil cooler tube leak/ rupture.

Note that the list shown is only a partial one. When applied by a knowledgeable PHA team during deviation brainstorming and with detailed process safety information, the process yields numerous causes to be studied in the PHA. Many causes—such as blocking a compressor discharge valve—can be evaluated in more than one deviation (e.g., "high pressure" or "no flow"). The cause does not need to be evaluated in each deviation, but there is no harm in discussing the scenario. Teams often "run out" of new causes as they move down the list of deviations because many causes were already studied as part of a previous deviation.

System Noding

Explanation

Noding is the process by which a large, complex system is broken down into smaller, more manageable portions so that the team can more easily direct attention. Noding is critical because it sets the boundaries and process limits for the study in its entirety, as well as rules for the deviation brainstorming process.

Application and Discussion

The team can break up a system into different nodes in various effective ways, but we provide some good rules of thumb that can help smooth the PHA process.

First, the size of each node should be manageable and somewhat easy for the PHA team to digest and understand. It should be selected so that it is easy to navigate and visualize for a group of experienced refrigeration professionals but complex enough not to oversimplify or overly complicate the study by adding too many superfluous or tiny nodes.

Second, grouping identical (or nearly identical) pieces of equipment into a shared node is often a very helpful approach and a significant time-saver. Many facilities have multiple high-stage compressors that are very similar, if not identical. These machines can be grouped together in a single node, reducing complexity and saving time. Condensers and evaporators are additional examples of commonly used and typically duplicated types of equipment. Where functionally identical or similar, the team may choose to group the equipment to reduce any redundancy during the study.

The example in Table 4 shows a relatively simple two-stage refrigeration system broken down into 11 nodes. This hypothetical system has five high-stage screw compressors, two low-stage booster compressors, one reciprocating pump-out compressor, a low-temperature recirculator, a high-temperature recirculator, a high-pressure receiver, three evaporative condensers, a plate and frame (P&F) condenser, ten shipping/receiving dock ceiling hung evaporators, six penthouse units for cooling a production space, and a tunnel freezer.

Table 4. Example ammonia system nodes

Node	Description
1. High-stage screw compressors (C1–C5)	Same manufacturer but varying in model and size. C5 is cooled via liquid injection.
2. Low-stage screw compressors (C6 and C7)	Same manufacturer. Thermosiphon-cooled.
3. Pump-out compressor (C8) and vessel	Reciprocating compressor, vessel, and associated pump-out header.
4. Evaporative condenser (EC1–EC3)	EC1—forced draft with no local sump. EC2 and CD3—induced draft with a shared local sump.
5. P&F condenser (E1)	
6. High-pressure receiver (V1)	Integrated high-pressure receiver and thermosiphon.
7. High temperature recirculator (V2)	Includes two centrifugal pumps.
8. Low temperature recirculator (V3)	Includes two centrifugal pumps.
9. Dock evaporators (AU1–AU10)	Air defrost only.
10. Production penthouses (AU11–AU16)	Hot gas defrost with gas-powered suction valves. AU15 and AU16 have motorized suction valves because they were installed later.
11. Tunnel freezer (TF1)	

In the above example, a few areas require further discussion and evaluation. Note that compressor C5 is cooled via liquid injection, and C1–C4 do not have their cooling mechanisms stated, but we can assume they are thermosiphon-cooled. Should C5 be separated into its own node because of the different cooling technology? The answer depends on the facilitator and the team's knowledge, but a good argument can be made to treat C5 as a separate node because the two types of compressors have several different potential causes.

A similar discussion can be made regarding the penthouse air units AU15 and AU16. They are similar but employ different styles of suction valves. Ultimately, the team decides, guided by the facilitator, if the difference is significant enough to warrant them being studied as part of a separate node.

Finally, the noding process should be applied flexibly. It is common to identify the nodes for a study during the preparation phase only to find that a piece of equipment was incorrectly noded or better fits in a different node. This can easily be addressed as the study progresses.

Identifying Causes

Explanation

Once a set of deviations is agreed upon and the study is broken down into manageable nodes, the HAZOP study can start in earnest. The first step is for the team to brainstorm causes. Typically, the facilitator probes the team to ascertain the causes of the deviation being studied.

"In this node, what could cause no flow?" is a typical starting point that helps the team to identify any valves (manual or automatic) or motors that, when malfunctioning or misaligned, can result in a low- or no-flow condition. This will be documented as a cause.

Initiating events or causes fall into one of three broad categories: external events, equipment failures, or human failures (Center for Chemical Process Safety, 2001). External events include natural weather or geographic phenomena, impacts from fires or explosions at adjacent facilities, and third-party interventions, such as an impact by a motor vehicle or other construction equipment. Equipment-related events include electrical and mechanical failures of equipment and controls, process control failures, software failures, or failures of rotating equipment due to vibration or lapses in proper maintenance. Lastly, human failure is the result of an operator executing the steps of a task improperly or not responding appropriately to a condition or other system prompt. When assessing a system for potential causes of deviations, all types of initiating events should be covered.

Of note, certain issues do NOT qualify as initiating events. Management systems, for instance, are not usually listed as an initiating event, even though they are often the root causes of human error. A poorly written or confusing procedure does not, in and of itself, cause a human error. In many cases, operators make good decisions despite poorly documented procedures. However, that same procedure is likely to contribute to an operator closing the wrong valve at the wrong time, which IS an initiating event that can cause a deviation.

The initiating event or cause must be well understood and specific to proceed through the HAZOP study effectively and, especially, to complete a successful LOPA later. Take, for example, one cause of high temperature in a compressor—loss of oil flow. Loss of oil flow as a cause of high temperature is too vague and does not work well when assessing risk later in the HAZOP study and LOPA. In fact, many variables can lead to a loss of oil flow: a continuous oil pump can malfunction, a manual valve may be inadvertently closed, or an oil filter can become blocked over time. Furthermore, even the location of a manual valve in the oil flow circuit of a compressor matters. In some instances, certain oil pressure and temperature safeguards may or may not be effective depending on valve arrangement and the mis-isolated valve. For this reason, being specific is important when defining initiating events. Not only can the effectiveness of safeguards vary depending on the location of components, but the frequency with which failures occur also varies depending on the initiating event.

Application and Discussion

Typically, after reviewing a few examples, PHA teams quickly catch on, and cause identification becomes easier. Often, some, if not many, causes end here. That is, we accept an operational inefficiency or see no process safety consequence at all. For example, the causes can result in increased energy consumption or necessitate an earlier replacement of wear components for a particular machine, but they are not expected to result in a release of ammonia from the system. The team should document this and move on.

However, some causes do lead to consequences with the potential to harm internal or external stakeholders. These are referred to as consequences of interest. Whether there is potential for minor, reversible injuries or catastrophic impacts to the surrounding communities, we must assess these *consequences of interest* and understand the risk associated with each. These consequences and their associated causes are called *cause–consequence pairs*, which are the focus of the remainder of this paper. Once a consequence of interest and its cause are identified, the real work of assessing risk can begin.

Assessing Risk

To first assess risk, we must understand what risk is. According to AIChE's website, risk is defined as "a measure of human injury, environmental damage, or economic loss in terms of both the incident likelihood and the magnitude of the loss of injury" (Center for Chemical Process Safety, 2024). Furthermore, it can be visualized as a combination of three factors:

- Hazard—What can go wrong?
- Consequence—How bad could it be?
- Likelihood—How often might it happen?

When evaluating risk as it relates to process safety in a HAZOP study or LOPA, we can express it as the product of the frequency of a certain event and the consequence of the event:

$$\text{Frequency (events/yr)} \times \text{Impact (impact/event)} = \text{Risk (impact/yr)}$$

Industries and companies have different measures for risk, and some have been previously discussed. The difference is primarily in the way the impact is measured, and one common method is quantifying the type or types of injury(ies) expected during an event. The combination of a potentially fatal impact with an event that has a 1% chance of happening each year leads to a risk value of one chance in 100 of a fatality happening in any given year.

The most used tool to help visualize the risk equation is a risk matrix. A risk matrix can be qualitative or semi-quantitative, and it incorporates the frequency and impact parameters of the equation on two axes.

Table 5. Example risk matrix

Frequency	**Frequent** *1/yr*	**Probable** *0.1/yr*	**Occasional** *0.01/yr*	**Rare** *0.001/yr*	**Improbable** *0.0001/yr*
Severity					
Catastrophic *(≥4 fatalities)*	IV	IV	IV	III	II
Major *(1–3 fatalities)*	IV	IV	III	II	II
Serious *(lost-time accident)*	IV	III	II	II	I
Minor *(medical aid)*	III	II	II	I	I
Insignificant *(first aid/no injury)*	II	II	I	I	I

Table 5 provides a simplified and generic example of a risk matrix similar to the one applied in our experience, with several critical features. First, across the top (x-axis) are order-of-magnitude frequencies for initiating events. The frequency for an initiating event can be viewed as the probability that the event will happen in a given year (e.g., a 1% chance of happening each year is 1 in 100 or 0.01/yr).

The vertical axis in Table 5 represents the severity of the consequence, ranging from *catastrophic* to *insignificant*. In this example, *catastrophic* severity is defined as an event resulting in more than four fatalities. From here, the consequences progressively get less severe ending with *insignificant*, where the worst credible consequence includes minor first aid or no injury.

The intersection of a row and column represents a requirement for risk reduction to determine acceptability. Approximately speaking, red is unacceptable, orange is rarely acceptable, yellow is potentially acceptable, and green is broadly acceptable. Because we cannot completely eliminate risk, we must define a level of risk that is acceptable. Many factors come into play when determining an acceptable level of risk, and each facility should closely evaluate what is acceptable and carefully consider the design of a process safety risk matrix before completing any PHA, especially any HAZOP study or LOPA. The establishment of an acceptable risk tolerance is outside the scope of this paper; however, it is a critical component of any process safety program. For simplicity and ease of understanding, the risk matrix in Table 5 accepts that a catastrophic event happens every million years and that a major fatality event occurs every 100,000 years. Further discussion on how we arrive at these probabilities is provided later in the paper.

Determining and Evaluating Consequences

Explanation

The next step in assessing process safety risk after a consequence of interest is identified is to ensure the consequence is well-defined. A well-defined and properly written consequence is essentially a detailed time-sequence description of events that can unfold when the initiating event occurs. It should be specific and use equipment tags where possible, and it should include a written description of the impact as assessed by the team. For example, for node 1 of our hypothetical system described in Table 4, the system's high-stage compressors are under the no/low flow deviation, as described in Table 6.

Table 6. Example HAZOP scenario—Step 1

Deviation	Cause	Consequence	Severity	Initiating Event Frequency
No/Low Flow	HSD-C01 manual valve inadvertently closed (compressor discharge isolation valve)	Loss of flow from compressor C1 resulting in increases in pressure and temperature in the compressor discharge. Potential to exceed the design parameters of the compressor, resulting in compressor damage and a release of ammonia into the machinery room. Potential for two operator fatalities.		

Note that both the cause and the consequence contain descriptions and equipment tag numbers for future reference. Additionally, the consequence must be described in the absence of any preventive or mitigative safeguards. This is usually a challenging aspect for teams to understand, and it arises frequently early on in any PHA.

Application and Discussion—Without Safeguards

According to both Process Safety Management (PSM) and Risk Management Program (RMP) regulations, “The process hazard analysis shall address: … Consequences of failure of engineering and administrative controls” (Occupational Safety and Health Administration, 1992). As a result, PHA practitioners typically adopt a worst-case approach and assume that a scenario can progress to its ultimate consequence without the intervention of any safeguards. This can be initially challenging to many and especially challenging to experienced refrigeration and maintenance mechanics who have spent their entire careers working on the very systems they must assume will fail. However, after training and working on a few examples, most catch on very quickly. Take the following discussion as an example, building on the consequence above:

Facilitator: If we block the compressor discharge valve, what will happen?

Operator: The compressor will shut down.

Facilitator: Why will it shut down?

Operator: Well, high pressure will cut it out pretty quickly!

Facilitator: Yes, that's right, but what if the high-pressure cutout wasn't working?

Operator: It'll still go down at high temperature.

Facilitator: Ok, that's a safeguard as well, and remember, we need to determine what might happen without any safeguards.

Operator: I suppose then the reliefs will lift…

Facilitator: Ok, so now we have an overpressure and a release, but those PSVs (pressure safety valves) relieve safely; what if they were also not there?

Operator: I guess at that point, something is going to break.

Facilitator: Ok, something breaks, and now we have a pretty serious ammonia release in the machinery room.

This scenario illustrates the typical back and forth as a team learns the rules of the road, but it also is an excellent foreshadowing of the safeguard identification step that is discussed later. Once we strip away the safeguards, we can obtain a clear picture of the worst-case outcome of the cause. In addition, we can assess what may happen if the applicable safeguards fail, satisfying regulatory requirements and setting ourselves up for a thorough analysis later in the study.

Application and Discussion—Worst Credible Outcome

Another question that frequently arises during PHAs is how to determine the worst outcome and what the difference is between the worst possible and worst credible outcomes. To determine risk, we must make a judgment on the impact of a potential incident. We need to answer "How bad can it be?". Furthermore, this risk assessment is required to properly complete the HAZOP study and LOPA.

Using our compressor overpressure and rupture scenario, we can ask ourselves what a *credible* outcome is, as opposed to what is *possible*. Is it possible that a large tour of company executives is in the machinery room at the time of the release? Is it *possible* that you have invited a dozen community volunteer firefighters for a walk-through of your facility when it happens? Yes, both are possible. However, it may be equally (or perhaps even more) likely that no one will be in the machinery room when the incident happens.

This is where a PHA team must determine what is credible. The incident occurring during one of the previously mentioned tours resulting in mass casualties is not a credible situation. These types of gatherings or special tours may happen once in the life of a facility, if at all. What is credible, however, is an operator or two being present in the machinery room, completing PMs, rounds, or any other task requiring their presence at any time. Furthermore, it is credible that a large release of ammonia can fatally impact an operator in the area or one entering the area without knowing a release has occurred. In this example, one or two fatalities are credible; therefore, the PHA team assesses risk based on this worst credible consequence.

Application and Discussion—Severity

With the *worst credible consequence* in mind, the PHA team can now determine its severity. Using Table 5 in conjunction with the discussion from the HAZOP consequence description, we can see the potential for two operator fatalities, which means, in our example, the severity can be classified as *major*.

Often, however, agreeing on the ultimate outcome and severity of a consequence can be difficult. Disagreement can arise anywhere along the consequence chain of events. Take the example in question: Would the overpressure cause a catastrophic rupture or maybe just blow out the compressor seals? Would the motor's overcurrent protection or the physical limitations of the motor even allow for the generation of pressure sufficient to rupture the compressor's casing? If there was a rupture, where would the rupture occur? Could it rupture downstream of the discharge check valve, creating an opening from the entire high side of the system? Another example that we frequently see is the operation of vessels below their minimum design metal temperature (MDMT) or their maximum allowable external working pressure (MAEP). Even with no knowledge of any failures, a team must consider the implications of operating a vessel outside its design limits.

In most cases, the worst credible consequence is exceedingly rare and usually outside the experience of any PHA team member. PHA teams often must deal with the unknown and, ultimately, come to a consensus on what result is credible.

Initiating Event Frequency

Explanation

After defining a consequence of interest for a discrete initiating event and its severity, the team makes the next risk judgment: the frequency with which the initiating event will happen. Because of the likely rarity of the worst credible consequence, assessing its frequency can be difficult. For this reason, PHA teams first assess the frequency of the initiating event, which they can typically understand more easily. While experience with major and catastrophic releases of ammonia is, thankfully, minimal, teams commonly have experienced failures that are actual initiating events. Often, operators describe experiencing high-level trips, control failures, and the like that were interrupted by a safeguard in the system.

For example, the team describes the annual problems they have with a condenser tank water level device that seems to break at least once a year. It has never caused problems; usually, someone catches the issue, and a few times, it goes unnoticed when the level in the tank falls low enough to starve the water pumps and the system shuts down on high head pressure. This example illustrates our point. The team has not experienced the entire consequence of having a major overpressure and toxic release, but they have started down that path multiple times over the last few years. They experienced the initiating event, but a negative consequence was, in fact, prevented by effective safeguards.

The process safety industry has spent significant effort gathering data from various sources to assist in determining equipment failure rates and initiating event frequencies. PHA teams and organizations should also consider company experience regarding equipment failure rates. In *Layers of Protection Analysis: Simplified Process Risk Assessment*, the Center for Chemical Process Safety (CPPS) provides typical frequencies of initiating events (Center for Chemical Process Safety, 2001), as shown in Figure 1. Note that ranges are given for each type of initiating event first, followed by the exact value a company or PHA team may choose to use in a LOPA or HAZOP study. CCPS provides additional guidance on typical frequencies in *Guidelines for Initiating Events and Independent Protection Layers in Layer of Protection Analysis* (Center for Chemical Process Safety, 2014).

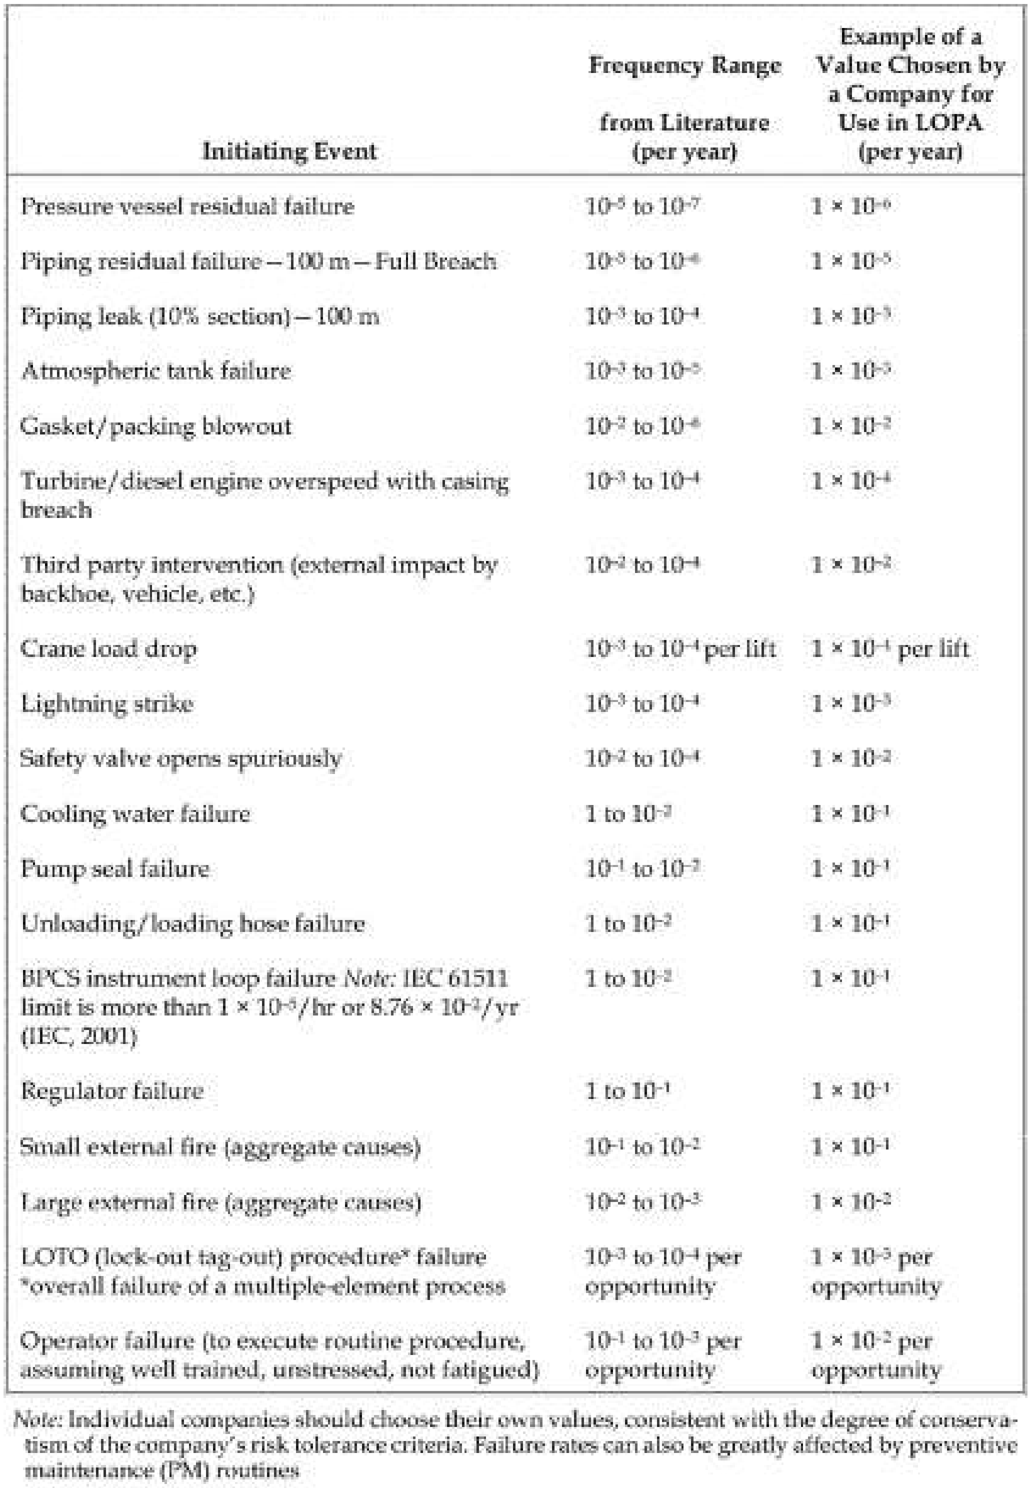

Initiating Event	Frequency Range from Literature (per year)	Example of a Value Chosen by a Company for Use in LOPA (per year)
Pressure vessel residual failure	10^{-5} to 10^{-7}	1×10^{-6}
Piping residual failure – 100 m – Full Breach	10^{-5} to 10^{-6}	1×10^{-5}
Piping leak (10% section) – 100 m	10^{-3} to 10^{-4}	1×10^{-3}
Atmospheric tank failure	10^{-3} to 10^{-5}	1×10^{-3}
Gasket/packing blowout	10^{-2} to 10^{-6}	1×10^{-2}
Turbine/diesel engine overspeed with casing breach	10^{-3} to 10^{-4}	1×10^{-4}
Third party intervention (external impact by backhoe, vehicle, etc.)	10^{-2} to 10^{-4}	1×10^{-2}
Crane load drop	10^{-3} to 10^{-4} per lift	1×10^{-4} per lift
Lightning strike	10^{-3} to 10^{-4}	1×10^{-3}
Safety valve opens spuriously	10^{-2} to 10^{-4}	1×10^{-2}
Cooling water failure	1 to 10^{-2}	1×10^{-1}
Pump seal failure	10^{-1} to 10^{-2}	1×10^{-1}
Unloading/loading hose failure	1 to 10^{-2}	1×10^{-1}
BPCS instrument loop failure *Note:* IEC 61511 limit is more than 1×10^{-5}/hr or 8.76×10^{-2}/yr (IEC, 2001)	1 to 10^{-2}	1×10^{-1}
Regulator failure	1 to 10^{-1}	1×10^{-1}
Small external fire (aggregate causes)	10^{-1} to 10^{-2}	1×10^{-1}
Large external fire (aggregate causes)	10^{-2} to 10^{-3}	1×10^{-2}
LOTO (lock-out tag-out) procedure* failure *overall failure of a multiple-element process	10^{-3} to 10^{-4} per opportunity	1×10^{-3} per opportunity
Operator failure (to execute routine procedure, assuming well trained, unstressed, not fatigued)	10^{-1} to 10^{-3} per opportunity	1×10^{-2} per opportunity

Note: Individual companies should choose their own values, consistent with the degree of conservatism of the company's risk tolerance criteria. Failure rates can also be greatly affected by preventive maintenance (PM) routines

Figure 1. CCPS initiating event frequency ranges (CCPS, 2001)

We can use the data in Figure 1 to estimate the initiating event frequency for our example from Table 6. Remember that the deviation of no flow was caused by the inadvertent closing of the manual discharge valve of a running compressor. The only way to close this valve is through human intervention, i.e., operator error. Figure 1 gives a range of 10^{-1} to 10^{-3} failures per opportunity for a well-trained and unstressed operator functioning with clear and documented instructions. The representative company chose a failure rate of 10^{-2}, which we will also apply for this example. This failure rate describes the rate of human failure per opportunity (as opposed to per year), so we will need to assess the failure further to determine its per year rate. For simplicity, we assume an operator is required to operate the discharge valve of any given compressor once per year. The following equation represents how we can arrive at failure rate per year:

$$\frac{1\,opportunity}{year} \times \frac{1\,failure}{100\,opportunities} = \frac{1\,failure}{100\,years}\ or\ 10^{-2}\ failures\ per\ year$$

With our initiating event frequency determined, we have a clear picture of the level of unmitigated risk for the consequence of interest being studied. Table 7 provides an updated picture of our HAZOP scenario:

Table 7. Example HAZOP scenario—Step 2

Deviation	Cause	Consequence	Severity	Initiating Event Frequency	Risk Reduction Required
No/Low Flow	HSD-C01 manual valve inadvertently closed (compressor discharge isolation valve)	Loss of flow from compressor C1 resulting in an increase in pressure and temperature in the compressor discharge. Potential to exceed the design parameters of the compressor, resulting in compressor damage and a release of ammonia into the machinery room. Potential for two operator fatalities.	Major	10^{-2}	III

Table 7 identifies a risk reduction value, which comes from the example risk matrix in Table 5 at the intersection of the consequence severity *major* and the occasional frequency of 0.01/yr (10^{-2}/yr), which we just determined. This risk reduction value is critical as we move into the next phase of our risk assessment process.

Application and Discussion

One specific area valuable to risk practitioners in ammonia refrigeration relates to our understanding and knowledge of industry-specific failure rates for the types of equipment common to our systems. While chemical industry failure rate data provide a solid starting point, additional details are helpful when HAZOP studies and LOPA are applied to refrigeration systems. Industry experience and failure rate data from manufacturers prove valuable in the following areas:

- Failure rates for refrigeration control systems (programmable logic controllers (PLCs))
- Failure rates for ammonia regulators and other mechanically controlled valves
- Failure modes and rates for rotating equipment (e.g., compressors, pumps, oil pumps, evaporator fans)
- Industry accident databases and types of failures

If available, these and other data specific to our industry provide valuable information to enable an accurate understanding of risk and improve risk mitigation across the industry.

Safeguards

Now that the work of assessing the unmitigated risk is complete, the PHA team moves on to the heart of a HAZOP study (and certainly the primary focus of a LOPA)—safeguards. First, we must consider the definition of a safeguard for our purposes. A safeguard is any device, system, or action that will likely interrupt the chain of events following an initiating event (Center for Chemical Process Safety, 2008).

With many different types of safeguards, some are more effective than others. A HAZOP study generally allows safeguards to be any measure effective at preventing the cause from leading to the consequence. LOPA, being semi-quantitative, requires quantification of the risk reduction provided by each safeguard. Since quantification is difficult or impossible for some safeguards, LOPA uses some simplifications that limit allowable safeguards to only those that meet the requirements of independent protection layers (IPLs). IPLs are safeguards that meet the IDEA principle—independent, dependable, effective, and auditable (Weber, 2024):

- Independent—Safeguards must be independent of each other and independent of the initiating event. Safeguards that share a common input (e.g., transmitter, detector) or output function (e.g., motor, valve) are not independent of each other.
- Dependable—Safeguards must be dependable. They must be maintained in such a manner that they will function as intended. A dependable safety cutout is one that has been tested per manufacturer and industry recommendations, maintained, and calibrated regularly.
- Effective—Safeguards must be effective at stopping the chain of events or mitigating the release to reduce risk. An effective pressure cutout device acts prior to exceeding the design pressure of any system component and stops the equipment (i.e., compressor) from creating the pressure.
- Auditable—Safeguards must have thorough documentation that demonstrates they meet the criteria above, i.e., they are designed, installed, and maintained to provide protection for the scenario being studied.

If a safeguard meets these criteria, it can be considered to provide credible IPLs. In the following sections, we explore the four IDEA principles more deeply and their application to ammonia refrigeration systems.

Independence

Safeguard independence is likely the most challenging and contested aspect of HAZOP and LOPA implementation for ammonia refrigeration systems. The two aspects of safeguard independence that must be true for the safeguard to be considered credible to reduce risk are as follows:

1. It must be sufficiently independent of the initiating event.
2. It must be sufficiently independent of other credible safeguards.

Take, for example, a scenario where a vessel might overfill as the result of a failed solenoid filling valve, e.g., the automatic valve fails in the open position. In this scenario, the vessel is equipped with a level transmitter that, upon detecting a high-level condition, closes the fill solenoid valve. While it may be present, the level control safeguard cannot be considered credible for this scenario because it is not independent of the initiating event. To stop the event progression, the level transmitter *depends* on closing the same valve that has failed in its open state. In this case, the safeguard is not independent.

Safeguards must also be independent of each other. Typically, IPLs cannot share the same input device, output function, or a common logic controller. Independence between IPLs and causes sharing a common control system is complicated, and often, simplified rules help a team address this independence. Figure 2 illustrates an example of guidance provided. These simplified flow diagrams represent common applications and discussions on how many credible safeguards can be assigned to each.

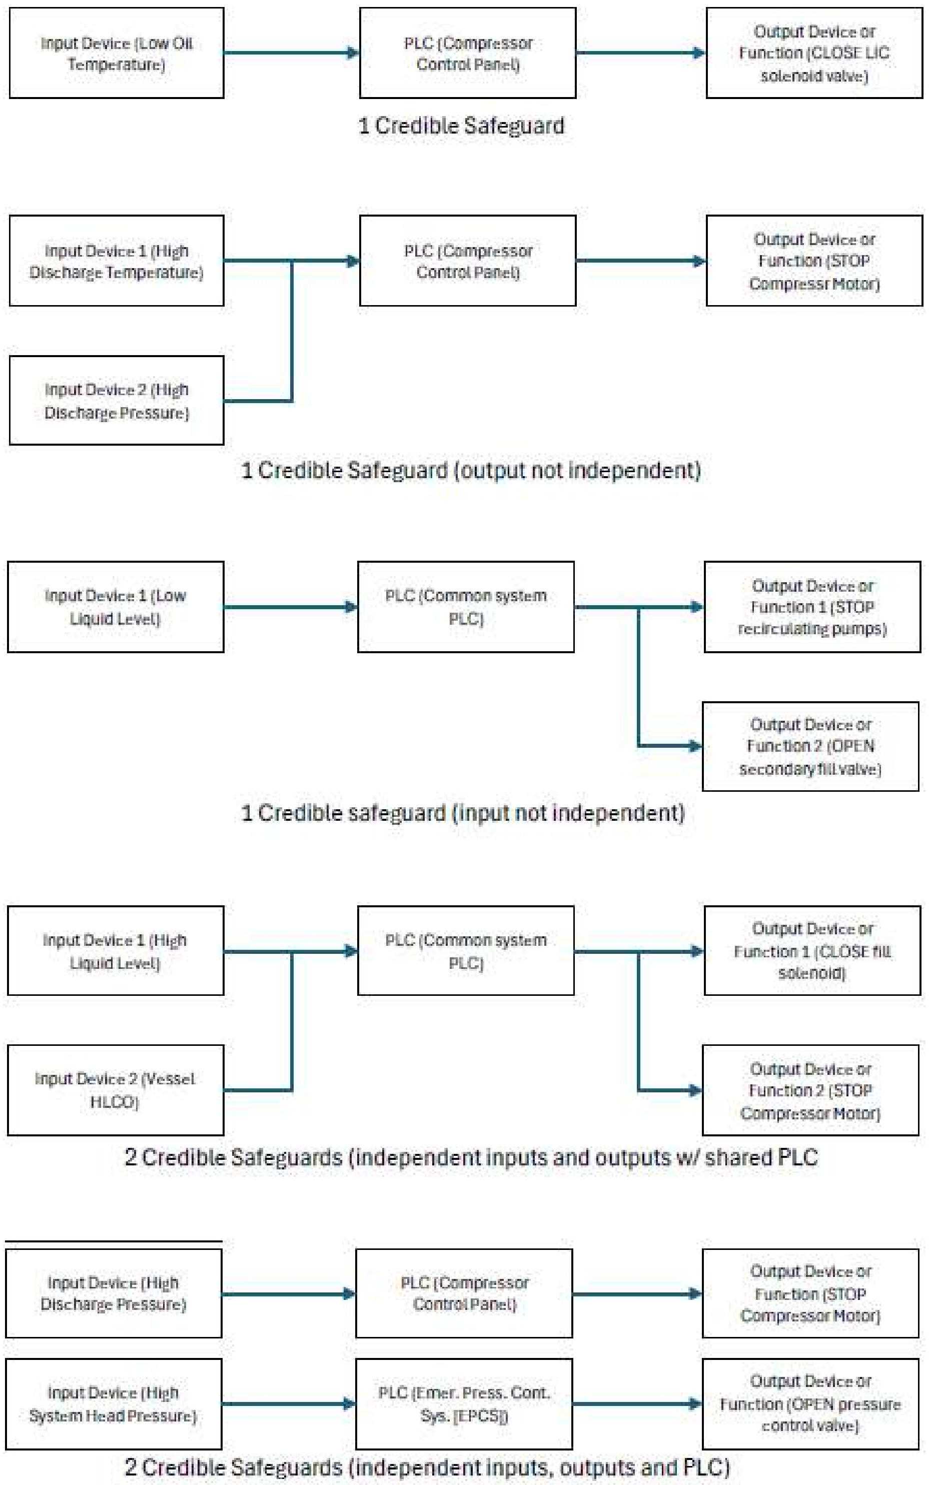
Input Device (Low Oil Temperature)
PLC (Compressor Control Panel)
Output Device or Function (CLOSE LIC solenoid valve)
Input Device 1 (High Discharge Temperature)
Input Device 2 (High Discharge Pressure)
PLC (Compressor Control Panel)
Output Device or Function (STOP Compressr Motor)
Input Device 1 (Low Liquid Level)
PLC (Common system PLC)
Output Device or Function 1 (STOP recirculating pumps)
Output Device or Function 2 (OPEN secondary fill valve)
Input Device 1 (High Liquid Level)
Input Device 2 (Vessel HLCO)
PLC (Common system PLC)
Output Device or Function 1 (CLOSE fill solenoid)
Output Device or Function 2 (STOP Compressor Motor)
Input Device (High Discharge Pressure)
PLC (Compressor Control Panel)
Output Device or Function (STOP Compressor Motor)
Input Device (High System Head Pressure)
PLC (Emer. Press. Cont. Sys. [EPCS])
Output Device or Function (OPEN pressure control valve)

1 Credible Safeguard

1 Credible Safeguard (output not independent)

1 Credible safeguard (input not independent)

2 Credible Safeguards (independent inputs and outputs w/ shared PLC

2 Credible Safeguards (independent inputs, outputs and PLC)

Figure 2. Safeguard diagrams

This independence requirement produces many challenges when performing a HAZOP study or LOPA on a refrigeration system. Because most of the risk associated with a refrigeration system comes from compressing the refrigerant, the compressors are a natural point of focus to automatically bring the system to a safe state. The compressors generate pressure, heat, and flow in the system, and turning off the compressor is the first action that can be taken to prevent the progression of the consequence. As a result, the compressor acts as the final output element of almost every safeguard for an ammonia system, which leads to an abundance of independence conflicts. In these cases, the output device or function, along with the PLC (the compressor control panel), is usually shared by multiple input safety functions. Figure 3 demonstrates the plethora of compressor safety devices that share a common PLC and output function.

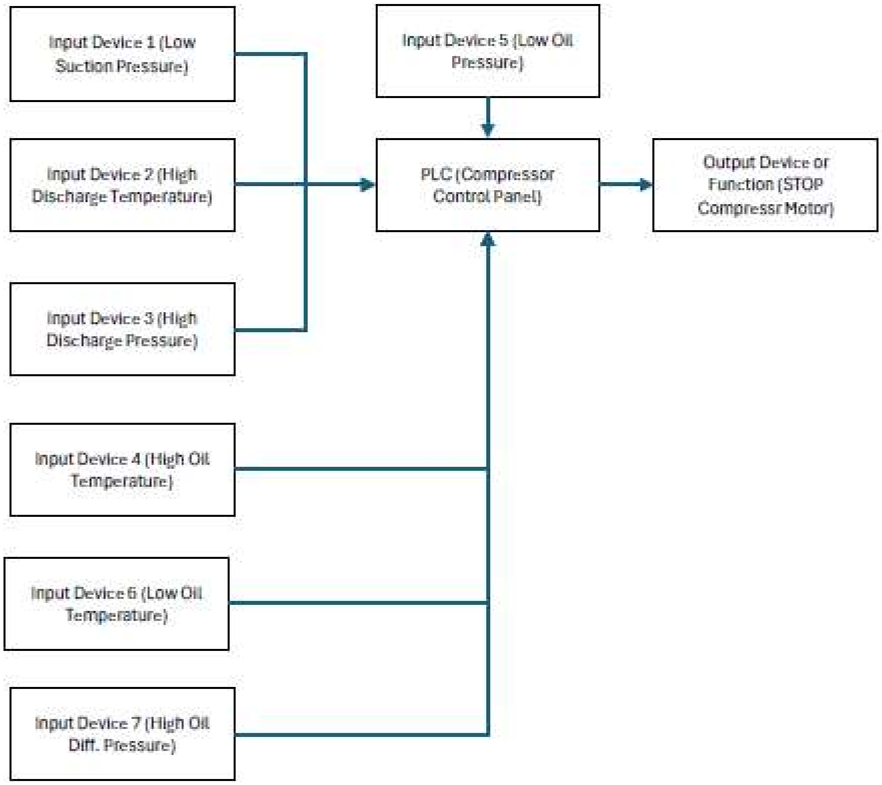
Input Device 1 (Low Suction Pressure)
Input Device 5 (Low Oil Pressure)
Input Device 2 (High Discharge Temperature)
PLC (Compressor Control Panel)
Output Device or Function (STOP Compressr Motor)
Input Device 3 (High Discharge Pressure)
Input Device 4 (High Oil Temperature)
Input Device 6 (Low Oil Temperature)
Input Device 7 (High Oil Diff. Pressure)

Figure 3. Compressor safeguard diagram

Although not inclusive of all compressor types, manufacturers, and designs, Figure 3 illustrates many typical safeguards on any given compressor. Clearly, typical rules for independence pose challenges when assessing an ammonia system with HAZOP and LOPA.

One mitigating factor described in industry involves the crediting of two safeguards associated with a shared logic controller as long as the input and output functions are independent. This approach assumes that if a safety loop fails, the component that induced the failure is more likely the detection device or the final control element, and a failure of the logic controller itself is less likely (Center for Chemical Process Safety, 2001). This approach is generally accepted in industry, with certain caveats and qualifiers, but it still requires the independence of both the input and output functions. For a typical ammonia refrigeration compressor, the PLC and the output function are shared. *This means that crediting even two compressor safety devices violates typical independence requirements for HAZOP and LOPA methods.*

To effectively employ HAZOP and LOPA methodologies in ammonia refrigeration systems, operators and designers must be aware of potential pitfalls in safeguard independence. Additionally, they must be prepared to deal with the discrepancies identified.

Dependability

Regarding dependability, can we count on our safeguards to act when necessary? Are we performing the proper maintenance, calibration, and testing of our safety devices? Was the safety device designed to reduce the risk of it failing when demanded? These are a few of the questions to ask when assessing a safeguard for dependability, and they provide clues to two essential factors regarding safeguard dependability:

1. How is the safeguard maintained?
2. How was the safeguard designed?

Safety device maintenance, including testing and calibration, is the first key element to ensuring dependable safeguards. Historically, the ammonia refrigeration industry has placed serious rigor and requirements around safety device maintenance. ANSI/IIAR 6 details many of the frequencies and methods for testing safety devices. Adherence to these requirements is a vital component to ensuring safeguards remain dependable. In our opinion and experience, the maintenance of safeguards is very well addressed and, in most cases, implemented effectively in ammonia refrigeration systems. Therefore, additional discussion in this regard is not included here.

The second consideration in dependability is the design of a safeguard or safety system. Generally speaking, the refrigeration industry excels at defining specific requirements for mechanical safety systems. For years, the International Institute of All-Natural Refrigeration (IIAR), as a RAGAGEP setting body, has fine-tuned requirements for critical ammonia refrigeration safeguards. Emergency mechanical ventilation is well-defined with clear requirements. IIAR 2 prescribes the various safeties required for compressors and includes system design conditions to reduce risk and make our systems safer. An understanding of and compliance with these industry standards are critical, which continues to serve the industry very well in efforts to reduce risk and improve safety.

One area where the refrigeration industry lags behind, however, is instrumented safety systems, or instrumented safeguards. As discussed in the prior sub-section, from an instrumented safety perspective, our systems rely heavily on many sensors and cutout devices that de-energize the compressor. These critical systems are our first line of defense against many potential catastrophic releases. High temperatures, high pressures, losses of lubrication, etc., can quickly lead to serious consequences down the road. While we maintain these critical safeguards with diligence, the opposite seems true in the design of these instrumented safeguards. Several years ago, at an industry conference, we performed a very informal and anecdotal survey of manufacturers. We asked if they offered "SIL-rated" refrigeration components and packaged equipment. Only one manufacturer even knew what SIL meant (safety integrity level).

An entire industry exists to deal with functional safety. IEC 61508 and IEC 61511 set forth many of the standards for safety instrumented systems (SISs) and SILs for electrical, electronic, and programmable electronic elements used to perform safety functions (IEC, 2010). The standards also include detailed design and safety requirements for all components of these systems. A full discussion and analysis of these requirements, their implications, and how they can be applied to refrigeration systems (as well as whether they should) is a topic for an entire paper in itself. Suffice it to say that, for this discussion and as far as we are concerned, an entire world of resources and potential improvements is available, which is an area where the ammonia refrigeration industry could grow significantly. Yes, these systems are costly, and yes, the risk reduction would be incremental, but the question of whether our systems are currently designed as safe as possible is worth asking.

Effectiveness

While the point of effectiveness may seem obvious, to be credible in stopping the event progression, a safeguard must be effective to do just that—prevent the ultimate consequence. The *effect* must *affect* the consequence. An effective safeguard must be able to accomplish two tasks:

1. Detect—It must first detect that an abnormal or unsafe condition exists.
2. Act—It must then act, in some way, to stop the event progression BEFORE the ultimate consequence occurs.

In many cases, this may be obvious. In a compressor dead-head scenario, the high-pressure cutout is a clear safeguard. The high-pressure device detects the high-pressure condition and automatically acts to cut out the source of the pressure, i.e., the compressor. However, in some cases, the detect and act requirements are much less clear.

For example, we have the same scenario, but instead of the high discharge pressure, let us consider the compressor's high oil temperature cutout. Because of the increasing pressure in the compressor, we would also expect the temperature to increase, which, in turn, increases the oil temperature of the compressor. However, the oil may continue circulating and cooling to some degree (though no longer within its design parameters). The oil will likely overheat, but when? What pressure might the compressor achieve before the oil temperature reaches a cutout level? Can the oil temperature sensor detect the increase in temperature? Certainly yes. Can it act in time? This can be debated; therefore, it may not be a credible safeguard for this scenario. Plausibly, the compressor can exceed its design pressure before the oil temperature exceeds its cutout limit. In this scenario, then, while the oil temperature sensor can detect the condition, it may not be able to act before the ultimate consequence. In cases where multiple safeguards *might* be applicable, we can apply the "first-out" principle: ask the team which safeguard or safety device would trigger first, second, third, etc., to determine which is the most appropriate.

Auditable

We despise the sentence, "If it's not documented, it didn't happen." Unfortunately, this statement offers a kernel of truth and is applicable to many aspects of life and work, including refrigeration safety and, specifically, safeguards. The final check for a credited safeguard in any HAZOP study or LOPA is auditability, the last component in the IDEA principle.

Each of the first three aspects—*independence*, *dependability*, and *effectiveness*—must be well-documented through the facility's process safety information and mechanical integrity files. This documentation serves to demonstrate that what we say in the PHA regarding safeguards is true and that any internal or external party can independently verify the information. The PHA team should discuss and address serious deficiencies in documentation. Some safeguards that are otherwise independent, dependable, and effective may not be creditable if they lack the appropriate documentation.

IDEA Discussion

Two additional points are worth discussing briefly here and more thoroughly in the future.:

- The independence and effectiveness of compressor oil separator pressure relief valves
- The ability of machinery room detection and ventilation to prevent fatalities due to toxic exposure

ANSI/IIAR 2-2021 permits compressor relief valves to be de-rated to the minimum flow rate of a compressor if the compressor is equipped with automatic capacity regulation and pressure-limiting devices. When applied, this means that the pressure relief valves (PRVs) on a compressor are no longer independent of the compressor's controller or its high-pressure cutout. Additionally, they are not capable of maintaining a safe pressure in the event of a high-flow failure (a runaway compressor). Because they lack independence and effectiveness in these situations, they cannot be credited as IPLs in many circumstances.

Because of the many variables and dynamics associated with a release and toxic exposure inside a machinery room, whether the room's detection and ventilation would prevent or mitigate such exposure is difficult to claim with confidence. Arguments can certainly be made for and against this statement, and further investigation is warranted to confirm or reject the ventilation system as an effective IPL.

Conclusion

Now that we have thoroughly reviewed safeguard and their efficacy, we can wrap up the discussion on safeguards by returning to our example, as updated in Table 8. Now, we can certainly include the high discharge pressure cutout device, but we raise some important questions regarding independence for the compressor discharge temperature safety and the effectiveness of the machinery room ventilation. Clearly, what initially might have been a long list of possible safeguards to prevent a compressor overpressure event is reduced to only three (as the PSV is not sized for the full flow of the compressor). This illustrates the importance of a thorough analysis of safeguards in a PHA.

Table 8. Example HAZOP scenario with safeguards

Deviation	Cause	Consequence	Severity	Initiating Event Frequency	Risk Reduction Required	Safeguards

No/Low Flow	HSD-C01 manual valve inadvertently closed (compressor discharge isolation valve)	Loss of flow from compressor C1 resulting in an increase in pressure and temperature in the compressor discharge. Potential to exceed the design parameters of the compressor, resulting in compressor damage and a release of ammonia into the machinery room. Potential for two operator fatalities.	B Major	10^{-2}	III	High discharge pressure;
						high discharge temperature;
						machinery room ventilation

Now that the HAZOP portion of this example is complete, we turn our attention to LOPA, where we perform even further analysis of several pieces of our HAZOP PHA scenario.

LOPA

Selecting Scenarios for LOPA

Once a HAZOP study is complete, a PHA team may choose to further study scenarios that include potentially catastrophic outcomes, such as multiple fatality events or events with the possibility of severe consequences to members of the surrounding community. LOPA is an excellent tool that provides a thorough, semi-quantitative, order-of-magnitude risk assessment that facilities and companies can use to compare residual risk to internal risk tolerance. The first step is to define which scenarios warrant LOPA and consistently apply this definition. While LOPA can be applied to all HAZOP scenarios, this effort will likely be resource-prohibitive.

Assume that we wish to apply LOPA to any scenario that can result in one or more fatalities. From our risk matrix in Table 5, we can see that this means any *catastrophic* or *major* scenario requires completing LOPA. Of note, we must perform LOPA one scenario at a time. If an initiating event can result in multiple fatality consequences, each scenario must be considered individually, i.e., an initiating event that can result in a toxic exposure fatality OR a machinery room explosion fatality requires two separate LOPAs.

The LOPA method requires a much stricter adherence to many of the rules already discussed. In a HAZOP study, the PHA team has a wide degree of flexibility with many of these guidelines. However, LOPA takes what can be informal analyses from the HAZOP and formalizes them into a structured, semi-quantitative analysis. Once scenarios are selected, the LOPA process can begin.

Enabling Events and Conditional Modifiers

One key difference with a LOPA, compared to a HAZOP study, is after the scenario selection and the determination of the initiating event frequency (see the section on initiating event frequency), when we apply enabling events/conditions and conditional modifiers. According to *Guidelines for Enabling Conditions and Conditional Modifiers in LOPA*, enabling conditions are associated with the part of an incident sequence leading up to a release, whereas conditional modifiers are associated with post-release portions of an incident sequence.

An enabling event or condition must be present for the initiating event or consequence to proceed. For instance, a common enabling condition in an ammonia refrigeration system is the outside weather conditions. Most refrigeration systems are designed to operate in variable climates, such as hot, humid summer days and cold winter months. Many systems operate without condenser water in the winter because it is not needed. A failure of a condensing water system in January when air temperatures are near freezing does not result in a dangerous overpressure situation; however, the same cannot be said for a failure on a hot, humid summer day, where the same failure quickly leads to high system pressure with a potential for catastrophic consequences. In this example, certain outside weather conditions are required for the consequence to proceed. The weather, then, is an enabling condition. Only during certain atmospheric conditions will the loss of condensing water result in system overpressure.

Conditional modifiers, on the other hand, affect the probability of certain consequences of concern or consequence impacts. One classic example of a conditional modifier for an ammonia refrigeration system is the presence of personnel in an impacted area, especially a machinery room. During the HAZOP study, we assume that an operator is present in the machinery room during a hazardous event. However, what if this is not the case? Might the impact be much less severe? The answer is yes, and this type of conditional modifier may be applicable. If a facility can demonstrate that operators are in the room during specific times, the probability that they are present during an incident can be factored into the overall risk determination.

Finally, we advise that enabling conditions and conditional modifiers be used with caution. According to CCPS's *Guidelines for Enabling Conditions and Conditional Modifiers in LOPA*, LOPA analysts must not apply enabling conditions or conditional modifiers when the team has insufficient knowledge or data. They must also ensure adherence to any company internal requirements regarding the application of these conditions and modifiers. Ultimately, the team must be able to demonstrate that any modifier taken will have the attributed effect of reducing scenario risk (Center for Chemical Process Safety, 2014).

Consider a scenario where operators rush to the machinery room during a release or potential explosion event. Although not present at the time of the release, alarms may sound, alerting operators to respond. In this case, is it reasonable to take a conditional modifier that operators are only present for a portion of the time? Not when certain conditions make it more likely that an operator will respond. In this case, perhaps the assumption that an operator is always present is appropriate.

Determining PFD

An IPL is a formal safeguard, and it follows many of the requirements in the IDEA principle. All IPLs are safeguards, but not all safeguards are IPLs. To be considered an IPL for a LOPA, it must first pass a series of rigorous tests to help determine and quantify the effectiveness of an IPL. This IPL effectiveness is described as its probability of failure on demand (PFD), i.e., the probability that an IPL will fail to perform a specified function when required (Center for Chemical Process Safety, 2001).

The tests described for the IDEA principles (in the previous section on safeguards) must be applied to safeguards to elevate them to IPL status. During the HAZOP study, teams may apply some or all of the rules listed with some flexibility. During a LOPA, however, these factors are critical to ensuring a quality outcome. If, during the HAZOP study, the PHA team applies the IDEA rules, the effort of determining IPLs in LOPA will be quite easy, as much of the work has already been done.

Next, the LOPA team must assess the overall reliability of identified IPLs by assessing their PFDs. For most LOPA practitioners, these PFD values can be obtained from industry literature such as CCPS's *Layer of Protection Analysis*. Some commonly used IPLs in ammonia refrigeration systems and their PFDs are given in Table 9 (Center for Chemical Process Safety, 2001).

Table 9. Typical IPL PFD values

IPL	Comments	PFD from Literature	Recommended PFD
Relief valve	Sized to prevent system exceeding specified overpressure1; installed in clean service.	10^{-1}–10^{-5}	10^{-2}
Basic process control system (BPCS)	Can be credited as IPL if not associated with initiating event; same as the compressor's PLC or a system control PLC.	10^{-1}–10^{-2} ($> 10^{-1}$ not allowed by IEC)	10^{-1}
Safety instrumented functions (SIL1 or higher)	Consists of dedicated safety-rated input devices, logic controllers, and output devices; SIL level determined by a competent individual.	Varies based on requirements, design, and maintenance. $1x10^{-1}$ – $1x10^{-3}$	Varies
Human response to BPCS indication or alarm with 40 min response time	Simple, well-documented action with clear and reliable indications that the action is required.	10^{-1} ($> 10^{-1}$ allowed by IEC)	10^{-1}

1—Relief valves must be independent and sized to prevent exceeding specified pressure to qualify as an IPL.

Calculating Scenario Risk

After understanding our scenario and detailing our IPLs and their PFDs, the LOPA team can now assess the overall mitigated risk for the scenario in question. The team first estimates the frequency with which the scenario can occur, which can then be compared to a company's risk matrix.

To ascertain the frequency, we multiply the initiating event frequency by the product of the IPL PFDs and the probabilities of the applicable modifiers and enabling conditions. The formula below represents this in formulaic notation:

$$f_i^C = f_i^I \times \prod_{j=1}^{J} PFD_{ij} \times p^{m1} \times p^{m2} ... \times p^{mn}$$

where f_i^C is the frequency for consequence *C* for initiating event *i*; f_i^I is the initiating event frequency for initiating event *i*, is the PFD_{if} of the *j*th IPL that protects against consequence *C* for initiating event *i*; andp^{mi}, p^{m2}, …, p^{mn} are the probabilities associated with all applicable conditional modifiers and enabling conditions.

While this equation may look complicated, in simpler terms, we multiply each probability together to determine the frequency of the consequence of interest. If we multiply the initiating event frequency by the PFD of all IPLs and by the probabilities associated with our chosen modifiers and enablers, we arrive at a mitigated event frequency. We can apply this to our working example, where f_i^C is the frequency of an operator fatality resulting from valve closure, $f_i^I = 10^{-2}$ is our initiating event frequency of 1 in 100 per year, and $p^{m1} = 0.1$ is the probability that an operator is present and unable to escape:

$$\prod PFD = (PFD\ for\ pressure\ cutout) \times (PFD\ for\ ventilation) = 0.1 \times 0.1 = 0.01$$

$$f_i^C = 0.01 \times 0.01 \times 0.1 = 0.00001 = 10^{-5}$$

Thus, the mitigated event frequency for our consequence of interest, considering all applicable IPLs and modifiers, is 0.00001/yr. This frequency can now be compared to the facility's or company's risk tolerance. If a company requires single fatality events to be mitigated to a probability of 0.00001/yr, this consequence is properly mitigated. However, if it requires a lower fatality probability, then additional risk mitigation is required.

In this LOPA example, several safeguards are omitted (one that was listed above in the HAZOP example), and one additional modifier was included. The compressor's PRVs are not credited as an IPL because they are not sized for the full flow of the compressor without the compressor's control system functioning. Additionally, the compressor's high-temperature cutout device is not credited because it is not independent of another credited safeguard, the compressor's high-pressure cutout. Furthermore, a modifier for operator presence is included to illustrate how these modifiers can work to reduce risk. This also demonstrates the importance and consideration needed when deciding whether to apply conditional modifiers and enabling conditions. Note that, in LOPA, these modifiers reduce risk by the same value as critical safety systems. The LOPA team must have a similar level of confidence that any modifier applied is valid and that adequate evidence and data exist to support this.

Benefits and Challenges of HAZOP Studies and LOPA

Moving from a what-if/checklist PHA to a more robust HAZOP/LOPA method comes with many benefits and challenges. First, we detail several challenges encountered in our journey to implement HAZOP and LOPA PHAs. We also provide some potential resolutions to eliminate or alleviate these problems. Then, a discussion of the benefits of applying HAZOP and LOPA methodology follows, along with how they can lead to a more effective and thorough reduction of process safety risk.

Challenges

Safeguard Independence and Design Dependability

The primary challenge to effectively applying HAZOP and LOPA techniques to refrigeration systems, in our opinion, is the lack of independence and dependability of safeguards and IPLs, bringing about many conflicts. This is a common concern with ammonia refrigeration safety systems and an area where improvement is warranted, particularly when it comes to safeguard design from a reliability perspective.

Industry P&ID Maturity

The maturity and detail included in ammonia refrigeration piping and instrumentation diagrams (P&IDs) also lead to challenges when completing HAZOP and LOPA studies. According to IIAR's *Refrigeration Piping Handbook, Appendix A: Guidelines for Preparation of Ammonia Refrigeration Diagrams*, P&IDs should *exclude* the following (International Institute of All-Natural Refrigeration, IIAR, 2019):

- Secondary heat transfer fluid piping, such as condenser water piping and brine loops
- Non-ammonia support equipment, such as condenser pumps and cooling water pumps
- Control loops (pneumatic and electrical lines) for the sake of simplicity
- Factory prefabricated equipment package details

Unfortunately, this list includes important information for completing an effective HAZOP PHA and LOPA. Condenser water piping, tanks, pumps, and other associated equipment can fail, resulting in consequences of interest. Since this equipment is not shown on a P&ID, the PHA team may miss multiple causes and consequences. The lack of instrumentation and control loops often leads to confusion regarding what safeguards are applicable and their effectiveness, given a certain cause. The lack of compressor detail on P&IDs (i.e., oil flows, pumps, filtration, safety device location) can result in missing or improperly assigning safeguards or IPLs or mistakenly crediting a device that may not be effective for a given scenario.

In our opinion, IIAR should consider a re-evaluation of P&ID guidelines or possibly the creation/adoption of a much more thorough, robust list of requirements for P&IDs that includes much of the information currently excluded. Organizations outside of the ammonia refrigeration industry have already developed robust P&ID standards and requirements. Process Industry Practices (PIP) is a self-funded consortium dedicated to harmonizing standards and practices for design, procurement, construction, and maintenance, and it has developed detailed guidelines for P&IDs. (Process Industry Practices (PIP), 2024) The PIP document *PIC001, Piping and Instrumentation Diagram Documentation Criteria* details the requirements of thorough and modern P&IDs, which include many elements omitted by IIAR. These guidelines are adopted broadly in other processing industries and can easily be applied and tailored to fit the needs of the refrigeration industry.

Repetitiveness

Ammonia refrigeration systems contain many pieces of similar, or in some cases, identical equipment. Assessing each piece of equipment can become repetitive. However, during the initial noding discussion and throughout the HAZOP study, teams can easily identify these similarities and group equipment together. For example, a typical system that contains 75 individual pieces of equipment can be broken down into 10–15 groups of identical pieces of equipment, each having a node, and assessed in groups rather than individually. Other synergies exist as well that lead to efficiency during the HAZOP and LOPA studies. Most recirculating vessels are designed similarly, and once one is studied, the same process can be applied to others with a few adjustments.

Applicability to Procedural or Human-based Tasks

We do not advise applying the LOPA methods described above to task-oriented activities or those relying heavily on human interaction. Because of their reliance on a single procedure or a few people performing a given task, many of the rules for safeguard independence fail when applied to procedural failures. Other techniques are available for assessing procedures and maintenance tasks; however, these are outside the scope of this paper. In these cases, a specific checklist is more appropriate. Tasks where checklists can be applied include the following:

- Ammonia unloading
- Ammonia pump-outs
- Ammonia line breaks
- Oil draining
- Compliance with internal or external minimum design standards

Resource Requirements

Completing a more thorough risk assessment inherently requires more resources. According to *Guidelines for Hazard Evaluation Procedures*, a HAZOP study for a small system can take 2 to 6 days, whereas a what-if/checklist can be completed in 1 to 2 days. This means longer time away from the facility for operators and mechanics in the PHA and additional time requirements for contract facilitators and facility engineering resources. In our experience, what-if/checklist PHAs are typically completed in 2.5 days on average, while a HAZOP PHA—complete with LOPA where appropriate as well as an additional checklist to review human factors, facility siting, previous recommendations, high-risk tasks, etc.—takes 5 days on average. This duration varies significantly depending on the size and complexity of the system being evaluated, with a small, single-stage system being completed in as few as 3 days and larger two-stage or more complex systems spread across multiple machinery rooms taking up to 10 days.

Benefits

Thoroughness

When completing a HAZOP PHA, we can clearly see that the assessment is very thorough. The method lends itself to the complete assessment of each piece of equipment in a very methodical way, and when applied correctly by an experienced facilitator and team, it is extremely effective at identifying risk in a process. Additionally, the HAZOP method presented in this paper, in addition to LOPA, provides a deep assessment of the safeguards on which we rely and allows refrigeration teams and risk practitioners to assess and dissect the efficacy of these crucial safety systems. A significant amount of time has been dedicated to discussing safeguards and IPLs in this paper, and we refer the reader to those sections for full analysis. However, applying the HAZOP and LOPA method leads to the identification of numerous opportunities for improvement of safety systems, which, ultimately, reduce overall risk. It offers teams a much clearer picture of risk and provides a much better understanding of whether a process is safe enough.

Reduced Reliance on Team Experience

As we continue to see and struggle with labor shortages in the ammonia refrigeration industry and face an exodus of experience, we clearly need to continue to effectively assess the process safety risk of our systems. Furthermore, over the next 10–20 years, the industry will continue to lose knowledge and experience. The HAZOP study was originally developed to anticipate hazards and operability problems for technology with which organizations have little experience (Center for Chemical Process Safety, 2008). Although the ammonia refrigeration industry is mature and well understood and we have thorough RAGAGEPs available, the experience available to facilities and companies when it comes to operating and maintaining these systems continues to decline.

The HAZOP method is a perfect fit, especially because of its value to less experienced teams. According to *Guidelines for Hazard Evaluation Procedures*, the application of what-if analyses alone is "a powerful technique if the staff is experienced; otherwise, the results are likely to be incomplete." Additionally, according to IIAR's *PSM Guidelines for Ammonia Refrigeration*, what-if/checklists "are very dependent on the experience and thoroughness of the team leader and the team" (International Institue of All-Natural Refrigeration, IIAR, 1998). These experience requirements, along with the continuing trend toward inexperience in the ammonia refrigeration industry, should lead all of us as risk professionals to seriously consider this important benefit of the HAZOP method.

Risk-based Prioritization and Resolution

The semi-quantitative nature of this HAZOP method and LOPA provides an opportunity for process safety professionals in ammonia refrigeration facilities to understand and prioritize risk resolution better. In a what-if or what-if/checklist PHA, depending on how they are applied, the level of residual risk to be mitigated by a given recommendation may be unclear. Safeguard PFDs are not quantified, initiating event frequencies may not be discussed, and other quantifiable information is not a typical output of these types of studies. HAZOP and LOPA PHA outputs contain all of this important risk information. Some recommendations may mitigate a single IPL gap, while others are intended to mitigate multiple gaps. Another recommendation may address a single low-risk scenario when others may tackle tens of high-risk scenarios.

This information regarding recommendations can be very helpful for facilities when prioritizing work to resolve recommendations, and it is available when HAZOP and LOPA are employed. In a world and industry where every hour and dollar is scrutinized, having concrete process risk data can help demonstrate the importance of making the improvements needed to operate and maintain a safe refrigeration system.

Conclusions and Recommendations

According to the EPA, ammonia is the most frequently released hazardous chemical from facilities regulated under Section 112(r). According to its FY 2024–2027 National Enforcement and Compliance Initiatives memorandum, the EPA intends to continue focusing on chemical accident risk reduction as an initiative with an increased focus on inspecting and addressing noncompliance at facilities using anhydrous ammonia. In the coming years, EPA also aims to "use all available enforcement tools to address violations of risk management requirements, including holding entities criminally responsible." Considering this increased regulatory scrutiny, the case for improved management of risk cannot be clearer, and one key tool we have for this is PHA.

To facilitate this change, to better understand our process safety risk, and to continue pushing the boundaries toward safer refrigeration systems, facilities, companies, and the industry at large should embrace a change for better PHAs. Implementing HAZOP and LOPA methodology more broadly is one way to accelerate this effort. In our experience, a well-thought-out and planned HAZOP PHA with LOPA applied to the most catastrophic scenarios can benefit facility risk managers and improve process safety. It allows facility teams to assess risk on a level that what-if/checklist PHAs cannot. It provides methods that are more useful to inexperienced operators and technicians. It demands a higher standard of process safety information. It holds critical safeguards to a level of efficacy that current methods do not. Moreover, while it does require more resources and effort to implement, it allows us to identify, prioritize, and mitigate or eliminate risk in a way that is more thorough and robust.

We owe it to our teams, our businesses, our owners, our families, and the communities in which we operate to ensure that our systems are safe. That our team members can go home at the end of each day in the same shape in which they arrived to work. That our neighbors do not have to worry about the possibility of a chemical release affecting their lives or, worse, injuring them. That regulators no longer feel the need to apply additional scrutiny to our industry. We owe it to ourselves, as professionals in risk reduction and mitigation, to do everything we can to protect our systems and the stakeholders that rely on them. When it comes to risk identification and mitigation, we should not settle on "good enough." We should push for improvement, continued risk reduction, and safer refrigeration systems. HAZOP and LOPA can be key tools in our effort to continually make our systems safer.

References

Center for Chemical Process Safety. (2001). *Layer of Protection Analysis - Simplified Process Risk Assessment*. New York: Center for Chemical Process Safety/AIChE (CCPS).

Center for Chemical Process Safety. (2007). *Guidelines for Risk Based Process Safety*. New York: Center for Chemical Process Safety/AIChE (CCPS).

Center for Chemical Process Safety. (2008). *Guidelines for Hazard Evaluation Procedures (3rd Edition)*. New York: Center for Chemical Process Safety/AIChE (CCPS).

Center for Chemical Process Safety. (2014). *Guidelines for Enabling Conditions and Conditional Modifiers in Layer of Protection Analysis*. New York: Center for Chemical Process Safety/AIChE (CCPS).

Center for Chemical Process Safety. (2014). *Guidelines for Initiating Events and Independent Protection Layers in Layer of Protection Analysis*. New York: Centter for Chemical Process Safety/AIChE (CCPS).

Center for Chemical Process Safety. (2024, December 26). Retrieved from AIChE.org: https://www.aiche.org/ccps/resources/glossary/process-safety-glossary/risk#:~:text=A%20measure%20of%20human%20injury,Risk%20%3D%20Frequency%20x%20Consequence).

Environmental Protection Agency, EPA. (2023, August 17). *EPA.gov*. Retrieved from https://www.epa.gov/system/files/documents/2023-08/fy2024-27necis.pdf

IEC. (2010, April). IEC 61508-1 Edition 2.0.

International Institue of All-Natural Refrigeration, IIAR. (1998). *Process Safety Management Guidelines for Ammonia Refrigeration*. Alexandria, VA: IIAR.

International Institute of All-Natural Refrigeration, IIAR. (2019). *The Refrigeration Piping Handbook*. Retrieved from https://evantage.gilmoreglobal.com/reader/books/01BOO-EN0202

International Institute of All-Natural Refrigeration, IIAR. (2024, December 26). *IIAR.org.* Retrieved from https://www.iiar.org/IIAR/iiar/about_ammonia_refrigeration/the_history_of_ammonia_refrgeration.aspx

Jordan, P. (2014). Layers of Protection in an Ammonia Refrigeration System. *IIAR 2014 Industrial Refrigeration Conference & Heavy Equipment Show.* Nashville: International Institute of Ammonia Refrigeration.

Occupational Safety and Health Administration. (1992, February 24). *OSHA.gov.* Retrieved from eCFR.gov: https://www.ecfr.gov/current/title-29/subtitle-B/chapter-XVII/part-1910/subpart-H/section-1910.119

Process Industry Practices (PIP). (2024, December 26). *PIP.org.* Retrieved from PIP.org: PIP.org

Weber, R. (2024). Process Risk Assessment Facilitator Training.

Acknowledgments

This paper could not have been written without the encouragement of my former colleague and current mentor, Rhonda Schmidt. In our short time working together, Rhonda impacted both Cargill and me personally in many ways, and she was a key driver and change champion for the adoption of HAZOP and LOPA PHAs in Cargill's Protein Business. She was also the one who initially planted the idea for a paper and presentation at IIAR. Finally, we must give credit where it is due for the clever name of our paper. The title *Livin' la Vida LOPA* was her very creative and fun idea.

Notes:

Notes: