

Maximising Control Effectiveness

A Systems Approach to Reducing Serious Risk in High-Hazard Operations



**Incident
Analytics™**

www.incidentanalytics.com.au

This fifth whitepaper in the series examines a fundamental challenge in critical risk management: control effectiveness is not determined by whether controls exist, but by whether they reliably reduce exposure under real operating conditions.

Drawing on research into human performance, system reliability, organisational drift and barrier management, together with large-scale incident analysis and field-based evaluation of critical risk systems, it explores the factors that influence control effectiveness and reframes it as a system-reliability challenge rather than a compliance problem.

Table of Contents

Prologue

SCALE®: Evidence-based methodology and analysis technology	6
Research study	8

Whitepaper | Issue 5

Executive summary	10
The problem with control presence	12
What the data shows	14
Why controls fail	16
What determines effectiveness	18
From compliance to reliability	24
What maximising control effectiveness looks like	26
Key recommendations for safety leaders	28
Conclusion	30

Copyright © 2026 Incident Analytics Pty Ltd. This document remains the intellectual property of Incident Analytics Pty Ltd and is protected by copyright and registered trademarks. No material from this document is to be reproduced or used in any format without express written permission.

Prologue

This whitepaper continues our deep dive into serious injury and fatality potential incidents, controls, and the factors that influence them.

We have completed an extensive analysis of more than 10,000 incidents, including a closer examination of 680 high-severity potential incidents. This work reflects several years of collaboration with our clients. We thank our clients, partners, and team for their contributions.

The result is a series of whitepapers that unpack data, examine research findings, and share insights drawn from real-world incident patterns. Together, these papers explore how organisations can improve the performance of critical controls in high-hazard industries.

Our series includes:

1. Our first whitepaper, **A Deep Dive into SIFp: Understanding the Misclassification of Serious Incidents**, examines the factors that contribute to the misclassification of serious incidents.
2. Our second whitepaper, **Making Critical Controls Work: Overcoming Barriers to Effective Implementation**, builds on these insights with a focus on critical controls.
3. Our third whitepaper, **Beyond the Hierarchy: A Modern Approach to Reliable and Effective Safety Controls**, expands on our evidence-based methodology for assessing the adequacy and reliability of controls.
4. Our fourth whitepaper, **Unmasking the Human Element: Cognitive and Behavioural Influences on Serious Safety Incidents** explores the cognitive and behavioural drivers behind serious incidents.

This fifth paper explores that control effectiveness should be judged by one question: does the control reliably reduce exposure in the conditions where work actually occurs? The answer depends not only on control design, but on context, detectability, human dependence, and operational reliability.



Introducing SCALE®: Evidence-based methodology and analysis technology

Using our [SCALE®](#) technology, we analysed the data from more than 10,000 incidents. Our proven approach takes the guesswork out of incident analysis. By using a structured, repeatable approach, we identify which incidents need deeper investigation. This systematic process has resulted in more effective solutions to prevent serious incidents in high-risk work.

Our comprehensive and repeatable fatality potential decision-tree system determines which incidents merit deep-dive analysis. The subsequent analysis leverages a proprietary controls database that helps inform strategies to improve the effectiveness of controls, including:

- Strengthening critical safety controls
- Addressing human behaviour patterns
- Fixing gaps in organisational systems

SCALE® analysis technology helps determine which incidents need deep-dive analysis. It is built on a robust methodology that enables organisations to better understand the effectiveness of controls and the operational, and system factors that contribute to the conditions for unplanned events to occur.





Research study

Incident Analytics analysed data from across 12 industry sectors spanning 12 countries and six continents. These industries included mining, road and rail transport, warehousing, utilities, and port operations.

Our applied research into serious injury and fatality prevention is based on an analysis of data sourced from many real-world operations.

We analysed more than 10,000 incidents and found there is inconsistent taxonomy and insufficient information to inform adequate learning. Where this wasn't an issue, we applied deeper focus (on 5,091 incidents) and of these, 680 had serious injury or fatality potential. These specific incidents were the subject of much deeper review and analysis using SCALE®.

We identified more than 1,000+ specific controls in high-risk work that were absent or weren't effective in preventing or mitigating SIFp incidents from occurring.



12 Industry Sectors	Utilities	Metals Manufacturing	Mining
	Engineering	Food Manufacturing	Agricultural Services
	Land Management	Fuel Distribution	Warehousing
	Rail Operations	Transportation	Port Operations

Executive Summary

Maximising Control Effectiveness: A Systems Approach to Reducing Serious Risk in High-Hazard Operations

Over the past decade, significant effort has been invested across high-hazard industries to define and assure the effective implementation of critical controls. Organisations have developed detailed control frameworks, performance standards, and verification regimes intended to prevent serious injuries and fatalities.

Despite this, serious incidents continue to occur in environments where controls are clearly defined, understood, and often verified as being in place.

This paper explores a fundamental issue:

Control effectiveness is not determined by whether controls exist. It is determined by whether they reliably reduce exposure in real operating conditions.



Control presence is not enough

A commonly held view on control effectiveness is that if they are functional, clearly defined, properly implemented, and regularly verified, then workers and the organisation should be adequately protected.

This approach assumes:

- Controls can be implemented as designed
- Workers can reliably follow procedures
- Documented controls represent actual exposure reduction


In practice, these assumptions do not consistently hold.

Analysis of real-world operations shows that:

- Controls are often difficult or impractical to implement
- Many controls are in fact support processes, not exposure-reduction mechanisms
- Control performance varies significantly under different operational conditions

This creates a persistent disconnect between work-as-imagined (the defined system) and work-as-performed (ie execution in real conditions when conditions change).

Work-as-Imagined (WAI) vs. Work-as-Done (WAD)¹



Work as Imagined (WAI) describes how work is expected to be performed according to procedures, standards, training and design assumptions, while Work as Done (WAD) reflects how people actually adapt and execute work in real operating conditions to achieve the required outcome. Consistent with this thinking, our focus is on whether controls are functioning effectively in the field and genuinely reducing risk exposure.

Organisations are often not managing control effectiveness – they are managing control presence.

We propose a shift to: “is the control actually reducing risk?”

¹ The effectiveness of controls can be undermined where there is a mismatch between work-as-imagined (how work is expected or prescribed to occur) and work-as-done (how work is actually performed in practice) (Hollnagel and Clay-Williams, 2022; Conklin, 2019)

What the data shows

Across a large dataset of serious incidents and high-potential events, clear patterns emerge regarding control performance:

Controls are frequently difficult or unworkable

A significant proportion of controls cannot be fully implemented in practice. In one analysis, over one-third of controls were difficult or unworkable. This is not non-compliance — it is system mismatch. It's just the way we've tended to do it around here.

Workers adapt controls to complete tasks

Where controls are impractical, workers adapt, steps are modified, and controls are partially applied. These adaptations are rational, often necessary, and more often system-driven, not behaviour-driven.

Human error occurs even in “fully enabled” environments


Unintended errors (slips, lapses, mistakes) occur even when controls are available. They increase under low attention and high familiarity and are driven by predictable cognitive limitations.

Workarounds emerge from control design issues

Intentional deviations occur when controls are difficult or incompatible with the task. They become normalised over time and tend to reflect system design limitations.


These patterns reveal a fundamental insight:

Control failure is not random. It follows predictable pathways driven by design, context, and human factors. When controls depend on perfect execution under variable conditions, their effectiveness degrades. When this degradation is not visible or managed, exposure persists.



1,000+
Controls Analysed

Errors contributing to control failures



Why controls fail

The failure pathway below shows how that unreliability develops over time — from design assumptions, through adaptation and drift, to undetected degradation and exposure.



Our own applied research into SIF or HiPo incidents reveals that the cycle of control failure begins as early as the design phase and is likely to be influenced by a combination of factors:

1. Dependence on human vigilance and judgement

A large proportion of controls rely on workers to:

- Notice deviations
- Interpret changing conditions
- Make correct decisions
- Execute precise actions

This places significant cognitive demand on individuals operating in environments that are often complex and dynamic. Under conditions such as time pressure, fatigue, or distraction, this model becomes unreliable. Risk accumulates in the gaps between perception, decision, and action, leading inevitably to drift.

2. Failure to manage adaptation and task change

Work rarely proceeds exactly as planned. Tasks evolve, conditions change, and unexpected variations occur. In many systems, control frameworks do not adequately respond to these changes. Controls are not reset, re-established, or re-validated when conditions shift.

As a result, adaptation becomes a key source of instability. Failure is most likely to occur during change, not during steady-state execution.

3. Reliance on assumed state rather than verified state

Many control frameworks depend on the belief that conditions remain safe once controls are implemented. Verification is either implicit or infrequent. In practice, system state changes continuously. Without explicit verification, controls degrade silently.

This creates a gap between what is believed to be true and what is actually true at the point of work.

4. Weakness in detection of control degradation

Control degradation is often gradual and difficult to observe. Where detection relies on human observation alone, early warning signals are missed or misinterpreted. This allows exposure to increase without intervention.

By the time failure becomes visible, recovery options are limited.

5. Over-reliance on downstream or protective controls

In many control architectures, insufficient emphasis is placed on preventative controls that eliminate or isolate hazards.

Instead, systems rely on protective or mitigative controls that operate after exposure has already occurred. These controls are inherently weaker because they depend on earlier controls having partially failed.

A consistent theme emerges:

Controls fail where they rely on assumption, adaptation, and human perfection. This is not a reflection of workforce capability. It is a function of how systems are designed.

Risk materialises exactly where controls depend on conditions that cannot be consistently sustained in real work.

What determines effectiveness

Control effectiveness is not a single attribute. It is a function of how a control interacts with the hazard, the operating environment, and the people required to implement it.

Incident Analytics’ deep analysis of serious injury and fatality potential (SIFp) incidents shows that five factors consistently determine whether a control will perform when it matters.

1. Control Purpose: Where the control acts in the hazard pathway

Controls vary fundamentally in how they influence risk. Some controls eliminate or isolate hazardous energy entirely. Others reduce the likelihood of exposure, detect emerging danger, or mitigate consequences after exposure occurs.

This distinction is critical. Controls that act earlier in the hazard pathway provide inherently stronger protection than those that rely on response after exposure has already occurred.

2. Control Design: Degree of reliance on human performance

Controls also differ in how dependably they deliver their intended function. At one end of the spectrum are controls that function autonomously, requiring minimal or no human interaction. At the other end are controls that depend entirely on human judgement, memory, and attention.

Human-dependent controls are inherently variable. Their performance is sensitive to fatigue, distraction, workload, and complexity. Even highly trained and motivated workers cannot consistently deliver perfect execution in dynamic environments.

As a result, controls that rely heavily on human cognition introduce predictable points of failure.

Human error probability

How likely a task is to be performed incorrectly, by task type.

Task	Generic Task Type	Error Probability (log scale)	1 in X
GTT1	Totally unfamiliar task, performed at speed	0.55	1 in 2
GTT2	Change system to new state, no supervision or procedures	0.26	1 in 4
GTT3	Complex task requiring high comprehension and skill	0.16	1 in 6
GTT4	Fairly simple task, performed rapidly / scant attention	0.09	1 in 11
GTT5	Routine, highly practised, low consequences	0.02	1 in 50
GTT6	Change system to new state, following procedures + checking	0.003	1 in 333
GTT7	Completely familiar, well-designed, highly practised	0.0004	1 in 2500

What makes error more likely

Error-producing conditions multipliers:



x11

Time shortage



x4

Misperception of risk



x3

Operator inexperience



x3

Inadequate information



x1.3

Emotional stress

Source: HEART (Human Error Assessment and Reduction Technique), J.C. Williams, 1985 ².

What the HEART research means for control effectiveness:

The question is not whether people will make mistakes. Decades of research show that under certain conditions they inevitably will. The more important question is whether the control system remains effective when normal human variability occurs. Controls that tolerate error are generally more reliable than controls that depend on error-free performance.

² Williams, J.C. (1985) HEART – A proposed method for achieving high reliability in process operation by means of human factors engineering technology. In: Proceedings of a Symposium on the Achievement of Reliability in Operating Plant. Birmingham: Safety and Reliability Society (SaRS), NEC.

3. Operating Context: Conditions under which the control must perform

Controls do not operate in isolation. Their effectiveness is shaped by the environment in which they are applied. Stable, engineered environments tend to preserve control reliability. In contrast, dynamic or unstable conditions—such as time pressure, environmental variability, or concurrent work—degrade performance.

Controls that appear robust under ideal conditions often become fragile when exposed to real-world variability.

4. Detectability: Ability to identify degradation or exposure change

Many control failures do not occur as sudden events. They emerge through gradual degradation, loss of separation, drift in conditions, or partial implementation. The ability to detect these changes early is a critical determinant of effectiveness.

Where detection depends on human vigilance or weak signals, failures can persist unnoticed until exposure escalates beyond recovery.

5. Operational Reliability: How the control actually performs in practice

The final and most critical factor is how consistently the control is applied under real conditions. Incident Analytics' research into the drivers of serious incidents shows that controls often degrade due to:



- Time pressure
- Competing priorities
- Equipment constraints
- Local workarounds
- Organisational drift

This creates a gap between expected and actual performance.

A control may appear strong by design, but if it is inconsistently applied or fails during actual events, its contribution to risk reduction is limited.

Taken together, these factors show that control effectiveness is not determined by design intent alone. It emerges from the interaction between design, environment, detection, and real-world use.

The five factors that determine control effectiveness:



From compliance to reliability

The limitations of traditional control frameworks are not due to lack of effort or intent. Most organisations invest heavily in defining controls, developing procedures, and verifying compliance.

The issue lies in how effectiveness is defined.

Control systems are commonly built on a compliance model:

- Define the control
- Communicate the requirement
- Verify adherence

This model assumes that if people follow the rules, exposure will be reduced. However, our research into SIFp incidents shows that compliance is not a reliable proxy for effectiveness.

Controls can be:

- Fully documented
- Clearly communicated
- Regularly verified

Yet still fail under real operating conditions.

The reason is simple. Compliance measures whether a control was followed, not whether it was capable of preventing exposure under variation.

Compliance model	Reliability model
Follow the rule	Tolerates variation
Behaviour focus	System-design focus
Verify presence	Verify performance

This distinction is critical.

Real work is dynamic. Conditions change, tasks evolve, and competing demands emerge. In these environments, systems that rely on consistent human performance are inherently unstable. Even where compliance is high, variability in execution inevitably introduces gaps.

This leads to a shift in how control effectiveness must be understood.

The key question is no longer: “Did people follow the control?” It becomes: “Would this control still prevent harm under normal human behaviour and real operating conditions?”

This reframing moves the focus from rule adherence to system reliability.

A reliable control system does not depend on perfect execution. It is designed to tolerate variation, detect degradation early, and prevent exposure even when conditions are not ideal.

This requires a shift in emphasis from:

- Documentation to functionality
- Compliance to performance
- Expectation to evidence

Ultimately, maximising control effectiveness is not about strengthening rules. It is about designing systems that are inherently more difficult to fail.

What maximising control effectiveness looks like

High performing control systems share six characteristics:



Unsafe state are hard to engage



Reliance on human performance is reduced



Exposure is minimised before controls are required



System state is explicitly verified



Adaptation is controlled, not unmanaged



Detection is strengthened and explicit

Characteristic	What it means
Unsafe state are hard to engage	Exposure is controlled through physical or system constraints rather than behavioural rules. Hazardous conditions are eliminated, isolated, or physically restricted so unsafe states are structurally harder to achieve and less dependent on worker vigilance.
Reliance on human performance is reduced	Effective systems minimise dependence on attention, memory, and judgement by shifting from procedural controls toward technical, exposure-triggered, and engineered controls. The worker manages system conditions rather than acting as the primary control.
Exposure is minimised before controls are required	Risk is reduced by design through minimising hazard interaction, shortening exposure duration, and moving work into lower-risk states. This reduces the burden placed on downstream controls.
System state is explicitly verified	Safe conditions are confirmed rather than assumed. Clear definitions of safe states, structured verification points, and distinction between verified and assumed conditions reduce the likelihood of hidden degradation.
Adaptation is controlled, not unmanaged	Work deviations and changing conditions are anticipated and managed through defined triggers, mandatory resets, and visible control re-establishment processes. Reliability is strengthened during transitions where failure is most likely.
Detection is strengthened and explicit	Effective systems actively detect control degradation, changing exposure, and emerging risks using clear signals, redundancy, and structured verification. Early detection provides the opportunity to intervene before exposure escalates.

The key takeaway is control effectiveness improves when systems become less dependent on perfect human performance and more capable of preventing, detecting and tolerating variation in real work.

Key recommendations for safety leaders

The following recommendations provide a practical framework for building more reliable and effective control systems:

1 Strengthen control design

Prioritise controls that act directly on exposure. Focus investment on controls that eliminate, contain, separate or shield people from hazards rather than relying on administrative safeguards alone.

Systematically reduce human dependence. Where possible, replace reliance on vigilance, memory and judgement with engineered, automated or system-enforced controls.

Build systems that tolerate human error. Design controls that remain effective when mistakes occur, rather than relying on perfect execution.

2 Design for real work

Reduce exposure before controls are required: Redesign work to minimise contact with hazards so that critical controls have less exposure to manage.

Design controls for real-world variation: Ensure controls remain effective under normal operational pressures, changing conditions and human variability.

Focus verification on controls most vulnerable to slips, lapses, or mistakes.

3 Make degradation visible

Implement effective mechanisms that reveal when controls are weakening, partially implemented, or no longer functioning as intended.

Evaluate controls based on performance, not presence. Measure whether controls reliably reduce exposure in real work, not simply whether they exist, are documented, or have been verified.

4 Shift assurance from compliance to performance

Replace assumed safe states with verified safe states. Critical conditions should be actively confirmed rather than assumed to remain safe over time.

Focus assurance on reliability, not compliance: Assess whether controls prevent harm in practice, not simply whether people followed the prescribed process.


Manage adaptation explicitly: Recognise that work changes and design systems that identify, control and respond to those changes.

Conclusion


Control effectiveness sits at the centre of serious incident prevention.

The evidence is clear. Controls often exist, are understood, and are applied, but still fail to prevent high-consequence outcomes. This is not due to lack of intent or effort. It is the result of systems that depend on conditions that cannot be consistently sustained in real work.

Human variability, changing environments, and operational pressure are not exceptions. They are normal features of work. Any control system that relies on perfect execution under these conditions will inevitably degrade.



Maximising control effectiveness requires a different approach



This shift reframes control effectiveness as a system design challenge. The objective is not to eliminate human error. It is to design systems that remain effective in its presence.

Ultimately, the goal is simple: to ensure that critical risks are managed by controls that reliably reduce exposure when it matters most, not just controls that exist, or controls that are followed under ideal conditions.

The organisations that succeed in this shift will not simply improve compliance outcomes.

They will build systems that are fundamentally more robust, predictable, and defensible, even under the conditions where failure is most likely to occur.

“All effective controls must function, but not all functioning controls are effective”

References:

J. Reason (1990): Human Error
J. Reason (1997): Managing the Risks of Organisational Accidents
S. Dekker (2011): Drift Into Failure
S. Dekker (2014): Field Guide to Understanding Human Error
E. Hollnagel (2014): Safety-I and Safety-II
E. Hollnagel (2009): ETTO Principle
E. Hollnagel (2004): Barriers and Accident Prevention
J.C. Williams (1985): Human Error Assessment & Reduction Technique (HEART)
ICMM Critical Control Management Guide
Energy Institute (2018): BowTie / Barrier Management guidance
Rasmussen (1997): Risk Migration model
Snook (2000): Practical Drift
Woods & Hollnagel (2006): Resilience Engineering
Todd Conklin (2019): The 5 Principles of Human Performance



**Incident
Analytics™**

Uncover hidden risks, improve controls & prevent serious incidents.

We are a risk management research & data analytics company supporting high-hazard industries including mining, utilities, and transport.

Our research and analysis technology gives senior leaders and safety professionals the business intelligence to help improve safety performance.

We hope our work helps to inform your safety strategies in the workplace. To learn more about Incident Analytics go to [our website](#), or [get in touch](#) to see what we can do for your business.

P: 1300 955 300
E: info@incidentanalytics.com.au
www.incidentanalytics.com.au

Copyright © 2026 Incident Analytics Pty Ltd. This document remains the intellectual property of Incident Analytics Pty Ltd and is protected by copyright and registered trademarks. No material from this document is to be reproduced or used in any format without express written permission.