Code-to-Cloud and AI Security

We find the problems and we fix them.

We scan your code and your cloud accounts, then open pull requests with the fixes for you to review.

AI Security

The AI agents in your AWS accounts, the data each one could damage, and whether that damage could be undone.

- What each agent could read, change or delete
- AI attack paths and the change that closes them
- Optional gate for agent actions (Enterprise plan)
- AI inventory and shadow AI in code

Code Security

Checks on your repositories for problems that would reach production when the code is merged.

- Secrets, Dockerfiles and AI usage in connected repositories
- Infrastructure as code, dependencies and application code in the free scan
- Cloud findings traced to the code that created them

Cloud Security

Your AWS, Azure and GCP accounts scanned continuously.

- Security detections and attack paths
- Drift from your IaC code
- Code-to-cloud mapping

Compliance

Every finding mapped to the frameworks your auditors and customers ask about.

- Frameworks such as SOC 2, ISO 27001 and NIS2, plus ATT&CK coverage
- Posture by framework
- Report with evidence of merged fixes

Real-time threat detection

AWS · Enterprise plan

Changes and threats in your AWS accounts as they happen, from AWS CloudTrail events and GuardDuty findings, with a live change log and an alert register for audits.

Why we built it

Security and engineering teams spend weeks on security issues that someone else introduced: reading the finding, working out which resource it means, and finding the team that owns it. Netlumi finds the issue and writes the change, so your team reviews a pull request instead of investigating a ticket.

Fixes as pull requests. When Netlumi can fix a finding, the fix arrives as a pull request.

AI Security

The AI agents in your cloud and what they can reach.

Netlumi lists every AI agent in your AWS accounts, the data each one can reach through its permissions, and which of its actions could not be undone. It recommends the single change that removes the most risk.

What Each Agent Could Do

Netlumi shows what each agent can read, change or delete.

Netlumi tells you which data an agent could change that you could not get back, and what to do first.

acme-support-assistant	What it could do
acme-customers DynamoDB table · delete	Cannot be restored
acme-invoices S3 bucket · change	Can be restored
acme-orders-db RDS database · change	Unknown

Precaution: **turn on point-in-time recovery** on acme-customers.

AI Attack Paths

Each path starts where untrusted text can reach an agent, such as a public endpoint, a document or a message. It runs through the agent and its tools to your data.

Netlumi shows the change that closes each path, starting with the one that closes the most. Where it can, the fix arrives as a pull request.

AI path to acme-customers — Cannot be restored

1. Untrusted input: the public support chat
2. Agent: acme-support-assistant
3. Runs as: acme-support-tools role
4. Crown jewel: acme-customers table

Cut here first: **remove the delete permission from acme-support-tools**.
Closes 3 paths.

Runtime Gate

Enterprise plan · opt-in

An optional gate can stop an agent action that would destroy data you cannot restore. Netlumi never reads the prompt or the conversation.

You choose which agents it covers, and you can run it in record-only mode first. Every allow and block is shown with its reason.

Gate decisions	Block · last 7 days
get_invoice acme-billing-agent	Allowed
delete_customer acme-support-assistant	Blocked
update_order acme-support-assistant	Allowed

Blocked: **the action would delete data that cannot be restored**.

AI Security · Runtime

Agent activity and decoys.

Netlumi follows what your agents do in your AWS accounts, and raises an alert when anyone uses a decoy.

Agent Activity

AWS · Growth and Enterprise plans

Netlumi follows what each agent's identities do in your AWS accounts. It does not read the agents' conversations.

The Activity page shows the week's alerts and gate decisions. An alert is raised only for high-severity events, such as a blocked action.

Activity	Last 7 days
Gate blocked delete_customer `acme-support-assistant`	Alert
Read items from acme-orders `acme-support-assistant`	Seen before
Wrote a report to acme-exports `acme-billing-agent`	Seen before

Decoys

Growth and Enterprise plans

Decoys tell you within minutes if an attacker finds them or an agent is being steered by something it read.

Each alert names the address the use came from.

Decoys	Last used
`ci-deploy-key` acme/platform	Never
`support-tools` acme-support-tools	Never
`kb-onboarding-doc` acme-help-kb	2 hours ago

kb-onboarding-doc used from 203.0.113.24. A critical alert was raised in Threats.

AI Security · Models and shadow AI

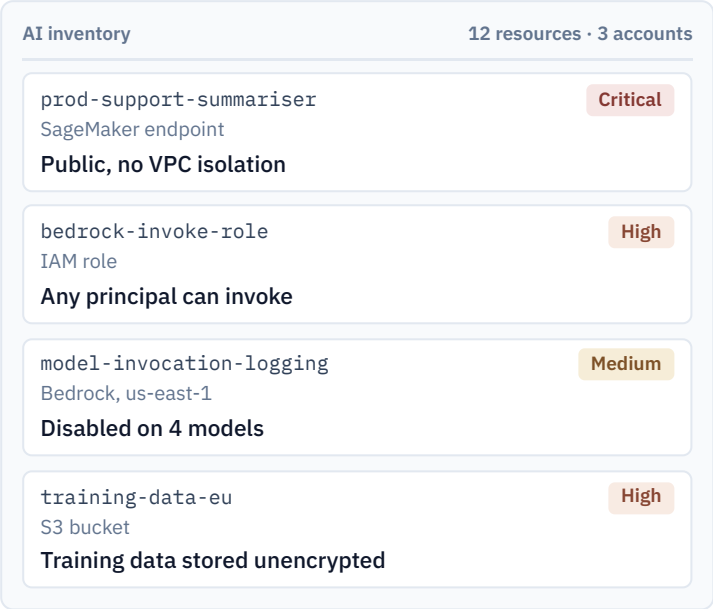
The AI models you run and the AI usage in your code.

Netlumi checks the models and endpoints already in your accounts, and finds new AI usage in your repositories before it reaches production.

Model and Endpoint Security

Bedrock and SageMaker are cloud resources and can be misconfigured like any other. Azure OpenAI and Google Vertex AI are detected in code.

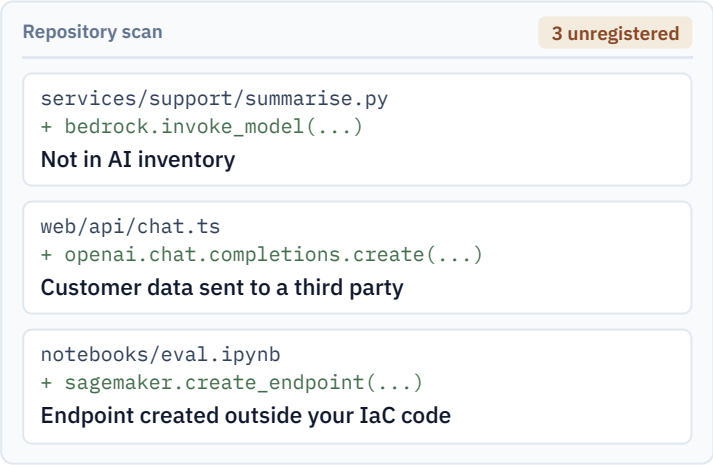
We check every model and endpoint you run for public exposure, roles that let anyone invoke them, invocation logging that is switched off, and unencrypted training data.



Shadow AI

Model calls and API keys are often added to services without a security review.

We read your repositories and find AI usage that is not in your AI inventory, with the file and line where it appears.



Code Security

Find problems in your repositories before they ship.

Connect GitHub and Netlumi checks your repositories on a schedule and as code changes, and opens a fix pull request when one is needed. The free scan of a public GitHub repository runs more checks today, with no sign-up.

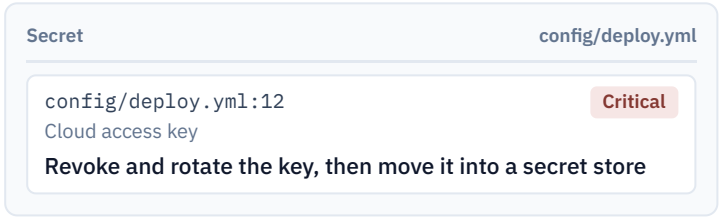
Infrastructure as Code

The free scan finds misconfigurations in your IaC code, Kubernetes manifests and Dockerfiles. In connected repositories, Netlumi checks Dockerfiles today and reads your IaC code to trace cloud findings to their file; infrastructure-as-code checks are coming.



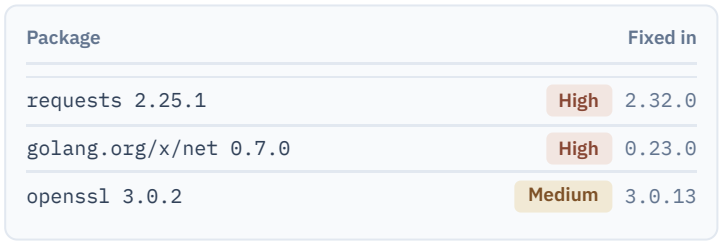
Secrets

AWS access keys, private keys and API keys for AI model providers committed to a connected repository. The free scan also finds many other common key and token formats. Each finding says what to do about it. Secrets in git history are coming.



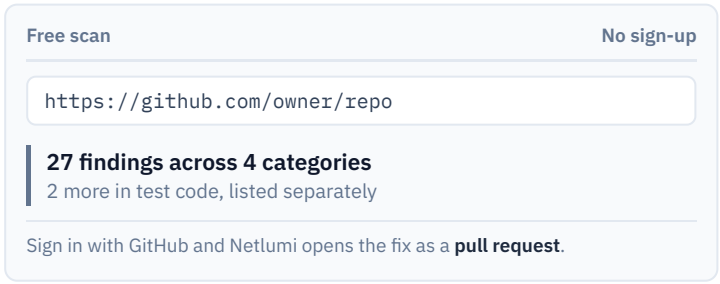
Dependency Vulnerabilities

The free scan finds known CVEs in lockfiles and manifests for JavaScript, TypeScript, Python, Go, Java, Rust, Ruby, PHP and .NET, and your dependency list is not sent to a third party. In connected repositories, Netlumi checks the base image of each Dockerfile for known high and critical vulnerabilities; dependency checks are coming.



Code Checks

The free scan finds security issues in application code for Python, Go, JavaScript and TypeScript. Findings in example, test and fixture code are listed separately and not counted. These checks are coming to connected repositories.



Cloud Security

What we look for in your cloud accounts.

One read-only connection per account, on AWS, Azure or GCP. Scans run on a schedule and when things change, so a new problem is found soon after it is introduced.

Security Detections

Continuous checks against good practice and CIS Benchmarks™ recommendations, at the same depth in every account you connect. Findings are mapped to compliance frameworks (see Compliance).

Latest scan	prod · all regions
RDS instance publicly accessible	Critical
Root account without MFA	Critical
Security group open to 0.0.0.0/0 on 22	High
CloudTrail log validation off	Medium

Attack Path Analysis

Netlumi maps the chain from an exposure, through identity escalation, to your data, so you can start with the paths an attacker could use rather than ordering work by severity label alone.

Path to prod-rds-01 — Critical

1. Public internet reaches the load balancer
2. Bucket policy on app-assets allows write
3. Lambda execution role is assumable
4. That role can read the production database

Drift Detection

Netlumi shows when live infrastructure stops matching its IaC code, what changed, who changed it, and whether the change went through your pipeline. The fix arrives as a pull request.

Resource	State
`prod-db-instance`	Drifted
`api-gateway-prod`	Console change
`s3-backup-bucket`	In sync

Changed by **j.doe@company.com** through the AWS console, outside the pipeline.

Code-to-Cloud Mapping

Netlumi traces every resource defined in your IaC code back to the file, line and commit that created it, so you know where a fix goes and which team owns it.

`terraform/aws/rds.tf` — Line 45

```
resource "aws_db_instance" "prod" {
  engine = "postgres"
  publicly_accessible = true
}
```

Committed by **John Doe** two days ago on `main`.

Real-time Threat Detection · AWS · Enterprise plan

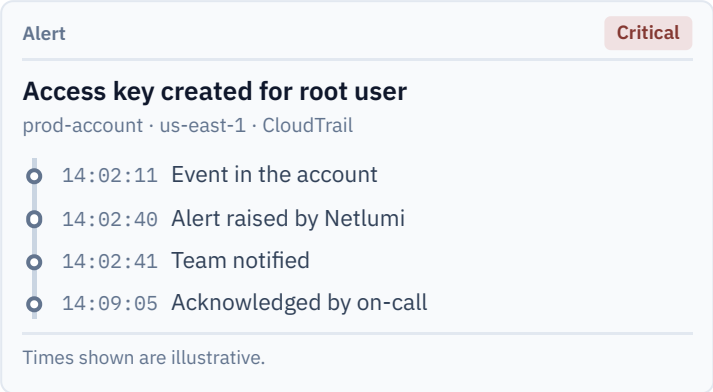
Alerts as changes happen in your AWS accounts.

Scheduled scans show the state of your accounts. Real-time detection shows changes as they happen. Your AWS accounts forward their CloudTrail events, and Netlumi raises an alert when one matches an alerting rule.

Live Changes and Threats

Netlumi's rules run on the cloud events your accounts forward, alongside GuardDuty, the AWS detector you already use. Results are shown in one place, with a live change log of who did what, where and when.

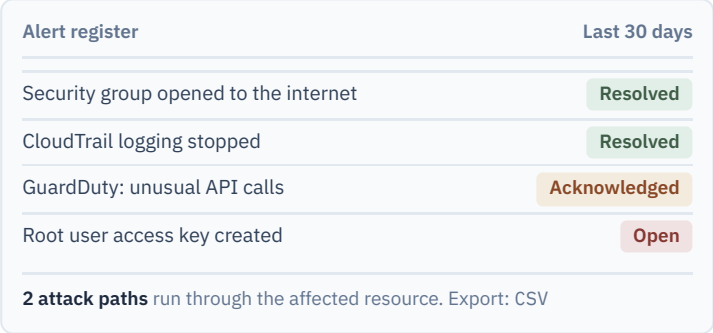
Only high-severity threats raise an alert. Everything else is shown in your risk view without notifying anyone.



Context and Audit Record

An alert on a resource shows the attack paths that run through it, so you can see at once whether a suspicious change opened a route to your data.

Every alert is resolved with an outcome and notes. The alert register records when each one happened, was detected, notified, acknowledged and resolved, and downloads as CSV for your auditors.



Real-time and scheduled scanning. Real-time detection covers AWS on the Enterprise plan. Everything else, in every plan and every cloud, is rescanned on a schedule and when things change.

Compliance · Paid plans

Your posture against each framework.

Every finding from code, cloud and AI is mapped to the frameworks below, so you can see your posture against each one. The fixes you merge are recorded as evidence.

Posture by Framework

See where you stand on SOC 2, ISO 27001, PCI DSS, HIPAA, NIS2 and the rest, which controls are failing, and which findings and resources are behind each one.

Framework	Controls passing
SOC 2	
ISO/IEC 27001:2022	
NIS2	
CIS AWS recommendations	

Security and Compliance Report

A report with a compliance section, plus PDF and CSV evidence of the fixes you merged: what was wrong, the pull request that fixed it, and when. Netlumi Lens answers questions about your framework mappings in plain English.

Evidence	PDF · CSV
S3 public access blocked	PR #284
RDS made private	PR #291
CloudTrail validation on	PR #302

Security and assurance

SOC 2 Trust Services Criteria	2017	ISO/IEC 27001:2022 Annex A	2022
PCI DSS	v4.0.1	HIPAA Security Rule	45 CFR 164 C
NIST SP 800-53	Rev. 5	NIST Cybersecurity Framework	2.0
NIST SP 800-171	Rev. 2	FedRAMP Moderate baseline	Rev. 4
CISA Cyber Essentials	2021	MITRE ATT&CK (cloud techniques)	Enterprise

European regulation and cloud standards

NIS2 Directive	2022/2555	GDPR	2016/679
BSI C5 (Germany)	2025	ANSSI SecNumCloud (France)	3.2
ENS (Spain)	RD 311/2022		

Cloud benchmarks

CIS AWS Foundations Benchmark™	v7.0.0	CIS Azure Foundations Benchmark™	v6.0.0
CIS GCP Foundation Benchmark™	v5.0.0	AWS Foundational Security Best Practices	1.0.0
Microsoft cloud security benchmark	v1		

This is Netlumi's own mapping of its checks to each framework. It is not an audit, an attestation or a statement of compliance. CIS Benchmarks are referenced as a mapping to their recommendations; Netlumi is not affiliated with the Center for Internet Security.

Remediation

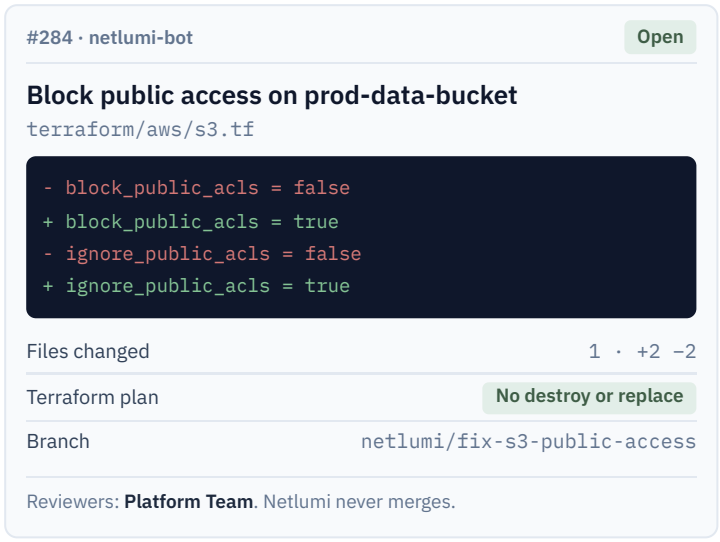
Fixes in your repository, and answers to your questions.

Fixes for code, cloud and AI findings arrive as a pull request in your repository, in the file that defines the resource. Netlumi Lens answers questions about your findings.

Fixes as Pull Requests

Netlumi writes the fix and opens a pull request in your repository. Your engineers review it like any other change.

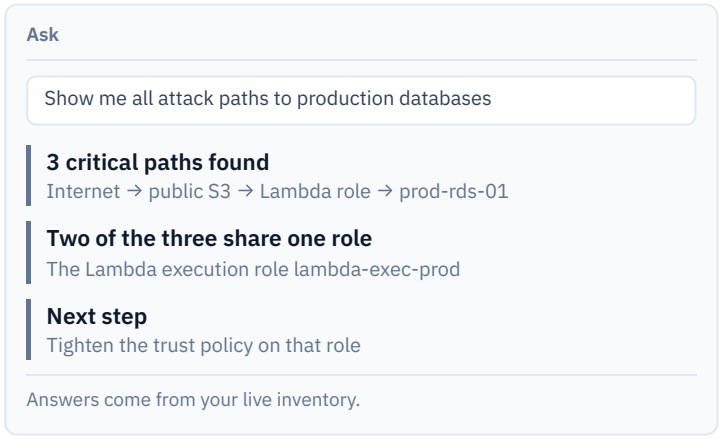
- Review mode by default: nothing is written to your repository until you approve the fix.
- Fix pull requests are checked before you see them, and never applied for you.
- Each pull request says what was checked and what was not.
- Fixes that would destroy or replace a resource are blocked.



Netlumi Lens

Ask questions in plain English about your findings, attack paths and framework mappings, such as who can read the customer bucket, what changed in the last day, or which paths reach production.

Answers cite the records they come from, so you can check them. You do not need a query language or a spreadsheet export.



Nothing changes without your approval. Netlumi has no write access to your cloud, never pushes to your default branch and never merges. Your team reviews the pull request and your pipeline applies it, as it would any other change.

Getting started

Connect your cloud accounts and your repositories.

Netlumi is agentless. You give us read-only access to the cloud accounts you want covered and connect your repositories. The only software you deploy is the optional AI agent gate.

01

Connect read-only

An AWS role assumed with an external ID for one-hour credentials, or workload identity federation for Azure and GCP. No stored access keys.

02

Connect your repositories

A read-only GitHub app for scanning. A separate fix app opens pull requests on the repositories you choose.

03

Netlumi scans

Cloud configuration, and secrets, Dockerfiles and AI usage in your code, on a schedule and when things change.

04

You review the pull request

Fixes arrive from a `netlumi/fix-*` branch. Your team reviews and merges.

Code	Connected repositories: secrets, Dockerfiles, AI usage and MCP configuration; cloud findings traced to your IaC code. The free scan adds infrastructure-as-code, dependency and code checks
Clouds	AWS, Azure and GCP
AI services	Cloud checks on Amazon Bedrock and SageMaker; Azure OpenAI and Google Vertex AI detected in code
AI agents	AWS. Optional gate for agent actions (Enterprise plan)
Compliance	We support compliance and show how you stand against frameworks such as SOC 2, ISO 27001, PCI DSS, HIPAA, NIS2 and GDPR, with MITRE ATT&CK coverage
Real-time detection	AWS CloudTrail events and GuardDuty findings (Enterprise plan)
Runtime protection	Growth and Enterprise plans: native detector coverage (GuardDuty, Defender for Cloud, Security Command Center), triage, evidence and decoy credentials
Scanning	On a schedule and on change, in every region of every account you connect
Integrations	GitHub
Your code	Working copies of connected repositories are kept on Netlumi's scanning hosts, encrypted at rest, while the repository is connected, and proposed fixes are kept for up to 30 days. The free scan does not keep the code: it is downloaded into a temporary, isolated environment that is destroyed when the scan ends, and its report (short excerpts, secrets masked) is kept for 90 days
Plans	Free, Starter, Growth and Enterprise, with unlimited users. Free covers one cloud account with daily scans. What each plan includes, and prices, at netlumi.com/pricing
Delivery	Cloud-managed SaaS. Netlumi Limited, Cork, Ireland

Netlumi

Code-to-Cloud and AI Security · October 2026

netlumi.com

Scan a public repository free at app.netlumi.com/free-scan, or book a walkthrough.