Utel

Utel AI and Security Report

2026

Criminals have upgraded. Have you?

From our position monitoring tier-1 operators across numerous countries, we watched fraudsters adapt faster than defensive measures in every case.

This isn't theory; these are attacks we have found, patterns we have detected, and criminal adaptations we have observed in real-time.

This report is written for decision-makers and operators responsible for protecting networks, revenue, and trust. It examines real-world cases where traditional approaches fall short and draws on operational intelligence to anticipate the security landscape of 2026.

In 2025, we experienced incidents that will shape how we view security going forward. In this report we will go into detail on four of them.

These incidents reflect a broader industry reality. Operators reported a rise in fraud attempts that increasingly evaded detection by blending into legitimate traffic and exploiting trusted identities and familiar patterns.

It is this change that will define the security challenges of 2026.

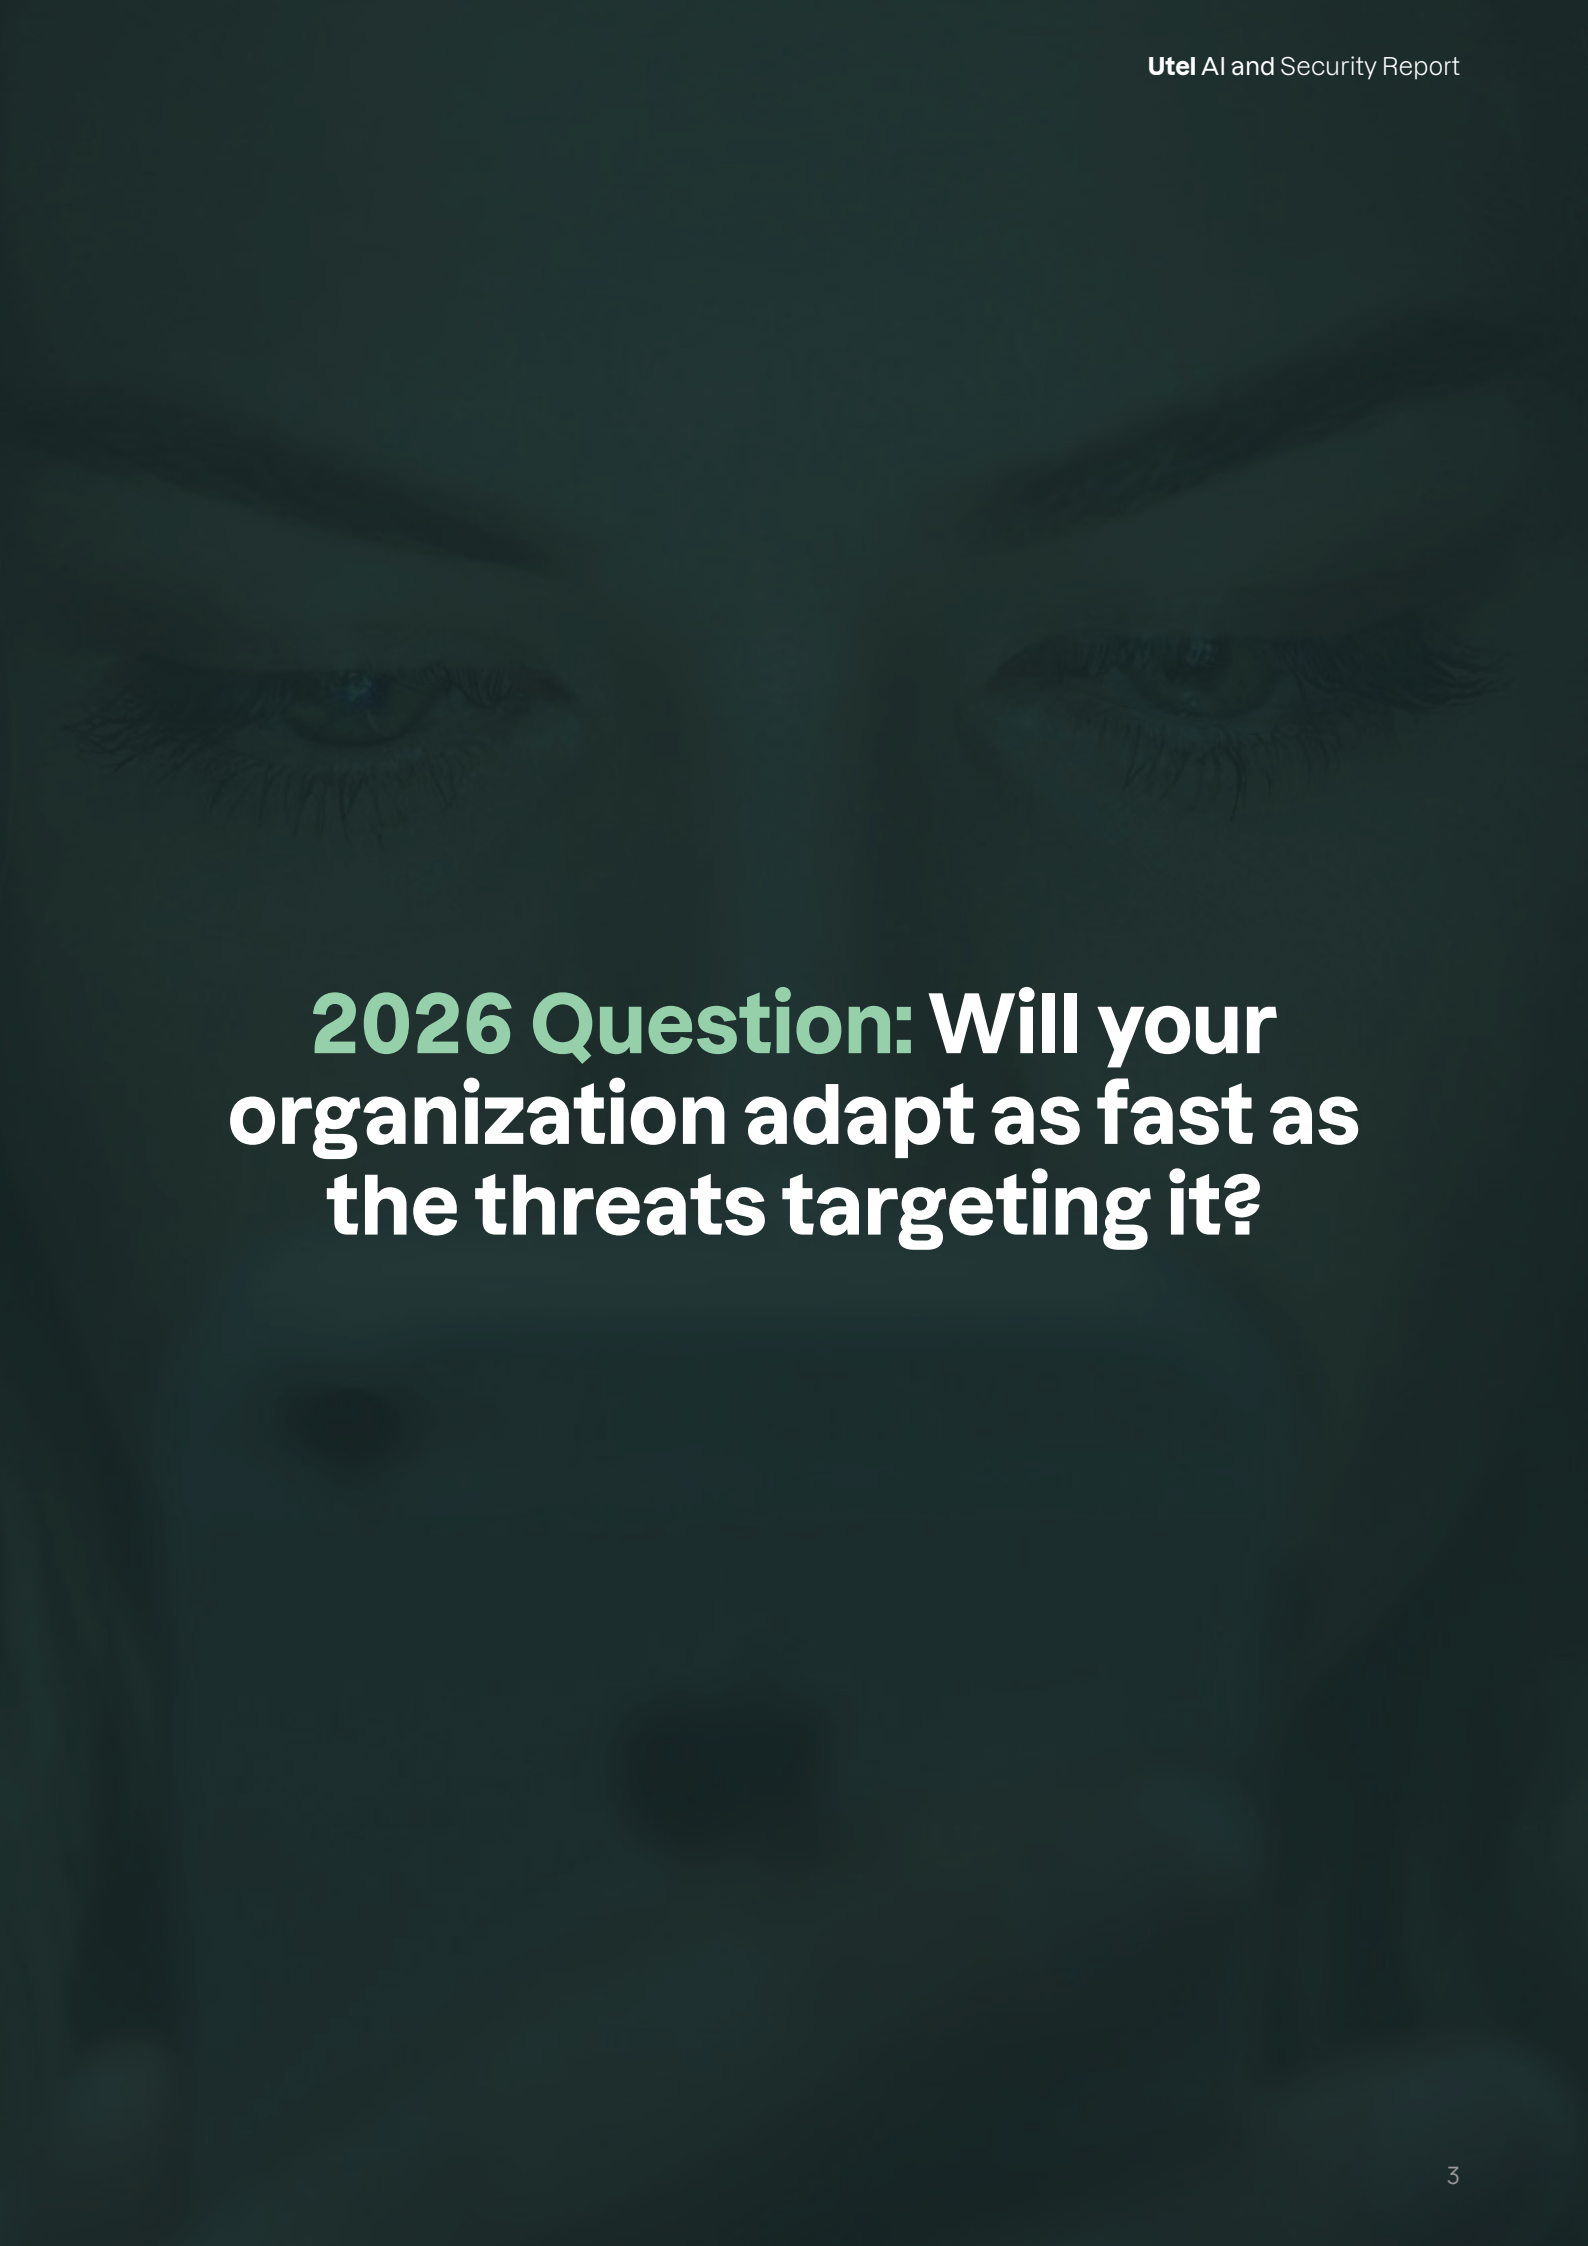

2026 Question: Will your organization adapt as fast as the threats targeting it?

Introduction

In 2025, we watched the fundamental assumption of telecommunications collapse: that you can trust who you're talking to. From our position, monitoring traffic across multiple countries and tier-1 operators, we see fraudsters adapt faster than defensive measures.

The Norwegian government's "digital shield" quickly stopped millions of spoofing attempts, but was simultaneously bypassed by some fraudsters within a couple of weeks.

When we deployed AI-based detection for a Central European operator, it uncovered a large money laundering operation within 72 hours that had been running undetected for months.

This report is based on operational experience and practical insights, rather than compliance requirements or external analysis. We made it to provide real insight, not general observations.

The cases in this report are threats and patterns we have actually detected and blocked, and the adaptations we have observed criminals make over the past year.

Are you ready for the threats of 2026?

We know how communication work.

As a Norwegian-based company, we draw on a long national tradition of being ahead of the curve in technology, security, and defense against fraud and crime.

Years of analyzing communication patterns across telecom networks have shaped how we see risk, equipping us to recognize new behaviors as they begin to form. That accumulated experience underlines this report, offering insight grounded in long-term observations of how communication evolves.

We all know that by now, AI has been implemented in many aspects of our everyday lives. AI agents and automated workflows are embedded in business and now shape how we work, make decisions, and communicate.

AI has enhanced performance but has also lowered the barrier to sophisticated fraud at a global scale. On average, fraud costs companies 3% of their total revenue[1].

Fraud is not only a security and financial issue. Fraud directly affects trust. It affects how people trust institutions, how customers trust brands, and how we all trust communication. This report is written because 2025 made it clear that trust can no longer be assumed.

When people cannot trust who they are speaking to, what they are seeing, or what messages are actually legitimate, the impact extends beyond financial loss into decreased trust in systems, institutions, and each other.

In 2026, organizations and regulators face a shared responsibility to rethink how trust is built, verified, and protected in an AI-driven world. Security is no longer only about stopping attacks but about preserving the conditions that allow our digital lives to function.

That's what 2025 broke and what 2026 must rebuild.

1 TransUnion (2025). Telecommunications Industry Fraud Report.

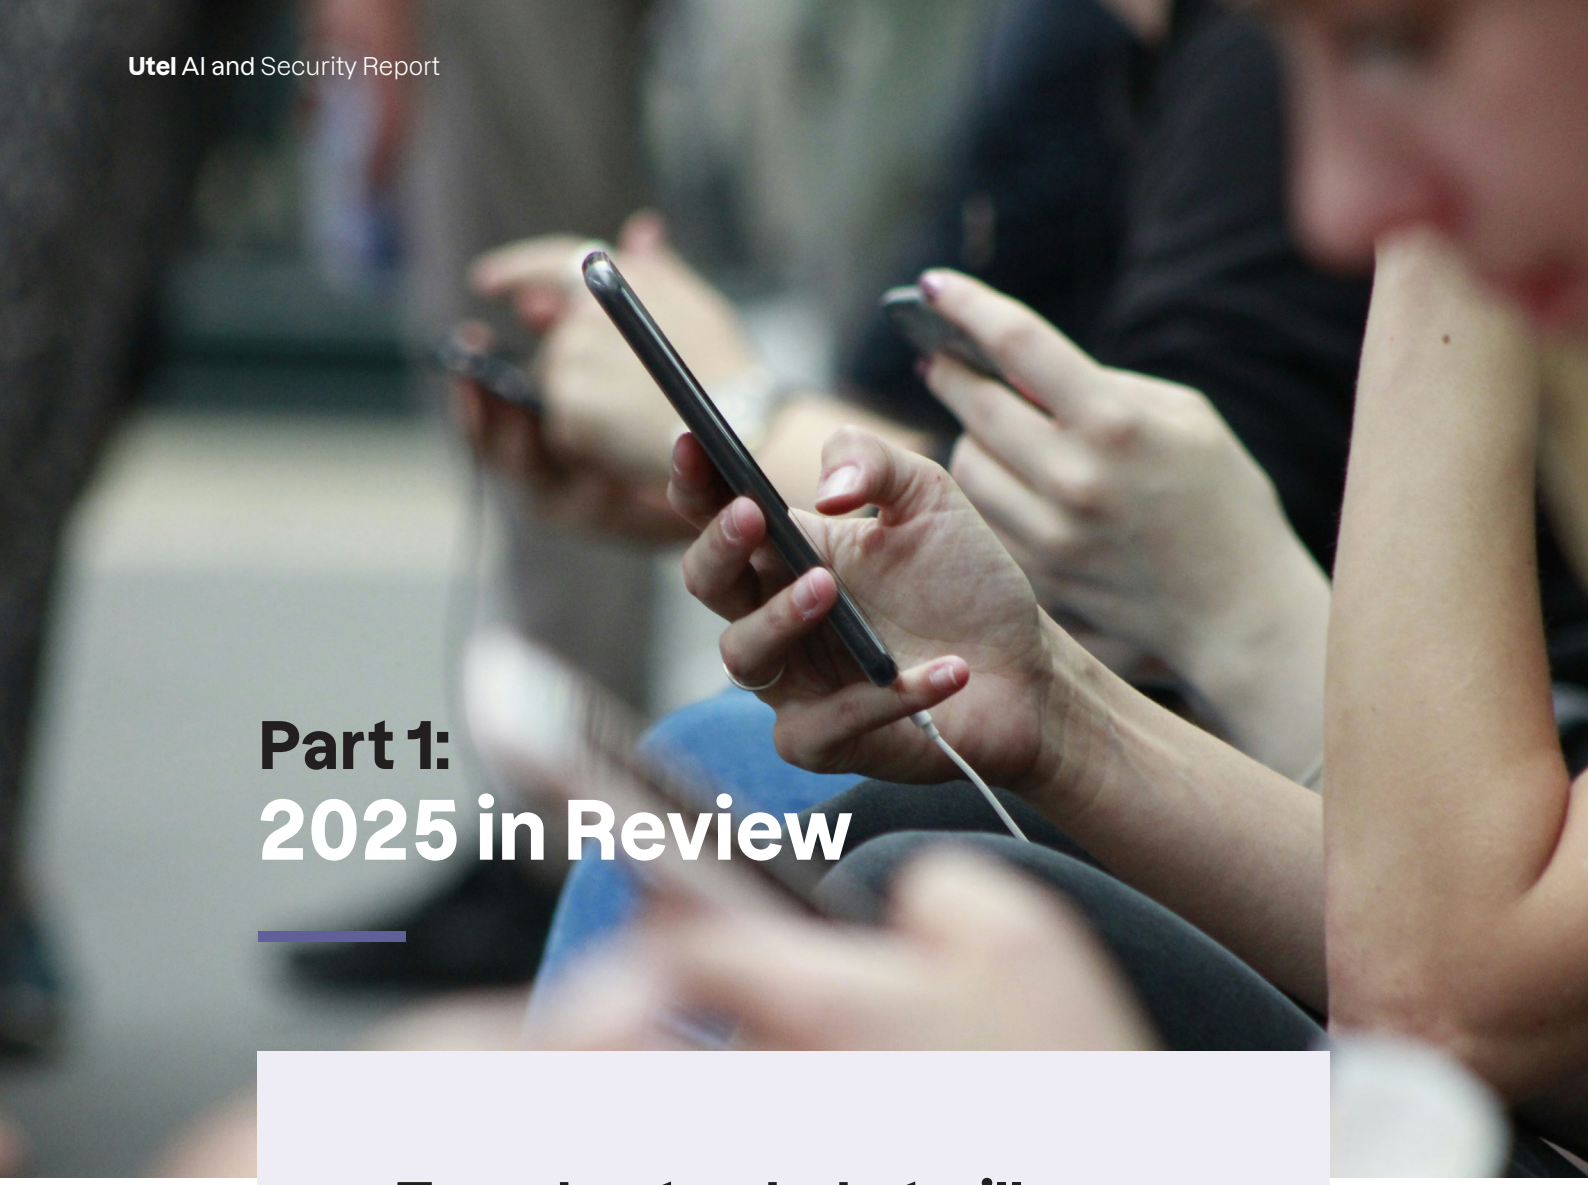

Part 1:
2025 in Review

To understand what will influence 2026, we will look back at the cases that marked 2025.

In 2025, we watched criminals adapt faster than defenses in every single case.
AI has been around for a long time, but we saw a major shift in capability and accessibility in 2025.

Highly advanced AI models suddenly became available to large enterprises, the public, and criminal organizations. The scale and speed of security threats surpassed what human-led processes could realistically manage. Traditional monitoring, manual review, and rule-based systems became insufficient and could no longer keep pace with attackers using AI-curated fraud.

In 2025, we experienced incidents that will shape how we view security going forward. These are some examples of what we have caught:

Intelligence Brief 01:
Breach of the Digital Shield

Threat Classification:
Adaptive Spoofing / International Bypass

Date:
Q1-Q2 2025

Target:
Norwegian telecommunications infrastructure

Sophistication Level:
Medium (existing infrastructure, simple modification)

At the very beginning of 2025, the Norwegian government, in collaboration with national telecom network operators, launched a digital shield to prevent criminals from spoofing Norwegian numbers. The shield proved effective, and by June 2025, it had stopped 61 million spoofing attempts[1]. However, fraudsters adapted their methods within less than a month.

Instead of spoofing Norwegian numbers, they now place automated calls from foreign numbers – primarily European or international ranges – using so-called "Wangiri" or robocall techniques.
In some cases, attackers simply modify Norwegian mobile numbers by adding a country prefix (e.g., adding "+" in front of a national number), causing the call to appear as an international call and thereby bypassing the national spoofing filters. Although the number appears foreign on the handset, many users still answer such calls, allowing fraud attempts to continue.

We were able to detect this because the traffic patterns were suspicious, even though the calling numbers appeared legitimate.

This demonstrates how rapidly fraudsters adapt to technical countermeasures and why detection systems must go beyond number-based filtering to behavioral and traffic-pattern analysis.

Lesson for 2026:
Static defenses create predictable attack surfaces, and fraudsters iterate faster than regulations. Behavioral analysis is now mandatory, not optional.

Wangiri

Wangiri is a Japanese term meaning "one ring and cut." It refers to a widespread phone scam where fraudsters trick you into calling back expensive, premium-rate international numbers.

1 Digitaliserings- og forvaltningsdepartementet (2025, June 11).

Your network isn't just a fraud target, it's potential money laundering infrastructure.

Intelligence Brief 02:
Exposing Money Laundering

Classification:
Money laundering via telecom infrastructure

Date:
Q2 2025

Location:
Central Europe

Sophistication Level:
High (requires coordination, infrastructure, sustained operation)

A mobile network operator in Central Europe deployed Utel Defense to replace a legacy, rule-based fraud-detection system. Shortly after activation, the AI-based anomaly detection identified suspicious voice traffic patterns.

The investigation revealed approximately 1,400 prepaid SIM cards generating continuous calls, 24/7, routed through one European operator, to a premium-rate number in another European country. The destination number was registered to a fake adult escort service that charged high rates for receiving calls.

The traffic pattern strongly indicated a money-laundering scheme in which criminals converted illicit funds into legitimate revenue through premium-rate calls. This case highlights how modern fraud and criminal organizations increasingly combine telecom abuse with financial crime, and how AI-based behavioral analysis is essential to uncover such complex schemes.

Lesson for 2026:
Your network isn't just a fraud target; it might be facilitating organized crime and serving as a potential money-laundering infrastructure. The critical question regulators will ask: "Did you deploy technology capable of detecting this?"

Ask yourself:

- Can your fraud detection correlate patterns across millions of SIM cards?
- Do you analyze aggregate behavioral patterns, or just individual violations?
- Can you detect "statistically impossible" consistency in usage?
- Would you notice if 1,400 cards exhibited identical behavior?

Intelligence Brief 03:
The "Nina" Scam

Classification:
Data-driven personalized fraud / Cross-border vishing

Date:
January 2025

Target:
Norwegian Mobile Subscribers

Attack Origin:
Primarily France (international numbers)

Sophistication Level:
Very High (requires data correlation + behavioral targeting)

The "Nina" scam is a targeted voice-fraud campaign recently observed in Norway, in which fraudsters place large volumes of scam calls to subscribers with the name Nina. A Norwegian operator identified a surge in such calls originating abroad, particularly from France.

The pattern suggests that attackers are using leaked or open-source data to match names with phone numbers, increasing the likelihood that recipients will trust the call. This represents a shift from random robocalling to more data-driven, personalized fraud attempts designed to improve success rates. This works because it addresses the victim by name, building trust and indicating that it must be legitimate.

Utel's AI model detected the scam within a day of deployment by identifying suspicious, irregular calling patterns. Because of the detected strange behavior from a grouping of faked calling numbers, our technicians were able to correlate the numbers being called to names and uncovered the scam method. Without detection, the scammers would have continued through the phone book, calling people in alphabetical order.

The "Nina" scam was a proof-of-concept. If a name + phone number combination can be weaponized for targeted fraud, should organizations limit exposure of this data pairing?

Lesson for 2026:
Attack doesn't require sophisticated hacking. It requires data shopping. If attackers can correlate "Nina" with phone numbers, what else can they correlate?

Ask yourself:

- Can your system analyze call patterns across population segments?
- Can you detect anomalies in "calls received by subscribers named [X]"?
- Would you notice 847 calls to "Nina" in one hour vs. a baseline of 3-5?
- Can you identify name-based targeting patterns in real-time?
- Could you detect this on Day 1, or would you find it weeks later in analysis?

Compliance is not security. Ticking boxes stops nothing.

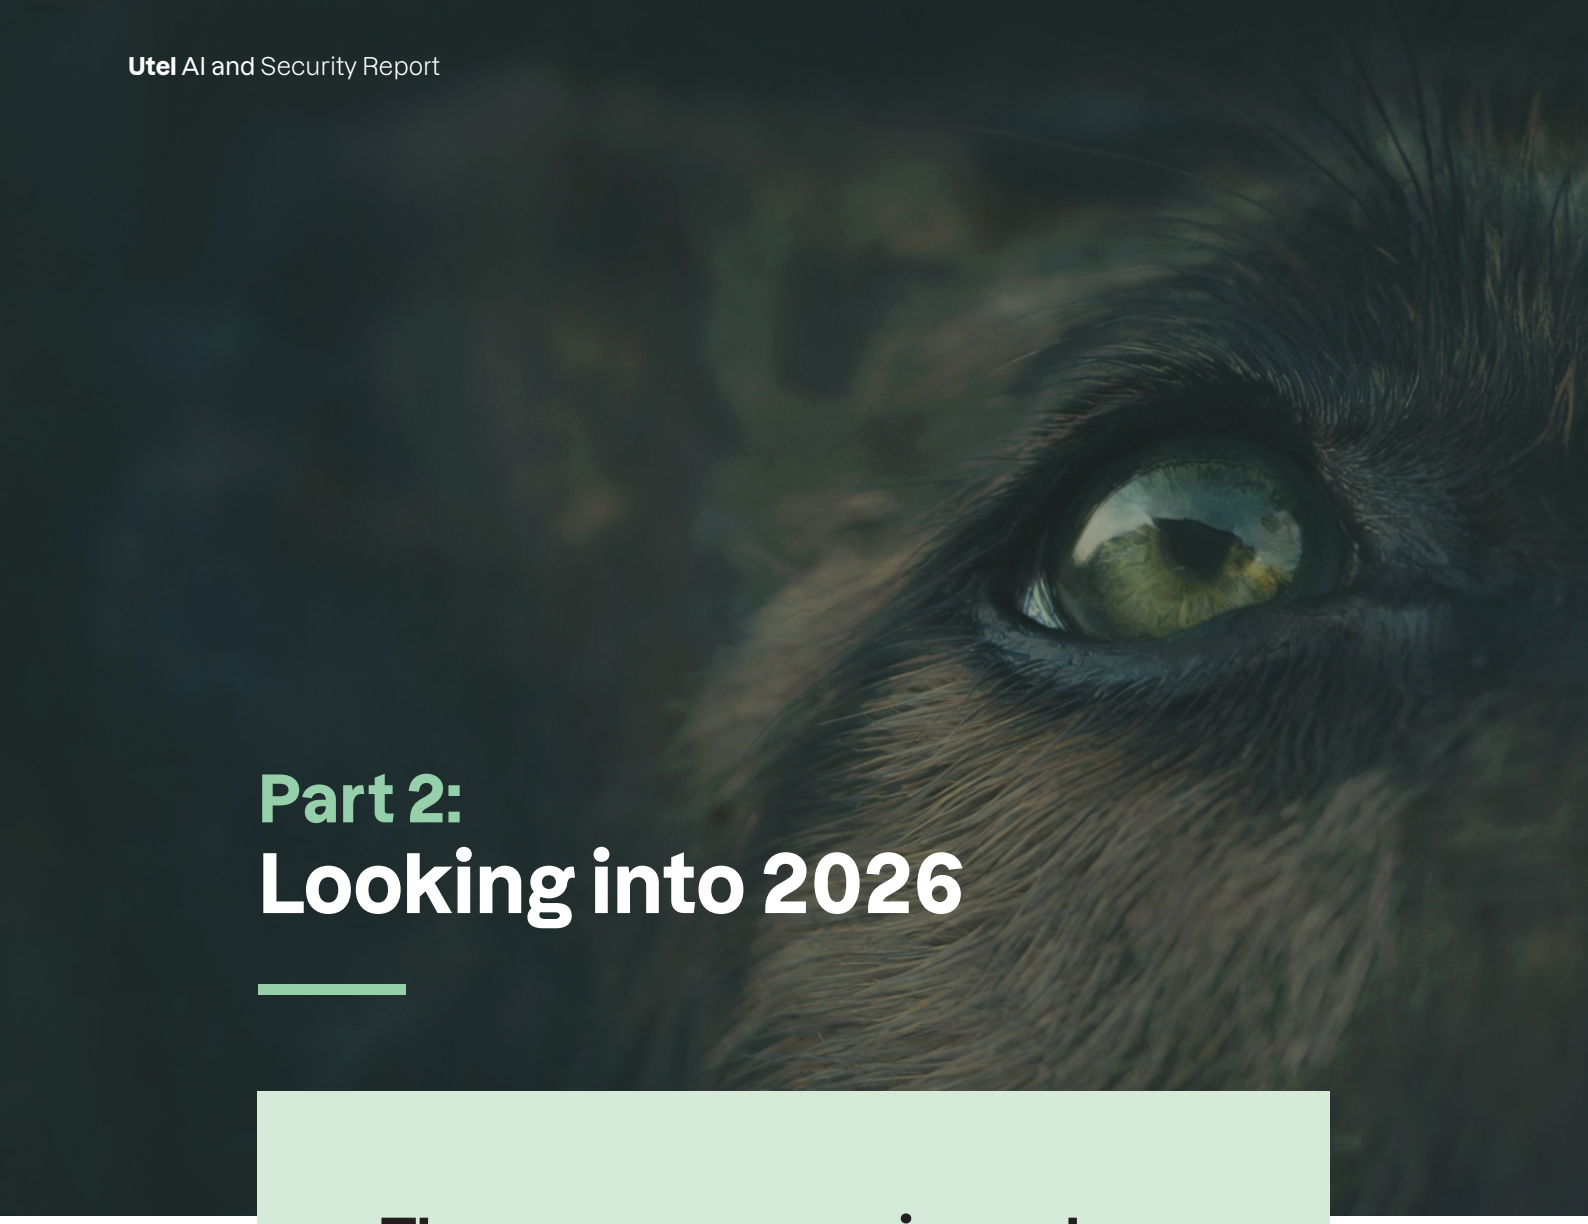

Part 2:
Looking into 2026

The cases we experienced in 2025 will only continue to evolve.

For 2026, we focus on three main security threats:

- Highly personalized scams with use of real-time AI-adaptions
- Enterprise-level fraud
- Hidden data leakage

All these trends have multiple layers and methods, and they have in common that they use AI as a means to bypass security systems and natural skepticism.

Let's consider the case of the "Nina" scam. If criminals can target by first name now, what's next?

The customers' data isn't just leaked, it's being actively weaponized against them through telecom networks. For 2026, the methods will utilize

- Age-based targeting uses scripts adapted to age, where seniors are more vulnerable
- Location-based targeting uses city- or address-specific scams, e.g., calling about property tax in a specific area where the target lives
- Latest Purchase-based targeting utilizes recent acquisitions and people who have made investments, e.g., car buyers being called about warranty, and calls or SMS regarding shipping or customs
- Life-event targeting addresses recent movers, job changers, etc., and exploits a possible vulnerable situation

All of this data exists and can be correlated with phone numbers. This data can be used for advanced social engineering to manipulate people and play on their trust.

We can no longer trust our senses as photos, videos, and voices can be faked quickly, and communication can be adapted to fit personal language and style.

At the same time as personalized fraud becomes more effective, we see a clear shift towards larger-scale fraud directed at organizations rather than individuals.

Multi-factor authentication can be bypassed by targeting employees, email threads can be infiltrated, and tools and software can introduce a risk of data leakage and security bypass.

AI and deep-fakes make scamming people much easier for criminals, but they can also play a major role in detecting attacks by providing extensive analysis within seconds.

Deepfake

Deepfake is a piece of media, usually a video or audio clip, that has been digitally manipulated by artificial intelligence to make it look or sound like someone is saying or doing something they never actually did.

AI in business has delivered clear value, but it has also introduced a new category of enterprise risk.

Managing AI has shifted from being a technical task to becoming a strategic responsibility that affects security, privacy, and compliance. As AI becomes integrated into workflows, the attack surface expands beyond systems to include people and communication itself.

The balance between security and privacy has become a managerial challenge, particularly when choosing AI-assisted tools and communication platforms. These decisions directly impact GDPR compliance, risk of data leakage, and exposure to espionage.

Regulatory initiatives such as EU Chat Control reflect the growing need to strengthen security while maintaining privacy protections.

The 2026 Shift:

In 2026 we expect to see an increase in attacks against organizations and fraud targeting employees rather than individuals and private funds.

As individuals, we have become increasingly more cautious with personal funds, and previous fraudulent activities are starting to affect private citizens' awareness. We might not be equally aware as employees or equally cautious with company funds.

The underlying driver for this shift is a change from volume-based to value-based fraud. Earlier scams targeted individuals, generating small gains per victim and forcing criminals to operate at high volume.

Targeting enterprises is more efficient: each successful attack delivers a higher payoff, requires fewer attempts, and creates less activity, making it harder to detect.

Managing AI effectively means protecting people as much as systems. This is how your business is targeted:

Shadow AI and Indirect Prompt Injection
Shadow AI represents a critical internal vulnerability where employees use public AI tools and Large Language models, like ChatGPT and Gemini, outside approved controls, uploading sensitive corporate data without IT oversight.

This creates hidden exposure points and immediate security blind spots, undermining existing security, compliance, and data protection measures.

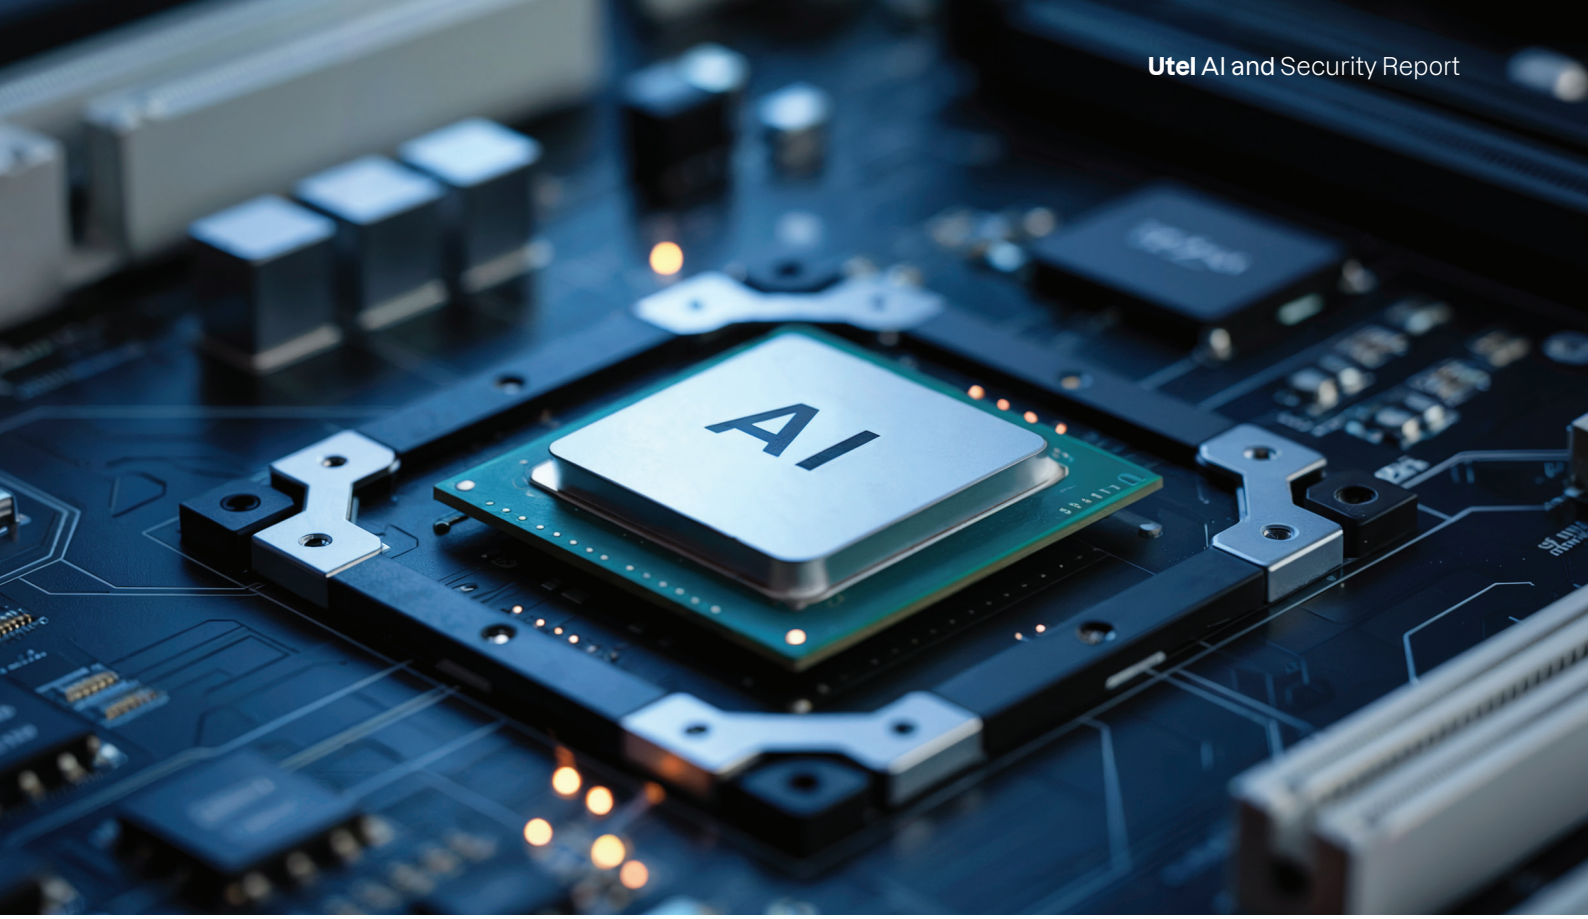

At the same time, external actors are exploiting increased reliance on AI through indirect prompt injection.

By hiding malicious instructions inside seemingly harmless documents, they can influence how AI tools behave when employees ask them to analyze or summarize a file, potentially extracting data or manipulating results without the user noticing.

The MFA Illusion:

In 2025, we began to see a shift: attackers stopped trying to steal passwords and instead bypassed them entirely. Now they steal "active sessions" instead.

While the extensive adoptation of Multi-Factor Authentication (MFA) has made simple password theft ineffective, attackers have adapted by targeting the user directly. Instead of guessing passwords, criminals trick employees into installing malware through social engineering, disguising it as routine updates or software patches.

This new generation of malware often uses AI to continuously rewrite its code, allowing it to slip past traditional antivirus scanners.

Once installed, the malware bypasses MFA entirely by stealing the "active session" from the user's browser, letting the attacker impersonate employees without ever needing to log in.

Ask yourself:

- Where does data go? (EU or foreign servers?)
- Who can access it? (Company only or vendor too?)
- What decisions can it make? (Read-only or take action?)
- How can it be manipulated? (Prompt injection vulnerabilities?
- What regulations apply? (GDPR, NIS2, DORA?)

Wangiri 1.0

Consumer Target - The Old Way

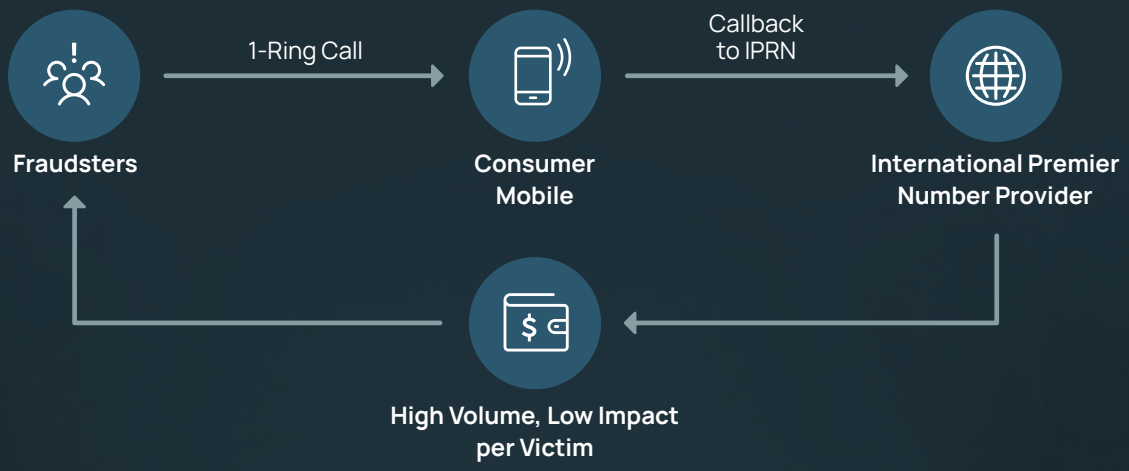

Wangiri 2.0

Enterprise Target - The 2026 Way

The threats of 2026 don´t look like attacks. They look like "Business as Usual".

Wangiri 2.0:

When "Call me Back" forms become traps.
This type of fraud has long been associated with missed calls to individual consumers, tempting them to call back expensive international numbers, but this tactic has evolved. Wangiri 2.0 marks a strategic shift from mass-consumer scams to targeted enterprise fraud.

Wangiri 2.0 is an updated version of the old "one ring" phone scam. Instead of calling random people, attackers now target businesses. Bots fill out contact forms and "call me back" requests on company websites, using international premium-rate phone numbers.

When sales or customer service follows up, the company ends up calling the expensive number itself. Everything looks normal until the phone bill arrives.

The losses accumulate quietly across numerous call-backs and typically surface only when finance reviews rising phone expenses, by which point the charges have already been settled, and the revenue shared with the attackers.

Thread Hijacking:

AI makes your Emails sound like you.
Thread Hijacking target business through email. Attackers monitor ongoing conversations and use AI to replicate language and style. When the timing is right, they jump into the thread with a fake message, e.g., asking for a payment or transfer, or expressing good news like a raise or bonus. With the help of AI, the message sounds like the real sender, and because it's part of an existing conversation and comes from a trusted address, it rarely raises alarms.

What makes these attacks so effective is that nothing feels suspicious. Employees are doing exactly what they're supposed to do: replying to emails, calling back leads, and moving business forward. The attack is hidden inside normal work. AI doesn't just make fraud possible; it makes fraud nearly indistinguishable from legitimate communication.

For telecom operators, this means that your business customers are the targets, and enterprise customers can blame the operators, not just the criminals. In 2026, we believe enterprise customers are looking for operators to provide security, not just connectivity. 2026 requires security that understands normal business behavior well enough to spot abnormal behavior hiding inside it. That requires behavioral AI, not rule-based systems, because rules can't detect if something is subtly wrong.

Smokescreen Attacks:

When your fraud detection becomes the weapon.
Let's consider this scenario: In Q4 2025, a European operator celebrated stopping a massive fraud attack. A high number of SIM cards were making obviously fraudulent international calls. Every alert triggered, and the fraud team mobilized and acted.

The attack was blocked, and all considered the scenario a success story. One month later, during routine reconciliation, they discovered €2.3 million in losses from undetected fraud that had run during the same period, as they were flooded with fraudulent calls. The SIM card attack wasn't fraud. It was a distraction.

A growing trend in telecom fraud is the deliberate generation of large volumes of easily detectable fraudulent traffic designed to trigger rule-based detection systems. By flooding networks with obvious scam activity that exceeds alert thresholds, fraudsters effectively create a "smokescreen" that consumes the attention and resources of fraud teams. While analysts are busy handling these high-volume alerts, more sophisticated, lower-volume fraud activities are executed in parallel and remain below detection thresholds.

While your fraud team is investigating the loud operation, the quiet operation runs undetected. By the time you finish investigating the obvious fraud, the real fraud is complete. This tactic exploits the limitations of static, rule-based systems. It underscores the need for AI-driven anomaly detection that can identify subtle deviations and correlated attack patterns across traffic types.

For telecom operators, this means fraud detection itself has become part of the attack surface. By 2026, sophisticated actors study detection logic, exploit limited team capacity, and deliberately overwhelm predictable thresholds. Responding effectively to AI-powered attacks like these requires AI defense that cannot be distracted or saturated, and that identifies behavioral patterns attackers create without realizing it.

SMS-Triggered Call-Back Fraud:

When your customers dial the scammers.
Traditionally, scammers have called the victims, tricking them into various actions, often resulting in a money transfer. A growing fraud trend involves attackers initiating scams via SMS rather than direct calls. Victims receive a text message claiming an urgent issue, such as a failed delivery, an unpaid customs fee, or an account problem. They are instructed to call a specific phone number to resolve it. The victim, not the fraudster, therefore initiates the actual scam call. For a telecom operator, this looks like normal behavior, and the scam completes successfully, leaving the customer out of pocket.

This reverses the traditional fraud pattern and makes detection more difficult, as the call appears to be a legitimate outbound call from the subscriber. Effective detection requires correlating SMS activity with subsequent outbound calls and analyzing behavioral patterns rather than relying solely on inbound call filtering.

For telecom operators, this means the direction of a call no longer signals legitimacy: inbound does not equal scam, and outbound does not equal safe. By 2026, effective protection requires cross-channel behavioral detection that correlates messaging content with subsequent voice activity. Without it, customers continue to fall victim to scams, lose money, and hold operators responsible for failing to intervene.

What can be done?

- Cross-channel monitoring (SMS and voice analyzed together)
- Content analysis (detect call-back triggers in SMS text)
- Temporal correlation (link SMS receipt to subsequent calls)
- Behavioral baseline (customer's normal calling patterns)
- Real-time processing (block call before it connects)

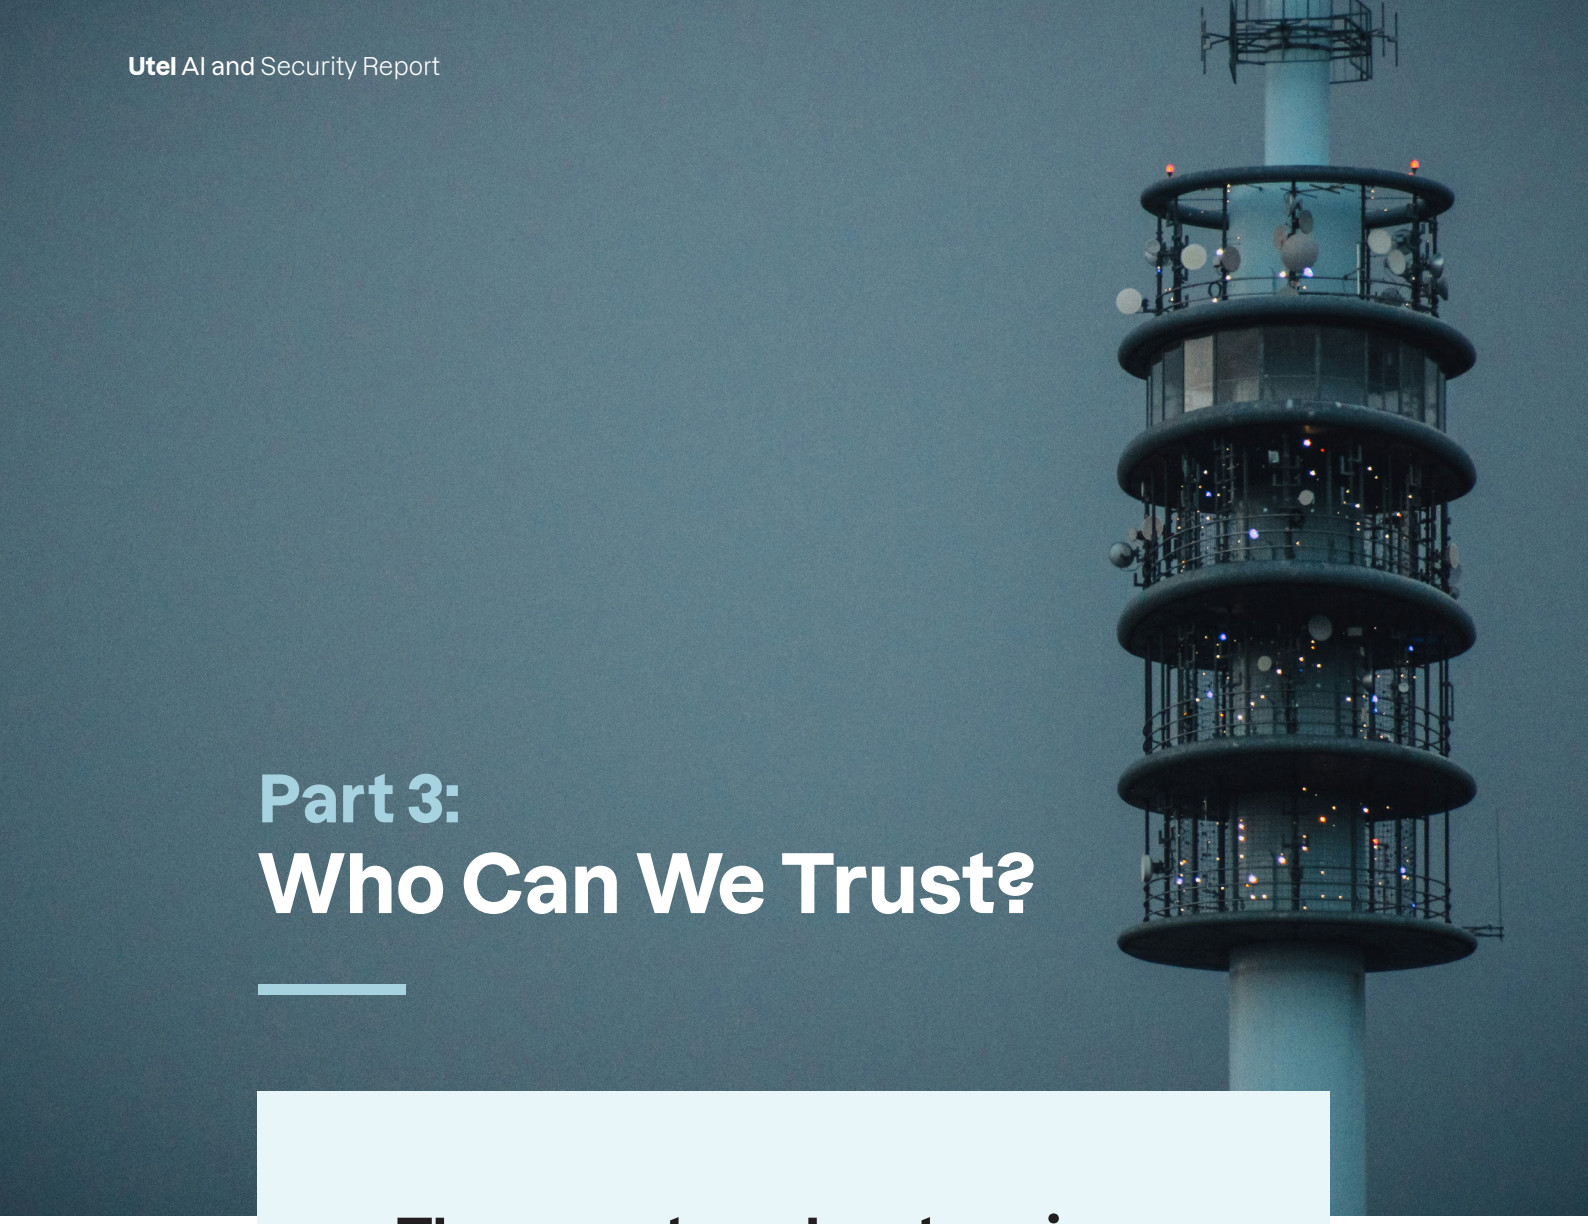

Part 3:
Who Can We Trust?

The operator advantage in the age of AI fraud and how you can win against voice over IP offered by the large tech companies.

Most people are uncritical about whom they trust, a familiar voice, a known face, and apps with seemingly harmless terms and conditions. For decades, trust in digital communication has relied on human judgment. Now that voices can be cloned, faces can be generated, and language can be replicated with near-perfect accuracy from data accessed on the platforms we use every day, AI has surpassed the limits of human perception.

Telecom operators are in a unique position to provide security and trust for their customers.

With collective regulations, agreements, and collaboration among parties, operators can secure their networks and protect customers from fraudulent traffic. This is a security for most tech companies with complete platform sovereignty that doesn't provide for their customers.

Traditional security models were not designed for this reality. Rule-based fraud detection is inherently reactive, relying on predefined directions, known patterns, and obvious irregularities.

AI-driven attacks easily surpass these filters. AI-attacks are becoming adaptive, context-aware, and capable of behaving correctly until the damage is done.

Security systems need to be predictive and proactive, and to provide real-time verification before any attacks reach consumers, so that our fundamental trust can be restored and maintained. While individual attributes can be copied, continuous behavioral patterns are far harder to fake.

AI models designed for security monitor normal behavior among users, systems, and networks, and detect subtle anomalies that indicate risk.

Why Telecom operators have a structural advantage:

Collective security: Industry agreements and regulatory requirements enable operators to collaborate on security. When one operator detects a new attack pattern, protective measures circulate across networks, while other tech platforms compete on data and intelligence. They believe that sharing threat intelligence reduces competitive advantage and that security is secondary to growth. We believe a strong focus on security, along with being upfront and proactive, provides a competitive advantage.

Infrastructure control: Operators control the network layer. You can analyze traffic patterns, detect behavioral anomalies, and block attacks before they reach endpoints. Tech platforms only see the application layer, and by the time they detect fraud, it's already reached the victim, if they ever detect it at all.

Regulatory accountability: Operators answer to telecom regulators, privacy authorities, and government agencies. There's legal liability for facilitating crime through your network. Tech platforms often operate in regulatory grey areas, and "Platform" vs "publisher" distinctions let them avoid responsibility for harmful content flowing through their systems.

2026 is about proving you can be trusted. You need to take a position.

The strategic choice

Businesses face a clear choice. They can compete on price with global platforms or differentiate on trust by owning security and protection as core elements of the brand.

Customers will increasingly choose companies that actively defend them, not merely connect them. Security must become a brand position, not just a technical feature. As digital dependence grows and fraud becomes more sophisticated, trust is emerging as a primary differentiator in the telecom market.

In practice, this means making protection visible and measurable. Publicly committing to outcomes, such as the number of fraud attempts blocked, signals accountability, regular updates improve transparency, and customer education builds trust. When security is clearly communicated and consistently delivered, it becomes a differentiator that strengthens customer loyalty.

Failing to prioritize security does not equal neutrality; it creates opportunity. When fraud goes undetected, organized crime can operate uninterrupted, using legitimate infrastructure to move money, exploit customers, and scale illicit activity. In this context, security is no longer just about protection, but about preventing networks from becoming enablers of criminal operations.

Over the past few years, the meaning of "we didn't know" has shifted dramatically. In earlier stages, not knowing what was happening in the network could be considered a reasonable explanation for not stopping criminal activity.

That threshold has since moved toward regulatory scrutiny and financial penalties, followed by reputational damage as customers and partners question an operator's role in enabling crime.

Looking ahead to 2026, a lack of awareness may no longer be treated solely as negligence, but as grounds for legal and criminal investigation in an environment where detection capabilities are known to exist.

The 2026 Security Mandate:

- **Own the Outcome**
  Make your battle against fraud transparent and evident. Include your fraud fighting metrics in your annual report.
- **End Plausible Deniability**
  Treat lack of network visibility as a legal risk. If you can't see it, you can't act on it. If you didn't know, you may be liable.
- **Compete on Trust**
  Competing on Price is a race to the bottom. Security is the only way up.

Moving forward, visibility, operational intelligence, and security must be a part of every organization's strategic plan.

Do you need help to strengthen your network intelligence and security? Get in touch with our experts.

Frode Gorseth, CCO

+47 41 21 31 84
fg@utel.tech

Rune Borge Kalleberg, CTO

+47 48 60 87 37
rbk@utel.tech

Utel