Binance, Kerviel and NotPetya: Recovery Beyond the Bad Actor

Julien Laurenceau | Published 2026-10-30 | Updated 2026-10-30

https://pepitedata.com/disaster-recovery-trusted-insider/

PepiteData | LinkedIn: Julien Laurenceau | X: @Julien_Lau



Conceptual illustration: identifying one faulty component does not explain why the surrounding controls failed.

On October 29, 2026, about $6.3 billion left Binance hot wallets, according to Bitquery and TRM Labs. Bitquery reports that Binance attributed the entry to a compromised engineer's laptop. The reported access path raises a question that extends well beyond crypto: how much damage can one trusted identity authorize before independent controls stop it?

Société Générale faced a version of that question in the Kerviel affair. Maersk faced another when NotPetya destroyed its online identity infrastructure. These are different incidents, with different responsibilities and remedies. Together they show why naming the person or machine involved cannot substitute for examining the controls, escalation rules and recovery dependencies that failed.

A recovery plan must work with hostile credentials and distrusted management systems: independent recovery access, protected backups, a tested dependency order and explicit authority to stop operations. A data restore can rebuild service. It cannot reverse stolen funds, replace lost capital or establish who committed a crime.

What do Binance transaction traces establish?

Binance-attributed wallets sent roughly $6.3 billion to attacker-controlled addresses on October 29, 2026, according to Bitquery and TRM Labs. Bitquery counts nine chains over nearly four hours; TRM counts ten and places most outflows within about 20 minutes. These are source-specific reconstructions,

not an audited final loss statement. Their chain coverage and observation times differ.

On Ethereum, TRM describes 50 pre-signed EIP-7702 authorizations pointing a Binance hot wallet at attacker-controlled code. A matching authorization let that code transfer the wallet's assets and sweep later deposits. TRM assesses a compromise of hot-wallet signing infrastructure and attributes the operation to North Korea's TraderTraitor with high confidence. That is TRM's assessment, not a completed public forensic report.

Bitquery reports a compromised laptop and no leaked private key, citing Binance. That distinction matters: an attacker who can make a signing service approve transactions may not need to extract its keys. Key custody, transaction authorization and the ability to revoke an existing delegation are separate controls.

Both reconstructions report about $2.7 billion frozen across issuers and Arbitrum's Security Council. A freeze is not the same as returned customer funds, and balances remaining with an attacker vary with observation time, swaps and prices. The engineering lesson does not require a single final dollar figure: control of an authorized signing path can be enough to move assets.

What does the skeptical Binance reading establish?

The skeptical Binance reading raises control questions but does not establish an inside job. A post by Hannibal999 on October 30 alleges that management staged the incident to conceal liquidity problems. The post provides no independently verified evidence for that allegation. Its interpretation must remain separate from the transaction reconstructions and TRM's attribution.

The post also contrasts a reported $6.3 billion loss with roughly $3.3 billion allegedly observed on-chain. That is not an established accounting discrepancy: the Bitquery and TRM reconstructions both trace roughly $6.3 billion. Before comparing totals, reconcile the address set, chains, timestamps, token prices, gross transfers, recoveries and frozen balances. A social post's lower total does not establish hidden insolvency.

The useful questions concern the authorization path. Which policies did signers enforce? Could one compromised endpoint request withdrawals across chains? Were limits independent of the affected service? Who could suspend signing or revoke delegations? Hardware security modules and multi-party computation can protect keys while still signing harmful requests if authorization rules or their administration are compromised.

I would test those boundaries without assuming either a malicious employee or management involvement. Attribution determines investigative and legal responses. Recovery must also address the narrower facts already observed: an untrusted signing path, potentially hostile credentials, dependencies on external freezes and a need for independent authority to contain the incident.

Why must recovery assume hostile credentials?

Recovery must assume hostile credentials because a malicious employee and an external attacker using a valid employee account can reach the same administrative systems. Those are different causes with overlapping technical consequences. If production administrators can also delete backups, change retention and approve recovery access, one compromised identity can damage both operations and the means of rebuilding them.

The Kerviel affair concerns unauthorized trading and failures of oversight. NotPetya concerns destructive malware, credential reuse and a compromised software supply chain. Neither is interchangeable with a cryptocurrency theft. The common design question is where trust can accumulate without an independent check, and whether the response still works once that trust is withdrawn.

Separate operational access from backup administration, recovery credentials and the authority to halt service. Preserve logs somewhere the suspected account cannot alter. Test the response when a privileged operator is unavailable, compromised or under investigation, including the case where normal management approval cannot be trusted.

What does the Kerviel affair add?

The Kerviel affair shows that individual criminal liability and institutional control failures can coexist. On March 19, 2014, France's Cour de cassation left Jérôme Kerviel's criminal conviction standing but overturned the civil award of the entire €4.9 billion loss. The court required the bank's own faults, which contributed to the fraud and its financial consequences, to be considered.

Société Générale's public account centers on a trader who concealed unauthorized positions of about €50 billion and put the institution at risk. That account should be read alongside the supervisory record. France's Finance Ministry summarized the Commission bancaire's 2008 sanction: a reprimand and €4 million fine for serious, persistent internal control failings that made the fraud possible.

There is therefore a defensible critical reading: Kerviel became the public face of a disaster whose institutional causes extended beyond his conduct. Treating the employee as the complete explanation lets governance and management escape scrutiny. Calling him a scapegoat for a concealed liquidity crisis goes further, however. The cited findings establish contributing control failures, not that he was innocent or that management orchestrated a cover story.

The funding distinction also matters. Société Générale's 2008 registration document records a €5.5 billion capital increase to reinforce its capital base after the fraud. Capital absorbs losses; liquidity meets payments when due. A recapitalization does not, by itself, prove a concealed liquidity shortfall. A continuity review should test both funding needs and capital constraints, without treating them as interchangeable.

For infrastructure teams, I would carry that lesson into the incident review: record the actor's actions and the organization's enabling failures separately. Identify who owned the limits, who received exceptions and who could stop the activity independently. For recovery, name an alternate decision-maker and protect evidence before restoring normal access. Replacing the employee while retaining the same permissions and approval chain leaves those failure paths open.

How did NotPetya disable Maersk?

On 27 June 2017 the NotPetya worm reached Maersk through a hijacked software update channel and encrypted machines across the company within minutes, leaving the world's largest container shipper with no working PCs, desk phones, email or booking systems across 130 countries. NotPetya presented a ransom demand, but paying was not a reliable recovery path: the malware caused destructive damage.

NotPetya spread through M.E.Doc, the Ukrainian tax accounting software that Russian military hackers had backdoored through its update servers months earlier. It combined the EternalBlue network exploit with Mimikatz, a credential dumping tool, so patched and unpatched machines both fell: stolen passwords reached the patched ones. Within hours it had crippled Merck, FedEx's TNT Express, Saint-Gobain, Mondelēz and Reckitt Benckiser, each with nine-figure costs.

At Maersk headquarters in Copenhagen, staff gathered at the IT help desk carrying laptops that read "oops, your important files are encrypted" and demanded $300 worth of bitcoin. Employees ran through hallways telling each other to unplug machines. Key-card gates were paralyzed. Disconnecting the global network took more than two panicky hours, and disrupted booking, communications and terminal operations across the company.

How did Maersk recover its identity data?

Maersk recovered its directory from a surviving domain controller in Ghana, disconnected by a power outage before NotPetya reached it. According to Andy Greenberg's WIRED reconstruction, the recovery team could find backups for nearly every individual server, three to seven days old, but none for its domain controllers. Reconstructing the directory and access rules would have added a separate recovery problem.



Conceptual illustration: recovery access that depends on the failed identity system creates a circular dependency.

Maersk ran about 150 domain controllers, and they synced with one another so that any single one could serve as the copy of the whole directory. Replication did not protect against a common destructive event affecting every reachable replica. The directory was needed to re-establish access to restored services.

The Ghana machine's link was too thin to move several hundred gigabytes in useful time, and no West African staffer held a British visa to carry it. Maersk flew the hard drive in a relay: one staffer to Nigeria, a handoff in the airport, then a six-and-a-half-hour flight to Heathrow with the recovery data in hand luggage.

Maersk's directory replication had provided availability under ordinary failures. It did not provide an independent backup for the common destructive event. A rehearsal should test both backup coverage and the time needed to move the required recovery data to a clean environment.

What did the Maersk rebuild cost?

Maersk chairman Jim Hagemann Snabe reported rebuilding 4,000 servers and 45,000 PCs in 10 days, according to WIRED. Some staff continued rebuilding software for close to two months. Snabe estimated the incident cost at $250 million to $300 million and the shipping-volume reduction at 20 percent during the outage. The initial rebuild and full business recovery were different milestones.

Five months later, Maersk chairman Jim Hagemann Snabe described the effort at the World Economic Forum in Davos. Employees who worked the emergency operation privately suspected the cost figure was low. The reinstallation ran out of a Copenhagen cafeteria turned assembly line, 20 laptops at a time on dining tables, with help desk staff walking the rows with USB sticks for hours.

WIRED reports that Maersk relied on staff improvisation and manual workarounds while systems were rebuilt. The surviving Ghana directory made that work possible. A recovery design should deliberately

provide protected identity data and usable logistics, rather than depend on an unrelated outage preserving a replica.

What must a destructive-attack recovery plan contain?

A destructive-attack recovery plan needs independent access, protected backups, a tested dependency graph and authority to halt operations. Assume production credentials and management systems may be compromised. Validate which data is clean and recoverable before resuming service. The recovery design must cover access, decryption, data movement and application checks, rather than assuming that an intact backup automatically provides a working business service.

What the plan must assume	What that forces you to build	How you prove it works
Every online credential and key may be in hostile hands	Independent break-glass access and a tested credential, key and delegation recovery procedure	Timed rotation drill inside the recovery rehearsal
Every reachable copy may be encrypted or destroyed	Immutable or offline backups of the identity layer and data stores, outside the blast radius	Restore from the immutable copy into a clean environment
Recovery is a dependency graph, not a list	A tested dependency graph starting from independent recovery access, then rebuilding identity, secrets, data and applications as their dependencies require	Full recovery rehearsal with the dependency order recorded and timed
Restore throughput constrains achievable recovery time	Restore targets sized from the measured rate, including shipping media when links are thin	Timed recovery plus validation of the latest clean, recoverable data point
The attacker may still be inside while you recover	A clean recovery network that trusts nothing from the compromised estate	Recovery executed in an isolated segment, not on the production network
Recovery may depend on third parties choosing to act	Counterparty state, custodial balances, stablecoin issuers, frozen funds, mapped into the blast radius and the recovery assumptions	Recovery rehearsal run with every centralized freeze unavailable
The suspected actor may normally approve the response	An alternate decision-maker, independent stop authority and externally preserved evidence	Exercise with the usual administrator and manager unavailable or under investigation

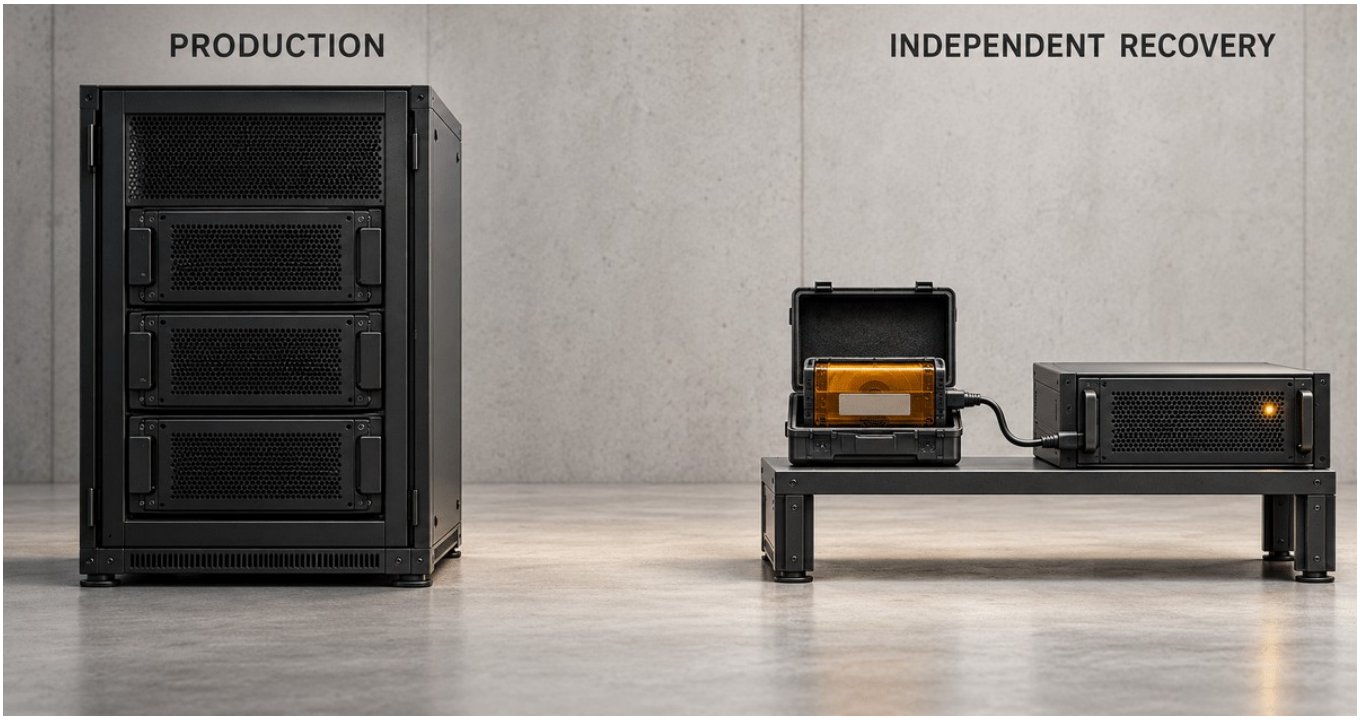

Separate hardware is only part of recovery isolation. Backup administration and recovery access must also be independent of production credentials.

Start by bootstrapping independent recovery access. Identity services, DNS, certificates, encryption keys and backup catalogs can depend on each other. Recover the minimum trusted services needed to unlock and validate backups, then rebuild the wider identity service and applications in the tested order. Production directory replication alone does not provide an independent recovery path.

Separation of duties applies to recovery too. Protect recovery credentials independently, with dual control for destructive or high-impact operations. Production access should not grant backup deletion, retention changes or unrestricted recovery approval. Rehearse the loss of the production directory and the absence of the normal approver together.

External freezes are a containment dependency. Bitquery and TRM report about $2.7 billion frozen across several issuers and Arbitrum's Security Council. A freeze can preserve a recovery opportunity, but return of funds may require separate legal and operational steps. Model unavailable funds, delayed counterparties and cash needs alongside the service recovery exercise.

How do you prove recovery works?

You prove a recovery plan by executing it: restore the identity layer from the immutable copy into a clean environment, time the throughput, and compare measured recovery time with the RTO, and the latest verified clean data point with the RPO. Include authentication, key access, data movement and application validation in the timed exercise. A successful backup job alone cannot establish that the service can resume within its objective.

The Maersk story shows which measurements decide survival. The Ghana copy survived by accident; the account does not establish a zero RPO or the exact age of its last usable directory state. One immediate recovery constraint was bandwidth: several hundred gigabytes could not cross the link to Ghana in time, and the fallback was a hard drive flown through two airports. A rehearsal can expose those dependencies before an incident.

In my resilience and disaster recovery work I measure four things: restore throughput from each backup tier, replication lag under partition, the dependency order of a full rebuild, and whether the stated RTO and RPO are reachable from those measurements. testing the recovery path is what converts a plan into

evidence. The full model of RTO, RPO and replication is in my data platform resilience guide.

What should a CTO check this week?

A CTO should check independent recovery access, protected backup coverage, credential revocation, measured restore time and authority to halt operations. Each check needs an executable test and a named owner. Rehearse with the production directory unavailable and the usual administrator or approver unable to act; otherwise the exercise may depend on exactly the trust it is supposed to replace.

1. When did the identity layer last restore from backup end to end, and was the restore rate measured?
2. Are the backups immutable or offline, and do they include domain controllers, KMS and secret stores?
3. Who holds the break-glass credentials, are they offline, and do two people control them?
4. Does the recovery runbook revoke or replace compromised credentials before reconnecting restored services?
5. Has the restore bandwidth been proven at the rate the RTO demands, including the offline media fallback?
6. Who can independently halt operations, authorize recovery and switch the business to manual workarounds, and when was that rehearsed?

A Resilience and Disaster Recovery Assessment is how I answer those questions with measurements: timed restore tests, replication lag under partition, dependency ordering and a verdict on whether your RTO and RPO are reachable. If the platform also carries performance or scaling questions, the data platform audit covers both in one review. You can also book a 15-minute intro call to discuss which one fits.

FAQ

Can disaster recovery really defend against a malicious insider?

Disaster recovery can reduce the consequences of insider sabotage if protected backups and independent recovery access survive. It cannot guarantee recovery or reverse stolen funds. Prevention, detection, evidence preservation and recovery require separate controls. Test the recovery path with production credentials and the usual approver treated as unavailable or compromised.

Was the Binance hack an inside job?

The sources cited here do not establish an inside job. Bitquery reports that Binance attributed entry to a compromised laptop; TRM attributes the operation to TraderTraitor with high confidence. A social post alleges management involvement without independently verified evidence. A compromised employee account does not, by itself, establish that the employee acted maliciously.

What is a post-crypto-lock restore?

A post-crypto-lock restore rebuilds affected services after ransomware encryption or destructive malware makes them unusable. It starts from immutable or offline backups that the malware could not touch, restores the identity layer first, revokes or replaces compromised credentials and validates the required recovery keys, and only then reconnects restored systems to the network.

How often should we test disaster recovery?

Set the test cadence from business impact, change rate and applicable obligations. As a starting point, I would plan a full rehearsal annually and after material changes, with more frequent restore tests for critical stores. The test must be timed end to end, because the restore rate is what turns your RTO from a target into a measured verdict. Set pass criteria for elapsed time, recoverable data age and application consistency before each exercise.

Are immutable backups enough for disaster recovery?

No. Immutability can protect retained versions against alteration or deletion within its configured scope. It does not establish that the data is clean, complete, decryptable or restorable. Validate backup contents, key availability, independent access, restore throughput and application consistency. Maersk's missing directory backup illustrates why coverage matters as much as retention.

What should we back up first for a full recovery?

Protect the prerequisites for recovery: independent access, directory data, backup catalogs, certificates, DNS, key-management recovery material and required secrets. Follow the supported recovery procedure for each system; some cryptographic keys cannot be exported. Establish a dependency graph that works without production identity, then test data-store and application recovery against it.

What determines achievable recovery objectives?

RTO and RPO are business objectives. Compare the RTO with end-to-end recovery time, including containment, clean infrastructure, data restore and application validation. Compare the RPO with the latest clean, consistent point you can actually recover. Restore throughput affects recovery duration; backup coverage, retained history and compromise timing constrain recoverable data age.

Sources

- Cour de cassation, March 19, 2014, appeal 12-87.416 (criminal conviction maintained; civil damages award overturned to account for contributing faults by the bank)
- French Finance Ministry, DAJ legal newsletter 36, July 10, 2008, Finance section (summary of the Commission bancaire sanction)
- Société Générale, Kerviel affair (the bank's account, an interested party's position)
- Société Générale, 2008 registration document (€5.5 billion capital increase)
- Andy Greenberg, The Untold Story of NotPetya, the Most Devastating Cyberattack in History, WIRED, 2018 (Maersk and NotPetya details in this article)
- Bitquery, Binance Hack: $6.3B Drained on 9 Chains, $2.7B Frozen, Gaurav Agarwal, 30 October 2026 (on-chain traces, timeline and freeze figures)
- TRM Labs, North Korea steals USD 6.36 billion from Binance in the largest hack in history, dated October 29, 2026, figures as of October 30 at 00:30 UTC (attribution, EIP-7702 analysis, totals)
- 0xBobdbldr, Binance transfer thread, X, October 29, 2026 (public reconstruction)
- Hannibal999, skeptical thread on the Binance hack, X, 30 October 2026 (the counter-narrative section; claims in the thread are unverified)

Figures come from public reporting and on-chain traces reviewed on 30 October 2026, while the investigation is still open and several trackers disagree on totals. I have no affiliation with any organization cited here, and none of them reviewed or endorsed this article.

Related posts

- Data Platform Resilience: RTO, RPO and Replication in Practice
- RTO and RPO for Data Platforms: Test the Recovery Path
- MinIO Site Replication: Sync Mode Will Not Save Your RPO
- MinIO Replication Drops Objects Into a Counter Nobody Reads

PDF snapshot of the public article. The online version is authoritative; links remain clickable.