DATA CENTER
SAFETY COUNCIL

Managing High-Risk Activities in Live Data Centers: A Shared Approach to Risk, Resilience, and Readiness

Contents

Acronyms

AI - Artificial Intelligence: used in the context of autonomous systems and emerging technologies in data centers.

DCSC - Data Center Safety Council: a collaborative industry body establishing health and safety standards common to live data center operations.

HOP - Human and Organizational Performance: a safety approach that focuses on system factors rather than individual blame.

HRA - High-Risk Activity: tasks where credible failure could cause a serious injury or fatality or catastrophic service loss.

HSE - Health, Safety, Environment: the discipline or area of work associated with the management of workplace health, safety and environment.

HV - High Voltage: electrical systems typically above 1000V AC or 1500V DC.

LOTO - Lockout/Tagout: safety procedure for controlling hazardous energy during maintenance.

MEWP - Mobile Elevated Work Platform: equipment used for work at height (e.g., scissor lifts, boom lifts, etc.).

NASA - National Aeronautics and Space Administration: referenced for SCAD HFES resilience research.

OSHA - Occupational Safety and Health Administration: US regulatory body for workplace safety (referenced in compliance context).

PDU - Power Distribution Unit: equipment distributing power to servers and IT equipment.

PTW - Permit-to-Work: formal authorization system for high-risk activities.

SCAD - Safety Culture Assessment and Development: NASA research program on resilience and adaptive coordination.

SIF - Serious Injury or Fatality: Life-altering harm or death; key threshold for HRA classification.

UPS - Uninterruptible Power Supply: battery backup system for maintaining power during outages.

Introduction

In a live data center, uptime and safety share a single goal: continuity. Both depend on teams recognizing, managing, and learning from high-risk work in environments where energy is constant and margins for error are small.

This white paper sets out a shared approach to identifying and managing high-risk activities (HRAs) across global data center operations. Data Center Safety Council (DCSC) member organizations aligned their definitions, validated them through peer review, and established a framework that focuses leadership attention where credible failure could result in serious injury, fatalities, or catastrophic service loss. Traditional compliance-driven safety systems often mistake paperwork for assurance. In contrast, effective HRA management depends on anticipation, control verification, and adaptive readiness. HRAs must be planned, reviewed, and led with clear intent, visible competence, and a live discussion of "what could go wrong?" Permits, risk assessments, and method statements serve as validation tools, not static records.

The paper also explores how organizations learn from "work-as-done," recognizing that successful outcomes do not always mean systems worked as they were designed. Authentic learning requires psychological safety, open dialogue, and feedback loops that translate field insights into system change.

The goal is simple but fundamental: to build a consistent, data-driven, human-centered approach that allows every organization, vendor, and technician to recognize high-risk activities in the same way, manage them effectively, and provide a better chance that every person goes home safe at the end of each day.

Why a Common Language Matters

In live data centers, not understanding what constitutes an HRA undermines resilience. Without a shared language, teams could misclassify risks, under-communicate hazards, or fail to monitor and control HRAs effectively. Consistency matters. When the same task is viewed differently across teams, vendors, or regions, assurance becomes unreliable and leadership attention drifts towards noise rather than consequence.

Developing a common language is not about enforcing uniform terminology. It is about building a shared mental model of risk. Across global projects, the same activity is often classified differently due to regulatory expectations, translation gaps, regional norms, and contractor-specific standards. Terms such as "live," "energized," "isolated," or "controlled" do not always mean the same thing between organizations or even between teams in the same organization. These mismatches widen the gap between work-as-imaged and work-as-done, making it harder to maintain consistent control over serious hazards.

A shared language removes that ambiguity. It allows teams to recognize HRAs in the same way, regardless of geographical locations or contractual relationships. This is essential in environments where high-energy, interdependent systems leave little room for misunderstanding.

A common language also strengthens learning. When teams use consistent definitions for HRAs, incident trends can be compared more accurately across sites, critical controls can be tested with the same assumptions, and lessons can be shared with less distortion. It increases the visibility of weak signals and early indicators of instability.

A shared language does not eliminate uncertainty, but it does make it discussable, and that is a foundation of resilient, high-performing operations.

High-Risk Activities in Live Data Centers

A high-risk activity (HRA) is any task where credible failure could cause a serious injury or fatality (SIF) or catastrophic service loss in a live data center. What makes an activity high-risk isn't how often it's done or how many controls are in place; it's the presence of high levels of energy that, if the task were to fail, could cause life-altering harm or worse, death.

Just because a task is complex or critical, does not necessarily mean it is high-risk. Many tasks carry low physical consequences, even if they are essential to uptime. If the worst-case scenario results in a delay, a reworked task, or a minor injury, this is not a HRA. It requires controls, but not at the same level of scrutiny as a HRA.

Creating and agreeing upon a distinction between "life-altering harm/death" and a "minor injury" is critical to ensure the spotlight shines brightly on tasks that pose the highest risk to worker safety.

Establishing a Common Ground

To reach a shared understanding of what qualifies as an HRA in a live environment, members of the DCSC conducted a structured comparison of all existing HRA lists across participating organizations. Each list was reviewed line by line to identify:

- activities consistently recognized as high-risk across member organizations;
- variations in definitions, terminology, or scope between organizations; and
- outliers that reflect unique operational or regulatory conditions.

It is not about avoiding failure, it is about ensuring that when failure occurs, its impact is minimized and contained.

This comparative process created a consolidated baseline. Where differences existed, they were debated and resolved through discussion among the HRA workstream members. Consensus focused on the energy present and the potential severity of consequences, rather than task frequency or complexity.

Once alignment was reached, the resulting list was shared for validation across the broader DCSC organizations. The peer review strengthened confidence that the definitions were relevant, operationally grounded and representative of real work performed in live data center environments. The following list represents a validated set of activities that DCSC member organizations agree constitute high-risk work in operational live data center environments.

High-risk activities include:

- Electrical Work
- Control of Hazardous Energy (Lockout/Tagout)
- Confined Space Entry
- Work at Height
- Hot Work (e.g., cutting, welding, grinding)
- Lifting and Rigging
- Ground Disturbance
- Mobile Plant and Equipment (Mobile Equipment)

These tasks involve high-energy sources and potential brittleness in systems. Recognizing energy is foundational to effective control. Tools like the Energy Wheel help visualize risk and reinforce awareness.

Examples of High-Risk Activities

The task itself does not define the risk. Two activities may look identical on paper but differ significantly in consequences depending on the energy involved, the environment in which the work occurs, and the system state at the time. Classification must be context-driven, not purely procedural. It is the magnitude and concentration of energy that determines whether a task crosses the threshold into high-risk territory. Recognizing the distinction allows teams to focus on assurance, supervision, and learning where failure would have irreversible outcomes, not where the work appears complex or critical.

The following examples are drawn from real scenarios across DCSC member organizations. They illustrate the distinction between activities with high-consequence potential and those that, while still controlled, pose limited physical danger.

	High Risk	*Not High Risk*
Electrical Work	Accessing a 415V live panel to test load during commissioning in a live UPS room, with direct exposure to high energy failure, could cause arc flash, serious injury or fire.	Plugging in a server to a pre-installed PDU, because it is low voltage, with no exposure to conductors, with minimal energy risk.
Hazardous Energy	Isolating multiple power sources feeding a chiller with electrical, pneumatic, and hydraulic energy before maintenance.	Isolating a lighting circuit for a build replacement. Low voltage, minimal risk.
Confined Space	Entering a sump pit for valve replacement in a chilled water system with limited entry/exit points and potential atmospheric hazards.	Working adjacent to a tank without entering it.
Work at Height	Installing a cable tray above a live switchgear using a mobile elevated work platform with the potential for falls and the proximity to live, energized equipment.	Using a step stool to access a low rack shelf.
Hot Work	Cutting steel supports in a battery room using an angle grinder creating a fire and exploding risk from off-gassing batteries in an enclosed space.	Grinding brackets in a dedicated fabrication zone, well away from any combustible materials.

	High Risk	*Not High Risk*
Lifting and Rigging	Hoisting replacement critical plant equipment onto a roof above live data halls with high loads, risk of dropping objects, and potential impact to the critical infrastructure of the building.	Moving a 20 kg (45 lb) crate with a pallet jack in an open space, as it is manual handling with no rigging in a low-risk environment.
Excavation and Ground Disturbance	Deep trenching near known underground HV electrical conduits with a risk of service strike, electrocution and utility disruptions.	Surface vacuum excavating soil away from known utilities, not into the subsurface or risk of cable strikes.
Mobile Plant and Equipment	Operating a telehandler to place equipment inside a live campus with pedestrians and vehicle interfaces with the chance of collision, crushing, and potential system disruptions.	Use of a floor scrubber in a cordoned-off corridor, where energy is low, is well controlled, and there would be minimal consequences if the activity were to fail.

A Look to the Future: Autonomous Equipment and Emerging Risks

As data center campuses expand, autonomous systems such as robotics, drones, and AI tools will become more common. While these tools bring efficiency, they introduce new high-risk dynamics, especially when interacting with human workers in shared spaces. These systems can behave unpredictably, lack situational awareness, and may not respond correctly to unplanned changes. Any autonomous or AI-powered machinery operating in a live environment should trigger caution during planning of deployment.

Regulations in this area are still developing, and there are no globally agreed-upon standards for autonomous systems in these environments. As a result, organizations must build their own governance frameworks, and they can draw on lessons from other high-reliability domains, such as aviation. The following principles remain the same:

- Verify reliability and failure modes.
- Treat updates and model changes as high-risk interventions.
- Test autonomy in controlled environments before deployment.
- Establish clear protocols for human-machine separation or shared space interaction.
- Train workers to understand the system limits, not just its capabilities.

Managing High-Risk Activities

In live data centers, planning is not just paperwork; it's a first line of defense against serious harm and operational uncertainty. Managing high-risk activities is not about ticking boxes. It's about anticipating what could go wrong and ensuring the right people, tools, practices, and safeguards are in place. Managing high-risk activities means:

- conducting advanced planning for high-risk activities to avoid task overlaps, energy conflicts, or operational clashes across multiple vendors;
- testing readiness by reviewing who is leading and who is supervising work, ensuring competence and clarity of roles matter more than the number of signatures;
- verifying critical controls before work starts to validate that controls exist and are effective, not just documented;
- identifying credible possible failures and ensuring that if one layer of defense fails, there is a recovery path that prevents serious harm;
- treating permits, method statements, and risk assessments as live artefacts that drive discussion and evolve as conditions change;
- sequencing and timing so that live tasks, such as switching, lifting, or hot work, are carried out under stable operational conditions, with sufficient support available if tasks fail;
- embedding pre-task reviews that encourage the work team to pause, confirm energy isolation, and discuss 'what could go wrong' before proceeding; and
- establishing clear communication and escalation procedures so everyone knows how the call to stop work is triggered and how restarts are managed.

In resilience terms, this is not control for control's sake, but control that anticipates brittleness and ensures adaptability when needed [Ref. Woods et al., Behind Human Error]

Effective management of high-risk activities shows up in three visible ways:

1. Clarity of intent - everyone involved can state the purpose of the task, the critical steps, and the worst-case outcomes if it fails.
2. Control verification - leaders verify not only that the controls are in place, but that they will work under load.
3. Adaptive readiness - teams know to stop safely, recover, and resume if conditions change.

Common Failure Patterns

During member reviews, several recurring failures appear in poorly managed HRAs:

- planning done in isolation from operations, resulting in energy overlap or unrecognizable system interdependencies;
- permits issued remotely or without verification;
- over-reliance on documentation, with no real discussion about credible failures;
- work sequences based on schedule pressure, not system readiness; and
- supervisors stretched across multiple concurrent high-risk tasks.

Acknowledging these patterns helps organizations spot weakness early and intervene before conditions align for a serious event.

Learning from Work-as-Done

Every high-risk activity is a stress test for your systems. On paper, everything looks controlled, but in the field, it rarely plays out that way. Teams adapt constantly to incomplete information, late changes, or competing operational pressures. Good outcomes don't always mean the system worked. They might just mean your people worked around it brilliantly. Recognizing that difference is the foundation of real learning [Woods et al. Behind Human Error].

If you want to know what really happens in a data center, you must create an environment where psychological safety and mutual trust is present at all levels. When workers see leaders respond constructively to what they share, it builds confidence to identify early warning signals before they escalate.

From Observations to Understanding

Leadership site walks and engagement provide opportunities to ask real questions that keep fatal and significant risks front and center. Curiosity is key, as is approaching the conversations with respect and acknowledgement of the experience of the people carrying out the work. Questions like:

- Where could this activity fail?
- If the control fails, what is plan B?
- Is there a better way of working, and if you could change anything at all, what would you change?

It's important for these site walks and engagement sessions to be seen not solely as information-gathering, but as a consistent activity that provides vital feedback to leadership. Feedback once given is acted upon to build trust with the people carrying out the work.

Learning from work-as-done only matters if the information gathered leads to change. Too often, engagement sessions produce good insight but no action. When people see that nothing has changed, they stop speaking up. Leadership must therefore treat field feedback as learning data, not commentary. Insights should be logged, reviewed, and acted upon in the same way that incident investigations are followed through. The key difference is timing. Learning from normal work happens before failure. [Ref - Howie - The Post-Incident Guide]

Site walks and engagement sessions should feed directly into:

- updating procedures that don't match real conditions;
- adjusting control strategies that prove impractical in live environments;
- reframing leadership assurance priorities based on what actually drives risk exposure; and
- informing the next cycle of HRA reviews.

What does good learning look like? You can see when an organization learns from work-as-done when:

- frontline workers feel safe raising issues without fear of repercussion;
- supervisors discuss risks in plain language, not acronyms;
- leaders act quickly on feedback and communicate what changed as a result; and
- high-potential near misses and observations are reviewed at the same level as incident investigations.

Learning from work-as-done is not a soft exercise; it's a hard requirement for resilient operations. In live environments, where systems are tight, interdependent, and energy-dense, the quickest way to strengthen safety is to understand how people succeed under real conditions and then design the system to better support them.

Building and Embedding Resilience

In data center operations, resilience is a measure of whether the system can keep performing safely when conditions change. Data center infrastructure already demonstrates it with backup generators, dual feeds, and UPS systems that hold the line when the grid fails. The same mindset needs to be applied to how high-risk activities are managed. Resilience in this context means designing human and organizational systems that can anticipate, absorb, and recover from disruption without serious harm. It is not about avoiding failure, it is about ensuring that when failure occurs, its impact is minimized and contained. Without this capability, every deviation becomes a potential loss event. With it, small disturbances are recognized early, and the system can adapt before people get hurt or operations are interrupted.

Managing HRAs in live environments is therefore less about control and more about readiness. Readiness comes from shared understanding, visible leadership, and systems that listen. It is built through dialogue, verification, and deliberate practice. When leaders walk the floor, validate the integrity of critical controls, and talk openly about what could go wrong, they strengthen the organization's ability to recover when something does go wrong.

True resilience is designed into the work itself. Procedures are clear enough to set boundaries, yet flexible enough to allow informed adjustments. Controls are layered so that a single failure cannot trigger a catastrophe. Information flows freely so weak signals and near misses are surfaced early. These are the human equivalents of redundant power supply and bypass circuits, parallel paths that prevent total collapse when one part of the system fails.

Resilience also relies on leadership behavior. Resilient leaders do not equate quiet periods with safety. They understand that systems drift towards failure when feedback loops are weak and assumptions go untested. They encourage curiosity and treat deviation as data. They focus less on fault and more on recovery, setting the tone that adaptability is a strength, not a weakness. As Weick observed in his work on sensemaking, resilience grows out of collective mindfulness. That is, people staying alert to what could happen, even when everything seems normal. Embedding resilience in data centers requires the same intent as designing an electrical redundancy system. It means rehearsing responses to credible failures, verifying critical controls actually work under load, and reviewing what happens when work does not go to plan.

This approach is not theoretical; it is validated across multiple high-reliability fields. NASA's SCAD HFES studies found that adaptive coordination prevents escalation more effectively than procedural compliance. Hollnagel's research in Safety in the Digital Age demonstrated that the ability to respond, monitor, learn, and anticipate defines resilient systems. Feltovich and colleagues confirmed

that coordinated improvisation enables control under surprise. Together, these findings show that resilience depends not on perfection, but on adaptation and collective sensemaking.

If you want a safer, more resilient operational environment, start by asking better questions, engaging with the field, and planning like it matters. Resilience isn't perfection. It's adaptation done well, done early and done together.

Conclusion

Some key dos and don'ts as takeaways to consider when planning your high-risk activities:

Do	*Don't*
Focus on the energy, not the task.	Assume compliance equals safe.
Build shared language with your teams and supply chains.	Stretch supervisors across multiple HRAs.
Verify controls at the point of work.	Rely on remote or paper-based assurance.
Ensure supervisors are present, competent, and available.	Treat HRAs as routine or administrative.
Treat credible failures as the starting point for planning.	Confuse quiet periods with safety.

High-risk work in live environments is predictable. The hazards are known, the energy is all around them, and the consequences of uncontrolled failure are serious. The difference between safe operations, with safe failures, and serious injuries and fatalities, is often down to the quality of planning, the strength of the control verification, and the competence and curiosity of the people leading the work.

Organizations that succeed treat HRA management as a shared responsibility. They build a common language for risk, they verify controls at the point of work, and they encourage teams to act on weak signals early.

If data center operations want to prevent life-altering harm and keep services stable, the focus must stay on serious energy, credible failures, and resilient systems that are built on shared understanding, clear roles, and prompt actions. These will outperform systems built on paperwork every time.

Every organization is at a different starting point with regard to HRA management. Recognize the work that carries real consequences and plan it with precision. Verify the controls that matter. Learn from the work your people actually do, and review your system as often as chances occur.

Resilience is built through deliberate practice. Start there.

References

This white paper was written based on the readings and research of some of the world's leading thinkers in this field. Below is a list of reading materials and references where you can read and learn more about work in this area.

- Feltovick, Hoffman, Woods, and Roesler - Mastering Complexity in System Design (2004) - Cognitive resilience and joint activity.
- Weick, K. E. (1995). Sensemaking in Organizations. Thousand Oaks, CA: SAGE Publications.
- Woods, D. D., Dekker, S., Cook, R., Johannesen, L., and Sarter, N. (2010). Behind Human Error (2nd ed.). Farnham, UK: Ashgate Publishing.
- Dekker, S. (2014). The Field Guide to Understanding 'Human Error' (3rd ed.). Farnham, UK: Ashgate Publishing.
- Hallowell, M. R. (2025). "Energy-Based Safety." CRC Press.
- Hollnagel, E. (2023). Safety in the Digital Age: The Fuzz and the Unpredictability in Socio-Technical Systems. London: Routledge.
- National Aeronautics and Space Administration (NASA). (2016). Safety Culture Assessment and Development (SCAD): Final Report. Human Factors and Ergonomics Society (HFES). Moffett Field, CA: NASA Ames Research Center. Available at: https://humanfactors.arc.nasa.gov/
- Conklin, T. (2012). Pre-Accident Investigations: An Introduction to Organizational Safety. Ashgate Publishing.
- Dekker, S. (2014). The Field Guide to Understanding 'Human Error' (3rd ed.) Ashgate Publishing.

About the Data Center Safety Council (DCSC)

Data Center Safety Council provides a unified approach towards ensuring safety and wellbeing of personnel in operating data centers. The industry has seen a period of rapid growth adding urgency to the need for standardization of best practices, knowledge sharing and collaboration, and defining resources for training and development. This industry group will provide advocacy for aligning health and safety regulations and addressing operational health and safety challenges in a consistent approach. More information at [www.datacentersafetycouncil.org.](http://www.datacentersafetycouncil.org)

DATA CENTER
SAFETY COUNCIL