



Defend the Missions the Public Depends On

The only AI SOC platform built for
machine-speed defense of public services

Public sector organizations sit at the top of attackers' target lists, from nation-states probing for espionage and disruption to criminal groups chasing ransom. The same organizations charged with advancing public and national interests run resource-constrained security teams that are expected to defend against some of the largest attacks on earth, and to do it under oversight rules that private companies never face.

That's the gap public sector security leaders have to close: machine-speed attacks on one side, human-speed response and manual compliance on the other. This paper lays out why the public sector is uniquely exposed, what separates a real AI SOC from the tools claiming the label, and how the Torq AI SOC Platform closes the gap in production today, with the transparency and audit-ready evidence that government agencies demand.



Federal Government



State Government



Local Government



Higher Education



K-12 Education



Critical Infrastructure

Why the Public Sector Is a High-Risk Target

The average public sector data breach climbed to \$3.50 million in 2026, a 22% jump in a single year and one of the sharpest increases of any sector.¹ And unlike a private company, what's at stake isn't only the balance sheet — a successful attack disrupts the services citizens depend on. The threat is accelerating on every front:

Ransomware attacks on government agencies surged 65% in the first half of 2025.³

U.S. schools were hit by 130 ransomware attacks in 2025, with K-12 districts targeted nearly three times as often as higher education.⁴

82% of K-12 schools experienced a cyber incident between July 2023 and December 2024.⁴

U.S. utilities faced 1,162 cyberattacks in 2025, a nearly 70% increase.⁵

What sets the public sector apart is what sits downstream of the breach. When ransomware locks a county's systems or takes a water utility offline, the damage isn't measured only in records exposed; it's measured in benefits and essential resources that don't reach citizens, emergency services that can't dispatch, and infrastructure that can't run. More than half of utilities hit in a recent year suffered permanent corruption or destruction of data and systems.⁵

Slow response compounds all of it. Across industries in 2025, organizations took an average of 241 days to identify and contain a breach.⁶ For a public agency, every day in that window is a day services stay at risk, and adversaries dig deeper into systems the public relies on.

Public sector SOC's can't hire their way out of this. Agencies compete for scarce talent against private-sector salaries they can't match, alerts pile up faster than lean teams can clear them, and every minute in the queue is a minute an attack moves closer to critical services. It's why 85% of security leaders now want a single platform they can trust to keep pace.⁷

WHAT AN AI SOC ISN'T — AND IS

Most tools marketed as an "AI SOC" stop at triage. They prioritize alerts, generate a recommendation, and hand the real work back to a human. The investigation still happens by hand, the response still means logging into another tool, and the case still closes on human time. The bottleneck moves; it doesn't disappear. Others simply bolt a chatbot onto legacy architecture or hide their reasoning inside a black box — a non-starter in a sector where every decision has to be explained to an auditor, an inspector general, or an oversight body.

A true AI SOC carries the alert through to the end. It investigates, correlates evidence, and takes containment and remediation actions across the stack, escalating to a person only when human judgment is genuinely required. Every step is grounded in context, logged, and reversible.

The test is simple: if it can't take action, it isn't an AI SOC. In the public sector, that distinction is the whole point, because the value of autonomy is closing the window between when an attack starts and when it's stopped, before it reaches the services citizens depend on.

The Torq AI SOC Platform for the Public Sector

The Torq AI SOC Platform runs the entire threat lifecycle — triage, investigation, response, and remediation — with agentic AI under your team's direction. It integrates with the stack you already own, produces the audit-ready evidence government frameworks demand, and is built for the distributed, high-oversight environments the public sector actually runs in.

Built-In Compliance and Auditability

The public sector operates under a dense, overlapping set of requirements: FedRAMP, FISMA, NIST 800-53, CMMC, CJIS, FERPA to name a few, and an evolving patchwork of state and federal mandates, where compliance isn't optional and doesn't slow down for attacks.

Torq builds audit-ready evidence into every agentic action. Torq HyperAgents™ log their full reasoning chain — every step, every verdict, every action — and native case management preserves a full chain of custody, turning daily operations into a continuous compliance record instead of a pre-audit scramble.

Transparency and Control

Government agencies can't run on black-box decisions. Torq HyperAgents operate under explicit direction: you define the agent's role, permissions, escalation boundaries, and authorized actions. Analysts, auditors, and oversight bodies can review the reasoning and evidence behind every action in real time, with override always available, and adjustable autonomy lets a team start conservatively and extend the platform's authority as confidence grows. The balance between agentic action and human control is a dial, not a switch.

Agentic AI for the public sector's most critical missions:

- ✓ **Ransomware containment:** Isolate infected endpoints and lock compromised accounts before an attack spreads across interconnected agencies, departments, and systems.
- ✓ **Nation-state threat investigation:** Ingest alerts from the entire security stack and investigate sophisticated, high-compliance federal threats with specialized AI Agents that close cases while preserving FISMA, CMMC, and executive-order auditability.
- ✓ **IT and OT threat detection:** Monitor across IT and OT environments in critical infrastructure, detect anomalous behavior before it cascades, and coordinate response with human-on-the-loop controls that keep operators in command.
- ✓ **Student data and PII protection:** Detect exposure of student records and personal data, launch investigations regardless of the alert's source, and protect institutional integrity across large, distributed campus and district environments.
- ✓ **Secure access management:** Grant time-limited, role-based access to staff, contractors, and third-party vendors, and automatically revoke it when the session ends.
- ✓ **Compliance and audit evidence:** Automated, audit-ready evidence collection for FedRAMP, FISMA, CJIS, and FERPA by default.

Integration Without Rip-and-Replace

Torq layers on top of the complex environments public sector organizations already run — from legacy on-premises systems to statewide networks and cloud infrastructure — with more than 400 integrations across tools like Microsoft, CrowdStrike, ServiceNow, and Wiz. Where an integration doesn't exist, teams prompt Torq Socrates™ Agentic Builder in natural language to easily create it.

Proven at Scale

Torq runs in demanding enterprise environments today, closing cases autonomously while keeping analysts in the loop for the calls that require judgment. The platform scales from a single agency to statewide networks spanning hundreds of departments, handling millions of daily events across distributed environments without adding headcount or rebuilding the security stack from scratch while running more than one billion automated actions a week across its customer base.

WHY CHOOSE TORQ?

- ✓ Agentic AI across the full threat lifecycle
- ✓ Socrates orchestrates specialized Torq HyperAgents from natural language prompts
- ✓ Torq Auto Triage delivers explainable verdicts that cut alert noise
- ✓ AI-native case management from investigation through case closure
- ✓ Adjustable autonomy with human-on-the-loop control over every high-impact action
- ✓ Transparent, auditable reasoning with immutable logs for every verdict and action
- ✓ Audit-ready evidence built in for HIPAA, HITECH, and SOC 2
- ✓ Enterprise-grade RBAC, approval workflows, and policy guardrails
- ✓ Grounded in the Torq Context Graph, with memory and learning that sharpen every verdict over time
- ✓ 400+ integrations with no rip-and-replace
- ✓ Coverage across EHR, medical devices and IoT, cloud, and remote care
- ✓ Proven at enterprise scale, running more than one billion automated actions a week

Gartner Peer Insights™

"What I really like about this product is the design, which works very smoothly, minimizing the disruption and training requirements. It also provides us with the best level of customization, which has empowered our team to design and implement according to our needs."



VP, Data and Analytics
Gov't/PS/ED <5,000 Employees

"As a data analyst, my overall experience with Torq has been positive. It aids in providing us with efficiency and a smooth approach for automating our workflow, which leads to data accuracy in analytical processes."



VP, Data and Analytics
Gov't/PS/ED <5,000 Employees

"Our experience with Torq has been positive, effectively shielding our organization from cybersecurity threats by automating complex tasks, along with providing smooth connectivity with other applications."



Director, IT Security and Risk Management
Gov't/PS/ED <5,000 Employees

Defend Public Missions with Security That Acts at Machine Speed

The Torq AI SOC Platform runs the entire threat lifecycle autonomously and audit-ready, so you protect the services and the trust the public depends on.

SEE IT IN ACTION

THE AI SOC CATEGORY LEADER



Leader in all four categories of the 2026 Leadership Compass: The Emerging AI SOC



The Company to Beat in AI SOC Agents for Threat Investigation, May 2026, Aug 2026



Leading Innovator in 2026 AI SOC Market Report

Recognized Across the Market

PeerSpot #1-ranked AI SOC platform

Gartner Peer Insights 4.8 Stars

References

1 Check Point Research, *Top Sectors Under Cyberattack in 2025*

2,6 IBM, *Cost of a Data Breach Report 2026*

3 Comparitech, *Government Ransomware Roundup*

4 2025 CIS MS-ISAC K-12 Cybersecurity Report

5 Forescout, *The State of Utility Cybersecurity in 2025*

7 Torq, *The 2026 AI SOC Leadership Report*

Gartner, AI Vendor Race: Torq Is the Company to Beat in AI SOC Agents for Threat Investigation, 25 May 2026. Gartner does not endorse any company, vendor, product or service depicted in its publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner publications consist of the opinions of Gartner's business and technology insights organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this publication, including any warranties of merchantability or fitness for a particular purpose. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.

About Torq

Torq is transforming cybersecurity with the Torq AI SOC Platform. Torq empowers enterprises to instantly and precisely triage, investigate, and respond to security events at scale. Torq's customer base includes major multinational enterprises such as Check Point Security, Chipotle Mexican Grill, Inditex (Zara, Bershka, and Pull & Bear), Informatica, Kyocera, Lego, Prudential, Rocket Companies, Telefónica, Uber, Valvoline, Virgin Atlantic, and Wiz.