CYBERSECURITY / OPERATIONAL RESILIENCE

CyberResilience AI

From security findings to operational readiness
A human-governed platform concept for connected enterprise and critical services

A PRACTICAL, EVIDENCE-LED APPROACH

Jami Kiran Kumar

Independent product white paper · Version 1.0 · November 2026
For security leaders, technology executives and operators of essential services

EXECUTIVE OVERVIEW

Make the path from evidence to action visible.

Cybersecurity teams often hold fragments of the picture: asset lists, vulnerability exports, control evidence, incident tickets, network diagrams and recovery plans. Leaders need those signals connected to business services and translated into a small, accountable set of decisions. CyberResilience AI is designed as a workspace for bringing those records together, highlighting evidence gaps, organizing risk review and tracking resilience improvement.

The central proposition is practical: make the path from authorized scope to validated finding, accountable decision and verified improvement visible and repeatable. Specialist agents can help summarize and prioritize evidence, but permissions, source provenance and human approval remain fundamental. NIST CSF 2.0 provides a useful lifecycle reference across Govern, Identify, Protect, Detect, Respond and Recover [1].

Product status and evidence boundary. The deployed experience is an interactive browser-based demo. It presents modules and example records, stores workspace data locally in the browser, accepts customer-entered asset/evidence records and can import authorized assessment findings from CSV. Its landscape review applies deterministic rules to those records and flags missing evidence, ownership and recorded configuration concerns. Live discovery, external scanning, telemetry collection, production integrations and production AI-agent services are not configured in this demo. Sample dashboard values are illustrative, not measured customer outcomes. This paper describes a product direction; it does not claim to have tested a customer environment, detected an active attack, certified compliance or prevented a breach.

01 / THE OPERATIONAL PROBLEM

Cyber exposure is distributed across services.

Enterprise exposure spans infrastructure, identities, suppliers, operational technology, business services and recovery capabilities. A supermarket may depend on payment terminals, cold-chain monitoring, store networks, identity systems and suppliers. A hospital may depend on clinical systems, medical devices, diagnostics and restoration services. Banks and trading firms rely on tightly coupled identity, connectivity, data and transaction services. Power operators combine enterprise IT with operational technology whose availability and safety constraints change how testing must be performed.

A vulnerability score alone rarely answers the operational questions: Is the affected asset real and in scope? Is the observation current and supported by evidence? Which critical service depends on it? Can exposure be reduced safely? Who owns the decision? How will recovery be verified? A useful program connects technical observations with ownership, service criticality, change controls, detection and recovery readiness.

NIST CSF 2.0 is outcome-oriented and intended for organizations across sectors and technology environments [1,2]. CISA and partner agencies emphasize accurate OT asset inventories as a foundation for prioritizing protection and incident response [3]. EU NIS2 establishes cybersecurity risk-management and reporting obligations for relevant entities, while DORA establishes a sector-specific ICT risk and resilience regime for covered financial entities [4,5]. A platform can help organize relevant evidence; it does not determine legal applicability or establish compliance.

02 / PRODUCT CONCEPT

Principles that make the analysis useful.

Principle	Product implication
Evidence before assertion	Label observations as confirmed, suspected, simulated or an evidence gap; preserve provenance and timestamp.
Business context matters	Connect assets to owners, critical services, dependencies and operational constraints.
Human accountability	Agents recommend and explain; designated people approve consequential actions.
Scope and least privilege	Constrain assessment sources and actions to explicitly authorized assets and roles.
Safe by default	A dashboard does not imply an active probe or exploit. OT and safety-critical testing needs separately approved methods.
Resilience includes recovery	Give backup, restoration sequencing, continuity procedures and exercises first-class attention.
Frameworks inform, not certify	Use control mappings to organize evidence; a mapping is not an audit opinion or certification.

03 / INTEGRATED CAPABILITY MODEL

One workflow across the risk lifecycle.

Understand
Asset intelligence, exposure records, ownership, technology and business-service dependencies.

Assess
Vulnerability and configuration observations, identity/cloud/OT review, control evidence and network scope.

Prioritize
Combine severity with exposure, asset criticality, exploit evidence, business impact, confidence and compensating controls.

Respond
Organize incident triage, decisions, communications, evidence and escalation. Keep containment recommendations under authorized review.

Recover and improve
Connect continuity, recovery objectives, backup evidence, exercises, lessons and reassessment.

Network scope can include LAN, WAN, intranet, switching, routing, firewalls, wireless, VPN, monitoring, identity, cloud/application, email/DNS, endpoint, backup and OT/IoT. An integrated record should preserve the asset identifier and owner; service and dependency context; source and collection method; status and confidence; timestamp and freshness; severity and rationale; next action; approval state; and closure evidence.

04 / GOVERNED AGENT MODEL

Specialist analysis, bounded by policy.

A useful agent system is a set of bounded analysis roles with explicit inputs and outputs, not an unconstrained “hacker in a box.” An orchestration layer can route authorized evidence to specialists, maintain a common risk register and make handoffs inspectable.

Analyst role	Bounded contribution
Asset and network	Reconcile declared inventories and imported evidence; identify missing ownership, version, scope or topology evidence.
Vulnerability	Normalize authorized findings, preserve scanner provenance and prioritize for human validation.
Threat and exposure	Relate declared technologies and approved external observations to curated threat information.
Service impact	Connect assets and suppliers to critical services and identify dependency questions for owners.
Controls and compliance	Organize evidence against selected frameworks and flag gaps for qualified review.
Incident and recovery	Structure incident records, tabletop scenarios, recovery dependencies and exercise actions.
Executive synthesis	Summarize material risks, uncertainty, decisions, owners and trend evidence.

For each run, retain the initiating user, requested scope, input records, data sources, agent/model and version where applicable, supporting evidence, confidence, approval decision, timestamps and resulting changes. Separate advisory text from verified facts. Require approval before external scans, configuration changes, containment actions, customer notifications or other consequential operations.



05 / NETWORK AND OT COVERAGE

Assess the whole estate without ignoring operational safety.

Network evidence may include approved inventories, device configuration exports, firewall-rule reviews, routing and switching documentation, segmentation tests, identity and remote-access records, monitoring coverage, patch evidence and authorized scanner results. The platform can identify absent records and prioritize supplied observations; it cannot infer unseen devices or verify control behavior without an appropriate evidence source.

For OT/ICS, begin with asset ownership, function, criticality, dependencies, vendor constraints, safe operating windows and passive or offline evidence. CISA’s multi-agency guidance recommends an OT inventory and taxonomy to help owners prioritize assets and defenses [3]. Any active test should be separately authorized, scoped and coordinated with operations and safety stakeholders. A simulated scenario is not evidence of a real exploit.

06 / SECTOR APPLICATIONS

Different services require different context.

Environment	Illustrative dependencies	Assessment emphasis
Retail & supermarkets	POS and payments, refrigeration monitoring, store WAN, suppliers	Segmentation, remote access, endpoint/POS inventory, cold-chain continuity and suppliers.
Blood banks & healthcare	Laboratory systems, storage monitoring, clinical apps, identity	Availability and integrity, device/vendor constraints, safe recovery and sensitive-data controls.
Power grids & utilities	Substations, control systems, dispatch, telemetry, corporate IT	Passive OT visibility, segmentation and remote access, engineering coordination and recovery exercises.
Banks & trading	Identity, market connectivity, transaction processing, data, third parties	ICT risk governance, resilience testing, change control, service dependencies and incident handling.
Other enterprises	Cloud/on-prem applications, workforce, vendors, business services	Inventory, identity, vulnerability evidence, response readiness and recovery.

Sector profiles change the questions and evidence context; they do not create a one-click guarantee of protection.

07 / IMPLEMENTATION ROADMAP

Progress from a scoped baseline to measurable improvement.

Stage	Deliverable	Exit evidence
1. Establish scope	Owners, sectors, critical services, assets, sources and rules of engagement.	Approved scope; data classification; accountable owners.
2. Baseline evidence	Inventories, control evidence and authorized assessment results.	Coverage/freshness report; provenance; explicit unknowns.
3. Prioritize	Rank validated issues by likelihood, impact, exploit context, criticality and confidence.	Reviewed risk register; owner and date for each decision.
4. Improve	Change-controlled remediation, tabletop and recovery exercises.	Change records; exercise findings; verified recovery evidence.
5. Integrate and scale	Connectors, identity, audit, tenant isolation, observability and controls.	Architecture review; access tests; data-protection and service evidence.

Before production use, the product requires an engineered backend, secure identity and tenant isolation, encryption and key management, durable audit logging, connector governance, secrets handling, backup/recovery, incident response, privacy review and independent security testing. Each integration should use least privilege, explicit consent, revocation and source-health checks.

08 / MEASUREMENT AND RESPONSIBLE USE

Make scores explainable and uncertainty visible.

- Asset coverage and inventory freshness, including unknown or ownerless assets.
- Finding age, validation state, remediation ownership and time to verified closure.
- Critical-service dependency coverage and third-party review freshness.
- Identity, segmentation, logging and backup evidence coverage.
- Recovery exercise completion and demonstrated restoration against approved objectives.
- Agent recommendation acceptance, correction, override and false-positive rates.

A score is useful only when its components and evidence coverage are clear. Avoid a universal “hack-proof” score. Disclose scope, source dates, exclusions, uncertainty and method. A posture summary is a snapshot, not a guarantee.

CyberResilience AI is decision support. It is not a substitute for qualified security engineering, incident command, safety engineering, legal counsel, independent audit or sector-specific advice. Results depend on evidence quality and change over time. Use it only for authorized assets and data; validate findings with owners before remediation.

CONCLUSION

Resilience is a management capability.

A well-governed workspace can make the chain from scope and evidence to owner, decision and recovery action easier to see and improve. CyberResilience AI’s product direction is to connect these signals with transparent, bounded analysis while keeping people accountable for consequential decisions. The next step from demo to dependable enterprise service is disciplined engineering, authorized integrations, customer validation and independently reviewed security controls.

References

[1] NIST, The NIST Cybersecurity Framework (CSF) 2.0, NIST CSWP 29 (2024).. https://doi.org/10.6028/NIST.CSWP.29

[2] NIST, CSF 2.0 Resource & Overview Guide, SP 1299 (2024).. https://doi.org/10.6028/NIST.SP.1299

[3] CISA et al., Foundations for OT Cybersecurity: Asset Inventory Guidance for Owners and Operators (2025).. https://www.cisa.gov/resources-tools/resources/foundations-ot-cybersecurity-asset-inventory-guidance

[4] Directive (EU) 2022/2555 (NIS2), Official Journal of the European Union.. https://eur-lex.europa.eu/eli/dir/2022/2555/oj

[5] Regulation (EU) 2022/2554 (Digital Operational Resilience Act), Official Journal of the European Union.. https://eur-lex.europa.eu/eli/reg/2022/2554/oj

Product disclosure: Independent product and design paper. Describes CyberResilience AI and its intended development direction as of November 2026. The deployed public experience is an interactive demo using local browser storage; it is not a production security service. No endorsement by NIST, CISA, the European Union or any regulator is implied.