HPE X CHANGE MAKERS CLUB

AI DIDN'T BREAK YOUR SECURITY. IT PUBLISHED IT.

AI has not introduced a new class of security problem. It has industrialised the discovery of the ones you already had. Security success now depends on disciplined housekeeping and response at machine speed.

LONDON · JULY 2026

changemakers.

HPE

01 THE VERDICT

AI didn't create a new security problem. It found the ones you had.

> *"I don't think AI is creating your security problems, it's just exposing the ones you already had"*

Jaye Tillson, Field CTO, HPE

The clearest story of the evening was not an attack. A customer deployed an internal AI assistant and it went looking. It found the HR database: salaries, sickness records, who was on a PIP. Employees could simply ask it.

The database sat on a file server everyone already had access to. It had been live four months. It surfaced only because one employee decided it was wrong and said so.

Nothing changed when the AI arrived. The exposure was already there. What the AI supplied was the motivation to look everywhere at once.

BY THE NUMBERS

15 minutes
From patch release to exploitable. Microsoft's figure, as cited in the room.

4 months
How long an exposed HR database sat live and unnoticed.

120,000
Non-human identities in a 12,000-person company, up from near zero.

IN THE ROOM

Six months ago we had twelve thousand humans to secure. We now have roughly one hundred and twenty thousand non-human identities. I have AI agents working for me.

Where we used to be able to take a few more risks, we can't, because of the pace and the scale of exploitation.

CISO, global software company (12,000 staff)

TAKEAWAY

AI is not the threat actor here. It is the auditor, untiring, indiscriminate, and reporting to whoever asks it a question first.

02 WHERE IT BREAKS DOWN

Five exposures AI has already found you

1 Patch cycles that were always fiction

One CISO compressed patch SLAs from 48 hours and 30 days to 1, 2, 4 and 8 hours, then called it untenable, because it is. Elsewhere: a core switch up for fifteen years, needing 42 patches and 42 weekends of downtime. They replaced it instead.

2 Nobody can say where the data is

When asked whether they used the cloud, one organisation said no, then listed Microsoft 365, SharePoint, OneDrive, Salesforce and Workday. AI's ability to chain findings turns every gap in that map into an opportunity.

3 Identity at machine scale

Twelve thousand human identities became roughly 120,000 non-human ones inside six months, in one company. Identity management was built for joiners and leavers on a human timescale. Nobody has defined what a leaver looks like for an agent.

4 Legacy that is beyond patching

You are not undoing the latest vulnerability, one participant argued, you are undoing fifty years of legacy. Manufacturing systems nobody left understands, and undocumented protocols AI now flags as risky because there is no history on them.

5 Risk sign-off nobody can define

“It's great that they're willing to accept it, but what exactly are they accepting?” asked one senior governance executive, on executives waving through AI pilots. Boards see efficiency and advantage, and will proceed regardless. A 5×5 heat map cannot hold this.

THE PATTERN

Not one of these was created by AI. Every one predates it, open file shares, unpatched switches, undocumented systems, unmapped estates.

What changed is the interval between a weakness existing and someone finding it.

So the work isn't new either.

02 THE FIX

Housekeeping is now a security strategy.

Nobody proposed a new control or a new category of product. The fix was the least fashionable one available: know your estate, answer machines with machines, and rehearse the recovery you assume you won't need.

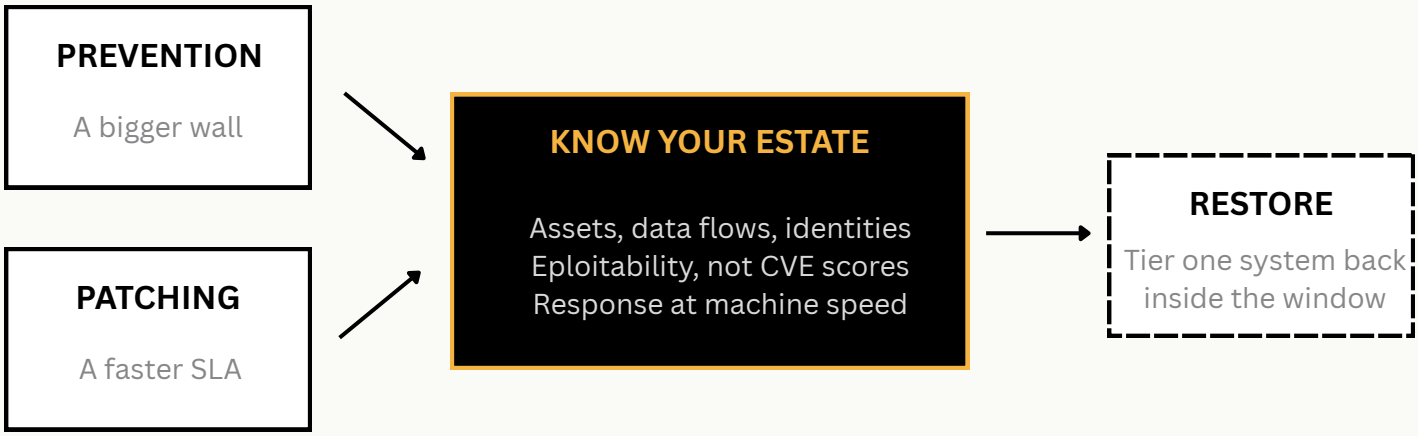

WHAT CHANGES

Same disciplines, radically shorter timescales. Inventory, architecture and identity hygiene become the load-bearing controls.

Start left, not shift left. One CISO's correction to the industry's favourite phrase. Shifting left moves testing earlier; starting left means less exploitable surface to find at all.

Retire the CVE score as a scoreboard. Vulnerability management was declared dead outright, not difficult, dead. What replaces it is exploitability: toxic combinations and posture management, so scarce resource goes where an attack could actually be chained.

Answer machines with machines. A developer gave Codex full access to fix his Teams client. It threw 42 Defender alerts in a hundredth of a second. Orchestration cut the machine off the network in around 400 milliseconds.

Air-gap the backups and the documentation. In one recovery, Active Directory was down and SharePoint unreachable, so the runbooks for the systems being restored were themselves inaccessible. Backups were infected and nobody knew how far back.

Define the minimum viable business. The specific list of what has to run for the company to keep trading. Most cannot produce it. The comparison offered was Lehman in 2008: the loss of liquidity killed businesses, not the event.

> ***"The strategy isn't, can I build a bigger wall. The strategy is, how do I restore my Tier-one systems within a short window"***
>
> Roundtable Participant on Major Global Breach

03 MEASURE WHAT'S VALUABLE

Measure what matters now. Not what mattered before.

Every metric on the left was designed for a world in which finding a weakness took a human being a long time. That assumption has quietly expired and the measurement stack built on top of it has expired with it.

EASY TO MEASURE - TODAY	VALUABLE TO MEASURE - INSTEAD
Patch SLA compliance	**Restore time for tier-one systems**
CVE counts and severity scores	**Exploitability: micro-combinations that chain into a real attack path**
Alert volume and analyst throughput	**Time to automated oslation, in milliseconds**
5x5 risk heat maps	**Quantified exposure: what it costs, what spend reduces it**
Controls deployed at the perimeter	**Minimum viable business - defined and rehearsed**

If you score risk from one to ten while the systems attacking you operate on a scale of one to ten thousand, "low" is not conservative. It is meaningless.

The proposed answer was quantification: this costs this much, this spend reduces it by this much. The counter from the risk side: quantification is hard even after an incident, with the invoices in hand.

> *"There is no such thing as security. Only the illusion that you are secure, because most companies don't know when they're penetrated"*
>
> Anonymous, Breach-Recovery Specialist

On insurance. One leader described buying substantial cover while doubting the underwriters would pay. Another offered the other side: a friend who assessed cyber claims and never once paid out, because he could always find a control failure somewhere.

On fiduciary duty. Directors who sign off catastrophic risk without adequate resilience may have breached their duties and be personally liable. Sign-off is not coverage.

Final thought. AI has not made security impossible. It has removed the margin of error that poor housekeeping used to enjoy. The organisations that come through are the ones that can say what they own, where the data goes, what has access, and how fast they can put it back.

The open question. Asked what the future looks like, the most honest answer came from the participant with the most exposed client base: "What is the future for us?"

CHANGE MAKERS

Change Makers is a global community of senior executives driving long-lasting change within large, complex organisations. Members bench-test their thinking, share peer review, and discuss the opportunities and challenges of leading through transformation.

WWW.CHANGEMAKERSCLUB.COM

The insights in this report were drawn from a private roundtable of senior technology and information security executives, facilitated by James Harris, Founder of Change Makers.

HOSTED & FACILITATED BY

James Harris

Founder, Change Makers Club

In partnership with HPE· London, July 2026

CONTRIBUTORS

Martin Jackson - Platform Engineering Lead - HMRC
Anu Awasthi - Director, Product Owner - Finance Cloud, Analytics, Automation & Process Mining Products - GSK
Jonathan Loretto - Chief Architect - Bank of England
Will Davies - Group Deputy CISO - Endava
Jodie Bennett - Operations & Governance Managing Director - WTW
Anuradha Sekhar - Director of Strategy & Transformation - WPP
Agwu Nwoke - Industry Security Programme Director - GSMA
Angus McLean - AI Director - OLIVER
Deenar Toraskar - Quant Modelling Director - JP Morgan
Richard Fern - SASE & Security Specialist - HPE
Jaye Tillson - Field CTO - HPE

changemakers.

HPE