



Where Awareness Fails

Why Cybersecurity Training Isn't Stopping Breaches

Reframing cybersecurity as a decision-making problem, not a knowledge problem



DOI: <https://doi.org/10.5281/zenodo.19360661>

Where Awareness Fails

Why cybersecurity awareness training does not consistently change behaviour in real work

Executive Summary

Cybersecurity awareness training has become a standard component of organisational defence. Most organisations invest in it, employees complete it, and baseline understanding of common threats has improved significantly.

Despite that progress, cyber incidents continue to appear in familiar patterns.

This paper examines the gap between what employees know and how they act in real working conditions. While awareness training improves recognition of threats in controlled environments, it does not consistently translate into secure behaviour during everyday tasks.

Through a series of real-world decision scenarios, this paper demonstrates that cyber risk rarely appears as something clearly malicious. Instead, it emerges through routine activities — processing payments, responding to colleagues, sharing information, or interacting with familiar systems. In these situations, actions feel legitimate and aligned with normal responsibilities.

As a result, decisions are not experienced as security decisions.

They are experienced as part of the job.

This paper shows that the issue is not a lack of awareness, but how decisions are made in context. Time pressure, familiarity, trust, and workflow demands shape behaviour more strongly than knowledge alone. When situations align with expectations, awareness is often not triggered, even when risk is present.

This creates a consistent pattern.

Individuals make reasonable decisions based on the information available to them. Those decisions are taken quickly, with the intention of maintaining progress and completing tasks. When viewed in isolation, they are justified. When viewed collectively, they create exposure.

The paper further explores why this gap persists. It is reinforced by how cybersecurity is structured within organisations — as a knowledge problem that can be addressed through training, measured through completion, and validated through compliance. These approaches provide clarity and accountability, but do not fully reflect how decisions are made in practice.

As a result, organisations may be compliant and well-trained, while still exposed to the same behavioural risks.

The cost of this misalignment builds gradually. Incidents repeat in different forms, teams experience disruption, and a false sense of security can develop as training activity is mistaken for effective risk reduction.

This paper reframes cybersecurity as a decision-making challenge rather than a knowledge problem. It highlights the need for a closer alignment between how training is designed and how work is actually performed, recognising that behaviour is shaped by context, not information alone.

Understanding this shift is critical.

Because until training reflects how decisions are made in real environments, organisations will continue to experience incidents driven not by ignorance, but by reasonable actions taken in situations that feel normal.

This paper defines the problem.

The next step — explored in *Beyond Awareness: Why Cybersecurity Training Must Become Behaviour-Led* — is to examine how organisations can begin to address it.

Introduction: Risk Doesn't Look Like Risk

A project manager reviews an invoice late in the afternoon. It comes from a supplier they've worked with for years. The format is familiar, the amount is expected, and the message explains that payment is overdue and needs to be processed the same day.

Nothing about the situation feels unusual. Processing it quickly feels like the right thing to do.

In another part of the organisation, a human resources specialist receives an email from a senior executive requesting an updated payroll list. The tone is direct, the request is time-sensitive, and the context makes sense. The information is compiled and sent without delay.

Elsewhere, a developer is asked by a colleague to share access details to help resolve an issue before a deadline. The request is informal, collaborative, and aligned with the pressure the team is under. Sharing access feels like the fastest way to keep things moving.

In each of these situations, the decision made is reasonable.

There is no obvious warning sign. No clear indication that something is wrong. The actions taken align with normal responsibilities, established relationships, and the expectation to work efficiently.

And yet, these are the kinds of moments where many cyber incidents begin.

Not in situations that clearly look like threats, but in situations that feel like legitimate work.

This is the challenge that sits at the centre of modern cybersecurity.

Because most organisations have already invested in awareness. Employees have been trained to recognise phishing emails, avoid suspicious links, and follow security policies. In isolation, this knowledge is often applied correctly.

But real work does not happen in isolation.

It happens under time pressure, across multiple systems, within ongoing conversations, and alongside competing priorities. Decisions are made quickly, often without the opportunity to step back and assess risk in a structured way.

In these conditions, awareness does not disappear. It becomes secondary.

Understanding this shift — from knowledge to behaviour — is critical. Because it reveals a gap that traditional approaches do not fully address.

A gap between what people know, and what they do in the moment it matters.

The Assumption Behind Awareness Training

Cybersecurity awareness training is built on a premise that is both logical and widely accepted.

If people understand the risks they face, they will make better decisions.

This assumption underpins most organisational approaches to human risk. Employees are taught to recognise phishing emails, identify suspicious links, and follow established security practices. Training is designed to increase knowledge, reinforce policy, and reduce the likelihood of human error.

In structured environments, this approach works.

When individuals are given time to assess a scenario, free from interruption or competing priorities, they are often able to identify risks correctly. The indicators are clearer, the context is simplified, and the decision can be made deliberately. In these conditions, awareness translates effectively into action.

For example, when shown a standalone phishing email during training, most employees can identify inconsistencies in tone, formatting, or sender details. The risk is visible, and the task is clear: assess whether the email is safe.

This is why awareness training is measurable. Completion rates are tracked, assessments are scored, and organisations can demonstrate that employees understand what they have been taught.

However, this model is built on an important assumption.

It assumes that the conditions in which knowledge is learned are similar to the conditions in which decisions are made.

In practice, this is rarely the case.

Real work does not present problems in isolation. It unfolds across conversations, systems, and tasks that are already in progress. Decisions are made while managing deadlines, responding to colleagues, and maintaining momentum. Information is incomplete, time is limited, and attention is divided.

In these moments, the question is not usually "Is this secure?"

It is something more immediate and practical.

Whether to respond now or later.

Whether to trust the request or question it.

Whether to keep work moving or introduce delay.

The presence of awareness does not disappear in these situations, but its influence changes. It becomes one factor among many, competing with urgency, familiarity, and the expectation to act.

This is where the original assumption begins to break down.

Because knowing what to look for is not the same as recognising risk in context. And recognising risk in context is not the same as acting on it under pressure.

The gap between knowledge and behaviour is not created by a lack of training.

It is created by the difference between how training is experienced and how work is actually carried out.

Until this distinction is understood, organisations will continue to rely on approaches that improve awareness, but do not consistently influence the decisions that matter most.

What Awareness Training Gets Right

Cybersecurity awareness training has played an important role in improving how organisations understand risk. Over the past decade, it has helped shift cybersecurity from a purely technical concern into something that involves the wider workforce.

Most employees today have a baseline understanding that did not exist previously. Concepts such as phishing, password security, and data protection are more widely recognised, and there is generally greater awareness that individual actions can have organisational consequences.

This shift is significant. It has reduced exposure to more obvious threats and created a shared language around cybersecurity that allows organisations to communicate expectations more clearly.

In structured scenarios, this awareness is clearly visible.

When employees are presented with a suspicious email in isolation, many are able to identify the warning signs. An unexpected request, a mismatched sender address, or a message that creates unnecessary urgency is more likely to be questioned than it would have been in the past. The task is clear, the context is limited, and the decision can be made deliberately.

This is where awareness training delivers measurable value.

It improves recognition of known threats.

It reinforces policies and expected behaviours.

It provides organisations with a consistent baseline across teams.

These outcomes are important, and they form a necessary foundation for any cybersecurity strategy.

However, the effectiveness of awareness training is closely tied to the conditions in which it is applied.

It works best when risk is clearly visible, when the individual has time to assess the situation, and when the decision can be made deliberately. In these moments, the connection between knowledge and action is strong.

The challenge is that these conditions do not reflect how most work actually happens.

In practice, employees are rarely stepping back to analyse a situation as a security problem. They are responding to emails, progressing tasks, collaborating with colleagues, and managing competing demands. Decisions are made within the flow of work, not outside of it.

This is where awareness reaches its natural boundary.

It prepares people to recognise threats when they stand out. It does not fully prepare them for situations where risk is embedded within something that appears normal.

Understanding what awareness training gets right is essential, because it highlights where it stops. And it is beyond this point that a different approach becomes necessary.

Where Awareness Breaks Down

In real working environments, cyber risk rarely presents itself in a way that clearly signals danger. It does not arrive as something obviously suspicious or out of place. Instead, it appears through situations that feel routine, expected, and aligned with the work that needs to be done.

To understand where awareness breaks down, it is necessary to look closely at how decisions are made in context.

Consider a finance team responsible for processing supplier payments.

An email arrives during a busy part of the day. It comes from a supplier the organisation has worked with before. The name is familiar, the format matches previous communication, and the message references an outstanding invoice that requires urgent payment to avoid disruption.

Nothing immediately stands out as unusual.

The request fits naturally into the flow of work. Invoices are processed regularly. Suppliers follow up on payments. Urgency is not uncommon, particularly around reporting deadlines or reconciliation periods.

The individual reviewing the email is not assessing it as a potential security threat. They are assessing it as part of their role. Their focus is on accuracy, timeliness, and maintaining the supplier relationship.

The decision is shaped by what makes sense in that moment.

The sender is recognised.

The request is expected.

The urgency creates a need to act.

There is no clear signal that something is wrong, and no obvious reason to pause. In fact, pausing introduces friction. It delays the task, adds complexity, and risks creating a problem where none appears to exist.

Processing the invoice feels like the correct decision.

A similar pattern can be seen in other parts of the organisation.

A member of the HR team receives an email from what appears to be a senior executive. The message is brief, direct, and requests an updated payroll file for an ongoing matter.

Again, nothing feels out of place.

Requests from senior staff are normal. The tone reflects authority. The urgency suggests importance. Responding quickly feels appropriate, while questioning the request may feel unnecessary or even uncomfortable.

The individual is not ignoring risk. They are interpreting the situation based on how the organisation typically operates.

They recognise the name.

They understand the request.

They respond in a way that supports the task.

In this moment, awareness is not absent. It is simply not triggered.

The same dynamic appears in collaborative environments.

A developer is working under pressure to resolve an issue affecting a project. A colleague asks for access credentials to help diagnose the problem. The request is informal, made through a familiar channel, and framed as a way to speed up resolution.

The situation aligns with how the team normally works.

Collaboration is expected. Sharing information is common. The priority is to restore functionality and keep the project moving.

Pausing to question the request introduces delay and disrupts progress. It may even feel unnecessary in a context where trust is already established.

Providing access feels like the most practical decision.

Across each of these scenarios, the individuals involved are acting in ways that are consistent with their roles, responsibilities, and environment.

Nothing in these moments feels like a clear security incident.

This is where awareness reaches its limit.

Awareness training prepares individuals to recognise threats when they are clearly distinguishable from normal activity. It helps people identify emails that look suspicious, messages that contain obvious warning signs, or requests that do not align with expectations.

However, when risk is embedded within something that appears legitimate, the decision is no longer about identifying a threat. It is about interpreting a situation.

And interpretation is shaped by context.

Time pressure encourages faster decisions.

Familiarity reduces the need for scrutiny.

Trust lowers the likelihood of verification.

Workflow prioritises completion over interruption.

These factors do not remove awareness, but they change how it is applied. The individual is no longer evaluating the situation as a potential risk. They are responding to it as part of their work.

This is why the same person who can correctly identify a phishing email during training may respond differently when faced with a similar situation in a live environment.

The difference is not knowledge.

The difference is context.

These examples are not isolated incidents. They reflect patterns that exist across organisations, roles, and industries.

Finance teams process urgent requests.

HR teams respond to senior staff.

Technical teams prioritise speed and collaboration.

In each case, the structure is the same. The details change, but the conditions remain consistent.

Over time, this creates a predictable outcome.

Decisions are made quickly because they need to be.
Requests are trusted because they appear legitimate.
Verification is skipped because nothing clearly signals a problem.

Each decision, taken in isolation, is reasonable.

But collectively, they create exposure.

This is how cyber incidents develop in practice — not through a single obvious mistake, but through a series of small decisions that make sense at the time they are made.

The Pattern Behind These Decisions

Across each of these scenarios, the details are different, but the underlying structure remains consistent.

The individuals involved are not acting randomly or carelessly. Their decisions follow a predictable pattern shaped by how work is experienced in real environments. What changes from one situation to another is not the nature of the decision, but the context in which it is made.

In each case, the situation appears familiar. The request aligns with expectations, and the action supports the task that needs to be completed. Nothing feels out of place, and there is no clear reason to stop and reassess.

Because of this, the decision is processed quickly.

There is no disruption to normal workflow, and no obvious signal that deeper scrutiny is required. The situation fits within an existing understanding of what “normal” looks like, and is therefore accepted as such.

This is how most decisions are made in everyday work.

Rather than analysing each situation in detail, individuals rely on recognition. They interpret what they are seeing based on previous experience, context, and the immediate demands of the task in front of them. This allows work to continue efficiently, without unnecessary interruption.

In most cases, this approach is not just common — it is necessary.

Without it, routine work would slow down significantly. Every request would require verification. Every action would need deliberate scrutiny. The ability to recognise patterns and respond quickly is what allows organisations to function at scale.

However, this same mechanism creates a predictable vulnerability.

When a situation is designed — or manipulated — to align with what is expected, it is far less likely to be questioned. The closer something resembles legitimate work, the more easily it is accepted.

This pattern is not limited to emails or direct requests. It extends into the systems and platforms people use every day.

For example, an employee receives a notification to log in to a familiar platform. The interface looks correct. The branding is consistent. The prompt appears as part of a normal workflow — perhaps following a session timeout or a routine access request.

There is no obvious reason to question it.

The individual enters their credentials and continues with their work.

In this situation, the decision is not perceived as a risk. It is perceived as a necessary step to proceed. The interaction is treated as part of the system, not something external that requires validation.

Again, the pattern holds.

The action fits the task.

The context feels legitimate.

The process supports progress.

This is not a failure of awareness.

It is a consequence of how decisions are made under normal conditions.

Understanding this pattern is critical, because it explains why incidents that appear different on the surface often share the same underlying structure.

The names change. The messages change. The tools and systems involved may vary.

But the decision process remains the same.

And it is this process — not just the presence or absence of knowledge — that determines how risk is recognised and acted upon.

Cybersecurity as a Decision-Making Problem

The patterns explored throughout this paper point to a fundamental shift in how cybersecurity risk should be understood.

It is not primarily a knowledge problem.

It is a decision-making problem.

This distinction directly challenges the assumption outlined earlier — that if individuals understand risk, they will act accordingly. That assumption holds in controlled conditions, where situations are simplified and decisions can be made deliberately.

However, as the previous sections have shown, real decisions are not made under those conditions.

They are made in context.

They are made while work is already in progress, while attention is divided, and while multiple priorities are competing for focus. In these moments, individuals are not stepping back to assess whether something is secure or insecure. They are responding to what appears to be a legitimate part of their role.

This is where the gap between awareness and behaviour becomes visible.

The decision to act is not driven by knowledge alone. It is shaped by interpretation — how the situation is understood in the moment.

Does this look familiar?

Does it fit with what I expect?

Does responding help me complete the task in front of me?

When the answer to these questions is yes, the decision to proceed feels justified.

This is why the examples explored earlier — processing an invoice, responding to a senior request, sharing access to resolve an issue — do not feel like security failures at the time they occur. Each action is consistent with the individual's responsibilities and the environment in which they are operating.

The issue is not that the individual does not know what to do.

It is that the situation does not present itself as one that requires that knowledge to be applied.

This is a critical distinction.

Awareness does not function as a constant filter applied to every decision. It is triggered selectively, typically when something appears clearly unusual, inconsistent, or out of place. When that trigger is present, individuals are more likely to pause, reassess, and apply what they have learned.

But when a situation aligns with normal expectations, that trigger is often absent.

The decision is then guided by other factors.

The need to maintain progress.

The expectation to respond quickly.

The assumption that familiar interactions are safe.

These are not irrational behaviours. They are the mechanisms that allow work to function efficiently. Without them, every interaction would require deliberate scrutiny, and normal operations would slow to a halt.

However, these same mechanisms create predictable exposure.

As demonstrated in the previous section, when risk is embedded within something that appears legitimate, it bypasses the conditions that awareness depends on. The individual does not recognise the situation as requiring caution, and therefore does not apply it.

This is why many cyber incidents are only recognised in hindsight.

Once the outcome is known, the indicators become easier to see. The request may appear unusual, the context may seem inconsistent, and the decision may be questioned. But this clarity is retrospective. It does not reflect how the situation was experienced at the time.

At the point of decision, the action made sense.

Understanding cybersecurity as a decision-making problem shifts the focus away from what people know, and towards how they interpret and respond to situations in real time.

It recognises that behaviour is not determined by knowledge alone, but by the conditions in which decisions are made. It acknowledges that awareness, while necessary, is inherently limited by the way human decision-making operates under pressure, familiarity, and workflow demands.

This does not make awareness irrelevant. It remains an essential foundation.

But it does mean that awareness, on its own, cannot consistently influence the decisions that matter most.

Because ultimately, cybersecurity is not defined by what people know.

It is defined by how decisions are made when knowledge, context, and pressure intersect.

Why This Gap Exists

The gap between awareness and behaviour is not accidental, and it is not the result of a lack of effort or investment.

Most organisations are already doing what would be expected. They provide training, communicate policies, and reinforce the importance of cybersecurity across the workforce. In many cases, employees demonstrate a clear understanding of the risks they face.

And yet, the same types of incidents continue to occur.

As explored in the previous sections, this is not because people are unaware of what to do. It is because the conditions in which decisions are made differ significantly from the conditions in which awareness is learned and assessed.

To understand why this gap persists, it is necessary to look beyond individual behaviour and examine the structures, assumptions, and measures that shape how cybersecurity is approached within organisations.

Several factors contribute to this disconnect.

Cybersecurity Is Treated as a Knowledge Problem

At the core of most organisational approaches to cybersecurity is a simple and widely accepted assumption: if individuals understand risk, they will act accordingly.

This assumption has shaped how training is designed, delivered, and evaluated. The focus is placed on increasing awareness, improving recognition of threats, and ensuring that employees understand what is expected of them. Over time, this has led to structured programmes that are consistent, scalable, and relatively easy to measure.

From an organisational perspective, this approach provides clarity.

Training can be delivered across teams and departments. Completion can be tracked. Assessments can confirm that individuals are able to identify known threats and understand key policies. These outcomes create a tangible sense of progress, reinforcing the belief that risk is being actively managed.

However, this model is built on the assumption that knowledge will translate directly into behaviour.

As explored earlier, this assumption does not fully reflect how decisions are made in practice. Understanding a threat in isolation is not the same as recognising it when it appears within a familiar and time-sensitive situation. Knowing what to do does not guarantee that it will be applied when the context does not clearly signal risk.

This creates a subtle but important limitation.

The organisation measures what people know, but risk is determined by what people do.

When this distinction is not recognised, the effectiveness of awareness training is often overestimated. The presence of knowledge is taken as evidence of reduced risk, even when the conditions required to apply that knowledge are not consistently present.

Over time, this reinforces a model that continues to prioritise information, while the underlying decision-making dynamics remain unchanged.

Security Is Separated from How Work Actually Happens

Cybersecurity policies, controls, and training are typically designed in structured environments. Risks are identified, scenarios are defined, and expected responses are established in a way that is clear and repeatable.

This structure is necessary. It allows organisations to standardise expectations and communicate them consistently across the workforce.

However, real work does not follow the same structure.

It unfolds across multiple systems, conversations, and tasks that are already in progress. Employees are managing deadlines, responding to requests, and balancing competing priorities. Decisions are made within this flow, often quickly and with limited opportunity to step back and reassess.

This creates a disconnect.

Security is defined in one context, while decisions are made in another.

When these two are not aligned, the responsibility for managing risk is effectively placed on the individual. They are expected to recognise when a situation requires caution, pause their workflow, and apply what they have learned.

As the earlier scenarios demonstrated, these moments do not always present themselves clearly.

A request that fits expectations does not signal the need to stop. A familiar interaction does not feel like a risk. The decision is interpreted as part of the task, not as a security judgement.

The result is not a failure of the individual, but a misalignment between how security is designed and how work is experienced.

Until these two perspectives are brought closer together, the gap between awareness and behaviour will continue to persist.

Success Is Measured Through What Is Easy to Track

Organisations rely on measurement to understand whether their efforts are effective. In cybersecurity, this often means tracking indicators that are straightforward to quantify.

Training completion rates provide visibility into participation.

Assessment scores demonstrate understanding.

Policy acknowledgements confirm that expectations have been communicated.

These metrics are useful, but they have limitations.

They measure activity and knowledge, not behaviour.

An employee can complete training, achieve a high score, and still make a decision that introduces risk in a real-world situation. This does not indicate a failure of the individual. It reflects the difference between understanding information and applying it under pressure, within context, and alongside competing priorities.

When success is defined by what is easy to track, it can create a misleading picture.

The organisation sees evidence that training has been delivered and understood. This reinforces the assumption that risk has been reduced. However, the conditions that shape decision-making remain unchanged, and therefore the likelihood of similar incidents persists.

This creates a form of false confidence.

Effort is visible. Progress appears measurable. But the underlying drivers of behaviour are not being addressed.

Over time, this can lead to repeated incidents that seem disconnected, but are in fact produced by the same conditions.

The Problem Is Only Visible in Patterns

One of the most challenging aspects of this gap is that it does not present itself clearly at an individual level.

Each decision, when viewed in isolation, appears reasonable. The context makes sense. The action aligns with expectations. There is no obvious indication that a mistake has been made.

This is why incidents are often difficult to identify as systemic issues.

They appear as isolated events, attributed to individual error or unusual circumstances. The focus is placed on what happened in that specific moment, rather than how similar decisions are made across the organisation.

However, when these incidents are viewed collectively, a different picture emerges.

Patterns begin to form.

The same types of requests are acted upon.

The same conditions are present.

The same decision-making process is repeated.

It becomes clear that the issue is not the individual decision, but the structure that produces it.

This makes the problem inherently difficult to address through traditional approaches.

If the focus remains on isolated incidents, the underlying pattern is missed. If the pattern is not recognised, the response remains reactive, addressing outcomes rather than causes.

Understanding this shift — from isolated events to consistent patterns — is essential. Because it is only at this level that the gap between awareness and behaviour becomes fully visible.

Cost and Practicality Shape the Approach

Cybersecurity awareness training is often shaped not only by what is effective, but by what is practical to deliver within the constraints organisations operate under.

Time, budget, and operational impact all influence how training is designed and implemented. Organisations need solutions that can be delivered across large groups, completed within limited timeframes, and integrated into already busy working schedules. Training that requires significant time away from day-to-day responsibilities, or that disrupts workflow, can be difficult to justify.

As a result, approaches that are structured, repeatable, and efficient are prioritised.

Awareness-based training fits these requirements well. It can be standardised across teams, delivered consistently regardless of role or department, and completed in a relatively short period of time. From an operational perspective, this makes it a reliable and scalable solution.

This practicality is part of its strength.

However, it also introduces a limitation.

As explored earlier, the decisions that lead to risk are shaped by context — by how work is actually experienced in real time. Addressing this requires training that reflects real workflows, explores decision-making under pressure, and engages individuals in a way that goes beyond information delivery.

This type of training is inherently more complex.

It often requires more time, deeper engagement, and a closer alignment with specific roles and environments. It is less easily standardised, and therefore less straightforward to scale.

This creates a natural tension between practicality and effectiveness.

The approaches that are easiest to implement and measure are not always the ones that most closely address how risk develops in practice. Over time, organisations tend to favour solutions that are efficient and manageable, even when they do not fully engage with the behavioural dynamics that have been outlined in earlier sections.

The result is not a lack of training, but a mismatch between what is delivered and what is required to influence real-world decision-making.

Compliance Requirements Reinforce the Model

In many sectors, cybersecurity training is closely tied to regulatory and compliance obligations.

Organisations are required to demonstrate that employees have received training, understand key risks, and are aware of their responsibilities. This creates a need for clear, auditable evidence — something that can be recorded, reported, and presented if required.

Awareness training aligns well with these requirements.

Completion can be tracked across the organisation.

Assessment scores provide measurable outcomes.

Policy acknowledgements confirm that expectations have been communicated.

These outputs are clear, consistent, and defensible.

From a compliance perspective, this is essential. Organisations must be able to show that they have taken reasonable steps to educate their workforce and reduce exposure to risk.

However, compliance frameworks tend to focus on evidence of activity, rather than the effectiveness of behaviour in real situations.

As long as training has been delivered and recorded, the organisation is able to demonstrate that it has met its obligations. Whether that training meaningfully changes how decisions are made under pressure, within context, and during everyday work is often outside the scope of what is assessed.

This distinction is important.

It means that an organisation can be fully compliant, while still being exposed to the same behavioural risks outlined earlier in this paper.

Over time, this reinforces the existing model.

The objective becomes ensuring that training is completed and evidenced, rather than examining whether it aligns with how decisions are actually made in practice. This creates a stable, defensible approach, but one that does not fully address the underlying conditions that drive incidents.

Organisational Change Moves Slowly

Even when the limitations of current approaches are recognised, changing how cybersecurity is addressed within an organisation is not straightforward.

Training programmes, compliance processes, and reporting structures are often built over time, becoming embedded within the organisation's operating model. These systems provide consistency, support governance, and align with regulatory expectations.

As a result, they are not easily replaced.

Adopting a different approach to cybersecurity — particularly one that shifts focus from knowledge to behaviour — often requires changes at multiple levels. It may involve redesigning training, rethinking how success is measured, and adjusting how responsibility for risk is distributed across teams.

This introduces complexity.

New approaches may be less familiar, harder to quantify, and more difficult to integrate into existing structures. They may challenge established assumptions about what training should look like and how effectiveness should be assessed.

In this context, organisations tend to favour continuity.

Even when current approaches are known to have limitations, they are often maintained because they are understood, accepted, and aligned with existing systems. Change requires not only recognising a problem, but also having the confidence and clarity to adopt a different model.

This is why the gap persists.

It is not due to a lack of awareness or intent. It reflects the reality that organisational change is gradual, and that established approaches are often sustained until there is a clear and compelling reason to move beyond them.

Taken together, these factors explain why the gap between awareness and behaviour continues to persist.

It is not the result of a single failure, or a lack of effort within organisations. It is the outcome of how cybersecurity has been structured, measured, and integrated into everyday work.

Training is designed around knowledge. Success is measured through completion. Risk is managed within systems that operate separately from how decisions are actually made.

Within this model, the presence of awareness is taken as evidence of protection. Yet, as the earlier sections have shown, awareness alone does not consistently influence behaviour in real-world conditions.

This is why the same types of incidents continue to occur, even in organisations where training is well established and widely understood.

The issue is not that people are unaware of risk.

It is that the conditions in which decisions are made do not reliably support the application of that awareness.

Understanding this distinction is critical.

Because it shifts the focus away from increasing knowledge, and towards addressing how decisions are shaped in practice.

And it is this shift that defines what effective cybersecurity training needs to become.

The Cost of Misalignment

When there is a disconnect between how cybersecurity is taught and how decisions are made in practice, the impact is rarely immediate or obvious.

From an organisational perspective, the right elements often appear to be in place. Training has been delivered, employees have completed it, and there is a clear record of awareness across the workforce. Policies exist, expectations are communicated, and compliance requirements are met.

On the surface, this suggests that risk is being managed.

However, as the earlier sections have shown, the conditions that shape real-world decision-making remain unchanged. Employees continue to operate within environments defined by time pressure, familiarity, and the need to maintain workflow. Decisions are made quickly, often in situations that feel routine and legitimate.

It is within this gap that misalignment begins to have an effect.

The cost does not appear in a single moment or a single failure. It develops gradually, through repeated decisions that make sense in isolation but introduce risk when viewed collectively. Each action feels justified at the time, which means the underlying pattern often goes unnoticed.

For example, a finance team may process a payment request that appears consistent with previous supplier communication. The amount is plausible, the context fits, and the urgency feels reasonable. The transaction is completed without issue being identified at the time.

In a separate instance, an employee may respond to a request for information from what appears to be a senior colleague. The tone, timing, and context all align with normal expectations, and the response is sent quickly to support the task.

In another case, a user may enter their credentials into what appears to be a legitimate system login page after receiving a routine-looking prompt. The interface is familiar, the request is expected, and the action allows them to continue their work without interruption.

Individually, each of these actions is understandable.

There is no clear moment where the individual perceives that they are making a security decision. The action is interpreted as part of normal work, and therefore completed in the same way as any other task.

Over time, this creates a cycle that is difficult to break.

Incidents occur and are investigated.

Training is revisited or reinforced.

Awareness is increased.

Yet the same types of situations continue to arise, because the conditions that shape those decisions have not fundamentally changed. The outcome may differ, but the structure remains the same.

This leads to repeated exposure that is not always recognised as systemic.

There are, of course, direct costs associated with this misalignment. Financial loss, operational disruption, and the time required to respond to and recover from incidents are often visible and measurable. These impacts are typically what draw attention to the issue in the first place.

However, the indirect costs are less visible, and often more persistent.

Time is spent investigating incidents that follow familiar patterns, even when the individuals involved have completed training and demonstrated understanding. Teams experience repeated disruption as similar issues reoccur, often in slightly different forms. Confidence in processes can begin to erode, particularly when problems arise despite the presence of awareness programmes.

In regulated environments, this misalignment can also surface in more subtle ways.

An organisation may be able to demonstrate that training has been delivered, that employees understand key risks, and that policies are in place. However, repeated incidents — particularly those that follow familiar patterns — can raise questions about how effectively these measures are translating into behaviour.

The organisation remains compliant, but not necessarily aligned.

Perhaps the most significant impact is the sense of reassurance that this model can create.

When training is completed, measured, and reported, it provides a clear signal that action has been taken. This can reduce the perceived need to question whether the approach itself is sufficient. The presence of activity is interpreted as progress, even when the underlying decision-making conditions remain unchanged.

This is what allows the gap to persist.

The cost of misalignment is not always visible in isolation. It is the accumulation of small, reasonable decisions, repeated across different roles and situations, shaped by environments that do not fully support the application of awareness in practice.

Addressing this does not require more information.

It requires a closer alignment between how cybersecurity is understood, how training is delivered, and how decisions are actually made within the flow of work.

Conclusion

Cybersecurity awareness has not failed. It has done exactly what it was designed to do.

It has improved understanding, increased recognition of common threats, and established a shared foundation across organisations. Employees today are more aware of cyber risk than they have ever been, and this shift has played an important role in reducing exposure to more obvious threats.

However, awareness alone does not determine behaviour.

As this paper has explored, decisions are not made in isolation. They are made within the flow of work — shaped by time pressure, familiarity, trust, and the need to maintain progress. In these conditions, risk does not present itself clearly. It appears through situations that feel legitimate, expected, and aligned with everyday responsibilities.

This is where awareness reaches its limit.

The scenarios examined throughout this paper show that individuals are not making careless or uninformed decisions. They are making reasonable decisions based on the context available to them at the time. The issue is not a lack of knowledge, but a gap between how risk is taught and how it is experienced in practice.

That gap is not incidental.

It is reinforced by the way cybersecurity is structured, measured, and embedded within organisations. Training is designed around knowledge. Success is measured through completion and understanding. Risk is managed through systems that operate separately from the environments in which decisions are actually made.

Within this model, awareness becomes something that exists alongside work, rather than within it.

As a result, the same patterns continue to emerge.

Incidents occur not because people do not understand risk, but because the conditions required to apply that understanding are not consistently present. Decisions are shaped by what feels normal in the moment, not by what is known in theory.

This is why increasing awareness alone does not resolve the problem.

Addressing this challenge requires a shift in perspective.

Cybersecurity must be understood as a behavioural and decision-making challenge — one that exists within everyday workflows, not outside of them. It requires a closer alignment between how training is designed and how work is actually performed. It requires an approach that reflects how decisions are made in real time, under real conditions.

This paper has focused on identifying and explaining that gap.

The next step is to consider what it means to address it.

This is explored further in *Beyond Awareness: Why Cybersecurity Training Must Become Behaviour-Led*, which examines how training can move beyond information delivery and begin to influence how decisions are made in practice.

Because ultimately, cybersecurity is not defined by what people know.

It is defined by how decisions are made when knowledge, context, and pressure intersect — particularly in the moments that feel routine, expected, and unremarkable.