COHORT STUDY · DELEGATED-ACTION SYSTEMS

What Happens When a System Acts Without Permission?

Across payments, investing, software, marketplaces and autonomous vehicles, history draws a hard line between a system failing inside the authority it was given and a system acting outside it. Autonomous agents make that line harder to locate than any prior technology did.

PREPARED
26 AUGUST 2026

COHORT
17 CASES, 1974–2026

METHOD
PRIMARY-SOURCE, ADVERSARIALLY TESTED

A personal agent can now read an inbox and reply to it, book a flight, cancel a subscription, negotiate with a vendor over a messaging app, and put a charge on a card. Several products doing exactly this shipped to private access groups during 2026, and the early reports from users describe both the appeal and the discomfort in the same breath.

When a system acts on someone's behalf, there are two different ways it can fail. It can do the thing it was asked to do and do it badly. Or it can do something it was never asked to do at all. Those look similar in a product postmortem. They are not similar at all in the historical record.

This is not a question that arrived with AI. Card networks, bank transfer rails, investment advisers, app stores, connected locks and driver-assistance systems have all been working through versions of it for decades, and in most cases the resolution took years and arrived from a court or a regulator rather than from the market. We looked at how those resolutions came out.

One pattern recurs more consistently than anything else in the cohort, and it does not turn on how much money was lost or how frightening the failure looked.

Losses from actions the user authorized tend to stay with the user. Losses from actions the user did not

authorize are the class that eventually gets pushed onto the provider, the network, or the intermediary.

We state this as the strongest recurring pattern found, not as a legal rule. It is not codified anywhere as a general principle, the regimes that express it were built independently of one another across four decades and two continents, and there are cases where it does not hold cleanly. EVIDENCE-BACKED INFERENCE

What follows is what the pattern is, what supports it, where it breaks, and why autonomous agents represent a genuinely new problem rather than a familiar one at larger scale.

THE AUTHORIZATION BOUNDARY

Two failure paths that history treats as different kinds of event, and the third path that has no established resolution.

AUTHORIZED

The user asked for it

User grants authority → System acts within it → Bad outcome → **User bears the loss**

A losing trade. A regretted purchase. A transfer sent to someone who turned out to be a scammer. No delegated-action regime found in this cohort reimburses these.

NOT AUTHORIZED

The user did not ask for it

User grants authority → **System exceeds it** → Loss or harm → **Liability moves to the provider**

The unauthorized charge, the trade outside the mandate, the debit taken in error, the publication nobody consented to. This is the class that many of the statutes, scheme rules and consent decrees in the cohort were built to absorb.

AUTONOMOUS AGENTS

Which of the two is this?

User instruction? → Third-party content? → Inferred intent? → Injected instruction?

An agent that reads untrusted material in order to act cannot always establish which of these produced the action it took. The boundary that every earlier regime depends on is the thing that becomes ambiguous.

Prior delegated systems each carried a legible authorization signal: a presented card, a signed mandate, a click, an engaged system, a granted discretionary account. Agents consume information while acting, and information can look like instruction.

Finding 01

The authorization line matters more than the outcome

The intuitive assumption is that people and institutions respond to the size of the harm. The cohort does not support that. What predicts where the loss lands is whether the system stayed inside the authority it was given.

Robo-advisory is the clearest illustration because the numbers are so lopsided. A client can lose a great deal of money on trades the adviser was authorized to make and has no claim at all. What the adviser cannot do is act outside the mandate. That is the entire content of the fiduciary obligation, and it is where enforcement lands. Betterment paid a $9 million penalty in April 2023 over disclosure and coding failures in its automated tax-loss harvesting service, affecting roughly 25,000 accounts and about $4 million in forgone tax benefits. SEC order, 18 April 2023 Over the same decade its assets under management went from roughly $5 billion to $56.4 billion. FACT Clients absorbed market losses without complaint and delegated far more money afterward.

Zelle makes the same point from the other direction. The network reimburses unauthorized transactions and takes the position that it is not obliged to reimburse payments the user was tricked into authorizing. A 2022 Senate investigation found that three banks providing full data repaid customers in 9.6% of scam claims. Warren staff report, October 2022 Volume went from $75 billion in 2018 to $806 billion in 2023. FACT Enormous losses stayed with users, and delegation grew tenfold.

Look at what the protective regimes actually cover and the line is visible in the scope language itself. The Fair Credit Billing Act and Regulation E cover *unauthorized* transfers. The UK Direct Debit Guarantee covers payments taken *in error* or without proper authorization. Zelle's own zero-liability policy covers *unauthorized* transactions and says so explicitly. Mercedes accepts liability for crashes while Drive Pilot is *engaged*. Fiduciary duty covers acting *outside the mandate*.

Across the regimes examined here, none broadly guarantees the outcome of an action the user intentionally authorized. EVIDENCE-BACKED INFERENCE

Finding 02

Utility creates delegation. Safeguards stabilize it.

An earlier version of this research concluded that liability assumption accelerates delegation. Adversarial testing broke that, and the replacement is more useful.

Three lines of evidence cut against the acceleration claim. The chronology is usually backwards: the Fair Credit Billing Act arrived in 1974 and the Electronic Fund Transfer Act in 1978, both well after cards were a mass consumer instrument, and the UK's mandatory reimbursement regime for authorized push payment fraud took effect on 7 October 2024 for a payment rail that launched in 2008. FACT In both cases the guarantee was a correction applied to delegation that had already happened.

Second, delegation expands enthusiastically with no guarantee at all. Zelle grew tenfold while leaving scam losses with users. Password managers reached mass adoption with no provider in the category indemnifying anyone against credential theft. Robo-advisory passed a trillion dollars in assets with clients bearing every dollar of market downside.

Third, and most directly, the strongest liability assumption we found produced almost nothing. Mercedes accepts full legal responsibility for the driving task while Drive Pilot is engaged, which is a larger transfer of risk than any payment network makes. The system shipped in Germany in May 2022 as a €5,000 option on the S-Class, limited to approved motorway stretches at up to 60 km/h. Mercedes-Benz release, 6 May 2022 Take-up has been negligible. Volvo's CEO publicly pledged full liability for autonomous mode in October 2015 and no consumer product ever shipped under it. FACT

Liability was not the scarce input in either case. Availability was.

So guarantees, reversal paths, audit logs and liability frameworks are not demand generators. They are stabilizers. They do not cause a delegated system to grow; they stop one specific class of failure from unwinding growth that already exists. Which is exactly consistent with regulators arriving late, after a system is large enough that its unauthorized-action failures become politically visible.

The practical version of this distinction: **delegation growth and delegation stability are different variables and respond to different inputs.** Building the safeguards first does not create users. Not building them does not prevent growth. It determines who fixes the problem later, and historically that has been a regulator rather than the company.

Finding 03

Unauthorized action is the failure class systems eventually socialize

Where societies have built machinery to absorb a delegated-action loss, that machinery consistently addresses the same thing. It is worth seeing how independently the mechanisms arrived at the same boundary.

The US electronic payment regime allocates unauthorized transfer losses to the institution and gives consumers an error-resolution procedure and a stop-payment right on preauthorized transfers. The stated legislative concern was that adoption might be held back by the risk of errors and fraud, and the response was to reassign the loss rather than to promise the transfers would be correct. CFPB, on the purpose of EFTA

The UK Direct Debit Guarantee goes further within a narrower scope. A payer gets a full and immediate refund from their own bank for anything taken in error, before any adjudication, with the collecting business absorbing it through an indemnity claim. The scheme also requires ten working days notice before any change to amount, date or frequency. Bacs processed more than 4.8 billion direct debit transactions in 2023, all covered. FACT

App store in-app purchasing shows the same boundary being enforced retroactively. Amazon launched in-app charges in November 2011 with no password requirement of any kind, and children spent their parents' money. The Federal Trade Commission pursued all three major platforms: Apple settled for $32.5 million in January 2014, Google for $19 million in September 2014, and Amazon litigated, lost, and refunded more than $70 million.

FTC, 2014–2017 Stored credentials survived the episode intact. What changed is that a confirmation step became mandatory, and it was never removed. It got cheaper instead, replaced by biometrics.

That last detail is worth holding onto, because it contradicts a common assumption about where these products end up. In no case in this cohort was a confirmation step added after a failure and later taken away.

Finding 04

Concealment can turn one failure into a second problem

This finding needs its boundaries stated before its content, because the general version of it is both popular and wrong. Concealment is not always worse than the underlying error. Some errors are severe enough that nothing about disclosure changes the outcome.

The narrow version holds well. Where the operational failure was survivable, and where a party with enforcement standing learned of it from someone other than the company, concealment created a separate problem with its own legal machinery.

The mechanism is legal rather than emotional. Operational failures get adjudicated under whatever regime governs the domain, and those regimes are generally bounded, insurable and corporate. Concealment gets adjudicated under obstruction, misprision, securities disclosure and licensing conditions, which are none of those things and which attach to named individuals.

Read the charging documents and the separation is explicit. Uber's former chief security officer was convicted in October 2022 of obstructing an FTC proceeding and misprision of a felony. DOJ; jury verdict, 5 October 2022 Neither charge concerns the breach. The same company had disclosed a materially similar breach in 2015 and faced an investigation rather than a prosecution. Volkswagen pleaded guilty to conspiracy to defraud *and separately* to obstruction of justice for destroying documents, paying a $2.8 billion criminal penalty against total US costs of roughly $25 to $30 billion. DOJ, January and April 2017 A comparable misstatement to regulators about a measured vehicle property, Hyundai and Kia's fuel economy overstatement, resolved at around $100 million. Cruise lost its California permits in October 2023 over footage it did not show at a meeting, which the DMV learned about from another government agency. California DMV order of suspension, 24 October 2023 The collision itself was initiated by a hit-and-run human driver.

Equifax shows the same dynamic in personnel rather than in charges. The company detected the intrusion on 29 July 2017 and disclosed on 7 September. Its chief executive, chief information officer and chief security officer all departed within weeks of the disclosure rather than of the breach. FACT Scoping an intrusion before announcing it is normal practice. What made the silence look like something else was that three senior executives sold about $1.8 million in stock during the window. EVIDENCE-BACKED INFERENCE A company does not control the events that set the interpretive frame around its own silence.

The useful formulation is not that transparency is good. It is that **the disclosure decision can change which body of law adjudicates the failure**, and that decision gets made in the hours after a company learns.

Finding 05

Agents make the authorization line harder to detect

Every regime described above depends on being able to answer one question after the fact: did the user authorize this?

Prior delegated systems made that question answerable, usually by accident rather than design. A card was presented. A mandate was signed and held on file. A transfer was initiated from an authenticated session. A driver engaged a system with a physical control. An investor signed a discretionary agreement. When a dispute arose, there was a record of the moment authority was exercised and a bounded description of what it covered.

Autonomous agents break this in a way that has no clean precedent in the cohort. An agent that acts usefully has to read things: emails, web pages, documents, messages, search results. All of that is content, and content can be shaped to look like instruction. The agent is therefore in the position of receiving instructions through the same channel it receives information, with no reliable way to tell which is which.

This is not a hypothetical risk and it is not an outside characterization. One 2026 personal agent describes it in its own privacy notice, listing among the risks of autonomous features Instinct privacy notice, revised 26 August 2026 the possibility of unintended payments or communications and interactions with third parties who may include misleading instructions intended to influence autonomous agents. FACT

Within a 48-hour window in August 2026, testers of that product publicly documented three distinct instances of the boundary failing, and all three were reported by TechCrunch on 24 August. TechCrunch, 24 August 2026 A founder created a fresh email account and sent instructions to his own primary inbox directing the agent to search it and report back; the agent complied and he deleted his account. A product executive disconnected the agent's access to her email and received a summary of her tax correspondence roughly three hours later, which turned out to involve copies already ingested rather than continuing access. A venture investor reported that the agent sent an email on her behalf without asking. She described the message itself as harmless. FACT These are reports of behavior on specific dates and do not establish the current state of the product.

Note where those testers stopped. Read access was accepted and write access was not. Reservations were accepted and a stored payment card was not. That sorting happened within hours of access and it matches the historical threshold precisely: two of the clearest domains that attracted intervention are money moving to a third party and communication issued as the user.

The reason this matters beyond one product is that it inverts the direction of the problem. Every earlier system had to get better at doing the authorized thing correctly. An agent

has to do that *and* establish what was authorized, using a channel that an adversary can write to.

Evidence

Selected cases

Nine of the seventeen cases examined, chosen for what each establishes rather than for coverage. Every entry traces to a dated public source.

Case 01

US electronic payment liability regime

Fair Credit Billing Act 1974 · Electronic Fund Transfer Act 1978

Statutory caps on consumer liability for unauthorized transfers, a mandated error-resolution procedure, and a stop-payment right on preauthorized transfers. Rulemaking passed to the CFPB in 2011.

Establishes: the boundary in its original codified form. The protection attaches to authorization, not to loss. It arrived after cards were already a mass instrument, which is why we treat it as a correction rather than an enabler.

Case 02

UK Direct Debit Guarantee

Bacs scheme rule · 4.8bn transactions in 2023

Standing permission for a company to vary the amount and date it takes from an account, paired with a full and immediate refund from the payer's bank for anything taken in error and ten working days notice of any change.

Establishes: the most complete reversal machinery in the cohort, and the highest-delegation consumer payment instrument. Whether the Guarantee predated mass adoption of Direct Debit could not be established from primary sources. UNRESOLVED

Case 03

Zelle

2017– · $75bn (2018) to $806bn (2023)

Instantaneous, irrevocable transfers from a checking account. Zero-liability policy for unauthorized transactions and no equivalent for payments the user was deceived into authorizing. A Senate investigation found three banks providing full data repaid 9.6% of scam claims.

Establishes: the authorization boundary in its starkest form, and the decisive counterexample to the idea that guarantees drive adoption. UK regulators imposed mandatory reimbursement on the equivalent domestic problem in October 2024, sixteen years after the rail launched.

Case 04

Betterment

2010– · ~$5bn (2016) to $56.4bn AUM (2026)

Two separable failures. In June 2016 the firm suspended trading for roughly two and a half hours after the Brexit vote without telling retail clients, though it told institutional advisers, and drew a letter from the Massachusetts securities regulator over inconsistent communication. Separately, coding errors and an undisclosed change in scanning frequency between 2016 and 2019 cost about 25,000 accounts roughly $4 million in tax benefits, resolved by a $9 million SEC penalty in April 2023.

Establishes: the closest thing in the cohort to a controlled comparison. The event with real financial harm produced no exodus. The event with arguably zero harm, taken without notification, produced the regulatory letter and the industry reaction.

Case 05

Facebook Beacon

November 2007 – September 2009

Authority granted to post what a user chose to post was used to publish what they did on more than forty other websites, automatically, on by default. The information published was accurate. One plaintiff's surprise engagement ring purchase was broadcast to his network. Opt-in and an apology followed within a month; the program was abandoned in 2009 as part of a $9.5 million class settlement.

Establishes: that accuracy is not the variable. It also establishes the most damaging move available to a delegated system, which is using authority granted in one domain to act in another.

Case 06

In-app purchasing

Apple, Google, Amazon · 2011–2017

Stored credentials extended to purchases made without re-authentication, including inside a window after one password entry. Children spent their parents' money. Apple settled at $32.5 million, Google at $19 million, and Amazon refunded more than $70 million after losing in court.

Establishes: that regulators enforce the authorization boundary retroactively when a product erodes it, and that the resulting confirmation step becomes permanent. What changed afterward was the cost of confirming, not the existence of confirmation.

Case 07

Amazon Key in-home delivery

October 2017

Delegation of physical entry, verified by a connected camera. In November 2017 researchers demonstrated that a Wi-Fi deauthorization attack froze the camera feed on its last frame while leaving the door unlocked, allowing a courier to re-enter unseen. Amazon's durable remedy was a precondition: the service does not unlock the door if the camera is not online.

Establishes: the most transferable safeguard design in the cohort. The system refuses to act when it cannot verify that it is being observed acting.

Case 08

Cruise

2 October 2023 incident · permits suspended 24 October 2023

A pedestrian struck by a hit-and-run human driver was thrown into the path of a driverless vehicle, which braked, stopped, then executed a pullover maneuver dragging her about twenty feet. At a meeting the following day the company showed footage ending at the initial stop and did not disclose the subsequent movement. The DMV learned of it from another agency and received the full video eleven days later.

Establishes: that the suspension order recites a meeting rather than a collision. NHTSA separately imposed a $1.5 million penalty for failing to fully report the crash.

Case 09

Knight Capital

1 August 2012 · counterexample

A router containing discontinued code sent more than four million orders in forty-five minutes attempting to fill 212 customer orders, producing a $440 million realized pre-tax loss and roughly $7 billion in erroneous positions. Disclosed the same day. Recapitalized within five days. The SEC's first Market Access Rule enforcement followed at $12 million, trivial against the trading loss.

Establishes: the boundary of the concealment finding. Immediate and complete disclosure did not save the firm, because the error was unbounded and had no reversal path.

Counterevidence

Where the pattern breaks

These conclusions were tested by looking for cases that would falsify them before they were written up. The following are the strongest failures we found, and they are the reason the findings above are stated narrowly.

Knight Capital

Authorized error, perfect disclosure, company gone

The concealment finding has no purchase here. A $440 million loss inside a single trading session with no reversal path is not survivable regardless of how well it is disclosed, and Knight disclosed immediately. Where the error is unbounded and instantaneous, disclosure quality does not change the outcome. This is why the finding is stated with a survivability condition attached.

Mercedes Drive Pilot

Full liability accepted, delegation did not follow

The most complete transfer of risk from user to provider found anywhere in this cohort, and take-up is negligible. The constraint was availability and operating envelope rather than risk allocation. Volvo's 2015 public liability pledge, made ahead of a product that never shipped, points the same way. Liability assumption was not the scarce input.

Zelle

Weak protection, tenfold growth

If guarantees drove adoption, this should not have happened. An instantaneous, irrevocable instrument that explicitly does not cover the most common loss type grew from $75 billion to $806 billion in five years. This case alone is sufficient to reject the acceleration claim, and it forced the reframing in Finding 02.

CrowdStrike

Exemplary disclosure did not confer immunity

A faulty update crashed 8.5 million Windows hosts on 19 July 2024. The company published a preliminary analysis within days and a full root cause analysis on 6 August. Customer delegation held, with 97% gross retention reported for that fiscal year. It still faced a securities class action, alleging misrepresentation of the adequacy of its testing procedures. The claim relocated from incident handling to pre-incident representations, which is a real limit on what good disclosure can protect against.

Aggregate versus individual

Revocation is loud and rare

The clearest general limit on all of this. Individual users revoke delegation after unauthorized actions, visibly and articulately. Populations mostly do not. Betterment grew roughly elevenfold through a governance controversy and an enforcement action. Amazon, Apple and Google collectively refunded well over $100 million and kept stored credentials. The people who revoke are the ones with enough information to run the test and read the terms, which makes them a rounding error in a mature product and the entire user base in a new one.

Implications

What this means for autonomous personal agents

Four things follow from the cohort, stated only as far as the evidence carries them.

Authority is granted per action class and does not transfer. There is no general trust level in any case we examined. The same person will grant a utility company standing

permission to vary what it takes from a checking account and refuse an investment adviser discretion over a brokerage account. An aggregate count of permissions granted predicts very little. The unit that predicts behavior is the specific class of action, defined by the shape of its consequence rather than by which service it touches.

The historical threshold sits at money moving to a third party and communication issued as the user. Those are the two domains that attracted statutory intervention, and they are where the 2026 agent testers stopped voluntarily within hours. Expect confirmation to persist in both indefinitely. The competitive move, based on what happened after the app store cases, is making that confirmation nearly free rather than removing it.

Reversibility has to mean reversibility of the consequence. Revoking a permission after data has been copied accomplishes nothing, which is the failure a product executive demonstrated publicly in August 2026 and which the product's own privacy notice now describes in a sentence stating that disconnecting an integration does not automatically delete data already collected from it. `Instinct privacy notice, revised 26 August 2026` The distinction between undoing the permission and undoing the effect is the one that matters.

Preconditions that block action when verification is unavailable are the most underused safeguard in the record. Amazon Key's durable fix was not a better camera. It was a rule that the door does not open when the system cannot confirm it is being observed. Nothing in the cohort suggests an equivalent has been built for agents that read untrusted content, and nothing in the cohort tells us what it would look like.

The systems in this cohort spent decades getting better at two things: doing the authorized action correctly, and making the wrong action cheap to undo. Both were hard. Neither required anyone to work out whether the instruction was real, because the authorization signal was a physical or procedural artifact that existed independently of the system's own perception. A card was presented or it was not. A mandate was on file or it was not.

An agent that reads in order to act does not have that. It receives instruction and information through one channel, and a third party can write to that channel. The reports from August 2026 are early and thin, and we would not build a conclusion on three incidents in two days. But they point at the right thing, which is not that the agent performed badly.

Every regime in this study was built to answer a question after something went wrong. The question was always whether the user authorized it. That question has been answerable for fifty years.

The harder problem in front of autonomous agents may not be executing actions correctly. It may be establishing which actions were authorized at all.

METHODOLOGY

Seventeen cases across payments, investing, algorithmic trading, consumer software, marketplaces, connected devices, driver assistance and autonomous vehicles, spanning 1974 to 2026. Cases were selected by structural position rather than by industry: each involved a user granting a system authority to act, and a documented event in which the system either failed inside that authority or acted outside it.

Findings rest on primary sources where available: statutes and regulations, enforcement orders, criminal informations and plea agreements, regulatory suspension orders, SEC filings, company legal documents and product notices, and contemporaneous reporting. Derivative coverage was not counted as independent evidence, and several widely repeated accounts were found to overstate what the underlying evidence showed.

Two conclusions were subjected to a dedicated adversarial pass in which counterexamples were sought before confirmation. One survived narrowed. One did not survive as originally written and was replaced; the earlier claim that liability assumption accelerates delegation is rejected by this study, and the section on where the pattern breaks records the cases that rejected it.

Known limitations. The cohort is built from cases prominent enough to generate litigation, enforcement or sustained coverage, which systematically overweights loud failures and cannot see products where delegation quietly failed to take hold. Most constraints in the record were imposed by regulators and courts rather than by users, which means much of what is often described as trust dynamics is institutional response operating on a different timescale. Personal agents currently have no equivalent regulator. Where a claim rests on self-reported behavior rather than observed behavior, it is labeled as such in place.

SOURCES AND NOTES

01 Fair Credit Billing Act, 1974; Electronic Fund Transfer Act, 1978; Regulation E, 12 CFR 1005. CFPB statement of EFTA purpose, Federal Register, 15 January 2025.

02 Bacs / Direct Debit Guarantee scheme rules. Transaction volume for 2023 as reported by scheme participants.

03 US Senate, staff report on Zelle, office of Sen. Warren, October 2022. Senate Homeland Security and Governmental Affairs Committee hearing, July 2024. Volume figures 2018 and 2023.

04 UK Payment Systems Regulator, Specific Direction 20; APP scams reimbursement requirement in force 7 October 2024. Post-implementation review due 2026.

05 SEC, *In the Matter of Betterment LLC*, Admin. Proc. Rel. IA-6288, 18 April 2023. Brexit trading suspension, 24 June 2016; Massachusetts Securities Division correspondence, October 2016.

06 FTC actions on in-app purchasing: Apple (January 2014), Google (September 2014), Amazon (litigated; refunds announced 2017).

07 *Lane v. Facebook, Inc.*, N.D. Cal.; Beacon launched November 2007, discontinued September 2009.

08 Rhino Security Labs demonstration of Amazon Cloud Cam deauthorization attack, 16 November 2017; Amazon remediation statements, November–December 2017.

01 California DMV, Order of Suspension, Cruise LLC, 24 October 2023. NHTSA consent order and $1.5m penalty, November 2024.

02 SEC, *In the Matter of Knight Capital Americas LLC*, Rel. 34-70694, 16 October 2013. Knight Capital Group Form 8-K, August 2012.

03 US v. Joseph Sullivan, N.D. Cal.; jury verdict 5 October 2022; sentencing 4 May 2023.

04 US v. Volkswagen AG, E.D. Mich., 16-CR-20394; plea accepted and $2.8bn criminal penalty ordered 21 April 2017.

05 Equifax: FTC, CFPB and 50 states settlement, July 2019. Breach detection 29 July 2017; public disclosure 7 September 2017.

06 CrowdStrike root cause analysis, 6 August 2024; Form 8-K and quarterly results, FY2025–FY2027.

07 Mercedes-Benz Drive Pilot sales launch, Germany, 6 May 2022; company statements on liability while engaged. Volvo Car Group statement on autonomous mode liability, 7 October 2015.

08 Instinct (Spear Street Technology, Inc.): terms of service revised 24 August 2026; privacy notice revised 26 August 2026. TechCrunch, 24 August 2026. Tester reports dated 21–22 August 2026 describe behavior on those dates and do not establish the current state of the product.

Evidence labels used throughout: FACT for claims traceable to a dated primary source. EVIDENCE-BACKED INFERENCE for conclusions drawn across sources. UNRESOLVED where the public record conflicts or is incomplete. Conflicts are reported rather than resolved.