

THE LAST ATTACK SURFACE

Why the future of threat will always be human

By Steve Wright

TABLE OF CONTENTS

03 Looking Beyond The Technology

04 The Race Between Attack and Defence

05 When Technology Becomes a Liability

06 The Human Domain

07 The Insider Threat Is Always There

08 Why Organisations Don't See the Human Attack Surface

10 The Paradox of Security

12 How Human Capability Accumulates

13 The Future Insider Threat

14 Questions for Organisation

15 Conclusion - The Last Attack Surface

LOOKING BEYOND THE TECHNOLOGY

For decades, organisations have held a simple belief: that with sufficient investment, innovation and control, risk can be engineered out of systems. From physical security to cyber defence, each technological leap has promised greater protection, visibility and, ultimately, greater control. Firewalls replaced fences, analytics replaced intuition, and automation reduced reliance on human judgement.

Artificial Intelligence represents the latest, and perhaps most powerful, expression of that belief.

We are rapidly approaching a point where AI-driven defence can identify anomalies in real time, predict malicious activity and respond at machine speed. At the same time, offensive AI use is advancing just as quickly, automating reconnaissance, enhancing deception and identifying vulnerabilities at scale.

The inevitable outcome is parity, not dominance.

Attack and defence evolve together. Each gain in capability is met with a countermeasure, and each advantage is temporary. Over time, this can produce a technological equilibrium in which systems become increasingly resilient, monitored, and self-defending.

But history suggests something important: when one method becomes too difficult to exploit, the adversary does not necessarily give up. Instead, they shift the problem.

This paper argues that the human being is different from every other attack surface because people remain dynamic. They respond to trust, pressure, relationships, incentives, fatigue and opportunity. They can be influenced in circumstances that no technical control can fully predict.

The insider threat should therefore not be viewed simply as what remains after every technical defence has failed. It exists alongside those threats because organisations must continue to operate through people, who must continue to have access, information and responsibility.

> *When technology reaches its limits, the focus does not disappear. It narrows.*
>
> Steve Wright

01 THE RACE BETWEEN ATTACK AND DEFENCE

Security has always been a race. One side develops a capability, while the other develops a way to detect, disrupt or defeat it. New controls emerge, and new methods are developed to bypass them, and the cycle continues.

Artificial Intelligence accelerates that contest. Defensive systems can process vast volumes of information that would overwhelm human analysts, identify patterns and respond rapidly. Offensive capability also benefits from the same acceleration: reconnaissance, deception and the search for weaknesses can all be conducted at greater scale.

That undermines the assumption that better technology will eventually produce a decisive advantage. If both sides improve, one side's progress may simply shift the other's baseline.

For organisations, this can create a false impression of success. While security metrics might look better, detection times may decrease and technical breaches may become less common, the adversary's core goal remains unchanged: to gain access, information, influence, or cause disruption.

WHEN THE TECHNICAL ROUTE GETS HARDER

If a network becomes extremely difficult to penetrate directly, the value of attacking it diminishes. If a communications channel is heavily monitored, the value of using it diminishes. If automated controls quickly identify a pattern, repeating it becomes inefficient.

The adversary can instead pursue an alternative route to the same objective. This is the strategic importance of the human domain.

A person is simultaneously a user, decision-maker, custodian, relationship, source of information, and potential point of access. These functions cannot simply be patched or switched off without changing how the organisation operates.



Takeaway: When the technical route becomes harder, the objective does not disappear. The method can change.

02 WHEN TECHNOLOGY BECOMES A LIABILITY

I have seen this principle in practice in Afghanistan in the early 2000s. Mobile phones offered obvious advantages for communication and coordination, but they also created opportunities for interception, location identification and tracking.

Behaviour adapted. Communication methods were simplified, reduced, or abandoned. In some circumstances, coordination reverted to face-to-face interaction or low-tech alternatives that offered less capability but greater survivability.

The lesson was simple: The most sophisticated method was not necessarily the most effective in that environment.

Not more advanced - More appropriate.

This is important because organisations can become so focused on improving the sophistication of their controls that they assume the adversary will continue to challenge them in the same technological space. But when technology becomes a liability, adversaries can simply bypass it.

THE PROBLEM WITH ASSUMING PROGRESS IS LINEAR

There is a tendency to view security development as a ladder, with each new technology assumed to sit above the previous one. From experience, threat environments do not necessarily work that way.

A lower-technology method can become effective precisely because the organisation has stopped looking for it. When attention is focused on sophisticated systems, something old and unfamiliar can be unexpectedly hard to detect or use.

Consider someone walking into a police station today carrying important evidence on a five-and-a-quarter-inch floppy disk or a Zip disk. The data may still exist, but the organisation may be severely delayed in accessing it because the hardware, software and expertise required to read it are no longer readily available.

That is a useful analogy for threats. The adversary does not have to use the newest tool. They only need to use the tool that works.

> *They regress to progress.*
>
> Steve Wright

03 THE HUMAN DOMAIN

The human being is difficult to secure for a simple reason: people are not static.

A technical system can be configured. A person changes. Their circumstances, relationships, priorities, perception of an organisation, and tolerance for pressure can change.

That variability creates an enormous range of possible combinations. A person who is trusted, knowledgeable and well connected may be a highly valuable employee. Add financial pressure, a grievance, a new relationship, fatigue or a new opportunity, and the risk profile can shift even if the underlying access remains unchanged.

The number of possible combinations is difficult to grasp. A deck of 52 playing cards has an extraordinary number of arrangements. Chess creates an even larger decision space because each move changes what is possible next. Human behaviour is even more complex, as there is no fixed board and no requirement for people to follow a fixed sequence.

The point is not to calculate a precise number of behavioural combinations. It is to recognise that an organisation expecting people to behave in a small number of predictable ways is building its security model on an assumption that reality does not support.

WHY UNPREDICTABILITY MATTERS

Humans are influenced by trust, pressure, loyalty, curiosity, fatigue, financial incentives, ideology and personal grievance. They are capable of discretion as well as error. They can be persuaded, manipulated, coerced or simply convinced that what they are doing is justified.

None of these factors needs to exist in isolation. Risk can arise from their interaction.

That poses a fundamental challenge to traditional security thinking.

- Technical controls can ask whether an action is authorised.
- Human risk requires a second question: why is an authorised person taking that action, and what has changed around them?

> *They are not a flaw in the system. They are the system.*
>
> Steve Wright

Takeaway: Human risk is not static. Access may remain unchanged, but shifts in circumstance, motivation and relationships can fundamentally alter how that access might be used.

04 THE INSIDER THREAT IS ALWAYS THERE

As organisations strengthen their external defences, adversaries may increasingly seek to exploit people with legitimate access. However, that is only one pathway through which insider threats arise.

Insider threat sits alongside every other threat because every organisation must rely on people. Networks need administrators. Information needs owners. Customers need service. Exceptions need to be managed. Problems need to be solved.

Every one of those requirements creates legitimate human access.

The human route therefore does not disappear when technical controls improve. In some respects, its value becomes clearer. If an external route becomes expensive, noisy or unreliable, an internal route may remain comparatively straightforward.

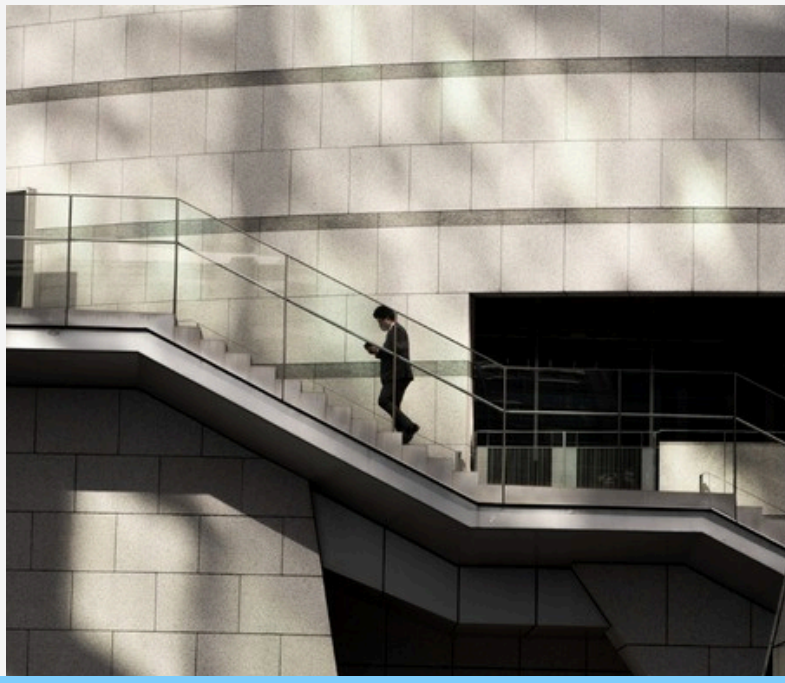
This is not because insiders are inherently malicious. The overwhelming majority are not. It is because the organisation cannot remove the human element without also removing much of its ability to function.

THE CHEAP AND SIMPLE FALLBACK

An adversary does not necessarily need to defeat a sophisticated defensive architecture if the same objective can be achieved through a person. They may seek information in conversation, exploit a trusted relationship, persuade someone to make an exception, influence a decision, or take advantage of a moment when normal judgement is impaired.

The method may be simple. The consequences may not be.

This creates an uncomfortable asymmetry. Organisations can spend significant amounts of money to make systems harder to penetrate, while the human pathway may require only time, patience and an understanding of people.



Takeaway: The insider threat is not waiting at the end of the security model. It is present beside the model because organisations must continue to operate through people.

05 WHY ORGANISATIONS DON'T SEE THE HUMAN ATTACK SURFACE

Organisations are drawn to tangible controls. A technical control can be purchased, implemented, configured and measured. A firewall can report attempted connections. An endpoint platform can report detections. An identity system can show privileged accounts and access events.

These measures produce what leaders naturally value: evidence of progress.

Human risk is different. The absence of an incident does not prove that the organisation has no exposure. It may simply mean that no one has attempted to exploit it, that any attempt has not yet succeeded, or that the organisation would not recognise the pathway if it were to occur.

That can make an insider threat difficult to justify as an investment. A programme may require training, investigative capability, behavioural understanding, leadership attention and changes to organisational practice, yet its most important success may be that nothing happens.

The response can become familiar: We trust our people. We hire well. We pay well. We have a good culture. We have policies. We already have security controls in place.

All of those things may be true. None removes the underlying issue.

THE MEASUREMENT PROBLEM

Organisations often prefer risks that can be quantified. Quantifying risks builds confidence because progress can be tracked over time.

But human behaviour does not lend itself to the same clean measurement. A reduction in technical incidents may be measurable, whereas a reduction in the likelihood that a trusted person will misuse access is much harder to demonstrate.

This can create an imbalance in investment. Technical security receives funding because its outputs are visible. Human security can be resisted because its value is harder to quantify and its implementation can feel uncomfortable.

There is also a cultural dimension. A new firewall does not feel like an accusation, whereas a new human-risk process can feel very different. Employees may interpret increased scrutiny as a lack of trust, while leaders may worry that acknowledging insider exposure signals a lack of trust in their workforce.



THE CONTROLS WE DON'T LIKE TALKING ABOUT

A mature insider-threat approach is not about assuming everyone is a potential offender. It is about recognising that legitimate access creates capability that can be misused, either intentionally or unintentionally.

The questions feel uncomfortable mainly because they focus on valuable employees:

- Who understands the workaround?
- Who can make an exception happen?
- Who is routinely called when something needs fixing quickly?
- Who can obtain information because people trust them enough to share it?
- Who has accumulated knowledge that is not written down anywhere?

Takeaway: The organisations most comfortable measuring technical risk can be the least comfortable discussing human risk, because the latter is harder to quantify and more difficult to separate from trust and culture.

06 THE PARADOX OF SECURITY

This creates a paradox.

The more secure organisations become technically, the more exposed they may be in human terms. Confidence in systems can breed complacency in behaviour. Investment in technology can outpace investment in understanding people. Controls can become so effective that they shift attention away from the very domain that cannot be controlled in the same way.

This does not mean technical investment is misguided. AI-driven defence will be essential for managing the scale, speed and complexity of modern threats. It will close gaps that humans alone cannot. It will reduce noise, enhance detection and improve response times.

But technical success should not be mistaken for risk removal. It changes the shape of risk.

WHEN THE SYSTEM IS SECURE BUT THE PERSON ISN'T

A well-known example is Edward Snowden's removal of highly classified information from the National Security Agency environment in 2013.

The significance of the example is not that a sophisticated organisation was simply defeated by a thumb drive. The more important point is that the technical environment still relied on authorised people performing authorised functions.

Snowden was working as a systems administrator for an NSA contractor. Contemporary reporting indicated that his role gave him broad access to parts of the NSA's networks and also provided legitimate reasons for him to use removable media. Investigators reported that he downloaded material from NSA systems and used a thumb drive to remove classified information from the facility.

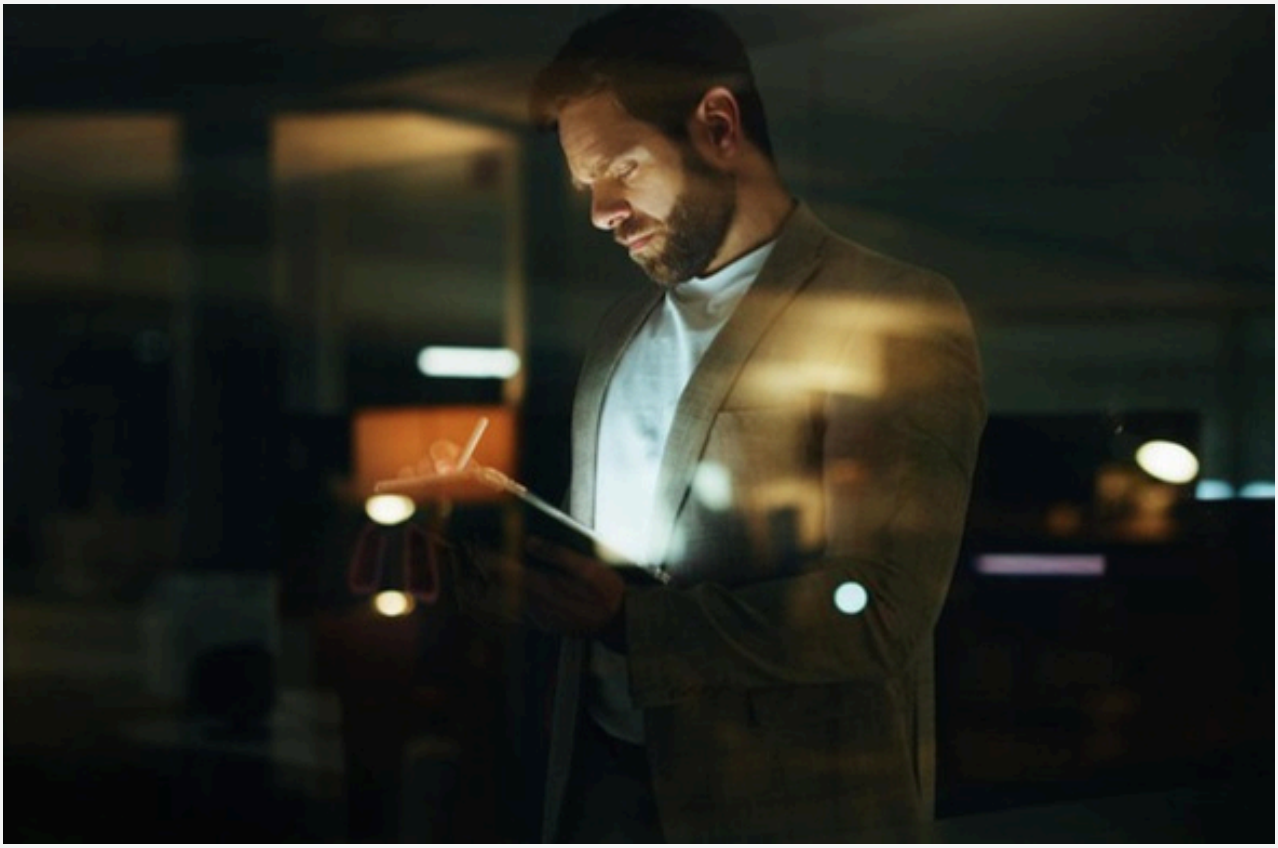
The systems could be highly secure. The networks could be tightly controlled. Removable media could be restricted. Activity could be monitored.

But somewhere within that environment was a person with legitimate access, legitimate responsibilities and a legitimate reason to interact with the systems.

A security system can be designed to recognise an unauthorised connection. It is much harder to determine whether an authorised person is using an authorised capability for an unauthorised purpose.

The distinction is subtle, but fundamental.

The system does not necessarily have to be defeated. Sometimes, the person simply has to use it.



The system does not necessarily have to be defeated. Sometimes, the person simply has to use it.

Steve Wright

Takeaway: Technical controls can constrain what people do. They cannot completely remove the need for people to have legitimate access, discretion and the ability to act.

07 HOW HUMAN CAPABILITY ACCUMULATES

The human attack surface becomes even more significant when viewed over time. A person rarely acquires consequential capability in one step. It accumulates.

Someone joins an organisation with limited access. They learn a process. Their responsibilities increase. Additional permissions are granted. They build relationships. They become trusted. They learn the exceptions that are not documented. They become the person others call when something needs to be resolved quickly.

None of these events is inherently concerning. They are usually indicators of a successful employee.

But collectively they can change what the person is capable of doing.

FROM ACCESS TO CAPABILITY

Access creates opportunity. Knowledge creates understanding. Relationships create pathways. Trust reduces friction. Experience reveals workarounds. Influence changes what others do. Opportunity brings those characteristics together.

The result can be significant capability, even if no single event appears unusual.

Organisations frequently review access as a snapshot. That can confirm that a person has the permissions appropriate to their role, but it may not reveal how much additional capability they have accumulated beyond those permissions.

A person may know who to call, which system is authoritative when two systems disagree, which manager will approve an exception, or how a process works in practice rather than as documented.

From a security perspective, that institutional knowledge can also be a capability.

> *Access can be gained. Capability can be developed. The person who seems unlikely to attract attention today may, under different circumstances, become highly significant tomorrow.*
>
> Steve Wright

Takeaway: Insider risk should be understood as a moving picture. The question is not only who has access today, but who is becoming increasingly capable over time.

08 THE FUTURE INSIDER THREAT

The future insider threat will not necessarily resemble the insider threat of the past.

As AI becomes embedded in security systems, organisations will become better at detecting technical anomalies, correlating activity and identifying patterns. That will be valuable. It will also change the environment in which adversaries operate.

The more effectively machines defend against other machines, the greater the incentive to seek out circumstances that machines cannot fully interpret.

That may mean influence rather than intrusion. It may mean access rather than exploitation. It may mean persuading a person to make a decision rather than exploiting a vulnerability in the system that records the decision.

The future insider threat may therefore be less about stealing a password and more about shaping what an authorised person does with the access they already have.

Technology does not remove human behaviour. It changes the environment in which human behaviour occurs. People will continue to experience pressure, form relationships, make mistakes, seek approval, become frustrated and make decisions.

Those are not security failures. They are human characteristics. The security challenge is understanding how those characteristics intersect with access and capability.



QUESTIONS FOR ORGANISATION

Understanding human exposure requires organisations to look beyond who has access and to consider what people can actually do with the access, knowledge, authority and influence they possess.

The following questions challenge leaders to examine whether their security model may provide confidence in technical controls while overlooking the human pathways through which harm can occur.

- **Where does our technical security end and our human exposure begin?**
- **Who has legitimate access that would be difficult for an outsider to obtain?**
- **Who understands how our processes actually work rather than merely how they are documented?**
- **Who knows where the exceptions, workarounds and informal pathways lie?**
- **Which people have developed capability over time without that capability ever being formally assessed?**
- **If a particular technical route proved too difficult to exploit, what human route would remain?**
- **Are we measuring insider exposure, or are we simply measuring the controls in place?**
- **Would our current security model recognise an authorised person using an authorised capability for an unauthorised purpose?**
- **Who could materially affect an outcome without appearing to hold significant authority?**

Takeaway: The most useful insider-threat question may not be 'Who is dangerous?' It may be 'Where does meaningful capability exist, and how could it be influenced or misused?'

CONCLUSION - THE LAST ATTACK SURFACE

This is not an argument against technological advancement - far from it.

AI-driven defence will be essential for managing the scale, speed and complexity of modern threats. It will close gaps that humans alone cannot. It will reduce noise, enhance detection and improve response.

But it will not remove the need to understand, anticipate and manage human behaviour, because that is where uncertainty remains.

For organisations, the temptation will be to continue moving towards greater automation, visibility and control. That trajectory is understandable. It is also inevitable.

The problem is that control has limits when the system being controlled includes people. A person can make a decision that was not anticipated. They can interpret a policy differently. They can form a relationship that changes what is possible. They can accumulate knowledge over years. They can be influenced by circumstances that no security dashboard captures.

The insider threat is therefore not merely the final layer of a security model. It is a permanent feature of the model. It exists because organisations need people, people need access, and access creates capability.

The future of threats will not be exclusively human. Technology will continue to shape the environment, and machines will increasingly defend against other machines. But when the technical contest reaches a point of parity, the remaining advantage will often lie where technology cannot fully control it.

It will be found in influence, access and decision-making. It will be found in moments when judgement overrides policy, when trust overrides caution, and when intent intersects with opportunity.

It will be found in people.

> *When technology reaches its limits – when it becomes too effective to exploit directly – the focus does not disappear. It narrows. What remains is the last attack surface: human.*
>
> Steve Wright

ABOUT THE AUTHOR



Steve Wright is an intelligence and organisational risk practitioner with more than 20 years' experience spanning defence, national security, counterintelligence, financial crime and operational leadership. A former Royal Navy professional, his career has included senior positions within UK Defence intelligence, specialising in counterintelligence, human intelligence and the identification, assessment and management of threats arising from individuals with legitimate access to sensitive information, systems and capabilities.

His experience encompasses both strategic and operational intelligence, including the management of complex intelligence operations, the development of intelligence capabilities and the provision of advice to senior government and defence stakeholders. Having worked across government, military and private-sector environments, he brings a practical understanding of how organisational vulnerabilities can emerge through human behaviour, informal influence, trusted relationships and the accumulation of access and capability.

Steve has a particular interest in challenging conventional approaches to insider-threat management, moving beyond a narrow focus on malicious intent to examine the wider relationship between access, capability, opportunity and human behaviour. His work seeks to translate intelligence principles and operational experience into practical approaches that help organisations better understand their human attack surface, identify overlooked vulnerabilities and strengthen their resilience against emerging and enduring threats.



THE LAST ATTACK SURFACE

Why the future of threat will always be human

By Steve Wright

Contact Us

www.insiderthreats.au

hello@insiderthreats.com.au

+61 2 6198 3381