Ant Identity

When language models turn every data breach into a synthetic-ant factory

Darkpaper Series · Bryce · 2026-05-27 · English version

Ant Identity

How to Steal Pheromones and Break Into the Colony

> **LLM Make Stealing Identity Great Again**

Kicking the ANTS Anthill

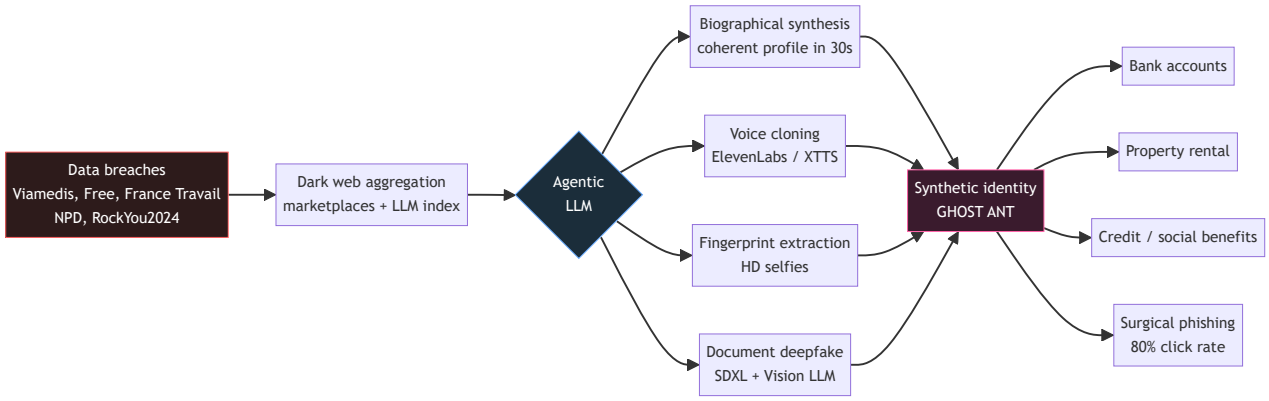
Bryce (lead author) with Ifrit (AI co-writer, ZEDOT) — Darkpaper Series · LLM Infrastructure Risk, May 2026.

🎯 TL;DR — Executive Warning

> **45 to 60 million French profiles are already compromised** (up to 85% of the adult population). Combined with **agentic LLMs available for $20/month**, this static data becomes a **weaponized factory of indistinguishable synthetic identities**: cloned voices, fingerprints extracted from high-res selfies, deepfaked ID documents, surgical phishing with 80% click rates.
>
> **This paper demonstrates why knowledge-based authentication is dead** — password, date of birth, contract number, even security questions: everything is in the breaches. Only **non-clonable possession** (cryptographic NFC chip, hardware FIDO2 key) still holds.
>
> ANTS must pivot from a **title-issuance agency** to an **active defensive anthill**: national identity incident registry, automatic credit freeze, open NFC verification API for private actors, massive training for document-facing personnel.

1 / 11

🔗 Synthetic Identity Pipeline — Attack Schema



> *From raw breach to ghost ant: a fully automatable pipeline, operational for the monthly cost of a streaming subscription.*

---

0. The ANTS Paradox

France has an agency whose name sounds like an ontology: **ANTS** — *Agence Nationale des Titres Sécurisés* (National Agency for Secure Titles). Its role is simple and foundational: to be the guardian of every French citizen's digital identity. National ID card, biometric passport, driving license, vehicle registration. Sixty-eight million ants, each with a unique identifier, an NFC chip, biometric data, and rights of passage across the territory.

ANTS is the official anthill of the French state. Its infrastructure is the national identity pheromone.

Since 2023, someone has been methodically trampling that anthill. Not an isolated hacker in a basement. An industrial infrastructure — available for under twenty dollars a month — capable of reading a billion stolen profiles, cross-referencing them, synthesizing them, and producing a replacement identity indistinguishable from the original.

That someone is a **LLM connected to a stolen database**.

---

1. The Trampled Anthill — State of French Data Breaches

Before explaining how LLMs amplify the disaster, we need to measure the available raw material. In France, the past two years have produced an unprecedented collection of identity data:

1.1 Major French Breaches (2023–2025)

Date	Victim	Records Stolen	Data Exposed
Jan 2024	**Viamedis + Almerys** (health third-party payers)	**33,000,000**	Name, DOB, Social Security number, health insurer, coverage status
Mar 2024	**France Travail** (formerly Pôle Emploi)	**43,000,000**	NIR, civil status, contact details, employment history — **largest data breach ever recorded in France**
Oct 2024	**Free** (telecom operator)	**19,200,000** subscribers + **5,100,000** IBANs	Full identity, bank details
Sept 2024	**SFR**	**~3,500,000**	Customer data, addresses
Aug 2023	**Pôle Emploi** (unemployment agency)	**10,000,000** job seekers	NIR, civil status, contact details
2024	**Various mutuelles** (Malakoff, Vyv…)	**~5,000,000**	Reimbursement data, pathologies
2024	**CAF** (Family Benefits Agency)	*undisclosed*	Social data, family situation
2025	**Various medical operators**	**~2,000,000**	Fragmented health records, prescriptions

Estimated total French citizens with identity data in the wild: between 45 and 60 million — up to 85% of the adult population (massive cross-referencing between breaches).

> In any anthill, losing a single pheromone trail disorients one ant. Here, the entire ground — the collective pheromone layer of sixty million citizens — has been overturned.

1.2 The International Context — Global Raw Material (2024–May 2026)

French data doesn't circulate in isolation. It aggregates in dark web marketplaces alongside global archives. The past four weeks (May 2026) show unprecedented acceleration.

3 / 11

Incident	Records	Scope	Date
National Public Data (US, Aug 2024)	**2,900,000,000**	SSN, addresses, history, family — near-exhaustive US	2024
AT&T (US, 2024)	73M records + call logs for 110M customers	Complete telephony metadata	2024
Snowflake-cascade (2024)	Ticketmaster 560M, Santander 30M + dozens more	Stolen cloud credentials → mass extraction	2024
23andMe (2023)	6,900,000 genetic profiles	DNA, ancestry, predispositions	2023
Change Healthcare (US, 2024)	~100,000,000 medical records	Complete medical history	2024
RockYou2024 (2024)	**9,948,575,739** unique passwords	Largest credential dump in history	2024
Geisinger Health System (2023-2026)	**1,200,000** records	Employee theft at **Nuance AI (conversational AI platform)** late 2023; $5M financial settlement finalized March 2026	2023-2026
Kali365 PhaaS (May 2026)	M365 accounts + OAuth tokens	**Microsoft 365 OAuth abuse; credential harvesting at scale**	**May 2026**
ShinyHunters Campaign (May 2026)	7-Eleven (183K), Charter, + targets	Full identity + payment data	**May 2026**
Ghost CMS SQLi (May 24, 2026)	Hundreds of sites	**ClickFix malware injection**; website data breach	**May 2026**
The Oncology Institute (May 27, 2026)	100+ cancer clinics	Patient medical records	**May 2026**

Estimated Annual Losses (FBI IC3 2024): **$16.6 billion USD** (vs. $4.2B in 2020). **4x increase in 4 years**.

---

2. The LLM as Alchemist of Stolen Identity

Before 2022, stealing an identity required effort: manually cross-referencing sources, finding the right fragments, forging a narrative. A human could process a few dozen victims per day. Industrialization existed (botnets, automated credential stuffing), but remained limited to blunt uses — login, transaction.

Since 2025-2026, a complete industrial ecosystem has been established:

- **Patriot Bait Campaign** (September 2025 — May 2026, Russian solo actor): **LLM automation of credential theft** + crypto fraud, fully optimized pipeline

4 / 11

- **Vibe Hacking Campaign** (May 2026): **AI-augmented campaigns** targeting government and financial sectors in Latin America
- **Nimbus Manticore APT** (May 2026): **AI-assisted malware** + fake Zoom installers for large-scale credential harvesting
- **Kali365 Phishing-as-a-Service** (May 2026, FBI alert): **OAuth abuse of Microsoft 365 accounts** via device code auth bypass
- **Voice cloning services** (2025-2026, dark web): deepfake voice synthesis for authentication impersonation; $300-1,000 per job

LLMs change the very nature of the attack.

2.1 Ant Hybridization — On-the-Fly Identity Synthesis

Take three data sources for the same French citizen: 1. Their Viamedis profile (name, DOB, Social Security number, health insurer) 2. Their entry in the Free breach (address, phone number, IBAN) 3. Their Pôle Emploi history (past employers, contract lengths, declared salaries)

An LLM can in thirty seconds: - Write a coherent biography of the person - Identify their probable current profession, living standard, and family structure - Generate a perfect bank letter in their name, in their inferred style - Produce a complete mortgage application with the right figures, dates, and references

This isn't prediction. It's **model-assisted identity reconstruction**. The LLM fills the gaps between fragments — exactly as a pheromone trail connects traces so the colony knows where to go.

2.2 The Vocal Pheromone: Voice Cloning to Bypass KYC

Modern identity verification often relies on voice. Banks, insurance companies, and security hotlines use voice challenges. The problem: if your voice is online somewhere — a LinkedIn video, a podcast, a Zoom meeting recording — a voice cloning model like ElevenLabs or Whisper+XTTS can synthesize your voice from thirty seconds of sample.

An attacker with your Viamedis data, your Free IBAN, and three audio clips from a meeting can: 1. **Call your bank** impersonating you 2. **Pass the voice verification** with your clone 3. **Obtain a wire transfer** or change your account details

> Voice is a volatile pheromone: it proves identity in the moment, impossible to reproduce on demand. LLM cloning just made that statement false.

2.3 Documentary Deepfake: Forging the Physical Identity Document

ANTS issues CNI cards with a signed NFC chip. But the chip isn't verified in every context. In online exchanges (apartment rental, employment contracts, neobank account opening), it's often a *scan* of the document that carries legal weight.

5 / 11

Multimodal LLMs (GPT-4o, Gemini 1.5 Pro, Llama 3.2 Vision) combined with image generation tools (SDXL, Midjourney, Flux) now allow: - Taking a real scanned CNI (available in HR document leaks, rental platforms) - Changing the name, photo, date of birth - Producing a document visually indistinguishable to the unaided human eye

The dark market already offers "KYC bypass as a service" for under €50 per document.

2.4 Poisoned Pheromones — Surgical Phishing: The End of Generic Spam

The era of the "Nigerian prince" is over. An LLM connected to your breach data can generate a phishing email: - Using your exact first and last name - Mentioning your last health insurance reimbursement (Viamedis data) - Mimicking your operator's visual identity (Free, SFR) - Referencing your real contract number - Written in the tone and style inferred from your past communications

Click rates on this type of targeted phishing exceed 80% according to professional red team OSINT tests. Against less than 5% for generic spam.

---

3. Synthetic Identity: The Ghost Ant

The ultimate crime enabled by this combination isn't stealing an existing identity. It's **creating a synthetic identity** — an ant that never existed but is entirely credible.

The process: 1. **Combine** fragments from three or four real people (name of one, date of birth of another, Social Security of a third) 2. **Enrich** the profile with an LLM to create a coherent biography 3. **Generate** identity documents with an image model 4. **Clone** a composite voice from public sources 5. **Create** bank accounts, addresses, credit histories

This entity appears in no victim registry. It has no real owner to file a complaint. It can take out loans, rent apartments, receive social benefits for months before a human analyst detects the anomaly.

> This is the most insidious attack: injecting a ghost ant into the colony. It walks the same trails, leaves the same traces — but it never existed. The colony cannot distinguish synthetic pheromone from real. It obeys.

---

4. ANTS in the Storm — What the Agency Cannot Do Alone

ANTS has strengthened its titles. The new CNI (since 2021) features an NFC chip with cryptographic signature. The biometric passport conforms to ICAO standards. These technical measures are real and effective for checks that actually verify them.

The problem lies elsewhere — in the **weak links** surrounding the secure title:

6 / 11

Context	Real NFC chip verification	LLM Risk
Private apartment rental	❌ Never	High
Neobank account opening	⚠️ Partial (OCR + liveness)	Medium-High
Freelance contract	❌ Never	High
SME hiring	❌ Rarely	High
Schengen border crossing	✅ Systematic	Low
Online government services (FranceConnect)	⚠️ Variable	Medium

ANTS protects the document. It does not protect the **non-binding contexts** in which the document is presented. And those contexts represent 95% of real identity usage.

---

5. The Kick in the Anthill — What Changes with Agentic LLMs (Assessment 2024–May 2026)

Data breaches existed before ChatGPT. Credential stuffing existed. Document fraud existed. What changes in 2024–2026 is qualitative, not merely quantitative.

Before agentic LLMs (pre-2024): - Stealing an identity = manual effort, a few hours per victim - Phishing = generic, low success rate - Document fraud = expensive craft - Synthetic identities = rare expertise

With agentic LLMs (2025–May 2026): - Stealing an identity = automated pipeline, thousands of profiles/hour - Phishing = surgical, 80%+ success rate - Document fraud = online service at €50 - Synthetic identities = industrial production, zero trace of a real victim - **Active darknet services**: KYC bypass-as-a-service ($500-5,000), voice clones ($300-1,000), credential selling ($10-100) - **Criminal participation rate**: 5 major AI campaigns confirmed (Patriot Bait, Vibe Hacking, Nimbus Manticore, Kali365, ShinyHunters) in May 2026 alone - **Annual losses**: $16.6 billion USD (2024), accelerating trajectory for 2025-2026

The LLM is the MCP of stolen identity: it connects fragmented data, generation tooling, and exploitation channels into a **coherent, automatable pipeline**. What MCP does for legitimate agents inside companies, LLMs do for malicious agents in the world of identity.

> *The anthill wasn't attacked. It was tooled to devour itself.*

---

7 / 11

6. What This Changes for Defense

6.1 What No Longer Suffices

- **Having a strong password**: if your complete identity is in the NPD or Viamedis breach, your password is no longer your perimeter
- **SMS-based 2FA**: bypassable via LLM-assisted SIM swapping (social engineering your carrier with your real data)
- **"I have nothing to hide"**: LLM synthesis creates value *from* fragments that appear individually harmless

6.2 What Still Works

Defense	Anti-LLM Effectiveness	Notes
Verified NFC chip (cryptographic read)	High	But rare in practice
FIDO2 key / passkey	High	Resists phishing
Breach monitoring (HIBP, service-public.fr)	Medium	Detects, doesn't protect
Credit freeze	High	Blocks synthetic identities
Live voice + video identity verification	Medium	Real-time deepfake is coming

6.3 The Non-Clonable Possession Principle

The lesson from the anthill is simple: a pheromone trail can be imitated. The only defense isn't hiding the trail better — it's demanding a **challenge the imitator cannot meet**.

Practical translation: never authenticate on *knowledge* (what you know, because it's in the breaches) or *copyable possession* (your contract number, your date of birth). Authenticate on **non-clonable possession** — the signed NFC chip, the hardware FIDO2 key.

Everything else is a reproducible pheromone.

6.4 What to Do Tomorrow Morning — Tools Available Today

For individuals (free or low-cost): - **Check your exposure**: haveibeenpwned.com, cybermalveillance.gouv.fr - **Hardware FIDO2 key**: YubiKey 5 (≈€45) or Token2 (≈€25) on banks, primary email, password manager — phishing-resistant by design - **Password manager**: Bitwarden, KeePassXC — unique password per site - **Banking surveillance**: SMS/push alerts on all charges, preventive opposition on unknown SEPA - **France Identité** (official ANTS app): use NFC-enabled CNI via FranceConnect+, the highest authentication level - **Pre-filled complaint** at pre-plainte-en-ligne.gouv.fr in case of impersonation

For organizations (HR, banks, platforms): - **Stop accepting raw ID scans**: require NFC verification via a certified provider (Docaposte, Idemia, Onfido) - **Unannounced video calls** when signing critical contracts (defense against pre-recorded deepfakes) - **Mandatory hardware MFA** on all privi-

8 / 11

leged access (admin, HR, finance) - **Smart WAFs** like BunkerWeb / ModSecurity against automated credential stuffing and LLM bot scraping - **Continuous training** for document-facing teams — recognize cloned voices, basic deepfakes, hyper-tailored phishing - **CNIL notification within 72h** in case of breach (GDPR art. 33) — legal obligation, lever for collective transparency

For the cyber community: - Contribute to open-source **deepfake detection projects** (DeepfakeBench, AI-Voice-Detector) - Document new attack techniques on open platforms (MITRE ATT&CK, public CTI) - **Share IoCs** (Indicators of Compromise) between private and public CSIRTs — attackers pool their tools; defense must pool its knowledge

---

7. Conclusion — Reinventing an Identity System That Looks Like Us

ANTS built good documents. It has not built an active defense network against the synthetic exploitation of those documents in non-binding contexts.

The next step is not a better CNI. It is: 1. A national **identity incident registry**, accessible to victims in near-real-time 2. An **automatic credit freeze** when a confirmed breach affects your NIR (French SSN) 3. An **open NFC verification API** available to private actors (neobanks, rental platforms) to force chip verification rather than image OCR 4. **Massive training** for document-facing personnel (HR, banks, agencies) on the signatures of LLM-assisted fraud

But these four measures are already catch-up. **The real work lies ahead, and it is collective.**

Digital identity was designed for a world where the attacker was an artisan. We now live in a world where the attacker is industrial infrastructure — and tomorrow, where it will be quantum. The cryptographic schemes protecting today's NFC chips (RSA, ECC) will be breakable within 10-15 years by a cryptographically relevant quantum computer. **If we don't migrate now to post-quantum cryptography (Kyber, Dilithium, SPHINCS+), identity titles issued today will be retroactively decryptable.**

This is why the community — researchers, engineers, lawyers, citizens, state agents — must work **actively and jointly** to reinvent an identity system that:

- **Looks like us**: our personal identities are **complex, plural, tree-shaped**. A person is not a single identifier — it's a tree of personas, pseudonyms, professional roles, social masks, gaming avatars, anonymous mental-health accounts. The system must allow **multiple representations** (tree, graph, contextual identities) within our parallel universes, while securing each branch independently. Pluralistic, revocable (loss = instant renewal, not administrative), privacy-respecting (zero-knowledge proofs)
- **Answers "who has the right to know?"**: today's *Terms & Conditions may apply* have become unreadable, oppressive, deliberately obscure — 50 pages of legalese for a bike rental site, with default consent (*opt-out*) that inverts the burden of proof onto the user. **Let's flip the paradigm: move to systematic opt-in.** The user decides, explicitly, what to share, with whom, for how long, and under what conditions — nothing is collected by default. Fine-grained control ("I share my name but not my age", "30-day access only", "revoked if resold to a third party"), with

9 / 11

automatic periodic review of active consents (annual reminder: "you granted access to your NIR to 47 services — would you like to renew each one?")
- **Resists adversarial evolutions**: today's agentic LLMs, tomorrow's quantum computers, the day after's general AI — crypto-agile by default
- **Is interoperable and open**: European standards (eIDAS 2, EUDI Wallet), auditable code, multi-stakeholder governance — no vendor-lock black box
- **Is usable by everyone**: identity must not be reserved for geeks. A one-legged, freemason, communist, deaf-mute-from-birth grandmother must be able to manage her 12 contextual identities as easily as a developer. UX-first, native accessibility (screen readers, plain-language, multimodal voice/text/visual support)
- **Puts the citizen at the center**: digital self-sovereignty (*self-sovereign identity*), decentralized identity wallet, granular and revocable consent

Beyond "crying wolf"

I want to address here the cybersecurity experts who have, for years, **been shouting to reduce the impact** of these abuses. Their voice is too often dismissed as yet another *"crying wolf"* — even though the damage is already documented, quantified, and the leaks pile up month after month. We are **sufficiently informed**; the problem is no longer information, it's **staging**.

By constantly shaking the anthill with anxiety and panic — alarmist headlines, apocalyptic metaphors, doom-laden threads without an action plan — **this is all the public retains**: a background hum of dread that ends up anesthetizing rather than mobilizing. And too often, the experts themselves **no longer know what to do**: shout louder? stay silent? produce yet another report? The result is a collective dissonance: everyone knows, no one acts, and the next breach is greeted with a shrug.

We must therefore **organize our communications as a crisis-management operation at the scale of the Country, and indeed of Europe**: a unified cell (ANSSI, CNIL, ANTS, ENISA, European counterparts) with a **shared communication doctrine**, messages calibrated by audience (citizen, executive, elected official, journalist), a **single public registry** of incidents and their remediation status, and **concrete action windows** published alongside every alert ("here are the 3 things to do within 24 hours"). No more panic without a plan, no more report without follow-up, no more isolated expert shouting into the void.

Real ant-work, if I may put it that way — patient, coordinated, distributed, where every grain-bearer counts, and where the column advances because it walks **together and in the same direction**.

Sixty-eight million ants. Forty-five to sixty million complete identity profiles already in the wild. Industrial LLM tooling available for twenty dollars a month. And behind the horizon, quantum computing that will obsolete our current signature schemes.

The kick has already landed. The anthill has been overturned. The question is no longer just *who rebuilds the ground* — it is **with what materials, what governance, and what vision of what it means to be a "digital ant" in the 21st century**.

10 / 11

This work will not be done by a single agency, a single government, or a single company. It will be done by **all of us, together, or it will not be done at all**.

Références / References: - CNIL — Violations de données notifiées 2023–2025 : cnil.fr/fr/violations-de-donnees - ANTS — Titres sécurisés : ants.gouv.fr - HaveIBeenPwned — Breach database : haveibeenpwned.com - Cybernews — National Public Data analysis (août 2024) - Emsisoft — Snowflake breach report (2024) - DINUM — FranceConnect threat model (2024) - Darkpaper — MCP : L'Entrée des Hackteurs (série, 2026) : article-mcp-douane-llm

---

Darkpaper · Statut: draft · 2026-05-27

11 / 11