WHITEPAPER

The Four Types of OT Threat Detection in The Dragos Platform

Lessons Learned from the TRISIS Malware

MATTHEW PAHL | SENIOR PRINCIPAL THREAT DETECTION ENGINEER

DRAGOS, INC

SEPTEMBER 2025

Cyber threats don't announce themselves; they slip in through the cracks, hidden in the noise of normal operations. A misconfigured system setting could weaken defenses. An anomaly – subtle but out of place – might be the first sign of trouble. Unusual behaviors can reveal intent before action, while indicators of compromise (IOCs) leave behind the fingerprints of an attack already in motion. Together, these dimensions of threat detection form a critical framework, enabling organizations to stay vigilant and proactive in the face of evolving cyber risks.

The Dragos Platform captures these four essential types of cyber threat detection – configuration changes, anomalies, adversary behaviors, and indicators of compromise (IOCs) - that are key to identifying potential threats to industrial control systems (ICS) and operational technology (OT). One such threat, the TRISIS malware, discovered by Dragos and reported in Dragos WorldView, was deployed against at least one victim in the Middle East and stands out as the first attack targeting safety systems in OT.

This whitepaper explores the methodology behind the analytics engineered for the Dragos Platform through the lens of the TRISIS malware, offering real-world examples of how threat intelligence is used to generate high-fidelity alerts and a reminder of the telemetry and complexities of each detection type.

What Is Cyber Threat Detection?

Cyber threat detection identifies potential security threats within an organization's digital environment, whether from malicious adversaries, system vulnerabilities, or unusual activities. It involves analyzing data from multiple sources, including network traffic, protocol commands, user behaviors, and system configurations, to recognize and respond to signs of compromise before they escalate into significant incidents.

Threat detection is not just about identifying threats but also about driving more thoughtful responses. By uncovering a threat's signals, defenders are prompted to ask critical questions: ***Why did this happen? How did the adversary gain entry? What can we do to stop it from happening again?*** These questions transform cyber threat detection from a reactive measure into a proactive strategy, enabling organizations to strengthen their defenses. Moreover, effective threat detection prevents material impacts beyond technical downtime. It safeguards against financial losses caused by disruption of mission-critical systems, mitigates risks to physical safety in critical infrastructure, and prevents environmental harm from attacks targeting industrial systems. In a world where every second counts, detection is the linchpin that connects awareness to action.

How Dragos Turns OT Threat Detection into Action

Within the Four Types of Threat Detection framework, Dragos translates theory into day-to-day outcomes through a unified loop of technology, intelligence, and services. Our model focuses on capturing signals earlier and reducing time-to-action: intelligence that updates quickly, analytics that trigger actions within OT context, and practitioners who can close the loop efficiently.

WorldView Threat Intelligence converts adversary research, malware analysis, and frontline OT investigations into operational content – adversary behaviors, prioritized IOCs, vulnerabilities, and investigation playbooks – delivered continuously as Knowledge Packs in the Dragos Platform.

The **Dragos Platform** operationalizes the Knowledge Packs while building OT context. Passive sensors perform deep protocol decoding to maintain a live asset inventory and communication map, and active agents provide targeted device/host level checks to verify configuration, user/process state, and to enrich detections – without intrusive scanning of industrial devices. Behavior-centric analytics execute against real devices, roles, and baselines, so alerts land with engineering detail and clean handoffs into existing SIEM/SOAR systems and SOC workflows.

Neighborhood Keeper introduces community-scale visibility: an opt-in anonymized collective defense network exclusive to Dragos Platform customers that produces a proprietary OT telemetry dataset at industry scale, and the largest of its kind, accelerating analytics updates and early warnings without exposing sensitive details.

Using the same content and context, **OT Watch** runs continuous hypothesis-led hunts and escalates only what's actionable. When teams want Dragos expert-led end-to-end execution, the **OT Watch Complete** service operates the Platform for customers – 24/7 security monitoring, configuration/tuning, asset enrichment, and seamless escalation – to ensure the best, fastest security outcomes from the Dragos Platform.

When incidents demand speed or depth, **Dragos Professional Services** applies the same telemetry, intelligence, and playbooks for readiness and response – closing the loop from content to confirmed action.

Outcome: Dragos delivers the #1 OT threat detection solution for industrial environments that adapts as adversaries evolve – earlier signals, less noise, faster decisions – and it plugs cleanly into the stack and workflows you already run.

Alerts can be consolidated into SIEM/SOC workflows to enable unified monitoring across all networks while enabling in-depth investigations within the Dragos Platform. They are categorized by severity and provide actionable insights into their potential operational impact. The Dragos Platform is complemented by continuous threat hunting from **OT Watch**, rapid response capabilities from **Professional Services**, and OT-specific threat intelligence from **Dragos WorldView**.

Dragos Platform detections are triggered by sensors that passively monitor industrial control system communications and conduct Deep Packet Inspection to capture asset information (vendor, family, series, model, firmware) and ICS protocol-specific operations for unmatched visibility. Lightweight collectors at edge routers, switches, and other critical remote locations provide visibility into cyber-physical systems across industrial environments.

Additionally, the Dragos Platform strengthens OT cybersecurity by managing the entire lifecycle of threats, including low-noise, high-fidelity detections driven by threat intelligence from WorldView. **Dragos Knowledge Pack (KPs)** updates regularly provide new IOCs, adversary behaviors, vulnerabilities, and tailored detection capabilities for OT systems to combat active and novel threats. WorldView IOCs integrated into existing security solutions empower SOC teams to identify early-stage adversary activity before it pivots to OT environments.

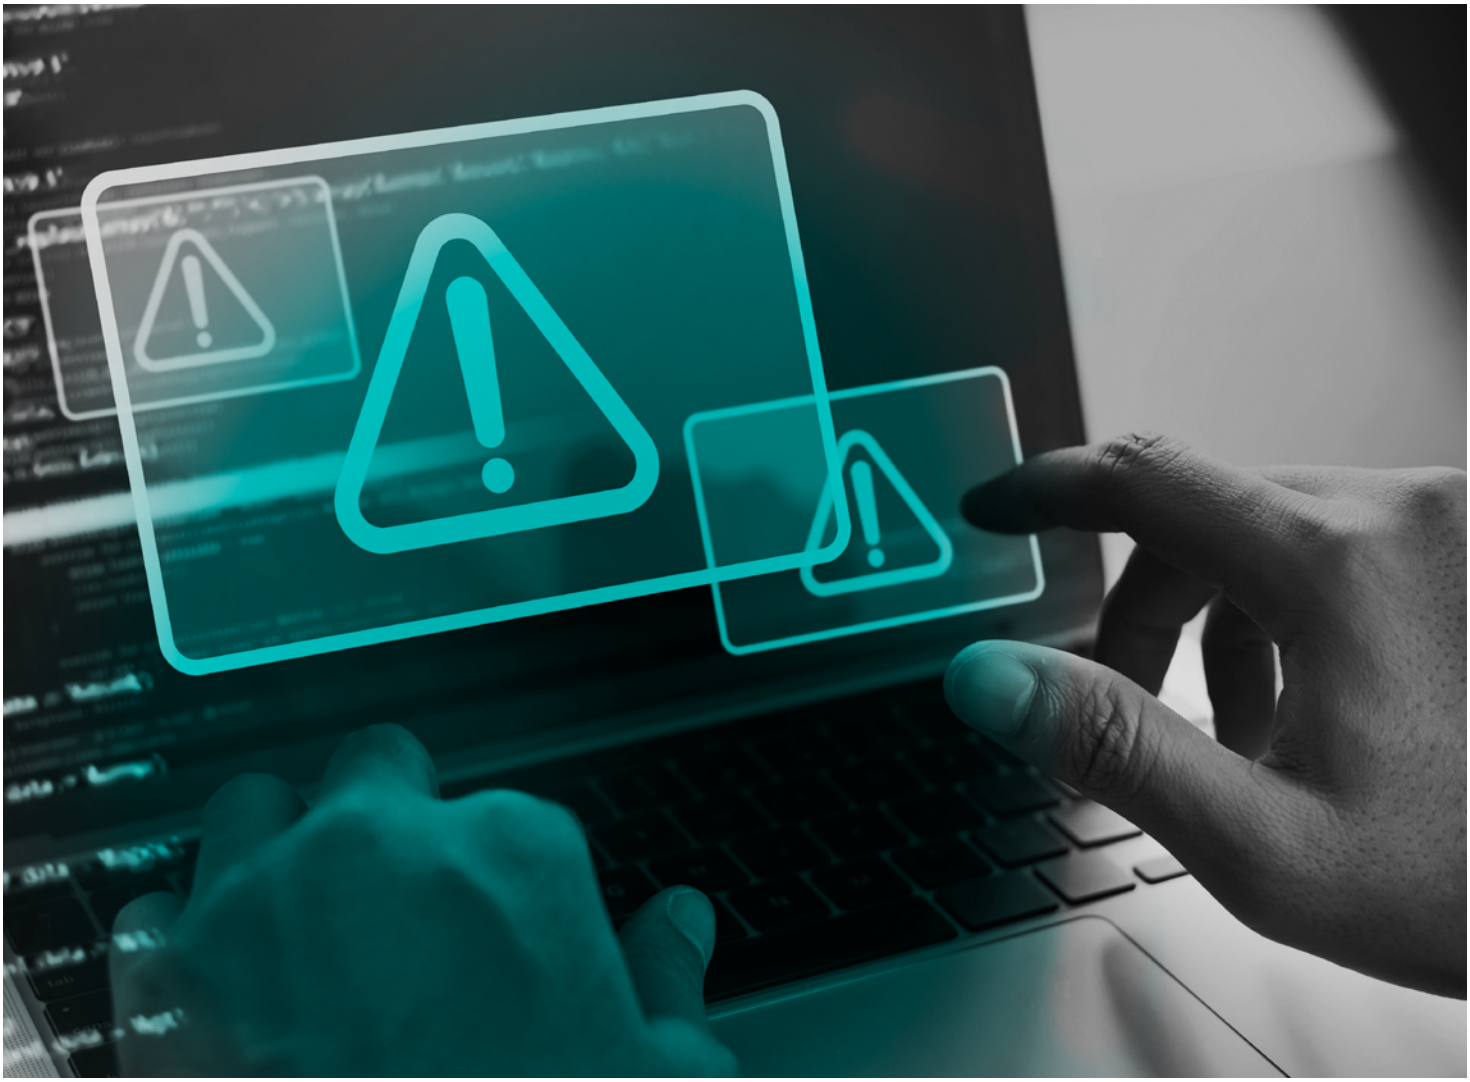

The Value of OT Cyber Threat Intelligence

Cyber threat intelligence specific to industrial control systems (ICS) and operational technology (OT) offers technical insight into the tactics, techniques, and procedures (TTPs) employed by cyber adversaries targeting these types of systems. As this technology diverges from traditional information technology (IT) setups, with unique protocols and operational requirements, specialized intelligence becomes indispensable. Domain-specific intelligence equips defenders with foresight to anticipate potential threats. It provides the information to build effective and targeted defensive strategies, ensuring robustness in the face of sophisticated attacks. While threat detections are foundational for signaling potential security anomalies, their true value is unlocked when paired with context-rich threat intelligence. Such enrichment amplifies a defender's understanding of the detected threats, illuminating adversaries' underlying intentions and modus operandi.

With this heightened perspective, defenders can prioritize responses more effectively because they understand the context of alerts related to the security landscape. In short, the defensive potential of detection and threat intelligence is more valuable than the sum of both.

The Four Types of Threat Detection

The four types of threat detection are:

Configuration Changes

Network & Protocol Anomalies

Indicators of Compromise (IOCs)

Threat Behaviors

Each category has advantages and disadvantages; no single detection type is the "best" for all environments. Instead, these categories work together to provide the most precise picture of events in an industrial control environment.

Configuration Changes

Configuration detections center on assessing and monitoring ICS components' setup, parameters, or states to ensure they align with security best practices or policies. They involve knowing system settings, user access rights, firmware versions, and other configurations beforehand. The goal is to identify and alert on deviations, misconfigurations, or vulnerabilities that can be attacked.

Configuration change detections help identify potential weaknesses in the system setup before exploiting them. That said, configuration detections may not detect active threats or attacks unless the detection alerts an attacker to change the target's configuration to exploit it. Usually, configuration detections alert on alterations to the state of a machine, network, or device from a known good state into a misconfigured state, which may represent a legitimate attack vector.

TRISIS Malware Analysis	TRISIS exploited the remote (REM) configuration mode of the safety instrumented system (SIS) controller, which allowed for remote (network) safety system logic changes to be made using its Triconex software to upload its safety logic. The attack would not have been possible if the controller's key position had been in RUN mode rather than REM mode. Regular monitoring for deviations from the established baseline might have flagged the active changes made by TRISIS. For instance, alerts could be generated for unexpected firmware changes or out-of-configuration access to the SIS controller. Detection engineers should implement detections that alert on these attempts, as they are in the Dragos Platform. Alerting on a change in configuration state is valuable. However, Dragos's assessment of this attack provided pertinent and comprehensive background for the configuration change.

Network & Protocol Anomalies

Networks and protocol anomalies rely on modeling a baseline or profile of expected system behavior and monitoring for states diverging from this baseline. By understanding the typical behavior of an ICS environment, this method flags anomalies or activities that fall outside the expected pattern. This can include unexpected communication between devices, uncharacteristic data transfers, or unexpected command sequences.

These detections effectively identify novel or zero-day threats that do not match known indicators. However, modeling detections may generate false positives if the model does not accurately capture all legitimate variations in system behavior. Additionally, the models assume that the baseline state is a "good" state, meaning that attackers are not already inside the environment. If the baseline is "trained" during a period in which the network is already compromised, then modeling detections will not likely identify attacker behavior.

TRISIS Malware Analysis	The TRISIS malware exhibited behaviors that would have been considered anomalous in a typical ICS environment, such as attempts to reprogram the SIS while the plant was operational. A system with a well-defined model of "normal" SIS controller behavior could have detected the anomalous activity of the TRISIS malware. For example, sudden or unexpected communication attempts between the engineering workstation and the SIS could be flagged as suspicious (outside of baseline). This also covers an attempt to load control logic that deviates from a baseline state. Dragos Intelligence highlighted methods for defenders to harden their networks from this attack. Dragos modeling detections that alert on this attack vector serve as a fundamental component of the recommended security posture.

Indicator of Compromise (IOC)

Indicator detections in OT cybersecurity focus on recognizing specific patterns, signatures, or attributes related to known threats or malicious activities targeting industrial control environments. This method identifies matches to previously observed or anticipated indicators of compromise (IOC), such as specific malware signatures, internet protocol (IP) addresses, or domain names associated with adversaries.

The advantage of these types of detections is that they are highly effective for detecting known threats that have previously been associated with adversaries. They are less effective at detecting unknown threats. However, with the correct prediction by a trained detection engineer, they can also be somewhat effective for detecting IOC-centric future threats. That said, indicator detections will not detect novel or customized attacks that do not match any existing or anticipated IOCs.

TRISIS Malware Analysis	After discovering the TRISIS malware, Dragos Intelligence extracted and published threat reporting that included indicators of compromise (IOC) related to the attack. These IOCs included specific malware file hashes, IP addresses, and an executable associated with TRISIS, e.g., compiled Python. The Dragos Platform operating in an ICS ecosystem is configured to alert if any network traffic or file activity corresponds to these known IOCs, signaling a potential intrusion and pointing the defender to WorldView reporting.

Threat Behaviors

Instead of focusing on known signatures or configurations, threat behavior detection emphasizes understanding and identifying adversaries' tactics, techniques, and procedures (TTPs). This method identifies patterns of malicious activity by identifying the broader behaviors and strategies employed by attackers, such as multi-stage attack sequences, lateral movement, or data exfiltration patterns. Threat behavior detections focus on what attackers must do to move between points in the ICS Cyber Kill Chain.

Threat behavior detections provide a broader perspective on threats, enabling the detection of complex, multi-step attacks. Because they are behavior-centric, not indicator-centric, they tend to be more effective at alerting on activity. However, they require deep expertise and understanding of the ICS environment, the current threat landscape, and attackers' past and potential tactics to be effective.

TRISIS Malware Analysis	TRISIS exhibited a multi-stage attack pattern: gaining initial access to the network, moving laterally to target specific systems, and finally, attempting to reprogram the safety systems. Dragos Intelligence detection engineers have worked hard to implement alerts that connect these behaviors and identify them as threats. Observing multiple stages of this threat sequence escalates the perceived severity to prompt immediate investigation. Notably, the detection team utilizes elements of past XENOTIME threat behavior to guide the creation of detections for potential future attacks, even if those attacks are unrelated to Triconex SIS controllers. However, defenders should familiarize themselves with up-to-date threat intelligence to understand how observed threat behaviors align with current threats.

The Role of Dragos Platform OT Detection Experts

Dragos Platform detections are much more nuanced than initially appear because OT detection engineers craft each detection by focusing on how it will live alongside existing alerts. We do this to ensure maximal coverage of behaviors targeted from different angles. We review existing capabilities versus requirements between each Knowledge Pack release; new detection capabilities are carefully selected for comprehensive coverage of the target activity.

For example, we might create relatively benign indicator detections, knowing that many disparate events centered around a single asset will be "rolled up" into a higher severity. An alert of this type would signal events happening together inside a window of time. The flood of indicators together is captured in a modeling detection that tells an operator, ***"Abnormal activity is occurring on this host. Investigate!"***

Detection Complexity

Dragos Platform detections fall into two categories: **atomic and composite detections**.

- **Atomic detections** often pivot off a single point of metadata or two from a single source. Indicator detections are frequently atomic detections.
- **Composite (or correlation) detections** involve two or more events from a single source or multiple sources. It is important to note that composite detections often comprise benign atomic detections when occurring in isolation but may reflect malicious behavior when combined in a certain way. Many of the best threat behavior detections are also composite detections.

Detection Telemetry

The primary sources of detection telemetry are **network events, host-based (Windows) events, and asset enrichment.**

- **Network events** reflect items visible on the network at OSI layer 2, which we can pivot from. This would include media access control (MAC) and IP addresses, certificate information, hashes, etc. Some payloads are encrypted, so we must get creative for those types of data streams and utilize detections based on traffic analysis. Packet capture (PCAP) is the source for network events.
- **Windows host-based events** provide another means of telemetry. These are produced as event logs (.log) sourced from Event Tracing for Windows and Sysmon from the host. These logs are fed into the Dragos Platform and parsed.
- **Asset enrichment** comprises a third data type that complements raw telemetry. Specifically, the Dragos Platform performs asset recognition and assigns values to hosts it sees on the network or referenced in host logs. This enrichment lets us know how the organization utilizes the asset (historian, workstation, etc.) or its location in relation to the Purdue model, among other things. Users can also arbitrarily assign their user-specified information to assets in the Dragos Platform.

These telemetry and complementary sources allow engineers to craft detections for highly unique threat behaviors. For example, we could write a detection that only fires when all the following conditions are true: we see a command from the interface on a Windows host, network traffic of a particular type originates from that host, that network traffic goes to a controller, then the controller does something unusual.

In Conclusion

Tracking and alerting on threats to OT/ICS environments is highly complex. Attackers can infiltrate a target's network, living off the land using legitimate processes, and then deliver a one-of-a-kind kill shot against systems that are rare outside of these environments. This means that there are a vast number of possibilities that defenders must consider. The Dragos team of detection engineers meets this challenge by selecting and layering detections that utilize one of the four detection types. Each detection type has its advantages and weaknesses; however, they are highly effective when applied to specific situations.

Efficacy is maximized by pairing the proper detection type with the appropriate telemetry source and crafting the detection's complexity to alert the desired activity in the proper severity. In this way, attackers' limitless complexity is met with the boundless flexibility afforded to the Dragos Intelligence detection team.

Combining threat intelligence reports with tailored detections dramatically enhances the effectiveness of cybersecurity defenses. Threat intelligence reports distill vast amounts of global threat data into nuanced insights about threat actors, their methodologies, and emerging attack vectors. The Dragos detection team has operationalized many of these insights into effective detections for the Dragos Platform, translating strategic intelligence into tactical actions.

Without the foundational understanding provided by threat intelligence reports like those in Dragos WorldView, detections risk becoming isolated alerts, providing only partial context and relevance. In tandem with alerting, they allow defenders to anticipate, identify, and counter threats with informed precision, ensuring that security resources are targeted where they are most needed and most effective. In this symbiotic relationship, threat intelligence provides the 'why' and 'how,' while detections provide the 'where' and 'what.'