VIEW POINT

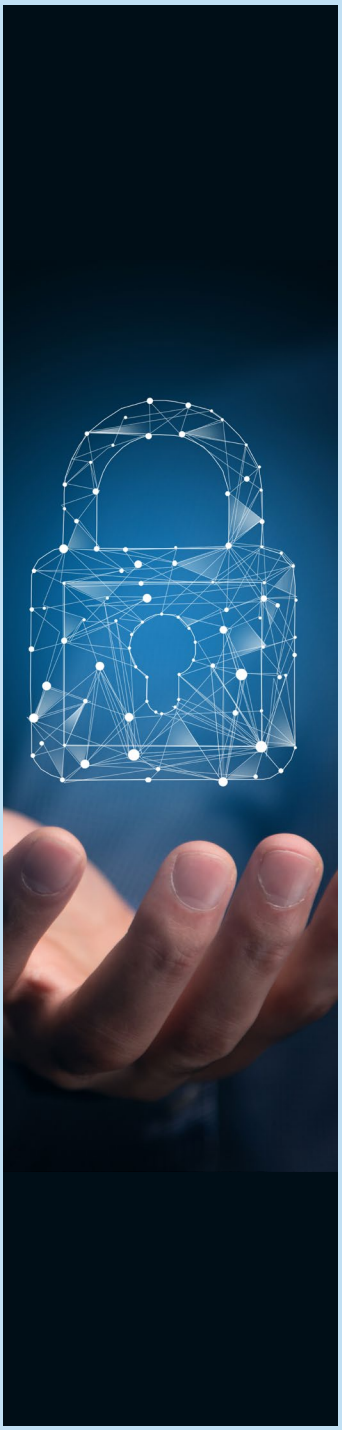

AI-DRIVEN OT SECURITY: BUILDING THE NEXT-GENERATION OT SOC

Infosys
Navigate your next

Executive summary

Operational Technology (OT) environments (often implemented as Industrial Control Systems (ICS)) have cybersecurity requirements that differ significantly from classic IT. In these settings, **safety, availability, and process integrity come first**, and many assets cannot be patched, scanned, or rebooted on typical IT timelines.

This shapes what "good" security automation looks like in OT: it must be consequence-aware, biased toward passive monitoring, and be tightly governed.

The strongest near-term value of AI for an OT SOC is not "autonomous response," but **operationally safe acceleration of fundamentals**: faster asset inventorying, better signal-to-noise through OT-aware correlation, improved triage consistency, and more scalable threat hunting across multi-site telemetry.

This aligns with the "next-generation OT SOC" thesis that convergence success depends on integrated visibility, shared datasets, and integrated workflows—not standalone tools.

Infosys' current service overview already emphasizes the right pillars asset inventory/discovery and classification/criticality, visibility & monitoring, incident response/playbooks, defensible architecture/segmentation, secure remote access/zero trust, risk-based vulnerability management, and supply-chain risk management—aligned to NIS2, CAF 4.0, and IEC 62443 concepts.

Infosys OT-IoT Security Services

- Asset Inventory/Discovery and Classification-know your assets, you can't protect what you don't know.
- Incident Response from OT perspective-what to do in case of an incident.
- Defensible Architecture-make the environment more resilient.
- Visibility and Monitoring-get insights what's going on in your factories.
- Secure Remote Access, Zero Trust Comes to OT and Identity Risk Management.
- Risk based Vulnerability management and Supply chain Risk management.

Together with our premier partners we offer:

- **Rapid Response**: Quickly deploy Dragos and Infosys responders during incidents for analysis, containment and remediation.
- **Incident Response Plan Development**: assistance with developing your IRP for OT.
- **Tabletop Exercises**: Carry out custom TTX's focused on OT that help identify gaps in current plans and procedures.

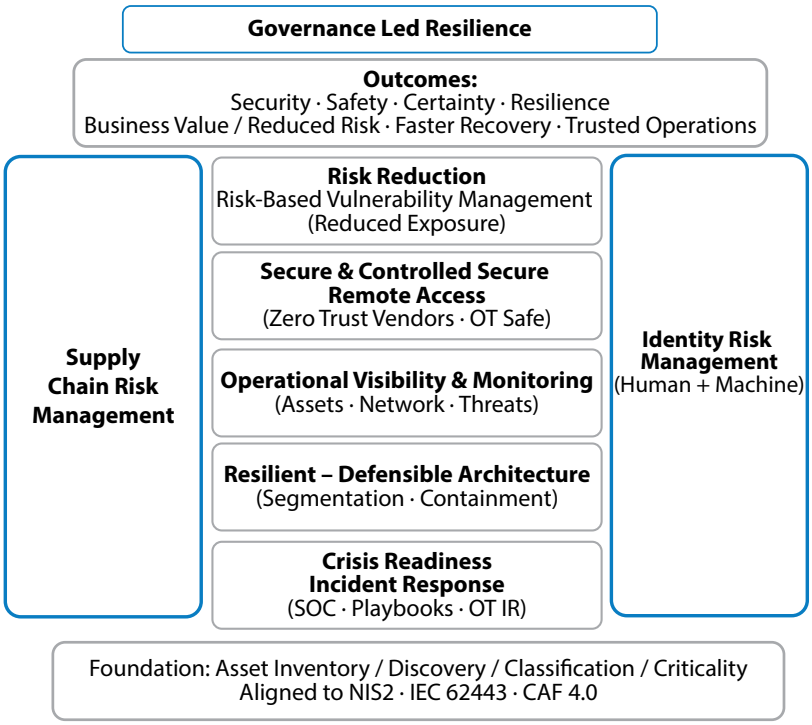

Operational Technology (OT) environments demand security solutions that enhance protection without disrupting critical industrial operations. Artificial Intelligence delivers the greatest impact when it is integrated directly into existing Security Operations Center (SOC) processes. Rather than functioning as standalone tools, AI-driven capabilities strengthen core SOC functions, improve detection quality, and reduce analyst workload while maintaining the reliability essential in OT systems.

AI brings six proven capabilities to OT security operations:

Asset Intelligence
Foundation for all OT security: passive discovery, automatic classification, and criticality ranking.

Assisted Triage
Highest immediate ROI: automated correlation, narrative building, and OT-safe response recommendations.

Guided Threat Hunting
ICS-specific hypotheses, signal clustering, and pattern identification across global sites.

Enhanced Detection
Cyberphysical awareness combining industrial protocols, process behavior, and identity context.

Cross-Domain Correlation
Single investigative timeline spanning IT, identity, cloud, and OT for modern attack detection.

Predictive Analytics
Dual-use capability bridging reliability engineering and security for integrity attack detection.

Asset Intelligence provides continuous passive discovery and classification of industrial assets, establishing a reliable and up-to-date foundation for risk assessment. **Assisted Triage** delivers the fastest return on investment by automating correlation and investigative narrative building, accelerating incident response and reducing manual analyst effort. **Guided Threat Hunting** uses data-driven pattern recognition and ICS-specific hypotheses to uncover advanced or emerging threats across distributed industrial environments.

AI also raises detection precision. **Enhanced Detection** fuses industrial protocol awareness, process behavior, and identity context to reduce false positives and strengthen cyber-physical visibility. **Cross-Domain Correlation** aligns signals across IT, cloud, identity, and OT into a unified investigative timeline — essential for understanding modern attack paths. Finally, **Predictive Analytics** bridges reliability engineering and cybersecurity, enabling early identification of integrity risks through long-term behavioral analysis.

Together, these AI capabilities create a more proactive, resilient, and efficient OT security ecosystem — one in which the SOC can respond faster, detect earlier, and operate with greater clarity across complex industrial environments.

OT SOC Reality Check: Today's Pain Points

OT SOC programs frequently struggle not because detection logic is impossible, but because the prerequisites for safe and accurate detection are incomplete: consistent asset inventory, reliable telemetry, segmentation/containment boundaries, and an operating model that connects IC/engineering expertise to SOC decision-making.

Key challenges that repeatedly show up across OT guidance and practitioner commentary include:

Foundational visibility gaps and "unknown asset" risk

OT networks often include legacy assets, specialized devices, and proprietary configurations; asset owners may lack accurate, continuously updated inventories and criticality mapping—yet these are prerequisites for prioritization and consequence-based response.

High cost of false positives and unsafe response

In OT, "blocking" or "shutting down" can create safety hazards, environmental impact, or production losses; therefore automation must be carefully bounded, and actions that could influence physical processes typically require human approval with engineering context.

Data fragmentation across sites and teams

Multi-site industrial operations tend to have site-specific configurations, vendor access patterns, and uneven data quality. The Industrial Cyber "next-gen OT SOC" perspective emphasizes that defenders need integrated visibility and shared datasets to correlate patterns across IT, OT, and IIoT systems.

Limited specialized skills and cultural divides

OT engineers and IT security teams often prioritize different outcomes (availability/safety versus confidentiality), leading to mistrust and slow adoption of SOC controls unless they are validated in high-fidelity environments and tied to operational impact.

Remote access and third-party connectivity pressure

OT connectivity is rising due to operational efficiency, remote monitoring, and predictive maintenance, but it expands the attack surface and increases the importance of hardened boundaries, logging, and isolation planning.

What remains unspecified includes existing SOC tooling (SIEM/SOAR), data retention and telemetry coverage per site, the staffing model and escalation paths, and the current OT engineering participation in SOC triage and decision-making.

IT/OT Convergence: How It Redefines the OT SOC Mission

A core claim is that next-generation OT SOC effectiveness is driven by **convergence in teams, workflows, and data**; not simply deploying more tooling. Convergence creates both opportunity and risk:

It increases detection power by enabling correlation of identity, remote access activity, asset criticality, process context, and IT-originated threat signals into a single investigative narrative reducing false positives and improving safe prioritization.

It also **expands the blast radius** of compromises: "IT risks become OT risks and vice versa" when corporate malware, credential abuse, or vendor remote access pathways can pivot into OT environments.

This pushes OT SOC services toward a few structural design principles:

Shared telemetry and a shared data model

The Industrial Cyber perspective stresses integrated visibility and a shared dataset for recognizing patterns and correlating threat intelligence across interconnected systems.

In practice, this implies normalized event schemas (identity, network flows, OT protocol events, engineering events) and consistent asset identity resolution across sites.

Workflow convergence without "one-size-fits-all response

Even if IT and OT incidents share a case management pipeline, response must remain consequence-driven: IT containment actions can often be automated; OT containment actions must default to human control when there is any risk to physical processes.

Connectivity governance as a first-class control

Recent guidance emphasizes a principles-based approach to designing and securing OT connectivity (including hardening OT boundaries, logging/monitoring all connectivity, limiting impact of compromise, and establishing isolation plans).

AI Capabilities that matter in OT SOC Operations

Practical AI use cases mapped to OT SOC outcomes

Infosys highlights that automation is currently most valuable in routine tasks such as inventorying assets, correlating anomalies across many sensors, and triggering containment for known-bad activity without interrupting safety-critical systems while not substituting human insight. That framing provides a useful filter: prioritize AI where it (1) reduces analyst toil, (2) improves signal quality, and (3) stays within safe operational guardrails.

Anomaly detection for cyber-physical and OT network behaviors

OT anomaly detection often needs to consider multi-variate time series (process variables), network states, and OT protocol semantics. Research demonstrates active work on time-series anomaly detection models and evaluation methodologies. For an OT SOC, the goal is not just "detect anomalies," but to produce actionable, explainable, consequence-ranked alerts.

AI-assisted triage and decision support (especially for multi-site SOCs)

AI-supported triage and decision support are near-term value drivers that improve consistency and reduce analyst workload. In OT SOC terms, this is best implemented as:

- Event summarization with evidence linking (what changed, where, when, on which assets)
- Enrichment (asset criticality, zone/conduit, known vulnerabilities, remote access identity)
- Recommended next steps mapped to OT-safe playbooks (observe, verify with operator, isolate conduit, collect forensics)

Asset discovery and continuous inventory enrichment

Because asset visibility is foundational, AI should first improve the completeness and timeliness of asset identity, role classification (PLC/HMI/engineering workstation), and communications baselines across sites. This aligns directly with your service emphasis on asset discovery/classification and operational monitoring/visibility.

Threat hunting aligned to ICS adversary behaviors

Because asset visibility is foundational, AI should first improve the completeness and timeliness of asset identity, role classification (PLC/HMI/engineering workstation), and communications baselines across sites. This aligns directly with your service emphasis on asset discovery/classification and operational monitoring/visibility.

Predictive maintenance as a dual-use OT capability

Predictive maintenance is often positioned as reliability engineering, but it intersects with OT SOC because many failure signatures look "anomalous" and because connectivity used for predictive maintenance expands attack surface. Surveys of predictive maintenance emphasize that ML can reduce unplanned stoppages and optimize maintenance planning, but success depends heavily on the right data, context, and methodology selection. A mature OT SOC should treat predictive maintenance telemetry as both: a security data source (process deviations, sensor integrity issues), and a governance driver (secure connectivity, logging, separation of duties).

AI brings six proven capabilities to OT security operations:

Core AI Capabilities for OT into SOC Services

AI delivers maximum value when embedded into existing OT security pillars rather than deployed as standalone tooling disconnected from operational workflows. The following list of capabilities represent proven applications where AI demonstrably improves security outcomes without introducing unacceptable operational risk.

Asset Intelligence
Foundation for all OT security: passive discovery, automatic classification, and criticality ranking.

Assisted Triage
Highest immediate ROI: automated correlation, narrative building, and OT-safe response recommendations.

Guided Threat Hunting
ICS-specific hypotheses, signal clustering, and pattern identification across global sites.

Enhanced Detection
Cyberphysical awareness combining industrial protocols, process behavior, and identity context.

Cross-Domain Correlation
Single investigative timeline spanning IT, identity, cloud, and OT for modern attack detection.

Predictive Analytics
Dual-use capability bridging reliability engineering and security for integrity attack detection.

Ensuring Trust: Explainability and Model Governance in OT Security

OT operators will not trust "black-box" models that cannot justify alerts in operational terms. The NIST AI Risk Management Framework (AI RMF 1.0) frames trustworthy AI through lifecycle risk management and emphasizes that organizations should understand context and impacts, not just model performance.

In OT SOC operations, explainability should be operationalized as:

- **Evidence-based explanations** (telemetry features, protocol events, process variable deviations),
- **Counterfactuals** ("this is unusual because this PLC has never issued this function code to that actuator during this shift pattern"),
- **Confidence + uncertainty** reporting
- P**laybook-safe recommendations** (never auto-suggest a disruptive action without explicit consequence labeling).

Engineering Constraints in OT that shape AI design

Unlike traditional IT, actions in OT can impact physical processes—production lines, safety systems, and critical infrastructure.

Hence, the SOC must operate with specific design constraints.

OT Security Implications for SOC Design

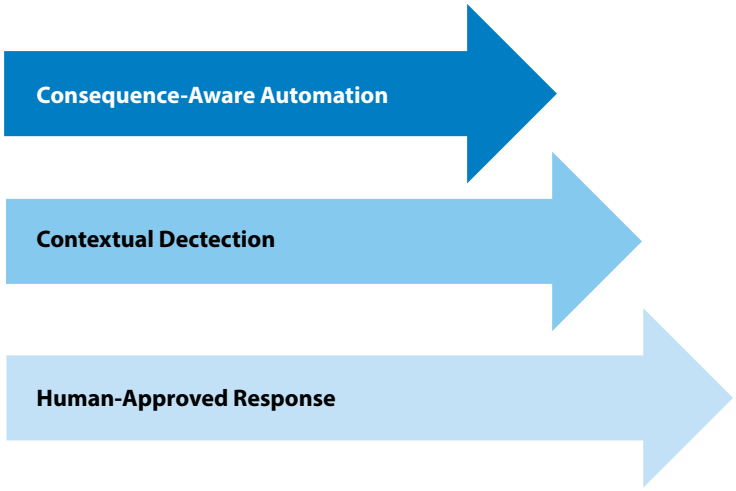

These constraints mean AI's role is not autonomous response but safe acceleration of detection, triage, and decision qality augmenting human expertise rather than replacing it.

1. Consequence-Aware Automation

The first principle is Consequence Aware Automation.

In OT, automation must understand the potential physical and operational impact of an action. This means we cannot simply lift IT playbooks or auto-response models and apply them to OT. A mis triggered automated action could stop operations or create safety risks.

Our automation here is cautious by design supportive, but never reckless.

2. Contextual Detection

Secondly, Contextual Detection.

OT detection only works well if it understands the industrial process behind the telemetry. We're not just looking at packets or logs, we're looking at how PLCs, sensors, and processes behave in their operational context. This reduces noise and helps us detect anomalies that are meaningful in the OT world, not just statistically unusual.

3. Human-Approved Response

Finally, Human-Approved Response.

OT changes the risk equation: response actions cannot be fully automated because of safety and operational continuity. Human validation remains mandatory before any response that could alter the physical environment. This keeps operators in control, especially when actions could impact safety or uptime.

Bringing this together, the role of AI in OT is not autonomous response. Instead, it accelerates detection, triage, and decision quality, while humans retain final authority.

AI augments our analysts and engineers; it doesn't replace them.

This balance is what gives us speed without compromising safety or reliability.

Safety, Reliability, and Real-time Requirements

Securing OT requires addressing unique performance, reliability, and safety requirements. Infosys reinforces this by warning that automated actions affecting OT systems (e.g., stopping PLCs) can cause damage or injury, therefore response decisions should involve both security teams and process/automation engineers.

A pragmatic OT AI should therefore adopt a **tiered autonomy model:**

- Tier A (safe automation): enrichment, correlation, deduplication, case summarization, evidence collection, and recommending OT-safe playbook steps.
- Tier B (guarded automation): triggering containment actions that do not alter physical processes (e.g., isolate a monitoring uplink, disable a specific remote access session) under pre-approved conditions and explicit rollback plans.
- Tier C (human-only): any action that could change control behavior, affect safety instrumented systems, or cause downtime; AI can recommend but not execute.

Tiered Autonomy Model: Ensuring Safe AI Automation

AI adoption in OT environments demands a rigorous safety framework that explicitly defines which actions can be automated and which require human judgment. The consequences of automation errors in cyber-physical systems are unacceptable, making a carefully designed autonomy model essential.

Tier A: Fully Automated

Scope: Actions with no physical impact or safety risk

- Alert enrichment and context gathering
- Signal correlation and pattern matching
- Case summarization and timeline building
- Threat intelligence lookups
- Asset inventory updates
- Automated reporting and documentation

Risk: Minimal — errors affect only information presentation, not physical systems.

Tier B: Guarded Automation

Scope: IT actions or OT actions with limited, reversible impact

- Revoking remote vendor access sessions
- Isolating compromised IT endpoints
- Blocking malicious network flows at IT/OT boundary
- Disabling compromised user accounts
- Initiating enhanced monitoring on suspicious assets

Risk: Moderate — actions may cause temporary service disruption but are reversible and don't directly impact production processes.

Control: Automated execution with immediate notification to SOC analysts and automatic rollback capability.

Tier C: Human-Only Decision

Scope: Any action directly affecting physical processes or production

- Stopping or modifying PLC logic
- Production line isolation or shutdown
- Process parameter changes
- Safety system modifications
- Emergency stop commands

Risk: High — errors could cause safety incidents, production loss, product quality issues, or equipment damage.

Control: Actions require explicit approval from qualified operations or engineering personnel.

This model ensures AI accelerates security operations where appropriate while maintaining human control over consequential decisions. Any action affecting physical processes remains under explicit human approval, preserving the safety oversight that OT environments require.

Risk, compliance, and the expanded attack surface of AI-enabled OT SOC

Adversarial ML and AI-specific security risks

AI creates new security failure modes: model evasion (inputs crafted to avoid detection), poisoning (corrupting training data), and broader system-level attacks. The NIST adversarial ML taxonomy and terminology work provides a structured vocabulary for these risks, which is useful for OT SOC threat modeling and vendor/security requirements.

In OT SOC environments, AI also expands attack surface through:

- New data pipelines and collectors
- Model endpoints (APIs) and management consoles
- Increased reliance on remote connectivity for telemetry aggregation and analytics

Regulatory and standards considerations

For an EU-context OT SOC, compliance pressures are now deeply intertwined with operational resilience:

- **NIS2 (Directive (EU) 2022/2555)** expands and strengthens cybersecurity risk management and incident reporting obligations for covered entities and includes explicit attention to supply-chain risk considerations.
- **NCSC Cyber Assessment Framework (CAF) 4.0** introduces a strengthened, outcomedriven cybersecurity assessment model emphasizing threatled risk management, secure software development, and enhanced monitoring capabilities. CAF 4.0 adds new contributing outcomes such as Understanding Threat and Secure Software Development and Support, along with expanded requirements for behavioral analytics and integrated threat intelligence.
- **ISA/IEC 62443** provides structured requirements and processes for securing industrial automation and control systems and explicitly bridges operations and IT considerations with a holistic approach.
- **NIST CSF 2.0** provides an outcome-driven taxonomy (including a "Govern" function) for structuring cybersecurity programs and communicating risk management.
- **CISA Cross-Sector Cybersecurity Performance Goals (CPGs) 2.0** provide a prioritized baseline of high-impact cybersecurity practices and are positioned as broadly applicable across critical infrastructure, including IT/OT contexts.
- **Supply-chain cybersecurity risk management** is addressed by NIST SP 800-161 Rev. 1, which provides guidance on identifying, assessing, and mitigating cybersecurity supply-chain risks across organizational levels.

AI Risk Matrix for OT SOC Adoption

The table below is a practical OT SOC-oriented AI Risk Matrix, grounded in adversarial ML concepts and lifecycle risk management expectations (NIST AI RMF).

AI risk in OT SOC	Why it matters in OT	Practical mitigations for an OT SOC service
False positives cause unsafe or costly actions	Downtime and safety impacts can outweigh security benefit	Tiered autonomy (AI recommends; humans approve), "OT-safe action" allowlist, rollback plans, site-by-site rollout
False negatives (missed detection)	Silent compromise may persist until disruption	Defense in depth (segmentation + monitoring), threat-hunting coverage aligned to ATT&CK for ICS, periodic purple-team validation
Data drift from process changes	Normal operations evolve (setpoints, product mix, seasonal ops)	Drift monitoring, retraining triggers, model "freeze windows" during changeovers, engineering sign-off for new baselines
Poisoning of training data	Attackers or misconfigurations corrupt baselines	Data provenance controls, separation of training vs production pipelines, integrity checks, restricted write access, secure logging
Evasion / adversarial inputs	Crafted traffic or sequences bypass detection	Ensemble methods, rule+model hybrids, robust feature engineering, adversarial testing in lab, monitor "confidence anomalies"
Model opacity reduces trust	Operators ignore alerts they cannot validate	Evidence-based explanations, "why now" summaries, link alerts to known process states and asset criticality
Expanded remote-access attack surface	Telemetry pipelines and model services increase connectivity	Secure connectivity principles, strict boundary hardening, monitoring all connectivity, isolation plans, least privilege
Vendor lock-in / tooling sprawl	Operational complexity undermines SOC effectiveness	Prefer open schemas, portable feature stores, clear integration contracts (SIEM/SOAR), exit criteria in procurement
Privacy/regulatory misuse of AI outputs	Logs may include sensitive operational data	Data minimization, retention policies, role-based access, compliance mapping to NIS2/IEC 62443 controls

Strengthening the OT SOC service with AI: Architecture and Integration Patterns

Architecture of an AI CISO SOC

Effective AI deployment in a converged CISO SOC requires a hybrid architecture that balances centralized intelligence with edge resilience. This design recognizes that OT environments must maintain operational capability even during network disruptions while benefiting from enterprise analytics and threat intelligence.

Edge Layer (Plant Level)

Deployed within each facility:

- Passive network monitoring capturing all OT communications
- Local behavioral baselining for immediate anomaly detection
- Feature extraction and dimensionality reduction before transmission
- Critical alert generation even during WAN outages
- Local data retention for forensic analysis

Benefit:
Operational resilience maintains security visibility even if connectivity to central SOC is lost.

Central SOC (Enterprise)

Enterprise Scale Cababilities:

- Cross-site AI analytics identifying global patterns
- Advanced model training on aggregated data
- Identity correlation across enterprise directories
- Threat intelligence integration and IOC matching
- Unified case management and workflow orchestration
- Executive reporting and compliance dashboards

Benefit:
Enterprise intelligence detecting sophisticated attacks that span multiple facilities or evolve over extended timeframes.

Governance Layer

Critical oversight and control:

- AI model lifecycle management and version control
- Safety approval workflows for new automation
- Model drift monitoring and performance tracking
- Explainability requirements and audit trails
- Exception handling and escalation procedures
- Regulatory compliance documentation

Benefit:
Trusted AI ensuring models remain accurate, explainable, and aligned with safety requirements.

This hybrid deployment model balances plant-level operational autonomy with enterprise intelligence sharing. Edge components provide resilience and low-latency detection, while central systems deliver the computational power and cross-site visibility needed for advanced analytics. The governance layer ensures AI remains trustworthy throughout its lifecycle.

This outlines how we architect an AI-enabled SOC that operates effectively across both IT and OT environments. The key message is that OT environments require resilience and autonomy, while the enterprise requires centralized intelligence and governance. A hybrid SOC design brings these capabilities together.

1. Edge Layer (Plant Level)

The edge ensures operational continuity, even when the network is down.

What happens at the edge:

- Passive monitoring of OT communications
- Local behavioral baselining for rapid anomaly detection
- Feature extraction and pre-processing before sending data to the center
- Local alerting even during WAN outages
- Local forensic storage

Why this matters:

Your plants maintain visibility and early detection even when disconnected from corporate networks. This is critical for cyber-resiliency in industrial environments.

2. Central SOC (Enterprise)

The central SOC provides the 'big-picture' analytics that no single site can see.

What the central SOC delivers:

- Cross-site analytics to identify global patterns
- Advanced AI models trained on aggregated OT data
- Correlation across multiple facilities
- Threat intelligence integration and playbook orchestration
- Unified case management and reporting

Why this matters:

This is where you detect sophisticated, multi-site attacks ,threats that evolve over long periods or span facilities. It's the enterprise brain.

3. Governance Layer

Governance ensures trustworthy, safe, and compliant AI across the SOC.

Governance responsibilities:

- AI model lifecycle management and version control
- Oversight of safety-critical workflows
- Monitoring AI drift and performance
- Documentation for audits and regulators
- Explainability requirements and exception handling

Why this matters:

AI in industrial environments must be explainable, controlled, and safe.

This layer ensures AI augments human decision-making without introducing new risk.

This hybrid model gives you the best of both worlds—local autonomy for resilience, and centralized AI and intelligence for scale and threat sophistication. It's the architecture that allows an AI-enabled SOC to be both safe and effective in converged IT/OT environments.

Deployment Architectures: Edge, Cloud, Hybrid

NIST OT security guidance emphasizes unique OT constraints (performance, reliability, safety), which often drives **edge or hybrid designs** for monitoring and analytics rather than "cloud-only".

A functional comparison for OT SOC AI looks like this:

Architecture	Best for	Key strengths	Key risks / constraints	Recommended OT SOC stance
Edge (site-local inference)	Low latency, strict segmentation, limited bandwidth	Keeps data close to process; resilient to WAN loss	Hardware lifecycle, patching constraints, limited compute	Use for feature extraction, baselining, and Tier A automation; push summaries upstream
Cloud (central analytics)	Cross-site correlation, heavy compute (training)	Scales ML training; enables fleet-wide patterns	Data sovereignty, connectivity risk, perceived "loss of control"	Use primarily for training, offline analytics, and cross-site pattern mining with strict connectivity governance
Hybrid (recommended baseline)	Most industrial enterprises	Balances site autonomy with central insights	Integration complexity	Make hybrid the default: edge collectors + central SOC correlation + governed AI services

Reference Architecture: AI Augmentation Layer for an OT SOC

Below is a suggested target architecture that keeps OT-safe boundaries while integrating AI into SOC workflows. The diagram is intentionally tool-agnostic; where your operating model references multi-site OT sensors and traffic routing via SPAN/TAP and aggregation, that fits naturally into the "OT Telemetry Edge" layer.

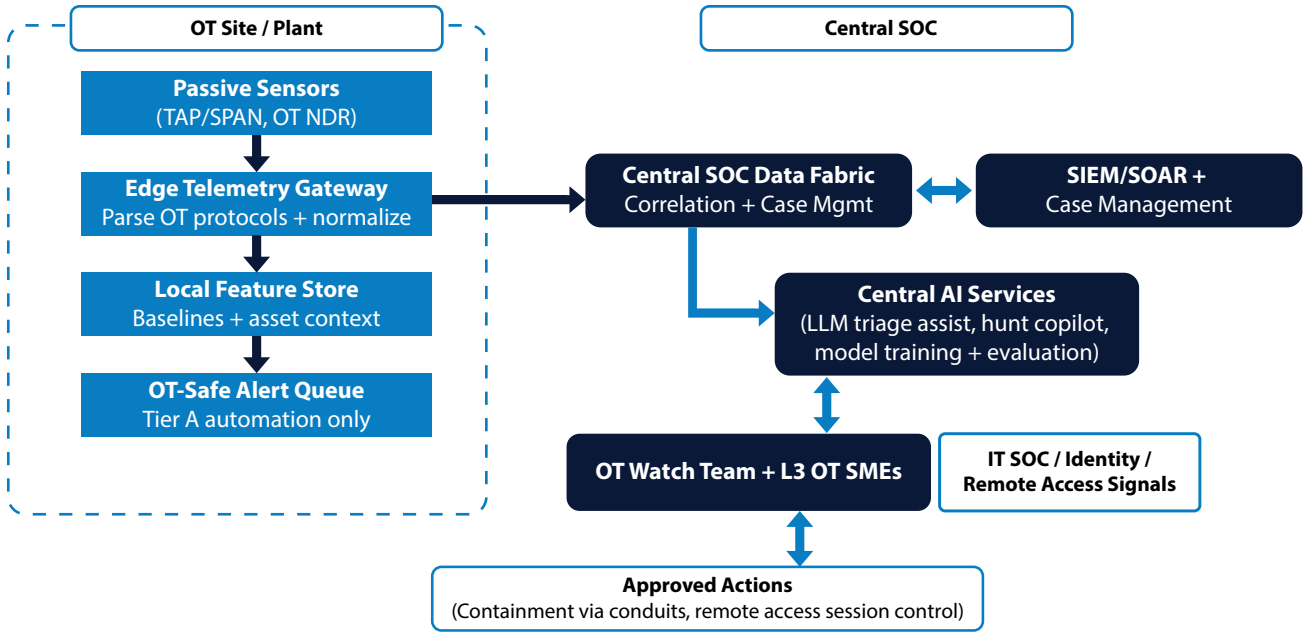

Integration patterns that strengthen services (actionable):

- OT-aware enrichment

Join every alert to asset criticality, zone/conduit, remote access identity/session, and known vulnerabilities before it reaches an analyst

- Case-first operations

Shift from "alert handling" to "case narratives," where AI aggregates weak signals into a single incident hypothesis (mapped to ATT&CK for ICS techniques) and attaches evidence links

- SOAR with "OT-safe actions" catalog

Automate only actions that cannot affect safety-critical operations (e.g., revoke vendor remote access token/session, isolate a non-control conduit), and require explicit OT engineer approval for anything beyond

- Model governance embedded into SOC QA

Treat models like detections—versioned, tested in lab, monitored in production, and reviewed in monthly service reviews

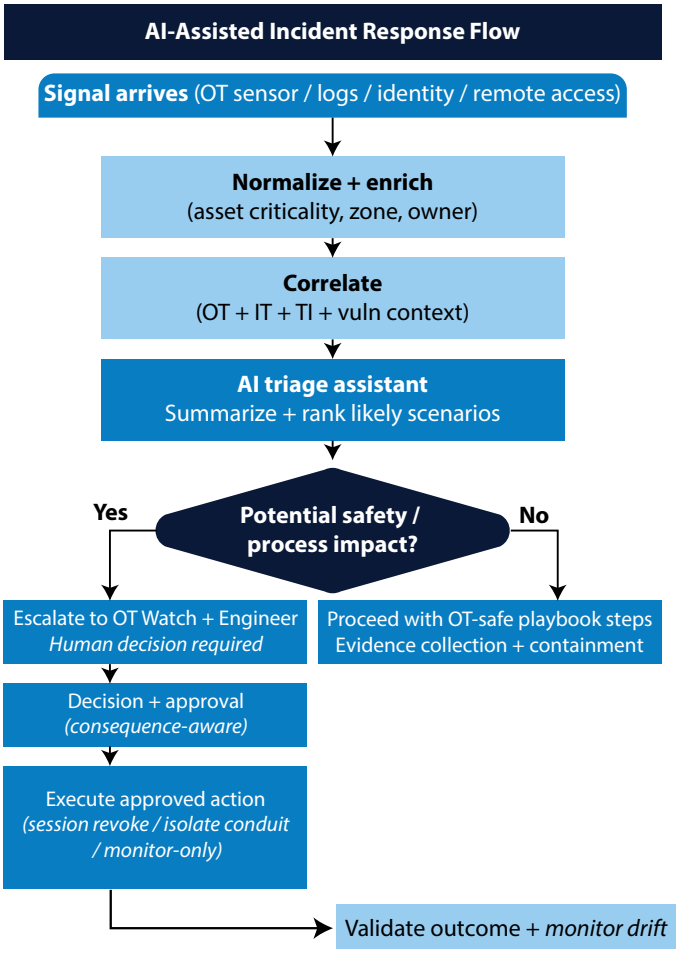

Real world examples of using AI in OT: The Dragos Platform

Model Context Protocol (MCP) Server

Enables customer enterprise AI tools (Copilot, Claude, ChatGPT, self-hosted models) to query Platform Data.

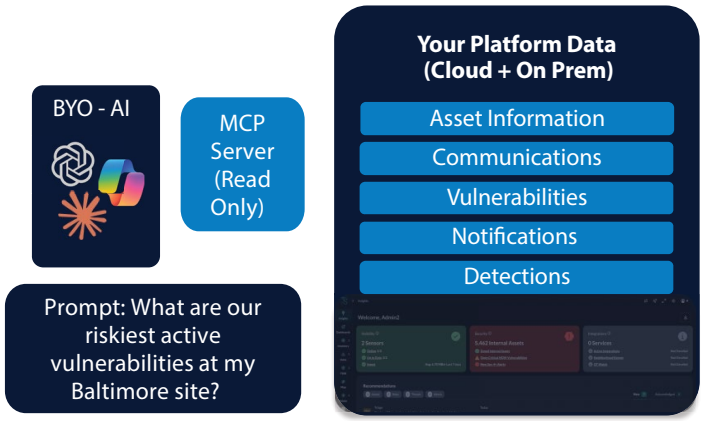

- **Open Standard** – Built on trusted MCP protocol, widely adopted across industry
- **Use Customer Enterprise AI** - Leverage tools already approved and invested in—no vendor lock-in
- **Automate Workflows** - Route AI query results to ticketing, chat, and orchestration tools
- **Full Control** - Customer chooses AI tools and keeps data on-premise for compliance

MCP stands for Model Context Protocol—it's an open standard that connects AI tools to data sources, like an API integration point.

Here's what this means for customers: they can connect their enterpriseapproved AI tools, including Microsoft Copilot, Claude, ChatGPT, and selfhosted models, directly to Dragos Platform data. No vendor lock-in. They use the AI they've already approved and invested in.

Once connected, analysts ask natural-language questions about assets, vulnerabilities, and notifications to understand their environment and triage faster.

Beyond that, customers can automate workflows by routing Dragos Platform data into email, Teams, Slack, Jira, SOAR platforms, turning fragmented manual processes into streamlined, intelligent response.

The key point here is that the customers maintain full control on their data, their AI tools, and their governance.

Analyst Assistant

AI-powered assistant that helps analysts explore, analyze, and explain OT data faster directly within the Dragos Platform.

- Fuses customer data with Dragos intelligence – Combines environment data with proprietary threat intelligence for informed answers
- Natural language queries – Ask questions about assets, vulnerabilities, threats, and operational data
- Accelerates investigations – No more navigating between multiple systems to piece together information
- Streamlines daily workflows – Supports reporting, documentation, and operational tasks

Ask Questions -> Get Answers

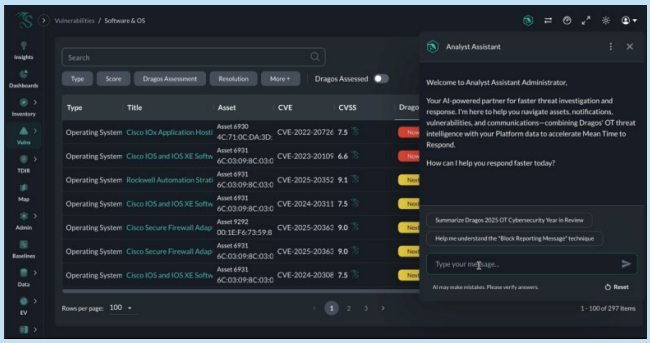

The Analyst Assistant is more than a chat interface. While it uses a conversational UI, its value lies well beyond simply retrieving information from documents.

It enables analysts to ask meaningful, operationally relevant questions such as: 'Which assets on the production line are affected by this vulnerability and pose the highest risk?' or 'What are some alternative mitigation steps if we can't afford downtime for this particular CVE?' It also handles daily workflow needs like 'I need to report this to my boss can you summarize the key findings?'

The Assistant fuses customer Platform data—assets, vulnerabilities, notifications with intelligence from the OT threat research team, making it all accessible through naturallanguage interaction within the Platform.

The outcome is faster, more confident investigations and smoother daytoday operations. Analysts no longer need to navigate multiple systems; instead, they can make wellinformed decisions directly within their workflow.

About the Author

Michel Bruggeman
Principal Consultant | EMEA IoT & OT Lead

Michel is a distinguished cybersecurity leader with 25+ years of experience, specializing in Operational Technology (OT) security, industrial cybersecurity architecture, and cyber resilience. He has held strategic roles across the insurance, manufacturing, energy, and semiconductor sectors, focusing on ICS/OT risk assessments, secure cloud adoption, and incident response readiness. Michel is a SANS /ISA Mentor and Microsoft Certified Trainer, with deep expertise in IEC 62443, DORA, NIS2, NIST, and ISO 27001.

For more information, contact askus@infosys.com

Infosys
Navigate your next