SAVE LIVES GLOBAL RESEARCH

Research Paper No. 003

Beyond Reliability – the Hidden Safety Risks of Liquid-Cooled AI Data Centers

A Call to Reevaluate SIF Prevention, Human Performance, and Operational Readiness

Simon Goncharenko

PhD, CSP, CSHO

July 2026

Save Lives Global Research

Research Paper No. 003 – SLG-RP-003

Version 1.0

First Edition

Publication Date: July 2026

Author - Simon Goncharenko, PhD, CSP, CSHO

Suggested citation

Goncharenko, S. (2026). *Beyond Reliability: The Hidden Safety Risks of Liquid-Cooled AI Data Centers*. Save Lives Global.

Website

SaveLivesGlobal.com

Research Series

Save Lives Global Research publishes practitioner-focused papers advancing human performance, Serious Incidents and Fatalities (SIF) prevention, contractor management, organizational resilience, and emerging safety challenges

Disclaimer

This white paper is intended to stimulate discussion and contribute to the advancement of safety, human performance, and critical risk management within the data center industry. The views expressed are those of the author and do not necessarily reflect the views of any current or former employer, client, professional organization, standards body, or committee with which the author is affiliated.

This publication is provided for informational and educational purposes only and should not be interpreted as engineering, legal, regulatory, or professional consulting advice. Organizations should evaluate the concepts presented in light of their own operational environments, applicable laws, regulations, standards, and technical requirements.

Executive Summary

The data center industry is undergoing one of the most significant technological transformations in its history. Driven by unprecedented demand for artificial intelligence (AI) computing, organizations are rapidly deploying liquid cooling technologies to support rack densities that would have been unimaginable just a few years ago.

In the process, the discussion surrounding liquid cooling has largely focused on thermal efficiency, sustainability, reliability, and uptime. These are important considerations. However, they are not the only considerations.

As an industry, we may be overlooking a critical question - have we adequately considered the safety implications of introducing liquid cooling into environments historically dominated by electrical and information technology risks?

Throughout history, new technologies have introduced new hazards. More often than not, organizations focus first on performance and later on safety. The result is frequently a reactive cycle where lessons are learned only after significant incidents occur. History provides numerous examples, including Piper Alpha oil rig explosion, Challenger Space Shuttle disaster, and Chernobyl Nuclear Reactor release.

The purpose of this paper is not to suggest that liquid cooling is unsafe. Nor is it intended to criticize the remarkable innovation occurring throughout the data center sector. Rather, it is intended to serve as a wake-up call and an encouragement for further dialogue.

The industry is moving rapidly toward widespread adoption of liquid-cooled infrastructure. Before that transition becomes fully normalized, safety professionals, operational leaders, engineers, and contractors should pause and ask a series of important questions:

- What new Serious Incident and Fatality (SIF) pathways are being introduced?
- How do liquid cooling systems alter the risk profile of the data hall?
- Are our emergency response capabilities prepared for these new failure modes?
- Have we adequately addressed the human and organizational factors associated with these systems?
- Are we asking the right questions before a significant event forces us to?

The central argument of this paper is straightforward – the data center industry has invested heavily in understanding the engineering and operational implications of liquid cooling. Comparatively less attention has been devoted to understanding the human, organizational, and SIF implications of these systems (partially illustrated by Figure 1).

The key premise of this paper is that while engineering controls remain essential, the industry's ability to prevent serious incidents and fatalities in liquid-cooled AI data centers will increasingly depend on human performance, organizational resilience, and operational readiness.

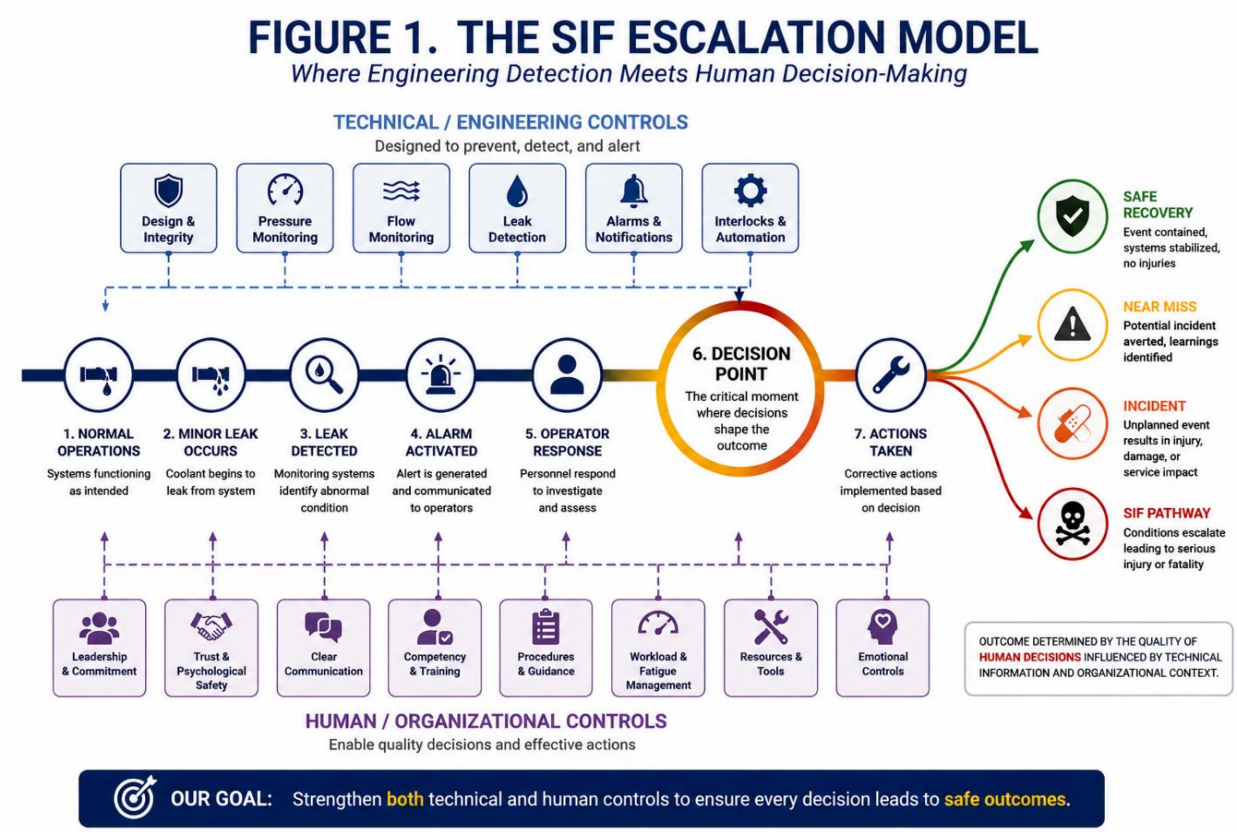

1. Introduction: The Pattern We Keep Repeating

Many industries have experienced periods of rapid technological advancement where innovation outpaced risk understanding. There are numerous examples of this from history, including but not limited to :

- Piper Alpha
- Challenger
- Columbia
- Deepwater Horizon
- Chernobyl
- Fukushima

In each case, organizations understood the technology far better than they understood the interaction between technology, people, and organizational systems. The lesson is not that innovation is dangerous. The lesson is that emerging hazards are often poorly understood until a major event exposes them.

The data center industry now finds itself at a similar inflection point.

It is important to distinguish between hazards and pathways. Hazards describe sources of potential harm. SIF pathways describe the sequence of technical, organizational, and human events that can transform a hazard into a catastrophic outcome. This paper focuses less on identifying individual hazards and more on understanding how those hazards interact with engineering controls, human performance, and organizational decision-making

2. The AI Revolution Is Reshaping Data Center Infrastructure

For decades, data center safety programs were built around a relatively predictable set of hazards:

- Electrical systems.
- Battery systems.
- Mechanical equipment.
- Material handling.
- Construction activities.
- Maintenance operations.

While these hazards remain, AI infrastructure is changing the nature of the environment itself. The exponential growth of AI, including continually growing GPU clusters, mind-blowing increases in rack densities, and unprecedented power demands have brought the industry to a place where traditional air-cooling strategies are reaching their practical limits.

As a result, liquid cooling is rapidly becoming a foundational element of modern data center design. Global bodies, like the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) are going as far as formally recommending immediate consideration of liquid cooling technologies in all future data center designs and retrofits, as it can provide a way out of the current power availability challenges (Emergence and Expansion).

Within what is possible with liquid cooling, there are numerous approaches, including direct-to-chip, rear door heat exchangers, immersion cooling, and a variety of hybrid approaches. The key impact of this creative engineering is that the final output is not merely an adaptation of a new cooling technology, but creation of an entirely new operating environment.

In other words, the implications extend far beyond thermal management. New cooling technologies introduce new equipment, new maintenance requirements, new competencies, and new operational challenges. As such, our questions should be updated as well. Perhaps, rather than asking whether liquid cooling will become commonplace, we should really be focusing on whether our safety management systems are evolving at the same pace as the technology they are intended to protect.

3. Reliability Has Dominated the Conversation

Most public discussion surrounding liquid cooling focuses on a common set of themes:

- Thermal efficiency
- Water conservation
- System reliability
- Hardware protection
- Energy consumption
- Scalability

These are important topics. Yet none of them concern the safety of people implementing the new technologies.

Whether it is cooling research from Microsoft, published by Catherine Bolgar in 2025, Nvidia's cooling announcements introduced in June of 2026 for its AI-focused Rubin GPU chips, or Google/Microsoft sustainability discussions, even a cursory review of public literature reveals extensive discussion about preventing coolant leaks, maintaining uptime, and protecting equipment. Far less attention has been given to:

- Worker exposure
- Human performance
- Emergency response
- Contractor competency
- Critical risk management
- Serious Incident and Fatality prevention in this space

This gap should concern us. Not because incidents are occurring at scale today. But because history suggests they eventually will. Figure 2 visualizes the above.

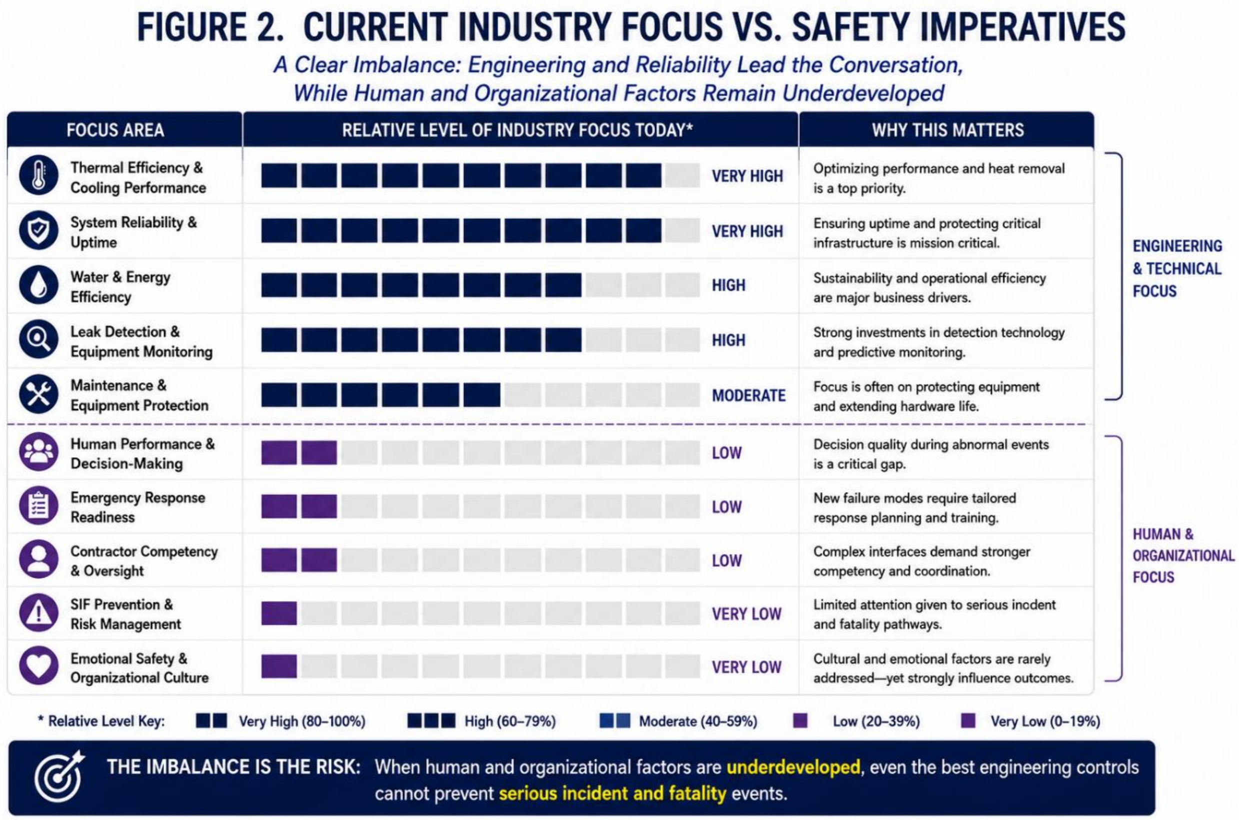

4. Understanding the New Hazard Landscape

Liquid cooling fundamentally changes the hazard landscape within AI data centers by introducing new equipment, new interfaces, and new interactions between existing hazards. While many individual hazards are not new, their proximity, complexity, and potential interactions create new pathways through which Serious Incidents and Fatalities (SIF) may occur (Figure 3).

For example, the increased interaction between energized electrical systems and circulating coolant requires careful engineering and operational controls. The proximity of energized busways, rack PDUs, switchgear, coolant distribution units (CDUs), manifolds, hoses, and quick-disconnects introduces combinations of electrical, hydraulic, mechanical, and human factors that have historically been far less common within traditional air-cooled data centers.

Figure 3 illustrates representative hazard categories associated with liquid cooling, including electrical interactions, pressurized systems, coolant chemistry, maintenance activities, contractor interfaces, commissioning, human factors, and emergency response. Individually, many of these hazards are familiar. Collectively, however, they create a more interconnected operating environment that requires new competencies, revised procedures, and a broader approach to risk management.

Emergency response planning also becomes more complex. Organizations should evaluate shutdown authority, isolation sequences, coordination between operations and maintenance personnel, spill containment, electrical boundaries, communication protocols, and emergency preparedness before abnormal conditions occur. As liquid cooling technologies continue to evolve, so too must the systems designed to respond when those technologies fail.


FIGURE 3. THE LIQUID COOLING HAZARD LANDSCAPE
An Interconnected Web of Hazards Introduced by Liquid Cooling Systems
ELECTRICAL
• Electrical shock
• Arc flash
• Short circuit
• Ground faults
• Equipment damage
COMMISSIONING
• System integrity
• Testing & validation
• Start-up hazards
• Documentation gaps
• Turnover readiness
PRESSURE
• Pressurized systems
• Hose/blowout failure
• High-energy release
• Fluid injection injury
• Stored hydraulic energy
CONTRACTOR
• Multiple trades
• Interface risks
• Varying competencies
• Work coordination
• Oversight & verification
LIQUID COOLING
New Technology.
New Interactions.
New Risks.
CHEMISTRY
• Coolant toxicity
• Spod exposure
• Corrosion & inhibitors
• Contamination
• Biological growth
EMERGENCY
• Leak response
• Evacuation challenges
• Fluid containment
• Fire interactions
• Coordination & roles
HUMAN FACTORS
• Decision-making
• Situational awareness
• Training & competency
• Workload & stress
• Communication
MAINTENANCE
• Complex systems
• Confined access
• Hot surfaces
• Heavy components
• System isolation
THESE HAZARDS DO NOT EXIST IN ISOLATION.
They interact in complex ways across the system lifecycle.
Understanding their connections is essential to preventing Serious Incidents and Fatalities (SIF).

5. What Has Actually Changed?

It is reasonable to ask whether liquid cooling truly introduces new hazards. After all, data centers have long operated chilled-water systems, cooling towers, and hydronic loops. The distinction is not the presence of water itself. Rather, it is the migration of circulating coolant from building infrastructure into the white space and, in many cases, directly to the compute equipment.

Historically, facility water systems were largely separated from energized IT equipment and routine technician interaction. Modern direct-to-chip and immersion cooling fundamentally change that relationship by introducing pressurized cooling circuits, quick-disconnects, manifolds, coolant distribution units (CDUs), and service activities immediately adjacent to energized equipment. This changes maintenance activities, operator interactions, alarm response expectations, contractor competencies, emergency procedures, and the potential pathways through which abnormal events develop (see Table below).

Traditional Facility Cooling	**Direct Liquid Cooling**
Water primarily in plant infrastructure	Fluid inside white space and racks
Facility technicians	IT + mechanical + electrical technicians
Equipment cooling is indirect	Cooling becomes part of compute architecture
Limited routine operator interaction	Frequent operator intervention and maintenance
Facility reliability concern	Human performance and SIF concern

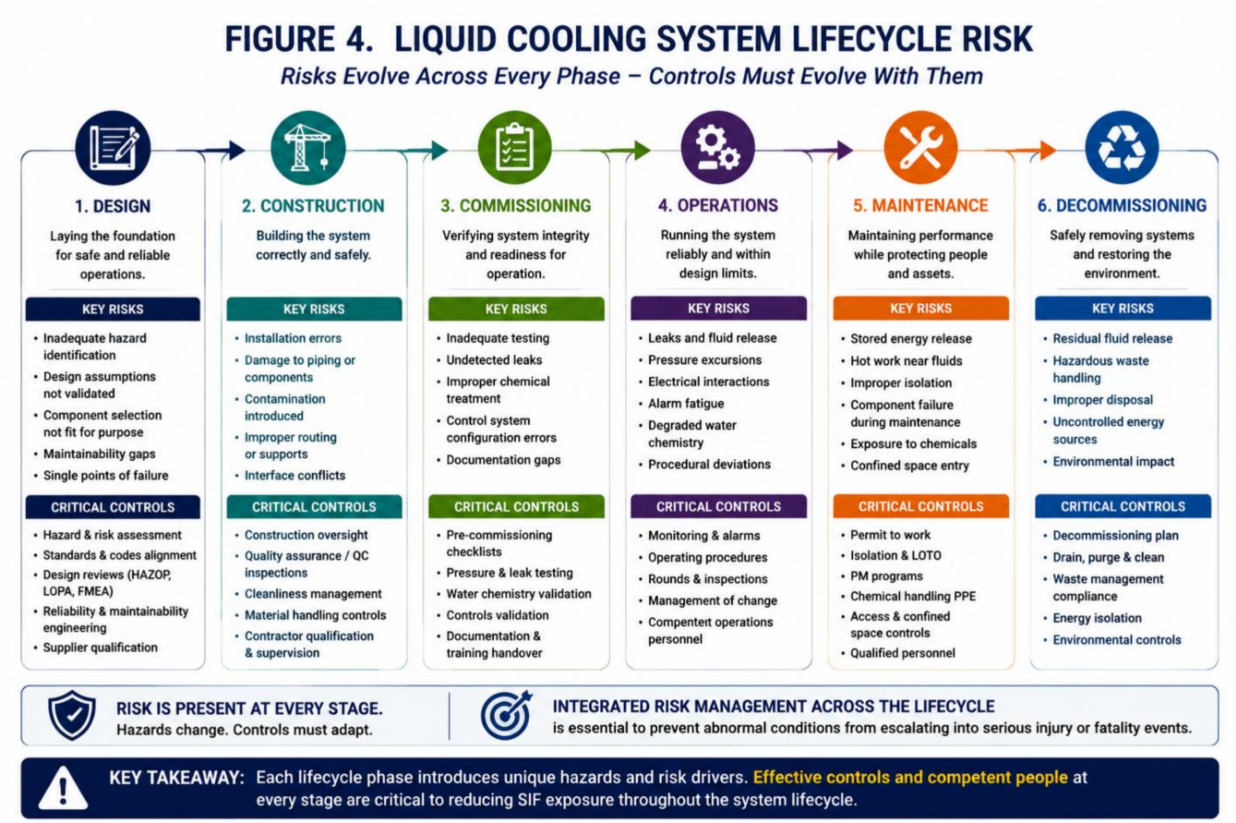

FIGURE 4. LIQUID COOLING SYSTEM LIFECYCLE RISK

Risks Evolve Across Every Phase – Controls Must Evolve With Them

6. Emerging SIF Pathways and Operator-Dependent Safeguards

Rather than viewing these examples as isolated hazards, they should be viewed as representative pathways through which technical failures, organizational pressures, engineering safeguards, and human decisions interact to create the potential for Serious Incidents and Fatalities (SIF). Similar pathways have been observed across numerous high-risk industries and are rarely the result of a single failure. Rather, they emerge from the interaction of multiple technical, organizational, and human factors.

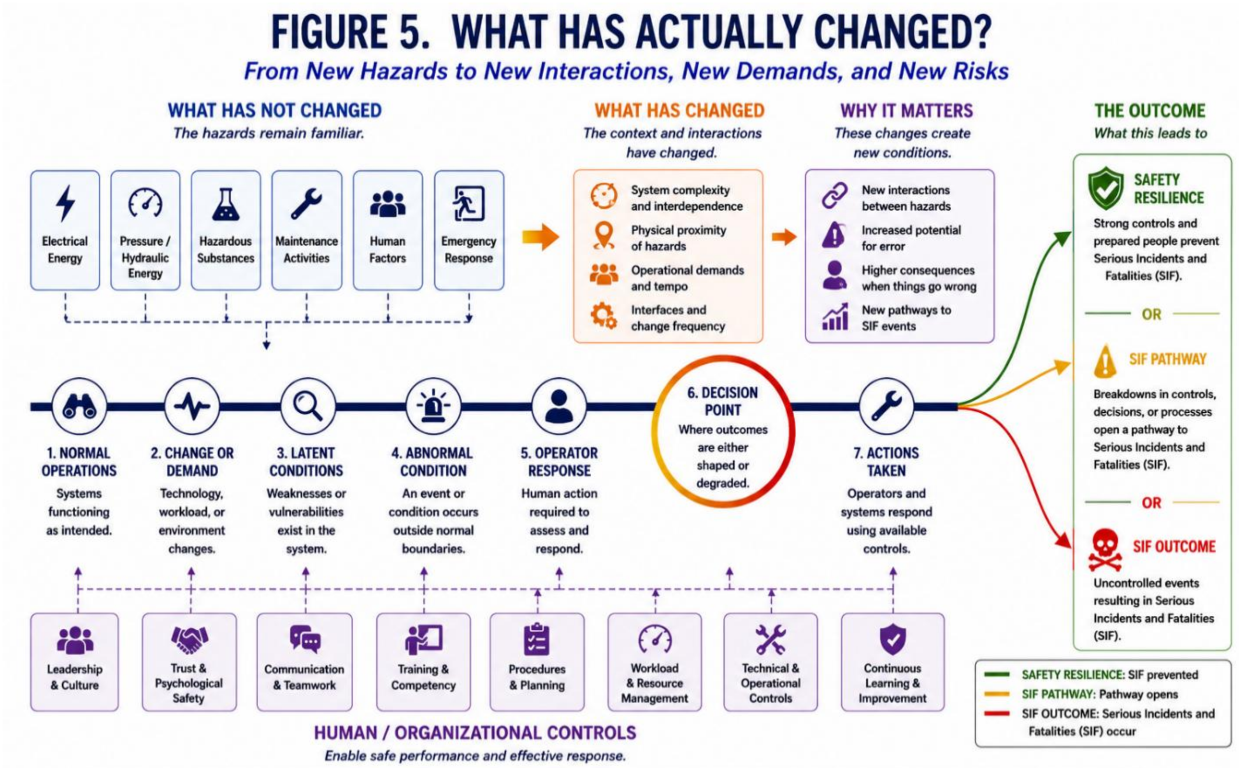

FIGURE 5. WHAT HAS ACTUALLY CHANGED?
From New Hazards to New Interactions, New Demands, and New Risks

One of the challenges associated with new technologies is that organizations often focus on direct hazards while overlooking the pathways through which serious events actually occur. In the case of liquid cooling, one representative SIF pathway is illustrated below:

- Coolant leak
- Alarm
- Production pressure
- Technician bypasses procedure
- Electrical exposure
- Potential SIF

Of course, there is no limit to the varieties of pathways for a serious incident or fatality in the context of liquid cooling, including the following example:

- Pressurized coupling fails
- Unexpected hose whip
- Worker struck
- Fall
- SIF

Or, this:

- Cooling failure
- Compute cluster at risk
- Pressure to restore uptime
- Improper isolation
- Unexpected energy release

These are not theoretical concerns. They are foreseeable conditions that deserve thoughtful evaluation before they result in significant consequences. However, identifying a hazardous event alone is insufficient. Organizations must also evaluate whether the safeguards intended to interrupt that pathway are capable of doing so under realistic operating conditions.

Figure 6 introduces an **Operator-Dependent Safeguard Evaluation Framework** that extends traditional hazard analysis by examining the effectiveness of operator response as a credited safeguard. When engineering controls cannot fully prevent or mitigate a hazardous event, organizations frequently rely on alarms and operator intervention. In lity of

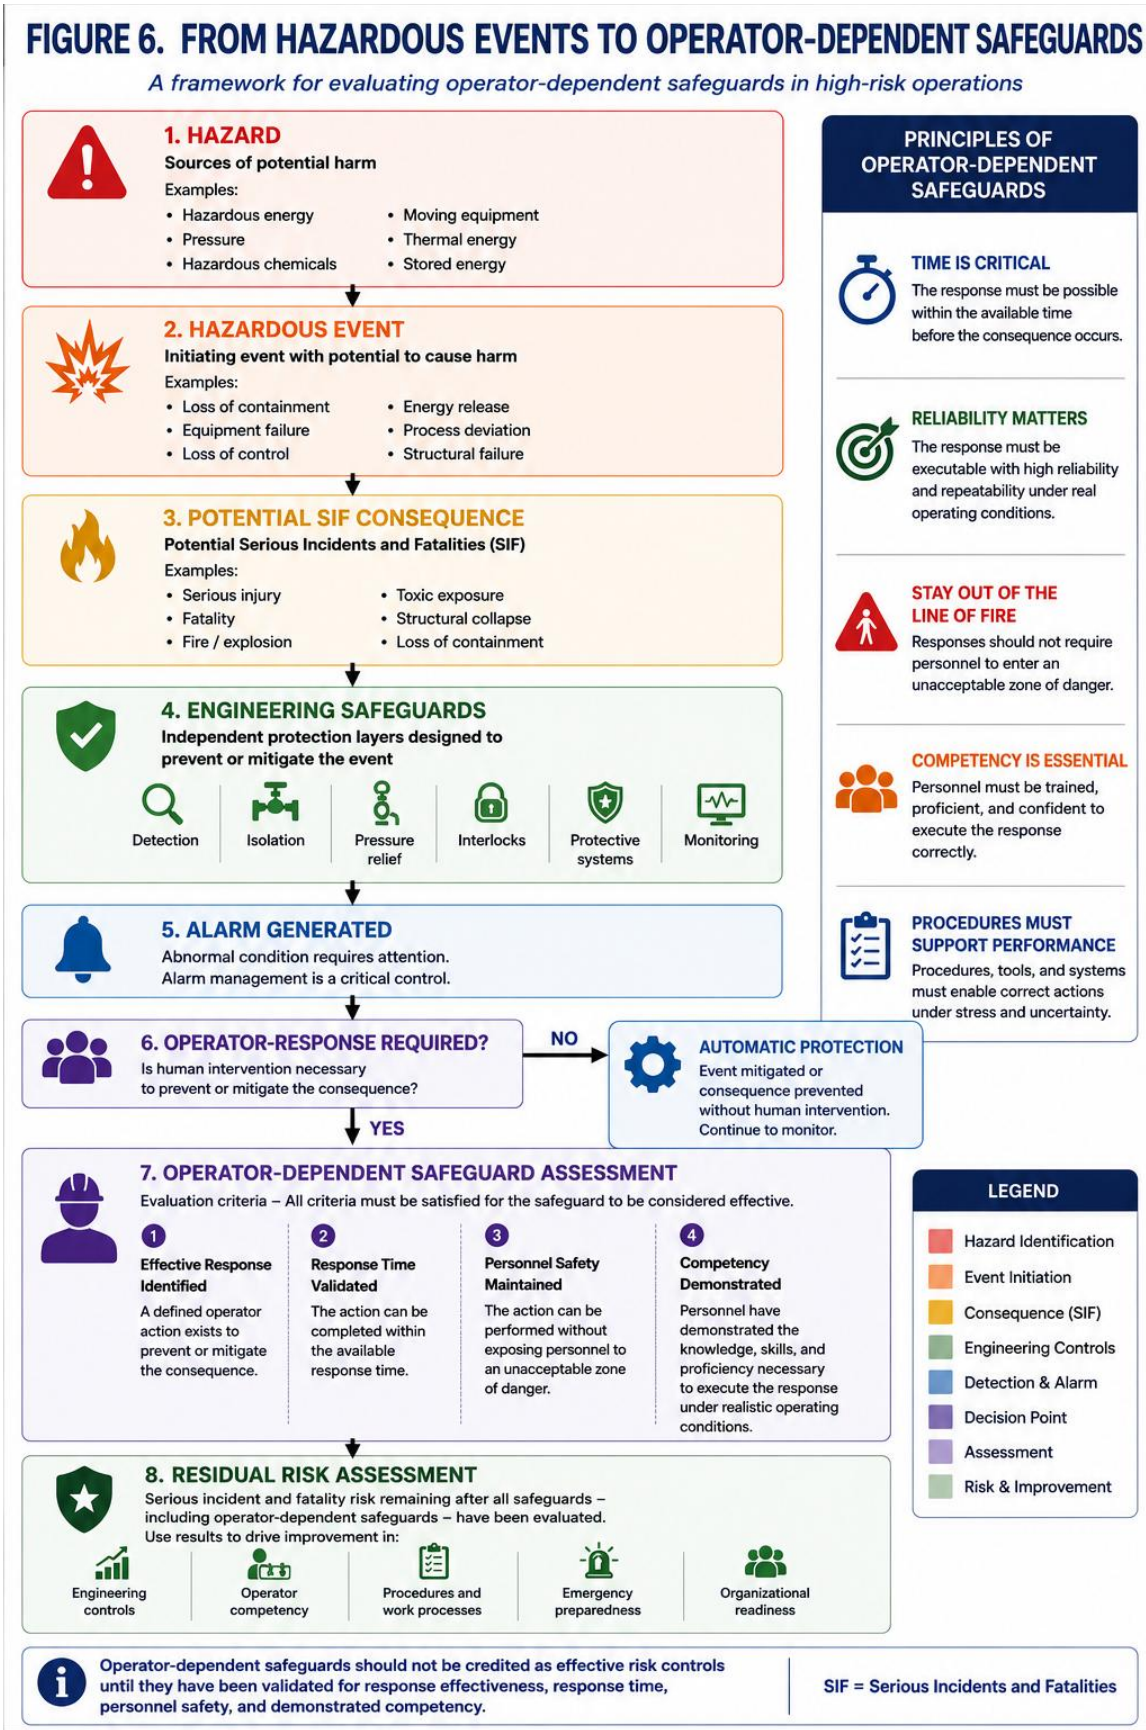
FIGURE 6. FROM HAZARDOUS EVENTS TO OPERATOR-DEPENDENT SAFEGUARDS
A framework for evaluating operator-dependent safeguards in high-risk operations
1. HAZARD
Sources of potential harm
Examples:
• Hazardous energy
• Pressure
• Hazardous chemicals
• Moving equipment
• Thermal energy
• Stored energy
2. HAZARDOUS EVENT
Initiating event with potential to cause harm
Examples:
• Loss of containment
• Equipment failure
• Loss of control
• Energy release
• Process deviation
• Structural failure
3. POTENTIAL SIF CONSEQUENCE
Potential Serious Incidents and Fatalities (SIF)
Examples:
• Serious injury
• Fatality
• Fire / explosion
• Toxic exposure
• Structural collapse
• Loss of containment
4. ENGINEERING SAFEGUARDS
Independent protection layers designed to prevent or mitigate the event
Detection
Isolation
Pressure relief
Interlocks
Protective systems
Monitoring
5. ALARM GENERATED
Abnormal condition requires attention.
Alarm management is a critical control.
6. OPERATOR-RESPONSE REQUIRED?
Is human intervention necessary to prevent or mitigate the consequence?
NO
AUTOMATIC PROTECTION
Event mitigated or consequence prevented without human intervention.
Continue to monitor.
YES
7. OPERATOR-DEPENDENT SAFEGUARD ASSESSMENT
Evaluation criteria – All criteria must be satisfied for the safeguard to be considered effective.
1 Effective Response Identified
A defined operator action exists to prevent or mitigate the consequence.
2 Response Time Validated
The action can be completed within the available response time.
3 Personnel Safety Maintained
The action can be performed without exposing personnel to an unacceptable zone of danger.
4 Competency Demonstrated
Personnel have demonstrated the knowledge, skills, and proficiency necessary to execute the response under realistic operating conditions.
8. RESIDUAL RISK ASSESSMENT
Serious incident and fatality risk remaining after all safeguards – including operator-dependent safeguards – have been evaluated.
Use results to drive improvement in:
Engineering controls
Operator competency
Procedures and work processes
Emergency preparedness
Organizational readiness
PRINCIPLES OF OPERATOR-DEPENDENT SAFEGUARDS
TIME IS CRITICAL
The response must be possible within the available time before the consequence occurs.
RELIABILITY MATTERS
The response must be executable with high reliability and repeatability under real operating conditions.
STAY OUT OF THE LINE OF FIRE
Responses should not require personnel to enter an unacceptable zone of danger.
COMPETENCY IS ESSENTIAL
Personnel must be trained, proficient, and confident to execute the response correctly.
PROCEDURES MUST SUPPORT PERFORMANCE
Procedures, tools, and systems must enable correct actions under stress and uncertainty.
LEGEND
Hazard Identification
Event Initiation
Consequence (SIF)
Engineering Controls
Detection & Alarm
Decision Point
Assessment
Risk & Improvement
Operator-dependent safeguards should not be credited as effective risk controls until they have been validated for response effectiveness, response time, personnel safety, and demonstrated competency.
SIF = Serious Incidents and Fatalities

procedures, but also on whether personnel can execute the required actions reliably, within the available response time, and without exposing themselves to unacceptable risk.

The framework presented in Figure 6 proposes four evaluation criteria for operator-dependent safeguards. First, there must be an effective operator response capable of preventing or mitigating the consequence. Second, the response must be executable within the available response time. Third, the action must not require personnel to enter an unacceptable zone of danger. Finally, competency must be demonstrated through realistic training and operational preparedness. If these criteria cannot be satisfied, organizations should reconsider whether operator intervention should be credited as an effective risk control or whether additional engineering safeguards, automation, or procedural changes are warranted.

In summary, risks evolve across every phase of liquid cooling implementation—from design and construction through commissioning, operations, maintenance, and decommissioning. As those risks evolve, so too must the methods used to evaluate both engineering and operator-dependent safeguards. Effective SIF prevention depends not only on identifying hazards, but also on validating that the controls intended to interrupt those hazards will perform as expected under realistic operating conditions.

7. The Human Performance Question

Figure 5 illustrates an important but often overlooked reality. Engineering systems can identify abnormal conditions. But what they cannot do is determine how people respond to them.

Technology does not fail in isolation. People interact with technology. They install it, maintain it, troubleshoot it, and respond when it does not behave as expected. While engineering controls are designed to detect abnormal conditions and protect equipment, it is ultimately people who determine how those conditions are interpreted and managed (Goncharenko, 2024).

As organizations introduce increasingly sophisticated liquid cooling systems, the human element becomes even more significant. New technologies inevitably require new competencies, new mental models, and new decision-making processes. Technicians who have spent years working in traditional air-cooled environments are now being asked to understand pressurized fluid systems, coolant chemistry, leak detection technologies, hydraulic components, and integrated mechanical-electrical systems that simply did not exist in many legacy data centers. Not only so, but they are also asked to learn how these new elements interact with one another in complex environments, as opposed to understanding their functionality in isolation and in simple linear trajectories.

This transition represents more than a technical challenge. It represents a human performance challenge (Hollnagel, 2018). Traditional engineering analyses often give comparatively less attention to these human performance considerations.

Unlike routine maintenance activities, abnormal events rarely follow a script. A technician responding to an unexpected coolant leak must rapidly gather information, determine the source of the leak, evaluate potential electrical hazards, assess whether equipment remains energized, communicate with operations personnel, and decide whether to isolate equipment, continue operations, or escalate the response. These decisions are often made under considerable uncertainty and time pressure.

Figure 5 illustrates this transition. Engineering systems can identify that an abnormal condition exists, but they cannot determine what happens next. Once an alarm is generated, the event transitions from an engineering problem to a human performance problem. The quality of the outcome is increasingly influenced by the decisions people make, the information available to them, and the organizational systems that support those decisions.

Decision-Making Under Operational Pressure

One characteristic shared by many serious incidents across industries is that decisions are rarely made under ideal conditions (Rasmussen, 1997; Dekker, 2011). Instead, they are made while competing priorities demand attention.

In an AI data center, a cooling system failure may place millions of dollars of computing infrastructure at risk while simultaneously threatening service availability for customers. Such circumstances naturally create operational urgency. Personnel may feel pressure to restore cooling quickly, minimize downtime, or prevent further equipment damage. While these pressures are understandable, they can also influence how risks are perceived and managed (Kahneman, 2011).

Questions organizations should begin asking include:

- How will technicians respond to a liquid cooling failure they have never encountered before?
- Do personnel understand both the electrical and mechanical hazards associated with these systems?

- Have realistic failure scenarios been incorporated into emergency response exercises?
- Are procedures designed around how work is actually performed during abnormal events?
- Have operators been empowered to stop work when uncertainty exists?

These questions shift the conversation beyond equipment design and toward operational resilience.

Competency Beyond Technical Training

Training should extend beyond equipment operation. Organizations should ensure personnel understand the interactions between liquid cooling systems, electrical infrastructure, hazardous energy isolation, emergency response, and organizational decision-making.

Competency is not simply knowing how a system works during normal operations. True competency is demonstrated when personnel can recognize changing conditions, adapt appropriately, communicate effectively, and make sound decisions during unexpected situations.

As liquid cooling technologies continue to evolve, maintaining that competency will require continual learning rather than one-time training (Conklin, 2019).

Human Performance as a Critical Risk Control

Many organizations view human performance as something to be managed after an incident occurs. Increasingly, leading organizations recognize it as a critical risk control in its own right.

Engineering controls reduce the likelihood of abnormal conditions. Human performance determines how effectively those conditions are managed once they occur.

For this reason, organizations should begin evaluating not only whether engineering safeguards function as intended, but also whether the systems that support human decision-making — including training, supervision, communication, procedures, workload management, and organizational culture — are equally robust (Goncharenko, 2024).

As liquid cooling becomes more prevalent throughout AI infrastructure, the industry's ability to prevent Serious Incidents and Fatalities will depend not only on advances in engineering, but also on advances in understanding and supporting human performance.

8. Emotional Controls and SIF Prevention

This may be the most overlooked aspect of the conversation. Many significant events do not begin with a failed physical control. They begin with degraded human performance (see Figure 7).

And degraded human performance is often influenced by emotional conditions, including:

- Fear.
- Uncertainty.
- Frustration.
- Pressure.
- Overconfidence.
- Complacency.

These factors influence how people perceive risk, communicate concerns, and make decisions.

Recent research in human performance, combined with the author's work on Emotional Safety and Emotional Controls, suggests that organizations should begin viewing these conditions through the lens of Emotional Safety and Emotional Controls, like:

- Leadership.
- Trust.
- Communication.
- Resources.
- Workload.
- Learning systems.

These organizational factors shape the emotional conditions under which people operate. When emotional controls weaken, decision quality often declines. When decision quality declines, exposure increases. When exposure increases, the potential for serious events grows. The future of SIF prevention may require us to evaluate not only physical controls, but emotional controls as well, as summarized in the Figure 7.

Organizations have long invested in engineering controls. The next frontier of SIF prevention may be developing organizational systems capable of producing consistently high-quality decisions under pressure.

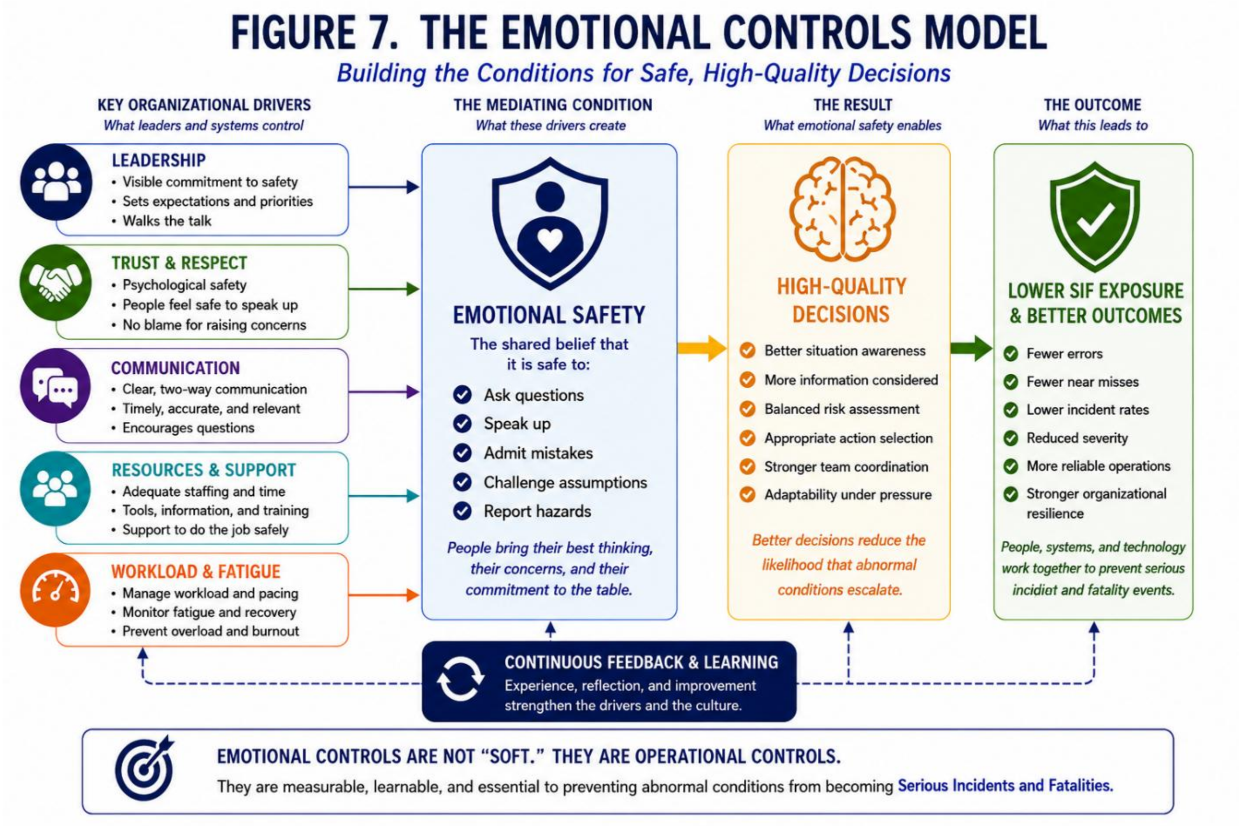

9. Contractor Competency

Data centers increasingly rely on multiple contractors for installation, commissioning, maintenance, and repair of liquid cooling systems. Unlike traditional electrical work, liquid cooling often requires mechanical, plumbing, controls, water treatment, and electrical expertise to interact simultaneously.

This creates organizational interfaces that have historically been uncommon within many data center environments. Ensuring that every contractor understands not only their own work, but also the hazards introduced by adjacent systems, may become one of the industry's most significant competency challenges (Goncharenko, 2026).

10. Questions Every Data Center Leader Should Be Asking

Before liquid cooling becomes fully normalized, leaders should consider the following questions:

- Have safety professionals been actively involved in design reviews?

- Have emergency response plans been tested against realistic liquid cooling failure scenarios?
- Do technicians understand both the mechanical and electrical hazards associated with these systems?
- Have critical risk controls been identified and validated?
- How might production pressure influence decision-making during a cooling system failure?
- What emotional and organizational factors could increase risk during emergency response activities?
- Have we considered all the layers of protection from Figure 8?

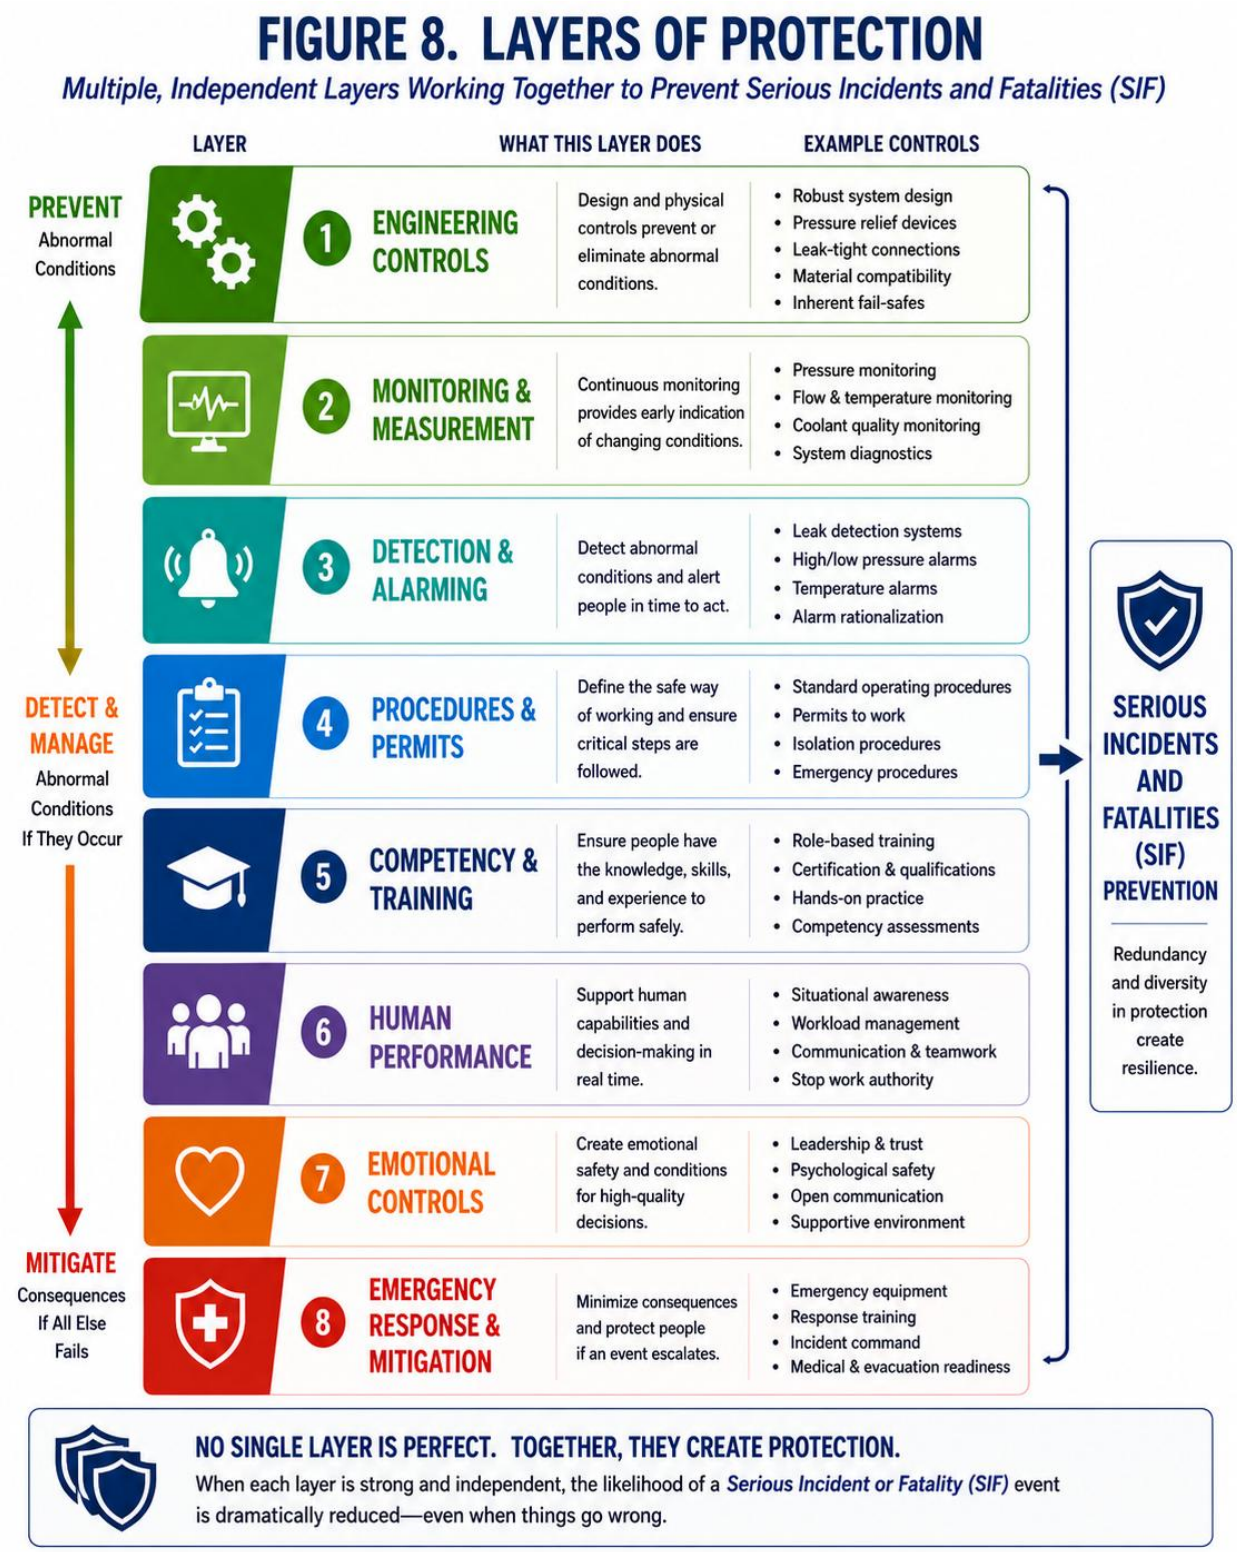

If these questions cannot be answered confidently, additional work may be necessary.

Conclusion

The data center industry has an opportunity that few industries receive. We are still early enough in the adoption curve to proactively identify and manage emerging risks. We do not need to wait for a catastrophic event to learn these lessons.

The technology is advancing rapidly and the investment is unprecedented. So, as the operational demands continue to grow, the question is whether our understanding of safety, human performance, and SIF prevention will advance at the same pace? I write this because the future of liquid cooling will not be defined solely by thermal efficiency or uptime. It will also be defined by how effectively we anticipate and manage the human and organizational risks that accompany it.

The time to ask these questions is before the industry is forced to answer them.

Bibliography:

American National Standards Institute (ANSI), & American Society of Safety Professionals (ASSP). (2019). *ANSI/ASSP Z10.0-2019: Occupational health and safety management systems*. American Society of Safety Professionals

American Petroleum Institute. (2021). *API Recommended Practice 754 (3rd ed.): Process safety performance indicators for the refining and petrochemical industries*. American Petroleum Institute.

ASHRAE Technical Committee 9.9. (2021). *Emergence and expansion of liquid cooling in mainstream data centers*. ASHRAE. https://www.ashrae.org/file%20library/technical%20resources/bookstore/emergence-and-expansion-of-liquid-cooling-in-mainstream-data-centers_wp.pdf

Bolgar, C. (2025). *AI chips are getting hotter: A microfluidics breakthrough goes straight to the silicon to cool up to three times better*. Microsoft Source. https://news.microsoft.com/source/features/innovation/microfluidics-liquid-cooling-ai-chips/

Center for Chemical Process Safety. (n.d.). *Guidelines for risk based process safety.* American Institute of Chemical Engineers. https://www.aiche.org/ccps/resources/publications/summaries/guidelines-risk-based-process-safety

Center for Chemical Process Safety. (2001). *Layer of protection analysis: Simplified process risk assessment*. American Institute of Chemical Engineers and John Wiley & Sons.

Clifford, R. (2026, May 12). *How direct-to-chip cooling is helping MSPs meet AI demand*. *ITPro*. https://www.itpro.com/technology/artificial-intelligence/how-direct-to-chip-cooling-is-helping-msps-meet-ai-demand

Conklin, T. (2012). *Pre-accident investigations: An introduction to organizational safety*. CRC Press.

Conklin, T. (2019). *The 5 principles of human performance: A contemporary update of the building blocks of human performance for the New View of Safety*. PreAccident Media.

Dekker, S. W. A. (2014). *The field guide to understanding "human error"* (3rd ed.). CRC Press.

Dekker, S. W. A. (2011). *Drift into failure: From hunting broken components to understanding complex systems*. CRC Press.

Goncharenko, S. (2024). *Save Lives: Pushing Boundaries of Human Factors*. Kindle Press.

Goncharenko, S. (2025). *Operationalizing Twenty First Century Safety: A Humancentric Practical Guide.* CRC Press.

Goncharenko, S. (2026). *Resilient Contractor Safety Management*. CRC Press.

Hollnagel, E. (2018). *Safety-I and Safety-II: The past and future of safety management*. CRC Press.

ICONPROCON. (2026, April 16). *How to detect coolant leaks in liquid-cooled data centers (CDU, direct-to-chip & immersion).* https://iconprocon.com/blog_post/how-to-detect-coolant-leaks-in-liquid-cooled-data-centers-cdu-direct-to-chip-immersion/

International Organization for Standardization. (2018). *ISO 45001:2018: Occupational health and safety management systems—Requirements with guidance for use*. International Organization for Standardization.

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

Leclerc, W. (2024, April 12). *Immersion cooling safety in a data center. Data Center Dynamics*. https://www.datacenterdynamics.com/en/opinions/immersion-cooling-safety-in-a-data-center/

National Fire Protection Association. (2024). *NFPA 70E: Standard for electrical safety in the workplace* (2024 ed.). National Fire Protection Association.

NVIDIA Corporation. (2025). *NVIDIA DGX SuperPOD: Data center design featuring NVIDIA DGX H100 systems* (Version 4). NVIDIA. https://docs.nvidia.com/dgx-superpod/design-guides/dgx-superpod-data-center-design-h100/latest/

NVIDIA Corporation. (2026, June 21). *Hotter than a hot tub: The 45°C breakthrough to cool AI's biggest machines*. NVIDIA Blog. https://blogs.nvidia.com/blog/liquid-cooling-ai-factories/

Open Compute Project. (n.d.). *Cooling Environments Project.* Open Compute Project Foundation. https://www.opencompute.org/community/cooling-environments

Open Compute Project. (2024). *OCP ACF reference design guidance* (Rev. 1). Open Compute Project Foundation. https://www.opencompute.org/documents/ocp-acf-reference-design-guidance-white-paper-r1-pdf

Open Compute Project. (2021). *Liquid cooling integration and logistics white paper* (Rev. 1.0.1). Open Compute Project Foundation. https://www.opencompute.org/documents/ocp-liquid-cooling-integration-and-logistics-white-paper-revision-1-0-1-pdf

Open Compute Project. (2025). *Modular Technology Cooling System (TCS).* Open Compute Project Foundation. https://www.opencompute.org/documents/ocp-modular-tcs-rev-1-final-2025-pdf

Occupational Safety and Health Administration. (n.d.). *29 C.F.R. Part 1910, Subpart S—Electrical*. U.S. Department of Labor.

Rasmussen, J. (1997). Risk management in a dynamic society: A modelling problem. *Safety Science, 27*(2–3), 183–213. https://doi.org/10.1016/S0925-7535(97)00052-0

Reason, J. (1990). *Human error*. Cambridge University Press.

Reason, J. (2016). *Managing the risks of organizational accidents*. Routledge.

Rahkonen, T. (2025, February 4). *AI and cooling: Methods and capacities*. Uptime Institute Intelligence. https://intelligence.uptimeinstitute.com/resource/ai-and-cooling-methods-and-capacities

Walicki, B. (2026, February 18). *Glycol under pressure: Managing risks in AI-scale cooling loops. Data Center Dynamics*. https://www.datacenterdynamics.com/en/opinions/glycol-under-pressure-managing-risks-in-ai-scale-cooling-loops/

Weick, K. E., & Sutcliffe, K. M. (2007). *Managing the unexpected: Resilient performance in an age of uncertainty* (2nd ed.). Jossey-Bass.

About the Author

Simon Goncharenko, PhD, CSP, CSHO, is Global Director of Environmental Health & Safety for Serverfarm, the author of Save Lives: Pushing Boundaries of Human Factors (2024) and Resilient Contractor Safety Management (2026), and founder of Save Lives Global Research. He has more than three decades of experience in construction, oil and gas, and

hyperscale data centers, and serves on national safety committees focused on human performance and contractor safety.

The views expressed in this publication are solely those of the author.

Appendix A – Sample Questions for Liquid Cooling Design Review

1. Has maintainability been evaluated?
2. Are emergency shutdown sequences documented?
3. Can cooling loops be isolated independently?
4. Are leak response procedures validated?
5. Have contractors demonstrated competency?
6. Has human performance been evaluated?
7. Have emotional controls been considered?

Save Lives Global Research

Advancing Human Performance

Preventing Serious Incidents and Fatalities

Building Resilient Organizations