OutThink

Adaptive Security Awareness Training Playbook

Drive Engagement and Enable Secure Behaviors

July 2024

v1.4

1. What is Adaptive Security Awareness Training (SAT)

The goal of this paper is to help security teams develop an **Adaptive SAT program** tailored for the modern, connected workplace. A program designed to address the growing cybersecurity threats posed by technology usage among employees. The most advanced security teams in the world incorporate adaptive learning techniques to cater to varying roles, levels of user knowledge and digital skills, ensuring that training is engaging, effective and relevant.

Adaptive SAT delivers **tailored learning experiences** that meet individual needs through personalized pathways and content. This approach emulates the one-on-one adaptability of tutors but scales to many students using technology. There are two main types: designed adaptivity, where educators set rules for responses, and algorithmic adaptivity, where algorithms determine the next steps based on user data and direct learner interactions. Adaptive learning technology enhances teaching, personalizes learning, and supports better **learner outcomes.**

The ability to deliver Adaptive SAT to drive engagement and enable secure behaviors is a cornerstone of Cybersecurity Human Risk Management, as per the recommendations and guidance developed by the **CHRM Forum**.

2. The Case for Adaptive SAT

Every company has, or should have, policies around acceptable use of IT assets and overall information security. It is the responsibility of the CISO to keep these policies up to date, aligned with the needs of the business, risk appetite and the realities of current threats. Typically, these verbose policies are put on an intranet wiki or in a shared folder and applied in the design and implementation of security controls. While this sounds straightforward, in reality security policies and controls apply differently to different stakeholders in the organization. Also, when it comes to users' ability identify phishing emails, report suspicious events and respond to **cyber-attacks,** such behaviors can't even be codified into policy. The translation layer between policy and desired behavior is often vague or missing. It mostly depends on context, interpretation, gut feeling and a bit of luck.

Interpretation of desired secure behaviors is complicated by the fact that humans have different perspectives when it comes to security controls, **threats and risk.** People are hard to read and predict. Some have a high commitment to being secure, some don't. Some have a level of confidence about behaving securely, others experience learned helplessness **("I'll never be able to create and remember complicated passwords!").** Some know how to read a URL, and some have no clue about the difference between a subdomain, domain and a page hanging off the root. Not to mention

specificities of various functional roles, specific tasks and associated risks, some of which have an outsize impact on the company's overall **security posture.**

The most advanced deployments of Adaptive SAT aim to bridge that gap between policy and each user in their current role, level of knowledge and digital skills. Training is honed to bring these policies out of the dusty digital binder to life in a comprehensible and interactive fashion, in the way that is best understood by each learner. Using this approach, **leading security** teams move their organizations from check-the-box compliance to systematic injection of knowledge and secure behaviors into company culture.

> **The best way to teach is the way that is most understood.**

3. The Rhetoric of Behavior Change

An informal survey of **300+** attendees of Infosecurity Europe in London and the Gartner Security & Risk Summit in Washington DC, confirmed that for those companies interested in getting beyond simple compliance the primary reason for running training is to shift **user behavior** toward greater security. To truly reach the intended audience requires significant marketing and rhetorical acumen, bordering on behavioral science.

The ancient Greek terms for rhetorical strategy that are most frequently cited are Logos, Ethos and Pathos, which stand for logic/data, character, and experience/empathy, respectively. The fourth element of rhetoric often left out in primary education is Kairos, which stands for **"right time" or "opportunity".** Every learner is on an individual journey towards understanding cyber risk and expected secure behaviors. Catching them in context, at the right time in that journey is a critical part of getting the information through and making it stick.

4. The Science of Behavior Change

The reader will see us use the phrase “below the waterline” in this paper. By this phrase we are referring to the applied psychology of each individual learner. What is happening inside their head that may not be visible to the casual observer, let alone the CISO or the company’s security team. This is the realm of behavioral science. Behavioral science offers some useful lessons in constructing an adaptive framework to influence behavior change. The key concepts relevant to this topic are Behavioral Segmentation, Concordance, Nudge Theory, and Self Efficacy. These elements are all summarized neatly in the HPE Awareness Maturity Curve, which is referenced in several academic papers on the topic of secure behavior.[1]

Steps to transforming security behavior (based on HPE Awareness Maturity Curve)

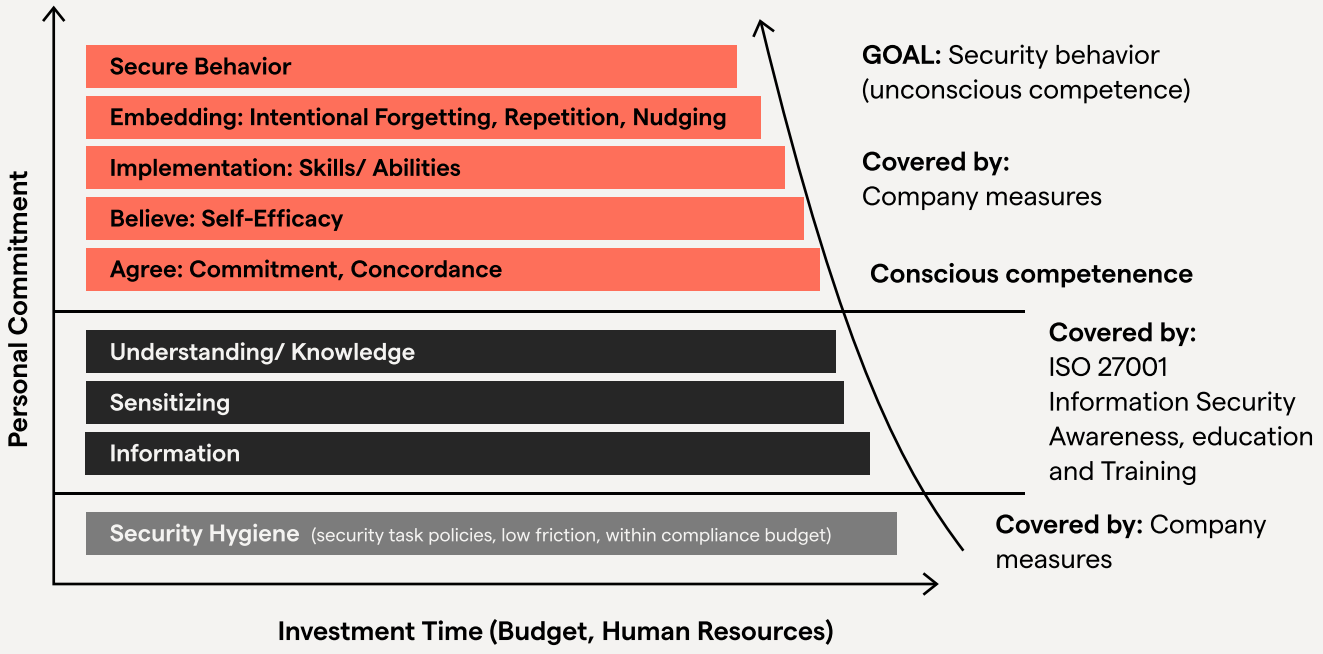

[1] Jonas Hielscher, Annette Kluge, Uta Menges, and M. Angela Sasse. 2021. “Taking out the Trash”: Why Security Behavior Change requires Intentional Forgetting. In New Security Paradigms Workshop (NSPW '21), October 25–28, 2021, Virtual Event, USA. ACM, New York, NY, USA 15 Pages. https://doi.org/10.1145/3498891.3498902

Nudge theory, developed by Richard Thaler, stems from **behavioral economics** and focuses on subtle prompts to influence behavior positively without restricting choices. Unlike traditional methods that often use fear or coercion, nudges guide people gently toward safer decisions.

Concordance is the practice of setting concrete goals, with a way to track progress towards those goals. For example, setting strong passwords and using a system to track these. The goal might be to set such passwords in one system at a time. The process then includes checking in with a coach to show progress towards said goals. This can be done with live coaches, line management or a virtual assistant through a security awareness platform. The idea is by encouraging the r**epetition of micro behaviors,** these become learned and tacit behaviors embedded in daily practice. Thereby reaching the stage of “unconscious competence”.

The concept of choice architecture is also useful in this context, configuring environments to influence decisions effectively. Which option comes on top? Which button is highlighted? Marketers use these concepts to drive traffic on the company website towards **“converting”** into sales. Sophisticated security teams use the same concepts to quietly nudge users towards more secure behaviors.

Adaptive SAT is timely, context-aware, and leverages cognitive biases, with behavioral segmentation offering a basis for application. By integrating these principles, organizations can create impactful, behavior-driven cybersecurity learning experiences that are engaging and effective in changing behavior.

5. The Playbook – Putting It All Together

That's a lot of theory, but how does one actually implement this in practice? The toughest part is taking the **first steps** towards bridging that chasm.

This section is a quick guide to the key elements of Adaptive SAT implementation, based on the small number of world-leading, Fortune **500 security teams** that have crossed that chasm and are reaping the benefits.

In order to get tailored and specific, the adaptive training program is built on a **solid footing** of data about each individual in the company. This doesn't have to cross the **"creepy line"** into monitoring, the personal details of their employment records or household drama – just the relevant human risk factors – level of access, role, attitudes, behaviors and risk understanding, as measured through interactions in the workplace. The way to put this into practice requires an understanding of behavioral science, combined with security system integrations, orchestration and data science.

> **In theory, theory and practice are the same. But in practice, they are not.**
>
> -Not **Yogi Berra**
> commonly misattributed

5.1 Foundational Data

Unravelling the data puzzle involves examining an individual's everyday actions. This includes their responses to phishing attempts or simulations, their diligence in installing necessary updates or lack thereof, their role in triggering malware incidents on the endpoint, or frequency of accessing sensitive data. All this information, linked to their user ID, initiates the construction of each user's human risk portrait. However, this is merely scratching the surface, akin to **analyzing the visible tip** of an iceberg.

The true essence of an individual lies below the waterline, submerged like the majority of the iceberg. Their attitudes, intentions, knowledge and risk perception are the vital human risk factors that drive poor security behaviors.

The section titled **'Human Risk Segmentation'** probes deeper into comprehending each individual's human risk posture. It explores beneath the waterline, shedding light on their applied psychographic segment, which is integral to any human risk assessment undertaking. A sample of this kind of human risk snapshot can be seen in the spider chart on the right.

All of this data is further combined with the **user's risk exposure and trust** or collaboration network – both the level to which their environment puts them in harm's way and the "blast radius" they might inflict by making a mistake.

When overseeing security for thousands of users in an organization, it is impossible to perform this analysis, on a continuous basis, without automation.

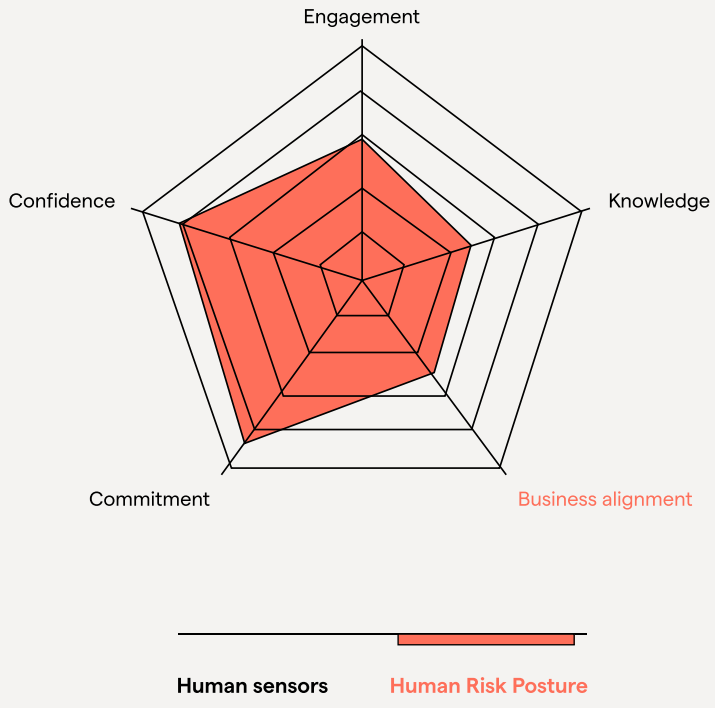

Once automated, machine learning algorithms can help **personalize training content** based on the individual user's risk portrait.

The adaptive training playbook inherently recommends a continuous feedback loop. This not only facilitates the evaluation of user progress and the effectiveness of the training material but also verifies if the content is genuinely resonating with the learners. Simultaneously, this feedback mechanism empowers the security practitioner to activate users as **a network of human sensors.** This process enables learning about the users and from them, identifying the risks and challenges they might face. Ultimately, this feedback also assists in refining the selection models for future training content.

5.2 Delivery Channel

In order for the most relevant content to be presented to the user at the right place and the right time, the training will frequently have to be delivered outside the **corporate Learning Management System (LMS)**. In practice it has proven useful to deliver nudges and mini training segments in between LMS campaigns, through a desktop agent, email, Slack or Teams virtual assistant, depending on context. For blue collar workers without an email address or a computer, the training content needs to be delivered via a shared device or terminal.

The nudge might be delivered in the email client, when trying to send an unencrypted file appearing to contain sensitive data. It might be in the browser when visiting sites outside the acceptable use policy. Reactions to these nudges should be stored, tracked and added to the user's risk dataset.

When engaging in a phishing simulation, the follow-up training should be **immediate, short and pertinent** to the learner's specific characteristics. It's punishing and condescending to force someone through the same generic content about phishing just because they clicked on a link. There could just be one aspect of phishing they haven't understood. **A 3-minute snippet** about the right aspect of phishing awareness based on automated root cause analysis (e.g., deception technique used, failure in interpreting a URL or analyzing suspicious sender information) often really hits the spot. It's still more training, but it's reinforcement not punishment. Being **tailored** to the learner, it reduces the wasted time and prevents mounting aggravation. Users might even learn something they would have otherwise been embarrassed to ask their colleagues!

> **True localization means the names and contact details of local infosec staff, links to department policy docs, and the tools provided for high-risk processes.**

5.3 Organization-specific Content

For the learning experience to be as relevant as possible, **it needs to be specific to the organization of the learner.** This is not simply about branding and styling the training content to make it look & feel on brand, although that is important too.

To make the training appropriate and an expression of the organization's security policy, you will want to avoid generic, off-the-shelf content. There are a number of items that will make the training **feel contextually relevant,** a few examples include:

- The number you want employees to call or the email address where they should write if they want to report a security incident or any suspicious activity
- The steps to take to report a phishing attack.
- The specific types of data deemed sensitive / confidential and the organization's internal information classification levels
- The specific password policy of the organization
- The approved platforms, applications or tools that teams can use to transfer large files
- The names and contact data of the key security personnel for their department or region
- Links to the policy documents on an intranet site or wiki relevant to their team
- The steps to take to request a security policy exception.

Simply put, the training content should be based on the **policies that pertain to each learner.** It makes no sense to train about badges and keycards if the learner works remotely. The more generic and "tone deaf" the training content, the more learners will turn away from it, and the security team overall.

5.4 Interactivity

The most rigorous implementations of adaptive training have a continuous component of interaction throughout the learning experience. Each user interaction serves as a datapoint for the security team to refine that user's human risk portrait and serve up more tailored, relevant content in the next touchpoint. As the learner trains, different scenarios might unfold, **making it truly a unique experience to each individual.**

When taken to its natural conclusion, the most adaptive SAT experience is one that takes in-flight responses from the learner to serve up the next scenario. Leveraging conversational AI you can allow a learning experience to go in whatever direction the user takes it. The use of conversational AI in security awareness training is an important new trend that should be explored. It's hard to get more engaging than having a conversation with an online security coach equipped with several relevant roleplay scenarios and that is able to understand the learner's job function, state of mind and level of knowledge.

Learners who demonstrate higher knowledge can be afforded the respect they deserve and served more advanced material. Security beginners get introductory material that would have turned off the high-knowledge learner right away. The more interactive the content, the more likely the learner to stay engaged. And of course, each interaction round trip is another datapoint for that user's risk dataset.

5.5 Room for Feedback

This might feel like scope creep, but in practice security teams that have gone **"adaptive" expand that "adaptive" thinking beyond just teaching.** With every training touchpoint comes an opportunity to collect unstructured comments from users. In an organization of thousands, this is the only scalable way for security teams to detect points of friction or other security challenges in their environment. When taken at scale, natural language processing (NLP) has been used surprisingly well to summarize the freeform information provided by thousands of users into a quickly relatable coherent picture.

> **Comments from thousands of learners in the context of adaptive training can be summarized by NLP tools to get a quick mapping where controls cause friction with business processes.**

Security teams that have deployed their awareness program with a feedback component have uncovered remote locations with poor physical security, business processes that could **benefit from a point solution** to remove security friction, departments with laggards still writing passwords on sticky notes or poor adherence to the organization's clear desk policy.

The security team is usually vastly outnumbered by the organization around them. The collective security knowledge lurking inside the enormity of the organization can often outweigh that of the security team and should be tapped.

This 'listening' approach, combined with a psychographic segmentation that identifies likely **Security Champions across theorganization** can further boost security efforts. Once recruited, Champions spread the word and reinforce the security policies and training within their organizations. In one example, a security champion placed sticky notes on colleagues' laptops whenever they left their laptops unattended and unlocked.

5.6 Getting Started

We recommend starting down **the Adaptive SAT journey** with an initial data feed that comes from the identity management system, like Active Directory or Okta, combined with learner self-declared information. The self-declaration can later confirm the intel coming from IT and security systems - any disparities will inform the user's risk dataset. The self-declaration can also enrich the basic data coming from Active Directory with information like the type of work people do, the nature of their role, threats and vulnerabilities they see in their part of the organization or the frequency with which they handle sensitive data.

This information is normally enough to dynamically allocate adaptive training in a targeted enough manner for the content to become immediately relevant to the learner. This way from the get-go the awareness training will be tailored to their role, knowledge level and directionally correct for their risk understanding.

6. Human Risk Segmentation

Humans are obviously not all alike. Even when specific behaviors are observed, like clicking on phishing links or credential sharing, that misses what each individual might be thinking or experiencing. Below the waterline. All users need to be viewed and treated as unique individuals. At scale, as a starting point, that means they should be viewed through the lens of human risk segments. Some effective segmentation dimensions have proven to be:

1. Level of knowledge – have they demonstrated a high level or low level of knowledge, have they shown understanding of specific secure behaviors. Have they ever been trained on the organization's security policies or are they completely new to the system.

2. Level of access – privileged access, do they have access to highly confidential information or handle sensitive data, such as PII

3. Role in the organization – are they in a role requiring specialized knowledge, e.g., system administration, domain controllers, engineering, finance, HR or vendor management. What are the unique tasks they carry out on a daily basis and the associated risks.

4. Workplace exposure – are they in an office with good physical security, or are they always working from shared spaces like cafes, their home or co-working offices. Do they travel frequently, having to work from public places such as airports. Do they have a VPN and is it fast, easy to set up, configure and use.

5. Psychographic segment – the Behavioral Segmentation Grid (BSG), developed by Dr. Angela Sasse and her team at University College London proposes a definitive list of 16 segments, ranging from Shadow Agents, Abdicators, to Rule Breakers to Champions. This is an important insight into the user's risk understanding and affective security (emotional response to security policies and recommendations) - both critical, below the waterline, human risk factors.

Champions, for instance, can be identified and recruited to help achieve positive security outcomes faster across the whole organization.

6. Attitudes and alignment – how high is their engagement, their intention to comply, are they dismissive towards security, how confident are they in ability to carry out specific security behaviors, do they perceive security controls to be in misalignment with their job, causing security policy friction.

7. Phishing resilience – are they a repeat clicker or one of those employees that are always quick to identify and report a phishing email. Which deception techniques are they most vulnerable to. Are they a detractor, that is someone deleting or just ignoring obvious phishing attacks. Or are they a defender, having demonstrated readiness to help the organization detect and respond to phishing attacks.

8. Exhibited behaviors – do they browse risky websites or use social media extensively, have they had multiple DLP violations, or perhaps they have frequent authentication failures. Do they use unapproved cloud applications or engage in shadow IT. Do they have overwhelming email volumes, or perhaps they leave confidential documents on their desk. Do they frequently cause malware events on the endpoint.

When thinking about Adaptive SAT, it's important to **look beyond just role-based training but look at all the other dimensions of human risk.** It is a desire of many organizations to deliver role-based training to their employees, and, with legacy SAT solutions, even that is not easy to do.

But as we now see, just having something role-based isn't enough to make the training timely and effective. When taken together, these eight segments of human risk can greatly increase the adaptability of the training, enabling you to deliver the right content to the right people at the right time. This results in high engagement, high knowledge retention and the highest possible level of derisking delivered to your organization.

7. Example Scenarios

It is impossible for us to **step through every single combination of human risk factors** under the eight dimensions in order to exhaustively explain how they are being used in adaptive training programs. With a rough average of 10 settings for each of the eight dimensions, that makes for 1049 permutations of scenarios. We can demonstrate the application of this model with just three such scenarios to demonstrate the logic and power of adaptive training. To go deeper would likely require a conversation with an adaptive security awareness expert.

> **Every instance of a non-secure behavior is an opportunity to adapt a micro training pointed specifically at the root cause of that behavior.**

Scenario 1 – Someone just clicked on a phishing simulation

Factors to consider for adaptive training response:

1. Are they a repeat clicker?
2. Do they have low intention to comply (commitment) or low confidence?
3. What is the root cause of why they clicked?
4. What social engineering tactic are they susceptible to(i.e., urgency, authority)?
5. What is their overall level of knowledge in relation to phishing?

Depending on the answers to these questions, the training response can **vary for different cases.** It could be the person doesn't know how to analyze a phishing link (URL). Security personnel don't realize that most people don't know the difference between a subdomain, a domain and the name of a page in the root directory of the URL. Some users might be too embarrassed to ask IT or security staff about such things.

Another person might be knowledgeable about URLs, but **gets 500 emails a day,** and doesn't have time to scrutinize anything before clicking, resulting in a low intention to comply with the 'take 5' / 'think before you click' advice. The content for these two individuals will be completely different.

If they are a repeat clicker and have low commitment or intention to comply, they might need to be shown something about the consequences of downloading malware, how that might affect the team around them and the company overall. Possibly even their own ego or livelihood.

Scenario 2 – Someone triggered an email data loss prevention (DLP) violation

Factors to consider for adaptive training response:

1. Are they a repeat offender?
2. Do they have access to highly confidential information?
3. Do they have low intention to comply or low confidence?
4. What is their overall level of knowledge in relation to phishing?
5. Did they report policy friction (low security-business alignment)?
6. Do they have access to an approved secure file sharing platform?

If the user has repeatedly sent unencrypted attachments that contain **highly confidential information** and they have low intention to comply, the adaptive training for this learner might focus on the repercussions of a data breach or recent examples of how bad actors exploited such behaviors. If the learner his high intention to comply but reported process friction in the past, their learning experience may include guidance on how to use the organization's approved secure file sharing platform (if one exists) or it might be a link to the intranet page to request a policy exception.

Scenario 3 – Someone frequently leaves their laptop unlocked while unattended

Factors to consider for adaptive training response:

1. Do they have low intention to comply or low confidence?
2. What is their risk understanding?
3. What is their knowledge level about this specific security behavior?
4. Do they work in a shared space or in a physically secure environment?
5. Do they travel frequently and are likely to leave their laptop unlocked and unattended in public places?

If they work in a highly secure office, or from home, this might be a lower risk. The likelihood of a bad actor accessing their laptop, applications and data is very low. If they are in a shared space or travelling, however, this is clearly a more material issue. If they have low intention to comply and low risk understanding, it may make sense to train them about data breaches that involved unattended devices and their consequences. If they have a high risk understanding and a high intention to comply, it may just be a matter of forgetfulness. In this case a local Security Champion can be activated to nudge the person to lock their laptop.

If all else fails, this user can either be instructed to set the timeout on their screen to one minute, or perhaps IT can force that setting on their machine.

Summary

Overall, the Adaptive SAT playbook aims to help security teams put the 'human' at the heart of security operations and create a more cyber resilient workforce. It's not just a "better approach", but a necessary evolution in the face of today's increasingly sophisticated threat actors.

We described an advanced adaptive training methodology that can be used in practice to deliver personalized, engaging, and timely training experiences that resonate with each user. By considering factors such as role, access, knowledge, attitudes, and exhibited behaviors, this approach provides a tailored learning journey for every learner. It bridges the gap between policy and practice, and is significantly more likely to succeed in equipping employees with the digital skills modern organizations require. Taking a tailored approach to security awareness is key to mitigating human risk – the risks associated with tech usage in the workplace and in personal settings as well.

The road to behavior change and effective human risk management begins with high levels of engagement. Engagement need not be a gimmick, like gamification or funny videos. The best way to achieve engagement is by delivering training that is organization specific, targeted and relevant to each learner, short, and interactive. This approach moves security awareness from being a missed opportunity that annoys users and wastes thousands of hours of effort in pursuit of mere check-the-box compliance. It turns the human element from a security weakness into a strength.

By advancing to state of the art Adaptive SAT, security leaders have made a beeline to true learner understanding, risk appreciation, and genuine attempts to enable secure behaviors and proactively defend their organization.

OutThink | Gartner Peer Insights™

The Highest Rated Cybersecurity Human Risk Management Platform (SaaS)

★★★★★

4.9 out of 51 reviews, As of March 2024

WHAT DO CUSTOMERS SAY ABOUT OUTTHINK?

> Truly incredible customer service and the only solution that gives accurate insights.

CISO, Retail Industry

> "It's not just better, it's different. Having used a large number of providers in this space over the last few years, OutThink stands head and shoulders above the rest on just about every criteria."

CIO, Finance Industry

Loved by 4+ Million People Worldwide

Danske Bank · Whirlpool · VINCI · London Gatwick · University of Notre Dame · Hollywood Burbank Airport · Epiroc

Pioneered
By CISOs for CISOs

100+ years industry experience
We make Security Teams Great!

Time to make a change

REQUEST DEMO ›

Customers	**FTSE100, Fortune 100**
Geographies	**Worldwide, 80+ Countries**
HQ	**London, UK**
Physical Presence	**60+ Staff, 8 Countries**

Recognised by

Gartner · ISF · Forrester · European Innovation Council · GIGAOM · UKRI Innovate UK · GCHQ

outthink.io