Article

Nov 2026

AI can find security vulnerabilities in minutes. Can you fix them as fast?

Why you need to automate vulnerability and patch management, and how to do it safely without overwhelming your security team



David Dove
Senior Information Security Consultant
Email david.dove@waterstons.com

AI is changing the economics of vulnerability discovery. Security researchers, software vendors and threat actors can increasingly identify and analyse weaknesses at greater speed and scale. October's PHP vulnerabilities show this is already a reality: two critical flaws, reportedly uncovered using AI, put millions of sites worldwide at risk and were being actively exploited within days of patches being released. That speed puts your patching process under pressure, because manual assessment, ticket creation, approval and deployment can't keep pace.

Hiring more people or asking your team to work faster won't fix it. Advice from the UK National Cyber Security Centre (NCSC) is to "update by default": automate repeatable, low-risk decisions, and keep human oversight for critical services, legacy systems, operational technology (OT) and changes with a significant business impact.

This isn't just a tooling decision. It's a governance one. You'll need clear ownership, good asset data, agreed risk tolerances and investment in automation you can show is safe and measurable.

Why vulnerability management has to change

The NCSC has warned organisations to prepare for a potential "vulnerability patch wave" – more software updates driven by vulnerabilities found across existing technical debt.

That isn't just a security problem. If you can't find and fix exploitable vulnerabilities quickly, you raise the risk of service disruption, data loss, regulatory scrutiny and unplanned operational expenditure. You can also be pulled into repeated emergency patching cycles that compete with your change programmes and other strategic priorities.

So don't aim to patch every vulnerability. Aim to reduce your exposure to the vulnerabilities most likely to cause material harm, while keeping your critical services available and stable.

Where should you automate?

Take a risk-based approach. Automate the low-risk work so your people can focus on the systems where a mistake would hurt most. The table below shows where each level of automation could sit.

Automation level	Suitable examples	Governance approach
Automatically remediate	Endpoint operating system/application patches, browser updates, endpoint configuration hardening, security agent updates and approved security baselines for standard endpoints	Pre-approved policy, phased deployment and automated rollback
Automatically prepare and recommend	Server operating system/application patching and configuration hardening, cloud configuration changes, application remediation actions and infrastructure hardening	Automated prioritisation and testing, followed by owner approval
Escalate for human decision	Critical production platforms, legacy systems, OT, major-versions and safety-critical environments	Formal risk assessment, testing, business approval and compensating controls

Six things to get right before you automate

1. **Clear ownership for assets and business services** Nothing improves until someone owns it. Every asset and business service needs a named owner before you start. Bring them into the planning early. They'll help set your risk appetite and flag where automation isn't right, whether because of risk tolerance or technical limits.

2. **Know what you have, and where you're exposed** You can only patch what you know exists. Manual inventories go stale fast, so automate discovery across endpoints, cloud workloads, SaaS, containers and shadow IT. Then look outside in: start with your external attack surface before working inwards across cloud and on-premises environments.

3. **Prioritise by risk, not by volume** Trying to patch everything is a fast route to burnout. A modern vulnerability programme prioritises automatically, based on:

- Exploit availability
- Active exploitation
- Internet exposure
- Asset criticality
- User privilege levels
- Presence of sensitive data
- Business service impact

4. **Decide where automation and AI add the most value in remediation** Review your process and decide where human judgement is genuinely needed, and where it isn't.

Tools can now automate deployment, phased rollout and rollback for Windows updates, third-party applications, browsers, endpoint agents, and standard server patching. As well as configuration enforcement, security baseline management and infrastructure hardening. Apply automation to the whole remediation process, not just patch deployment.

That leaves human approval for what matters most: critical production systems, OT, legacy applications and major-version changes.

AI earns its keep by answering "what should I patch first?" It can correlate threat intelligence, analyse exploit trends, flag vulnerable internet-facing assets and help predict what's most likely to be exploited.

5. **Build security into the development lifecycle, not just production** Don't wait until software is live. If you build or buy software, check security at every stage, not just at release. AI-assisted tools can help, but validate their findings, and fold them into your existing secure development practices.

6. **Verify remediation, and track what good looks like** Don't just patch and hope. Confirm the patch applied, the vulnerability has gone, and the service still works, then automate rescans to prove it. Set a board-level KPI to show the automation is working, for example the percentage of actively exploited critical vulnerabilities remediated within target, or the percentage of internet-facing assets under continuous monitoring.

Patching won't always be enough. Some vulnerabilities can't be fully fixed by a software update, and end-of-life or legacy technology may not be able to receive updates at all. In those cases, reduce risk another way: disable unnecessary services, tighten configuration, use security hardening measures, network segmentation or access restrictions, or add other compensating controls. If the technology sits on your external attack surface, you may need to upgrade it, replace it or bring it back into vendor support.

AI won't replace skilled security and technology professionals. It'll make their judgement matter more, because there is more to decide and less time to decide it.

Final thoughts

The organisations best placed to respond won't be the ones attempting to patch everything manually. They'll be the ones that understand their assets, prioritise vulnerabilities by business risk, automate repeatable decisions and save human expertise for the systems and circumstances where judgement matters most.

If you're a CIO, CTO or CISO, the question is no longer whether to automate vulnerability and patch management. It's whether you can introduce that automation safely, measurably and quickly enough to keep pace with the changing threat environment.

Unsure where automation is safe in your estate? We'll help you map ownership, priorities and risk tolerance, and build a plan you can defend to your board. Talk to us.


Digging i
Producti
& Automa