threatdefence_

MODERN SECURITY OPERATIONS WITH MICROSOFT SENTINEL

A practical guide to implementing an agentic SOC on the Microsoft security stack.

TABLE OF CONTENTS

THE NEW PACE OF SECURITY OPERATIONS 3

STRONG FOUNDATION

Starting With the Microsoft Stack 4

What to Plan For 5

What an Agentic SOC Adds 6

TRANSFORMING YOUR MICROSOFT SENTINEL SOC

A Roadmap for Your Microsoft Sentinel SOC 7

Phase 1: Optimize Cost and Coverage 8

Plan Coverage Across the Environment 10

Phase 2: Bring AI Into the SOC 11

Set Access and Guardrails 12

What Can Be Automated on Microsoft Sentinel 13

Phase 3: Expand Autonomy Under Control 14

Define What Agents Can Do 15

Phase 4: Measure and Scale 16

Enablement for People and Process 17

HOW THREATDEFENCE HELPS

Managed SOC for Microsoft Sentinel 18

A Proven Blueprint for Microsoft Sentinel 19

Auto Triage for Every Alert 20

AI Agents for Microsoft Sentinel Incidents 21

Cut Microsoft Sentinel Costs 22

Coverage Beyond the Microsoft Stack 23

Getting Started 24

NEXT STEPS 25

The New Pace of Security Operations

Security teams have less time to investigate and contain threats. A compromised account can give an attacker access to email, applications, and sensitive data. AI adds to the pressure by helping attackers create convincing lures and adapt their methods.

Microsoft provides a strong foundation for detecting and investigating this activity. Entra ID records identity activity, Defender connects supported endpoint and email signals, and Sentinel brings in evidence from the wider environment. Making that evidence useful takes more than collection. Someone must investigate the findings, understand their significance, and respond at any hour.

AI agents help close that gap. They can begin investigating when an incident is raised, gather evidence across connected sources, and assess what happened. Analysts receive the findings and recommended actions, giving them more time to resolve uncertainty and coordinate response. Permissions, operating policies, and approval rules govern what agents can do.

For organizations already using Microsoft security tools, the starting point is familiar: existing data, detections, and response controls. The next step is to connect those capabilities into an operating model where agents and analysts work together, then expand automation as results support it.

> **Your existing Microsoft investment can provide the foundation for an agentic SOC.**

5.3 hours

Average time before exposed cloud workloads were attacked[1]

52.2%

Valid account intrusions that involved follow-on credential theft[1]

80%

Reduction in incident response effort with Microsoft Defender, including Sentinel SIEM[2]

[1] Microsoft Digital Defense Report 2026

[2] Forrester, The Total Economic Impact of Microsoft Defender, June 2025

STRONG FOUNDATION

Starting With the Microsoft Stack

For many organizations, the foundations of a SOC are already in place. Depending on licensing and deployment, Microsoft provides endpoint protection through Defender for Endpoint, identity security through Entra ID, and email protection through Defender for Office 365. The opportunity is to turn these capabilities into a coordinated operation that investigates incidents and responds around the clock.

Start With the Data You Already Have

Defender XDR correlates alerts across supported Microsoft products. Sentinel extends that view with additional Microsoft and third-party evidence, supporting investigation through the Defender portal.

Built-in connectors simplify collection, and selected sources, including supported Office 365 audit logs and Defender alerts, have no ingestion charge. Eligible Microsoft 365 E5 and E7 arrangements also provide up to 5 MB per user per day for specified sources. Review these benefits against the data collected and the actual bill.

Use Ready Content and Built-In Response

Content hub solutions provide a starting library of analytics rules, workbooks, hunting queries, and playbooks. Select the content relevant to your environment, check its dependencies, and tune it against real activity.

Response can use familiar Microsoft controls: isolating a device through Defender for Endpoint, revoking sessions through Entra ID, or removing malicious messages from mailboxes. These actions still need the appropriate permissions and agreed response authority.

Grow Into a Full Security Operations Program

Sentinel can bring third-party firewall, cloud, identity, and SaaS evidence into the investigation. Effective use of that evidence requires working connectors, consistent identifiers, suitable detections, and maintained investigation procedures.

Microsoft also provides a path into AI-assisted operations. Eligible Microsoft 365 E5 and E7 customers receive Security Copilot capacity through a phased rollout. Native agents and Sentinel MCP tools support defined workflows, subject to their access requirements, available data, and capacity limits.

> **Microsoft provides an integrated foundation. The security program connects its data, detections, people, and response into a working service.**

What to Plan For

A strong Microsoft foundation still needs operational ownership and a clear view of coverage. Address both when planning the service.

Data Costs

High-volume firewall, DNS, proxy, and flow data can drive significant consumption. Assess its investigation value, then apply suitable filtering, aggregation, storage, and retention choices. Include query, automation, and AI costs in the operating budget.

Investigation Capacity

Incidents need attention outside business hours, while specialist Microsoft knowledge may be concentrated in a few people. Plan for continuous investigation, escalation, engineering support, and response.

Ownership

Assign responsibility for connectors, rule health, custom content, data quality, and retention. An MDR service does not necessarily include full Sentinel management; confirm the division of responsibility.

Limits of Built-In AI

Native agents support specific scenarios. Their coverage, permissions, capacity consumption, and preview status need to be understood before they become part of the operating model.

Network Visibility

Endpoint and identity evidence may provide limited visibility into unmanaged devices, internal communications, IoT, and OT. Dedicated NDR sensors can add the network observations needed to investigate those areas.

External Exposure

Discovering internet-facing assets and assessing vulnerabilities require dedicated capabilities. Review any existing Microsoft exposure and vulnerability tools, then add discovery or scanning products where the required coverage is missing.

Brand and Dark Web

Impersonation sites, lookalike domains, and exposed credentials require external monitoring, supported by procedures for validation, escalation, and remediation.

What an Agentic SOC Adds

An agentic SOC makes investigation capacity available across more of the evidence the organization already collects. Agents gather context, follow investigation steps, and prepare conclusions and recommended actions. Analysts direct complex cases, resolve uncertainty, and coordinate response.

The change is most useful in three areas:

Skills



Apply Microsoft investigation expertise on every shift

Maintained investigation skills capture which evidence to examine, how to test possible explanations, and when to escalate. Agents apply those methods consistently, while experienced analysts review results and improve the guidance.

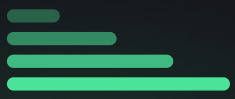
Speed



Shorten the path from incident to action

Agents can begin investigating supported Sentinel and Defender incidents as they arrive, at any hour. They assemble evidence and prepare response actions, such as session revocation or device isolation, for approval or execution within explicitly authorized limits.

Scale



Investigate more of the available evidence

Agents can handle concurrent investigations and examine cases that might otherwise wait. They connect identity, email, endpoint, and integrated third-party evidence, allowing analysts to focus on the findings that need their judgment.

> **An agentic SOC puts your Microsoft investment to work at greater speed and scale, with people setting its authority and overseeing the outcomes.**

TRANSFORMING YOUR MICROSOFT SENTINEL SOC

A Roadmap for Your Microsoft Sentinel SOC

Transforming a Microsoft Sentinel SOC works best in stages. Start with data costs and coverage, introduce AI into daily investigations, then expand autonomy as results support it. Measure progress throughout, whether you build in-house, use Microsoft's agents, or work with a partner.

1

Optimize Cost and Coverage

Align Sentinel spending with security value, assess visibility across the environment, and add the sources and products needed to support your security priorities.

2

Bring AI Into the SOC

Prepare the data, context, access, and investigation guidance agents need. Begin with selected incident types, with analysts reviewing conclusions and proposed actions.

3

Expand Autonomy Under Control

Move validated workflows from recommendations to approval-based execution, then allow specific actions within agreed limits.

4

Measure and Scale

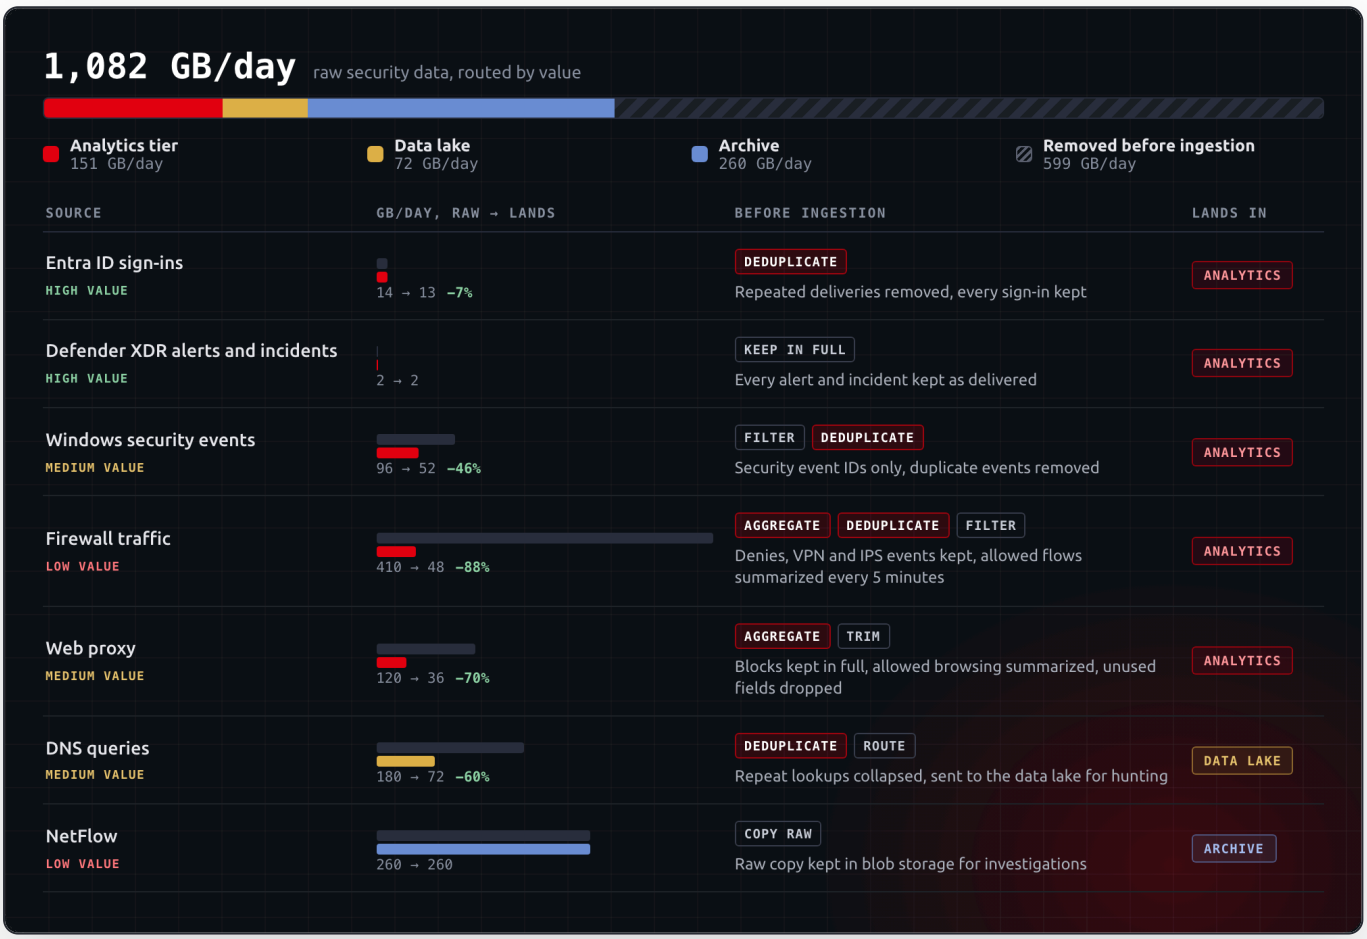
Expand to more incident types and sources when results support it. Use the capacity gained for hunting, detection engineering, and complex investigations.

Phase 1

OPTIMIZE COST AND COVERAGE

1

Start by understanding what Sentinel costs and what the SOC can investigate. Review the sources driving consumption alongside the assets, threats, and response requirements they support. This establishes where data can be reduced, where more evidence is needed, and which additional products belong in the design.



Choose a processing method that supports the intended reduction. Standard ingestion transformations and upstream pipelines have different capabilities, particularly for aggregation. Moving a table to the data lake also changes its availability for real-time analytics.

Reduce Data Costs by Design

Apply cost controls source by source, preserving the evidence needed for detection, investigation, and retention requirements.

Filtering	Remove unnecessary events and fields after reviewing their dependencies. Use supported collection transformations or an upstream processing pipeline.
Deduplication	Eliminate confirmed repeated deliveries before ingestion. Preserve distinct events that happen to contain similar values.
Aggregation	Summarize suitable high-volume telemetry while retaining the identifiers, timing, counts, and traffic volumes investigations need. Test the summaries against representative cases.
Data Tiers	Keep evidence required for real-time detections in the analytics tier. Use the data lake for suitable historical workloads, with a tested retrieval path for analysts and agents.
Commitment Planning	Size commitment tiers around expected consumption after the changes. Check eligible no-charge sources and Microsoft 365 ingestion benefits against actual usage.

> **Validate savings against detection and investigation needs, then monitor coverage and consumption after each change.**

Plan Coverage Across the Environment

Assess coverage against the organization's asset inventory and security priorities. Confirm that licensed Microsoft capabilities are deployed and producing useful evidence, then identify where additional sensors, discovery tools, scanners, or monitoring services are needed.

Area	Starting point	What to verify	Planned improvement
Identity	Entra ID and Defender for Identity	Relevant sign-ins, directory activity, and on-premises identity systems are monitored	Extend collection and sensor coverage across the agreed identity estate
Endpoint	Defender for Endpoint	In-scope workstations and servers are onboarded and reporting	Reconcile deployment with inventory and monitor sensor health
Email	Defender for Office 365	Protection policies and investigation capabilities cover the intended users	Review policy scope, exclusions, and response permissions
Cloud and SaaS	Defender for Cloud, Defender for Cloud Apps, and service logs	Required accounts, subscriptions, applications, and event categories are represented	Add supported integrations and validate their coverage
Network	Firewall, DNS, and proxy logs	Internal communications and unmanaged devices are sufficiently visible	Add NDR sensors at appropriate network collection points
External attack surface	Existing discovery tools, including Defender EASM where deployed	Internet-facing assets are discovered and changes are tracked	Extend attack surface management to the required domains and services
Vulnerabilities	Existing Microsoft or third-party assessment tools	Relevant assets receive suitable vulnerability assessments	Add scanning coverage and assign remediation ownership
Brand and dark web	Dedicated external monitoring	Impersonation, lookalike domains, and relevant credential exposure are monitored	Introduce brand protection and dark web monitoring with response procedures

Sentinel collects evidence from connected sources; asset discovery and vulnerability assessment require the underlying capabilities to be deployed. Microsoft offers products in these areas, and additional tools can extend coverage where needed.

Bring relevant findings into the same case-management and response process. Evidence may be ingested into Sentinel or accessed through an approved integration, depending on its purpose, cost, and investigation requirements.

Finally, define how source health will be monitored. Use expected delivery intervals, reporting assets, event categories, and volume patterns to identify collection failures. Assign someone to investigate those failures so that missing evidence receives attention.

Phase 2

BRING AI INTO THE SOC

2

With the foundations in place, introduce AI into daily investigations. Give agents access to the evidence and context they need, define their operating limits, and evaluate their findings alongside the team. This phase establishes investigation quality before automatic closure or containment is introduced.

Make the Environment Readable for Agents

Map entities so agents connect records to the right users, devices, applications, and resources. Keep business context in maintained watchlists, each with an owner and a review date: privileged and emergency-access accounts, service identities, approved tools, VPN ranges, critical assets, and change windows.

Test that agents can retrieve longer history from the data lake through their actual integration, and document investigation methods as reusable skills that analysts can review and improve.

Introduce AI Triage

Begin with selected alert classes from Microsoft Sentinel and Defender XDR, then expand toward every alert within the agreed scope. Agents gather evidence, assess possible explanations, and propose a verdict, priority, and next steps.

Start in recommend-only mode. Analysts continue to make closure and response decisions, compare their findings with the agent's assessment, and record material differences. Review both agreements and disagreements so that apparent consistency does not hide a shared mistake.

Start With Quick Wins

Choose workflows that consume meaningful analyst effort and have clear evidence requirements.

Alert triage

Investigate frequent alert classes and distinguish expected activity from findings requiring further attention.

Identity investigation

Connect sign-ins, authentication changes, application consent, and mailbox activity when assessing suspected account compromise.

Phishing review

Examine messages, destinations, delivery, and recipient activity to assess the threat and its reach.

Case documentation

Prepare evidence-based summaries, timelines, and recommended next steps for analyst review.

> **Reliable AI investigations depend on accessible evidence, maintained context, and clear investigation methods.**

Set Access and Guardrails

Separate investigation, incident handling, deployment, and response permissions. Give each service component the access its task requires, and test the complete workflow with the identity that will run it.

Read for Investigation
Grant query access at the intended workspace scope. Defender hunting and data-lake access may require separate permissions, depending on the integration.

Incident Handling
Microsoft Sentinel Responder supports incident management without granting general rights to change analytics rules or connectors.

Response Actions
Grant Defender and Entra ID response permissions separately, limited to approved operations and targets.

Service Identities
Use dedicated identities with clear ownership. Separate investigation, deployment, and containment responsibilities so their permissions and activity remain traceable.

Untrusted Input
Treat email bodies, filenames, web pages, and alert descriptions as evidence. Enforce tool and action restrictions independently of any instructions found in that content.

Administrative Control
Keep subscription ownership and access administration with your administrators. Agents should not be able to expand their own permissions.

For each investigation, retain the evidence retrieved, searches performed, failures, conclusions, and proposed actions. Apply limits to query scope, data volume, retries, and execution cost. Any later response workflow should also record its target, approval, execution result, and verification.

> **Enforce permissions and approval rules through the platform and connected tools, outside the AI model.**

What Can Be Automated on Microsoft Sentinel

Choose autonomy by task. Evidence gathering, drafting a case summary, closing an incident, and isolating a device have different consequences and need different controls.

The table below provides a starting point. Approve any expansion of authority for a specific workflow, supported by testing and ongoing review.

SOC function	What agents do	What people do	Starting mode
Alert triage and prioritization	Investigate alerts and propose verdicts and priorities	Validate findings, priorities, and closure decisions	Analyst review before closure
Investigation and correlation	Gather evidence, build timelines, and test explanations	Resolve uncertainty and direct further investigation	Read-only investigation
Reporting and case documentation	Draft summaries, timelines, and case updates	Validate conclusions and approve external communications	Internal drafts; approval before external release
Detection engineering and tuning	Draft detections and propose rule changes	Test effectiveness and approve production releases	Evaluation before deployment
Response and containment	Recommend actions and identify likely impact	Approve execution and define permitted automation	Approval before execution
Threat hunting	Run approved hunts and investigate anomalies	Set objectives, assess findings, and direct follow-up	Read-only access within agreed scope
Incident declaration and notification	Assemble evidence and draft notifications	Declare incidents and approve external notifications	Human decision and approval
Governance and strategy	Summarize performance, coverage, and operational trends	Set priorities, policy, and autonomy limits	Human authority

Agent work

Agents gather evidence, correlate activity, and document findings, adapting their investigation as new information emerges. They record uncertainty and escalate when evidence is insufficient or an operating limit is reached.

Human decisions

People direct complex investigations, authorize production changes and response, and remain responsible for incident declarations and external communications. They review performance and decide which workflows are ready for greater autonomy.

Phase 3

EXPAND AUTONOMY UNDER CONTROL

3

Once agents have proved their value, give proven workflows more authority, one action type at a time. Microsoft's own controls make each step concrete: an agent can revoke Entra ID sessions, remove an inbox rule, or isolate a device through Defender, but only when the workflow has earned it.

Shape Agents Around Your SOC

Configure agents around your investigation methods, business priorities, and response procedures. Define how they handle critical assets, privileged accounts, approved activity, and incidents that need specialist review. Successful methods become shared skills that analysts and agents use the same way.

Move From Approval to Action

Begin with actions held for approval: the agent prepares the containment and an analyst approves it in the case. Move a specific action to act within limits only when its evidence requirements, exclusions, and recovery steps have been tested. Closing a known-benign incident class and isolating a device are separate decisions, each with its own threshold.

Coordinate With Microsoft's Agents

Check what each built-in Microsoft agent does once it is enabled. The Phishing Triage Agent, for example, resolves alerts it classifies as false positives on its own, and existing playbooks may already act on incidents. Decide which component owns each decision so that agents, playbooks, and people never act on the same incident without knowing it.

> **Autonomy is earned one action type at a time, with people accountable for the decisions that carry risk.**

Define What Agents Can Do

The platform should let you decide which work agents carry out independently and where people must approve. Operating modes make that division explicit.

Mode	How it works
Recommend only	Agents investigate and recommend next steps. Analysts decide the outcome and initiate response.
Approval required	Agents prepare actions with supporting evidence. An authorized person approves before execution.
Autonomous within limits	Agents execute specifically authorized actions when defined conditions are met. Exceptions go to an analyst.

Move proven workflows into approved execution first. Before authorizing autonomous actions, test their target checks, expected impact, failure handling, and recovery procedures. A correct verdict alone does not establish that an action is safe.

Enforce Boundaries Through the Platform

Permissions and execution limits must operate outside the model's control:



Scoped access

Isolated runtimes, approved tools, protected credentials, and tenant boundaries limit what agents can reach.



Controlled actions

Target exclusions, approval rules, and scope and rate limits constrain execution.



Untrusted evidence

Content from alerts, emails, and logs cannot grant permissions or authorize actions.



Visible decisions

Queries, results, actions, and approvals remain available in the case for review.

Authorized analysts need controls to stop actions, reduce permissions, or suspend individual workflows. Test these controls, including how work returns to analysts when a source fails or automation is suspended.

Record changes to autonomy and assign a named owner. If an unintended action occurs, preserve the evidence, investigate the cause, and validate corrections before restoring the workflow.

Phase 4

MEASURE AND SCALE

4

Turn early success into sustained improvement. Compare results with the baseline established before deployment, and review them by incident type, source, and severity. An overall figure can hide a struggling workflow or a recurring data problem.

Measure	What it shows
Investigation coverage	Incidents eligible for agent investigation, completed investigations, and incomplete or unsupported cases
Verdict agreement	How often analysts agree with agent conclusions within the reviewed sample
Overrides	Changes to verdicts, priorities, or proposed actions, with the reasons recorded
Material errors	Incorrect closures, missed malicious activity, unsupported conclusions, and inappropriate response recommendations
Investigation and response times	Time to investigation, approved containment, and closure, measured separately
Data cost by source	Consumption and charges by source and tier, assessed alongside detection and investigation value

Verdict agreement helps identify consistency, but it does not establish accuracy. Review a representative sample of agreements as well as disagreements, include cases with known outcomes, and examine reopened incidents. Record sample sizes so that a small number of successful cases does not become the basis for broad autonomy.

Scale What Works

Use the results to decide which workflows can expand and which need improvement. A reliable identity workflow may support additional incident types, while inconsistent endpoint findings may point to missing telemetry or unclear investigation guidance.

Extend coverage where the security program requires it. Additional NDR sensors, attack surface management, vulnerability scanning, or brand protection can provide evidence beyond the current deployment. Connect those findings to maintained investigation methods and response procedures.

Use the capacity gained for regular threat hunts, detection improvements, and work on recurring causes of incidents. Look for sustained outcomes: smaller backlogs, more complete investigations, useful proactive findings, and consistent quality as the environment grows.

Review the full operating cost alongside those results, including data, AI consumption, supporting products, engineering, and analyst effort.

Enablement for People and Process

Enablement runs alongside every phase. Analysts need time and practice to move from collecting evidence toward reviewing agent findings, resolving uncertainty, and directing response. Senior staff maintain investigation methods, improve detections, and oversee agent performance.

Use real cases from your environment to develop these skills. Review how Defender, Entra ID, and Sentinel evidence supports a conclusion, where an investigation was incomplete, and what should change. Turn useful lessons into maintained skills and runbooks with clear owners, approval points, and escalation paths.

Assign responsibility for autonomy decisions and workflow changes. Hold regular reviews of overrides, closures, errors, and response outcomes. Improvements should pass through a defined evaluation and release process before reaching production.

Keep stakeholders informed about what agents handle, which actions require approval, and how to raise a concern or suspend a workflow. Service reviews should connect technical performance with coverage, workload, cost, and the priorities of the organization.

> **Lasting improvement comes from developing the people and processes alongside the agents.**

HOW THREATDEFENCE HELPS

Managed SOC for Microsoft Sentinel

Microsoft Sentinel and Defender XDR provide a strong foundation for security operations. ThreatDefence adds the AI SOC platform, operational expertise, and 24x7 team to run it as a managed service.

We maintain detections, connect relevant sources, investigate incidents, and respond within your approved authority. Our agents work alongside experienced analysts, threat hunters, and incident responders, bringing agentic investigations to the Microsoft environment you already use. Your Sentinel workspace stays in your subscription.



Build on Microsoft

Put Sentinel, Defender XDR, Entra ID, and relevant Purview capabilities to work through coordinated investigations and response.



AI on Every Incident

Agents investigate incidents in scope, resolve supported benign findings under approved policies, and escalate suspicious or uncertain activity with evidence.



24x7 SOC and Response

Our team investigates, hunts, and responds around the clock, following your procedures and approval requirements.



Lower Sentinel Data Costs

Intelligent filtering, aggregation, and storage tiering reduce unnecessary consumption while preserving the evidence your SOC needs.

A Proven Blueprint for Microsoft Sentinel

We bring detections, playbooks, automation, data design, and operating procedures developed through our 24x7 SOC. Your service starts with a tested foundation that we adapt to your licenses, sources, priorities, and risks.

Detection Library
Detections for Microsoft and third-party sources, mapped to relevant MITRE ATT&CK techniques and refined through real investigations.

Response Playbooks
Tested workflows for session revocation, device isolation, inbox rule removal, and sender blocking, configured around your approval requirements.

Automation Rules
Incident routing, enrichment, and appropriate suppression to reduce repetitive handling and bring useful context into each case.

Data Design
Source-by-source collection, processing, storage, and retention choices that support detection and investigation at a sustainable cost.

SOC Processes
Runbooks, escalation paths, response approvals, and reporting based on the procedures our analysts use every day.

We review planned changes with your team before deployment and validate the content against your environment. The deployed configuration, exceptions, and dependencies become part of the service record.

Auto Triage for Every Alert

AI alert triage investigates every alert in scope, gathering evidence from connected tools and producing an assessment with a verdict, confidence score, severity, and supporting investigation record.

Begin in recommend-only mode while ThreatDefence analysts review the results. As workflows are validated, supported benign activity can close under approved policies. Suspicious findings, missing context, and conflicting evidence reach analysts for further investigation.



Reduce Repetitive Triage

Agents assemble the evidence needed to assess an alert, giving analysts more time for complex cases, threat hunting, and response.

Follow the Evidence

Agents query connected tools, pursue related activity, and build timelines as the investigation develops.

Prepare for Response

Cases include recommended containment steps. Approved actions can execute within defined limits, while other actions wait for authorization.

Improve Investigation Methods

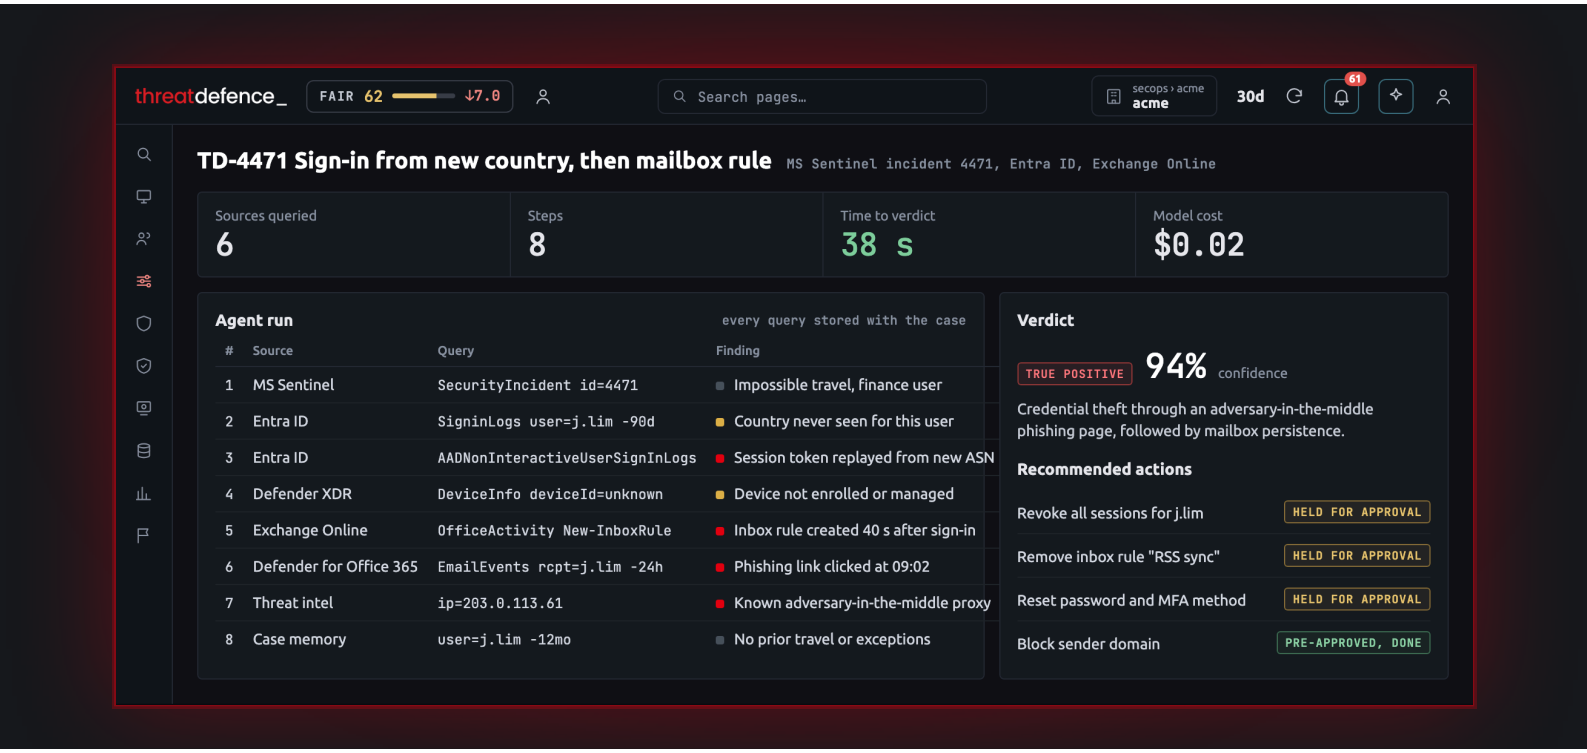
Analyst corrections and response outcomes inform reviewed updates to skills, business context, and detection content.

Confidence scores help route cases for review. Supporting evidence and your operating policies determine whether a finding can close or an action can proceed. Missing or conflicting information remains visible in the case.

AI Agents for Microsoft Sentinel Incidents

An incident can involve several alerts, accounts, devices, and applications. ThreatDefence agents investigate those relationships across Sentinel, Defender XDR, and other connected tools.

They examine sign-ins, device activity, mailbox changes, and related events to establish what happened and which systems may be affected. The investigation develops as evidence emerges, with each step and its findings recorded in the case.



Your team and our SOC work from the same investigation record: timeline, affected systems, evidence, conclusions, and action history. Analysts can validate the work, add business context, direct further investigation, and approve response without rebuilding the case.

Cut Microsoft Sentinel Costs

We assess ingestion, retention, data tiers, and supporting services against the detections and investigations they enable. The aim is to reduce unnecessary processing and storage while keeping useful evidence available.

Through intelligent filtering, aggregation, and appropriate tier placement, Microsoft Sentinel data costs typically fall 50 to 60%. The achievable saving depends on your starting configuration, source mix, retention, and coverage requirements.



Recommendations account for processing, queries, storage, and retrieval charges. We validate detection dependencies and representative investigations before changing collection or tier placement, then monitor coverage and consumption after deployment.

The free cost and coverage review identifies the opportunities in your environment and the priorities for implementation.

Coverage Beyond the Microsoft Stack

Where your security program needs additional visibility, ThreatDefence can add capabilities through the same AI SOC platform. Their findings contribute to shared investigations and the 24x7 response service.

Network Detection and Response

Investigate suspicious communications and lateral movement, including activity involving unmanaged devices, IoT, and passive OT monitoring.

Brand Protection

Detect lookalike domains, impersonation, and phishing sites, with investigation and takedown coordination.

Dark Web Monitoring

Assess relevant credential and information exposure, connect findings to affected identities, and coordinate account protection.

Attack Surface Management

Discover internet-facing assets and monitor exposed services and changes to your external footprint.

Vulnerability Scanning

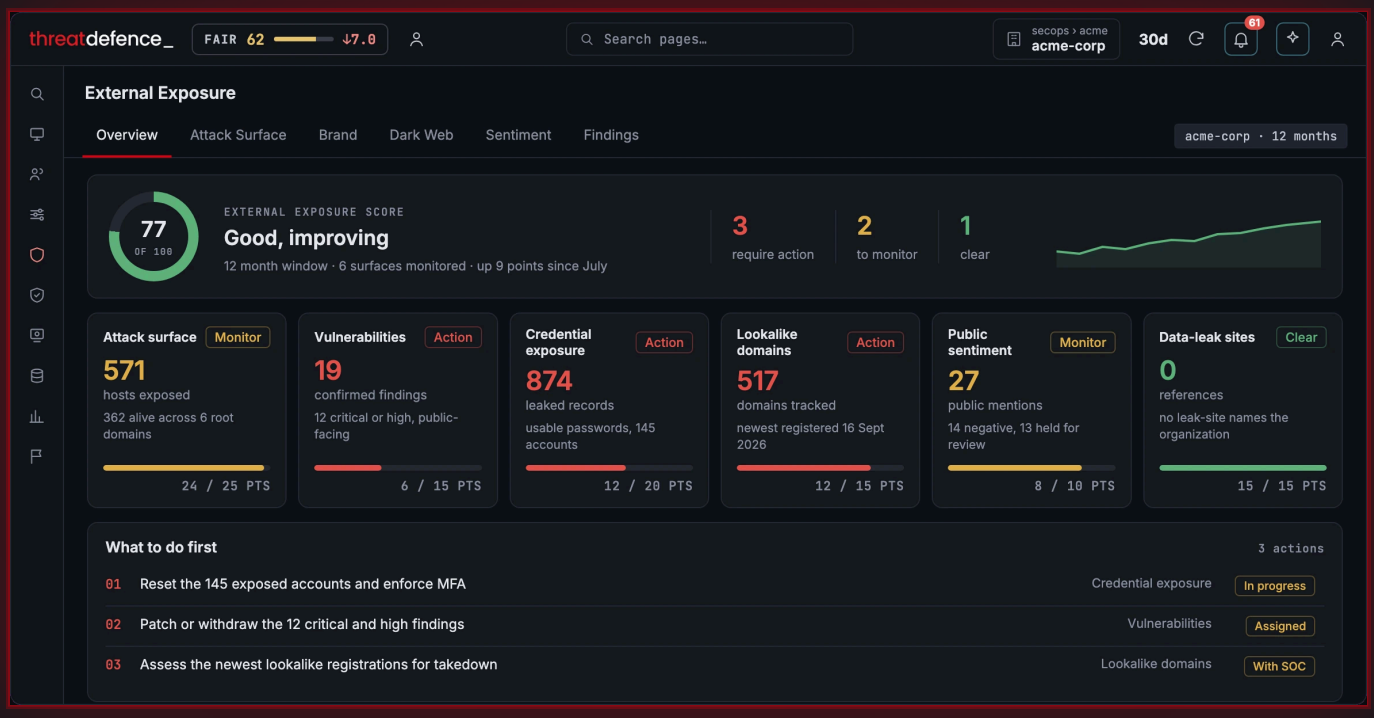
Assess agreed assets for weaknesses and prioritize remediation using exposure, threat intelligence, and business context.

AI Usage and Governance

Identify observed AI service use and review relevant activity against your policies.

Deception

Use decoys to surface unauthorized interaction and provide additional investigation leads.



Getting Started

Start with a free Microsoft Sentinel cost and coverage review. You receive a clear assessment of your current deployment, potential savings, and priorities for stronger security operations. If you proceed, that assessment becomes the starting point for implementation.

1

Connect

Your administrator grants access for the agreed review. We confirm its scope and permissions with your team.

2

Review

We assess recent incidents, ingestion, retention, detection health, source coverage, and available licenses. You receive recommended priorities and potential savings, at no cost and with no obligation.

3

Prepare

Agree the service scope, responsibilities, response procedures, and approvals. We adapt the blueprint, address agreed deployment needs, and validate investigations with your team.

4

Operate

Our AI SOC platform and 24x7 team run the service, supported by ongoing hunting, detection engineering, and regular reviews of coverage, performance, and cost.

threatdefence_

GET A FREE MICROSOFT SENTINEL REVIEW

See where your Microsoft Sentinel costs can come down, where detection gaps remain, and what AI agents and our 24x7 SOC would change in your environment. The review is free and comes with prioritized recommendations.

BOOK A FREE REVIEW

threatdefence.com/microsoft-sentinel-mdr
sales@threatdefence.com



threatdefence_

ABOUT THREATDEFENCE

ThreatDefence is an Australian cybersecurity company. Our AI SOC platform helps security teams investigate alerts, contain threats, and strengthen detections across their existing tools. AI agents handle repetitive investigation work, while analysts direct response and make the decisions that matter.

Our 24x7 SOC provides continuous monitoring, investigation, and hands-on response. We also deliver managed SIEM, threat hunting, incident response, and digital forensics, alongside brand protection and AI governance services.

Learn more: threatdefence.com

Request a demo: threatdefence.com/book-a-demo

Contact us: sales@threatdefence.com

© 2026 ThreatDefence. All rights reserved.