

Protecting your operations against AI-enabled threats


Including the threats from using AI tools

A guide for Chief Information Security Officers



Introduction

The rise of AI-powered threats has brought cybersecurity to an unprecedented level of prominence, in the media and in the boardroom.



A recent report by Gartner found that 93% of non-executive directors now see cybersecurity as a threat to shareholder value.

Regulators too have been spurred into action, issuing advisories and directives drawing attention to the fact that adversaries, including state-backed groups, are now using AI to increase the speed and intensity of their attacks. They are also urgently revising compliance frameworks, driving a transition from a check-box approach to a risk-based one.

At the core of this is a recognition that, with AI at their disposal, attackers are likely to breach networks somewhere – and that

once inside traditional flat networks, they are able to move very quickly to access key assets. So beyond the (still essential) efforts to stop those breaches, there's a need to adopt measures that can minimise the damage these AI-enabled threats can cause.

This guide examines what those responses should look like in practice. It sets out our view of the core steps that all organizations should consider to give them defence in depth against AI-powered attacks.

There are multiple ways to deliver these in practice. Some CISOs I've met are prioritising restructuring their operational teams to bring together previously separate skillsets, so they can respond faster and more effectively to the AI threat. Others are focusing on their processes, from proactive monitoring to breach response.

Both are valid and important. Our perspective is that the right tech can also help, giving you the insights to identify and close down the highest priority risks. So in boxes throughout this guide, we

explain how Nipper OmniSight, our network exposure assessment platform, can support you in protecting your operations.

And crucially, this can also help with a different AI challenge businesses are wrestling with: how to safely permit their teams to use AI tools in their work.

If reading this gives you food for thought, and you would welcome support in refining your approach to AI-enabled threats, please do get in touch.



Andrew Woodford
Chief Technology Officer,
Titania

Quickstart guide

Need to get answers for the Board right away? Here's where to find them.

How is AI changing the threat landscape?

AI hasn't yet led adversaries to create brand new attack vectors – because it means they don't need to. Instead, AI is dramatically increasing the speed and ease with which attackers can find and exploit existing vulnerabilities. Plus there's another risk to be conscious of: employees using AI tools in an uncontrolled way in their daily work.

See [How AI is \(really\) changing the threat landscape](#) for more.

Are we at risk?

Yes: the AI threat is present right now and almost every organisation is vulnerable, because existing approaches are too slow and reactive. This is discussed in [How AI is \(really\) changing the threat landscape](#).

How are we exposed?

That depends on how your networks are structured – are they segmented? – and how up-to-date your vulnerability management is. Your first step should be to analyse this exposure to give the Board clear answers.

Step 1: Analyse and prioritise

How can we shut the risks down?

You can't shut them down completely, but by switching to a more proactive approach and taking foundational security measures, you can substantially reduce both the attack surface and the impact of being breached. See [Step 2: Harden](#) and [Step 3: Segment](#).

Importantly, these steps will also reduce the risks associated with colleagues using AI tools.

How much will this cost us?

That depends on the size and complexity of your network and what steps you have already taken. But when the impact of breaches can be measured in the millions, and there are higher penalties for non-compliance, investing in foundational security – including a focus on recoverability when you are breached – can not only save you money: it could save your business.

What are the immediate steps?

Analyse the network, harden it to shut down known vulnerabilities, segment it to add layers of protection for key assets and restrict unnecessary access (including for AI tools!). Then check that all devices are following the policies and rules correctly. That's steps 1-5 of this guide!

What's the long-term plan?

These threats aren't going away, so once you've implemented a defence-in-depth strategy for your network, you need to monitor it on an ongoing basis.

See [Step 6: Monitor](#).


How AI is (really) changing the threat landscape

The impact of AI is everywhere, and cyber security is no exception – but it may not be in the way many expect.

To date, there are no signs of adversaries using AI to create wholly new methods of attack. Instead, they're using it to intensify existing approaches: accelerating their probing for weak passwords, insecure default configurations, unprotected devices and control gaps.

Because so many of these exploitable vulnerabilities exist, attackers don't need new methods.

With AI automation, they can make many more attempts to find a way in... much, much faster. And once they have breached the network, they can then apply the same approach to open the next locked door, accelerating lateral movement from the point of breach (such as a forgotten router) to high-value targets.



The sheer scale that AI provides was perfectly demonstrated by the discovery of the “Fortibleed” dataset – an online list of the credentials more than 75,000 compromised Fortinet firewalls and VPN devices. Analysts concluded that these credentials had been largely gained through automated trial-and-error techniques, using previously compromised credentials and known vulnerabilities. Some reports suggested that more than a billion such attempts were made – indicating a highly automated attack, almost certainly underpinned by AI.

Why conventional defense is no longer enough

When attackers can operate at these sorts of speeds, conventional defensive approaches of periodic assessment, detection and patching cannot keep up.

- Detection remains essential, but AI-powered attacks can move laterally across a network within minutes. By the time a breach is identified, critical assets may already have been exposed.
- Patching is still vital, but it's reactive; a means of covering up a weakness that has already been discovered. And the reality is that most organisations are already behind on their patching. As software issues continue to be exposed by AI, patches will be released even faster – and few will be able to keep up with. This has led the US Cybersecurity and Infrastructure Security Agency (CISA) to issue a press release urging organisations under its regulatory ambit to “patch smarter, not harder.”

Responding to the risk of AI-enabled threats demands a different approach, where efforts to stop breaches are combined with – or even secondary to – work to minimise the potential damage that attackers can cause. This is pragmatism, not defeatism; organisations need to tackle known vulnerabilities in a risk-prioritised

way, focusing their resources on the biggest risks to their operations and taking steps to prevent lateral movement when they are exposed.

That requires network segmentation: introducing barriers that restrict movement within the network and limit access key resources, tools and data. Segmentation must then be enforced by robust access controls, starting from a principle of least privilege access to the most important parts of the network.

With these crucial steps in place, organisations can significantly reduce the havoc that AI-enabled threats can wreak. The subsequent sections of this guide explain more.



Recognising and responding to the threat from colleagues using AI

While organisations are rightly concerned about attackers harnessing AI, there's also a risk closer to home.

Staff want to use AI to help them in their work. Where it drives efficiency and productivity, few leaders are going to object. But there are security risks to this, even where staff are using the tools constructively and responsibly.

Ask AI a simple reporting question, and it will look at all the data that the user is permitted to access to inform its response. Yet that could include commercially sensitive information, or customer data that shouldn't be shared.

All of a sudden, data is unwittingly exposed to the AI training models, as well as to others who view the report.

Organisations may also want to limit where AI tools users can be used. For example, AI-based fault detection systems can be incredibly valuable, but you may not want to permit them in an Operational Technology environment, because of what it could do with any faults it uncovers.

Clearly, effective policies are a key part of this, determining what can be safely allowed while monitoring use where possible. But the principles of segmentation and access control are critical to this area too: they provide the guardrails to limit what AI can view inside the organisation.




What's on your mind today?

+ Ask anything



Step one

Analyse and priorities



The fact that AI can move so fast has changed the rules for defence. Instead of trying to plug every vulnerability, organisations need to adopt a risk-based approach, where you address the most urgent issues first. Here's why, and how.

Using automation, attackers are finding vulnerabilities long before defenders do, and exploiting them before they can be remediated.

Already, the number of publicly identified Common Vulnerabilities and Exposures (CVEs) was increasing at a rate faster than most organisations could address. During

the first three months of 2026, numbers were nearly one-third higher than the corresponding period in 2025.

It's not just the sheer volume of threats that is overwhelming. The Cloud Security Alliance reported that in 2025, 32% of new exploits were "Zero Day" – that is, there was no prior warning for them, meaning there were no patches available... so no way for businesses to patch their way to safety.

When breach is consistently faster than remediation, no organisation can sustain a policy of seeking to shut down every vulnerability. Instead, businesses need to prioritise their defensive effort, based on the real risks they face.

Unmask the inside threats

One of the most important places to start is the network itself: the routers, switches, and firewalls that control traffic, enforce segmentation and manage access.

These are now prime targets for attackers: a network device, such as a router or firewall, will process large volumes of traffic and connect to many areas of the network.

Compromise an endpoint, and you're a long way from sensitive data and key assets: compromise a router, and you're on the fast track towards them.

Yet repeated studies have found that network devices are not protected as well as they should be.

In its "Riskiest Connected Devices of 2025" report, Forescout concluded that routers represent more than half of today's most vulnerable devices.

The reasons lie in foundational network management issues. Routers often have open ports used for administrative roles. Some have been overlooked when patching was rolled out. Others are still set up with default configurations and weak

passwords.

Despite these risks, most security tools don't look at device configurations and don't validate whether controls are in place.

Analyse your network

So to prioritise the defensive effort, first analyse your network. Complete a full inventory; review configurations and spot where devices haven't been hardened.

You can then delve deeper to understand potential attack paths: if Router X was breached,

Such analysis enables you to identify the critical risks to your operations and make informed decisions about how to resolve them. It can help you pinpoint quick wins: a particular device type that needs updating network-wide, or a strategically placed router where changing access rights will shut off several potential paths.

What could attackers then access?

What routes could they take to reach vital operational tools or customer data?

What barriers do you have in place to close these routes, and are they sufficient?



How Nipper OmniSight can help

Nipper OmniSight is Titania's network exposure assessment platform, designed to help enterprises adopt a proactive, risk-based security approach.

Applying the same methodology used by Nipper InfraSight but on a network-wide level, Nipper OmniSight can help identify exactly where the most critical risks are. It does so at speed, highlighting where:

- **Devices are misconfigured or not correctly hardened**

- **Access policies are not being enforced correctly, and**
- **Configurations would jeopardise compliance with key regimes.**

It then applies a risk-scoring method to list the most important issues first – along with clear guidance on how to resolve them.

In short, it can do the prioritization for you, based on the level of potential exposure and ease to fix.

Check	Findings
Log config changes enabled	CiscoOS15.sample.device was configured to not log configuration changes.

Remediation - CiscoOS15.device

The following will configure the device to log configuration changes.

Risk Level - Safe

```
archive
log config
logging enable
exit
exit
```

Step two

Harden



Hardening network devices will slow down adversaries and reduce the attack surface. On its own, it won't stop all attacks, but it will close down the easiest routes for attackers to exploit.

Network analysis invariably identifies foundational security issues that organisations can address to reduce their likelihood of being breached and make it harder for attackers to move laterally.

These will be tasks like:

- **Replacing weak passwords**
- **Updating default device configurations**
- **Applying relevant patches**
- **Strengthening encryption processes**
- **Addressing control gaps or rule conflicts, that result in exploitable exposures.**

Studies repeatedly find that these basic issues are the key targets for attackers. And while it's almost impossible to address them all, the analysis can pinpoint where you should start, to rapidly close down critical vulnerabilities.

Supporting compliance

Doing so can also help with compliance – especially as more regulatory regimes follow the CISA recommendation to patch smarter, not harder. For example, the Cyber Operational Readiness Assessment (CORA) introduced for US military and the UK's National Cyber Security Centre Cyber Assessment Framework (NCSC CAF) are focused on outcomes, rather than pure inputs. Addressing device misconfigurations and weak credentials, for example, are key to the CAF objective of Managing security risk.

Both the NIST cybersecurity framework and NERC CIP (Critical Infrastructure Protection), which applies to the North American energy sector, demand clear evidence of hardening.

Network hardening: An ongoing journey

Hardening is not a one-off task. Passwords need revisiting. New threats will emerge. New patches will be released.

Configuration drift can create control gaps. And in all but the smallest of networks, the

initial hardening effort will have only dealt with the most urgent issues.

So it pays to be systematic and check regularly for new risks and exposures, using trusted independent sources like MITRE ATT&CK.


 **NIPPER** OmniSight

How Nipper OmniSight can help

Nipper OmniSight provides a network view of risk, helping understand which steps are urgent –because you are non-compliant, because you are at significant risk, or both. It prioritises fixes that can achieve the most (e.g. patching a particular class of router that is widely used on the network) and reports against MITRE ATT&CK and other sources of information about observed tactics and techniques.




Once you have followed the remediation guidance, analyse your network with Nipper OmniSight again to confirm that the issues have been successfully addressed.

Devices assessed	Findings	Critical and high
11	735	156



Step three

Segment



Segmenting your network can close down attack paths and limit the blast radius of any breach. But it needs to be balanced against performance and operational considerations – and enforced through access rules that meet business and security requirements.

Steps 2 and 3 can happen in either order – depending on your priorities and current situation – or simultaneously if resource allows. Hardening can easily be undertaken by network engineers, while leaders focus on the segmentation strategy.

The same analysis used to underpin targeted hardening should also inform network segmentation. This is a deliberate process of dividing the overall network into smaller sub-networks, with access strictly controlled, to provide additional layers of defence for critical assets such

as operational technology, cyber physical systems and the most sensitive data.

If attackers breach the perimeter, the 'blast radius' of an incursion is restricted: attackers are contained within the segment they have penetrated.

Segmentation is widely viewed as an effective method of protecting key systems. It is integral to the International Society of Automation IEC 62443 standards for secure industrial automation and control systems: these specifically recommend a "demilitarized zone" being put in place between public-facing segments and the automation and control systems, to help ensure that there are at least two further layers of security in place.

The EU's Digital Operational Resilience Act (DORA) meanwhile requires banks, insurers, and investment firms to "design the network connection infrastructure in a way that allows it to be instantaneously severed and shall ensure its compartmentalisation and segmentation, in order to minimise and prevent contagion" in the event of an attack.

Effective segmentation may require network changes, additional hardware purchase and provisioning, plus changes in policy and routing rules. So it's not a small task, and often requires extensive input from different teams to assess whether

proposed segmentation would unduly restrict access, or create barriers between parts of the operation (for example, between the OT system and the supply chain management platform).

Macro vs micro-segmentation

There are two segmentation approaches:

- **Macro-segmentation, which effectively breaks the entire network down into functional sub-networks, and**
- **Micro-segmentation, which (as the name suggests), reduces it to smaller segments still.**

A useful analogy is that micro-segmentation is locking the door of every important room in the building; macro-segmentation is restricting access to entire secure corridors.

The optimal answer is to do both. That way, attackers not only have to break down a locked door each time, but they also can't identify which 'rooms' are in which corridors, so don't know where to target.

Macro-segmentation should be the first step for those starting from scratch, as it provides more extensive protection sooner. But for organisations that have already begun with micro-segmentation, applying macro-segmentation too will rapidly strengthen defences.

Create segmentation

Name

Baseline segment

Labels

100-devices ✕

1000-devices ✕

The role of access controls

Segmentation requires clear and robust access policies to function effectively. The default position should be that only those who expressly need access to a segment should be permitted it: essentially least privilege access (LPA) applied on an organisation-wide scale.

Adopting LPA as a default can bring the security team into conflict with colleagues

in other functions, who would prefer “just in case” access. But from the security management perspective, it’s far more prudent to grant access on a short-term basis, where strictly needed: this increases control, and makes it far less likely that users might be forgotten when access rules change, or that those same users could expose critical segments through their own inconsistent security behaviour.

How segmentation supports AI use in the business

As well as protecting the business from AI-enabled threats, segmentation can play a vital role in permitting the safe use of AI tools.

In the scenario outlined earlier, we highlighted that AI will use whatever

data it can access to answer queries. Segmentation supported by LPA means that unauthorised users cannot access sensitive data; therefore, any AI tools on their desktops will not be able to either.

 **NIPPER** OmniSight


How Nipper OmniSight can help

Nipper OmniSight enables attack path mapping: visualising the routes attackers can take from different entry points to key

assets. That understanding can inform segmentation decisions: where the barriers should be placed for maximum benefit.

Step four

Validate



Once you have taken steps to improve protection against AI-enabled threats, reassess your network to ensure the measures are working as intended.

This is a relatively simple step, but an essential one nonetheless. You cannot take for granted that your network hardening and segmentation actions have achieved the desired outcomes.

Changes can inadvertently create new problems or rule conflicts. Individual devices may not be enforcing the access policies as intended. Plus there's the possibility that, since your initial assessment, new vulnerabilities have been identified, new devices have been added and new users granted access.

That's why it's vital to validate your changes. This does not need to wait until an entire project is complete; it can take place after each round of remediation.

Validation should not take long – especially if you have a solid grounding from the initial analysis – because it will focus on what has changed. If there are issues, you can address them promptly. If all is going well, enjoy the win.

 **NIPPER** OmniSight

How Nipper OmniSight can help

Just as it can help deliver a rapid initial analysis of network device configurations, Nipper OmniSight can also be used on a regular basis to validate the changes. It will confirm whether devices are enforcing segmentation rules correctly and assess configuration against the latest

independent risk data.

Then, like with the initial analysis, it will generate a detailed report, with risks clearly prioritised. This should show the positive impact of the measures you have taken – which can also be helpful for internal reporting.

Step five

Prepare for breach



The sheer prevalence and pace of AI-powered attacks means that, even after hardening their networks, most organisations are likely to be breached at some point. The challenge is minimise the impact.

That requires a realism across the business and practical responses, from the boardroom to individual employees, who may be the first to see signs of a breach. Clear and well-tested policies are essential. From the IT or security side, that might include knowing exactly what devices and

connections should be shut down to stop AI-enabled threats proliferating or moving laterally between segments.

It would also include recovery policies: how do you restore the network to a controlled operational level and then back to full operation?

An up-to-date configuration management database can help with this process, as it means that instead of having to reconfigure every device based on corporate memory, you can restore to the last known good configuration.

Beyond technology

There are also people and process considerations. Who needs to be informed – regulators, partners, customers even? There may be different thresholds and timescales around this. For example, letting operational partners know there's an issue as early as possible can help manage production or mitigate losses along the value chain, whereas reporting to regulators may follow a different pattern.

It's also vital to have forums that allow rapid decision-making during a breach or similar crisis. Changes cannot go through a complex gateway approval process.

 **NIPPER** OmniSight

How Nipper OmniSight can help

Nipper OmniSight (Integrated) and (Continuous) can both help organisations keep their CMDB up-to-date, which is vital to a rapid recovery.

Step six

Monitor... and keep monitoring



Steps 1-5 strengthen your network defences against AI-enabled threats. But as most organizations are acutely aware, protecting your network is not a one-off process. Instead, it's a ongoing battle, which makes monitoring a vital front-line defence.

Where traditionally organizations may have monitored for "indicators of compromise", the more effective approach now is to monitor all changes on the network. These could be the first signs of a breach, but frequently the initial change is simply configuration drift.

For example, permissions may have been changed in the short term to allow specific access; these need to be changed back, before they can be compromised.

A new device could have been installed, but not hardened correctly: again, creating a potentially exploitable vulnerability.

Effective monitoring tracks these changes, alerting security teams to unplanned exposures and to check any planned alterations have been completed correctly.

This kind of proactive monitoring enables you to keep the attack surface as small as possible – and to respond promptly to anything potentially more damaging.



Smart monitoring also keeps a close eye on threat databases – KEV and MITRE ATT&CK for instance – so you can get ahead of any newly discovered vulnerabilities or tactics.

How often do you monitor?

The optimal solution is continuous monitoring of the network. However, not all organizations have the resource to commit to this – either in terms of the staffing to keep an eye on alerts, or being able to invest in new software that offers 24x7 monitoring.

Instead, the decision becomes a trade-off between risk and resource. Perhaps you run a monthly analysis, reflecting internal

change cycles, or opt to analyse the network configuration each weekend, when colleagues aren't working.

The issue with either approach is that AI-enabled threats could easily penetrate your network in between – so you're reliant then on the vigilance of your breach detection systems... and the effectiveness of segmentation and access controls to prevent the threat moving laterally.

 **NIPPER** OmniSight

How Nipper OmniSight can help

Nipper OmniSight (Continuous) is a solution designed to enable continuous network monitoring. It maintains a comprehensive CMDB, as the basis for spotting every change to the network as soon as it happens.

For organisations that can't justify continuous monitoring, Nipper OmniSight (Standalone) and Nipper OmniSight (Integrated) both enable scheduled network analysis; Nipper OmniSight (Integrated) does this against the CMDB it has populated, so again spots every change but does not do so in real time.



The role of AI tools in defence

With organizations increasingly – and rightly – focusing on AI-enabled threats, there's a wholly understandable line of thought that AI should be integral to your defences too. But it's not quite that simple.

AI is likely to play a critical role in cyber defence, but organisations should be cautious about granting it unrestricted authority. Human oversight remains essential for high-impact decisions and for balancing security with operational requirements.

The strength of AI in security terms is monitoring behaviour patterns on the network, and detecting anything unusual in the vast volumes of traffic.

That can provide insights far sooner than relying on human detection or even existing intrusion detection tools. However, the alerts it generates need to be monitored by humans to determine whether they are actual causes for concern, and how best to respond.

One reason for this is that most AI systems identify threats through pattern recognition, statistical inference and probabilistic

reasoning rather understanding threats in context. This was amply demonstrated by Claude Mythos AI identifying thousands of vulnerabilities, including some Zero Day vulnerabilities, within code – but being unable to determine whether each vulnerability was actually exploitable, or if they were closed or isolated systems, with no external access.

The problem is that businesses have to examine all the flagged vulnerabilities to work out which are relevant: a sizeable task, to reach a potentially very small number of genuine exposures.

Secondly, there's the simple possibility of AI getting things wrong: whether over-permissive or over-protective, the AI tool could cause significant collateral damage.

Even when AI correctly identifies a threat, determining the appropriate response often requires business judgement. Automatically shutting down a critical production

system may reduce cyber risk but create unacceptable operational disruption.

Finally, AI needs to be strictly controlled; in particular, in terms of who can access its analysis of vulnerabilities. If the AI solution has detailed information about your network, you clearly do not want this available to others.

So while AI should be part of the defence, it should be deployed strategically and within a clear framework governing what actions it can take without deferring to a human.

Most AI systems identify threats through pattern recognition, statistical inference and probabilistic reasoning rather than understanding threats in context.

Possibility of AI getting things wrong.

AI needs to be strictly controlled; in particular, in terms of who can access its analysis of vulnerabilities.



Defence in depth: The only answer

Because AI-enabled threats can attack on many fronts simultaneously, defence also requires multiple layers and facets.

Organizations should continue with their detection (potentially using AI tools), patching and vulnerability scanning, but complement that with segmentation and well-enforced controls.

Together, these build resilience and readiness: they reduce the attack surface that AI-enabled threats can target, and restrict what these threats can access if they do penetrate the network.

Monitoring remains a crucial part of the overall defence, to identify and close down

exposures before they are exploited. And finally, recoverability is critical: organizations need to be ready to act decisively to shut down threats, and find the safest way to restore operations swiftly.

Technology like Nipper OmniSight is invaluable for this, but it's not the whole solution: processes need to be adapted to the new threat landscape, and organizations need sufficient resource to prepare the network, monitor it and respond to alerts in a timely way.

With the spotlight firmly on AI-enabled threats, there is a substantial opportunity for security leaders to secure board level backing to help reduce risks.

This isn't a future issue. It's a live challenge right now: AI-enabled threats are already breaching organizations, disrupting operations and exfiltrating data. Businesses cannot afford to wait and see. You need to start building defence in depth today.

Next steps

To find out more about how Nipper OmniSight can help you prepare for AI-enabled threats, or see a demo, contact us today.

[Book a demo](#)



USA

Suite 600,
2451 Crystal Dr, 6th Floor,
Arlington, VA 22202
enquiries@titania.com

UK

167-169 Great Portland Street,
London, England,
W1W 5PF
enquiries@titania.com