ORDR

THE PROACTIVE CONNECTED-DEVICE AI SECURITY MODEL

A practical guide to using AI agents to achieve complete connected-device visibility, turn intelligence into protection, and coordinate containment and remediation in seconds.

ORDR THOUGHT LEADERSHIP GUIDE

Security that acts before risk does.

VISIBILITY DOES NOT REDUCE RISK. ENFORCEMENT DOES.

WHY THIS OPERATING MODEL MATTERS

Connected devices are essential to care delivery and business operations. Many cannot run endpoint agents, and some cannot be patched quickly or at all. When Security cannot see what is connected, understand what matters, and limit what each asset can reach, one compromised device can become a path to widespread disruption. The goal is a trusted path from device truth to safe protection.

THE THREE OUTCOMES THAT MATTER MOST

1. KNOW EVERY CONNECTED ASSET

Create continuous visibility across managed and unmanaged IT, IoT, OT, and IoMT without requiring an agent on the device.

2. REDUCE EXPOSURE BEFORE IT BECOMES AN INCIDENT

Prioritize meaningful risk using identity, behavior, vulnerabilities, communication paths, business context, and compensating controls.

3. CONTAIN AND COORDINATE RESPONSE IN SECONDS

M-Trends 2026 reports that the window to intervene has collapsed from hours to seconds.[1] Stop compromised devices from moving laterally or turning a contained event into enterprise disruption while preserving human control.

WHERE AI FITS

AI agents accelerate the path from identification to protection so every asset is understood, exposures are prioritized, and containment can be implemented in seconds. Humans approve only the decisions and actions they have predefined as critical.

HOW ORDR ENABLES IT

Deep device intelligence, governed AI orchestration, and segmentation work as one system. Every asset moves from identification and profiling through prioritized risk and approved protection in seconds, with human approval only at predefined critical steps.

THE FOUR STAGES OF PROACTIVE DEVICE SECURITY

A SELF-ASSESSMENT

The operating model advances when each stage shortens the distance between an unknown asset, a meaningful risk, and a safely enforced outcome.

1 DISCOVER

Continuously identify every managed and unmanaged connected asset, including devices existing tools miss.

2 UNDERSTAND

Build behavioral, vulnerability, ownership, dependency, location, and business context for each device.

3 PROTECT

Model, validate, and enforce least privilege policies through existing network and security infrastructure.

4 ORCHESTRATE

Use AI to coordinate investigation, tickets, containment, approvals, validation, and reporting across systems.

A mature program does not stop at inventory. It continuously observes change, decides what matters, applies the safest available control, and proves that protection worked.

HOW ORDR ENABLES IT

One continuously updated device truth powers every stage. Governed AI agents coordinate discovery, investigation, policy, approval, segmentation, and validation so protection keeps pace with change.

FIVE QUESTIONS EVERY SECURITY LEADER SHOULD ANSWER BEFORE NOON

DECISION READINESS

If answering these requires multiple consoles, a custom query, or a week of email, the problem is not the team's expertise. It is the way device intelligence and action are connected.

1 What is connected that our controls miss?

Identify unmanaged, agentless, newly connected, and unidentified assets across IT, IoT, OT, and IoMT.

2 Which devices create the most meaningful exposure?

Prioritize by exploitability, behavior, communication, criticality, compensating controls, and business impact.

3 What can communicate with our critical systems?

Reveal observed flows, dependencies, unnecessary access, and attack paths that cross trust boundaries.

4 Which protections can we enforce safely?

Model least privilege policy, validate operational impact, and choose the right enforcement point.

5 What changed, and what requires action?

Surface new assets, abnormal behavior, policy drift, exposures, incidents, and unresolved exceptions.

THE TEST: Can the team answer, act, and prove the outcome with current device-level evidence?

HOW ORDR ENABLES IT

Ask ORDR IQ: “Which unmanaged devices can reach critical systems?” Its agents use current identity, behavior, vulnerability, flow, and policy data to explain the answer and prepare a report, ticket, or approved response.

VISIBILITY TO PROTECTION IS A CONTINUOUS LOOP

HOW THE MODEL WORKS

A point in time inventory cannot protect a changing environment. The loop must continuously absorb new devices, behavior, exposures, and business context, then translate that intelligence into the safest available control.

DISCOVER

Find every connected asset and observe new or changed devices.

CLASSIFY

Identify device type, owner, role, location, operating profile, and criticality.

PRIORITIZE

Focus on exposure that is exploitable, reachable, abnormal, and consequential.

MODEL

Create least privilege policy from observed behavior and business need.

VALIDATE

Simulate impact, review exceptions, and obtain the appropriate approval.

ENFORCE

Apply controls through existing infrastructure and continuously verify the result.

A recommendation is not protection. Durable protection carries device identity and expected behavior through policy, approval, enforcement, validation, and rollback instead of relying on a fragile IP address or one-time spreadsheet.

HOW ORDR ENABLES IT

Continuous device intelligence learns how each asset should communicate. Governed AI agents identify unnecessary paths, model least privilege policy, coordinate approvals, enforce segmentation through existing infrastructure, and verify that protection worked.

AN AI AGENT IS NOT A CHATBOT WITH ACCESS TO A DASHBOARD

THE RESPONSIBLE SECURITY MODEL

A generic assistant can summarize an alert. A trusted security agent must reason over current device evidence, protect customer data, expose its logic, preserve role-based access and audit history, and keep consequential actions human-approved. Machine speed and accountability must coexist.

FIVE REQUIREMENTS FOR AI SECURITY CAN TRUST

GROUNDED	Answers use current identity, behavior, exposure, communication, policy, and relationship data, not generic internet knowledge.
ROLE-AWARE	Security, network, IT, and operations receive the context and permitted capabilities appropriate to their work.
EXPLAINABLE	Users can inspect the devices, evidence, policy logic, time range, and confidence behind a recommendation.
BOUNDED	Each agent has a defined purpose, approved systems, permitted actions, and clear escalation rules.
HUMAN-LED	Consequential changes preserve review, approval, rollback, and audit history.

HOW ORDR ENABLES IT	Role-aware agents ground every recommendation in current device evidence. Ephemeral prompts, zero data retention, no customer-data model training, and audit logging protect enterprise data. Segmentation executes only bounded, approved actions.

THE SECURITY AGENT TEAM: FIVE TASKS WORTH AUTOMATING

WHAT AI SHOULD ACTUALLY DO

The highest value agents remove repetitive investigation, coordination, and reporting work so responders can focus on judgment, consequence, and recovery. Start with the task, evidence, permitted output, and human owner.

ASSET TRUTH AGENT

Reconciles discovered devices with existing systems; flags unknown, unmanaged, duplicated, moved, or stale assets; assigns the right owner.

EXPOSURE PRIORITIZATION AGENT

Combines vulnerability, behavior, reachability, criticality, and controls to rank the exposures most likely to matter.

THREAT INVESTIGATION AGENT

Assassembles a device timeline, peers, flows, anomalies, related alerts, and evidence; recommends containment and next steps.

SEGMENTATION POLICY AGENT

Builds least privilege policy from observed behavior, simulates impact, identifies exceptions, and prepares approval ready changes.

INCIDENT COORDINATION AGENT

Creates tickets, assigns owners, tracks containment and remediation, validates recovery, and produces leadership updates.

HOW ORDR ENABLES IT

Governed AI agents use current device intelligence to investigate risk, model segmentation, coordinate owners and approvals, initiate protection, and validate outcomes. Repetitive work moves at machine speed while people retain control of critical decisions.

TWELVE PROMPTS WORTH STEALING

A PRACTICAL AI STARTER KIT – PART ONE

Good prompts begin with a security decision, not a request for more data. Specify scope, time, evidence, desired output, control boundaries, and approval owner.

1. Show unmanaged or unidentified devices that connected in the past 24 hours; group by location, owner, and criticality.
2. Rank connected assets with exploitable vulnerabilities by reachability, abnormal behavior, business criticality, and compensating controls.
3. Map communication paths from guest, clinical, and IoT networks to critical systems; identify unnecessary access.
4. Investigate device [ID]: summarize identity, timeline, peers, flows, anomalies, exposures, alerts, and policy changes.
5. Find devices communicating with newly observed external destinations; explain why each pattern is unusual.
6. Prepare a containment plan for affected devices, including safest control point, operational impact, owner, and rollback.

HOW ORDR ENABLES IT

Natural-language requests activate specialized agents grounded in current device evidence. They investigate risk, reveal communication paths, recommend segmentation, and coordinate approved protection without requiring analysts to navigate multiple consoles.

TWELVE PROMPTS WORTH STEALING

A PRACTICAL AI STARTER KIT – PART TWO

Use these as starting points. The strongest workflows connect investigation to a bounded, reviewable action, and include validation after enforcement.

7 Draft leastprivilege policy for device class [X] from 30 days of observed behavior; flag exceptions for review.

8 Identify assets affected by advisory [ID]; show active use, location, reachability, owner, control status, and next action.

9 Create incident tickets for devices in scope; attach evidence, assign owners, set deadlines, and track approval status.

10 Generate an executive incident brief: scope, business impact, containment, remediation, blockers, and devicelevel evidence.

11 Validate that containment is effective and that required clinical or operational communications still function.

12 Explain what changed in our connected device risk posture since last week and which devices drove the change.

CREATE YOUR ORDR IQ ACCOUNT

An account is required to access the sandbox, but there is no fee.

TRY IT >

HOW ORDR ENABLES IT

Natural-language requests become investigations, reports, tickets, policy proposals, containment plans, executive briefs, or validation packages, all grounded in current device intelligence and connected to governed segmentation workflows.

A 90-DAY PATH FROM DEVICE TRUTH TO ENFORCED PROTECTION

DEPLOYMENT BLUEPRINT

Move fast without agents on connected devices, new enforcement hardware, network redesign, or rip and replace. Narrow the scope, prove the evidence, enforce and validate one safe control, then scale the governed pattern.

DAYS 0-30 | PROVE DEVICE TRUTH

Discover and classify one highvalue environment. Reconcile existing sources. Validate identity, ownership, behavior, criticality, and current controls with Security, Network, IT, and operations.

DAYS 31-60 | REDUCE ONE EXPOSURE PATH

Prioritize a meaningful risk. Model least privilege policy. Simulate impact, review exceptions, approve the control, enforce through existing infrastructure, and verify the outcome.

DAYS 61-90 | OPERATIONALIZE RESPONSE

Use ORDR IQ to coordinate investigation, tickets, owners, approvals, containment, validation, and reporting. Measure response time, exposure reduced, manual touches removed, and policy coverage.

MEASURE WHAT CHANGED

Unknown assets resolved | Critical exposures reduced | Attack paths closed | Mean time to contain | Manual touches removed | Policy coverage | Exceptions aging

HOW ORDR ENABLES IT

Start with one exposure path. Establish trusted device intelligence, use governed agents to coordinate investigation and approval, enforce one safe segmentation control, verify the result, then repeat the pattern at scale.

BUILD THE LEADERSHIP CASE AROUND EXPOSURE, RESPONSE, AND RESILIENCE

EXECUTIVE CONVERSATION

Leadership does not need another feature tour. It needs a defensible plan showing which material risk will change, how operations remain protected, who controls consequential actions, and how value will be measured.

THE ONE-PAGE BUSINESS CASE

CURRENT EXPOSURE	Define the assets, communication path, vulnerability, operational consequence, existing controls, and evidence gaps.
PROTECTIVE WORKFLOW	Describe discovery, prioritization, policy, approval, enforcement, validation, escalation, and rollback.
EXPECTED VALUE	Quantify exposure reduced, response time compressed, analyst effort returned, incidents contained, and disruption avoided.
CONTROL MODEL	Document roles, evidence, approvals, permitted actions, system boundaries, audit history, and exception handling.
90-DAY PROOF	Choose one environment, one exposure path, one owner, one baseline, one target, and one executive readout.
HOW ORDR ENABLES IT	Governed agents translate current device, exposure, policy, and response evidence into a leadership-ready security plan, budget, and validation package that connects investment to faster, safely enforced protection.

ONE SOURCE OF DEVICE TRUTH. EVERY TEAM IN SYNC.

WHY ORDR MAKES THIS MODEL POSSIBLE

Protection breaks down when teams work from different inventories and incident context. ORDR continuously discovers and identifies connected assets, understands their behavior and relationships, and makes the same trusted device intelligence usable across Security, Network, IT, and operations. ORDR IQ gives each role the questions, evidence, and workflows it needs without forcing everyone to become a dashboard expert.

FOR SECURITY

Prioritize meaningful exposure, investigate threats, coordinate containment and remediation, and prove the outcome with devicelevel evidence.

FOR NETWORK & IT

Understand identity, ownership, location, dependencies, and expected communication before enforcing or troubleshooting policy.

FOR OPERATIONS

Validate device role and critical workflows so protective action preserves care delivery and business continuity.

One device truth makes handoffs seamless across Security, Network, IT, and Operations. The model is proven at enterprise scale, with 500+ customers and 100M+ devices secured.[2]

CREATE YOUR ORDR IQ ACCOUNT

An account is required to access the sandbox, but there is no fee.

TRY IT ›

CONTINUE WITH ORDR

FROM INSIGHT TO ACTION

See how ORDR turns trusted device intelligence into safe, continuous protection, and how ORDR IQ helps teams coordinate the work at machine speed with human control.

ORDR CONNECTED-DEVICE SECURITY

Explore continuous visibility, meaningful risk prioritization, and safe enforcement across IT, IoT, OT, and IoMT.

ORDR IQ

Ask connected asset questions in natural language and use specialized agents to investigate, report, coordinate, and validate governed workflows.

CREATE YOUR ORDR IQ ACCOUNT

An account is required to access the sandbox, but there is no fee.

TRY IT >