Trust and transparency in Al-driven CX

A practical guide for customer experience leaders





Trust is the new currency in CX

Al is driving CX advancements, with 70% of CX leaders recognizing it as a business imperative. From 24/7 Al-powered chatbots, intelligent routing to streamline interactions, and sentiment analysis that enables personalized responses, Al isn't the future – it's here and now. But as Al takes on more customer-facing tasks, one thing matters more than ever: trust.

While AI promises efficiency and scale, many CX leaders are still navigating its complexity, unpredictability, and risks. The gap between what AI can do and what customers are comfortable with is real – and widening.

While most use AI, less than half of people (46%) globally are willing to trust AI systems. And the problem is only growing: people have become less trusting and more worried about AI as adoption has increased.

Trust is the foundation of great CX. Without it, customers opt out, churn, or worse, lose faith in your brand.

In this guide, we'll share strategies to help you build ethical, transparent, customer-first AI systems that deepen loyalty, reduce risk, and unlock new opportunities.

What you'll learn:

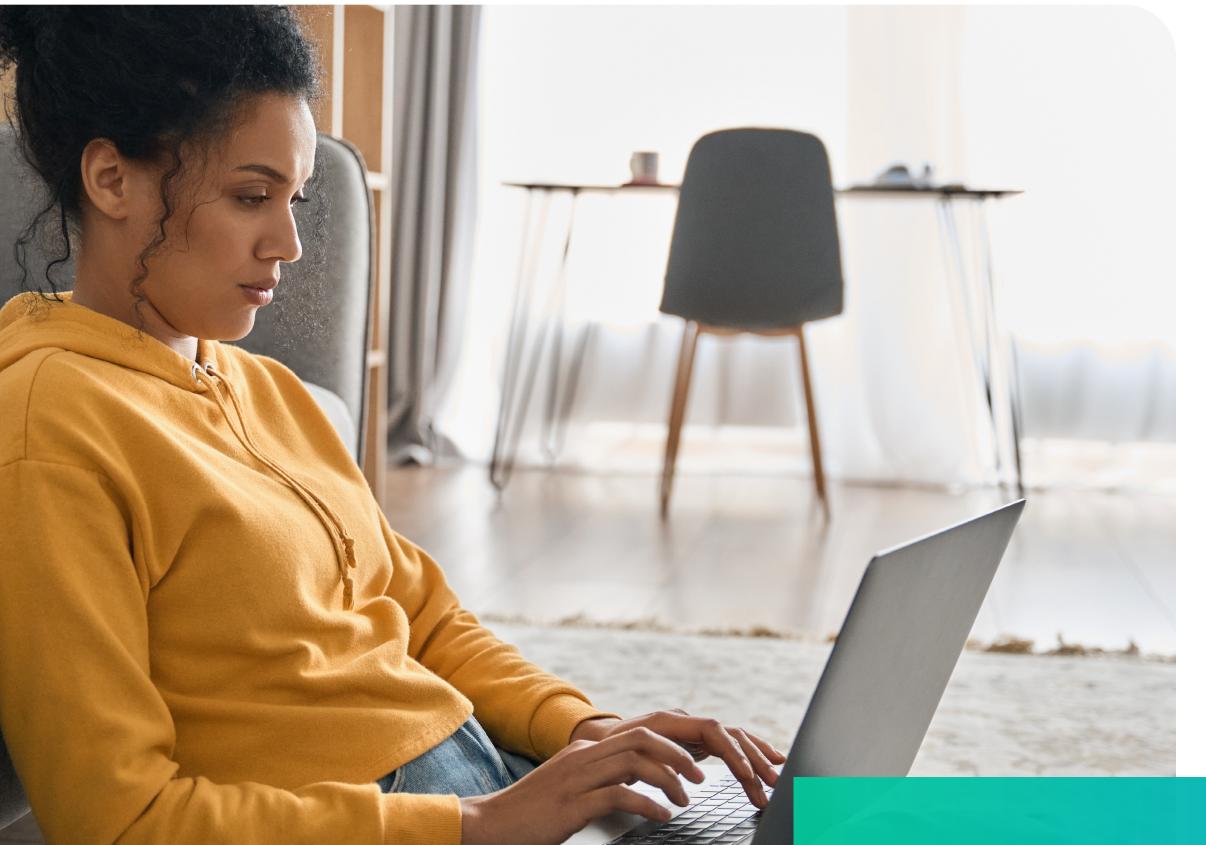
- Why trust is essential for Al adoption in CX
- Common mistakes that erode customer trust
- Step-by-step tactics to build transparency and ethical AI
- How human-guided AI builds lasting confidence
- Real-world examples and next steps





Less than half of people globally are willing to trust Al systems

Why trust is essential in Al customer experiences



Whether customers realize it or not, they're engaging with AI every day – through intelligent chatbots, virtual voice assistants, personalized product recommendations, and predictive support flows that anticipate their needs before they even articulate them.

From resolving billing issues to getting help navigating a healthcare portal, AI is showing up at critical moments in the customer journey. But while the technology is advancing, customer confidence hasn't fully caught up.

According to Salesforce's latest <u>State of the</u> <u>Connected Customer report</u>, 60% of consumers say that the rise of Al makes trust more important than ever before. And in a <u>Genesys poll</u>, 64% of consumers said they believe Al will improve the customer experience, but only if it's used responsibly.

That last part is key. When customers don't trust how AI is being used, the consequences are immediate and damaging:

- They opt out of automated flows or abandon self-service tools entirely.
- Skepticism grows, making them more likely to question outcomes, request human support, or even disengage from your brand.
- Worst case, a breach of trust leads to reputational damage, public backlash, or regulatory scrutiny.

Why does this happen? Because trust in Al isn't automatic. Customers need to know when they're interacting with Al, how it's making decisions, and that they can talk to a real person if something goes wrong. Without transparency, control, and accountability, Al can feel like a black box – and that uncertainty erodes confidence.

On the flip side, when companies design AI experiences that are clear, ethical, and backed by human support, they not only avoid pitfalls; they build brand equity. Trust becomes a differentiator. And in a world where products and pricing can be easily replicated, trust is one of the few competitive advantages that truly lasts.

2

64% of consumers said they believe Al will improve the customer experience, but only if it's used responsibly

Common trust breakers to avoid

Even well-intentioned AI can erode trust when it's not thoughtfully deployed. These common missteps may seem small in isolation, but together they create an experience that feels impersonal, disjointed, and even deceptive. Let's unpack the most frequent trust breakers and how to recognize them before they damage the customer relationship.



1. Al use is hidden

Customers often don't realize they're interacting with AI until something goes wrong. Maybe the bot loops endlessly or gives irrelevant answers – only then does it become obvious that no human was ever involved.

Why it matters:

When customers discover AI "by accident," it feels sneaky. That perception undermines confidence, even if the technology is working well behind the scenes.

2. No clear human fallback

If there's no clear way to reach a human agent – or if it takes jumping through hoops – customers feel trapped.

Why it matters:

Failing to offer a human option signals that the company values efficiency over empathy. And when issues get complex, automation alone can't always resolve them.

3. Bias and opacity

Customers often don't understand how AI makes decisions – especially if those decisions feel biased, unfair, or inconsistent. Whether it's a denied claim, rejected application, or skewed recommendation, a lack of transparency quickly breeds mistrust.

Why it matters:

Al systems trained on incomplete or skewed data can reinforce real-world biases, and if customers can't get a clear explanation, it undermines fairness and accountability.

4. Rigid automation

Sometimes automation is too strict. Customers get funneled through one-size-fits-all flows that don't accommodate unique needs or edge cases.

Why it matters:

When people feel like they're talking to a machine that doesn't listen – or worse, one that blocks their path – it causes frustration and drives them to abandon the interaction entirely.

5. Bots pretend to be human

Some AI systems are designed to sound like real people, and many customers don't realize they're talking to a bot until something goes wrong. Conversational AI isn't the problem, but it should be clear when a customer is and isn't talking to a real human.

Why it matters:

Misleading customers about who (or what) they're speaking to violates psychological safety. Once revealed, it creates a breach of trust that's hard to repair.

Each of these pitfalls introduces friction, confusion, or a sense of betrayal – none of which help your brand. But the good news? Transparency, control, and human connection are powerful antidotes. When companies are upfront about where and how AI is used, offer clear off-ramps to human help, and design with empathy, trust becomes not just possible, but a competitive advantage.





How to build trust with ethical, transparent AI CX

Al has the power to make customer experiences faster, smarter, and more personalized – but only when customers feel secure, respected, and in control. Ethical design isn't just about compliance; it's a chance to build lasting trust and show customers that you put their experience first.

Here's how to design Al-powered interactions that feel supportive, not suspicious.





Be transparent about Al-driven conversations

Transparency is the foundation of trust. Customers should never have to guess if they're speaking with a bot or a human.

What to do:

- Clearly label AI assistants in chat, voice, or email interactions.
- Let customers know when and why Al is being used (e.g., "Our assistant can help you faster during peak hours.").
- Reinforce the benefit: explain how AI
 helps them get answers quickly, avoid
 wait times, or receive personalized
 suggestions.

Why it matters:

When customers know what they're dealing with, they feel respected, and are more likely to stay engaged.



One of the fastest ways to erode trust is trapping users in automated loops.

Customers should know they can reach a real person if needed. These <a href="https://hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hybrid.com/hy

What to do:

- Make the "talk to a person" button or link highly visible – not hidden in a menu.
- Use intelligent escalation rules to detect frustration signals like repeated inputs, long pauses, or negative sentiment and proactively offer human support.
- Don't penalize customers for choosing human help (e.g., by making them repeat information).

Why it matters:

Knowing a human is just a click away reduces anxiety and improves tolerance for automation.

3 Embed bias checks in Al

Al systems learn from data, and if that data reflects real-world biases, those biases can show up in how customers are treated.

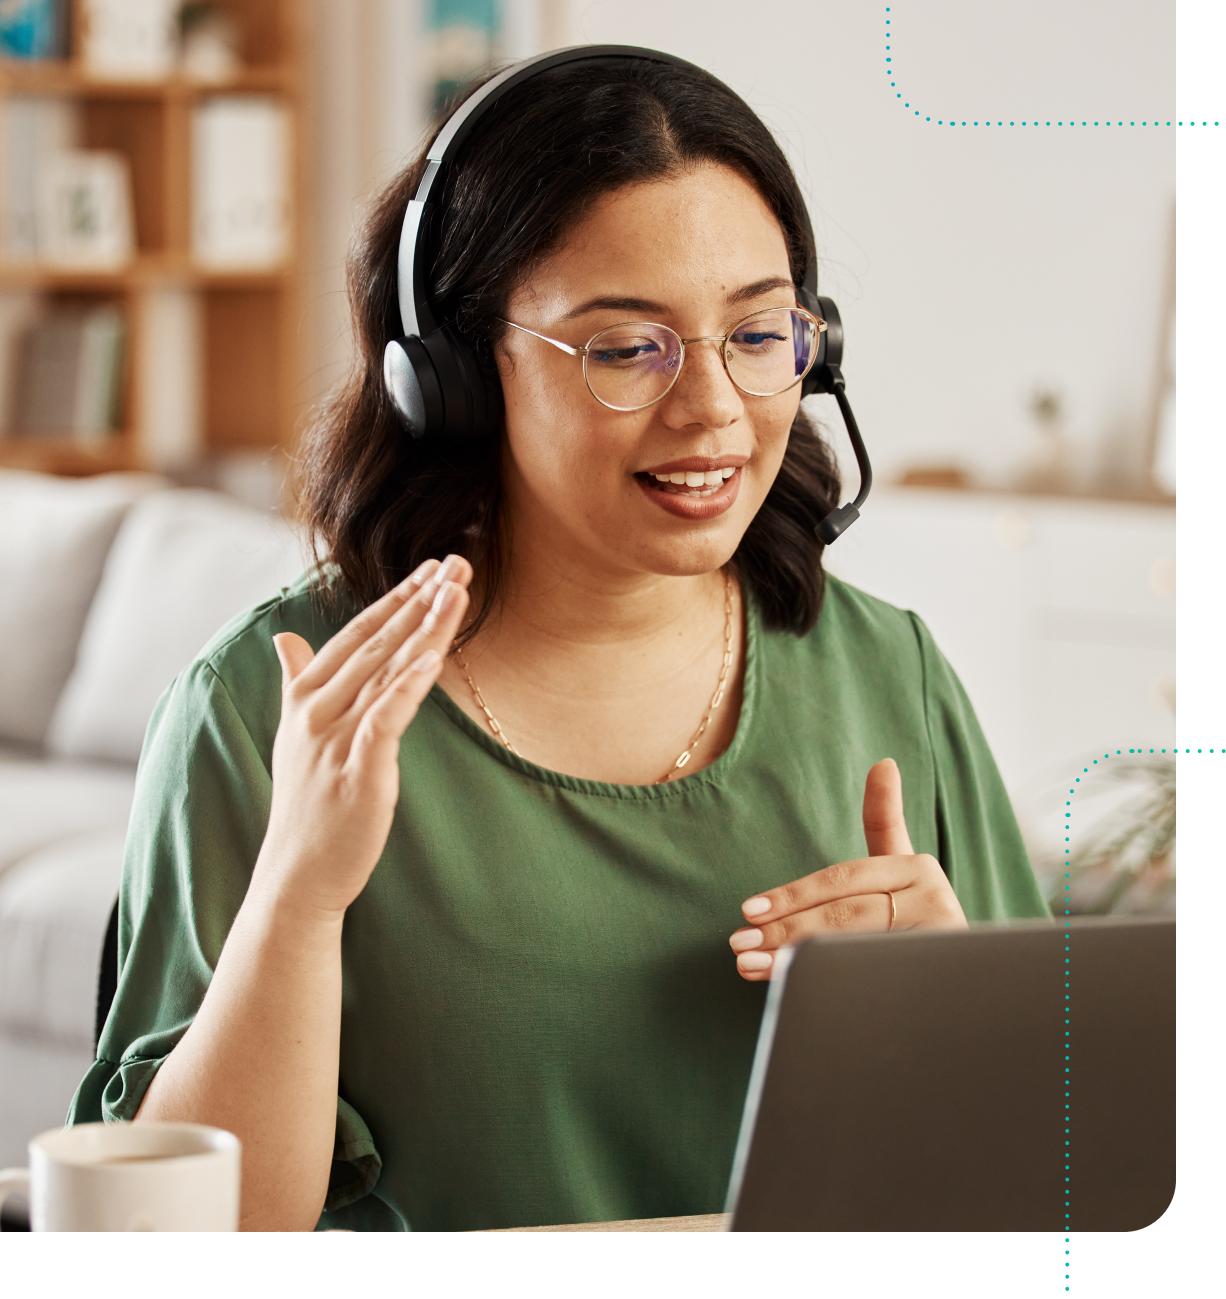
What to do:

- Train AI on diverse, representative datasets.
- Before deploying a chatbot, test its responses across five common customer personas. Include accents, tone variance, and emotional states to flag unintended bias or rigidity.
- Run regular fairness audits and test for biased outcomes across gender, race, age, location, and language.
- Involve ethics and compliance leaders in your model evaluation processes (not just your technical teams).

Why it matters:

Biased AI can alienate users and expose your company to legal or reputational risk. Proactively addressing it in your AI governance framework builds fairness and confidence.





4 Train agents for Al escalations

Al and human agents should work together seamlessly. When handoffs happen, the transition should feel smooth – not like starting over.

What to do:

- Educate frontline staff on what the AI does, how it makes decisions, and what common issues to look out for.
- Empower representatives to override or course-correct AI decisions when needed.
- Include AI tools and workflows in onboarding and ongoing training so agents feel confident, not blindsided.

Why it matters:

A well-informed agent can salvage a bad Al experience and restore trust in seconds.

5 Prioritize data ethics, transparency, and compliance

Al runs on data – but in industries like finance and insurance, that data must be handled with precision, transparency, and regulatory compliance.

What to do:

- Clearly communicate what data is collected, how it's used, and who can access it especially when AI is involved.
- Offer customers meaningful control, such as opting out of certain uses, adjusting personalization settings, or requesting data deletion.
- Align with applicable regulations (e.g., GDPR, CCPA, HIPAA, GLBA) and industry standards for data handling, encryption, and secure storage.
- Review AI models for bias or misuse of sensitive personal data, especially when used for decision-making.

Why it matters:

Responsible data practices are a regulatory requirement. Customers expect transparency and control, and regulators are watching. Building Al-powered experiences on a foundation of ethical, compliant data use protects your business and builds long-term credibility.

Building trust with human-guided Al

Al is powerful, but it's not perfect – and customers know it. That's why the most effective Al strategies aren't fully automated. They're humanguided. When Al is combined with real-time human support, the result is a more transparent, empathetic, and trustworthy customer experience.

Think of it as the best of both worlds: AI handles routine tasks quickly and efficiently, while human experts step in for high-stakes moments, emotional conversations, or complex decisions.

When customers can **see what's happening**, ask questions, and interact with a real person – even after AI has done its part – they feel informed, heard, and in control.

At the end of the day, AI becomes less intimidating, and more helpful, when customers know there's a human there to help. It's not just about automation. It's about assurance. And that's what builds real, lasting trust.

Human-guided AI in practice





A **chatbot** answers simple questions and triages requests, but when an issue gets more complicated, it hands off to a live agent who **cobrowses** with the customer to solve the problem together.





An Al tool suggests a **loan product**, but a human rep walks the customer through the terms via **video call**, helping them understand the fine print and make a confident decision.



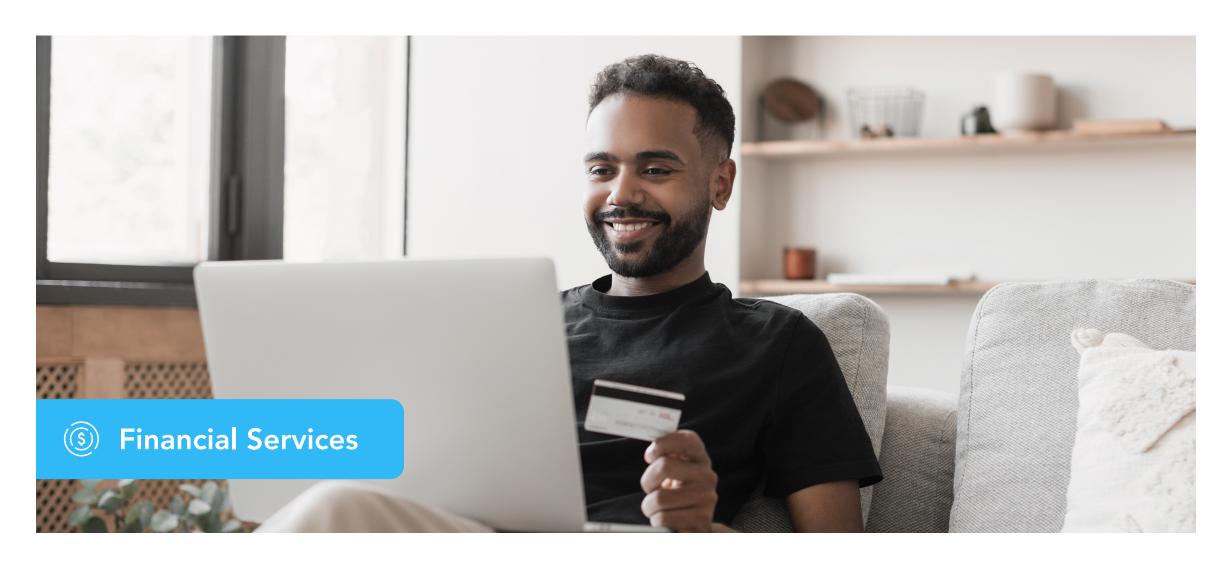


A **policy change** is triggered automatically based on customer behavior, but an agent joins in real time to explain what's happening and why, preventing surprise or confusion.

Tips to use humanguided AI effectively

- Let Al surface common issues or suggest solutions, but ensure agents are ready to guide resolution with empathy and expertise.
- Keep humans in the loop for anything requiring context, judgment, or emotional intelligence.
- Use visual tools like cobrowsing and video to simplify digital tasks, walk customers through complex workflows, and foster connection.
- Train agents to work with AI, not around it. When agents understand how AI fits into the journey, they can step in seamlessly when needed.

Real-world success stories





Transparent AI in action

UK-based bank, NatWest, uses an Al-powered digital assistant named Cora to answer bank queries 24/7. Using Generative Al, Cora can handle a variety of inquiries. However, when intervention from a customer service agent is needed, the customer's conversation with Cora is summarized and provided to the human agent, enabling them to quickly understand what support is needed.

This seamless handoff ensures the customer doesn't have to repeat themselves. It's also very clear to the customer when they are talking to an AI chatbot and when they're talking to an agent. And customers are responding positively: Cora handled 11.2 million customer conversations in 2024, and has helped drive a 150% improvement in customer satisfaction levels.

Al supports cancer patients

Unfortunately, doctors can't be by a patient's side 24/7. But a chatbot can. Memorial Sloan Kettering Cancer Center uses an Al chatbot to support chemotherapy patients. The bot gives personalized tips for managing side effects, tracks symptoms, and – here's the key – alerts doctors when something seems off.

This approach combines the convenience of AI with the care of real clinicians. Patients get quick help when they need it, and providers step in when it matters most. This has resulted in <u>fewer ER visits</u> and <u>higher patient satisfaction</u>.

Even if Al can "do it all," trust comes from knowing someone's there when it matters.



The bottom line: Ethical Al is good business

Trust is the difference between an Al-powered journey that delights and one that derails.

When you design customer experiences that are transparent, ethical, and human-guided, you:

- Reduce churn and complaints
- ✓ Increase loyalty and satisfaction
- Protect your brand's reputation

Al doesn't replace empathy. It amplifies it – when guided well.

The future isn't Al-only. It's Al-with-humans – designed with clarity, consent, and backup built in. At Glance, we help companies create **transparent**, **human-first Al customer experiences** through secure cobrowsing, video, and guided CX tools. Trust is a top priority to us, which is why **we adhere to strict security guidelines**.

Learn more and schedule a demo at Glance.cx





www.glance.cx

1 888 945 2623 +1 781 316 2596 (International) hello@glance.cx