


Analysis of Gartner Maverick Insights note released in June 2026 in relation to exercise effectiveness

Why cyber resilience must move from discussion-based rehearsal to connected, measurable validation across IT and OT

STAFFORD, VA, UNITED STATES, October 5, 2026 /EINPresswire.com/ -- Sheffield (UK), August 2026 – Cyber exercises are everywhere. Boards are convening tabletop sessions. National and sector-wide drills are expanding. Red and blue teams are being put under pressure. [Crisis simulations](#) are appearing in resilience programmes that did not exist a few years ago. It reflects an overdue recognition that resilience cannot be established by a policy, a compliance map or a maturity score alone.



As CISO at global [cyber range](https://cyberranges.com) vendor CYBER RANGES (<https://cyberranges.com>) and a cyber technology consulting veteran for both vendors and end-user organisations, Mr. Ed Bisceanu addresses here the matter of exercise effectiveness in response of a provocative Gartner Maverick research paper on this topic released in June 2026.

A graduate of the National Intelligence Academy and Military Communications Institute, and a former Col. Signals (ret.), Mr. Bisceanu has made significant contributions to the development of national cybersecurity programs and led the operationalization phase of the first civilian government CERT team in his country. Eduard has worked for several renowned end-user organizations, e.g. UniCredit Bank, Microsoft CE Europe, PwC, ProActive Defense, CEC National Bank.

Ed comments that “what makes the Gartner Maverick Insights note released in June 2026 so provocative – and so useful – is the concern that If every exercise succeeds, the organization may be rehearsing confidence rather than testing resilience.

“The provocation is about evidence, not vocabulary. Cyber resilience is not necessarily a new

discipline beyond cybersecurity. It is the observable result of cybersecurity, cyber-risk management, continuity and recovery executed properly. However, a separate resilience programme can become a fresh label for old work – segmentation, identity protection, recovery ownership, tested backups and threat-informed prioritization – which was already understood, funded or deferred. An organization does not manufacture resilience in the middle of an attack. The attack reveals decisions made much earlier. The Gartner note's most valuable challenge is therefore not semantic", states Ed Biscanu.

"It is evidential. Too many resilience tests are institutionally comfortable because meaningful failure carries a price." If an exercise shows that recovery does not work, a risk owner must act, a budget must move or someone must formally accept the exposure. A reassuring exercise produces a board update. An honest one may force a decision. That is the difference between resilience work and resilience theatre. An exercise can be useful without being assurance.

Ed observes that the word 'exercise' describes an event, not the strength of the evidence it produces. Before judging an exercise, organizations should be explicit about its purpose.

A learning exercise builds understanding. Participants may be briefed, coached and given time to prepare. Success means they leave better able to perform their roles.

A rehearsal tests coordination: authority, escalation, communications, hand-offs, dependencies and the sequence of decisions. Success means the organization can execute a known process coherently.

An assurance exercise asks whether a defined capability meets explicit criteria under specified conditions. Success must be supported by observable behaviour and evidence, not only by participant confidence.

An adversarial validation exercise deliberately challenges assumptions, allows unexpected paths and treats a consequential failure as a valuable result – provided it leads to remediation and re-testing.

The tabletop nature of an exercise is not the problem. A separate 2026 Gartner infographic describes tabletop exercises as an effective way to test response strategies and offers practical advice on executive participation, scenario choice, business impact, third-party dependencies and after-action learning. Its recommendations make sense for a learning or decision-rehearsal objective: prepare participants, avoid overwhelming executives with technical detail and focus the discussion on decisions that matter.

The same infographic shows where the assurance gap begins. In its survey of 75 CISOs, 47% ranked simulating a real-world incident among their three largest tabletop challenges, while 39% did the same for measuring effectiveness. These figures do not prove why the problem exists, but they identify the difficult territory: realism and measurement. But it has a real ceiling.

A tabletop can test the quality of a decision. It cannot prove that the system will obey it.

The response should not be to discard tabletop exercises. It should be to stop asking them to provide evidence they cannot produce alone. The mature model is not tabletop versus cyber range. It is one consequence-led scenario moving through a connected exercise continuum.

There are several layers of exercise each answering a different question and contributing a different kind of proof:

Executive tabletop: Can leaders set priorities, exercise authority, make business trade-offs and coordinate communications and external dependencies?

Functional or command exercise: Can teams execute procedures and hand-offs while operating with incomplete information, degraded communications and time pressure?

Instrumented cyber-range exercise: Can people, processes, tools and controls detect, investigate, contain and respond to safely emulated adversary activity in a representative environment?

Recovery and controlled operational validation: Can the organization restore the affected service, validate integrity and capacity, and meet the recovery objective under appropriately governed conditions?

Remediation and re-test: Did the failed assumption produce a changed decision or capability – and can the improvement survive the same challenge a second time?

The important design principle is connection. Executive choices should create observable operational consequences. Technical discoveries should force business decisions. Recovery should be performed, not narrated.

The same critical service, threat path and consequence should follow the exercise from the Boardroom into the operational environment and back again. Evidence should accumulate across all the layers: decision records, telemetry, control performance, recovery results, assigned actions and retest outcomes.

Three takeaways:

- 1) match the exercise to the evidence required
- 2) connect executive decisions to technical consequences
- 3) finish only when remediation survives re-test.

Not every tabletop needs to become a large live exercise. The level of validation should be proportional to the consequence being tested. But assumptions protecting a critical service

should not remain permanently protected by discussion.

The Gold Eagle Initiative in the USA makes the same point at national scale. Announced by the White House on 14th July 2026, the initiative is intended to coordinate vulnerability work across federal agencies – including US Treasury, DHS/CISA and the Department of Defense/War (DoD / DoW) – critical-infrastructure operators, open-source maintainers, frontier AI developers and industry partners. This initiative is still at an early stage and its governance and operational effectiveness remain unproven. Its relevance here is in the model, not a claim of success. That model shifts attention from finding more weaknesses to validating exploitability, prioritizing impact, assigning responsibility, coordinating remediation and verifying that the correction works without creating unintended functional effects.

Ed points out that “the parallel for exercises is direct: finding a gap is only the start. The exercise produces resilience evidence only when the weakness is owned, corrected and challenged again. Operational Technology makes the distinction physical, making this argument impossible to treat as an abstract debate. NIST's [OT](#) security guidance stresses that digital and physical risk are intertwined, that safety directly shapes engineering and operational decisions, and that changes should be tested before deployment where operational impact is possible.”

Ed continues: “No responsible operator should turn a live production plant into an unrestricted adversary laboratory. That legitimate constraint explains some reliance on tabletop exercises. It does not justify ending validation at the table. The right answer is graduated fidelity within an explicit safety envelope. NIST makes the bridge explicit: organizations should consider tabletop exercises or simulations when they need to reduce assessment impact on production OT, and they should use automated assessment tools carefully enough to avoid adversely affecting the system.”

An OT tabletop can clarify authority among operations, engineering, functional safety, IT, security, vendors and executive management. It can explore decisions around plant isolation, loss of view or control, remote-access compromise, manual operation, production shutdown and public or regulatory communications. The US Cybersecurity and Infrastructure Security Agency (CISA) has specifically recommended discussion-based exercises involving loss-of-visibility and loss-of-control scenarios in industrial environments.

Selected assumptions can then move into an isolated, representative environment like those that OT WORLDS by CYBER RANGES enables building. Depending on the objective, this may combine virtualized control networks, emulated or simulated PLC, HMI and SCADA behaviour, historians, engineering workstations, remote-access paths, security monitoring and the necessary enterprise dependencies. Where justified, representative hardware or hardware-in-the-loop can add to fidelity. NIST's control catalogue for OT includes incident-response training using simulated events and automated training environments, alongside incident-response testing.”

Mr. Biscéanu advises that: “Fidelity must be declared rather than assumed. A cyber range is not automatically a perfect digital twin, and a visually convincing replica is not necessarily operationally meaningful. The environment should reproduce the dependencies, behaviours and decision consequences relevant to the test – and be candid about what it does not reproduce.”

The exercise must validate the state of the process, the integrity of control logic and configurations, the restoration sequence and the authority for a safe return to service. The useful measures are no longer exercise attendance or the number of injects delivered. They include detection and escalation time, decision latency, containment without creating a less safe condition, ability to sustain authorized degraded operation, integrity of controller logic and engineering configurations, restoration sequence, recovery time and successful retest.

Ed points out that “in OT a technically successful containment action can still be an operational or safety failure. That is precisely why cyber defenders, engineers, operators and safety specialists must exercise together. MITRE ATT&CK for ICS can inform credible adversary behaviours, but operational consequences and safe response criteria must come from the asset owner and engineering context.”

However, Ed warns that “The hard gap is orchestration: tabletops, cyber ranges, recovery tests and OT assurance often sit under different owners, budgets and evidence models. A platform cannot compensate for weak scenario design, insufficient fidelity or reluctance to expose failure.” Thus, orchestration with access and integration – not simply technology availability. OT adds a genuine availability and maturity constraint. Representative testing may require process models, industrial protocols, vendor-specific logic, controllers, safety boundaries and specialized engineering support. A generic training lab will not automatically provide assurance for a particular plant or process.

Ed advises Executive Boards to count corrections, not exercises. Exercise count, attendance, scenario novelty and the length of the after-action report (AAR) are activity metrics. They do not establish resilience. A more useful measure is the time from observed failure to accountable, verified correction. The questions that matter are more demanding:

- Which assumption failed?
- Which decision changed?
- Which attack path or dependency was corrected?
- Did recovery work under pressure?
- Was the issue assigned, funded and remediated?
- Did the correction survive a re-test?

Failure discovered during an exercise is not the opposite of resilience. It is one of the most valuable outputs a resilience programme can produce. Failure found in a safe exercise is still under the organization’s control. Failure found first by an adversary is not.

Ed notes that “organizations need cyber range environments in which failure can be safe, observable and useful.” and concludes that:

A cyber range is not another resilience framework. It is where the framework is forced to meet consequences.

Anthony Munns
CYBER RANGES LTD

This press release can be viewed online at: <https://www.einpresswire.com/article/932455046>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.