Report

Customer Identity Trends Report 2025

Securing customer trust in the age of AI

Contents

Customer identity in the crosshairs

Securing customer trust in the age of AI

Welcome to the 2025 edition of the Customer Identity Trends Report.

Most companies today — no matter how big or small, in business-to-consumer (B2C) and business-to-business (B2B) — compete in a truly global marketplace. Accordingly, the importance of understanding and delighting customers has never been higher.

And that importance puts customer identity and access management (CIAM) in the crosshairs.

CIAM enables your customers to *sign up* for and *sign in* to your digital properties. Consequently, it powers a highly visible and frequently used touchpoint that plays a significant role in building trust, obtaining consent, and acquiring the first- and zero-party data so vital to achieving growth and revenue objectives. For these reasons, modernizing CIAM — and customer identity, in general — is a top priority for many organizations.

But CIAM is also in the crosshairs of threat actors, who regard the login box as a path to the information, privileges, and benefits reserved for account holders. Through the vital functions of authentication and authorization, CIAM strengthens trust by preserving customer privacy and helps to protect organizations from the consequences of breaches. These contributions to a strong security posture place CIAM squarely in the sights of CISOs, CIOs, and compliance officers.

Finally, the rush to develop and deploy generative-AI-powered agents has created a new opportunity for CIAM to deliver unique value — as the means to control and help secure the access afforded to AI agents. Customers need to know they can trust AI agents with their personal data — otherwise, the transformative potential of these agents won't be realized. If CIAM wasn't already important to CTOs and other technology leaders, emerging AI-related applications are forcing reconsideration.

Drawing upon a global survey of 6,750 consumers and operational telemetry from the Auth0 platform, this report explores a number of themes that all pertain to trust — and how it can be built and maintained in a rapidly changing digital landscape.

- Part one examines **customer experiences and attitudes**, primarily around identity, authentication methods, and customer journeys.
- Part two shines a light on **threats against customer identity**, serving as a sobering reminder of the importance of strong security measures.
- Part three goes where **AI meets customer identity**, uncovering important insights that may influence how organizations roll out AI agents.

But before we dive into the details, let's start with some highlights.

Key findings

Customer experiences and attitudes

Trust and fraud are top of mind for customers

Security and trustworthiness top quality and value
When deciding whether or not to create an account with a service provider, more users regard the company's reputation/trustworthiness (cited by 74% of survey respondents) and the company's security measures (72%) as important factors than report the same for the overall quality and value of the company's products or services.

Identity fraud is a top-of-mind issue
Across the entire respondent population, 64% of users indicated that they are concerned about identity fraud, while only 10% expressed little or no concern. While the specific numbers varied by respondent cohort, the sentiment itself was universal.

Long forms at signup or login frustrate users
Having to fill out long signup or login forms stands alone as the most frequently cited source of signup or login frustration for users, selected by 62% of respondents — well ahead of needing to provide private or sensitive information (52%).

Signup and login friction harms conversions
Nearly a quarter of respondents report either *always* (6%) or *often* (17%) abandoning an online purchase due to issues with signup or login processes, and a further 40% report *sometimes* doing so — pointing to a fairly universal problem.

Passwords remain ever present, but reuse increases risk
While respondents regard passwords as the most convenient authentication method, 68% of users report reusing passwords across multiple accounts — largely because remembering unique passwords is hard. This poor password hygiene increases risk for customers and organizations by enabling brute force identity attacks.

Key findings

Threats against customer identity

The rising risks in a digital-first world

Fraudulent registration attempts are commonplace
In 2024, across the entire Auth0 platform, the median proportion of registration attempts that met the criteria of a signup attack stood at 46.1%. And the scale of these brute-force attacks is truly staggering — during a multi-month sustained attack against the Retail/eCommerce sector (which endured the highest proportion of attack events among the top 10 industries on the Auth0 platform), the number of signup attack events exceeded the number of legitimate signups by a factor of 120.

Account takeover (ATOs) present a constant threat
Although some ATO attempts target individuals, most employ brute-force login attack techniques against password-based authentication to compromise as many accounts as possible. In 2024, across the entire Auth0 platform, the median proportion of login attempts that exhibited clear malicious behavior was 16.9%. Once again, Retail/eCommerce was in the unenviable position atop the charts, with 22.2% of login attempts exhibiting clear malicious behavior.

MFA abuse remains common, but may be declining
In 2024, across the entire Auth0 platform, the median proportion of MFA events detected as being malicious was 7.3%. This single-digit figure continues a downward trend first noted in Okta's State of Secure Identity Report 2023. While a tiny portion of these MFA events can be attributed to legitimate MFA failures (i.e., a genuine user repeatedly failing an MFA challenge) and a likely somewhat larger portion results from threat actors being stymied by MFA, the majority of such events are almost certainly associated with MFA fatigue attacks and SMS pumping (toll fraud).

Key findings

AI meets customer identity

People prefer people — for now

Users strongly and universally prefer interacting with humans over AI agents

Across the full respondent base, 70% favored interacting with humans — compared to only 16% who favored interacting with an AI agent. In fact, every cohort examined in this study — spanning demographic generations, attitude toward new technology, and country of residence — expressed this preference. Users who prefer human representatives do so largely because humans understand their needs, while those who prefer AI agents point to faster resolutions, the absence of human interaction, and a belief in progress.

Users are more likely to employ AI agents for tedious and rules-based tasks

Users are most likely to employ an AI agent to accomplish tasks that are somewhat tedious and objective (as opposed to tasks requiring subjective consideration). At the other end of the spectrum, users expressed comparatively little desire to have AI agents handle more personal responsibilities.

User resistance largely comes down to a lack of trust

The leading reason customers who don't use AI agents feel this way is "*I don't trust AI agents with my personal data.*" More broadly, 60% of survey respondents reported being either very concerned or concerned about AI's impact on the privacy and security of their digital identities — and these concerns were expressed by every cohort examined in this study.

Oversight, transparency, and ethical guidelines would increase trust

While customers want humans to remain in the loop — 38% of survey respondents indicated that human oversight to review or approve decisions made by an AI agent would increase trust — this may not always be practical. Fortunately, users indicated that transparency, ethical behavior, and accountability would also increase trust.

Customer experiences and attitudes

Survey-powered demographic insights into trust, digital identities, and customer journeys

What *really* matters to today's customers?

For many organizations, the rate at which users create accounts is a key performance indicator (KPI), both as a measure of the success of ongoing efforts and as a leading metric.

Product owners and marketing leaders may view each such conversion as a relatively simple act — one that signals that a user sees sufficient value in your offering and is ready to take their relationship with your organization to the next level.

All they need to do is provide some basic information, *et voilà*.

But, from the user's perspective, is deciding to create an account really this straightforward?

Signup decisions start with trust and security

When a user signs up through an organization's website or app, they

- Signal their intent to start or continue a relationship
- Trigger the creation of an authoritative customer identity
- Potentially unlock valuable zero- and first-party data — both of which are becoming important alternatives to third-party cookies

Consequently, marketers, product owners, and user experience (UX) designers orchestrate entire campaigns and refine interfaces to nudge users toward creating an account.

Security and trust drive signup decisions

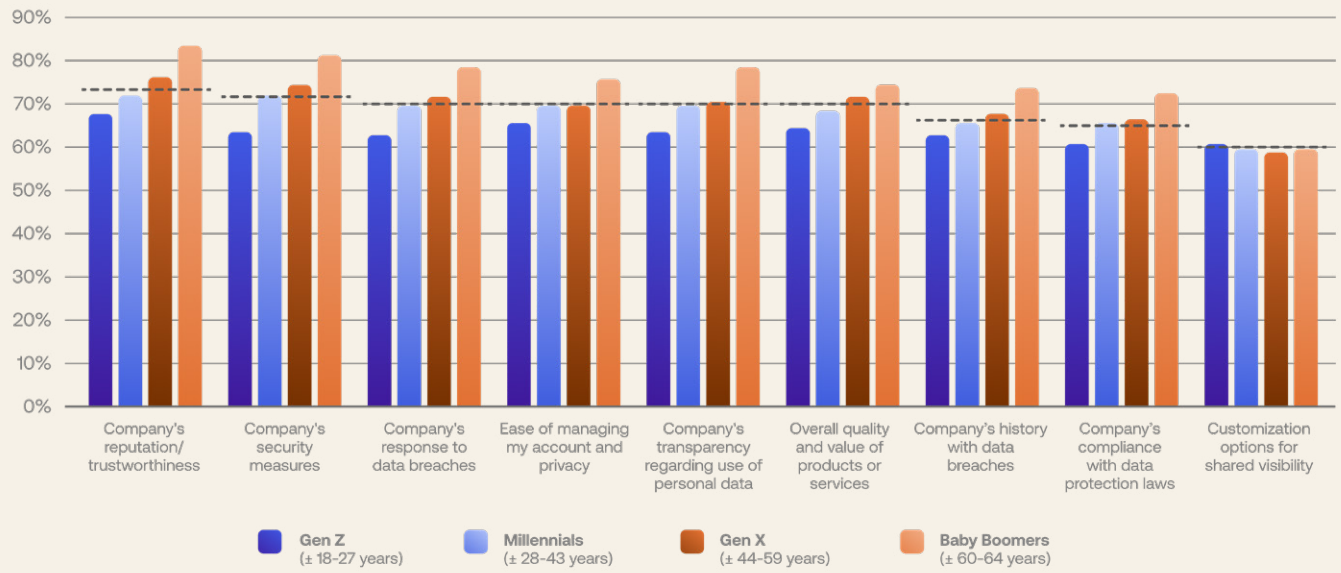

Influence of factors on signup decision, by respondent generation (sum of "very important" and "important")

"How important are the following factors when deciding to create a personal online account with a company?"

The value you provide isn't the most important factor users consider when deciding to create a new personal account

Respondents as a whole cited the company's reputation/trustworthiness and security measures as being *very important* or *important* more frequently than they did so for the overall quality and value of the company's products or services.

When offerings are similar, the brand with stronger security measures and privacy controls wins

Today's customers are concerned about identity fraud, which means they care about security, pay attention to breach headlines, and want to exert control over their own data.

Older or more tech-inclined users are even more likely to care about these factors

As a respondent's age or their affinity for new technology increases, the more likely they are to consider these factors — from trustworthiness to transparency — to be important when deciding to create a personal account.

Customers worry about identity fraud — and with good reason

Digital transformation has led to an age of convenience. Unfortunately, the same underlying technologies that contribute to efficiency, scale, and interconnectivity are regularly abused by malicious actors.

Today, identity fraud (the unauthorized use of another person's information) and its typical precursor, identity theft (the unauthorized acquisition of personal information), are all too common. Both are partly enabled by lax security and privacy controls that can result in data breaches, which can make it easy for criminals to impersonate other users.

Concern about identity fraud is universal among all users

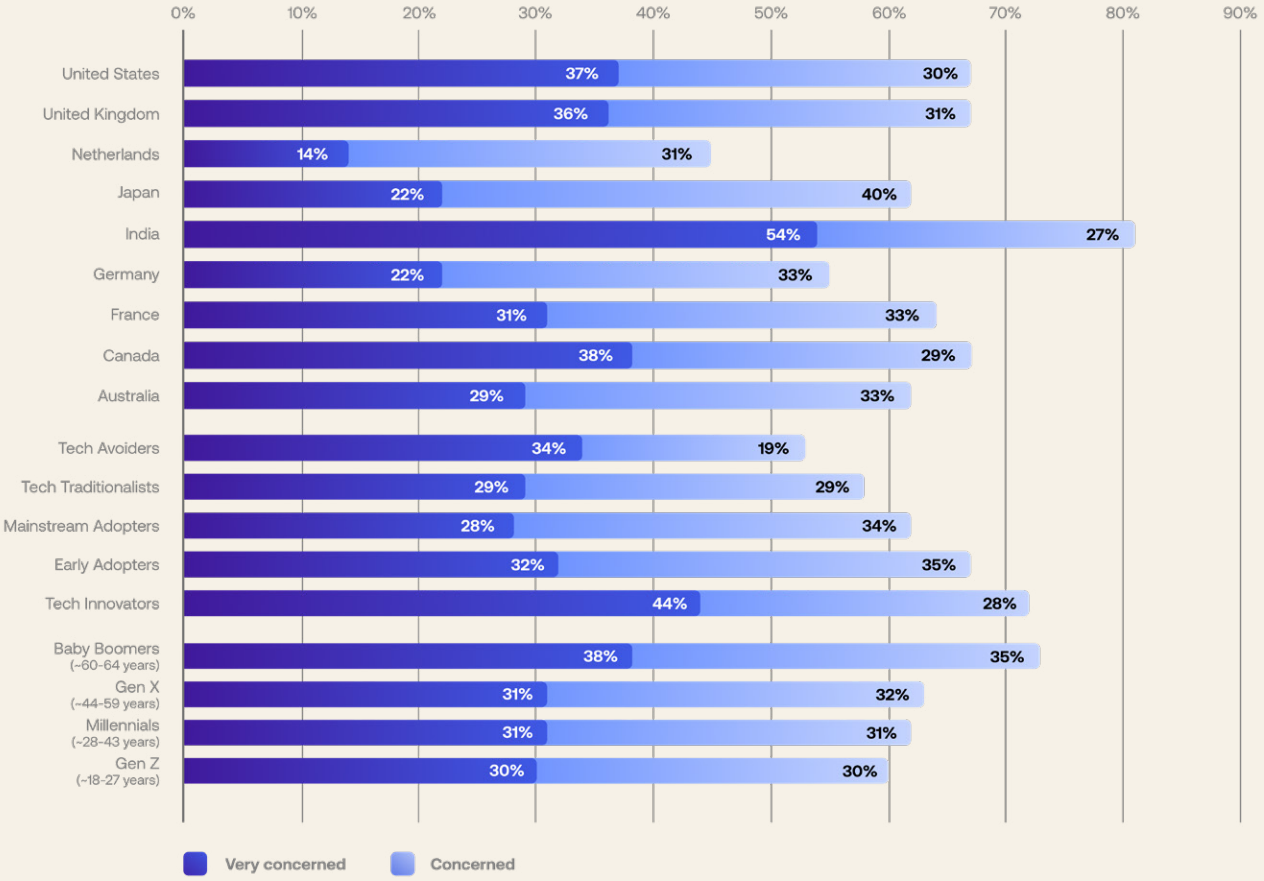

Concern about identity fraud, all cohorts

"How concerned are you about identity fraud?"

The strong majority of users are concerned about identity fraud
Across the entire respondent population, 32% of users indicated that they are *very concerned* about identity fraud; an additional 32% are *concerned*. With 27% as concerned as not, that leaves only 10% of respondents with little or no concern.

Concern about identity fraud is universal
As is clear from the figure, every demographic cohort is concerned about identity fraud. Respondents from the Netherlands expressed the least amount of concern — but even among this group, 45% are *concerned* or *very concerned*. Those from India have the most concern, with 54% being *very concerned* and a further 27% *concerned*.

Tech Innovators and Baby Boomers are also especially concerned
By generation, Baby Boomers express considerably more concern than the other cohorts. The differences aren't as stark when we group by attitude toward new technology, but Tech Innovators still stand out — largely on the strength of 44% (the second-most of any cohort) being *very concerned*.

Trust in data security depends on the player

A consumer's level of trust in a company drives revenue-generating behaviors — including their likelihood to purchase again, preference for a brand over competitors, and willingness to share personal data — according to Forrester's 2023 Consumer Trust Imperative Survey.

Per PwC's 2024 Trust Survey, the vast majority (90%) of business executives think customers trust their companies, but the truth is only 30% of consumers actually do — a staggering gap.

With identity fraud emerging as a top concern, customers are paying close attention to how institutions protect their personal data.

Which institutions users trust with their data

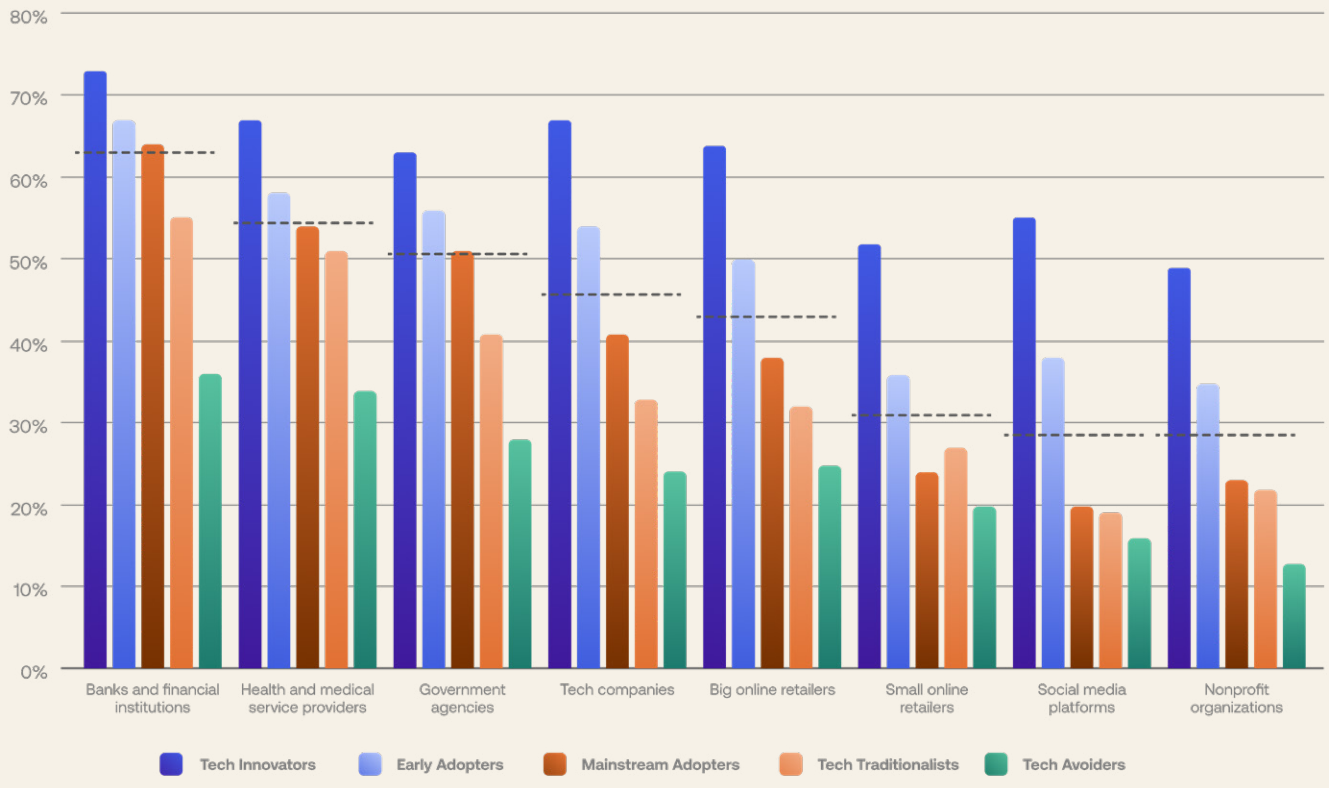

Trust in different institutions, by respondent attitude toward new technology (sum of "completely" and "rather completely")

"How much do you trust the following institutions to secure your personal data?"

Note: The dashed lines show the mean across the entire respondent population

Customer trust varies enormously by institution type

At the top end, 63% of respondents trust banks and financial institutions to secure their personal data. At the bottom, only 29% say the same about social media platforms and nonprofit organizations.

Those who embrace technology place more trust in institutions

Tech Innovators are especially trusting — 18 percentage points higher across all institutions, on average, than the full set of respondents. Tech Avoiders? Pretty much the exact opposite, averaging 19 percentage points lower than the full population.

Banks and financial institutions enjoy an enviable position

Every generation cohort placed the most trust in banks and financial institutions. For all other institution types, younger generations have more trust, and small online retailers and social media platforms are especially prone to this disparity.

Lengthy signups may be costing you customers

Customer experience often exists in some tension with security controls and the organization's desire to learn about their users.

For instance, strong safeguards are needed to provide confidence that each new account corresponds to a real user (rather than a bot), and that each login is being performed by the account's true owner (again, rather than a bot).

Likewise, organizations are eager to gather information that will help them deliver better overall service — including the high degree of personalization that many users now expect.

However, care must be taken to ensure that safeguards and information requests don't overly harm the experience. Better yet, organizations can implement solutions that serve the customer experience and business needs, such as progressive profiling and passkeys.

Long forms are the most frustrating signup or login hurdle
Having to fill out long sign-up or login forms stands alone as the most frequently cited source of signup or login frustration for users, selected by 62% of respondents.

Sources of frustration vary by country of origin
Respondents from Japan find most of these issues to be especially frustrating — and are particularly annoyed by complex password requirements and account verification.

Feelings are quite consistent across generations
Relative to the other cohorts, more Baby Boomers are frustrated with long forms and fewer Baby Boomers are frustrated by verification checks — but other than those differences, the generations are in agreement.

Top frustrations during signup and login

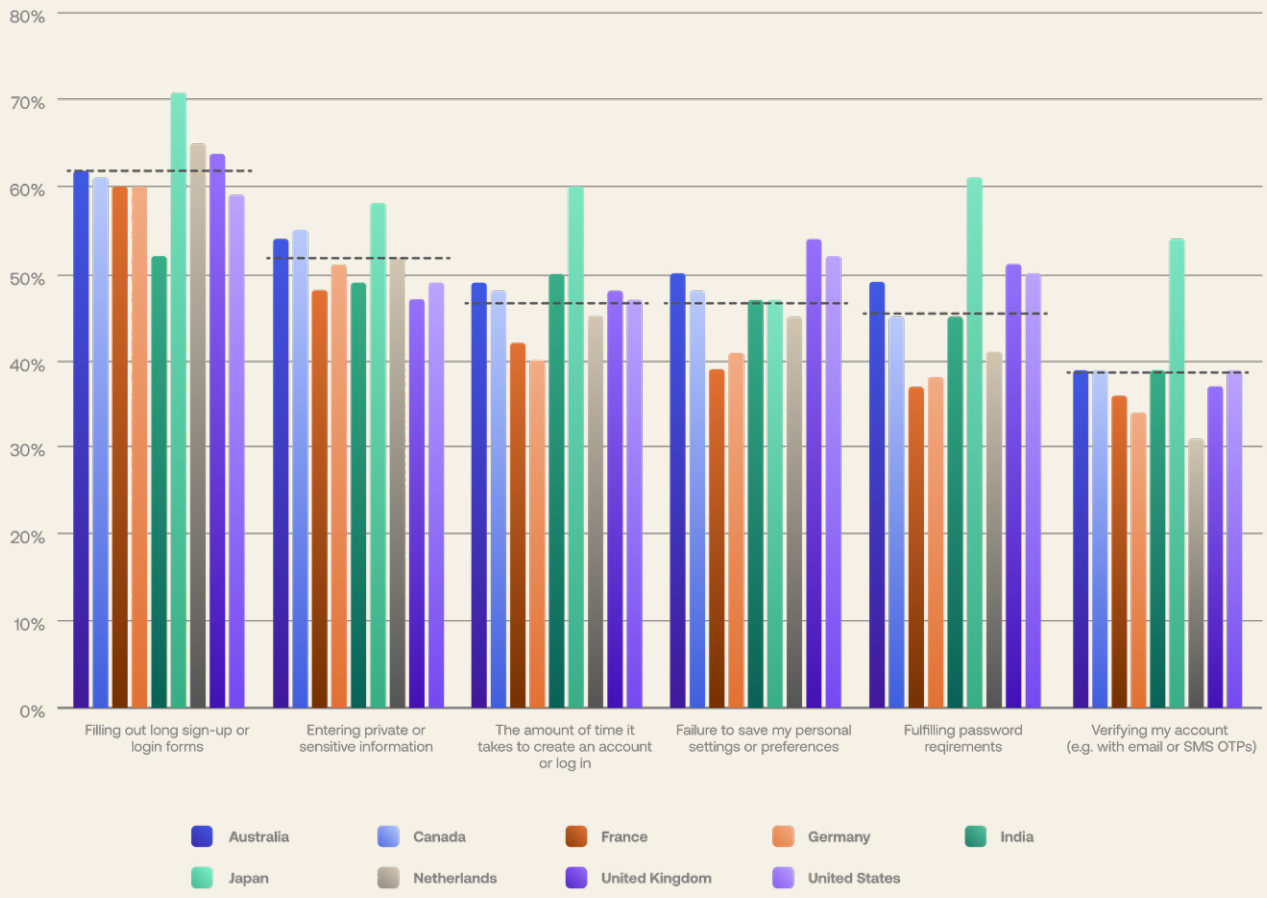

Factors that cause frustration during signup or login, by respondent country of residence (sum of "very frustrating" and "frustrating")

"How would you rate the following issues when signing up or logging in to a personal account?"

Note: The dashed lines show the mean across the entire respondent population

Customer identity and customer journeys

Today's companies must enable their customers to engage with their apps or services at any time, from any device.

In a customer identity context, these interactions may include (but are certainly not limited to) a user:

- Signing up for your service / registering an account with your organization
- Logging in to their existing account
- Providing you with consent to collect and use their data
- Updating their information and preferences
- Completing a transaction
- Resetting their password

Optimizing these flows and experiences — by simplifying and speeding up authentication, pre-filling forms, reserving MFA challenges for privileged access, being transparent about why data is being collected and how it will be protected and used, etc. — can have a direct impact on your overall conversion rates.

Identity flows are fundamental components of the customer journey and, as such, strongly influence conversion rates

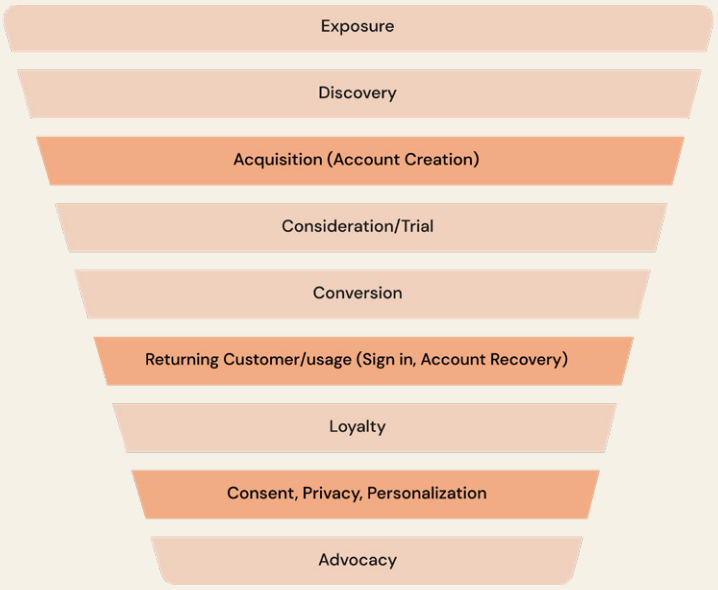

Abandon cart: Signup friction is a deal breaker for customers

In the digital world, *friction* refers to anything that slows down or otherwise impedes a user's interactions with your service.

While a "Goldilocks" (*just right*) amount of friction is required — to establish trust and implement the security controls that protect customers and companies — unnecessary friction creates poor customer experiences, lowers conversion rates, and can undermine your efforts to gather the data needed to build 360-degree profiles.

Signup and login issues drive customers away

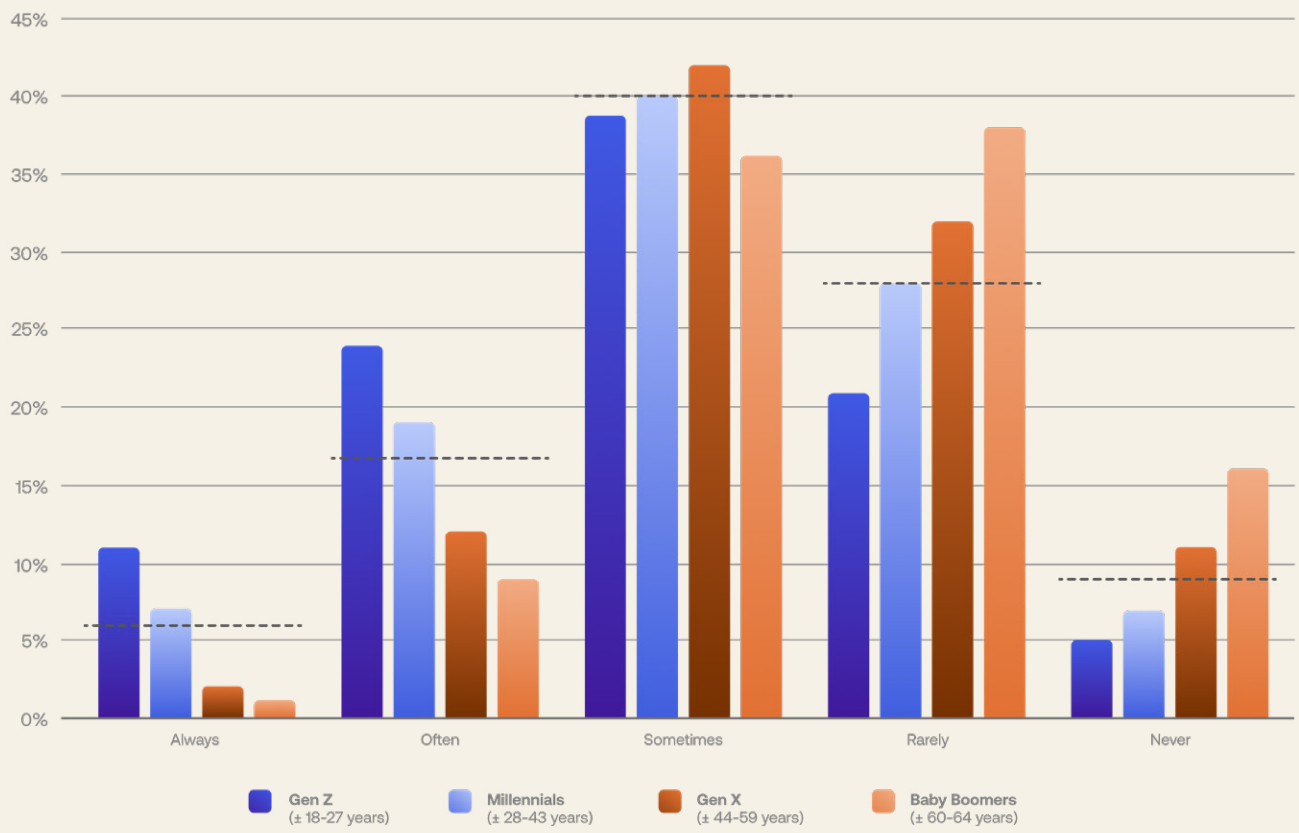

Frequency of purchase abandonment due to signup or login issues, by respondent generation

"How often have issues with the signup or login process led you to abandon an online purchase?"

Note: The dashed lines show the mean across the entire respondent population

Signup and login friction are the enemies of conversions

Nearly a quarter of respondents report either *always* (6%) or *often* (17%) abandoning an online purchase due to issues with signup or login processes, and a further 40% report *sometimes* doing so — pointing to a fairly universal problem.

Gen Z and Millennials are especially intolerant of signup and login friction

Younger cohorts — typically highly valued by businesses — report higher-than-average rates of purchase abandonment. Among Gen Z respondents, 11% *always* abandon an online purchase when faced with signup or login friction, and 26% *often* do so.

Tech Innovators are the least tolerant of all

A whopping 43% of Tech Innovators either *always* (17%) or *often* (26%) abandon an online purchase when faced with signup or login friction — all the most of any cohort under study. Clearly, members of this group know what a great experience looks like, and they expect brands to deliver it.

How convenience and familiarity shape login habits

When it comes to login habits, the saying "stick to what you know" seems to hold true. There's a clear and strong link between how often people use an authentication method and how convenient they think it is. The correlation coefficient[1] between reported usage and perceived convenience is 0.92 (out of 1), indicating a close relationship between the two.

Convenience and familiarity drive login habits

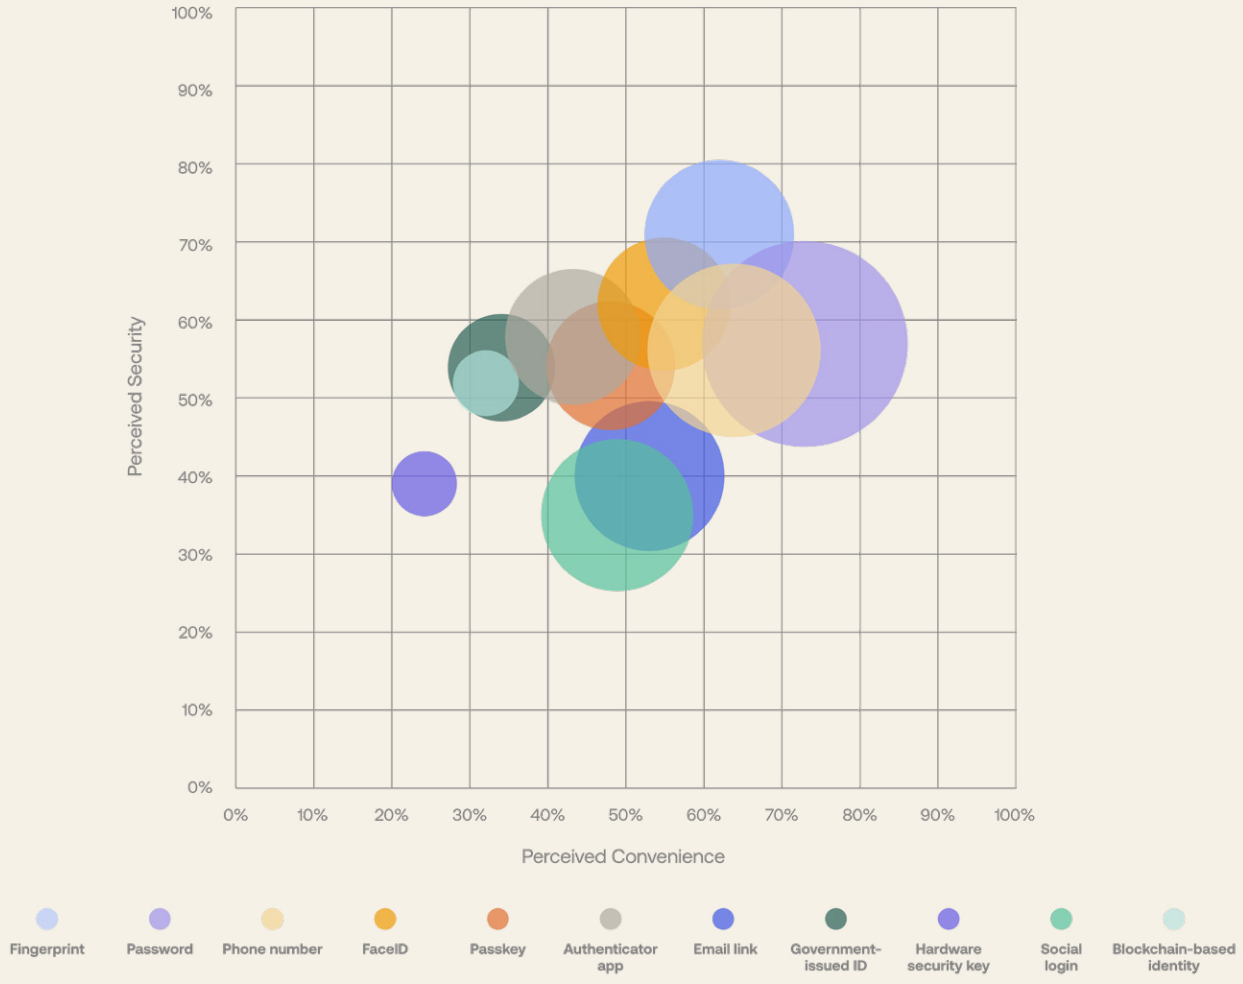

Perceptions and usage of authentication methods, all respondents

Note: This chart combines answers from three questions pertaining to the usage (bubble size), perceived convenience (horizontal axis, sum of "very convenient" and "convenient"), and perceived security (vertical axis, sum of "very secure" and "secure") of different authentication methods

[1] A numerical measure of the statistical relationship between two variables, ranging from –1 (strong negative correlation) to +1 (strong positive correlation), with 0 representing no correlation.

We believe the strong correlation is likely bidirectional:

- Using a method more frequently makes it feel more familiar — and greater familiarity influences opinions regarding convenience.
- Methods perceived as being more convenient are used more frequently.

In contrast, the correlation coefficient between perceived security and convenience is only 0.45, and between usage and perceived security is a very weak 0.19.

Users consider passwords convenient, flaws and all

Authentication can leverage three different factor types to confirm a user's identity:

- *Knowledge*, which presumes only the real user knows something (e.g., password)
- *Possession*, such as authenticator apps, government-issued ID, hardware security keys, and verification links and passcodes sent to a verified destination
- *Inherence*, as in biometric identifiers

Passwords' lingering appeal: perceived convenience

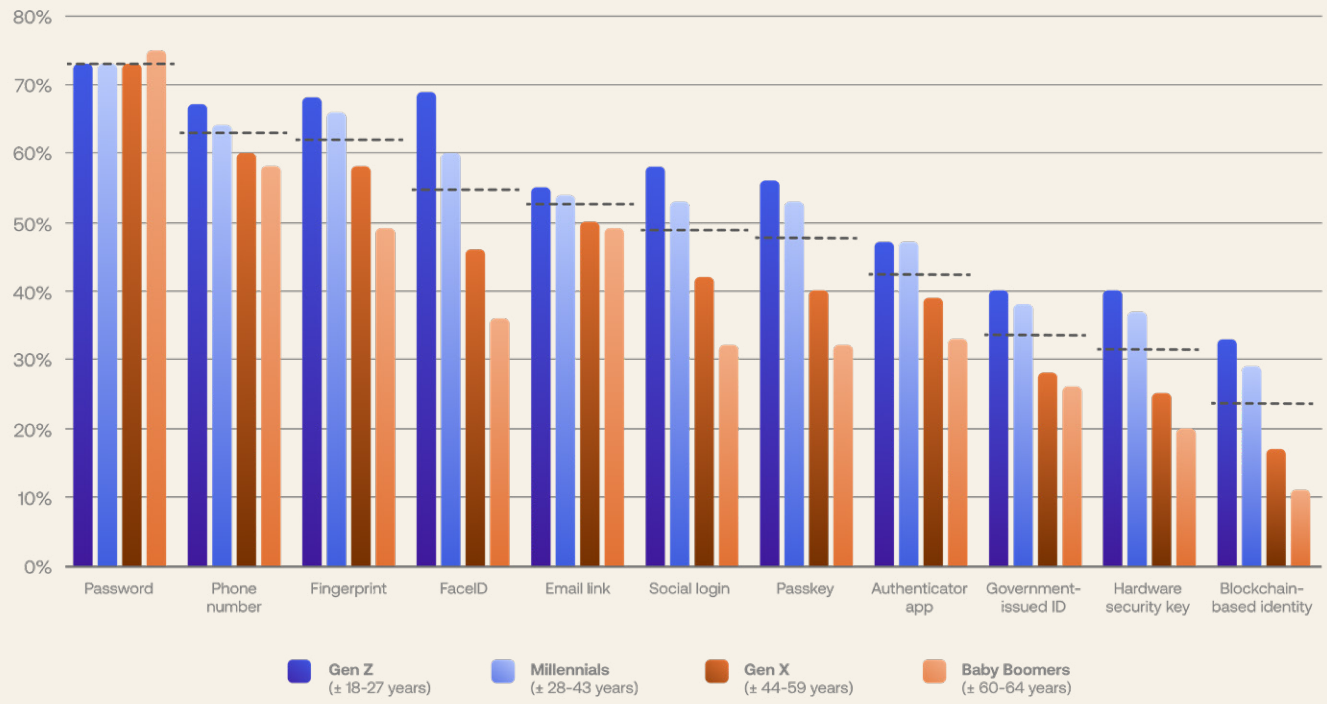

Perception of authentication method convenience, by respondent generation (sum of "very convenient" and "convenient")

"How would you rate each of the following methods for confirming your identity when signing up or logging into a personal account?"

Note: The dashed lines show the mean across the entire respondent population

Authentication security can be strengthened by combining two or more types of factor, as with multi-factor authentication (MFA). However, users may prefer a single-factor approach due to its simplicity and familiarity (e.g., passwords).

Passkeys bridge this usability gap by combining a possession factor with a knowledge or inherence factor in a single, convenient method.

Users perceive passwords to be the most convenient signup or login method

73% of respondents rated passwords as being either *very convenient* (40%) or *convenient* (33%) — fully 10 percentage points above the second most convenient method (phone number).

Biometric authenticators expose strong divisions

There's a significant generational divide in how fingerprint- and FaceID-based authentication are perceived — and the same pattern appears when respondents are grouped by their attitude toward new technology.

Passkeys are poised for growth

While passkeys are new enough to be unfamiliar to many users, more than half of Gen Z and Millennial respondents consider them to be convenient — suggesting that this more secure authentication method has a bright future.

Despite all the warnings, people are still reusing passwords

As attackers became adept at brute-forcing weak passwords and taking advantage of widespread password reuse, requirements about complexity evolved. This trajectory has resulted in ever-more complex requirements (e.g., for length, special characters, numbers, and combinations of upper and lowercase letters).

Users, for their part, have wholeheartedly embraced these requirements — and security hygiene in general — making password-based attacks a legacy of the earliest days of the digital age. Or have they? (They haven't.)

More than two-thirds of respondents reuse passwords

At this point, most every user has been told not to reuse passwords — yet 68% admit to either using the same password for every account (17%) or reusing only a small set of passwords (51%).

Tech Avoiders and Baby Boomers are the least likely to reuse passwords

These groups are doing their part to combat password-based attacks, with 42% of respondents within each cohort using a unique password for every personal account.

Passwords are too hard to remember

The most commonly cited reason for reusing passwords? More than half of respondents (53%) indicated that unique ones are "*too hard to remember*" — 22 percentage points higher than the second-leading reason, that using unique passwords "*takes too much time*."

Password reuse stubbornly persists

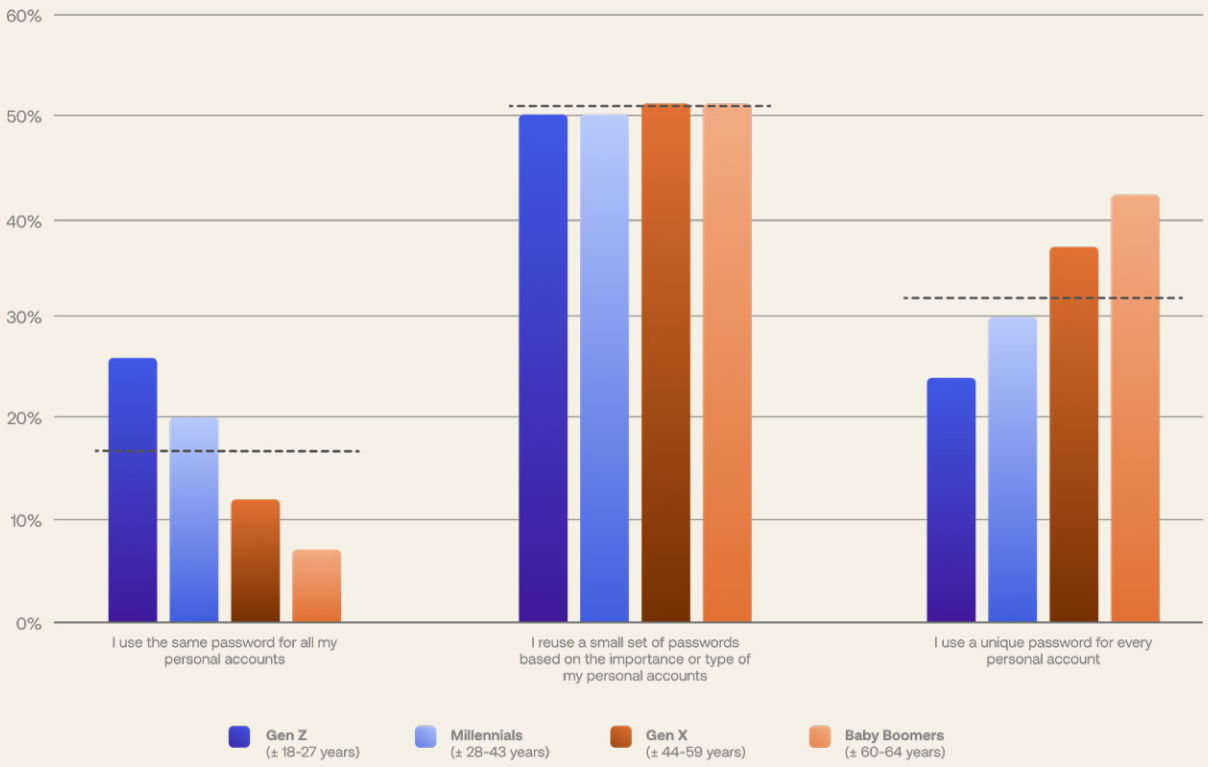

Password hygiene, by respondent generation

"Which statement about your password behavior applies to you the most?"

Note: The dashed lines show the mean across the entire respondent population

Customers point to fingerprints as the most secure login method

To be effective in customer applications — and especially with everyday consumers — authentication techniques must balance security and usability.

While legacy authentication techniques often imposed a tradeoff, modern approaches combine phishing-resistant security with the convenience of a fingerprint or facial scan, or the tap of a button on an authenticator app.

Biometrics seen as most secure login choice

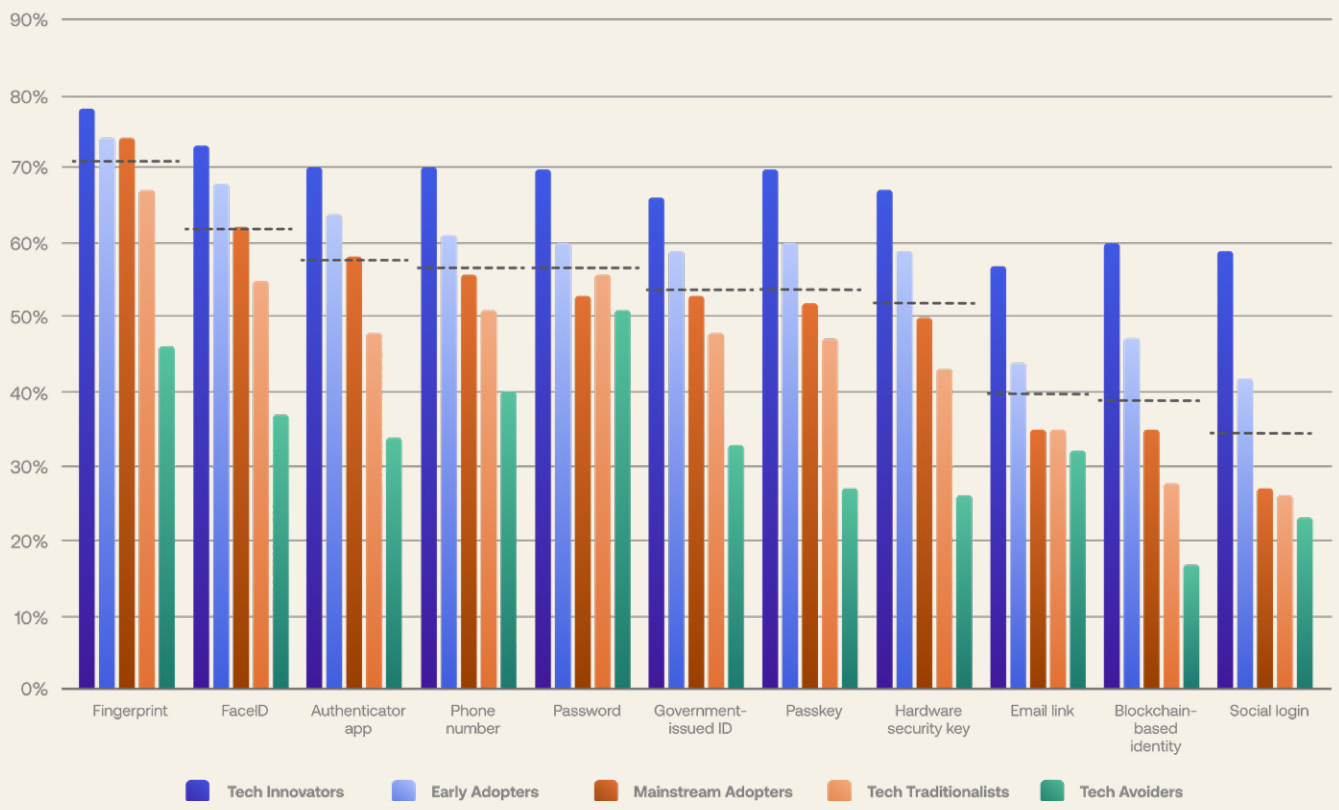

Perception of authentication method security, by respondent attitude toward new technology (sum of "very secure" and "secure")

"How secure would you rate each of the following methods when signing up or logging in to a personal account?"

Note: The dashed lines show the mean across the entire respondent population

Biometrics top the charts

Across every generation, fingerprint (chosen by 71% of users, overall) and FaceID (62%) finished one-two. The same was almost true for segmentation by attitude toward new technology — the only exceptions being Tech Avoiders' higher trust in passwords and phone number verification.

Social login is perceived as comparatively insecure

Only 35% of respondents regard social login as being *very secure* (13%) or *secure* (22%) — a worrying sign for marketers trying to nudge users in this direction, due to the demographic data that it can often provide.

Perceptions vary widely based upon a user's attitude toward new technology

As is clearly illustrated in the figure, Tech Innovators and Early Adopters are much more likely to regard a particular authentication method as being *very secure* or *secure* than Mainstream Adopters, Tech Traditionalists, and Tech Avoiders.

Personalized, private, and protected: Building trust in digital relationships

Considered as a whole, the results of the survey point to a paradox at the heart of online interactions and the digital Identities that enable them:

- Customers want frictionless, personalized, and instantaneous experiences when logging in to apps and making purchases;
- At the same time, they want to control what data they share, and they want appropriate security controls in place to protect that data.

Complicating matters for brands, a range of factors — including country of residence, age/generation, and attitude toward new technology, as we've seen — influences preferences and imposes specific requirements. Instead of a singular approach to customer identity, organizations should employ a combination of techniques to meet the personalization, privacy, and security needs of digital consumers.

It's imperative for brands to be mindful that digital relationships are formed and progress in the same way relationships do in real life: *over time*. The burden of establishing trust in a digital relationship will be on the service provider, and this trust must always be earned, respected, and protected.

Providing users with more secure authentication options

The survey findings repeatedly illustrated that different users have different perceptions and preferences, suggesting that brands should give users a choice between different secure authentication options.

By moving away from the traditional username-and-password combination, brands can concurrently strengthen security and create more convenient user experiences. A number of techniques are available to achieve the optimal balance appropriate to different scenarios.

For example, passkeys are vastly more secure than passwords and, while not as familiar to users as passwords, provide a more convenient authentication experience.

Other ways to reduce friction while strengthening security include:

- **Social logins**: Essentially single sign-on (SSO) for consumer apps, social logins streamline account authentication and reduce the risk that users will encounter problems when trying to log in to your services; user concerns about the security of this approach — concerns that could very well originate from the lack of friction — can be addressed with a short explanation of how it works (not so much the technical details, but the main takeaway that it's as secure as the authentication of the social identity provider)
- **Biometric authentication**: Increasingly supported by consumer devices and already regarded by customers as highly secure, biometrics replace cumbersome password entry with the convenience of a fingerprint or facial scan
- **Adaptive MFA**: A tool that only engages secondary factors when a user interaction is deemed risky based on behavioral data (e.g., an impossible travel scenario or a login from a new device)
- **Step-up authentication**: An approach that adapts authentication to the importance of the resources being accessed (e.g., a user may be prompted for additional authentication when attempting to alter account information or retrieve a sensitive document)

Acquiring personal data in a privacy-centric age

Countless surveys (including this one) have revealed that users find long forms frustrating and are questioning what brands are doing with their personal information.

To acquire the zero- and first-party data needed to craft the highly personalized experiences many subscribers expect, brands must adopt customer-friendly solutions like **progressive profiling**. This technique gradually asks the user for information as they experience more value from the service, reducing signup friction and avoiding triggering user concerns about providing a substantial amount of information to an unfamiliar brand.

Additionally, organizations should **be transparent with users about how digital identities are managed**, including why data is needed, how it will be used, and what security measures are in place to protect user accounts and the private data within them.

Threats against customer identity

A brief exploration of today's most common and most dangerous identity attacks — and how to safeguard applications against them

Securing customer authentication

Already, our digital identities control access to an ever-growing number of applications and services, impacting — and to some degree governing — many aspects of modern living. Over time, their importance will only grow.

Unfortunately, legitimate users aren't the only ones interested in what's behind the login gateway. There's money to be made for those who can break in, and economic forces have led to the emergence of an entire ecosystem of technologies, services, and other resources to enable such abuse.

Across industries, attacks against entities large and small continue. As cybercriminals direct more effort and expertise into getting past the login box, protecting it requires ever-more layers of ever-more sophisticated defenses — making authentication, authorization, and CIAM in general vital to preserving trust, security, and privacy.

No rest for the wicked: Fake signups don't take the day off

The easiest way for a malicious user to access the privileges, services, and information behind the login box is to create puppet accounts under their control from day one.

Especially when performed at scale, these fake signups can create significant problems and lead to unnecessary expenses. For example, fake users may

- Act as a foothold to be used later as aged accounts to bypass controls
- Enumerate/discover existing user accounts
- Consume rewards like signup bonuses
- Be used to execute denial-of-service attacks by consuming resources and exceeding rate limits

Fake signups are a constant threat

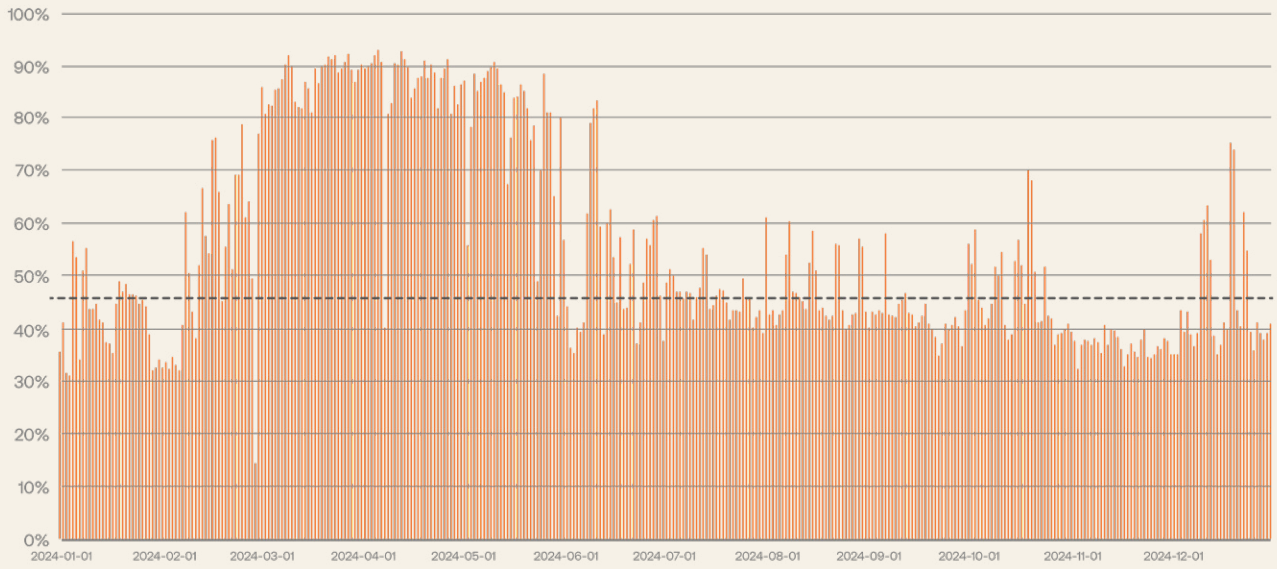

Suspected signup attacks, by day (January 1 to December 31, 2024)

Note: Each column shows the proportion of signup attempts on the Auth0 platform, on a given day, that met the criteria of a signup attack; the dashed line shows the median (46.1%) of these daily signup attempts across the entire platform; for the definition of a signup attack, see the Methodology

Additionally, entire conversion flows are often optimized based upon how users interact with the service. Fraudulent registrations pollute this data, complicating analysis and potentially leading to expensive clean-up projects.

Fraudulent signups remain an everyday problem

- Even just a glance at the chart above shows that fake registrations are a constant menace.
- In 2024, across the entire Auth0 platform, the median proportion of registration attempts that met the criteria of a signup attack stood at 46.1%. This reverses a downward trend and is consistent with a surge in signup attacks reported by other large tech companies — with some attributing the increase to AI-enabled attack workflows.

The day-to-day threat varies widely

- On April 6, 92.5% of registration attempts met the criteria of a signup attack.
- In contrast, on February 29 only 14.4% did so; this stands as a major outlier, as no other day had a signup attack proportion below 30%.
- We see a mix of short-term spikes throughout the year.
- Perhaps most evident of all, we see a sustained surge in malicious behavior running from mid-February through the end of May.

Fraudsters zero in on large retailers, FinServ companies

Deeper inspection of the underlying data reveals that fraudulent registration attempts are unevenly distributed.

In particular, organizations in different industries and of different sizes experience different rates of attack.

Retail/eCommerce companies are the top targets of fake signups

- In 2024, signup attacks accounted for nearly 70% of registration attempts with Retail/eCommerce companies — the highest proportion of the 10 industries with the most representation on our platform.
- Many online retailers offer signup incentives and member-only exclusives, and these programs may be attracting the attention of threat actors.

Fraudsters focus on retail and finance

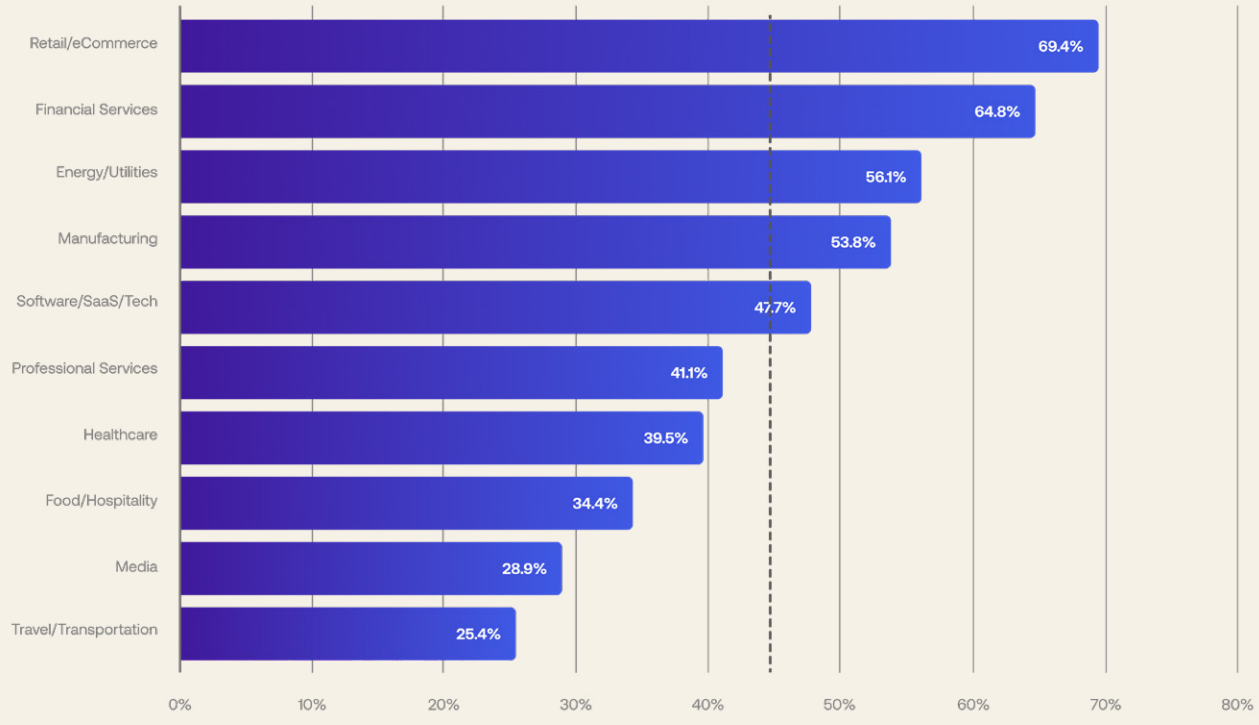

Suspected signup attacks, by industry (January 1 to December 31, 2024)

Note: Each bar shows the median daily proportion of signup attempts, for a given industry, that met the criteria of a signup attack, in 2024; the dashed line shows the mean (44.5%) of per-industry daily medians across all industries on the Auth0 platform (i.e., not just the top 10)

Financial Services (FinServ) companies are also in the crosshairs

- Nearly 65% of attempted registrations with FinServ companies were associated with signup attacks.
- Notably, this category includes many cryptocurrency startups, which often offer coins/tokens as welcome gifts.
- Accounts with more traditional FinServ organizations may also be desirable, as they can help to facilitate money laundering and synthetic identity fraud.

Enterprises attract the most unwanted attention

- Enterprises endured the highest proportion of signup attacks, at 64.3% of registration attempts.
- Mid-market organizations fared considerably better, with only 18.2% of registration attempts meeting the criteria of a signup attack.
- Small businesses fell right around the mean, with 43.3% of signup attempts displaying clear malicious behavior.

To illustrate, let's briefly examine the Retail/eCommerce sector, since it topped the chart on the previous page.

From the Legitimate Signup Events chart, we can see that the number of genuine registrations handled by the Auth0 platform is quite steady.

In obvious contrast, the Signup Attack Events chart exhibits much greater variation. We can also see that the Retail/eCommerce sector was the target of the multi-month surge seen two pages ago.

During this sustained attack period, **the number of signup attack events exceeded the number of legitimate signups by a factor of 120.**

The massive scale of signup attacks

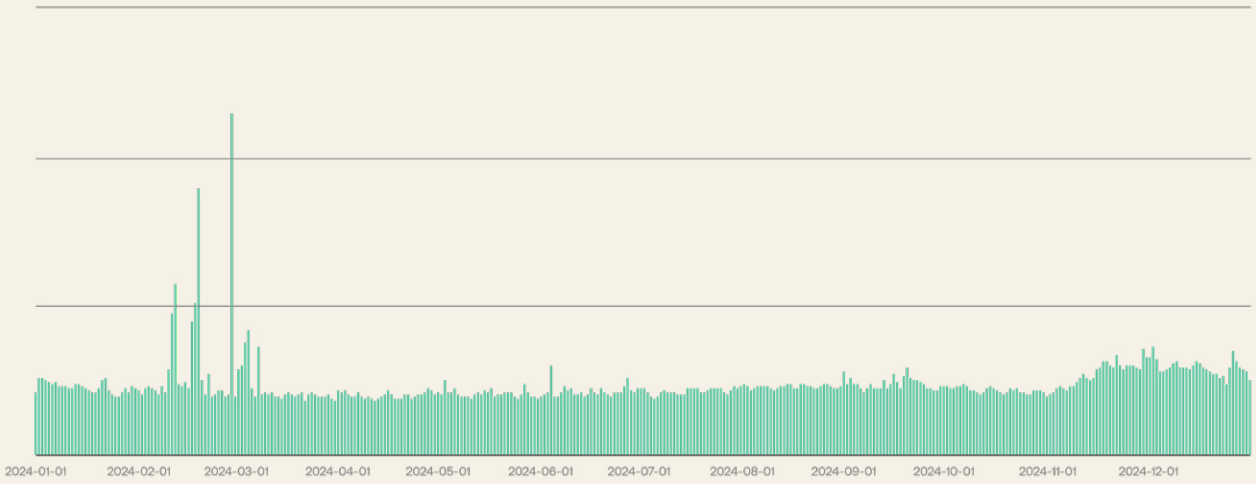

Legitimate Retail/eCommerce signups, by day (January 1 to December 31, 2024)

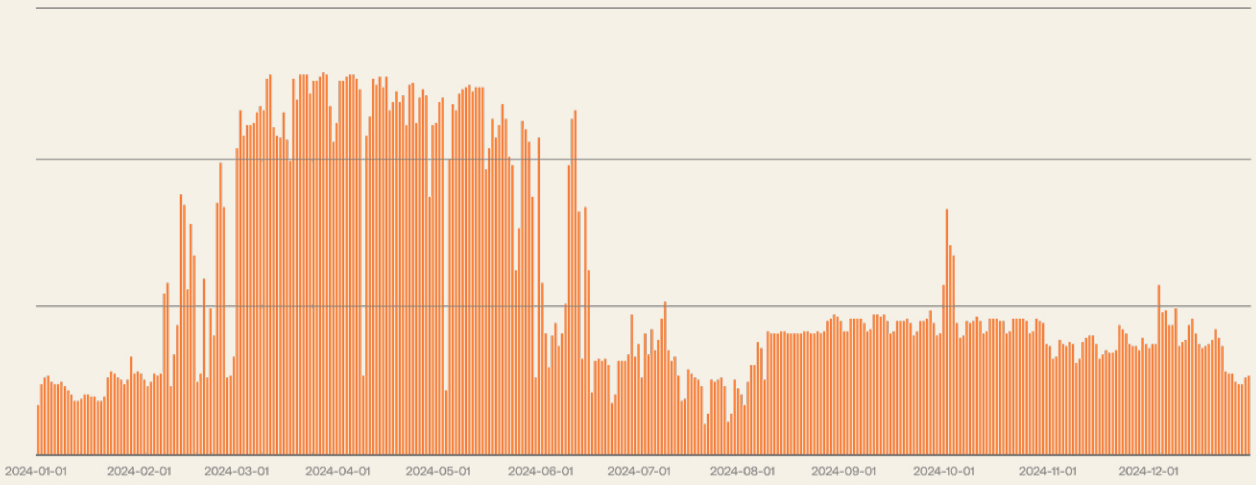

Suspected fraudulent Retail/eCommerce signups, by day (January 1 to December 31, 2024)

Note: Unlike most other graphs in this report, these two show absolute counts (rather than relative proportions); to enable easy visual comparison between them, these charts use the same logarithmic vertical access, which has been truncated by many orders of magnitude

How to defend against login attacks

Detecting threat actors as early as possible and taking action to remove them from the registration pipeline reduces system load and limits their ability to perform reconnaissance (e.g., by receiving and analyzing error messages).

To that end, a number of defensive measures exist across different layers of the identity infrastructure:

Host-layer defenses against brute-force attacks

Note: *These defenses apply to any of the brute-force attacks examined in this report (i.e., not just signup attacks).*

To prevent abuse of the services they host, hosting providers (e.g., Cloudflare, Microsoft Azure, Amazon Web Services) apply a range of defenses. In the context of customer identity, these defenses are upstream of the CIAM functionality, and typically include

- **Distributed Denial of Service (DDoS) mitigation:** Protections help your CIAM application remain available to legitimate users, even in the face of large-scale attacks (particularly at the TCP/UDP layer).
- **Bot management:** An initial layer of bot filtering is typically based upon a combination of behavioral analysis, threat intelligence, and feedback loops.
- **Rate limiting:** Controls help protect against DoS attacks, brute-force strategies, and API abuse by imposing restrictions on the rate at which a particular entity can access the CIAM platform/application.

Platform-layer defenses against brute-force attacks

Note: *These defenses apply to any of the brute-force attacks examined in this report (i.e., not just signup attacks)*

Your CIAM platform should provide you with an array of defensive capabilities to combat signup attacks. However, when determining the appropriate response to malicious behavior, it's critical to consider the trade-offs involved — in particular, how much friction your users will tolerate during the registration process.

As is often the case in security, layered solutions are most effective. The more of these techniques you can implement, the safer your customers and your organization will be.

1. **Employ bot detection**. Platform-layer bot detection capabilities typically analyze telemetry (of which more is available than at the host layer) to predict when a signup attempt is likely coming from a bot. Above a certain prediction/confidence threshold, the account registration flow presents a countermeasure intended to be easy for a legitimate user to satisfy but difficult — and costly — for a bot.
2. **Enable or increase CAPTCHA requirements**. If your applications aren't using CAPTCHA today, you should strongly consider enabling this functionality. Although CAPTCHAs do increase end-user friction, many users are accustomed to them and understand why they are in place. A balanced approach is to only show a CAPTCHA when a risk threshold has been reached; if you go this route, consider always showing a CAPTCHA if there are indications of a large-scale signup attack campaign. Also, be mindful that no CAPTCHA is perfect, and that a dedicated attacker will always be able to bypass one eventually. The goal is to not to make your signup process impervious to abuse, but to make it sufficiently hard to abuse that an attacker will decide to find an easier target.
3. **Tighten brute-force and suspicious IP thresholds.** Both of these approaches limit the number of allowed connections, and should not present any issues to genuine customers.
4. **Block abusive IPs.** Your CIAM platform should allow you to implement access control list (ACL) rules to outright block abusive IPs.
5. **Block malicious activity using Web Application Firewall (WAF) rules at edge.** If you are using an edge provider or a sufficiently equipped CIAM platform, consider blocking particularly abusive IPs, ASes, geographic locations, TLS clients, or other HTTP header elements (e.g., User-Agent strings) being used by attackers.

Signup attack-specific defenses

In addition to the brute-force attack defensive layers listed above, there are several signup attack-specific techniques that can be applied to reduce fraudulent registrations:

Perhaps the most effective is to **encourage users to sign up using a passkey**, as the cryptography behind them makes passkeys extraordinarily difficult to abuse as part of a signup attack. While beyond the scope of this report, **analysis of Auth0 platform telemetry shows attackers are not yet performing mass-scale attacks against passkey signup or login flows.**

Other defenses include

- **Pre-signup rules and actions** (e.g., enforce a challenge, require more information) to further reduce the chances that a new user is illegitimate
- **Social login** to "outsource" prevention of fraudulent signups
- **Identity proofing** when risk is perceived to be particularly high
- **Validating contact information** (e.g., email address, phone number), for example through a one-time passcode or magic link

Crucially, intelligence gained by studying attacks can inform refinements.

To illustrate this point, let's return to the charts presented previously contrasting legitimate and fraudulent Retail/eCommerce signups. Notice that a handful of spikes of permitted signups coincide with the beginning of the sustained period of high-volume attack.

These short-lived spikes represent attacks in which some malicious traffic evaded detection. In addition to underscoring the importance of preventative controls like bot detection and CAPTCHA, they also highlight the importance of reviewing logs and security dashboards to gain insight into the evolving tactics employed by threat actors.

In this case, analysis enabled the Auth0 security team to quickly refine the platform's defenses, resulting in the legitimate signup rate returning to normal despite a months-long heavy bombardment of fake signup attempts.

The ongoing battle against account takeovers (ATOs)

While fraudulent registrations are (at a minimum) an expensive nuisance, account takeovers (ATOs) pose a markedly greater threat to security and privacy.

In a B2C context, attackers may gain access to resources (e.g., loyalty points), privileges (e.g., ability to make purchases, especially of products in limited supply), and valuable demographic and personally identifiable information (PII).

In a B2B context, an attacker who successfully compromises a user account may use it to access highly sensitive data, resulting in a breach with severe regulatory and contractual penalties for the targeted organization.

Although some ATO attempts target individuals, most employ brute-force login attack techniques against password-based authentication to compromise as many accounts as possible.

Account takeovers: a persistent threat

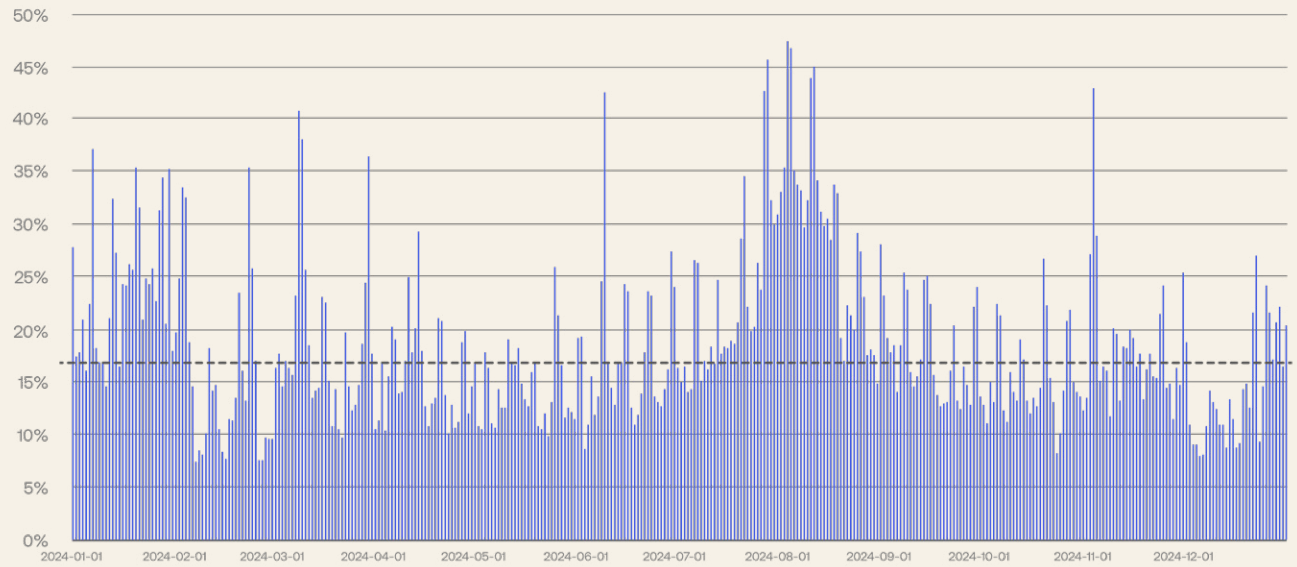

Suspected login attacks, by day (January 1 to December 31, 2024)

Note: Each column shows the proportion of password-based authentication events, on a given day, that met the criteria of a login attack; the dashed line shows the median (16.9%) of these daily login attempts across the entire platform; for the definition of a login attack, see the Methodology

Login attacks make ATOs a constant threat

- In 2024, across the entire Auth0 platform, the median proportion of login attempts that exhibited clear malicious behavior was 16.9%.
- This figure continues a downward trend first noted in Okta's State of Secure Identity Report 2023.

As with fake signups, the average masks tremendous day-to-day variation

- Across the full year, the highest daily rate of login attacks was 47.4%, while the lowest was 7.6%.
- We also see the same mix of short-term spikes (albeit many more than was the case with fake signups) as well as sustained multi-week periods with high login attack rates.

Top ATO targets? Online retailers and eCommerce

Diving more deeply into the data reveals that threat actors target particular high-value industries with their account takeover attempts.

Retail/eCommerce companies are the main targets of login attacks

- Just a quick glance at the chart (below) reveals Retail/eCommerce as a statistical outlier, with companies in this group having 22.2% of login attempts exhibiting clear malicious behavior — more than twice the average across all industries.

Professional Services and Financial Services organizations are also under fire

- Both of these industries have login attack rates right around the all-industries mean-of-medians rate.
- Another thing they have in common: Companies in these groups are likely to store highly sensitive personal information, including financial details, attractive to threat actors.

As with fake signups, enterprises attract the most unwanted attention for login attacks

- The enterprise segment endured the highest rate of malicious logins, with 24.9% of attempted logins meeting the criteria of a login attack.
- Nevertheless, login attacks remain an ongoing reality for mid-market organizations (7.6%) and small businesses (9.4%), merely to a much lesser extent.

Where suspected ATO attacks are concentrated

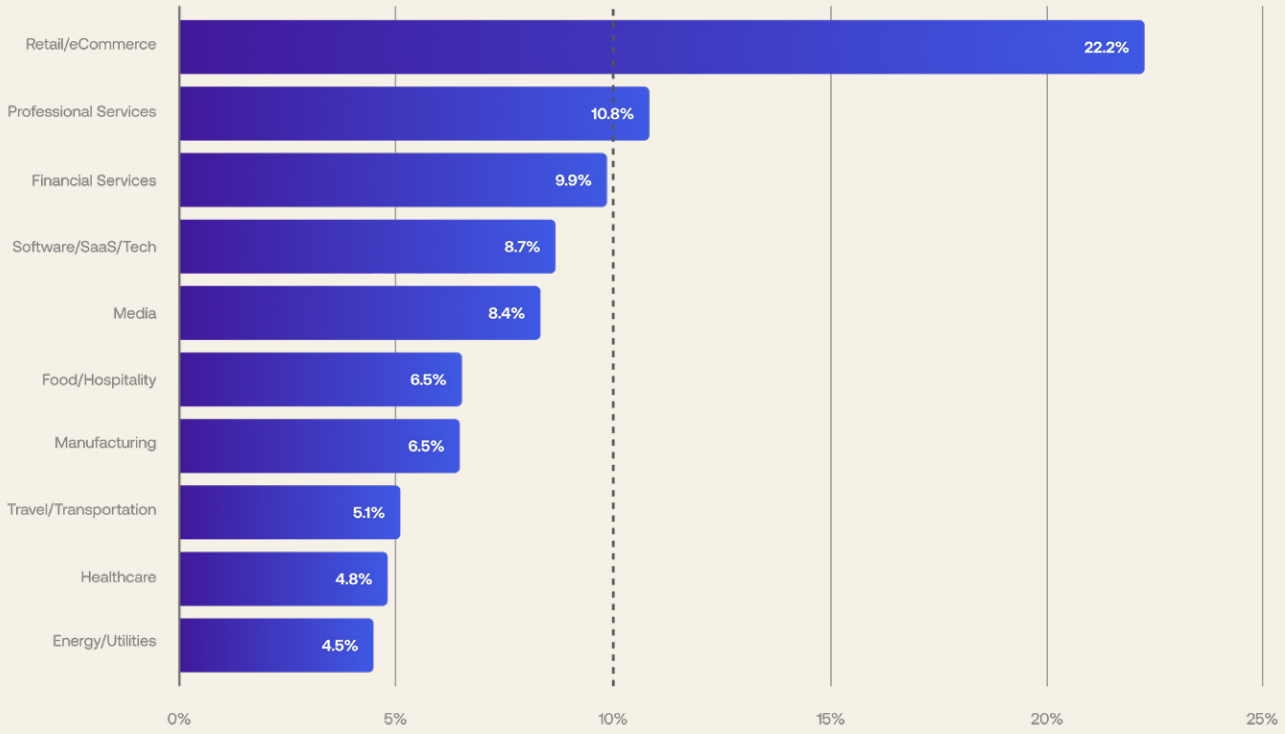

Suspected login attacks, by industry (January 1 to December 31, 2024)

Note: Each bar shows the median daily proportion of login attempts, for a given industry, that met the criteria of a login attack, in 2024; the dashed line shows the mean (10.0%) of per-industry daily medians across all industries on the Auth0 platform (i.e., not just the top 10)

ATO defenses in action: A tale of two charts

As Retail/eCommerce organizations endured the highest login attack rate in 2024, let's once again look more deeply at this segment.

The Legitimate Password Authentication Events chart shows the steady state of legitimate login activity.

The Login Attack Events chart isn't nearly as consistent, however. It varies widely through the year and shows a sustained attack campaign running from mid-June through mid-September.

On a normal day free of large-scale malicious behavior, genuine password authentication events outnumber login attack events by roughly 10 to 1.

However, **during the prolonged attack campaign, login attack events outnumbered legitimate password authentications by more than 62 times.**

A closer look at Retail/eCommerce authentications

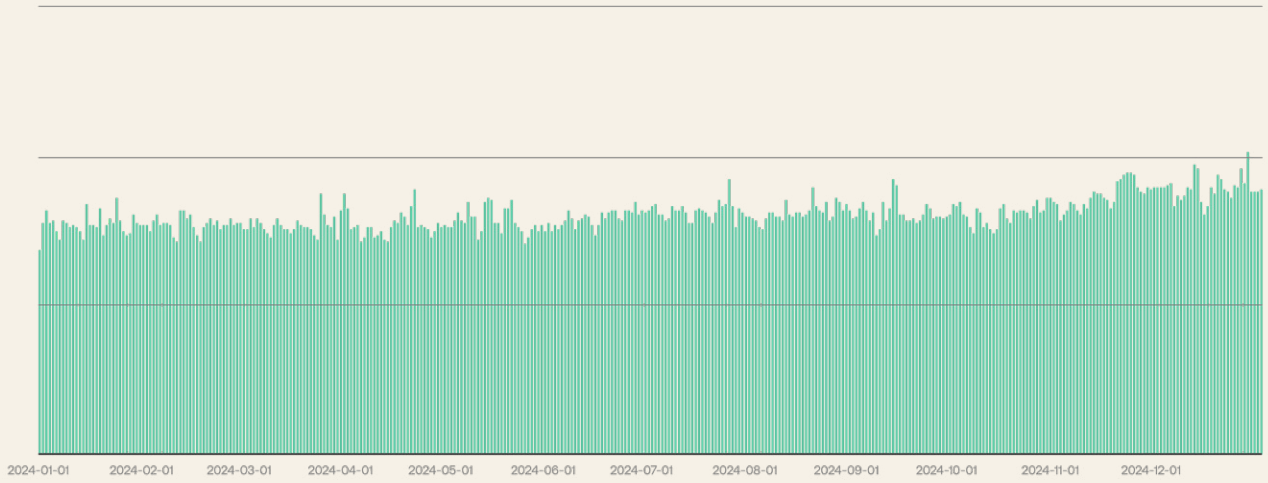

Legitimate password authentications, Retail/eCommerce sector, by day (January 1 to December 31, 2024)

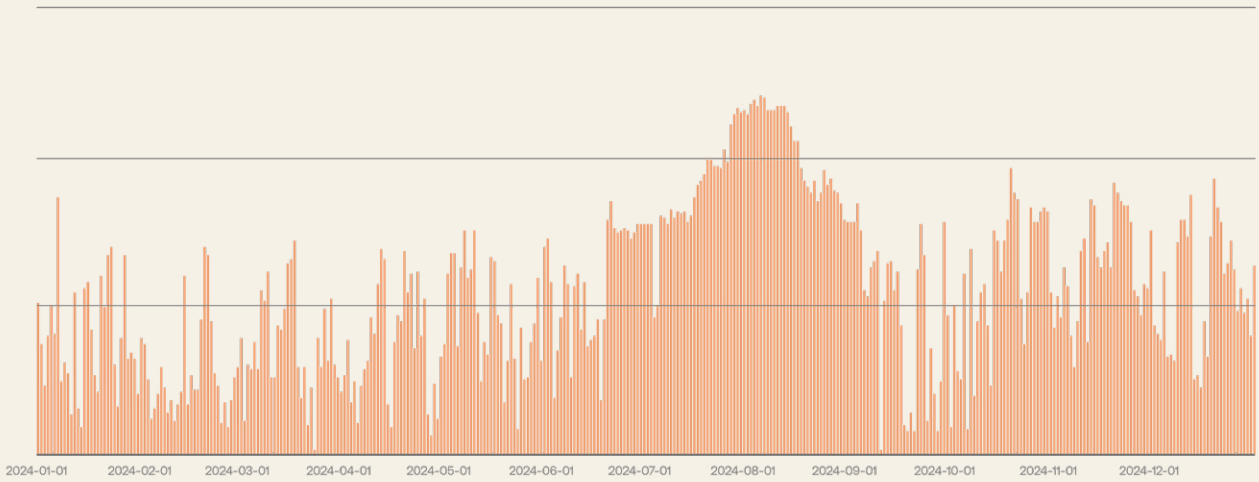

Suspected malicious password authentication attempts, Retail/eCommerce sector, by day (January 1 to December 31, 2024)

Note: Unlike most other graphs in this report, these two show absolute counts (rather than relative proportions); to enable easy visual comparison between them, these charts use the same logarithmic vertical access, which has been truncated by many orders of magnitude

How to defend against login attacks

Brute-force login attacks — most notably credential stuffing, password spraying, and password guessing — take advantage of poor password habits and policies, including

- **Password reuse:** Recall from earlier that 68% of users report reusing passwords across multiple accounts.
- **Minor modifications:** Many users who create a unique password for each account do so by slightly modifying a small set of passwords (e.g. using the same password with a different number at the end).
- **Short passwords:** Broadly, the longer a password the more effort required on the part of an attacker to guess it.
- **Long-lived passwords:** Combined with password reuse, failing to change a password makes an account more susceptible to compromise.

Such habits dramatically reduce the cost and effort associated with launching login attacks. For example, a small number of optimizations — including leveraging lists of breached passwords and dictionaries of words that are frequently used within passwords (or passphrases) — can dramatically improve the likelihood of trying the correct password (or, more accurately, of trying a password that hashes to the same value as the correct password).

Unfortunately, the barrier to launching login attacks is very low, and threat actors employ a number of tactics to try to evade defenses. For example, an attacker may intersperse known valid credentials — perhaps from fraudulent accounts already under their control — into the login stream to carefully manage the failure rate.

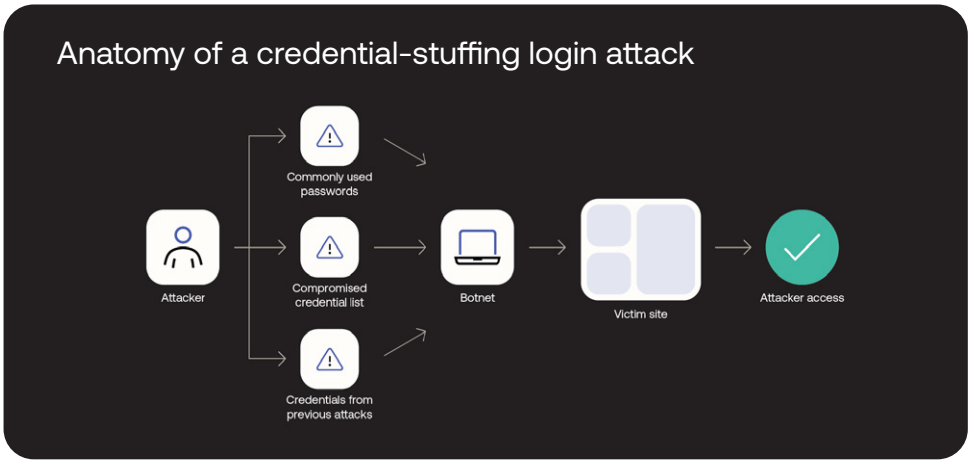

Platform-layer defenses against brute-force login attacks

Note: *These defenses should be applied in addition to the brute force attack defenses covered previously.*

Your CIAM platform should provide you with an array of defensive capabilities to combat login attacks. As is the case with signup attacks, when determining the appropriate response to malicious behavior, you must consider the impact to legitimate users.

Again, layered solutions are most effective, and the more of these techniques you can implement, the safer your customers and your organization will be.

- **Require users to reset breached passwords.** Presuming some form of breached password detection is available, customers whose credentials have been posted online should be forced to go through the password reset process. Note that this approach does not protect against dictionary attacks, but those can be dealt with using tighter brute-force thresholds.
- **Disable unneeded/unused features:** Unused endpoints or functionality can unnecessarily broaden your attack surface and allow for attackers to bypass controls (e.g., bot protection); unless a feature is absolutely required for your use cases, we recommend disabling it to avoid abuse by threat actors.
- **Block logins in impossible travel scenarios.** Block login attempts originating from a geolocation that would be impossible to reach within the time that has passed since the previous permitted login.
- **Require MFA for compromised accounts.** This approach avoids unnecessary friction by reserving MFA for users whose accounts are known to be compromised.
- **Implement adaptive MFA.** This approach avoids unnecessary friction by reserving MFA challenges for login attempts that exceed a predefined risk threshold.

- **Require strong, phishing-resistant MFA.** When introducing MFA, prioritize authenticator apps and WebAuthn-based methods; if you've already supported MFA for a long while, make an effort to migrate existing users away from legacy approaches and in favor of stronger secondary factors.

But perhaps the most effective defense against password-based ATOs is to **move away from passwords altogether** — a prospect that became much more realistic (especially in consumer markets) with the arrival of passkeys.

MFA abuse is common, but there are signs of decline

MFA challenges a user to prove their identity via two or more factors.

Unfortunately, while strong MFA is an effective defense against ATOs, threat actors routinely abuse MFA implementations.

As explained in the Methodology, the malicious MFA events presented in this report include

- Attempts to bypass MFA via bombing/fatigue attacks
- Attempts to commit toll fraud by abusing MFA to trigger phone or SMS messages
- Instances when a threat actor repeatedly fails an MFA challenge
- Instances when a legitimate user repeatedly fails an MFA challenge (These will represent only a tiny fraction of the events and will not skew the analysis.)

MFA abuse is common

- In 2024, across the entire Auth0 platform, the median proportion of MFA events detected as being malicious was 7.3%, with the majority likely attributable to MFA fatigue attacks and SMS pumping (toll fraud).

...but may be on the decline

- For context, Okta's State of Secure Identity Report 2023 showed a multi-year decline in the rate of MFA abuse and noted that for the first half of 2023, 12.7% of MFA attempts overall were considered malicious.
- In comparison, the 2024 mean was 7.8% — and dropped under 7% over the last half of the year.

Threat actors encounter and abuse MFA

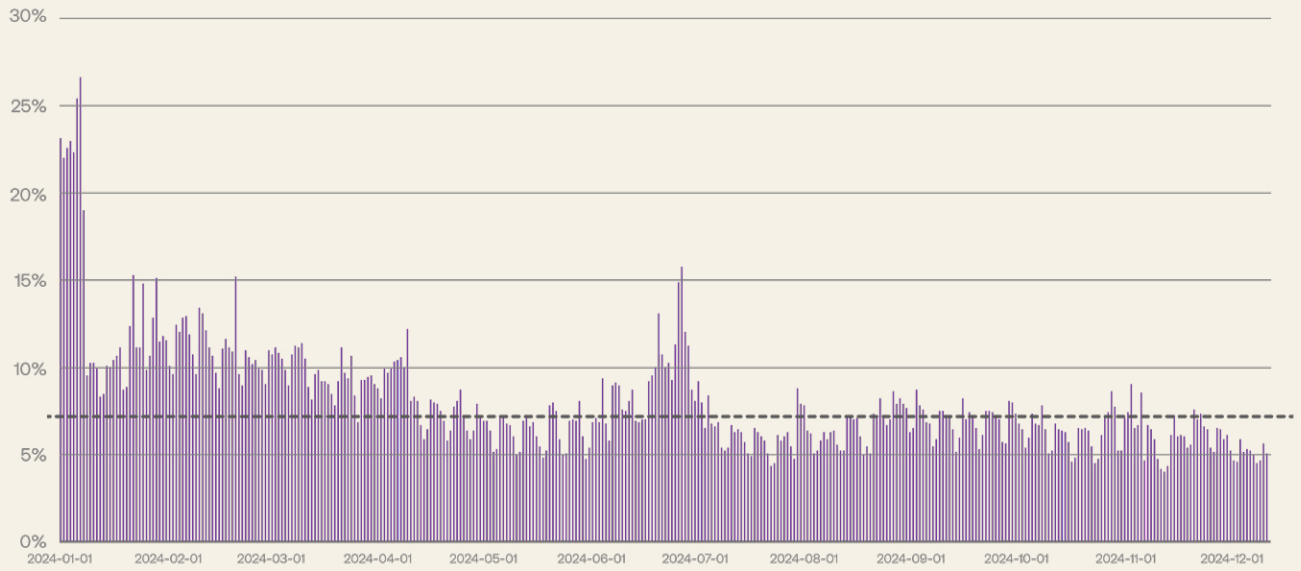

Suspected malicious MFA events, by day (January 1 to December 31, 2024)

Note: Each column shows the proportion of MFA events, on a given day, that met the criteria of malicious MFA activity; the dashed line shows the median (7.3%) of these daily MFA events across the entire platform; for the definition of a malicious MFA event, see the Methodology

Lights, camera, attack: Media leads in malicious MFA events

Breaking out malicious MFA events by industry reveals enormous variation.

Media companies continue to have the highest rate of malicious MFA events

- In 2024, more than 20% of MFA events associated with the Media industry exhibited malicious behavior.
- The Media industry also topped this chart in Okta's State of Secure Identity Report 2023, with 12.8%.

Professional services organizations and manufacturers also have higher-than-average rates of malicious MFA events

- In the professional services sector, 17.6% of MFA events were associated with malicious behavior.
- Consistent with Okta's State of Secure Identity Report 2023, manufacturers once again placed third, with nearly the same rate of malicious MFA events — 8.2% in 2024 versus 7.8%.

Enterprises experience the highest proportion of malicious MFA events

- In 2024, 11.6% of MFA events associated with enterprise customers exhibited malicious behavior — nearly double the overall average.
- Mid-market organizations (3.1%) once again showed the lowest rate of abuse, while small businesses (6.0%) were right on the average.

MFA is put to the test in the media industry

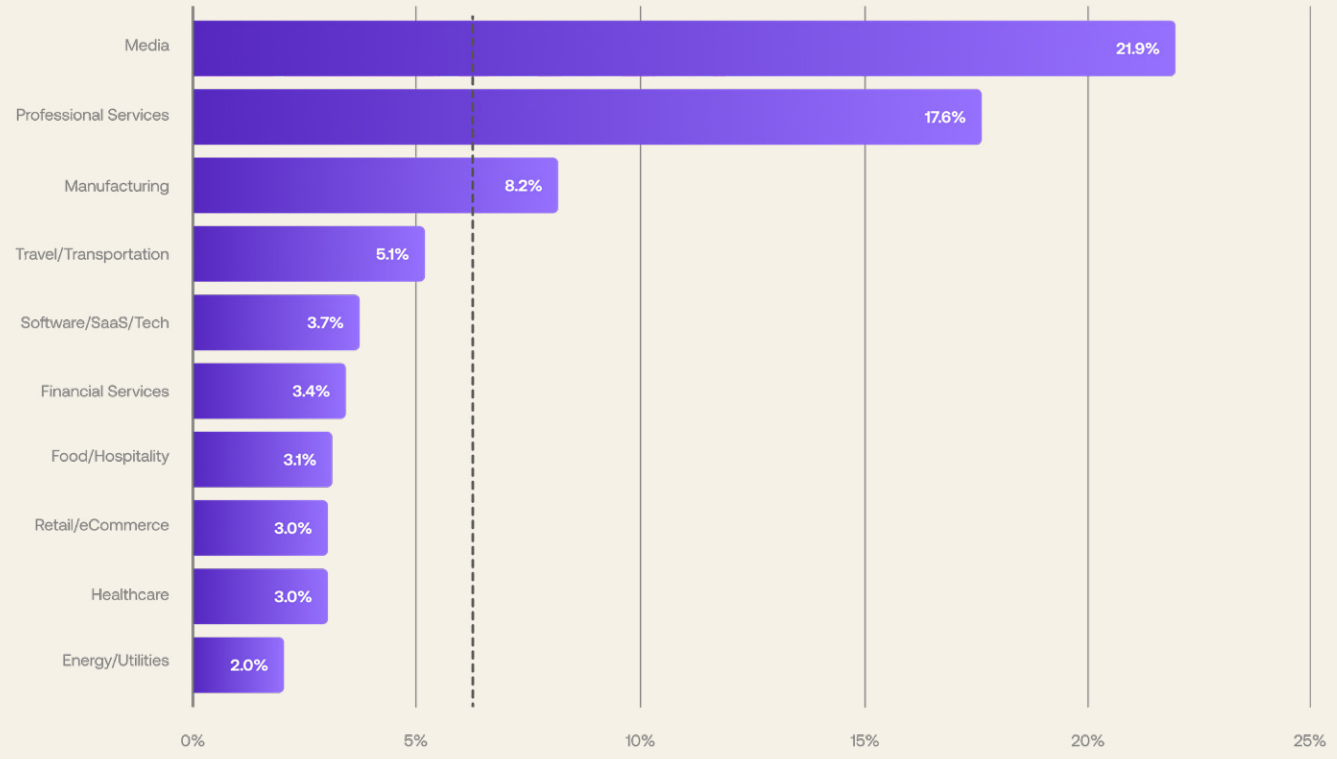

Suspected malicious MFA events, by industry (January 1 to December 31, 2024)

Note: Each bar shows the median daily proportion of MFA events, for a given industry, that met the criteria of malicious MFA activity, in 2024; the dashed line shows the mean (6.0%) of per-industry daily medians across all industries on the Auth0 platform (i.e., not just the top 10)

Attackers don’t take commercial breaks

Sticking with the approach we’ve used prior, let’s zoom in on the Media industry.

One characteristic that immediately stands out is the obvious weekday/weekend variation. In fact, the proportion of MFA events that exhibited clear malicious behavior also follows this pattern:

- Weekdays are consistently in the mid-20% range.
- Weekends fall to below 10%.

Plus — and unlike the industry-specific graphs shown earlier — we can also note the absence of any especially large-scale and/or prolonged attacks.

All of which suggests customer accounts at Media companies are subjected to a fairly consistent rate of login attacks — and MFA defenses are constantly put to the test.

Unpacking media sector MFA events

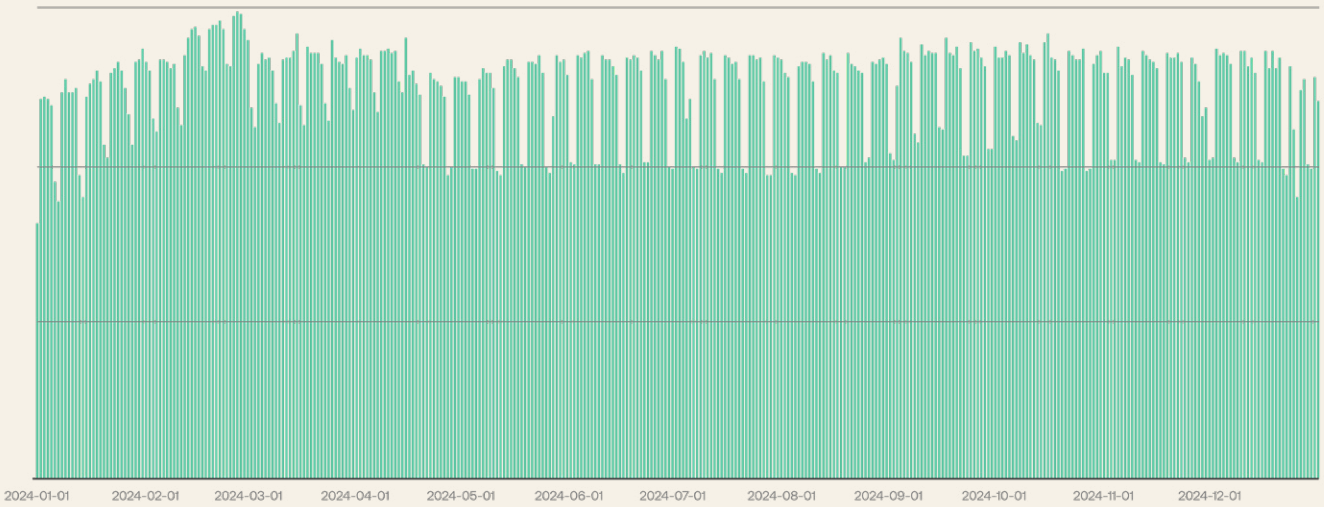

Legitimate MFA events, media sector, by day (January 1 to December 31, 2024)

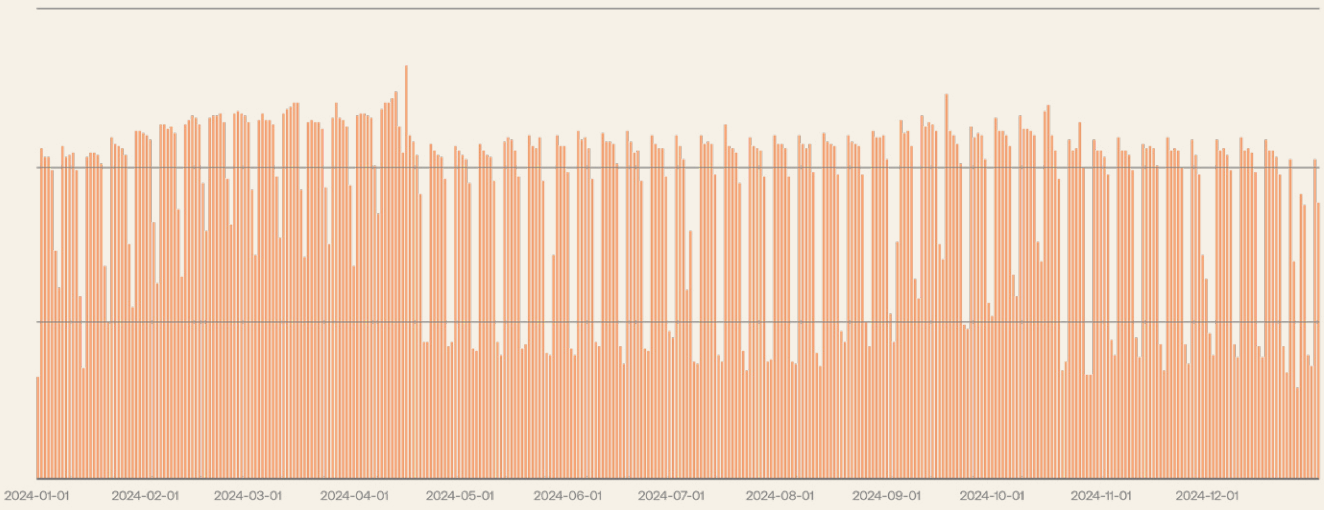

Suspected malicious MFA events, media sector, by day (January 1 to December 31, 2024)

Note: Unlike most other graphs in this report, these two show absolute counts (rather than relative proportions); to enable easy visual comparison between them, these charts use the same logarithmic vertical access, which has been truncated by many orders of magnitude

How to defend against MFA abuse

Given the dangerous and rapidly evolving threat landscape, it's essential that your MFA

- **Is implemented properly:** Gaps and workarounds (e.g., to support legacy authentication or for administrators to bypass MFA) will be exploited.
- **Uses strong secondary factors:** MFA bypass techniques generally target older factors (e.g., those that rely on SMS), and brute-force attacks still focus primarily on knowledge-based authenticators — so using authenticators based on possession or biometric factors can dramatically reduce the likelihood of a brute-force attack being successful.

As already noted, technologies that are effective in consumer applications must balance security and usability. Fortunately, while legacy authentication methods once forced a tradeoff between these two characteristics, the tradeoff is disappearing:

- **Adaptive MFA** is a flexible, extensible MFA policy that can help prevent ATOs without increasing friction for real users. It does so by assessing potential risk during every login transaction, and then prompting the user for additional verification only when necessary.
- **New MFA methods are both more secure and convenient:** MFA methods based on WebAuthn-enabled device biometrics (e.g., Apple Face ID, Apple Touch ID, Windows Hello) or WebAuthn-enabled security keys (e.g., YubiKey, Feitian, Titan) simultaneously deliver more security (threat actors hate WebAuthn) and high usability (recall that survey respondents rated fingerprint and FaceID as highly convenient).

While it remains unlikely that consumers at large will adopt dedicated security keys, biometric capabilities are becoming much more common within affordable devices.

AI meets customer identity

AI agents have arrived. Are users ready to trust them with decisions and personal data?

“Yes, I can help with that. What’s your credit card number?”

ChatGPT’s debut in late 2022 marked a watershed moment in the relationship between humans and AI. A few months later, it reached 100 million monthly active users faster than any technology in history.

Generative AI (GenAI) had gone mainstream, and GenAI agents were quick to follow. These non-human identities (NHIs)

- Execute actions on behalf of users and organizations
- Are autonomous, goal-seeking, and unbound by if/then logic
- Need access to multiple systems to fulfill their roles
- Are becoming common: Our survey revealed that **37% of customers encounter AI agents in online platforms or websites either *very frequently* or *frequently***

The benefits of AI agents are vast, but their actions require sensitive, valuable, and possibly even confidential information from businesses and consumers.

Above all, their actions require trust.

Humans aren't sold on AI agents — yet

The always-available and infinitely scalable nature of AI agents makes them very attractive for companies, who envision a profitable combination of superior customer service and reduced costs.

But are users ready for this brave new world, or do they prefer to handle things the old-fashioned way — by interacting with a fellow human?

It's a simple question, with much hinging on the answer.

Most users still prefer humans

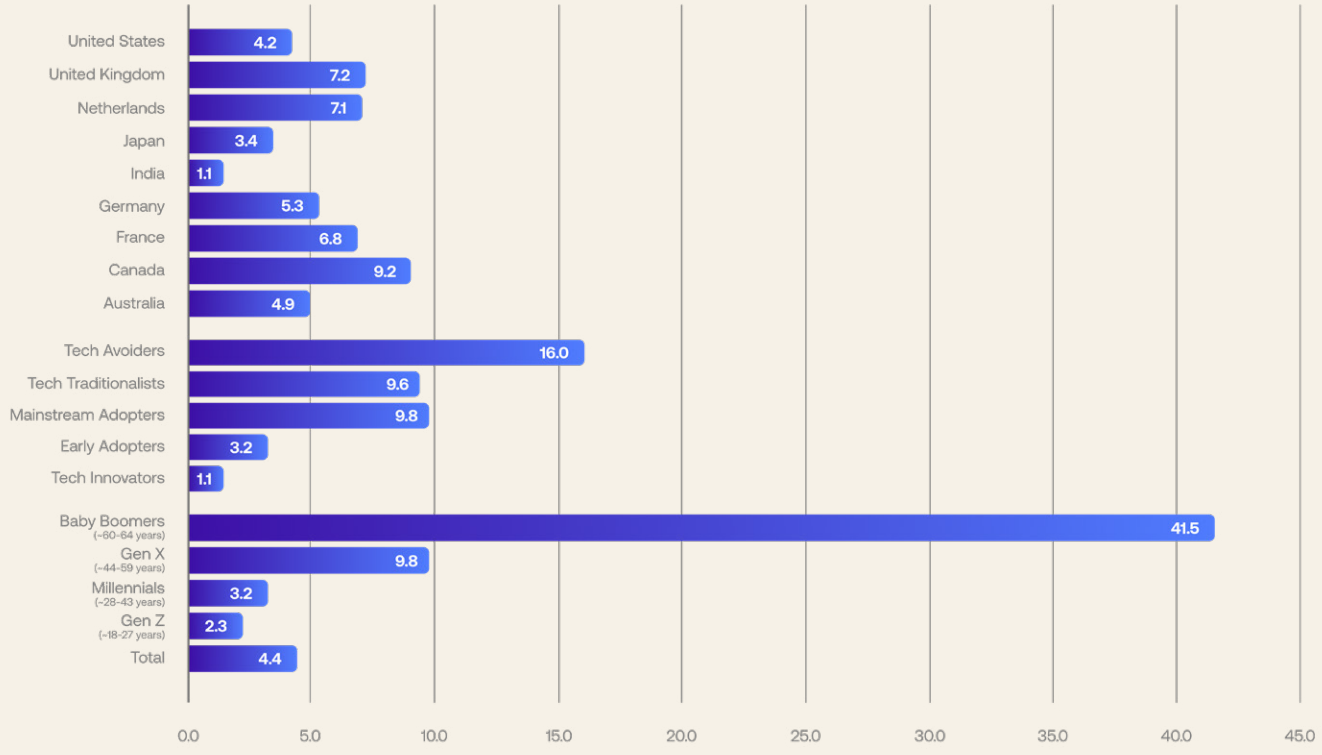

Ratio of preference for interacting with a human agent over an AI agent, all cohorts

"Which of the following best describes your preference for interacting with a company's AI agent versus a human representative?"

Users strongly and universally prefer interacting with human representatives over AI agents

Across the full respondent base, 86% expressed a preference one way or the other, with 70% favoring interacting with humans and 16% favoring AI — a ratio of 4.4 to 1; in fact, every cohort examined in this study preferred interacting with a human representative over a company's AI agent.

Respondent age and attitude toward new technology influence their preferences

Tech Innovators prefer human representatives by a slim 1.1 ratio, while Tech Avoiders are 16 times more likely to do so; similarly, Gen Z users prefer interacting with a human representative by a ratio of only 2.3, while 83% of Baby Boomers prefer humans compared to only 2% favoring AI agents — a ratio of 41.5!

Preferences are fairly consistent by country of origin

Respondents from India (who prefer human agents by a 1.1 ratio) and those from Canada (9.2) bookend the country findings, but also represent outliers — the remaining seven countries all fall between 3.4 (Japan) and 7.2 (the UK).

Why users prefer human help

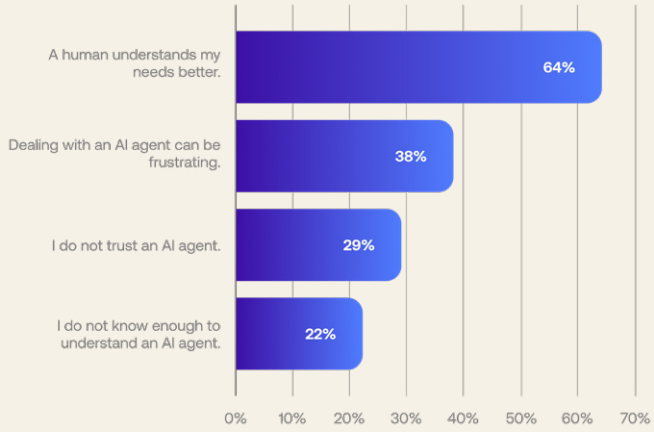

Reasons for preferring to interact with a human, all respondents

"What are your reasons for preferring to interact with a human representative?"

Note: Asked only to those respondents who prefer interacting with a company's human representative

Why users prefer AI agents

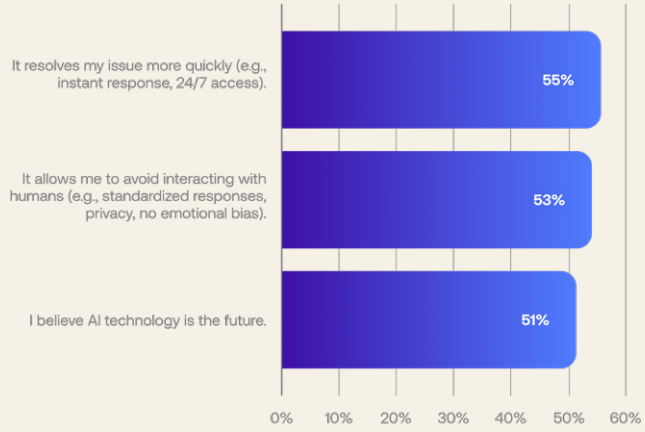

Reasons for preferring to interact with a company's AI agent, all respondents

"What are your reasons for preferring to interact with an AI agent?"

Note: Asked only to those respondents who prefer interacting with a company's AI agent

Users who prefer human representatives do so largely because humans understand their needs

Nearly two-thirds of respondents (64%) who prefer interacting with humans indicated that "*A human understands my needs better*" — pointing to a potential challenge, at least in the short term, for companies rolling out AI agents.

Frustrating experiences and lack of trust in AI agents also contribute to a preference for human representatives

38% of users reported that "*Dealing with an AI agent can be frustrating*" and 29% simply don't trust AI agents — barriers that can be overcome (at least theoretically) with an efficient and transparent implementation.

Users who prefer AI agents cite faster resolutions, the absence of human interaction, and a belief in progress

Interestingly, no single reason stood out, statistically, as being the most common reason why some users prefer interacting with AI agents.

- 55% pointed to quicker issue resolutions — including the around-the-clock availability of automated systems.
- 53% implied that they prefer avoiding interacting with humans — at least when a standardized response is preferred or when there are concerns about privacy or emotional bias.
- 51% are committed to keeping up with the times, and regard AI technology as the future.

Taskmaster, not decision-maker: The AI agent trust line

The potential applications of AI agents are seemingly limitless — but that only matters if customers are willing to use them.

As it happens, not only do users seem hesitant to trust AI agents to perform tasks or activities on their behalf, but the level of trust varies from one task to another.

AI agents: Earning trust one task at a time

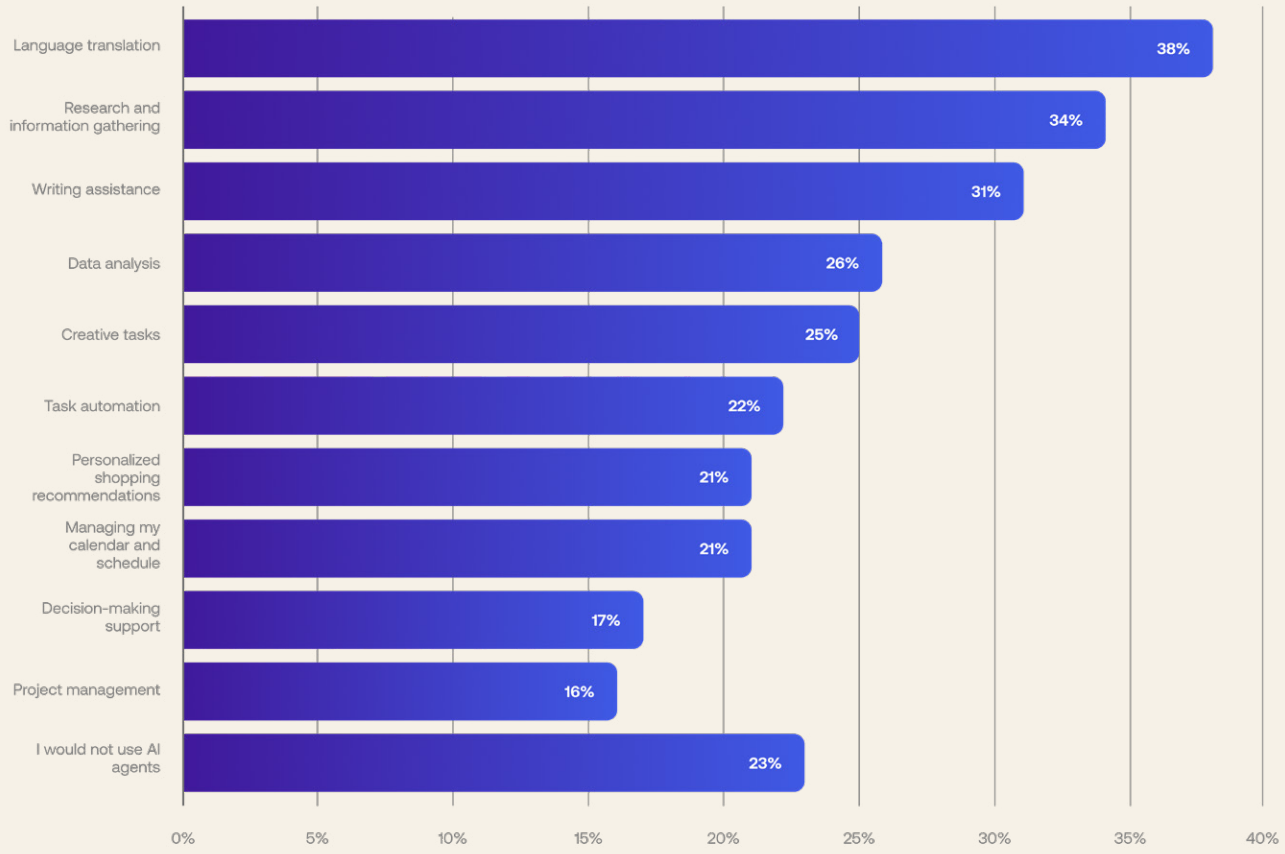

Comfort using AI agents, by task, all respondents

"In which everyday tasks or activities do or would you use AI agents?"

Note: "I would not use AI agents" was an exclusive option

Users remain broadly hesitant to employ AI agents

Before we look at the specific tasks for which customers would use AI agents, it's worth noting that the leading selection among respondents (language translation) was only selected by 38% of users — suggesting hesitance or skepticism on the part of customers.

Users are more likely to employ AI agents for tedious and rules-based tasks

The four tasks for which users are most likely to employ an AI agent (language translation, research, writing assistance, and data analysis) share the common characteristics of being somewhat tedious and objective — quite consistent with how computers have been traditionally utilized.

Users are less likely to employ AI agents for subjective and personal tasks

At the other end of the spectrum, users expressed comparatively little desire to have AI agents handle more personal responsibilities including creative tasks, personalized shopping recommendations, managing calendars/schedules, and providing decision-making support.

The trust gap between users and AI agents

From the chart on the previous page, we can see that **nearly a quarter of users have no intention of using AI agents**, with 23% of respondents selecting the exclusive "*I would not use AI agents*" option.

This number was driven largely by Gen X (32%), Baby Boomers (42%), Tech Traditionalists (43%), Tech Avoiders (a remarkable 73%), and respondents from Japan (37%).

The reasons behind these responses can inform strategies companies should use when introducing and educating about AI agents.

The path to customers using AI agents has its hurdles

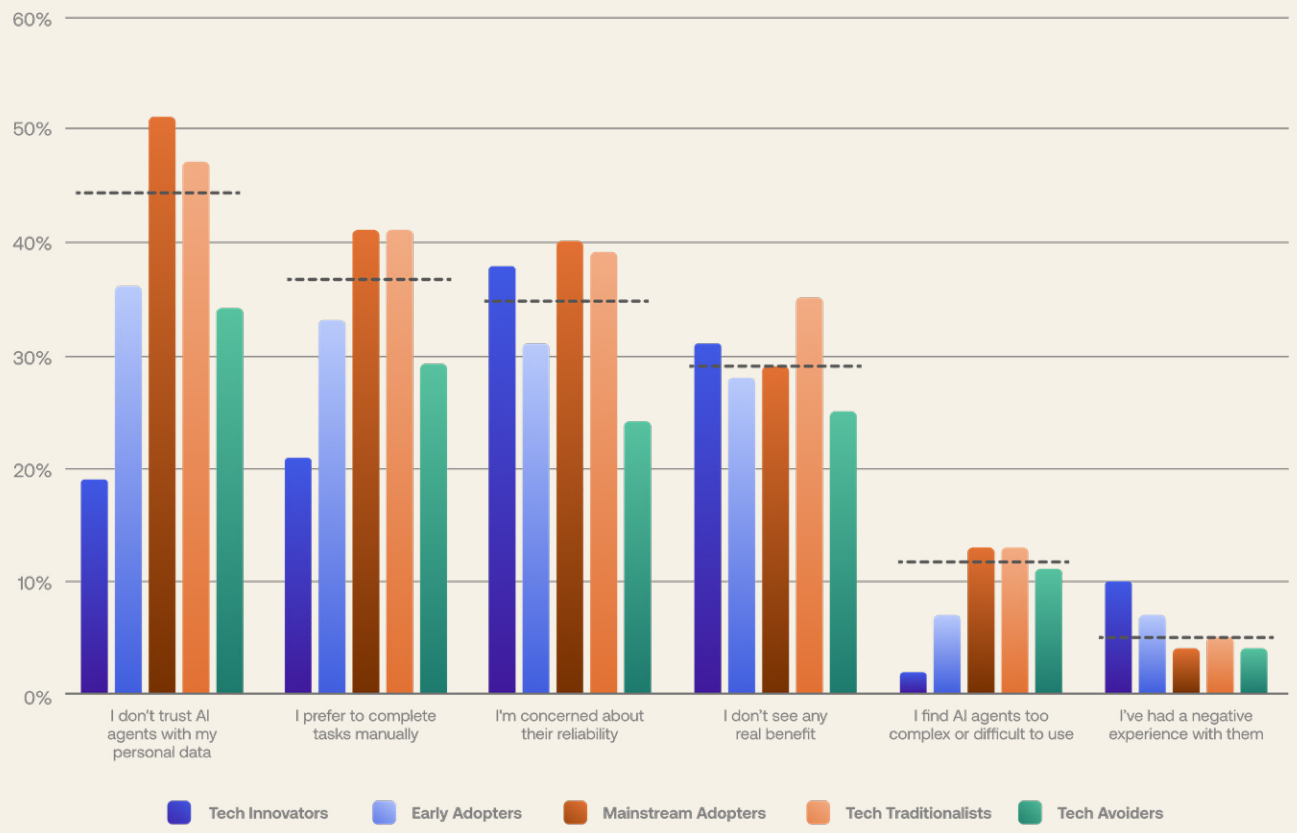

Reasons for not using AI agents, by respondent attitude toward new technology

"What are reasons why you would not use AI agents?"

Note: Asked only to those respondents who indicated "I would not use AI agents"; the dashed lines show the mean across the entire respondent population

Among customers who would not use AI agents, a lack of trust is the primary barrier

The leading reason why customers who don't use AI agents feel this way — selected by 44% of those respondents — is that "*I don't trust AI agents with my personal data*."

Concerns about reliability present a secondary hurdle

More than a third (35%) of customers who won't use AI agents cited concerns about their reliability; as demographic transitions occur, and as implementations mature, this issue should wane.

The good news for service providers is that complexity and negative experiences are essentially non-issues

While the leading issues cited by customers for not using AI agents are somewhat speculative, the survey shows quite concretely that comparatively few users find AI agents too complex or have had a negative experience.

A universal concern: AI's impact on privacy and security

As we've seen, privacy concerns are stopping some users from using AI agents. However, it's important to note that — even if not an outright blocker to use — such concerns are widespread.

The majority of users are concerned: Fully 60% of survey respondents reported being either very concerned or concerned about AI's impact on the privacy and security of their digital identities; in contrast, only 9% of respondents indicated little or no concern.

These concerns are universal: While there is some variation, in every cohort examined in this study, a majority of respondents expressed concern.

Across cohorts, worry about AI privacy and security prevails

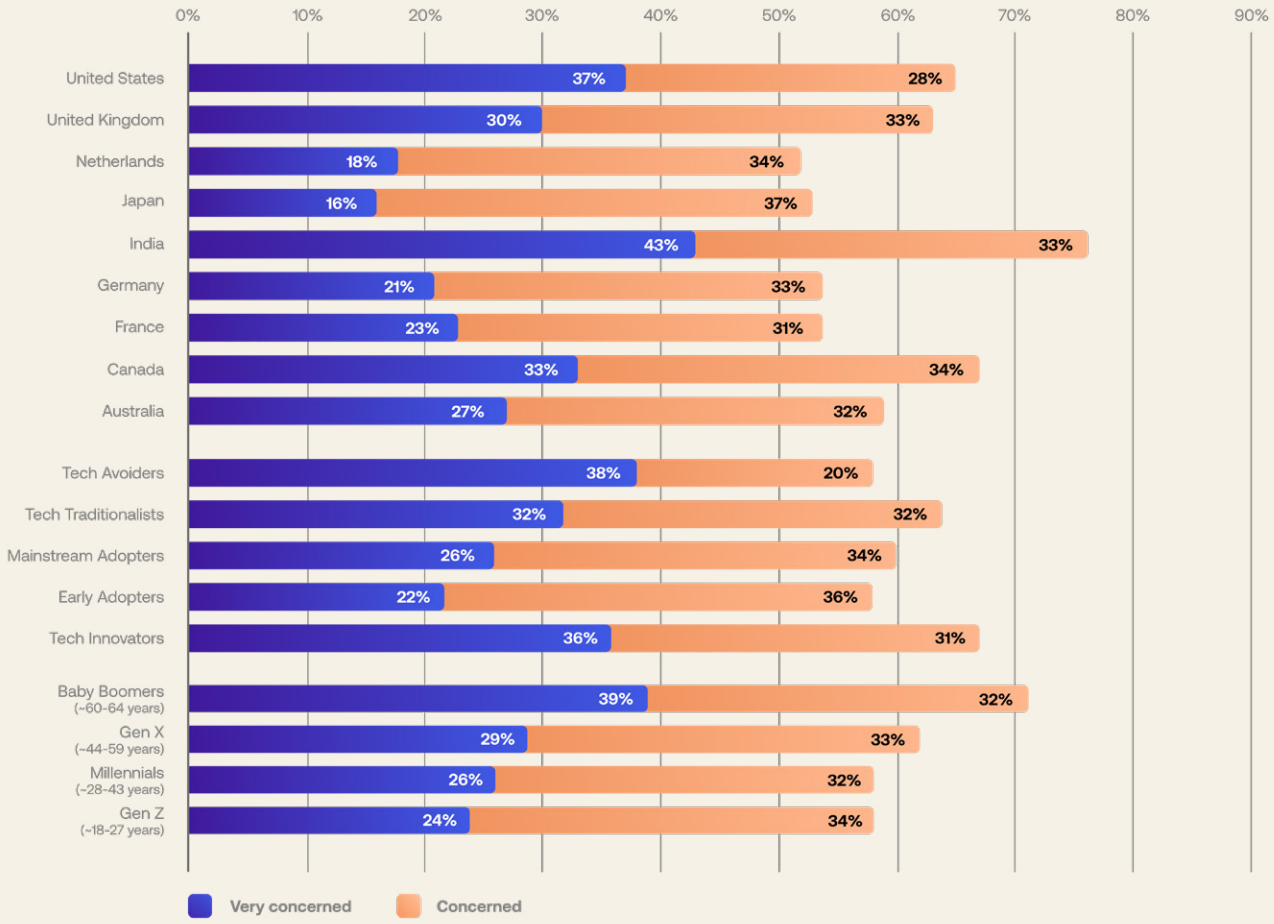

Concern about AI's impact on the privacy and security of digital identity, all cohorts

"How concerned are you about AI's impact on the privacy and security of your digital identities?"

Users' willingness to share personal info with AI agents is all over the map

Taking the long view, AI assistance remains in its infancy. Already, though, many use cases can only be satisfied if the AI agent has access to personal data — whether by asking the user to provide it, or by receiving a user's approval to access data already housed by the AI agent's company.

Willingness to share info with AI agents varies

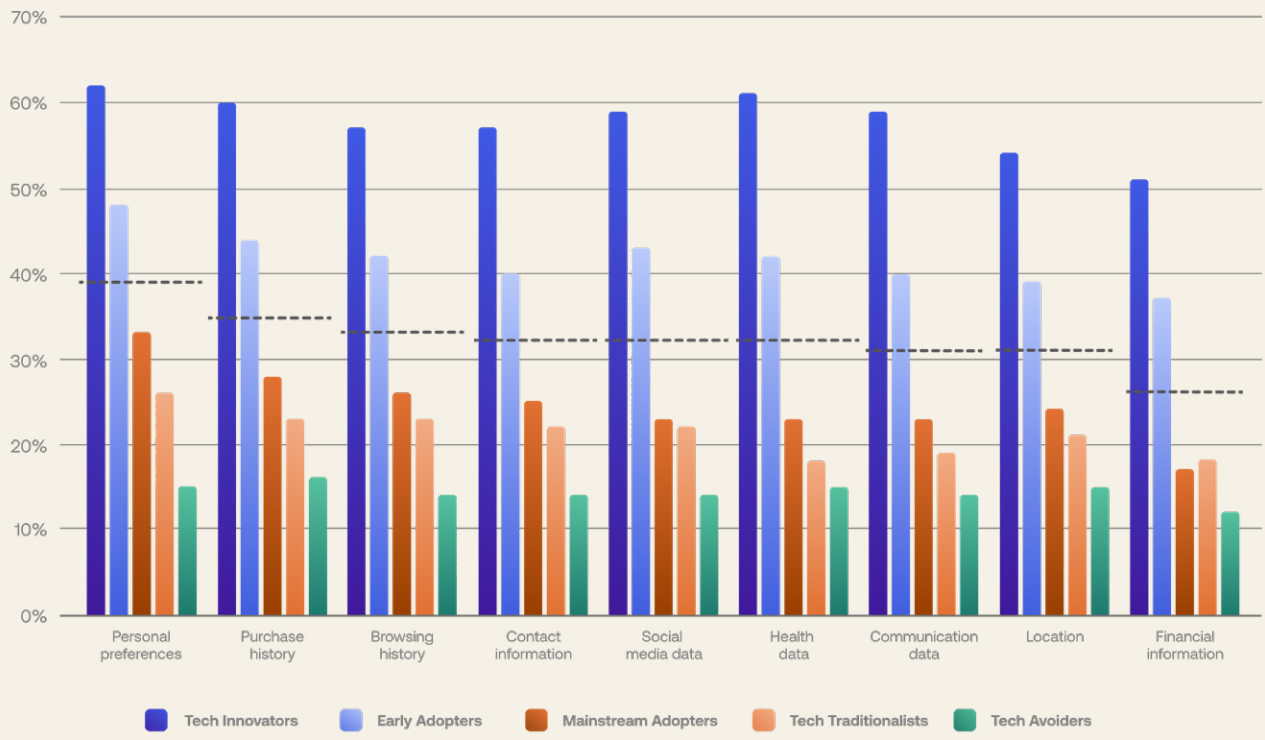

Likelihood of sharing personal information with an AI agent, by respondent attitude toward new technology (sum of "very likely" and "likely")

"How likely would you be to share different types of personal information with a company's AI agent?"

Note: The dashed lines show the mean across the entire respondent population

By and large, users consider themselves unlikely to share personal information with a company's AI agent

Even personal preferences (like a favorite color or sports team) — the most innocuous of the nine options available — are *very likely* or *likely* to be shared by only 39% of survey respondents.

The likelihood varies enormously by age and attitude toward new technology

As the figure vividly illustrates, the likelihood of sharing personal information with a company's AI agent strongly correlates with their attitude toward new technology; a similar, albeit slightly less extreme, pattern is also evident across generations.

Respondents from India are significant outliers

Across all nine types of personal information, respondents from India are nearly twice as likely as those from other countries to share personal information with a company's AI agent.

Trust is the missing feature in today's AI agents

Mutual trust isn't the default state between parties lacking a shared history. Rather, it's typically earned over time.

But AI agents are so new — and, at least to some, entirely unfamiliar — that many users are still forming opinions.

What can organizations do to create a foundation of trust?

Ways to build trust in AI agents

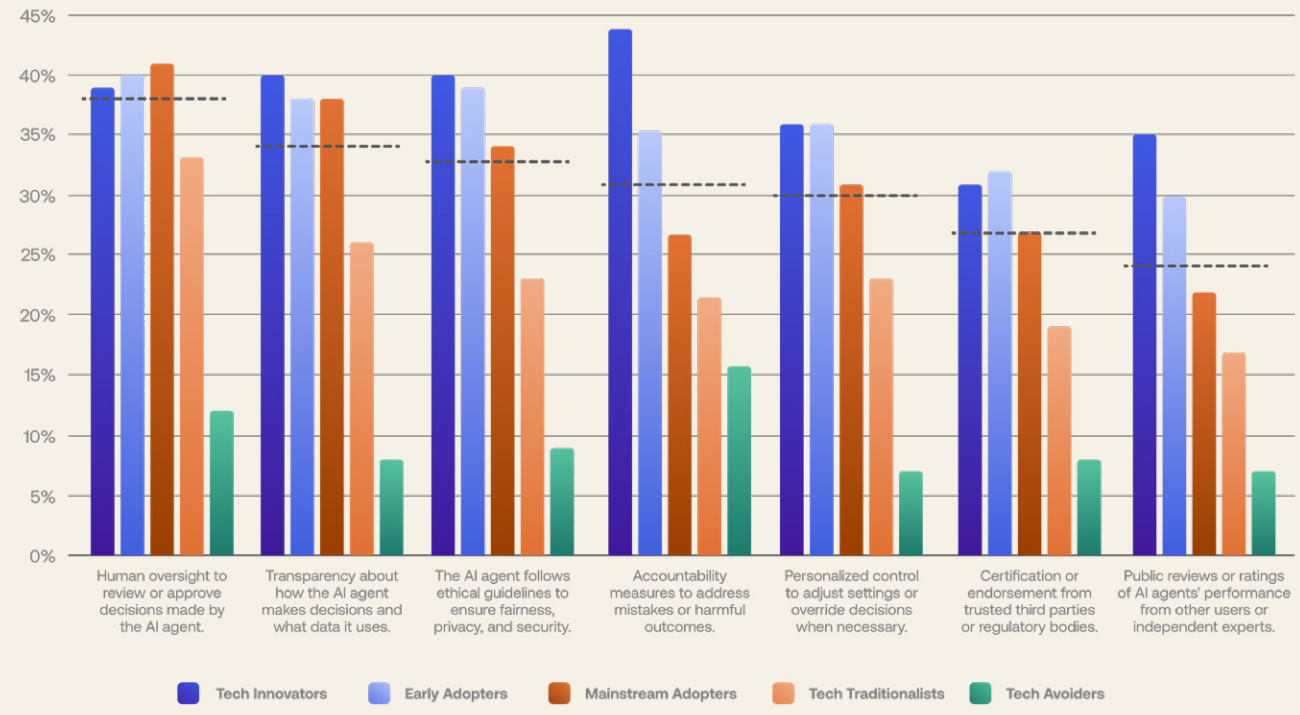

Factors that would increase trust in AI agents, by respondent attitude toward new technology

"What would increase your trust in AI agents to take action or make decisions on your behalf?"

Note: The dashed lines show the mean across the entire respondent population

Customers want humans to remain in the loop

38% of survey respondents indicated that human oversight to review or approve decisions made by an AI agent would increase trust; in fact, this option was the most frequent choice for 14 of the 18 demographic groups examined.

Transparency, ethical behavior, and accountability measures are also popular ways to grow trust

The popularity of these choices points to widespread concern about how an AI agent makes decisions and what recourse is available to users.

Gen Z customers value ethics above all

Gen Z was the only age cohort not to have human oversight at the top of their list; instead, "*The AI agent follows ethical guidelines to ensure fairness, privacy, and security*" led the way with 37% support.

Securing AI agents from the start

Clearly, when it comes to interacting with AI agents, users have trust issues. For many, these issues likely stem from concerns about handing over personal data and the top-of-mind issue of identity fraud.

For what it's worth, these trust issues aren't without justification — unfortunately, in the rush to deploy AI agents, security is being left behind:

- Apps that leverage GenAI, like chatbots or AI agents, employ user interaction and authentication patterns that are different from those used by web and mobile apps.
- As developers are under immense pressure to get AI agents to work, AI applications are being built and deployed without identity and Access Management (IAM) controls

Omissions or oversights can result in AI agents having access to the wrong data, with the ability to perform sensitive actions beyond their intended purpose. Should an attacker find a way to take control, the potential for abuse is immense.

And once these AI agents are live, securing them becomes much more difficult. To help developers secure their AI agents from the start, we've pinpointed four critical requirements where identity is crucial. While these requirements themselves aren't new, they became especially relevant with GenAI apps.

Authentication

For AI agents to operate more securely, they must be able to authenticate users just like any other application.

The agent needs to confirm who the user is before providing access to data or making decisions. This could mean verifying a customer's identity before confirming a purchase, or verifying a patient's credentials before giving access to medical records.

And just like any other app, authentication must be seamless and secure.

Calling APIs

AI agents interact with applications and backend systems on behalf of users, and they leverage APIs to do so.

Without strong identity controls, AI agents could access APIs they shouldn't, leak sensitive data to unauthorized sources, or be completely unable to perform tasks on behalf of users.

To implement such functionality securely, access tokens for the AI agents must be vaulted and secured — not hard-coded.

Asynchronous user confirmation

Unlike traditional applications or a human customer service agent, AI agents don't always complete tasks instantly. Some actions — like data processing, transaction approvals, or decision-making — can take minutes, hours, or even days, which means the AI agent might need to perform a task long after a session has ended.

But security systems today aren't built for long-running, asynchronous workflows. Securing AI agents requires an approach that allows them to authenticate just-in-time, when they need to act, without leaving the door open for attackers.

Authorization

Not every AI Agent should have the same permissions — some should only retrieve data, others should execute commands, and some should make high-risk decisions (e.g., approving a loan or processing a refund).

But without the right access controls, AI Agents can overstep their boundaries.

Just like human users, AI Agents should only get the permissions they need — nothing more. And those fine-grained permissions need to be dynamically updated to reflect changing business rules, compliance requirements, and risk levels.

Conclusion

Establishing trust while keeping pace with changing customer needs

The more things change...

Technology is rapidly evolving, bringing new ways to deliver applications, stronger ways to safeguard those services, and — unfortunately — more and cheaper ways for malicious actors to make things miserable for customers and organizations. Yet, in the face of all this change, some things remain the same: Customers want convenient and secure experiences; they care about their privacy; and they value the flexibility and understanding that only a fellow human can provide.

Crafting trustworthy and convenient experiences

A range of modern authentication methods — including passkeys, social logins, biometrics, adaptive MFA, and step-up authentication — enable signup and/or login experiences that are simultaneously convenient and more secure. Coupled with a restrained approach to gathering first- and zero-party data — and transparency about why personal data is needed, how it will be used, and how it is protected — these modern approaches can help allay customer concerns about identity fraud.

Defending against common identity threats

The most common and largest-scale identity attacks target older, vulnerable forms of authentication, namely passwords and legacy MFA techniques. The simplest and most effective way to safeguard your applications is to adopt — and require — phishing-resistant authentication, with passkeys and biometrics as the most practical examples. Echoing what was said earlier, the goal is not to make your signup and login flows impervious to abuse, but to make them hard enough to abuse that an attacker will decide to find an easier target.

Safely introducing AI agents

At best, customers seem hesitant to embrace AI agents — an attitude largely shaped by concerns about privacy and a perceived lack of options if the agent does something wrong or unexpected. When introducing such functionality, it's important to prioritize security from the start. Pay special attention to the IAM aspects, as AI agents use new and unfamiliar interaction and authentication patterns, and tell your users about how your AI agents are built with security in mind. Finally, don't overlook the value of human agents — if one thing is clear from the survey, it's that users still value that human connection.

Methodology

How we created this report

This edition of the Customer Identity Trends Report draws upon two primary data sources:

1. A global survey of consumers
2. Operational telemetry from the Auth0 platform

Consumer survey

Commissioned by Auth0, Statista conducted a global survey of 6,750 consumers, with 750 from each of nine countries: Australia, Canada, France, Germany, India, Japan, the Netherlands, the United Kingdom, and the United States.

Data was collected in February 2025 using an email invitation and an online survey. All participants were at least 18 years old.

We refer to this survey as “our survey” and “the survey” throughout this report, and refer to the people who completed the survey as “survey respondents” or “respondents.”

When we refer to “cohorts,” we are referring to groups of respondents with a shared:

- Country of residence (nine cohorts);
- Age/generation (four cohorts); or
- Attitude toward new technology (five cohorts).

Respondents by generation

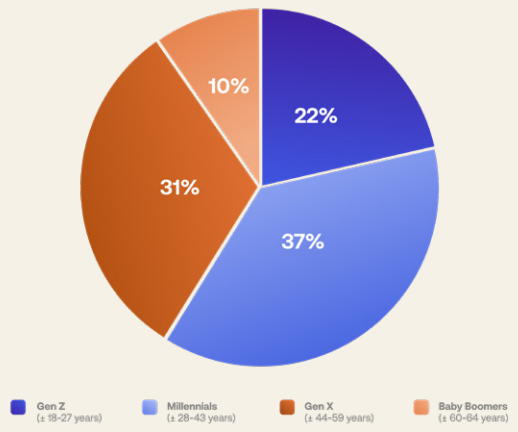

Respondents by attitude toward new technology

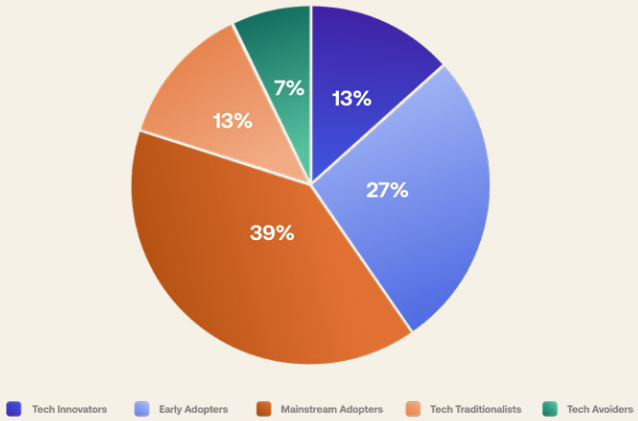

Respondents provided their age, which was used to determine their birth generation:

- Generation Z (~18 to 27)
- Millennials (~28 to 43)
- Generation X (~44 to 59)
- Baby Boomers (~60 to 64)

Respondents were also asked which of the following statements best describes their attitude toward new technology, which was used to assign a cohort label as indicated:

- I actively seek out and try new technologies before others do (Tech Innovators)
- I'm interested in new technologies and try them shortly after they are released (Early Adopters)
- I wait until new technologies are widely adopted and proven to be reliable (Mainstream Adopters)
- I prefer to stick with familiar technologies and rarely try new ones (Tech Traditionalists)
- I avoid new technologies unless absolutely necessary (Tech Avoiders)

Operational telemetry

Portions of this report use operational telemetry from the Auth0 platform, which provides CIAM functionality for thousands of organizations around the world.

For the period of January 1, 2024, through December 31, 2024, (inclusive), the report sums daily event logs associated with legitimate activity and detected threats (see definitions, below), allowing for the meaningful normalization of threat trends and controlling for ongoing changes in customer composition.

Where such information is available, event data is joined with a customer's industry (self-selected) and size (e.g., Small Business, Mid-Market, Enterprise), before being anonymously aggregated. Customers for which certain information is not available are omitted from the associated aggregations.

Because this report is based on real production deployments, it captures the actual activity on the Auth0 platform, and therefore is shaped both by the products and features each customer has enabled (as well as their configurations), and by the evolving capabilities of these products and features.

To accurately portray the true state of identity security, we deliberately chose not to omit outliers or to filter extreme events (e.g., large-scale attacks) from our analysis. However, due to the skewing effect such events can exert on averages, unless otherwise noted, we chose to show the arithmetic median value rather than the mean. When aggregating across days (e.g., for a yearly stat) or across industries, we take the mean.

Signup attacks

We define a **signup attack (or fraudulent registration attack)** as when an individual IP address has 10 or more failed signups in a single day.

A failed signup may result from:

- The identifier (username, email, phone, etc.) is already taken
- Identifier validation errors
- The signup rate limit is being exceeded
- Failure of custom database scripts

Note that the definition of a signup attack does not include user abandonment of a signup form.

Login attacks

Three types of brute-force attack comprise the **login attacks** examined in this report.

- **Credential stuffing:** A threat actor tries known credentials (i.e., from a breach/dump) across other sites and services.
- **Password spraying:** A threat actor tries a comparatively small list of the most common passwords across many different accounts.
- **Password guessing:** In this somewhat cruder approach, a threat actor tries many passwords across any number of accounts

We define a **login attack** as when an individual IP address triggers more than 10 events related to failed logins.

These events include failures such as

- Invalid username or password
- Trying to log in using breached credentials
- The IP is blocked from logging in by other attack protection features (e.g., Brute Force Protection, Suspicious IP Throttling, etc.)

Malicious MFA events

We define **a malicious MFA event** as when an individual IP address triggers 10 or more of the following MFA-related events within an hour:

- Email, SMS, or Push MFA notification sent
- MFA authentication failed or rejected
- A user exceeds the limit for OTP code failures
- A user enters an invalid recovery code or exceeds the limit for recovery code failures

In practice, this definition will capture

- Attempts to bypass MFA via bombing/fatigue attacks (e.g., triggering enough notifications to a user that they approve the challenge simply to stop the onslaught)
- Attempts to commit toll fraud by abusing MFA to trigger phone or SMS messages to premium numbers — driving up costs for the application provider while the threat actor takes a share of the proceeds
- Instances when a threat actor repeatedly fails an MFA challenge
- Instances when a legitimate user repeatedly fails an MFA challenge (albeit these represent a tiny fraction of the MFA abuse count)

How to cite the Customer Identity Trends Report 2025

We love it when people share Customer Identity Trends Report insights. Here's how to properly cite data, statistics, and any other information found in the Customer Identity Trends Report 2025:

- **Give us credit**
  Please cite the source as "Auth0 Customer Identity Trends Report 2025" when referencing any content.
- **No modifications**
  Content must be cited exactly as it appears in the report. If you wish to paraphrase, please contact us for approval.
- **Please share**
  If you'd like to share the report with others, please provide a link to our download page:
  auth0.com/customer-identity-trends-report

We appreciate your helping us keep our insights accurate and accessible to everyone.