CYBER PHYSICAL RISK ACADEMY

MANAGE RISK, SECURE YOUR FUTURE

Cyber Incident Strategy for the Process Industry

Guest Lecture | Technical University Darmstadt

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

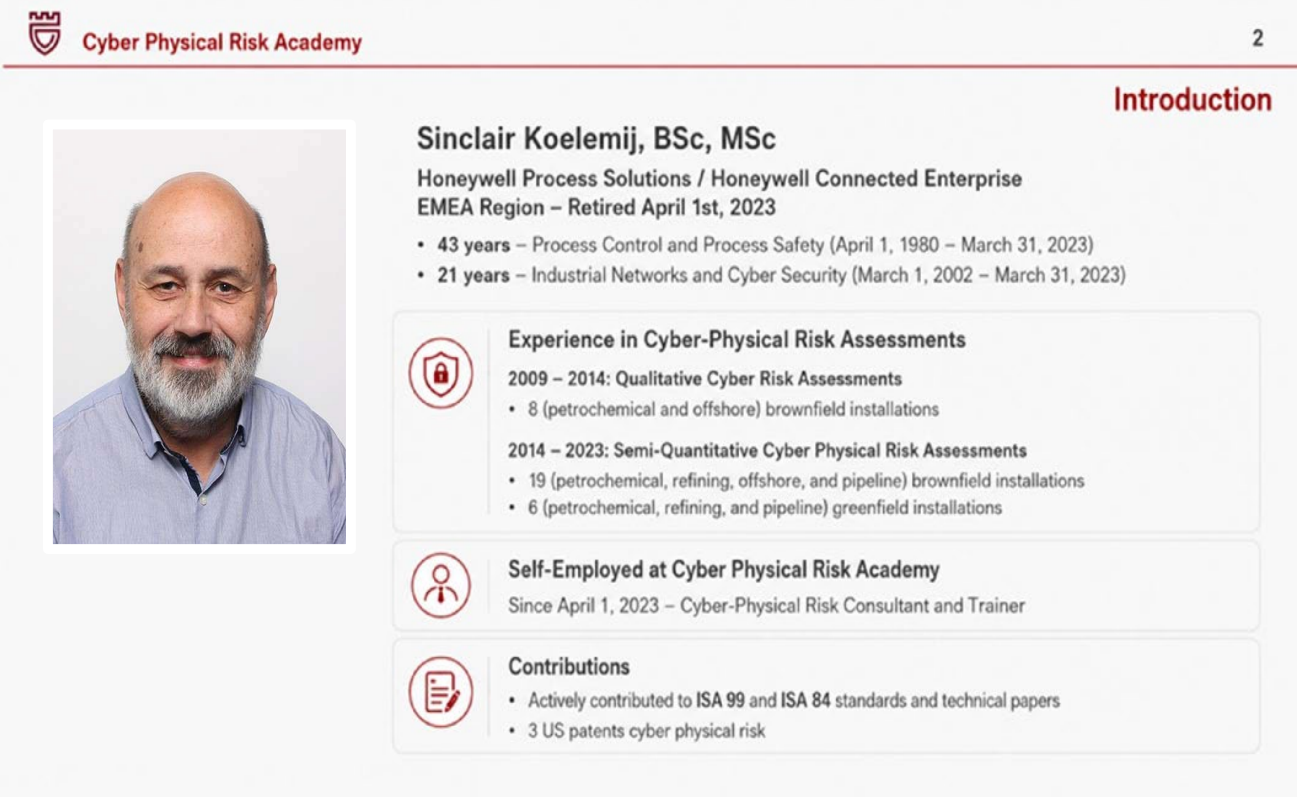

Cyber Physical Risk Academy

Abstract: This article explains why incident response in process automation cannot be treated as a normal Information Technology incident response activity. In a cyber-physical system, response actions can affect a live physical process. Containment, eradication, recovery, and restoration must therefore be judged against their process safety impact.

The central argument is simple: a cyber-attack may start the incident, but a poorly executed response may finish it. The objective is not only to remove the attacker or restore clean systems. The objective is to maintain or reach a safe, observable, controllable, and valid process state while the incident is investigated, contained, and recovered.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

Copyright notice

Copyright © 2026 Sinclair Koelemij. All rights reserved.

Unless otherwise stated, the text, structure, diagrams, slide concepts, and original visual material in this article are the intellectual property of Sinclair Koelemij and may not be copied, reproduced, distributed, modified, translated, published, or used for commercial purposes without prior written permission.

Where third-party material, public standards terminology, referenced frameworks, trademarks, or externally sourced elements are used, copyright and ownership remain with their respective rights holders. Such material is used only for reference, commentary, education, or illustration where permitted.

The author claims copyright only to the extent legally permitted, including the original selection, arrangement, explanation, interpretation, wording, diagrams, and visual presentation developed for this article.

Cyber Physical Risk Academy and associated visual branding are used as author/publisher identification. No transfer of rights is implied.

For permission requests, citation use, or reproduction requests, contact the author.

What is a cyber-physical system?

A cyber-physical system combines digital automation with a physical process. Sensors measure the process, controllers calculate actions, actuators change the physical state, and operators supervise the result.

In process automation, the cyber part is not only computers and networks. It includes control logic, alarms, safety functions, historian data, engineering tools, and operating procedures. The physical part includes equipment such as pumps, valves, vessels, compressors, reactors, columns, furnaces, and pipelines.

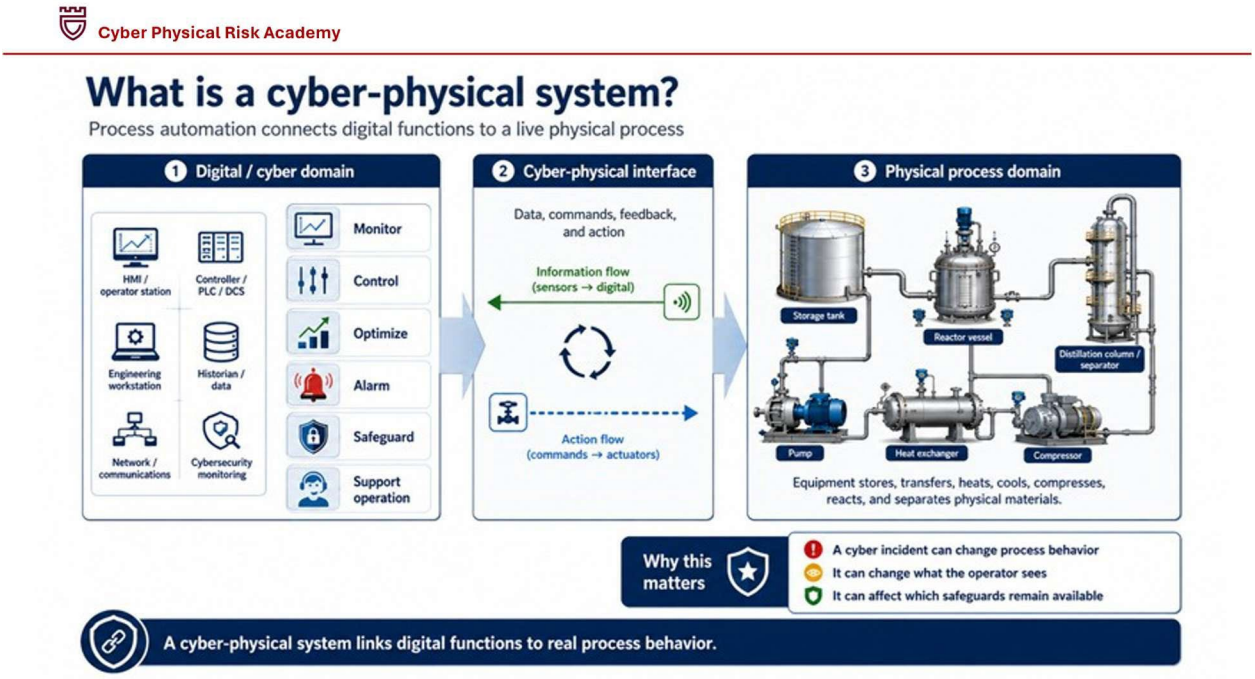

The key difference from a normal Information Technology (IT) system is the consequence path. In IT, a failure usually affects data, applications, business services, or privacy. In a cyber-physical system, a digital failure can change the physical process. That can lead to loss of control, loss of observability, equipment damage, environmental release, production loss, injuries, or fatalities.

This also changes the regulatory context. Process industries operate under permits, licenses, safety cases, and regulatory obligations that are directly connected to accepted risk criteria. These criteria may include limits for individual fatality risk, societal risk, environmental impact, and major accident potential. A cyber incident that changes the process state can therefore become more than a security incident. It can challenge the basis on which the plant is allowed to operate.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

What can happen when a cyber-physical system is attacked?

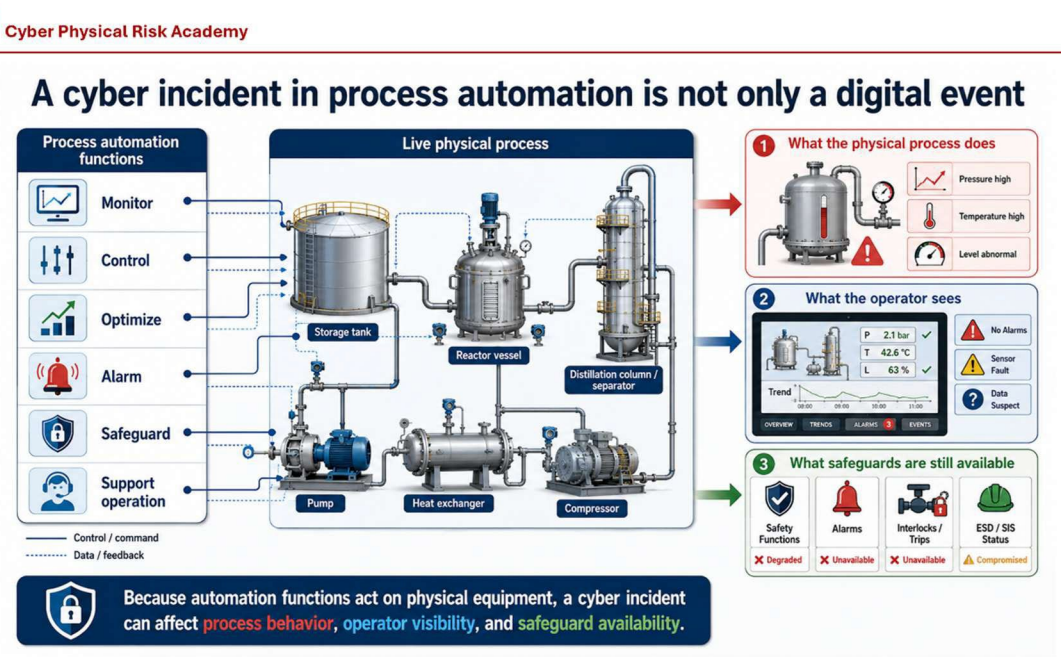

In process automation, the relevant functions do more than handle information. They monitor, control, optimize, alarm, safeguard, and support the operation of a live physical process. Those functions are connected to equipment such as tanks, pumps, heat exchangers, compressors, reactors, and separation units that store, transfer, heat, cool, compress, react, or separate physical materials.

That connection is what makes process automation different from a normal digital environment. A cyber incident does not remain confined to the digital domain. It can influence what the physical process actually does, what the operator is able to see and understand, and which safeguards are still available to prevent escalation.

This is why the slide distinguishes three consequences. First, the process itself may change: pressures, temperatures, levels, flows, compositions, or operating sequences may move away from their intended state. Second, operator visibility may be degraded or misleading: values may be frozen, manipulated, delayed, inconsistent, or presented without the alarms and context needed for correct judgment. Third, the availability or effectiveness of safeguards may be reduced: alarms may be suppressed, interlocks may be bypassed, and Safety Instrumented System (SIS) functions may be unavailable, degraded, or untrustworthy.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

So, the core message is simple: in process automation, a cyber incident is not only an information security problem. It is a cyber-physical event that can simultaneously affect process behavior, operator observability, and safeguard integrity. That is why incident response in this domain must always be judged against one central question: what does this mean for the safety and controllability of the live process?

Most people understand cyber incident response as a digital activity. Something is compromised, the incident team investigates, contains the attacker, removes persistence, restores systems, and brings services back online.

That logic is useful, but it is incomplete in process automation.

In process automation, the digital system is connected to a live physical process. It monitors, controls, alarms, optimizes, and sometimes safeguards equipment that contains energy, pressure, temperature, hazardous materials, rotating machinery, or chemical reactions. What happens in the digital environment can therefore change what happens in the physical process.

That changes the role of incident response.

A response action is not only a cybersecurity action. It may become a process action. Rebooting a controller, isolating a network segment, restoring an engineering configuration, disabling an alarm path, blocking operator visibility, or interrupting a safeguard function may all be technically reasonable from a narrow cybersecurity perspective. From a process safety perspective, the same actions may be dangerous.

Cyber Physical Risk Academy

Incident response in cyber-physical systems: first, do not become part of the evidence.

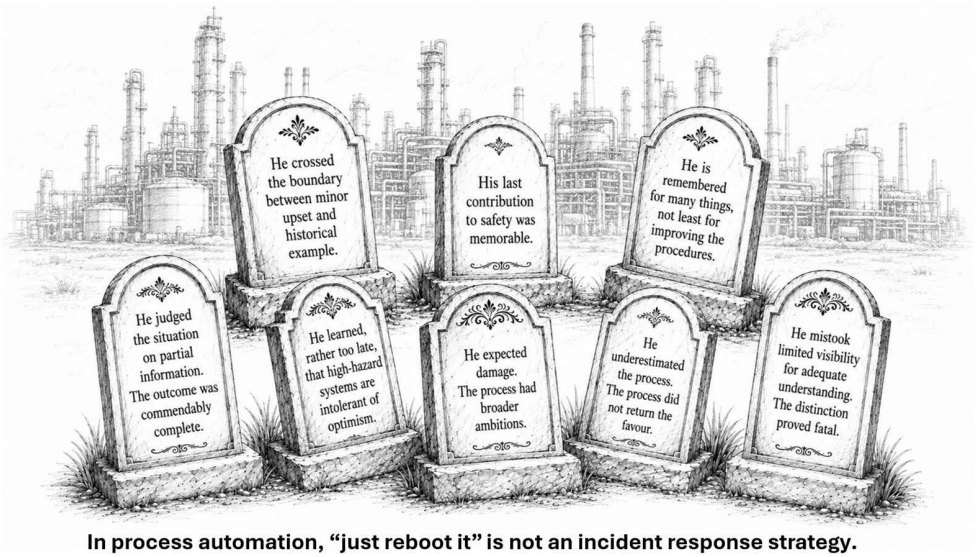

In process automation, "just reboot it" is not an incident response strategy.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

This is the uncomfortable starting point: in process automation, incident response can become dangerous not only because of the attack, but also because of the response itself.

A cyber-attack may push the process toward an unsafe state. A poorly executed response may push it further, faster, or in a less visible way.

That is why the rational decision in process automation is not always the fastest digital containment action. The rational decision is the action that reduces total risk: cyber risk, process safety risk, personnel exposure, environmental impact, and recovery risk.

It also changes the meaning of proportionality. In a normal Information Technology (IT) incident, proportionality is often judged by how effectively the attacker is contained, systems are cleaned, and services are restored. In process automation, that is too narrow. A response is only proportional if it helps maintain or restore a safe process state.

This is why “just reboot it” is not an incident response strategy in process automation.

Detection is not the end of the incident

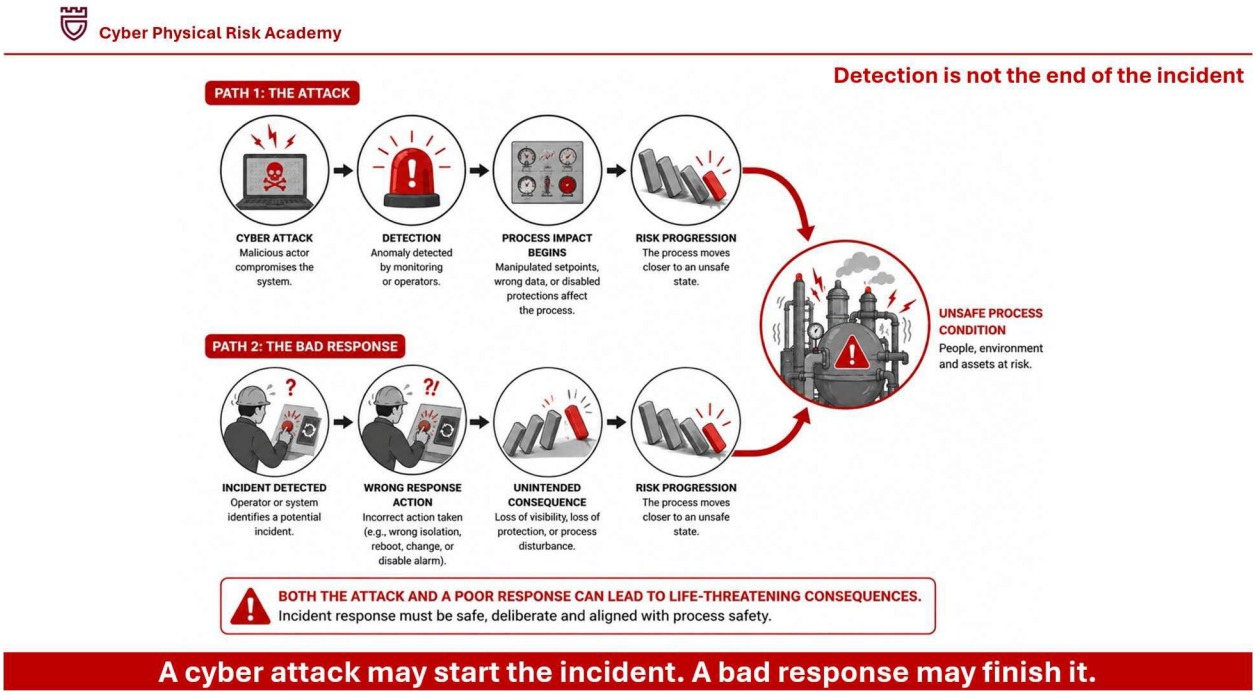

Detection often feels like progress. Something suspicious has been found. A security monitoring system may have detected abnormal communication. An operator may have seen unexpected behavior. A controller may have reported a diagnostic alarm. At that point, the organization may feel it has moved from uncertainty into response.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

But in process automation, detection is not the end of the incident. It is the moment where the response team becomes part of the causal chain.

The attacker may already have manipulated setpoints, logic, alarm behavior, device configuration, network communication, or operator information. The process may already be moving toward a less safe condition. At that point, the wrong response can remove the remaining barriers that still prevent escalation.

For example, a cyber-attack may disturb the automation function. But a response team may unintentionally make the situation worse by:

- isolating a network path that still carries critical operator visibility;
- rebooting a controller while the process is in a narrow operating window;
- disabling alarms to suppress noise while those alarms are still needed;
- restoring a configuration that no longer matches the actual plant state;
- blocking communications needed by a safety-critical diagnostic or monitoring function;
- forcing a shutdown sequence without understanding the current process conditions.

The attack may start the incident. A bad response may finish it.

That sentence sounds dramatic, but it is technically accurate. In a live process, response actions have physical consequences. The priority is not simply to remove the attacker. The priority is to maintain or reach a safe process state while investigating and containing the incident.

Incident response in process automation must therefore be understood as controlled intervention in a live physical process, not merely as restoration of a digital environment.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

Detection may not come from where we expect

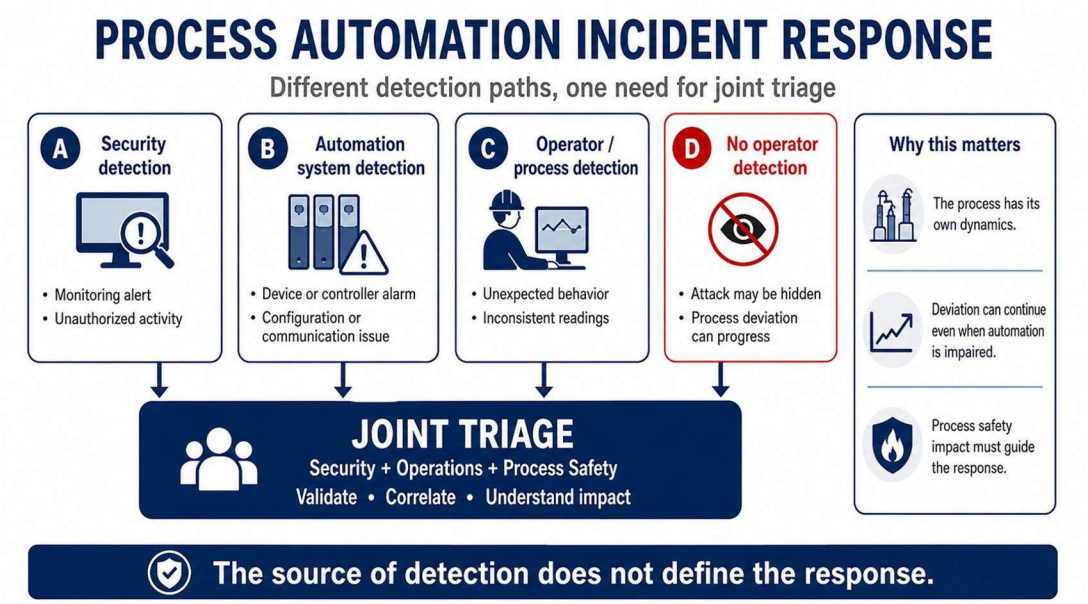

A common mistake is to assume that a process automation incident will be detected neatly by the cybersecurity team, the automation system, or the operator. That is too optimistic.

In high-hazard industries, the response logic must always consider that protection may fail, may be unavailable, or may not act as expected. This is not pessimism. It is normal process safety thinking. Protection layers are designed, tested, and credited because failures are possible. The same logic must apply during a cyber-physical incident. Detection may be late, operator visibility may be misleading, alarms may be suppressed, safeguards may be bypassed or degraded, and redundancy may no longer be available without being recognized.

A cybersecurity team or Security Operations Center (SOC) may detect an attempted attack, an incomplete attack, or an attacker that makes enough noise to trigger monitoring. Examples include suspicious remote access, abnormal network traffic, unauthorized engineering activity, unexpected protocol use, or unusual asset communication.

But a successful attack may have succeeded precisely because it remained below the detection threshold of the cybersecurity function. It may have used legitimate credentials, trusted engineering tools, expected communication paths, or slow manipulation that does not look malicious from a conventional monitoring perspective.

Therefore, the absence of a cybersecurity alert must not delay process-aware triage when the process behavior, automation behavior, operator visibility, or safeguard state becomes questionable.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

The reason is simple: in a cyber-physical incident, the first useful indication may not come from the cybersecurity monitoring function. It may appear first as an automation symptom, as an operator observation, as an unexplained process deviation, or it may remain hidden while the process condition is still reported as normal.

This means that the organization may need to move forward with incomplete information. That does not mean taking blind containment action. It means escalating early into structured joint triage, even when the available evidence is incomplete and even when the process still appears normal. The purpose of triage is to determine whether the apparent normal condition can still be trusted.

An automation symptom does not automatically mean equipment failure. Controller alarms, device diagnostics, communication failures, unexpected mode changes, abnormal command sequences, hidden or sabotaged redundancy, or control parameters that no longer match the actual process condition may all indicate that the automation state is no longer aligned with the live process state. In such cases, the automation system may still be running, but its control actions may no longer be appropriate for the actual process condition.

An operator observation does not automatically mean a normal operational upset. Unexpected process behavior, inconsistent readings, alarms that do not match the physical situation, weak or delayed control response, or actuator behavior that does not match the command may indicate that the operator is seeing symptoms of automation manipulation, degraded observability, or loss of control integrity.

The most difficult case is where the deviation is not visible to the operator at all. Trends may appear stable while the physical process is drifting. Safeguards may be bypassed, inhibited, degraded, or unavailable, while their actual state is not visible on the normal operator displays. The operator may still believe the process is inside its normal envelope while the actual process has already started to deviate toward an unsafe state.

Therefore, early escalation is not based on proof that a cyber-physical incident is occurring. It is based on uncertainty about whether process behavior, automation behavior, operator visibility, control effectiveness, or safeguard status can still be trusted. Joint triage may conclude that the situation is normal, but in high-hazard operations that conclusion should be validated, not assumed.

That is why the response should not wait for a confirmed cybersecurity alert or a clear operator alarm. When process behavior, automation behavior, operator visibility, or safeguard status becomes questionable, the organization must move into joint triage. Security, process

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

operations, process safety, and automation engineering must correlate their evidence and determine what the condition means for the live process.

That last case is often the most dangerous. An attack may blind the operator from relevant digital sensor values. Values may be frozen, delayed, filtered, substituted, or manipulated. Alarm behavior may be suppressed. Trends may appear stable while the physical process is drifting. Safeguards may be bypassed, inhibited, degraded, or unavailable, while their actual state is not visible on the normal operator displays. The operator may therefore believe that operation remains normal and protected, while the actual process has already started to deviate toward an unsafe or uncontrolled state and protective functions have been weakened, bypassed, or removed.

This is one of the key differences between digital incident response and process automation incident response. In process automation, an anomaly that cannot be explained by known process behavior, equipment failure, maintenance activity, operator action, or approved automation change must be treated as a potential indication of malicious interference until it is understood.

This places a different expectation on the process operator than on the average office employee in a normal digital environment. An office employee is usually expected to recognize obvious cyber warning signs, such as suspicious emails, unusual login prompts, or abnormal system behavior. A process operator must be alert to a different class of indicators: process behavior that no longer fits normal failure patterns, control actions that do not match the process condition, alarms that do not fit the physical situation, unexpected mode changes, abnormal actuator response, or safeguard status that appears inconsistent. The operator does not need to diagnose the cyber cause, but must recognize when the process or automation behavior is no longer credible and escalate early.

That does not mean every anomaly is an attack. Most anomalies will still have normal operational or technical causes, such as equipment degradation, process disturbance, maintenance activity, operator action, or an approved automation change.

Every anomaly should have a root cause. In normal operation, that cause is often found in the process, the equipment, control settings, recent maintenance activity, or operator action. In a cyber-physical incident, however, malicious interference may present itself through exactly the same kind of symptoms. It may look like a technical failure, an operational upset, a communication problem, or an unusual process disturbance.

That makes root cause analysis more important, not less. Not knowing the root cause immediately does not automatically mean that production must stop. But it does mean that the anomaly should not be dismissed, normalized, or closed without explanation. The organization

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

must continue to look for a credible cause, and that cause set must include malicious interference when the anomaly cannot be explained by normal process, equipment, maintenance, operator, or approved automation conditions.

Waiting for a clear cybersecurity alert may be too late. At the same time, recognizing malicious interference can be difficult because the first symptoms may look like normal process or automation problems.

For that reason, unexplained deviations must remain open until a credible cause is found. A strange control response, an unexpected mode change, inconsistent readings, degraded visibility, abnormal actuator behavior, an unsuitable command, or an unexpected safeguard status does not prove an attack. But if the deviation cannot be explained by normal process behavior, equipment condition, maintenance activity, operator action, or approved automation change, malicious interference must remain part of the possible cause set.

At that point, the issue should move into process-aware triage. The purpose is not to prove cyber compromise immediately. The purpose is to determine whether process behavior, automation behavior, operator visibility, control effectiveness, and safeguard status can still be trusted.

This requires a change in awareness. In modern process automation, malicious interference should not be treated as a distant or unlikely explanation that is considered only after all other causes have been exhausted. It is now a normal credible cause category for unexplained anomalies, especially when the behavior affects control, observability, alarms, safeguards, engineering changes, or operator trust. The organization does not need to assume that every unexplained anomaly is an attack, but it must keep malicious interference in the active cause set until a credible non-malicious cause has been confirmed.

That is why unexplained deviations matter. An attack against process automation may express itself through normal-looking symptoms: an unexplained control response, an unexpected mode change, inconsistent readings, a command that does not fit the process condition, degraded visibility, abnormal actuator behavior, or a safeguard status that does not match what should be available. None of these symptoms proves an attack by itself. But if they cannot be explained by credible process, equipment, maintenance, operator, or approved automation causes, they should remain under investigation and may need to trigger process-aware triage.

That triage is needed because the process has its own dynamics. Temperature, pressure, flow, reaction rate, level, material inventory, heat input, heat removal, and equipment stress continue to change even when the automation layer is unavailable, compromised, or misleading. Feed may continue, inventory may accumulate or deplete, heat may build up faster than it is

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

removed, pressure may rise, and equipment may move closer to its mechanical or process limits.

Therefore, unexplained anomalies must not be dismissed only because the cybersecurity monitoring function is silent or because the operator display still appears normal. An unexplained anomaly is not proof of an attack, and it does not automatically require a shutdown, trip, or move to degraded operation. It is a reason to keep the anomaly open, search for the root cause, and determine whether the process and its automation support can still be trusted.

The purpose of process-aware triage is to determine what is happening, what is affected, what is still trustworthy, what could become unsafe, and whether the remaining intervention window is being preserved.

This triage cannot be improvised from nothing during the incident. The relevant cyber-physical scenarios should be identified, analyzed, and documented beforehand as part of the cyber-physical risk analysis. This includes scenarios where malicious interference can affect the process state, the trustworthiness of automation support, or the ability of operations to intervene safely. Predefined scenarios help the response team understand which functions matter, which dependencies are critical, which symptoms may appear first, which protections may be affected, and which actions could make the situation worse.

This means that incident response requires several knowledge disciplines at the same time. Security must investigate the digital and process automation environment. Process operations must maintain or restore a safe process state. Process safety must assess hazards, safeguard availability, escalation potential, degraded modes, and safe-state requirements. Automation engineering must determine whether the automation functions, control modes, diagnostics, alarms, communication paths, and engineering changes can still be trusted.

These disciplines must not work as separate response tracks. They must be closely aligned and coordinated around the priorities imposed by the live process. The process dynamics determine the available time, the acceptable actions, and the order in which decisions must be made. Security investigation remains necessary, but it must operate within the time and action constraints set by process operations and process safety.

Traditional organizations are often still organized in separate pillars: security, process operations, process safety, automation engineering, and maintenance. That structure is useful for competence and accountability, but it is not sufficient for cyber-physical incident response. If escalation remains pillar-based, each discipline may optimize its own response while missing the combined effect on the live process.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

Emergency response and crisis management should therefore be understood as connected horizontal processes, not as isolated organizational silos. Emergency response starts with the immediate need to protect people, the environment, and the installation, and to stabilize the site. As the incident develops, emergency response extends into crisis management. Crisis management then coordinates the wider organizational response across multiple teams, including priorities, resources, internal and external communication, regulatory interfaces, business impact, recovery sequencing, and longer-term decision making.

Modern organizations therefore need an escalation process that brings the relevant disciplines together early, aligns them around the actual process state, and coordinates decisions according to the priorities imposed by process dynamics, personnel safety, and process safety. The practical priority is therefore clear: first preserve or restore a safe process state, then investigate, contain, eradicate, and recover within that operational constraint.

The source of detection does not define the response

A process deviation may have many possible causes. It may result from equipment failure, process disturbance, operator action, maintenance error, automation failure, or malicious interference. It may be detected by the operator, by the automation system, by process safety indicators, by cybersecurity monitoring, or it may not be detected early enough to preserve a comfortable intervention window.

That uncertainty is the reason for triage.

In process automation, the organization should not wait for a confirmed Indication of Compromise before involving the right disciplines. A confirmed Indication of Compromise is useful evidence, but it is not the only valid escalation trigger. Escalation should also occur when the process behavior, automation behavior, operator visibility, or safeguard state becomes questionable and the condition cannot be explained by a credible operational or technical cause.

The first question is therefore not:

“Who detected it first?”

The first question is:

“Could this condition, or our response to it, affect the ability to maintain or reach a safe process state?”

If the answer may be yes, joint triage is required.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

Joint triage is the first structured assessment by the disciplines needed to understand both the digital condition and the live process condition. Security investigates the digital and process automation environment. Process operations assess the live process state and decides how to maintain or restore safe operation. Process safety assesses hazards, safeguard availability, degraded modes, escalation potential, and safe-state requirements. Automation engineering assesses whether the automation functions and operator information can still be trusted. Detection source, by itself, is only context. The response category should be determined during joint triage, based on potential process safety impact and the trustworthiness of the process automation support.

Joint triage must determine whether the process condition is stable or becoming more critical, whether the automation state is trustworthy, whether safeguards and fallback capability are still available, whether there is still a safe operator intervention window, and whether proposed response actions could remove the remaining ability to maintain or reach a safe process state.

The response should be driven by process safety impact, process dynamics, and the trustworthiness of the process automation support, not by the source of the first signal.

Process safety impact defines the response

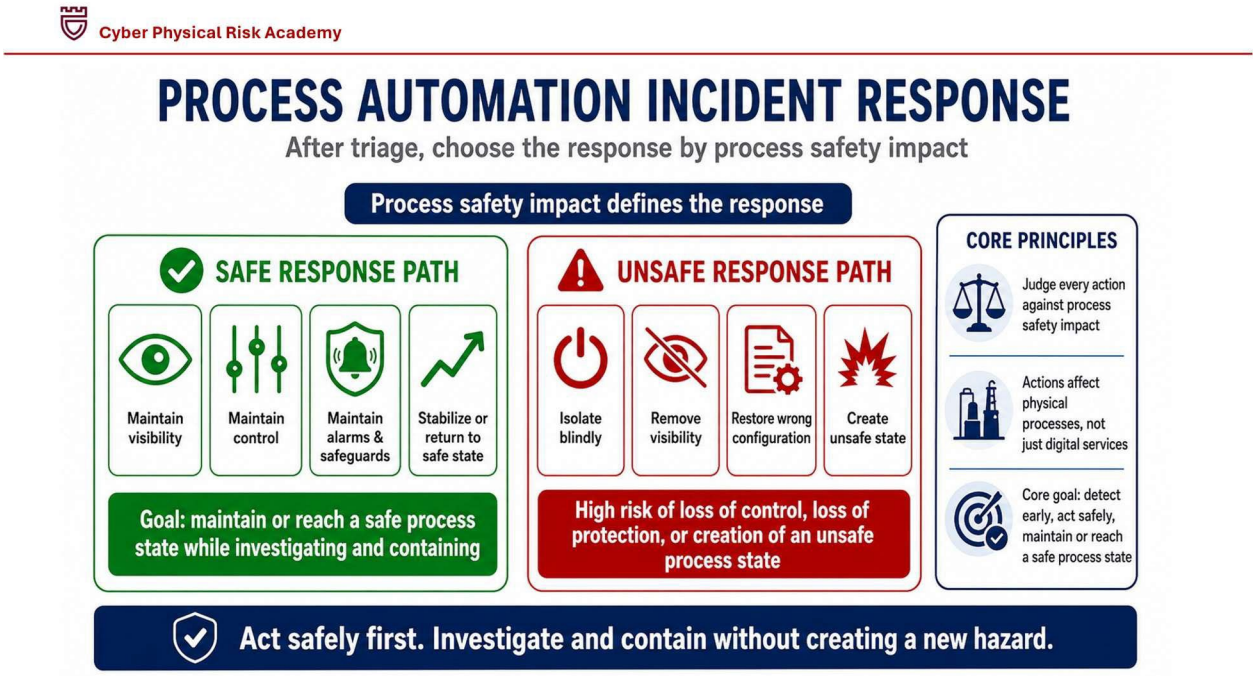

In IT incident response, a technically clean system may be a valid recovery objective. In process automation, that is not sufficient. A restored system can be cyber-clean and still be operationally wrong.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

The key question is whether the restored digital automation state still matches the actual physical process state.

Batch processes make automation-state misalignment easy to see because sequence state, recipe phase, material state, and equipment status must remain synchronized. A batch may have advanced, stopped, partially completed, or deviated while the automation system was impaired. If the system is restored to an earlier recipe step, wrong phase, stale permissive state, or outdated equipment status, the automation may issue actions that no longer fit the actual vessel condition.

This is where Recovery Point Objective (RPO) must be interpreted in process terms. It is not only about how much digital data loss is acceptable. It is also about how much operational state history can be lost before confidence in the actual process state is lost. That includes sequence state, batch history, operator actions, setpoint changes, mode changes, permissive states, override states, and process data.

Many control systems periodically store snapshots of the operational state, often called checkpoints. These snapshots may be taken every 30 minutes, every hour, or at defined operational milestones. A checkpoint can help restore the automation state closer to the latest known process state, but it is not automatically safe to use. The restored checkpoint must still be validated against the actual vessel condition, equipment status, material state, and any manual interventions that occurred during the incident.

For some processes, the safer recovery strategy is not to restart from the latest checkpoint. It may be better to return the process automation to a predefined safe initial state that has been recorded and validated specifically for restart or reconstitution. This choice should not be improvised during the incident. It should be defined in the disaster recovery plan and linked to the process safety basis, batch strategy, operating procedures, and reconstitution criteria.

The same principle applies to continuous processes. Stored controller parameters, setpoints, alarm limits, override states, equipment modes, constraints, or advanced control outputs may no longer be valid after a disturbance, manual intervention, partial shutdown, communication loss, or process drift. Here, Recovery Time Objective (RTO) also needs a process interpretation. It is not only about how quickly a server, workstation, controller, or historian is technically restored. It is about how quickly the minimum trusted automation capability must be available to support safe stabilization, verification, and process reconstitution.

In both batch and continuous processes, recovery is incomplete until the restored automation state is validated against the actual physical process state. Restoring the automation function from a clean backup does not prove that its assumptions still fit the live process.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

So recovery is not only a cybersecurity question:

“Is the system clean?”

It is also an operational and process safety question:

“Is the restored automation state still aligned with the actual process state?”

A clean system is useful, but it is not sufficient. It is not a success if it restores control logic, parameters, sequence states, or operating assumptions that no longer match the plant.

That is why process automation incident response needs a different decision hierarchy. Containment, evidence preservation, eradication, and system recovery remain important, but they must operate within the constraints set by the live process. They do not outrank the need to maintain a safe and valid relationship between the automation function and the physical process.

A safe response path preserves the ability to observe, control, alarm, safeguard, stabilize, and verify the process. It contains the attacker without removing the functions needed to keep the process safe.

An unsafe response path does the opposite. It isolates blindly, reboots blindly, removes operator visibility, disables alarms, restores stale states, restarts from the wrong batch phase or control condition, or creates an automation state that is clean but no longer valid for the process.

The difference is not whether the action is “cybersecurity correct.” The difference is whether the action is safe and valid for the live process.

The practical rule is simple:

Act safely first. Investigate, contain, eradicate, and recover without creating a new hazard or restoring an invalid automation state.

That does not mean doing nothing. It means response actions must be deliberate, process-aware, and aligned with process safety.

- Before isolating systems, determine what visibility, control, alarm, and safeguard functions depend on those systems.
- Before rebooting, determine whether the process can tolerate loss of that function at that moment.
- Before restoring a configuration, determine whether it matches the actual physical state of the plant.
- Before restarting a batch sequence, determine whether the recipe step, material state, equipment state, permissives, and process conditions still match.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

- Before reloading controller parameters, determine whether the stored values are still valid for the current operating condition.
- Before disabling noisy alarms, determine whether the noise is the symptom of the attack, the process deviation, or the loss of signal coherence.
- Before declaring success, determine whether the process is safe, controllable, observable, and operating within acceptable limits.

In process automation, incident response is not only about removing the attacker. It is also about avoiding a second incident caused by the response.

Three timelines run in parallel

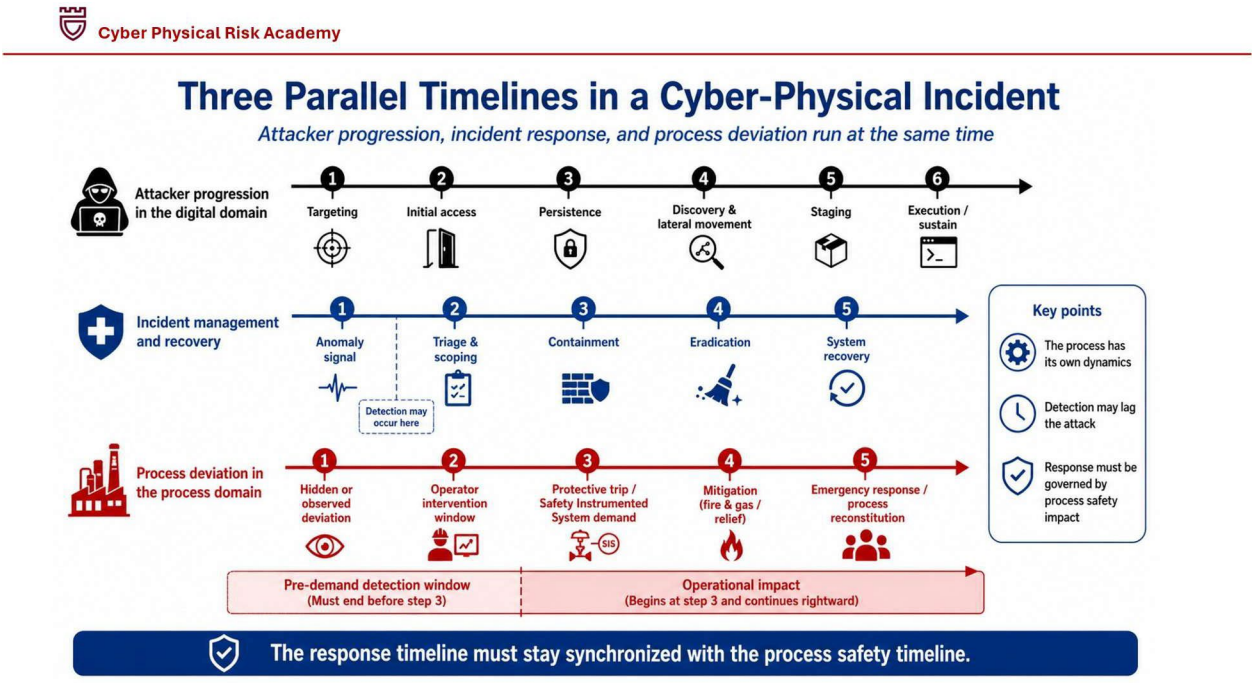

The previous sections explained why individual response actions must be judged against process safety impact. The next step is to make the timing explicit.

A cyber-physical incident does not follow a single timeline. At least three timelines run in parallel:

- attacker progression in the digital domain;
- incident management and recovery;
- process deviation in the process domain.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

The first timeline is **attacker progression in the digital domain**. This is the attacker's path through the digital environment toward an action that can affect the process. We have the following steps:

1. **Target selection** means the attacker decides what to influence. This may be operator supervision, process control, sequence control, alarm handling, safeguard logic, engineering authorization, diagnostic monitoring, maintenance intervention, field measurement, actuator behavior, process analysis, or another function that provides useful leverage over the process. Field-level examples include smart sensors, valve positioners, process analyzers, intelligent drives, or other equipment functions that affect what is measured, how actuators respond, or how the process condition is interpreted.
2. **Initial access** means the attacker obtains a first usable entry point. This may involve stolen credentials, remote access misuse, compromised engineering laptops, exposed services, infected portable media, supplier access, or another path into the digital environment.
3. **Persistence** means the attacker creates a way to remain present. The purpose is to survive disconnects, reboots, password changes, or partial cleanup actions. Persistence may be technical, such as malware or hidden accounts, but it may also rely on legitimate tools, scheduled tasks, configuration changes, or abused remote access mechanisms.
4. **Discovery and lateral movement.**

   Discovery means the attacker learns the environment. In process automation this is not only network discovery. It also includes learning which systems perform which automation functions, which controllers manage which equipment, where engineering tools are located, which displays operators use, which alarms matter, and which communication paths connect the automation functions.

   Lateral movement means the attacker moves from the first compromised function toward functions with higher control proximity. **Control proximity** expresses how close a compromised function is to causing physical process impact. The attacker may move from supporting functions toward functions such as operator supervision, process control, sequence control, alarm handling, safeguard logic, advanced control, optimization, diagnostic monitoring, or maintenance intervention. These functions matter because their compromise can more directly change process behavior, mislead operators, inhibit intervention, suppress safeguards, or accelerate escalation.

5. **Staging** means preparing the attack before the intended process effect occurs. In process automation this may include positioning tools, preparing scripts, preparing a sequence of commands, changing parameters for later use, weakening alarms, preparing false displays, creating unauthorized engineering access, or sabotaging an automation or safeguard function so that it fails when later demanded.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

This last case is critical. A staged modification does not always need a separate activation step. The malicious change may already be present, but remain unnoticed because the affected function is not continuously exercised. For example, a safeguard function may only reveal the sabotage when a process demand occurs.

Staging is important because it may still look like preparation in the digital domain, while it already determines how the physical process will be affected later. It can also determine whether the attacker can sustain the effect, inhibit operator intervention, suppress safeguards, or amplify escalation once the process starts to deviate.

6. **Execution** means the attacker triggers or allows the prepared action to affect the automation function or the process. This may involve changing setpoints, forcing outputs, modifying logic, suppressing alarms, blocking communication, manipulating displayed values, disabling safeguards, or creating conditions that move the process toward an unsafe or invalid state.

The second timeline is **incident management and recovery**. This is the organizational response path. It starts only when something becomes observable: a security alert, an automation anomaly, an operator observation, or another signal that triggers investigation.

1. An **anomaly signal** is an indication that something may be wrong. It may come from security monitoring, automation diagnostics, operator observation, inconsistent process behavior, failed communication, unusual engineering activity, or a mismatch between expected and observed process behavior.
2. **Triage and scoping** determine whether the anomaly is an incident, what is affected, what is still trustworthy, what could become unsafe, and which response actions are acceptable. In process automation, triage must include cybersecurity, operations, and process safety because the response may affect the live process.
3. **Containment** limits the attacker's ability to continue, spread, or influence the process. In IT this often means isolation, blocking, disabling accounts, or shutting systems down. In process automation, containment must be constrained by process safety. Blind isolation may remove visibility, control, alarms, safeguards, or engineering support that is still needed to maintain or reach a safe process state.
4. **Eradication** removes attacker access, malware, persistence, unauthorized changes, compromised accounts, or manipulated configuration. In process automation, eradication must be coordinated with operations because removing or restoring digital elements can change the automation state, available functions, or operator view.
5. **System recovery** restores trusted digital functions. This may include restoring systems, accounts, networks, configurations, backups, engineering tools, communication paths, monitoring capabilities, and where applicable, recent operational checkpoints. But system

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

recovery is not the same as process reconstitution. A clean digital system, backup, or checkpoint may still contain logic, parameters, sequence states, or assumptions that no longer match the actual process state.

The third timeline is process deviation in the process domain. This is the physical timeline. It is governed by process dynamics, not by ticket queues, forensic procedures, or malware analysis.

1. A **hidden deviation** is a physical process deviation that has already started but is not yet visible, recognized, or correctly interpreted as abnormal. The deviation belongs to the process domain, not to the digital detection domain. It may remain hidden because the physical change is slow, because it is still inside apparent normal variation, because available measurements do not expose it clearly, or because displayed values, alarms, or trends are manipulated.
2. An **observed deviation** is a physical process deviation that has become visible or recognizable through process behavior, measurements, alarms, diagnostics, equipment response, or operator observation.
3. The **operator intervention window** is the period in which operators may still be able to stabilize the process through manual or operator-initiated actions before a protective demand should occur. This may include reducing energy input, stopping feed, changing operating mode, initiating a controlled shutdown, requesting field verification, or executing emergency operating procedures. The phrase “should occur” is intentional: the process may reach the condition where protection is required even if the protective function has been bypassed, inhibited, modified, or sabotaged.
4. A **protective trip or Safety Instrumented System (SIS) demand** should occur when the process reaches a condition where a protective function is required to prevent escalation toward a hazardous event. Such a trip may be correct, necessary, and safe, but it is already an operational event. It may stop production, isolate equipment, depressurize a section, trigger downstream effects, or force the process into another operating mode. If the protective function has been compromised, the process may pass this point without the expected protective action.
5. **Mitigation and emergency response window** This is the phase in the process timeline where preventive control is no longer sufficient or has failed. The process has moved beyond normal control and beyond the point where preventive action should have kept it inside its safe limits. The response objective now shifts from preventing escalation to limiting consequences.
   Mitigation limits the effects of an escalating or realized hazardous condition. Emergency response protects people, the environment, and the installation once the situation can no longer be handled only as a controlled process intervention.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

6. **Process reconstitution** is the controlled return of the process and its automation functions to a safe, valid, and operationally meaningful state. It is not the same as digital recovery. It requires confirming that the physical process state, equipment status, control modes, alarms, safeguards, parameters, recipes, permissives, and automation logic are again aligned.

The important point is that these three timelines do not wait for each other.

The attacker may already be far into the attack before the first anomaly becomes visible. The incident response team may still be qualifying the event while the process is moving closer to a protective demand. The operator may still have an intervention window, or that window may already be closing.

That is the technical reason why process automation incident response must be synchronized with process safety. The incident management timeline must not be managed as an isolated digital workstream. It must be coordinated with the physical process timeline.

The key question becomes:

Can we respond fast enough to preserve the operator intervention window, but carefully enough not to create a new hazard or invalid automation state?

That is the central tension in process automation incident response.

Detection is only valuable if it is early enough

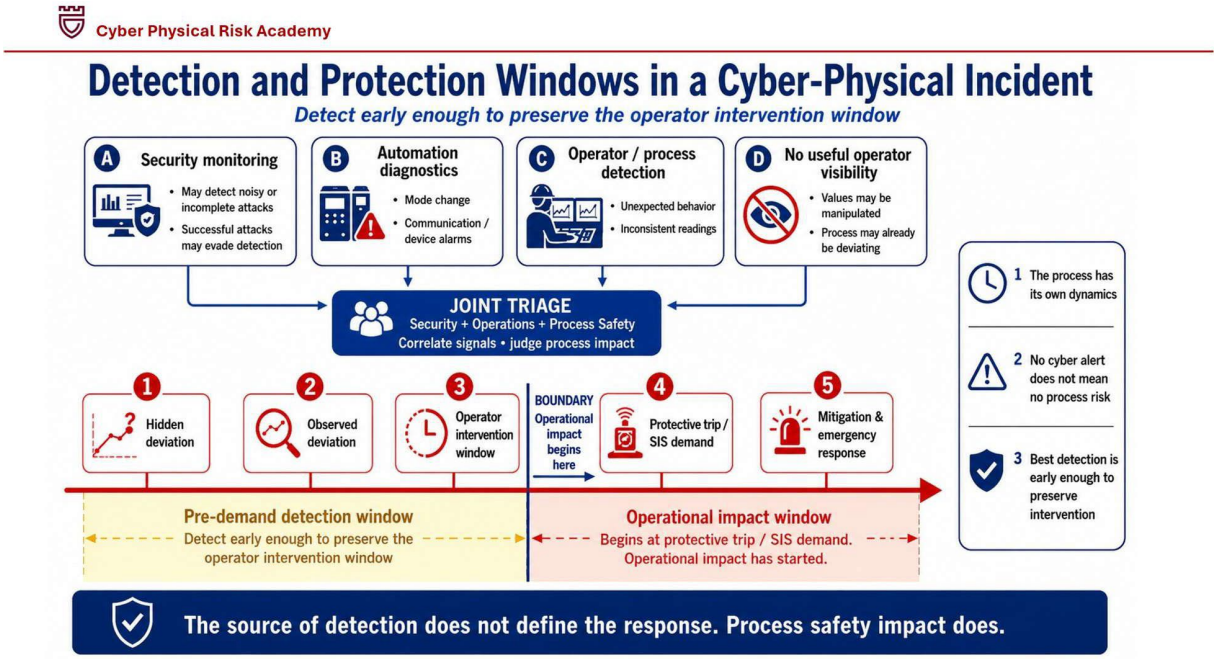

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

Detection is important, but detection alone is not enough. In process automation, the value of detection depends partly on whether it occurs early enough to preserve meaningful intervention options. A cyber alert before the point where a protective demand should occur can support prevention, stabilization, and preservation of the operator intervention window.

If the alert arrives after process impact has started, it is still important. It may no longer be pre-demand detection, but it can help determine whether the threat actor is still active, whether further escalation is possible, and which containment actions are safe under the current process conditions. The attacker may also change behavior once discovery becomes likely or confirmed. This may include removing forensic evidence, changing access paths, accelerating manipulation, inhibiting operator response, or attempting to increase damage before containment is effective. The objective then shifts from preventing initial impact to stopping further manipulation, limiting escalation, supporting safe stabilization, preserving evidence where this is safe, and preparing for controlled recovery.

The **pre-demand detection window** is the period before the process reaches the condition where a protective demand should occur. This window is valuable because it allows intervention before the process depends on a protective trip or SIS demand.

A trip is not just another detection point. A trip is an operational event. It may be correct, necessary, and safe, but it already affects operation. It may stop production, isolate equipment, depressurize a section, trigger downstream effects, or start a different operating mode.

From that point onward, the incident has entered the **operational impact window**. The operational impact window is the period in which process impact is no longer only potential. Operation has already been affected, or must now be affected, to prevent further escalation. This may be the result of a protective trip, a SIS demand, a controlled shutdown, a mitigation action, emergency response, or required process stabilization.

Before the incident reaches mitigation or emergency response, the organization may still have an opportunity to intervene. But no single detection signal is sufficient by itself.

A Security Operations Center (SOC) may detect noisy or incomplete attacks, but a successful attack may evade conventional cyber monitoring. The automation system may report diagnostics, but those diagnostics may only show symptoms. The operator may see unexpected behavior, but may also be blind if sensor values, alarms, trends, or displays are manipulated.

Therefore, detection needs joint triage. Security, process operations, process automation engineering, and process safety must correlate the available signals and answer two practical questions:

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

What does this mean for the live process?
And:

Do we still have a safe intervention window, or has the incident already moved into operational impact?

If a safe intervention window still exists, the response should focus on preserving visibility, control, alarms, safeguards, and operator decision space. The objective is to stabilize the process before it depends on protective action or mitigation, while security determines whether the threat actor is still active and what containment actions are safe under the current process conditions.

If the incident has already moved into operational impact, the priority shifts, but security does not disappear. Process operations and process safety must protect people, stabilize the process, verify whether protective functions have acted as expected, and determine whether mitigation or emergency response is required. At the same time, security and process automation engineering must determine whether the attacker is still active, whether further manipulation is possible, whether evidence can be preserved safely, and which containment actions can be taken without reducing visibility, control, safeguards, or operator intervention.

So the distinction is not “before impact equals security response” and “after impact equals operations response.” The distinction is the response priority. Before impact, the objective is to preserve intervention. After impact, the objective is to protect people, limit escalation, stabilize the process, stop further manipulation, and prepare for controlled system recovery and process reconstitution.

Therefore, we say that the source of detection does not define the response. **Process safety impact does.**

1. **Mitigation and emergency response**

   Mitigation limits the consequences after the process has moved beyond normal control or after prevention has failed. Examples include pressure relief, blowdown, fire and gas actions, deluge systems, ventilation, containment, inventory isolation, or spill control. A deluge system is a fixed fire protection system that rapidly releases a suppression or cooling medium, most commonly water, over equipment or a process area. Depending on the hazard, the medium may also be foam-water, water spray, water mist, or another suitable agent. Its purpose is to cool surfaces, limit fire escalation, reduce heat exposure, suppress vapor formation, or control the consequences of a release or fire. These actions do not restore normal operation. They reduce harm.

   Emergency response is the organizational response when an incident may threaten people, the environment, the process installation, or the continuity of safe operations. In the

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

diagram, it starts immediately after the initiating event and partly overlaps with activation, notification, crisis management, recovery, and later process reconstitution.

Emergency response is not a single action. It is a structured response process with overlapping phases.

Incident Response processes in Process Automation

Above the timeline, the figure in the next slide shows the organizational response: business continuity management, including emergency response, crisis management, and business recovery. Below the timeline, it shows the technical and operational recovery side: activation of the response organization, restoration of digital systems, reconstitution of the process, and the return to normal operations.

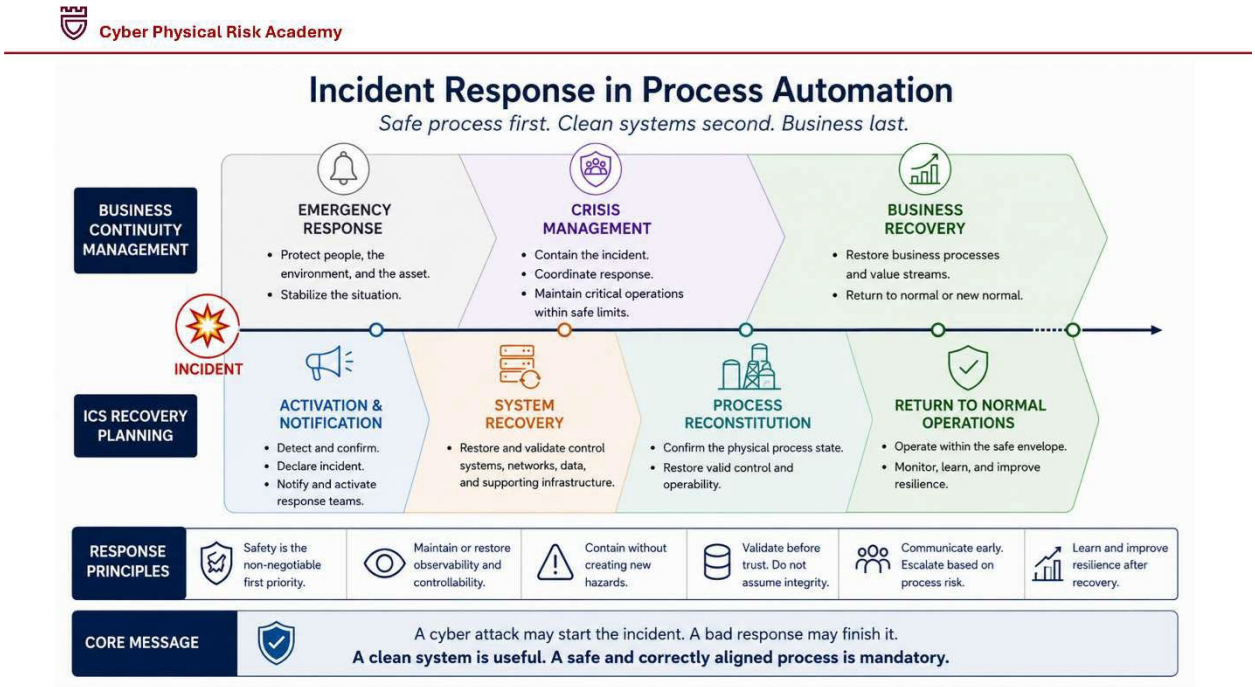

Activation and notification should be separated from site-wide emergency alerting.

Site-wide emergency alerting is the warning and instruction function for people who may be exposed to the hazard. In many process installations, the **Public Address and General Alarm (PAGA)** system is the primary site-wide warning function. It provides audible alarms, visual alarms, and, where configured, spoken instructions for evacuation, muster, shelter-in-

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

place, or other emergency actions. Its purpose is to protect people by making sure they receive clear and timely instructions.

Activation and notification, in the incident response sense, is the **mobilization of the response organization**. After initial triage, or when escalation criteria are met, the relevant roles must be notified and activated. This may include security personnel, operations, process safety, emergency response leadership, site management, external specialists, regulators, or public emergency services where required. Its purpose is to make sure the people who must assess, decide, coordinate, and act are engaged quickly enough to control escalation.

The distinction matters. PAGA tells exposed site personnel what to do. Activation and notification brings the response organization into action. Both may occur close together in time, but they serve different purposes and must not be confused. They can also collide operationally. External responders or specialists may be mobilized, but site access may be restricted at the gate during an emergency. If emergency management does not explicitly authorize their entry, critical expertise can remain outside the fence while decisions are being made inside the control room. Activation and notification must therefore include not only who is called, but also who authorizes access, how the gate is informed, and where external staff report once admitted.

It then moves into **immediate response actions**. In a cyber-physical incident, these actions must run in parallel across two connected workstreams.

The **security team** starts its initial activities: confirming the alert, preserving volatile evidence where possible, identifying affected accounts, systems, communication paths, remote access sessions, engineering changes, malware indicators, or unauthorized activity. These activities are constrained by **process operations** and **process safety**. Process operations has the lead because it is responsible for the live process state and the immediate safety of personnel on site. Process safety supports the assessment of hazards, safeguards, escalation potential, and safe-state requirements. Security supports by determining what is happening in the digital and automation environment.

The security team must therefore not isolate, reboot, block traffic, disable accounts, or remove systems without operational authorization and without understanding whether those actions affect visibility, control, alarms, safeguards, or operator intervention.

At the same time, **process operations and process safety personnel** focus first on the safety of people on site. Their immediate concern is not root cause. Their first concern is whether personnel are exposed to a hazardous condition, whether evacuation or muster is needed, whether firefighting, rescue, medical response, environmental response, emergency

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

shutdown coordination, or external emergency services are required, and whether the process is moving toward further escalation.

These workstreams must therefore operate hand in hand, but they are not equal in authority during emergency response. **Process operations has the lead** because it is responsible for the live process state, the safety of personnel on site, and the decision to maintain, stabilize, shut down, or move the process to a safe state. Process safety supports operations by assessing hazards, escalation potential, safeguard status, safe-state requirements, and possible consequences of response actions. Security supports operations by determining what is happening in the digital and process automation environment and by containing the attacker under operational constraints.

A Security Operations Center cannot lead emergency response for an industrial site. A SOC may detect the first signal, provide critical evidence, and support digital containment, but it does not own the process, the personnel exposure, the operating envelope, the safeguards, or the emergency response decisions on site. During the immediate response phase, all disciplines must support the operational command structure led by process operations.

Neither view is sufficient alone. A technically correct security action can be unsafe for the process, and a normal operational response can be wrong if the automation information is compromised. The primary objective during this phase is to protect life, limit environmental impact, prevent escalation, and keep or bring the process into a safe state. Digital containment remains important, but it does not outrank personnel safety or process safety.

After the immediate threat is controlled, the focus shifts toward **system recovery** and **process reconstitution**.

System recovery restores the digital and process automation functions needed to support safe operation. In a cyber-physical incident, this must be a **staged activity**, not an uncontrolled technical rebuild. The recovery sequence must be defined in a disaster recovery plan and agreed with process operations, automation engineering, process safety, and security.

The first system recovery objective is to make available the **minimum set of trusted functions** that process operations needs to perform its process reconstitution task as soon as possible. This may include a limited number of operator consoles, selected controller functions restored unit by unit, engineering access under strict control, historian or trend access where needed, and safety-related functions.

The term **safety-related functions** needs careful interpretation. Depending on the criticality of the process and the site safety philosophy, some safety functions may be mandatory before the process is allowed to start or restart. For example, a SIS or specific Safety

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

Instrumented Functions (SIFs) may be required to be active, tested, and available before feed, heat, pressure, rotation, or other hazardous process conditions are introduced. In that case, the process must not be restarted until those functions are confirmed to be in their required state.

In other cases, a safety-related function may be allowed to remain offline for a short, explicitly approved time window, but only under tightly controlled conditions. This normally requires a documented risk assessment, compensating measures, operational restrictions, defined authorization, increased monitoring, and a clear time limit for restoring the function. Such a condition is not “normal operation.” It is a closely guarded degraded mode.

Once this minimum operating capability is restored and validated, additional functions can be brought back in a controlled sequence, such as extra operator stations, engineering workstations, advanced control, reporting, maintenance interfaces, historians, remote access, or production management interfaces.

The important point is that system recovery must follow the safety requirements of the process, not the convenience of the digital restoration sequence. Some functions are useful for efficiency. Some are useful for analysis. Some are mandatory for safe operation. Those distinctions must be defined before the incident, not improvised during recovery.

This sequencing matters. Restoring a non-essential function too early can increase exposure, reintroduce compromise, overload the response team, or distract from stabilizing the process. Restoring a critical function too late can delay operator action and prolong unsafe or unstable conditions.

Process reconstitution goes further than system recovery. It is the controlled return of the process, the process automation functions, the safeguards, and the operating organization to a safe, valid, and operationally meaningful state.

Process reconstitution can only start when two conditions are met.

First, **process automation engineering must release the relevant system scope for operational use**. This release point is a distinct decision in the recovery process. It means that the restored system, or the agreed part of the restored system, is considered sufficiently trusted, stable, and functional for process operations to start using it again.

Second, **process operations and process safety must jointly confirm that process reconstitution can be performed safely under the defined conditions.** Operations confirms that the live process state and operating organization are ready: equipment status, process conditions, operating mode, staffing, field verification, procedures, operator workload, and available control room functions. Process safety confirms that the safety basis remains

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

acceptable: safeguards, alarms, bypasses, operating limits, degraded modes, compensating measures, and safe-state assumptions. Together, they decide whether the intended reconstitution step is safe to execute.

This release and acknowledgement may apply to a **partial scope**. That scope can be partial from a process automation perspective and from a process perspective.

From a **process automation perspective**, disaster recovery planning may define that process operations can start reconstitution when a minimum set of trusted functions is available. This may include selected controller functions, a limited number of operator consoles and operator stations within the console, essential alarms, controlled engineering access to allow for modifications requiring an engineer level authorization, required safety-related functions, necessary trend or historian data, and visual information from cameras where this is needed to verify field conditions, equipment status, access restrictions, leaks, smoke, fire, flooding, flare activity, or the presence of personnel in affected areas.

Examples include camera views for flare monitoring, tank farm observation, loading areas, compressor buildings, process unit access points, emergency response routes, and areas where field verification may be unsafe or delayed.

From a **process perspective**, the release may apply only to a defined process unit, train, section, package, or equipment group. For example, one utility system, one storage area, one compressor train, one reactor section, or one distillation unit may be released for reconstitution while other parts of the installation remain unavailable, isolated, shut down, or under investigation.

In that case, process reconstitution starts **unit by unit, function by function, or both**, under defined operational restrictions. The release boundary must be explicit: which automation functions are trusted, which process units may use them, which interfaces remain unavailable, which safeguards are active, which compensating measures apply, and which actions are still prohibited.

In other cases, disaster recovery planning may explicitly state that process reconstitution is not allowed until the full operational scope is available. That may be required when the process is highly interdependent, when batch sequence integrity depends on the full automation state, when one process unit cannot be safely operated without upstream or downstream functions, when operator workload would become unacceptable with only partial functionality, or when the site safety case assumes full availability of specific automation and safeguard functions before restart.

In a cyber-physical incident, this distinction is critical. A digital function may appear technically recovered before the physical process is ready to use it safely. Likewise, one

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

process unit may appear ready while its dependencies are not. Process values, controller states, sequence states, safety functions, bypasses, alarms, procedures, field equipment, communication paths, inter-unit dependencies, and operator trust in the system all need to be validated before restart or return to normal operation.

That is why emergency response must not stop at restoring systems. Emergency response is the first response to protect people, limit escalation, and stabilize the site. But as the incident develops, emergency response gradually steps into the wider **crisis management** process.

Crisis management has a broader responsibility. It coordinates the overall organizational response, sets priorities, allocates additional resources, manages internal and external communication, involves authorities or regulators where required, coordinates with corporate management, and prepares for business recovery. Emergency response deals with the immediate threat. Crisis management deals with the wider consequences and the sustained response.

Within this wider process, recovery must still follow a sequence that supports process operations first. Process operations should not start process reconstitution until process automation engineering has released the relevant automation scope for use, and process operations and process safety have jointly confirmed that the relevant process scope can be reconstituted safely.

Reconstitution must then confirm that the restored automation state is aligned with the actual physical process state for the specific unit, train, section, or installation scope being returned. Only then can that scope be considered demonstrably safe, controllable, observable, and ready for operation.

Business recovery starts from a different question. Process reconstitution asks whether the installation, the process automation functions, the safeguards, and the operating organization are ready to return to a safe and valid operating state. Business recovery asks whether the organization can restore the business function that depends on that operation.

In the diagram, business recovery sits above the technical and operational recovery activities because it has a wider organizational scope. It includes customer commitments, supply chain impact, product availability, contractual obligations, financial exposure, regulatory reporting, insurance, public communication, reputation management, staffing, logistics, and prioritization of production capacity. It may also include decisions about running at reduced capacity, restarting only selected process units, using alternative production sites, rerouting feedstock or product flows, delaying shipments, or accepting temporary loss of efficiency to preserve safe operation.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

Business recovery must not force process recovery. A business need to restart production does not prove that the process is ready, that the automation state is valid, or that safeguards are available. The sequence must remain disciplined: emergency response protects people and stabilizes the site; system recovery restores trusted technical capability; process reconstitution confirms that the physical process and process automation state are aligned; business recovery then determines how the restored capability is used to resume business commitments.

This is where crisis management becomes essential. Crisis management must balance business pressure against process safety constraints. It may decide what products, units, customers, or supply obligations receive priority, but it should not override the technical release by process automation engineering or the safety confirmation by process operations and process safety. Business recovery can only proceed within the safe operating envelope that has been re-established.

In a cyber-physical incident, this distinction prevents a second incident. Restarting too early to satisfy business pressure can reintroduce unsafe operating conditions, overload operators, bypass incomplete recovery steps, or reconnect compromised interfaces. Business recovery is therefore not the same as “getting production back.” It is the controlled restoration of business capability after the site has regained a safe, trusted, and operationally meaningful basis for operation.

Digital recovery is not process reconstitution

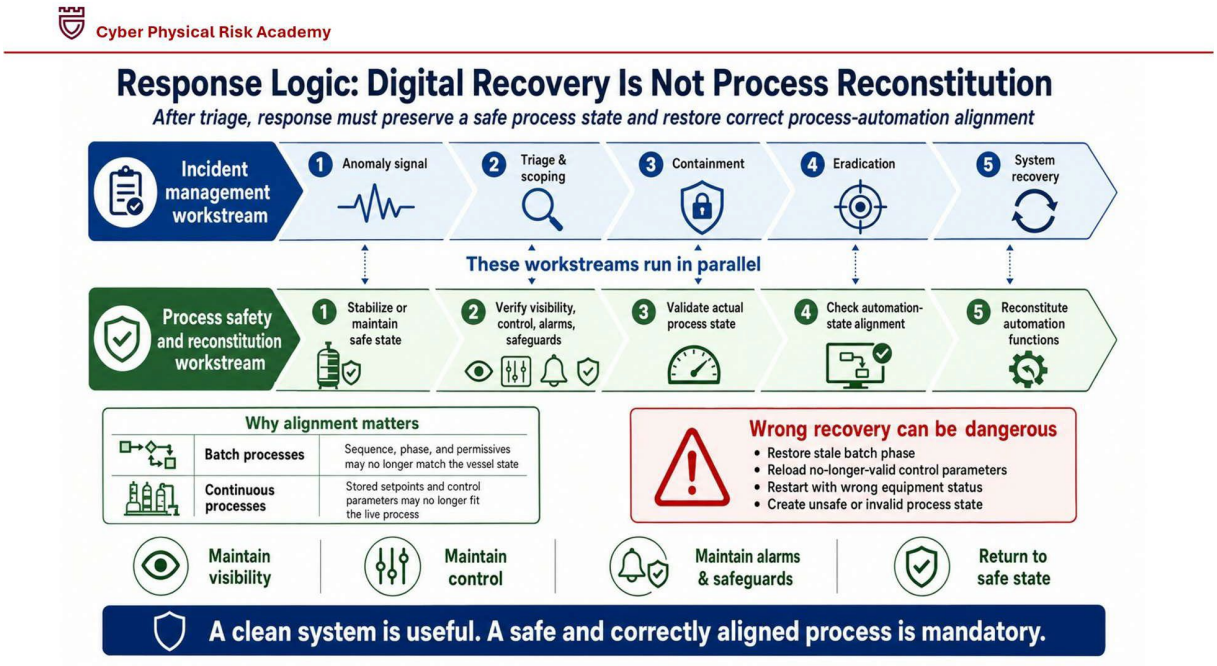

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

This is the final takeaway from the response logic. The final mistake is to assume that digital recovery means process recovery.

It does not.

Digital recovery restores trusted digital and process automation functions: systems, accounts, configurations, networks, backups, engineering tools, controller logic, operator consoles, communication paths, and supporting services.

Process reconstitution asks a different question: can the restored automation state be safely used for the actual physical process state?

That question cannot be answered by cybersecurity alone. Process automation engineering must confirm that the relevant automation scope is trusted and fit for operational use. Process operations and process safety must confirm that the relevant process scope can be reconstituted safely under the defined conditions.

The central issue is alignment.

In batch processes, the batch may have advanced, stopped, partially completed, deviated, or been manually intervened during the incident. Restoring an earlier recipe step, wrong phase, stale permissive state, or outdated equipment status can cause the automation system to act on assumptions that no longer match the material, temperature, pressure, level, or reaction condition in the vessel.

In continuous processes, stored setpoints, controller modes, alarm limits, override states, advanced control outputs, or operating parameters may no longer fit the actual process after a disturbance, partial shutdown, manual intervention, communication loss, or process drift.

A clean system can still be operationally wrong.

That is why the final recovery question is not:

“Is the system clean?”

It is:

“Is the system clean, trusted, and still valid for the actual process state?”

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

Until that answer is yes, recovery is incomplete.

A clean system is useful. A safe and correctly aligned process is mandatory.

Conclusion

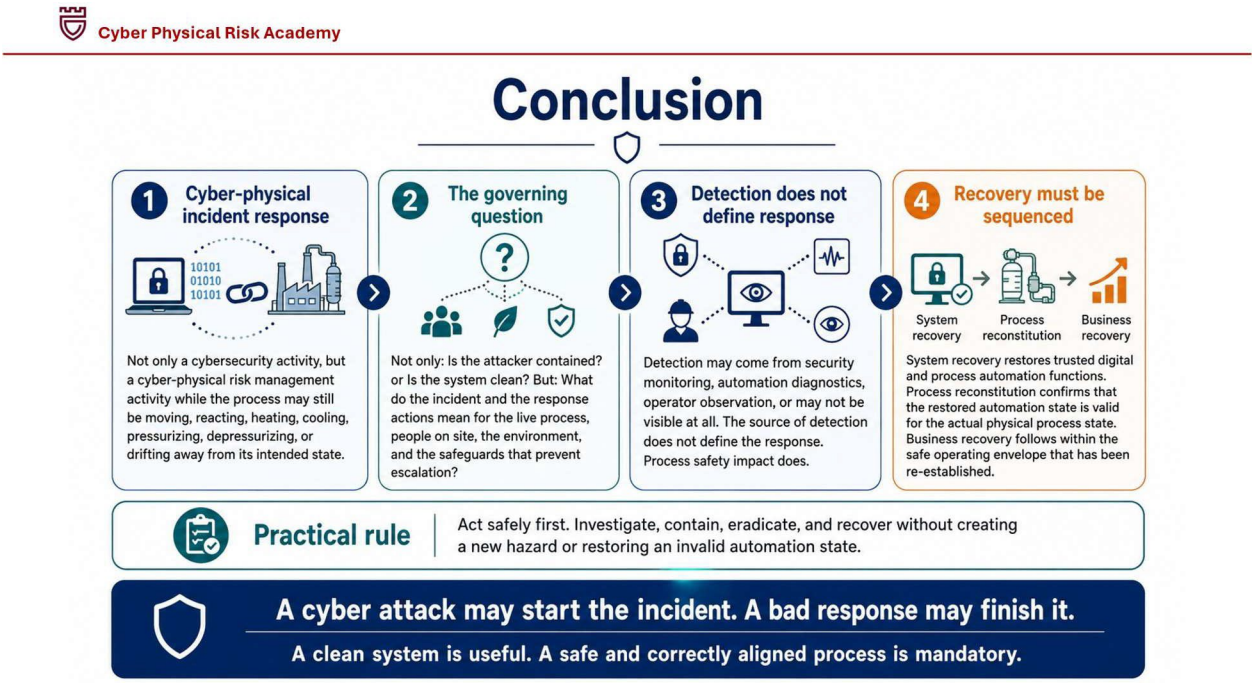

Incident response in process automation is not only a cybersecurity activity. It is a cyber-physical risk management activity carried out while the process may still be moving, reacting, heating, cooling, pressurizing, depressurizing, or drifting away from its intended state.

That changes the response logic. The first question is not only whether the attacker has been contained or whether the system is clean. The first question is what the incident and the response actions mean for the live process, the people on site, the environment, and the safeguards that prevent escalation.

Detection may come from security monitoring, automation diagnostics, operator observation, or it may not be visible at all. The source of detection does not define the response. Process safety impact does.

Recovery must therefore be sequenced differently. System recovery restores trusted digital and process automation functions. Process reconstitution confirms that the restored automation state is valid for the actual physical process state. Business recovery can only follow within the safe operating envelope that has been re-established.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.

The practical rule is simple: act safely first. Investigate, contain, eradicate, and recover without creating a new hazard or restoring an invalid automation state.

A cyber-attack may start the incident. A bad response may finish it.

A clean system is useful. A safe and correctly aligned process is mandatory.

Copyright © 2026 Sinclair Koelemij. All rights reserved. Text, diagrams, slide concepts, and original visual material are protected to the extent legally permitted. Third-party material, trademarks, standards terminology, and referenced frameworks remain the property of their respective rights holders.