

01 · EXPOSE — ASSET INTELLIGENCE & RISK ANALYSIS

From Blind Spots to Verified Posture

Asset Intelligence & Risk Analysis for Critical Infrastructure

Most organizations don't have an asset inventory problem — they have an asset intelligence problem. The data exists, scattered across tools that don't agree. What's missing is an authoritative, continuously maintained record of every asset's actual state, and a risk posture built from that record — no estimation, no black-box scoring, no manual assembly. ConsoleWorks delivers both.

THE PROBLEM

Inventory Without Intelligence Is Just a List

Every critical infrastructure security program starts with asset inventory. And nearly every one of them has the same underlying problem: the inventory reflects what the tools discovered, not what the environment actually contains. Passive network discovery tells you a device exists — not what firmware it's running, not whether its credentials have been rotated, not whether it carries a known high-severity CVE. The result is a picture that's broad, stale, and wrong exactly where the risk is highest.

This matters differently depending on who's reading the data. For IT security teams, a stale inventory means measurement and risk scoring are built on a flawed foundation. For OT operations teams, it means critical field devices — PLCs, RTUs, protective relays — exist in a separate, manually maintained spreadsheet that was accurate 14 months ago. For compliance teams, it means the evidence base for NERC CIP-010, IEC 62443, and NIST 800-82 audits is assembled on demand, from whatever data happened to exist at audit time.

The most common failure modes in critical infrastructure asset programs:

- Multiple tools each maintain their own asset view — none agree, and none are authoritative. Six tools discovering the same device produces six different records, not one.
- Passive discovery captures network presence, not configuration state. Firmware version, running configuration, and installed software are inferred from traffic — not retrieved from the device.
- New assets appear in the environment without being assessed, classified, or added to compliance scope. Discovery lag creates coverage gaps exactly when they're most dangerous.
- Inventory and risk exist in separate tools. Operators see a failed check in one system, cross-reference asset data in another, and open a ticket in a third. No direct path from exposure to remediation.
- Risk scores cannot be traced. When the score drops, no one can identify which measurement drove it — or which asset is responsible. The score becomes a number without a source.

The IT/OT boundary makes all of this harder. Engineering workstations, data historians, and HMIs running standard operating systems look like IT assets — but a historian with a direct connection to a PLC network is not an enterprise IT asset. It is an OT attack surface with an IT form factor. A conventional IT-only security posture misses the assets that carry the most operational risk. A conventional OT-only tool misses the IT-adjacent layer that attackers use to reach the control plane.

Inventory without configuration state is not an asset record. It is a device list — and device lists don't close gaps.

THE CONSOLEWORKS APPROACH

Two Capabilities. One Continuous Loop.

ConsoleWorks addresses the asset intelligence problem at its root through two integrated solutions — each sold separately, each designed to function as a complete capability, and each built so that together they form a closed, self-reinforcing loop from visibility to verified posture.

PART 1 · THIS SECTION

Asset Intelligence

Builds the authoritative, continuously maintained record of every asset — from passive tool aggregation and active direct collection. Runs Measurement Questions against every asset on every cycle. Produces the binary Pass/Fail results that drive everything downstream.

Foundation: the inventory is the accuracy of the score.

PART 2 · NEXT SECTION

Risk Analysis

Takes those Pass/Fail results and aggregates them through the SCF control hierarchy — producing a continuously updated, fully traceable risk posture across three operational lenses: security, compliance, and operations. Surfaces prioritized gaps and connects each one to a direct fix path.

Output: a posture you can defend, and a fix path you can act on.

Neither capability requires the other to function — but each is most powerful when both are in place. An asset inventory without risk scoring is a list. A risk score without an authoritative inventory is an estimate. Together, they are the foundation of the ConsoleWorks platform and the starting point for every other capability: Secure Remote Access, Credential Management, Configuration and Change Management, and Intelligent Event Monitoring all operate on what Asset Intelligence and Risk Analysis expose.

PART 1 OF 2 · 01 · EXPOSE

Asset Intelligence: Building the Foundation

An authoritative, continuously maintained record of every asset — built from every source you have, or from ConsoleWorks alone if you're starting from scratch.

HOW THE INVENTORY IS BUILT

Two Layers. One Authoritative Record.

ConsoleWorks Asset Intelligence is built on a fundamental architectural principle: the most trusted data should always win. That principle is implemented through two complementary collection layers that feed every asset record — passive aggregation from existing tools, and active collection directly from the device itself.

Most environments start with existing security and operational tools — discovery platforms, vulnerability scanners, firewalls, PAM systems, and operational CMDBs — each of which knows something about the assets in the environment. The challenge is that each tool knows something different, reports it differently, and cannot see what the other tools see. ConsoleWorks treats all of this existing data as an input layer — not a replacement for it.

Layer 1 — Passive Breadth from Existing Tools

Tool Data Collectors (TDCs) connect via API to existing security and operational tools — Tenable, Palo Alto Panorama, Dragos, firewalls, CMDB exports, and more. Each TDC pulls structured asset metadata from its source tool and delivers it to the ConsoleWorks normalization engine. Wide coverage from day one. No production impact. No changes to existing tools. If your environment already has discovery and vulnerability data, ConsoleWorks consumes it — and adds what those tools cannot provide.

Layer 2 — Active Authority from the Device Itself

Where passive collection captures what existing tools already see, ConsoleWorks Continuous Configuration Monitor (CCM) collects configuration state directly from each managed device via its native protocol. Firmware version, running configuration, active accounts, installed software — retrieved directly from the source. This is not inferred from network traffic, not approximated from scan data. It is the authoritative state of the device on every collection cycle. Active collection data takes precedence wherever it exists.

Active Collection vs. Passive Inference

Passive discovery tools infer device state from what they observe on the network. ConsoleWorks CCM connects directly to the device via native protocol and retrieves what's actually running. No estimates. No approximations. The difference matters most for OT field devices — where passive inference is unreliable and the cost of acting on wrong data is measured in operational downtime, not IT tickets.

SUPPORTED DATA SOURCES

The ConsoleWorks TDC framework ingests data from a broad range of source types — organized by collection method and trust level. Organizations start with whatever sources they have and build from there. ConsoleWorks does not dictate which tools must be in place; it integrates with what exists and supplements where gaps remain.

Active Collection (Highest Trust)

ConsoleWorks native configuration collectors interact directly with managed assets to retrieve configuration files, device settings, system baselines, and account inventories. Data is timestamped, tamper-evident, and authoritative. Takes precedence over all passive sources for conflicting fields.

Secure Remote Access Logs

ConsoleWorks captures detailed session records for all remote interactions with assets — access frequency, user actions, commands issued, and adherence to access control policies. Essential for answering Measurement Questions related to privileged access monitoring and audit trail integrity.

Passive Asset Discovery Tools

Network discovery and asset inventory tools detect devices connected to the environment. ConsoleWorks ingests this data to create new asset records, validate existing inventory, and provide metadata that drives downstream risk analysis and compliance measurement.

Vulnerability Scanners

Scan data provides critical insight into known software flaws, CVE exposure, and misconfigurations. Feeds patch status, configuration hardening, and vulnerability-related Measurement Questions. Cross-referenced against the asset record to surface risk in operational context.

Intrusion Detection Systems (IDS)

IDS alerts related to suspicious activity, potential breaches, or unauthorized behaviors are ingested and correlated to specific asset records. Security teams can assess exposure, verify whether compensating controls are in place, and trigger SRA sessions to investigate directly from the alert context.

Spreadsheets, CMDB & Manual Entry

In environments where data still resides in structured files or must be manually entered, ConsoleWorks ingests and normalizes this data alongside automated sources. Manual records are user-attributed and timestamped for traceability — treated as evidence, not just background context.

THE NORMALIZATION ENGINE

Mapping Rules: Turning Multiple Sources into One Record

Collecting data from multiple sources is only half the challenge. The harder problem is reconciling those sources into a single, authoritative asset record when they disagree — and they always disagree. Tenable may report a firmware version that Dragos reports differently. A manually maintained spreadsheet may have an asset name that doesn't match the hostname the device returns via active collection. Six tools discovering the same dual-homed device may produce six different records for what is physically one asset.

ConsoleWorks resolves these conflicts through Mapping Rules — a configurable transformation engine that defines exactly how each tool's data maps to the unified asset schema, and which source wins when they conflict. Priority is set per-field: a Mapping Rule for firmware version might favor CCM active collection data over Tenable scan data, which is in turn favored over a manual spreadsheet entry. The result is a single record that always reflects the most authoritative data available, from the tool best positioned to provide it.

- 1 Collect — Ingest from all configured sources**
 TDCs pull structured data from every connected tool on a configured schedule. Spreadsheet imports and manual overrides supplement automated collection where automated sources cannot reach.
- 2 Correlate — Resolve multi-source, dual-homed conflicts**
 Key Scripts determine whether incoming data creates a new asset record or updates an existing one. A device with multiple IP addresses — common in OT environments where assets span zone boundaries — is correlated into a single record. Six tools discovering the same device produces one record, not six.
- 3 Map — Apply priority-ordered field mapping**
 Mapping Rules write each source field to the appropriate destination field in the unified asset schema. When multiple sources report the same field, Priority settings determine which source wins. CCM active collection data carries the highest default trust level.
- 4 Normalize — Standardize values across sources**
 Translation Scripts transform source values into consistent destination values. "Microsoft Windows Server 2019 Standard," "Win2019," and "Windows Server 2019" from three different tools all become the same normalized value in the asset record. No manual reconciliation.
- 5 Deliver — One authoritative record per physical asset**
 The resulting record is the unified view of the asset — built from the most authoritative data available across every connected source. It updates automatically on every collection cycle, without operator intervention. Critically, ConsoleWorks preserves data source attribution throughout — every field value retains a record of which tool provided it and when, making every assertion about asset state auditable and attributable. Native active collection data is timestamped and tamper-evident, providing the strongest foundation for compliance verification and drift detection.

THE ASSET RECORD

What ConsoleWorks Knows About Every Asset

Every device in the ConsoleWorks inventory is represented as a structured asset record — a unified view built from the most trusted data available across all connected sources. The record captures not just what the device is, but what state it's actually in: hardware identity, firmware version, running configuration, active accounts, installed software, network interfaces, patch status, antivirus coverage, and open ports. Each subcomponent is populated by whichever connected source has the most authoritative view of that particular field.

Assets are classified through a Category → Type → Function hierarchy — Computer, Controller, Network Device, Plant Device — with specific types and operational functions defined within each. A historian classified as Computer → Server → Historian carries different compliance obligations than a PLC classified as Controller → RTU → Production Control, even if they share an IP subnet. This classification is how ConsoleWorks scopes measurement, compliance mapping, and risk weighting to the operational reality of each asset — not a generic IT model applied uniformly across the environment.

The asset record is not a flat list of discovered attributes. It is structured around a trust hierarchy: when multiple sources report conflicting data for the same field — one tool reports firmware v3.1.1, another reports v3.1.2 — ConsoleWorks applies a configurable precedence order to resolve the conflict. Active collection via native protocol takes highest precedence. The result is a record that reflects what is actually running on the device, not what a passive scan most recently observed or what a spreadsheet was last updated to show.

Every field in the asset record carries a collection timestamp. The record does not reflect a single snapshot taken at a point in time — it reflects the most recently verified state for each individual attribute, updated as frequently as collection runs. A firmware version collected 12 minutes ago is treated as current. A patch status last confirmed 18

days ago is flagged for staleness. The record makes its own freshness visible, so operators and auditors know exactly how current each attribute is without asking.

The completeness of the asset record is the ceiling for everything that follows. A Measurement Question that asks whether AV definitions are current can only return a reliable answer if the record contains a verified, recent AV definition timestamp. A risk score that weights patch status can only reflect operational reality if the record contains actual patch data – not an estimate or a default. This is why ConsoleWorks begins with inventory: every capability that follows is only as accurate as the record it operates on.

ASSET RECORD — AMARILLO-RTU-04 · SCHNEIDER ELECTRIC ECOSTRUXURE · CONTROLLER / RTU / PRODUCTION CONTROL			
FIRMWARE v3.1.2 — Drift Detected	PATCH STATUS 2 Critical Unpatched	ANTIVIRUS Definitions 47 days old	DEFAULT CREDENTIALS Detected
SITE Amarillo Plant — Zone 3	LAST ACTIVE COLLECTION 14 min ago	UNAUTHORIZED SOFTWARE 1 Detected	ZERO TRUST POLICY Least Privilege — Review Required

[Open SRA Session](#)

Measurement Results — AMARILLO-RTU-04		
AV_INSTALLED	Is antivirus software installed and recorded?	PASS
AV_DEFS_CURRENT	Are AV definitions current within the required window?	FAIL
PATCH_CURRENT	Patched within the required compliance window?	FAIL
DEFAULT_CREDS_REMOVED	Default credentials removed from all accounts?	FAIL
SESSION_RECORDED	All privileged sessions recorded against an audit record?	PASS
FIRMWARE_BASELINE	Firmware matches approved baseline for this device class?	FAIL

Measurement Questions are fully configurable to your organization's controls framework, compliance requirements, and operational priorities. The results shown above are examples of the measurement categories available – not a fixed set.

MEASUREMENT QUESTIONS

The Inventory That Measures Itself

Most asset inventory tools stop once they have a record. ConsoleWorks goes further: every asset in the inventory is continuously evaluated against the Measurement Questions applicable to it, producing a binary Pass or Fail result on every collection cycle. This is not a point-in-time assessment. The status of every asset always reflects current state as of the last collection run.

Measurement Questions are configurable to the organization's specific controls framework – NERC CIP, NIST 800-82, IEC 62443, TSA pipeline directives, or any combination. Questions can target any field in the asset record: firmware version against an approved baseline, AV definitions currency against a configured maximum age, password policy

compliance against an organizational standard. The binary result removes subjectivity from the measurement — a device either meets the control or it doesn't. No partial credit. No judgment calls.

For assets managed through ConsoleWorks Secure Remote Access, a failed measurement is not just a flag — it is an action item. Operators can open a direct session to the affected device from the asset record, apply the fix, and the measurement re-runs on the next cycle. Verified closure — not ticket closure. The gap is confirmed resolved when the measurement returns Pass, not when someone marks the ticket done.

Zero Trust Starts at the Inventory Layer

A Zero Trust architecture assumes no device is trusted by default. ConsoleWorks enforces that principle from the inventory layer: every asset is continuously measured against Least Privilege controls — credential management, session authorization, access scope — and any deviation surfaces immediately on the asset record. Assets that have not been measured, classified, or brought under SRA management are not assumed to be in compliance. They are flagged as unknown scope — which is a finding in itself.

CONTROL TRANSFORMATION

From Abstract Policy to Measurable Question

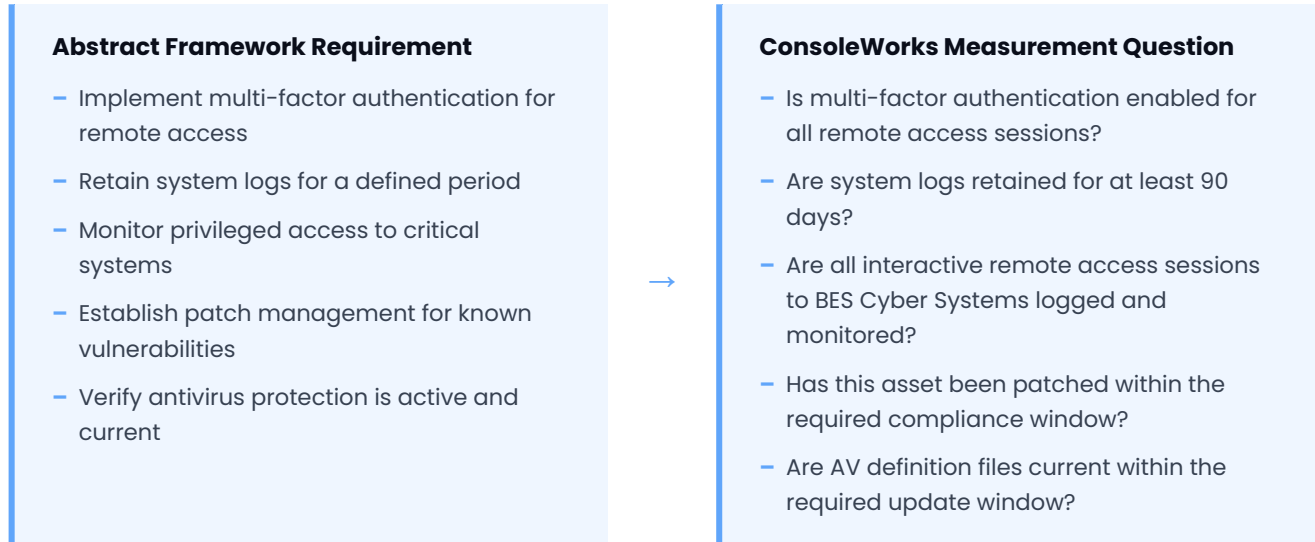
Every successful security and compliance program begins with a harder problem than most platforms address: translating abstract regulatory requirements into statements that can actually be evaluated. NERC CIP, NIST, and IEC 62443 contain requirements like "establish access control mechanisms" or "monitor for unauthorized activity" — directives that leave room for interpretation and have historically been addressed through checklists, self-attestation, or periodic manual assessment.

ConsoleWorks formalizes this translation through Measurement Question design. Each selected control requirement is expressed as a binary, testable question — one that can be evaluated automatically against actual asset data, answered definitively on every collection cycle, and mapped to the specific asset or assets it applies to. The binary format removes ambiguity: a control either passes or fails. No partial credit. No interpretation. No self-attestation. The system evaluates, the asset answers, and the record reflects the result.

This transformation is what makes ConsoleWorks a compliance enforcement platform rather than a compliance tracking tool. The control is not merely documented — it is operationalized. It is measured against every applicable asset, on every cycle, with a traceable evidence chain from the requirement to the collection event.

The translation from policy to question is not a one-time exercise. Organizations operating under NERC CIP, TSA pipeline directives, or IEC 62443 face evolving regulatory language, scope expansions, and periodic standard revisions. Measurement Questions are versioned and configurable — when a regulation changes, the question is updated once and the change propagates to every asset in scope. There is no manual rework of assessment spreadsheets or audit narratives. The updated measurement runs on the next collection cycle and produces updated evidence automatically.

For OT environments, this translation carries additional weight. Many OT-specific standards — NERC CIP-007, IEC 62443-3-3, NIST 800-82 — were written with IT-centric assumptions about what can be measured, how frequently, and through what mechanism. ConsoleWorks extends the Measurement Question model to the full OT asset landscape: PLCs, RTUs, protective relays, HMIs, and historians are measured through the same question framework as IT endpoints, using native OT protocols rather than agent-based mechanisms. The binary result applies equally — a device either meets the control or it doesn't, regardless of whether it runs Windows or a proprietary embedded operating system.

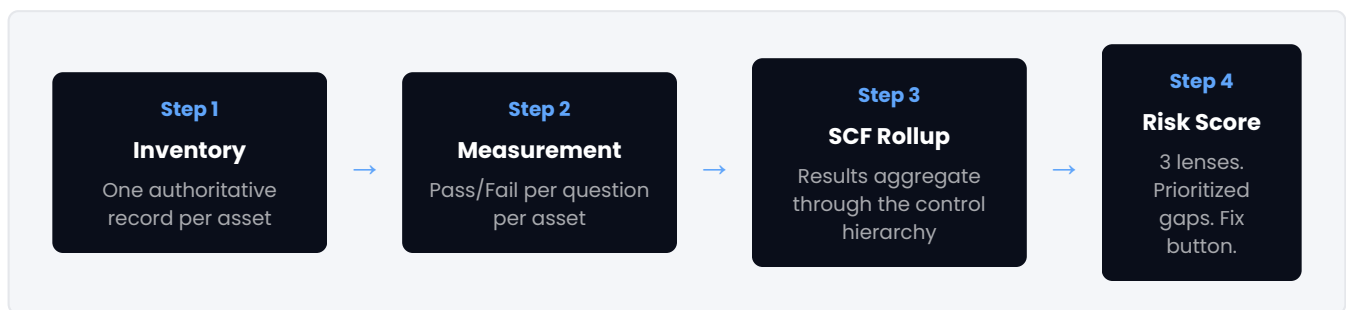


This question-based format offers three structural advantages over checklist-based compliance approaches. First, it eliminates ambiguity – the control is either met or not based on actual system state, not someone’s interpretation of the requirement. Second, it enables automation – ConsoleWorks evaluates each question against known asset data as frequently as the organization requires, not just during audit periods. Third, it provides traceability – every question links directly to its originating control requirement, enabling organizations to demonstrate not only that they addressed the regulation but exactly how, on which assets, and at what time.

FROM ASSET INTELLIGENCE TO RISK POSTURE

Part 1 Built the Foundation. Part 2 Scores It.

At this point in the ConsoleWorks lifecycle, three things are in place: a continuously maintained asset inventory built from every available source, a set of Measurement Questions configured to your organization’s controls framework, and a stream of binary Pass/Fail results – one per question, per asset, per collection cycle. These results are the input to Risk Analysis.



Risk Analysis does not introduce new data. It does not run separate assessments or consult external threat feeds to produce its scores. Every score it generates traces directly back to an inventory record and a specific collection event. The accuracy of the score is the accuracy of the inventory – which is why Asset Intelligence is the foundation, not a prerequisite.

PART 2 OF 2 · 01 · EXPOSE

Risk Analysis: From Data to Organizational Posture

The asset inventory is built. The Measurement Questions are running. Every asset returns a Pass or Fail. Risk Analysis is what ConsoleWorks does with those results — aggregating, weighting, scoring, and surfacing them as an actionable posture for every stakeholder that needs to act on it.

HOW SCORING WORKS

From Binary Results to Organizational Risk Posture

Risk Analysis does not score risk from estimates, questionnaires, or vendor-supplied threat intelligence. It scores risk from actual measurements — the binary Pass/Fail results that ConsoleWorks produces on every collection cycle for every asset in scope. Those results are the inputs. Risk Analysis is the processing engine that transforms them into a continuously updated posture.

The processing engine is built around the Secure Controls Framework (SCF) — a comprehensive controls framework that maps to over 100 global regulations and standards. Every Measurement Question is mapped to an SCF sub-control. Sub-controls roll up to controls. Controls roll up to domains. Domains roll up through the organization's asset group hierarchy: individual asset → site → region → organization → fleet. Every level recalculates automatically on every measurement cycle.

The weights applied at each level are configured by the organization. Asset criticality, site classification, and control domain priority all influence how a single FAIL at the device level propagates upward through the hierarchy. ConsoleWorks applies those weights continuously — the risk posture reflects the organization's actual priorities, not a generic model's opinion of what matters.

FLEET	All Organizations · All Sites	82%
ORG	Acme Energy Corp	79%
SITE	Site Alpha · Plant Operations Reporting: 47 / 50 assets	78%
DOMAIN	Endpoint Security (ES) NIST: SI	71%
CONTROL	ES-03 · Malicious Code Prevention NIST: SI-3	63%
SUB-CONTROL	ES-03.1 · AV Currency	41%
MEASUREMENT	AV_DEFS_CURRENT · FAIL	—

One FAIL at the measurement level propagates upward through every level on the next collection cycle — recalculating sub-control, control, domain, site, organization, and fleet scores automatically. The chain is unbroken and fully traceable in either direction.

Every Score Has a Source

Fleet score → domain → control → sub-control → measurement → asset → collection event. The chain is unbroken. No estimates. No black-box algorithms. If the score dropped, ConsoleWorks can tell you exactly which measurement drove it, on which asset, at what time. That traceability is what makes the posture defensible — for the board, for regulators, and for the operations team trying to close the gap.

One Measurement. Every Framework.

Critical infrastructure organizations operate under multiple compliance frameworks simultaneously — NERC CIP for bulk electric, TSA directives for pipeline, IEC 62443 for industrial control systems, NIST 800-53 for federal systems, NIST CSF as a baseline. Each framework has its own control language, its own audit evidence requirements, and its own reporting obligations. Most organizations manage these separately — maintaining separate assessment processes, separate evidence libraries, and separate gap analyses for each framework they're subject to.

ConsoleWorks eliminates that redundancy through the Secure Controls Framework (SCF) — a framework designed specifically for this crosswalk problem. Every Measurement Question is mapped to an SCF sub-control once. The SCF sub-control carries its mappings to every regulation and standard it satisfies. When a measurement runs, its result simultaneously satisfies — or fails — the requirements of every framework it touches. Configure the measurement once. ConsoleWorks handles the crosswalk everywhere.

Measurement Question	NERC CIP	NIST 800-53	IEC 62443	NIST CSF	SOC 2
AV_DEFS_CURRENT	CIP-007 R3	SI-3	SR 3.2	PR.DS-1	CC6.8
PATCH_CURRENT	CIP-007 R2	SI-2	SR 2.4	PR.IP-12	CC7.1
DEFAULT_CREDS_REMOVED	CIP-007 R5	IA-5	SR 1.5	PR.AC-1	CC6.1
SESSION_RECORDED	CIP-005 R2	AU-14	SR 2.8	PR.AC-3	CC6.3
FIRMWARE_BASELINE	CIP-010 R1	CM-6	SR 7.6	PR.IP-1	CC6.7

The audit evidence generated by these measurements is not assembled pre-audit. It is generated continuously — stored, timestamped, and traceable to the source collection event — ready to produce on demand for any framework, any scope, any time window.

REPORTING SCOPE

What You Don't Know About What You're Missing

Risk scores are only meaningful in the context of what they're based on. A posture score calculated from 847 of 1,200 expected assets is a fundamentally different number than one based on 1,200 of 1,200 — but most risk platforms don't make that distinction. They score what's reporting and leave the rest invisible. The unscored assets don't drag the number down. They simply don't appear.

ConsoleWorks surfaces Reporting Scope alongside every score in the rollup hierarchy. Every site score, domain score, and fleet score shows the number of assets actively returning measurement data versus the number expected to be in scope. Assets that are in the inventory but not reporting measurement data are flagged — not ignored. The posture reflects the full picture of what is known and what is not yet covered.

Reporting Scope — Site Alpha · Plant Operations

47 of 50 assets reporting · Score: 78% · 3 assets in inventory scope not returning measurement data — flagged for review

This transparency is particularly important for Zero Trust programs, where the assumption is that every asset must be in scope and accounted for. An unscored asset is not a low-risk asset. It is an unknown asset — and unknown assets are the ones that appear in incident post-mortems.

THREE LENSES ON ONE DATA SET

The Same Measurement. Three Different Conversations.

One measurement result means something different depending on who's reading it and what decision they need to make. The CISO needs to know whether the posture is real and improving. The compliance team needs to know whether it satisfies the framework. The operations team needs to know which device is broken and how to get to it. ConsoleWorks surfaces the same underlying measurement data through three operational lenses — so each team sees exactly what they need, without translating from a platform built for someone else.

Security Posture	Compliance Mapping	Operational Gaps
<p>CISO & SECURITY LEADERSHIP</p> <ul style="list-style-type: none"> ✓ Continuous posture — not quarterly assessments ✓ Fleet-to-asset drill-down from org-level score ✓ Trend indicators: improving, stable, or degrading ✓ Gaps ranked by organizational impact ✓ No black-box algorithms — every number has a source ✓ Board-ready trend data traceable to actual remediation 	<p>COMPLIANCE & AUDIT TEAMS</p> <ul style="list-style-type: none"> ✓ Compliance score per framework from the same data ✓ NERC CIP, NIST, IEC 62443, TSA scored simultaneously ✓ Framework crosswalk applied automatically via SCF ✓ Continuously current — reflects latest measurement cycle ✓ Audit evidence already generated — not assembled on demand ✓ Trend history: demonstrable improvement over time 	<p>OPERATIONS & ENGINEERING TEAMS</p> <ul style="list-style-type: none"> ✓ Failed checks visible at the individual device level ✓ Ranked by operational impact — most critical first ✓ Fix button opens SRA session directly from the gap view ✓ Score updates on the next cycle after fix — verified closure ✓ Operational risk dimension: stability and availability framing ✓ Least Privilege controls surfaced per device

THE FIX PATH

From the Score to the Session — Without Leaving the Platform

Most risk platforms stop at detection. They surface the gap, assign it a severity, and leave the remediation workflow to a separate tool — a ticketing system, a PAM platform, a terminal emulator. The result is a broken chain: the operator sees the gap in the risk view, opens a ticket, logs into a different system, connects to the device, makes the change, and closes the ticket. The risk score updates the next time the scan runs — which may be days away. And the ticket status, not a measurement result, is the evidence of closure.

ConsoleWorks closes this chain. The Fix button in the risk view is not a link to a ticketing system — it opens a direct SRA session to the affected device, from within the same platform where the gap was identified. The operator connects via ConsoleWorks SRA using its Zero Trust protocol-break architecture: every session is established through ConsoleWorks (Session 1: operator to ConsoleWorks; Session 2: ConsoleWorks to device). The operator never has a direct network path to the endpoint. Every keystroke is logged. Every GUI session is recorded as a full screen capture. Every change is tied to the session record and the asset record.

When the fix is applied, the Measurement Question re-runs on the next scheduled collection cycle. The score updates when the measurement returns Pass — not when someone closes the ticket. That verified closure is what makes the posture defensible. It is also what enforces Least Privilege at the remediation layer: access to the device is granted only for the duration of the session, through a verified identity, with a complete audit record attached to the asset that carries the finding.

The Closed-Loop Remediation Model

What the Closed-Loop Model Guarantees

1	Control Failure Identified Measurement Question returns FAIL on a specific asset. Surfaces in the risk view ranked by organizational impact.
2	Click-to-Remediate Operator selects Fix from the risk view. No tool switch. No ticket. No copy-paste of device address.
3	Secure Session to Asset Initiated ConsoleWorks SRA establishes a Zero Trust, protocol-break session. User identity verified. Session authorized for this device. No direct network path to endpoint.
4	Remediation Action Executed Operator applies the fix: patches the system, disables unauthorized accounts, updates configuration, restarts affected services, rotates credentials, or updates certificates.
5	Control Re-Checked & Closed Measurement Question re-runs on the next scheduled cycle. Score updates when the measurement returns PASS — not when the ticket is closed.
6	Evidence Logged to Audit Trail Every keystroke, command, and system response from the session is captured. The remediation action is tied to the user identity, the asset record, and the control finding that initiated it.

- **Traceability** — From the control question to the root cause to the remediation action to the re-check. Every step is documented without manual entry.
- **Audit readiness** — Complete documentation of what was done, by whom, and when — attached to the asset record and the finding that prompted it. Not assembled. Already there.
- **Reduced dwell time** — Operators resolve issues without waiting for external escalations, tool integrations, or ticket workflows. Detection and remediation happen in the same session.
- **Compliance resilience** — Operational response is embedded directly into the measurement and reporting cycle. The posture reflects what was actually fixed — not what was marked done.
- **Zero Trust enforcement** — Least Privilege is applied at the session layer: access granted for the specific operation, to the specific device, for the duration of the session, with a complete record attached.

ConsoleWorks is not a passive monitoring system. It is an active compliance enforcement and remediation platform — the same environment used to identify a control failure is also used to remediate it. This eliminates the inefficiencies and risk of hand-offs between siloed tools and teams.

Zero Trust Architecture at the Remediation Layer


ConsoleWorks SRA implements a protocol-break architecture that enforces Zero Trust principles at the access layer: no user ever has a direct network path to a managed endpoint. Session 1 terminates at ConsoleWorks. Session 2 originates from ConsoleWorks. The user's credentials never traverse the OT network. Every session is authenticated, authorized, recorded, and tied to both the user identity and the asset record. Least Privilege is enforced by session — access is granted for the specific operation, to the specific device, at the specific time. This architecture applies to every remediation action initiated from the risk view, not just to IT-adjacent assets.

THE PLATFORM LOOP — CLOSED

From Exposure to Evidence: One Unbroken Chain

At this point the complete Asset Intelligence and Risk Analysis lifecycle is in place. Asset Intelligence continuously maintains the inventory. Measurement Questions run on every cycle and surface Pass/Fail results at the asset level. Risk Analysis aggregates those results through the SCF hierarchy, scores the posture across three lenses, ranks the

gaps by organizational impact, and provides a direct fix path to every managed device. The fix is applied through a Zero Trust SRA session. The measurement re-runs. The score updates. The evidence was generated throughout — not assembled after the fact.



What the next section documents is why this architecture is distinct from every other approach in the market — and what it means for the organizations running the environments where failure has consequences that go beyond a security incident.

The table below shows what each team can now do — and what they no longer have to do — once the full loop is operational.

SECURITY & CISO	COMPLIANCE & AUDIT	OPERATIONS & ENGINEERING
<p>A continuously updated posture score across every domain, site, and asset — not a quarterly assessment. Every gap is ranked by impact, traceable to the control failure, and fixable from the same screen without a ticket.</p> <p>No more: Manual gap analysis, tool switching, or assembling evidence after a review.</p>	<p>Every Measurement Question maps to one or more framework controls via the SCF crosswalk. A single pass of collection produces evidence for NERC CIP, NIST 800-53, IEC 62443, TSA, and more simultaneously — no parallel audit programs required.</p> <p>No more: Separate control frameworks, manual evidence collection, or relying on screenshots as proof.</p>	<p>Inventory that reflects the operational state of every device — not just what was discovered at installation. Patch status, configuration, firmware, and open ports are current because they are collected on every cycle, not polled on demand.</p> <p>No more: Spreadsheet reconciliation, stale CMDB data, or unknown device drift between audit windows.</p>

The Evidence Standard

Continuous collection means continuous evidence.

Every collection cycle generates a timestamped record of what was found, which tool supplied it, and whether the control passed or failed. Evidence is native to the measurement — not assembled from screenshots, exports, or notes after the fact.

Tamper-evident by design.

Native collection through CCM captures configuration and session data directly from the device over a Zero Trust SRA session. The record cannot be modified without breaking the chain — making it defensible for NERC CIP, NIST, and IEC 62443 audits without supplemental documentation.

WHY CONSOLEWORKS

What Makes This Different

A risk score is only as good as the data behind it, the traceability through it, and the action it enables. ConsoleWorks is designed around all three — from the inventory layer through the score to the fix path.

COMPETITIVE DIFFERENTIATION

Not All Risk Scores Are Created Equal

Risk and inventory platforms in critical infrastructure tend to fall into three categories: discovery tools that build asset lists without intelligence, risk platforms that score from estimates rather than measurements, and GRC tools that aggregate compliance posture from manual inputs. ConsoleWorks is none of these. It is the only platform that builds the inventory with active collection, scores it from binary measurements, exposes the score through three operational lenses, and connects the gap directly to a Zero Trust remediation path — on the same platform, without a tool switch.

Capability	Discovery / Inventory Tools	Risk Platforms / GRC Tools	ConsoleWorks
Asset collection method	Passive only — network-based inference	Passive or questionnaire-based	✓ Passive aggregation + active device collection
Single authoritative record	Per-tool view — no normalization	Not applicable — risk only	✓ Multi-source normalization, conflict resolution by priority
IT + OT unified inventory	IT strong; OT field devices limited	IT-centric; OT requires separate tool	✓ Single inventory spanning IT and OT assets
Risk score source	No risk scoring	Estimates, questionnaires, or scan summaries	✓ Binary measurement results — Pass/Fail per asset per control
Score traceability	Not applicable	Partial — score to control, not to specific asset measurement	✓ Fleet → domain → control → sub-control → measurement → asset → collection event
Multi-framework compliance	Limited or manual mapping	Framework-specific tools; redundant effort per framework	✓ 100+ frameworks via SCF crosswalk — one measurement satisfies all applicable controls
Reporting Scope transparency	Scores what's reporting — gaps invisible	Rarely surfaced alongside score	✓ Reporting Scope displayed at every rollup level — unscored assets flagged
Continuous posture	Not applicable		✓ Updates on every measurement cycle —

		Periodic — scan cadence or annual assessment	reflects current state, not last quarter
Fix path from score	Detection only	No remediation capability	✓ SRA session opens directly from the gap — no tool switch, no ticket
Verified closure	Not applicable	Ticket-based — no measurement verification	✓ Measurement re-runs on next cycle — score updates when Pass is confirmed
Zero Trust architecture	Not included	Not included	✓ Protocol-break SRA — no direct path to endpoint, Least Privilege enforced by session
Agents required	Often required for deep collection	Often required for endpoint measurement	✓ Zero agents on any managed endpoint — ever

WHO IT SERVES

Built for the Regulatory Reality of Critical Infrastructure

ConsoleWorks is deployed across the sectors where the stakes of getting asset intelligence wrong are measured in operational consequence — power generation and transmission, oil and gas pipeline operations, water and wastewater systems, nuclear facilities, defense and federal infrastructure, and critical healthcare. In each sector, the regulatory frameworks are different. The asset types are different. The operational constraints are different. The common thread is that discovery tools designed for enterprise IT cannot handle the OT environment, and risk platforms designed for compliance cannot surface the operational context that operations teams actually need.

Sector	Primary Frameworks	ConsoleWorks Application
Energy & Utilities	NERC CIP-005 / 007 / 010 / 013	CIP-010 baseline inventory built continuously from active collection. CIP-007 R3 AV and patch evidence generated on every measurement cycle. Single inventory spanning BES Cyber Assets and IT assets in the ESP — no separate tools, no manual reconciliation.
Oil & Gas Pipeline	API 1164, TSA Pipeline Directives	SCADA and OT asset inventory built from field device collection. Vendor session documentation tied to asset records via SRA. TSA security directive compliance evidence generated continuously from the same measurement data.
Water & Wastewater	AWIA 2018, NIST 800-82, EPA	AWIA risk assessment inventory built without agents on treatment system controllers and SCADA endpoints. Configuration monitoring across heterogeneous OT environments — PLCs, HMIs, historians, and supervisory systems in a single inventory.
Nuclear	10 CFR 73.54, NRC RG 5.71	Agentless Critical Digital Asset (CDA) monitoring. NRC-accepted session logs via ConsoleWorks SRA. Inventory

		and measurement covering CDAs and non-CDAs in a single platform with Zero Trust access controls enforced by session.
Defense & DIB	CMMC 2.0, NIST 800-171, DFARS	Access control documentation, privileged session logging, and Least Privilege enforcement via SRA. CMMC Practice documentation generated continuously from measurement results — not assembled at assessment time.
Healthcare & Federal	HIPAA, FISMA, HITRUST	Deployed in 150+ active military hospitals. Medical device and clinical system inventory alongside IT infrastructure — single platform, unified posture, continuous compliance evidence.

THE PLATFORM STORY

Asset Intelligence and Risk Analysis Don't Stand Alone

Asset Intelligence builds the inventory that defines measurement scope. Risk Analysis scores it through the SCF hierarchy and surfaces the gaps. But the gaps have to be closed — and closing gaps in OT environments requires access to devices that don't accept standard remote access tools, in environments where every connection must be authorized, recorded, and audited.

ConsoleWorks Secure Remote Access provides the fix path. The same connection that CollectsConfiguration data for active collection is the connection that remediates the gap — no tool switch, no separate PAM platform, no vendor VPN. Credential Management rotates credentials directly on the device — including PLCs, RTUs, and protective relays that conventional IT PAM tools cannot reach — via native protocol through SRA. Configuration and Change Management captures device configuration before and after every session automatically, generating a field-level diff tied directly to the session record and the asset that carries the finding.

Intelligent Event Monitoring maintains a continuous log stream from every managed device. IEM events surface in the context of the asset record — correlated with the device's current measurement status, its access history, and any open gaps in the risk view. When an event warrants immediate investigation, an SRA session to the affected device is one click away — from the same platform that generated the alert.

The Enforce mandate sustains what Expose surfaces. Measurements re-run on schedule. Scores update automatically. Compliance evidence accumulates continuously — already there when the auditor asks, not assembled the week before. Gaps that reopen after remediation surface on the next measurement cycle. The posture is not a snapshot. It is the current state of the environment, continuously verified.

The ConsoleWorks Platform: EXPOSE. ELIMINATE. ENFORCE.

- **01 • Expose** — Asset Intelligence builds the inventory. Risk Analysis scores it. Everything starts here — because you cannot measure, remediate, or report on what you don't know you have.
- **02 • Eliminate** — Secure Remote Access, Credential Management, and Configuration & Change Management close the gaps that Expose surfaces — without agents, without disruption, and with a complete audit record of every action taken.
- **03 • Enforce** — Continuous Measurement re-runs every check on schedule and updates the score when the fix is verified. Compliance Reporting generates audit-ready evidence for any framework, any time, without assembly. The 100+ framework crosswalk means evidence that was never collected manually is available on demand.

GETTING STARTED

See Every Asset. Know What State It's In. Fix What Matters Most.

The place to start is not the risk score. The place to start is the inventory — because the accuracy of the inventory is the accuracy of everything that follows. ConsoleWorks can build that inventory from existing tools your environment already has, from active collection alone, or from any combination of both. The configuration required is a set of Tool Data Collector connections to your existing security and operational tools, and a set of SRA device profiles for the assets you want to bring under active collection. From there, the Mapping Rules engine normalizes the data into a unified asset record, Measurement Questions run against every asset in scope, and Risk Analysis surfaces the posture for every stakeholder who needs to act on it.

No agents. No disruption to operational technology. No separate compliance tools. No manual evidence assembly at audit time. One platform — from blind spot to verified posture.

KEY TAKEAWAYS

- Asset Intelligence is not a discovery problem — it is a normalization and authority problem. ConsoleWorks solves both through passive aggregation and active collection in a single unified record.
- Risk scores are only defensible when they trace all the way from the fleet level to the specific measurement on the specific asset that drove the number. ConsoleWorks maintains that chain unbroken.
- Compliance is not a point-in-time activity. The frameworks that govern critical infrastructure don't allow for it. ConsoleWorks generates evidence continuously — it is already there when regulators ask.
- Zero Trust is not a network architecture concept added on top of the platform — it is the operational model that governs every session, every credential, and every access event. Least Privilege is enforced by design, at the device layer.
- The fix path closes the loop. When a Measurement Question returns FAIL, ConsoleWorks provides a direct path to the device through SRA — and the score updates only when the measurement confirms the gap is resolved.
- IT and OT are one inventory. Engineering workstations, historians, HMIs, PLCs, RTUs, and protective relays live in the same asset record, measured against the same controls, visible to the same teams — without separate tools or manual reconciliation.

About TDi Technologies

TDi Technologies has been securing critical infrastructure for over 25 years. ConsoleWorks is deployed across energy, oil & gas, water, nuclear, defense, and federal healthcare — trusted by organizations where operational failure is not an option and compliance is continuous, not seasonal.

ConsoleWorks Platform

- 01 • Expose** — Asset Intelligence & Risk Analysis
- 02 • Eliminate** — Secure Remote Access, CCM, Credential Management, IEM
- 03 • Enforce** — Continuous Measurement & Compliance Reporting

Contact & Resources

consoleworks.com
 info@consoleworks.com
 +1.800.695.1258

[Learn More](#)
consoleworks.com/how-it-works/