

REGULATORY CONCURRENCY

Why Cyber Incident Reporting Has
Become a System-Scale Problem

2026




Table of Contents


EXECUTIVE SUMMARY	4
THE PROBLEM: REGULATORY CONCURRENCY AS A SYSTEM FAILURE	6
From "More Regulations" to a Different Kind of Problem	6
Three Dimensions of Regulatory Concurrency	7
Why Manual Coordination Fails	8
CONVERTING REGULATORY COMPLEXITY INTO MEASURABLE OUTPUTS	9
Incident Selection Criteria	9
Data Sources	10
Defining a "Reporting Obligation"	10
Obligation Counting Criteria	10
Key Metrics	11
Assumptions	11
Limitations	12
CHANGE HEALTHCARE (UNITEDHEALTH GROUP)	13
Incident Overview	13
Obligation Count Derivation (Step-by-Step)	15
Disclosure Artifacts	15
Fact Evolution	16
Downstream Reporting Multiplier	16
Concurrency Analysis	17
Concurrency Pattern	17
Key Finding and Observations	17
SNOWFLAKE CAMPAIGN (UNC5537)	18
Horizontal Concurrency: Distributed Reporting Across an Ecosystem	18
Incident Overview	18
Obligation Count (Representative Model)	20
Disclosure Artifacts (Per Organization)	20
Fact Stability (Distributed Instability)	21
Key Finding	22
SALESFORCE OAUTH CAMPAIGN (SALESLOFT/DRIFT)	23
Cascading Concurrency: When One Incident Creates Many More	23
Incident Overview	23
Obligation Count Derivation (Cascading Model)	25
Disclosure Artifacts	25
Fact Stability (Dynamic Expansion)	26
Key Finding	27

700CREDIT	28
Definition-Driven Concurrency: When Timing Ambiguity Becomes the Risk	28
Incident Overview	28
Disclosure Artifacts	31
Fact Stability and Definitional Risk	31
Key Finding (Case-Level)	33
SALT TYPHOON	34
Constraint-Driven Concurrency: When Disclosure Limits Become the Coordination Problem	34
Incident Overview	34
Why This Case Is Different	35
Timeline (Publicly Visible Layer)	35
Obligation Count Derivation	36
Disclosure Artifacts	37
Fact Stability Under Secrecy Constraints	37
Dependency and Coordination Model	38
Key Finding (Case-Level)	39
CROSS-INCIDENT ANALYSIS: THE FIVE MODELS OF REGULATORY CONCURRENCY	40
From Five Incidents to One Underlying Pattern	40
The Five Models of Regulatory Concurrency	41
The Unifying Framework: Concurrency Across Three Dimensions	42
A More Useful Mental Model: Concurrency as State Management	43
WHAT THIS MEANS FOR ORGANIZATIONS: WHERE CURRENT WORKFLOWS FAIL	45
The Central Operational Reality	45
The Five Common Failure Points	46
Why This Becomes a Leadership Problem, Not Just a Legal One	49
The Hidden Cost: Defensibility Erosion	49
What This Means in Practical Terms	50
THE OPERATING MODEL SHIFT: FROM MANUAL COORDINATION TO SYSTEM-DRIVEN REPORTING	51
The Core Shift	51
Why Incremental Improvements Are Not Enough	51
The Capabilities Required for Modern Reporting	52
What This Looks Like in Practice	54
Why Automation Is Not Optional	54
CONCLUSION	55
What the Evidence Shows	55
Why This Matters for Leadership	55

EXECUTIVE SUMMARY

Regulatory Concurrency — Why Cyber Reporting is No Longer Human-Scalable

Managing a breach has shifted from a simple checklist to navigating hundreds of overlapping regulatory clocks and unstable facts.



Cybersecurity incident reporting has undergone a structural shift. What was once a compliance exercise governed by discrete rules is now a system-scale coordination problem defined by overlapping regulatory obligations, compressed timelines, and unstable facts.

BreachRx conducted extensive research on five major cybersecurity incidents, **including Change Healthcare, Snowflake (UNC5537), Salesforce (OAuth campaign), 700Credit, and Salt Typhoon**. Based on combined analysis across all five, this research demonstrates that organizations are no longer managing "a breach." They are managing dozens to hundreds of concurrent reporting requirements, each with different triggers, timelines, and evidentiary expectations.

The Findings Are Unambiguous

THE CONCURRENCY EXPLOSION

100+ Concurrent Obligations

(Change Healthcare)

200+ Concurrent Obligations

(Snowflake)


300+ Concurrent Obligations

(Salesforce)

Even mid-scale incidents can exceed 100 obligations due to definitional ambiguity alone (700Credit).

In constrained environments, concurrency persists even with fewer obligations due to external disclosure limitations (Salt Typhoon).

THE SYSTEM FAILURE TRIANGLE



Organizations are no longer managing "a breach"—they are managing hundreds of concurrent obligations.

Modern incident reporting is not human-scalable.

Taken together, this illustrates the challenge of regulatory concurrency that comprises multiple reporting clocks, triggered by different definitions, requiring consistent outputs from evolving facts, under severe time pressure.

The conclusion is clear: Modern incident reporting is not human-scalable using manual coordination models.

Organizations that rely on email, spreadsheets, and static playbooks are structurally exposed to missed deadlines, inconsistent disclosures, regulatory scrutiny, and loss of defensibility months after the incident.

To operate effectively in this environment, incident reporting must be treated as an orchestrated enterprise business process, not a legal checklist.

THE PROBLEM: REGULATORY CONCURRENCY AS A SYSTEM FAILURE

From "More Regulations" to a Different Kind of Problem

Most organizations interpret the current environment as, "There are more regulations than before."

This is incomplete and dangerously misleading. Regulatory obligations are no longer independent requirements. They are overlapping, conditional, time-bound, and interdependent.

The Concurrency Challenge — Overlapping Triggers & Deadlines


EVOLVING REGULATORY TRIGGERS & DEADLINES (EXAMPLES)		
REGULATION	TRIGGER	DEADLINE
SEC Disclosure (Item 1.05)	Materiality determination	4 business days
HIPAA Breach Notification	Discovery of breach	60 days
State Breach Notification Laws	Residency-based identification	30–60 days (varies by state)
FTC Safeguards Rule	Notification event	30 days

Each of these starts at a different moment, requires a different threshold, and demands a different narrative.

In other words, it is a synchronization problem under uncertainty.

Three Dimensions of Regulatory Concurrency

The Regulatory Concurrency Crisis: Beyond the Checklist



From the five incidents we analyzed, regulatory concurrency consistently manifests across three dimensions:

Regulatory Concurrency Dimensions




TEMPORAL CONCURRENCY	DISTRIBUTED CONCURRENCY	THE HIDDEN CONSTRAINT
Organizations must satisfy multiple deadlines simultaneously (e.g., 4 days, 30 days, 60 days). Different clocks start at different times based on different definitions (e.g., SEC materiality vs. HIPAA discovery), creating synchronization challenges.	In supply chain or platform incidents, multiple organizations must independently assess and report. This results in parallel reporting systems operating without central coordination, as seen in the Snowflake and Salesforce events.	Organizations are required to produce stable narratives from unstable facts. Early disclosures must be made when breach scope, attribution, and data impact are often uncertain, creating long-term defensibility risks.

Why Manual Coordination Fails

Traditional incident response assumes linear workflow, single decision timeline, and stable facts before disclosure. None of these assumptions hold.

Teams must track dozens of regulatory obligations simultaneously, coordinate across legal, security, communications, and executives, maintain consistency across multiple disclosure outputs, update positions as facts evolve, and document decisions for future defensibility. Most do this using email threads, spreadsheets, and static playbooks.

This creates several issues.



CONVERTING REGULATORY COMPLEXITY INTO MEASURABLE OUTPUTS

The objective of this research is to convert the qualitative challenge of regulatory concurrency into quantifiable, defensible metrics that demonstrate:

- The scale of reporting obligations triggered by modern incidents
- The structure of those obligations across time and organizations
- The operational burden imposed on incident response teams


Rather than describing complexity, this methodology is designed to measure it in a way that can be compared, validated, and challenged.

Incident Selection Criteria

We selected the five incidents based on three criteria.

Incident Selection Criteria & Concurrency Types

INCIDENT SELECTION CRITERIA	INCIDENT	CONCURRENCY TYPE
Scale of Impact: Millions of individuals or hundreds of organizations	Change Healthcare	Vertical
Regulatory Relevance: Triggers multiple time-bound regimes	Snowflake	Horizontal
Structural Diversity: Represents distinct concurrency archetypes	Salesforce	Cascading
	700Credit	Definition-driven
	Salt Typhoon	Constraint-driven



Data Sources

All findings are derived from publicly available primary and secondary sources, including:

- SEC filings (e.g., Form 8-K disclosures)
- Regulatory guidance (HHS OCR, FTC, FCC)
- State Attorney General breach repositories
- Government statements (FBI, CISA)
- Threat intelligence reports (e.g., GTIG, Mandiant)
- Company disclosures and notices

Where possible, primary filings are prioritized over summaries; dates are taken from signed or published documents; and counts are taken from official disclosures or regulator records.

Defining a "Reporting Obligation"

A core challenge is defining what counts as a discrete obligation. For this analysis, a reporting obligation is defined as a distinct, time-bound requirement to notify a regulator, authority, or affected party under a specific legal or contractual framework.

To determine this, we included SEC disclosures (e.g., Form 8-K), federal regulatory filings (HHS, FTC, FCC), state-level regulator notifications (e.g., AG offices), consumer notification obligations (by jurisdiction), and media notification obligations (where required). We excluded internal communications, duplicate recipients of the same obligation, and voluntary disclosures not tied to a legal requirement.

Obligation Counting Criteria

For U.S.-based incidents with national impact, each state is treated as a distinct regulatory obligation. This includes consumer notification requirements and state Attorney General (or equivalent) notifications.

The standard model is 50 states + Washington, D.C. = 51 jurisdictions. Therefore, consumer and regulator notices each equal up to 51 obligations.

For Federal obligations, each applicable federal regime is counted separately.

For incidents involving multiple organizations (e.g., Snowflake, Salesforce) each affected organization is treated as an independent reporting entity, obligations are calculated per organization, then aggregated.

This is critical as concurrency scales multiplicatively, not additively.

As an example, Change Healthcare had **107 obligations** (or more as explained below), which is modeled as:

- 1 SEC obligation
- 1 HIPAA (HHS) obligation
- 1 HIPAA individual notification obligation
- 1 HIPAA media notification obligation
- 1 HIPAA (business associate notification to covered entities)
- 51 state consumer notifications (50 states + Washington, D.C.)
- 51 state regulator notifications (50 states + Washington, D.C.)

Key Metrics

To move beyond raw counts, several derived metrics are used.

Total Reporting Obligations – The total number of discrete, time-bound notifications required per incident. This is the primary metric used throughout the paper.

Concurrency Structure – Each incident is classified based on how obligations scale:

- Vertical – obligations over time (Change Healthcare)
- Horizontal – obligations across organizations (Snowflake)
- Cascading – one incident creates others (Salesforce)
- Definition-driven – ambiguity drives concurrency (700Credit)
- Constraint-driven – external limits shape reporting (Salt Typhoon)

Downstream Reporting Multiplier – The number of secondary entities required to perform their own reporting based on the incident.

Fact Stability Lag – The time between initial disclosure and stabilization of core facts.

Concurrency Density – Number of active reporting clocks divided by the time to factual stability. This metric is directionally useful but limited by incomplete public timestamps and ambiguity in "stability" definition. It is used qualitatively rather than as a primary quantitative output.

Assumptions

To enable consistent modeling, the following assumptions are applied.

Nationwide Exposure – For large-scale U.S. incidents, it is assumed to trigger obligations across all 50 states and Washington D.C, which is justified by reported scale (millions of individuals) and evidence of multi-state filings.

Obligation Independence – Each regulatory requirement is treated as independent and non-substitutable, even if the same facts are reused and the same team produces the output.

Minimum Defensible Counts – Where uncertainty exists, counts are presented as "Modeled" (based on strong inference) or "Lower bound" (based only on confirmed data).

Limitations

This analysis is intentionally conservative, but several limitations remain. This includes:

- 1. Incomplete Public Data** – not all filings are publicly available and internal reporting is not visible
- 2. Jurisdictional Ambiguity** – exact resident distribution is often unknown and state-level modeling is inferred from scale
- 3. Regulatory Interpretation Variability** – definitions may vary and organizations may interpret obligations differently
- 4. Multi-Entity Overlap** – in distributed incidents, some obligations may overlap across entities while operational burden remains multiplicative

Despite limitations, the findings are meaningful because they are grounded in primary sources, use consistent modeling logic across all cases, favor conservative estimates where uncertainty exists, and are directionally validated across multiple incidents.

Most importantly, the exact number is less important than the order of magnitude. Whether an incident triggers 120 or more than 200 obligations, the conclusion remains the same:

The problem exceeds human-scalable coordination capacity.

With the methodology established, we now apply this framework to each incident in detail—starting with the most illustrative case of vertical concurrency.

CHANGE HEALTHCARE (UNITEDHEALTH GROUP)

Executive System Blueprint

AT-A-GLANCE SUMMARY

Incident:

Change Healthcare

Concurrency Type:

Vertical

Obligations:

107

Duration:

17+ months

CONCURRENCY PATTERN VISUAL: EXPANDING TIMELINE

Day 1:
High
Uncertainty

HIPAA

SEC

State

Media

17
Months

OBLIGATION COUNT TABLE

Category	Count
Federal:	5
State:	102
Total:	107 (Conservative Model)

KEY INSIGHT



The most consequential disclosures were made when the least information was available.

Incident Overview

The Change Healthcare cyberattack represents the clearest example of regulatory concurrency within a single organization over time.

- **Discovery:** February 21, 2024
- **Public Disclosure:** February 22, 2024 (SEC Form 8-K)
- **Final Reported Impact:** ~192.7 million individuals
- **Incident Duration (reporting lifecycle):** February 2024 through July 2025

The incident began as a system intrusion and operational disruption but rapidly evolved into one of the largest healthcare data breaches in U.S. history.

What makes this case uniquely important is not just scale, but how the regulatory burden evolved over time while facts were still developing.

Timeline of Key Events

DATE	EVENT	REGULATORY IMPLICATION
Feb 21, 2024	Unauthorized access identified; systems isolated	Potential start of HIPAA "discovery" clock
Feb 22, 2024	SEC Form 8-K filed	Materiality determination → 4-day SEC clock satisfied
Apr 2024	Public updates on data impact	Early-stage scope uncertainty
Jul 19, 2024	OCR breach report filed (initial 500 placeholder)	HIPAA reporting obligation formally triggered
Aug 2024 – Jan 2025	Rolling notifications and updates	State + individual notification expansion
Jul 31, 2025	Final update (~192.7M individuals)	Closure of reporting lifecycle

The first disclosure occurred **within 1 day**. The final scope was not known for **over a year**. This gap is the essence of vertical concurrency. Early deadlines must be met, while facts continue to evolve for months.

Regulatory Mapping (Full View)

REGIME	TRIGGER CONDITION	CLOCK	APPLIES
SEC (Item 1.05)	Material cybersecurity incident	4 business days	Yes
HIPAA – Covered Entity Notification	Discovery of PHI breach	60 days	Yes
HIPAA – HHS OCR	≥500 individuals	60 days	Yes
HIPAA – Media Notification	≥500 residents (per jurisdiction)	60 days	Yes
State Data Breach Laws	PII exposure (by residency)	30–60 days	Yes
Contractual / Partner	Business associate obligations	Variable	Yes

Each of these clocks starts at a different moment, requires a different definition of "knowledge", and cannot be synchronized.

Obligation Count Derivation (Step-by-Step)

This is where the case becomes quantitatively defensible.

Federal Obligations

- SEC disclosure = 1
- HIPAA (HHS OCR) = 1
- HIPAA (individual notification) = 1
- HIPAA (media notification) = 1
- HIPAA (business associate notification to covered entities) = 1

Subtotal = 5

State-Level Obligations (assuming nationwide impact)

- Consumer notifications = 51
- State regulator (AG) notifications = 51

Subtotal = 102

Additional Obligations

- Substitute notice (for unidentified individuals) = 1
- Downstream delegated notifications (covered entities relying on Change) = implicit multiplier (not double-counted)

Final Modeled Count

Total = 107 concurrent reporting obligations (conservative model excluding some derived obligations)

Disclosure Artifacts

This incident required a minimum of 10 distinct artifact types. Each required legal review, fact consistency, and version control.

1. SEC Form 8-K (initial)
2. SEC amendments
3. OCR breach report
4. OCR updates (multiple)
5. State AG submissions (multiple states, multiple updates)
6. Consumer notification letters
7. Substitute notice mechanisms
8. Public communications (website updates, press releases)
9. Covered entity communications
10. Internal executive/board reporting

Fact Evolution

When the incident was first reported, the placeholder estimate was that **500 individuals** were affected. The final reported number was approximately **192.7 million individuals**.

At the time of early disclosures, scope is unknown, data mapping is incomplete, and deduplication is unresolved. Yet organizations are required to make SEC disclosures, begin regulatory coordination, and prepare for downstream reporting.

The key insight with this incident is that the most consequential disclosures were made when the least information was available.

Downstream Reporting Multiplier

Change Healthcare operated as a business associate to many covered entities. From Washington state evidence, at least **55 distinct data-owner** entities were identified in a single filing. OCR explicitly notes that covered entities remain responsible for their own reporting, and their timelines depend on information from Change Healthcare.

The dependency level was high. Reporting obligations did not stop at Change Healthcare, but rather propagated across the healthcare ecosystem.

Judgment & Decision Pressure

Critical decisions made under uncertainty:

FACTOR	DECISION/QUESTIONS
Materiality (SEC)	☉ Determined within 24 hours ☉ Based on incomplete scope
Scope Inclusion	☉ Which individuals were affected? ☉ How to count duplicates? ☉ When to update totals?
Delegation Model	☉ Who reports: Change vs. covered entities? ☉ When do downstream clocks start?
Notification Timing	☉ When is "discovery" achieved? ☉ When is information sufficient to notify?

The key risk here is that these decisions must be defensible months later, after facts have changed.

Concurrency Analysis

Concurrency Type: Vertical

- Single organization
- Multiple overlapping clocks
- Long-duration reporting

Characteristics

 DIMENSION	 OBSERVATION
Time	Extended (17+ months)
Organizations	Centralized + downstream
Fact Stability	Very low early, improves slowly
Coordination Burden	Extremely high

Concurrency Pattern

- **Early phase:** high uncertainty and fast deadlines
- **Mid phase:** expanding scope plus parallel obligations
- **Late phase:** reconciliation plus defensibility risk

Key Finding and Observations

The Change Healthcare incident triggered approximately **107 (or more) concurrent reporting obligations** across federal and state regimes, with **disclosure beginning within 1 day of discovery** and **continuing for over 17 months as core facts evolved**—demonstrating a level of time-based regulatory concurrency that exceeds the capacity of manual coordination.

It illustrates that:

1. **Speed does not reduce complexity** – Early SEC compliance did not reduce downstream burden.
2. **Facts evolve faster than reporting frameworks** – Reporting continued for over a year.
3. **Centralized systems create dependency bottlenecks** – Downstream entities relied on Change’s data.
4. **Concurrency is sustained, not momentary** – It persists long after initial disclosure.

Change Healthcare shows how concurrency expands over time within a single organization. Next, we examine the opposite structure through the Snowflake incident, where concurrency expands across organizations simultaneously.

SNOWFLAKE CAMPAIGN (UNC5537)

Horizontal Concurrency: Distributed Reporting Across an Ecosystem

Executive System Blueprint

AT-A-GLANCE SUMMARY

Incident:

Snowflake (UNC5537)

Concurrency Type:

Horizontal

Obligations:

209 (Lower bound from 3 victims)

Impact:

~165 independent organizations

OBLIGATION COUNT TABLE

Obligations:

103 (Live Nation / Ticketmaster)

Obligations:

103 (Advance Auto Parts)

Obligations:

3 (AT&T)

Obligations:

209 (Total Modeled)

CONCURRENCY PATTERN VISUAL



KEY INSIGHT



Shared incidents do not produce shared reporting—each organization must independently generate its own timeline and narrative.

Incident Overview

The Snowflake campaign represents a fundamentally different class of incident. There was not a single breach, a single reporting entity, or a single timeline. Instead, it is a shared-root-cause event, where credential compromise (often via infostealers) allowed threat actors to access multiple independent Snowflake customer environments to exfiltrate data.

Unlike Change Healthcare, there is no central authority coordinating reporting. Each organization needed to investigate independently, determine materiality, and report independently.

Because this is a distributed incident, there is no single timeline. **Instead, we reconstructed a composite timeline:**

Timeline of Key Events

DATE	EVENT	IMPLICATION
May 2024	Mandiant begins notifying Snowflake customers	Awareness begins across multiple orgs
May 20, 2024	Live Nation detects unauthorized access	Organization-specific discovery
May 31, 2024	Live Nation files SEC 8-K	Early disclosure (11 days)
June–July 2024	Multiple organizations disclose incidents	Staggered reporting
July 2024	AT&T files SEC 8-K after DOJ delay	National security constraint example

There was no synchronized "start." **Each organization entered the reporting process at a different time and needed to independently evaluate:**

REGIME	TRIGGER	APPLIES
State Data Breach Laws	PII exposure	Yes (most cases)
SEC Disclosure	Materiality	Yes (public companies)
FTC / Sector Rules	Depending on industry	Case-specific
Contractual Obligations	Customer/partner	Yes

The same underlying incident creates hundreds of independent regulatory mappings, not one.

Obligation Count (Representative Model)

Because a full campaign-wide count is not possible (no complete victim list + data per victim), we use a defensible lower-bound model for **3 of the affected organizations out of approximately 165 impacted**.

LIVE NATION/ TICKETMASTER	ADVANCE AUTO PARTS	AT&T
51 state consumer notifications	51 state consumer notifications	1 SEC disclosure
51 state regulator notifications	51 state regulator notifications	1 telecom regulatory obligation (FCC-CPNI)
1 SEC disclosure	1 SEC disclosure	1 customer notification obligation
103 obligations	103 obligations	3 obligations
209 TOTAL OBLIGATIONS		

209 obligations is not the total—it is a floor. The true concurrency level is orders of magnitude higher.

Separately, Mandiant and Snowflake notified approximately 165 potentially exposed organizations. We treat that as evidence of large-scale coordination burden, not as part of the core regulatory obligation count because the public record does not establish 165 discrete legal reporting duties for Snowflake itself.

Current evidence does not support counting Snowflake as a direct regulatory reporting victim of UNC5537, because public sources attribute the compromises to customer credentials rather than a breach of Snowflake's platform; however, Snowflake later disclosed lawsuits, regulatory investigations, and lawmaker inquiries arising from these incidents.

Disclosure Artifacts (Per Organization)

Each organization must independently generate:

- SEC filings (if applicable)
- State notification letters
- Regulator submissions
- Customer communications
- Internal reports
- Legal documentation

Each organization was responsible for 6–10 artifact classes. Across 165 organizations, this equaled potentially 1,000+ independent disclosure artifacts.

Fact Stability (Distributed Instability)

Unlike Change Healthcare (central instability), Snowflake introduces distributed fact instability. In this case, each organization faces different access patterns, data exposure, timelines, and evidence. One company may confirm exfiltration quickly, another may take weeks, while another may never fully confirm.

The result is that there is no single "truth state" across the incident. Each organization independently decides:

FACTOR	DECISION/QUESTIONS
Materiality (SEC)	☉ Is the incident material? ☉ When does the 4-day clock start?
Scope	☉ Was data accessed? ☉ Was it exfiltrated? ☉ Which customers are affected?
Attribution	☉ Is this part of a broader campaign? ☉ Is attribution required for disclosure?
Notification Timing	☉ When is investigation "complete enough" to notify?

The key risk here is that inconsistent decisions across organizations create inconsistent narratives for the same campaign.

Concurrency Analysis

Concurrency Type: Horizontal

 DIMENSION	 OBSERVATION
Time	Moderate per org (days-weeks)
Organizations	Extremely high (~165)
Fact Stability	Variable across orgs
Coordination Burden	None (fully decentralized)

The concurrency pattern is characterized by parallel investigations and reporting timelines with no shared coordination layer. Unlike Change Healthcare, there was no central dependency, no shared reporting authority, and no unified communication.

Therefore, we classify the dependency level as low (but fragmentation risk is high).

This incident demonstrates that:

1. **Concurrency scales multiplicatively, not additively** – Each new organization adds a full reporting workload.
2. **Decentralization creates fragmentation** – No coordination across entities.
3. **Shared incidents do not produce shared reporting** – Each organization must act independently.
4. **Fact consistency is impossible across organizations** – Different narratives emerge.

Key Finding

The Snowflake UNC5537 campaign **generated at least 209 concurrent reporting obligations** across just **three affected organizations**, while **approximately 165 organizations were notified of potential compromise**—demonstrating a level of organization-driven regulatory concurrency that scales beyond centralized coordination.

Snowflake shows how concurrency expands across organizations. Next, we examine a more complex evolution by covering the Salesforce OAuth Campaign, where concurrency not only spreads across organizations, but creates new incidents through cascading effects.

SALESFORCE OAUTH CAMPAIGN (SALESLOFT/DRIFT)

Cascading Concurrency: When One Incident Creates Many More

Executive System Blueprint

AT-A-GLANCE SUMMARY

Incident:

Salesforce OAuth Campaign

Concurrency Type:

Cascading

Obligations:

300+

Systems:

Expanding dynamically

OBLIGATION COUNT TABLE

Metric	Detail
Per SaaS Customer:	~103
Confirmed Victims Mapped:	3
Cascade Multiplier:	Applied across secondary systems
Total Modeled:	300+

CONCURRENCY PATTERN VISUAL: BRANCHING CASCADE

Primary Access Event

Compromised AWS

Compromised Snowflake

Compromised Tokens

KEY INSIGHT



In a multi-tenant SaaS environment, a single identity compromise can trigger cascading secondary incidents across the ecosystem, exponentially multiplying reporting obligations.

Incident Overview

The Salesforce OAuth campaign represents a third-order concurrency problem where one incident generates additional incidents—each with its own reporting obligations.

In this incident, threat actors compromised OAuth tokens tied to third-party integrations (e.g., Salesloft/Drift). These tokens allowed access to Salesforce environments and data exfiltration across multiple customers.

Critically, attackers searched exfiltrated data for secrets (API keys, credentials, tokens). The incident did not stop at Salesforce. It propagated into AWS environments, Snowflake instances, and other connected systems. It became an incident generator.

Timeline (Multi-Layered)

DATE	EVENT	IMPLICATION
Aug 8–18, 2025	Campaign active	Initial compromise window
Aug 20, 2025	Tokens revoked	Containment event
Aug 26, 2025	GTIG advisory	Public disclosure begins
Sept 4, 2025	PagerDuty disclosure	Victim-level reporting
Sept 12, 2025	FBI FLASH	Government-level guidance

Containment (Aug 20) occurred before full visibility. This creates immediate tension as reporting clocks may start before scope is known, while secondary exposure is still unfolding.

Regulatory Mapping (Per Victim Plus Cascades)

Primary Incident (Salesforce access)

REGIME	TRIGGER	APPLIES
State breach laws	Data exposure	Yes
SEC disclosure	Materiality	Yes (subset)
Contractual obligations	Customer impact	Yes

Secondary Incidents (Credential reuse)

Each discovered credential may trigger:

REGIME	TRIGGER
Cloud provider breach obligations	Unauthorized access
Additional state notifications	New affected individuals
Additional SEC disclosures	If material

The regulatory surface area expands as investigation progresses.

Obligation Count Derivation (Cascading Model)

For this incident, we cannot calculate a precise total due to the unknown number of affected organizations and secondary incidents. Instead, we model a defensible lower bound.

SINGLE ORGANIZATION	MULTI-ORGANIZATION LOWER BOUND	CASCADE FACTOR
Typical SaaS customer impact:	Assume:	Each victim may generate:
51 state consumer notifications	3 confirmed/well-documented victims	Additional incidents (e.g., AWS key exposure)
51 state regulator notifications		Additional reporting obligations
(Plus optional SEC)		
~103 obligations per organization	~300+ obligations	

While Snowflake multiplies across organizations, Salesforce multiplies across organizations AND systems.

Disclosure Artifacts

Per organization:

- Customer notifications
- Regulator submissions
- Incident reports
- Vendor coordination communications
- Security remediation notices

Because of cascading effects, multiple disclosure streams may exist per organization, with each tied to the primary incident and secondary exposure. The estimated artifact load is 6–10 artifacts per incident per organization, with potentially multiple incidents per organization.

Fact Stability (Dynamic Expansion)

Unlike previous cases, facts do not just evolve; they expand. For example, the initial understanding was that Salesforce data was accessed. The later discovery was that AWS keys and Snowflake tokens were exposed and additional systems were compromised.

The result is that the scope of the incident grows as investigation progresses. In other words, the system boundary of the incident is not fixed.

Organizations must decide:

WHAT CONSTITUTES "THE INCIDENT"?	WHEN TO DISCLOSE?	HOW TO SCOPE REPORTING?	RESPONSIBILITY BOUNDARIES?
Salesforce access only?	After primary investigation?	Single unified disclosure?	Vendor vs. customer
Or all downstream impacts?	After cascade analysis?	Multiple disclosures per system?	Platform vs. integration

The key risk is that mis-scoping the incident leads to incomplete or inconsistent reporting.

Concurrency Analysis

Concurrency Type: Cascading

 DIMENSION	 OBSERVATION
Time	Moderate (days–weeks)
Organizations	High
Systems	Expanding
Fact Stability	Decreasing during investigation

Concurrency pattern:

1. Initial breach
2. Data exfiltration
3. Secret discovery
4. Secondary system compromise
5. Additional reporting obligations

This shows that concurrency increases during investigation, not after. There is partial dependency on vendor (Salesforce/Salesloft), yet independent responsibility for customers with new dependencies introduced via compromised systems.

Therefore, we classify the dependency level as medium (with expansion).

Salesforce demonstrates that:

1. **Incidents can create new incidents** – Not a single reporting scope.
2. **Concurrency can grow during response** – Not just at the start.
3. **System boundaries are unstable** – Scope is not fixed.
4. **Identity is the new propagation vector** – Tokens and credentials extend impact.

Key Finding


The Salesforce OAuth campaign **generated over 300 concurrent reporting obligations** across a limited set of affected organizations, **with additional obligations created through cascading secondary incidents**—demonstrating a form of regulatory concurrency that expands dynamically during investigation.

We now move from time-driven concurrency (Change Healthcare), organization-driven concurrency (Snowflake) and cascading concurrency (Salesforce) to a subtler, but equally dangerous case in 700Credit, where concurrency is driven not by scale, but by ambiguity in regulatory definitions.

700CREDIT

Definition-Driven Concurrency: When Timing Ambiguity Becomes the Risk

Executive System Blueprint



Incident Overview

The 700Credit incident is a strong example of definition-driven concurrency—a form of reporting complexity created not by sheer size or ecosystem sprawl, but by ambiguity in the underlying legal triggers.

At a high level, the incident involved:

- Unauthorized access to personal information collected through 700Credit’s web application
- Exposure of highly sensitive data elements, including names, addresses, dates of birth, and Social Security Numbers
- Total affected population of approximately 5.8 million individuals

On scale alone, 700Credit is smaller than Change Healthcare. But from a reporting-workflow perspective, it reveals a different and equally important failure mode:

The difference between incident identification, breach determination and notification milestones can determine whether reporting is timely or late.

That distinction is not academic. It directly affects FTC reporting deadlines, state notification timelines, internal escalation logic, and defensibility months later.

Timeline and Trigger Ambiguity

The public record indicates three different dates that matter:

DATE	EVENT	WHY IT MATTERS
Oct 25, 2025	Incident identified; compromised third-party API tied to 700Credit web application	Earliest publicly reported incident-identification date and potential clock-start reference point
Nov 21, 2025	700Credit began notifying impacted dealership clients	Key downstream notification milestone
Dec 22, 2025	Written notices to affected individuals scheduled to begin	Start of consumer-facing notification

The gap between incident identification (Oct 25) and dealership-client notification (Nov 21) is approximately 27 days. The regulatory issue is that different reporting regimes may look to different trigger points, including identification of suspicious activity, confirmation that personal information was accessed or copied, determination that a notification event occurred, or the start of downstream notifications. If a rule start is tied to the earliest identification date, the reporting clock may look very different than if the organization treats a later breach-determination or notification milestone as operative.

In many incidents, teams assume the reporting problem begins once they "know enough." But regulators often use discovery, reasonable belief, notification event, or similar standards. These do not always align with internal incident milestones.

Regulatory Mapping (Full View)

700Credit appears to trigger multiple categories of obligations:

REGIME	TRIGGER CONDITION	CLOCK	APPLIES
FTC Safeguards Rule	Notification event involving 500+ consumers	30 days	Yes
State Data Breach Laws	PII exposure by residency	30–60 days	Yes
Consumer Notice	Affected individuals by jurisdiction	Varies	Yes
State Regulator / AG Notice	Resident thresholds met	Varies	Yes
Law Enforcement Coordination	Investigative / criminal referral	Variable	Yes

Unlike Change Healthcare or Snowflake, the complexity here is not primarily number of organizations and systems or duration of the reporting lifecycle; it is the need to reconcile multiple timing milestones across legal frameworks: incident identification, breach determination, dealer/client notification, FTC notification, AG notification, and consumer notice.

This case supports a relatively clean modeled count.

FEDERAL OBLIGATIONS	STATE-LEVEL OBLIGATIONS
	For a nationwide incident affecting 5.8 million individuals, it is reasonable to model:
FTC reporting obligation = 1	51 state consumer notification obligations
Law enforcement notification / coordination = 1	51 state regulator / AG notification obligations
Subtotal: 2	Subtotal: 102
FINAL MODELED COUNT: TOTAL = 104 CONCURRENT REPORTING OBLIGATIONS	

This is lower than Change Healthcare or Salesforce, but still far beyond what most organizations can manage reliably with ad hoc workflows.

104 obligations sounds manageable relative to a higher-volume case. It is not because the key burden here is not only the total number of outputs; it is the uncertainty surrounding when those outputs are due.

That means every obligation is wrapped in a second question: **When did the clock actually start?** This doubles the coordination challenge.

Teams must align around the incident timeline, the legal meaning of “discovery,” the evidence supporting that interpretation, and the downstream implications if that interpretation changes later.

Disclosure Artifacts

At a minimum, the incident required:

1. FTC reporting submission
2. State regulator notices
3. Consumer notification letters
4. Law enforcement notification / documentation
5. Dealer / client communications
6. Public-facing notice artifacts
7. Internal legal and executive decision records

As in other incidents, these artifacts are not simple duplicates. They require tailoring by jurisdiction, internal review, and synchronization of facts and dates.

Fact Stability and Definitional Risk

Most cyber incidents struggle with scope uncertainty, including number of affected individuals, data types compromised, and whether exfiltration occurred.

700Credit adds a different layer: uncertainty over which incident milestone should control each reporting clock. The public record includes at least three relevant dates: Oct. 25, when the incident was identified; Nov. 21, when impacted dealership clients began receiving notice; and Dec. 22, when written notices to affected individuals were scheduled to begin. The defensibility challenge is explaining which date controlled each obligation and why.

That means the organization is not only trying to answer what happened and who was affected. It is also trying to answer when it becomes reportable.

This creates three specific risks:

1. **Retroactive lateness** – A regulator may later determine that the Oct. 25 incident-identification date, rather than a later breach-determination or notification milestone, started the applicable clock.
2. **Inconsistent timelines across disclosures** – One filing may use Oct 25; another may use Nov 21.
3. **Defensibility failure** – Months later, the organization may be unable to explain why one date and not the other was treated as operative.

700Credit also presents a meaningful, though smaller, dependency problem. The company's role in the dealership ecosystem means that reporting is not purely internal. There is a layer of downstream coordination involving dealerships, clients, and delegated or centralized notification models.

A vendor-led notification model can reduce duplication, but only if all parties agree on the operative timeline, rely on the same facts, and use the same jurisdictional assumptions. If not, the model creates conflicting notices, staggered state submissions, and inconsistent representations to affected individuals.

Therefore, we classify the dependency level as medium–high, not because of scale equal to Change Healthcare, but because the reporting workflow sits across organizational boundaries.

Judgment & Decision Pressure

The 700Credit case turns on a small number of deceptively difficult decisions.

DECISION	CONSIDERATIONS
Which date starts each reporting clock?	☉ Did the clock begin when the incident was identified on Oct. 25, when the breach was confirmed, when dealership clients were notified on Nov. 21, or under another rule-specific trigger?
When does the FTC clock begin?	☉ At awareness? ☉ At reasonable confirmation? ☉ At formal determination?
Who owns downstream notice timing?	☉ 700Credit? ☉ Dealers? ☉ Both?
How should the timeline be represented publicly?	☉ Single unified date? ☉ Or multiple milestone dates?

The key risk is that these are not merely legal interpretation questions, but rather they are workflow design questions. If an organization cannot make and document these decisions clearly in real time, it will struggle to defend them later.

Concurrency Analysis

Concurrency Type: Definition-Driven

 DIMENSION	 OBSERVATION
Time	Moderate (with 27 days from incident identification to dealership-client notification and 31 days from dealership-client notification to scheduled consumer notice)
Organizations	Moderate (vendor plus downstream dealer ecosystem)
Fact Stability	Mixed
Legal Trigger Clarity	Low

The pattern of complexity shows fewer total obligations than larger incidents, but higher uncertainty around clock start and greater risk of retrospective scrutiny on timing. This case shows that concurrency is not always created by size. Sometimes it is created by ambiguity in the legal definition of readiness.

700Credit demonstrates that:

- 1. Timing ambiguity can be as dangerous as scale** – 104 obligations are hard enough; 104 obligations with disputed start dates are much harder.
- 2. Definitions create operational burden** – Incident identification, breach determination and notification-event standards are not just legal terms—they determine workflow deadlines.
- 3. Defensibility depends on timestamped reasoning** – The organization must be able to explain why one date controlled.
- 4. Mid-scale incidents can still exceed manual capacity** – The reporting challenge is not reserved for the largest breaches.

Key Finding (Case-Level)

The 700Credit incident triggered approximately **104 concurrent reporting obligations** across federal and state regimes, while a **27-day gap** between incident identification and dealership-client notification created substantial pressure around clock-start determination—demonstrating a form of regulatory concurrency driven by timing interpretation and defensibility rather than size alone.

We have now examined four distinct forms of concurrency (vertical, horizontal, cascading, and definition-driven).

The final case (Salt Typhoon) adds a fifth model where concurrency is shaped not by growth, ambiguity, or scale, but by external constraints on what can be disclosed at all.

SALT TYPHOON

Constraint-Driven Concurrency: When Disclosure Limits Become the Coordination Problem

Executive System Blueprint

AT-A-GLANCE SUMMARY

Incident:

Salt Typhoon

Concurrency Type:

Constraint-Driven

Obligations:

~10 (Visible minimum)

Systems:

Highly restricted

OBLIGATION COUNT TABLE

Metric	Detail
Federal/Telecom Constraints:	Yes
Public Visibility:	Intentionally incomplete
Total Modeled:	~10 (visible minimum)

CONCURRENCY PATTERN VISUAL: RESTRICTED VISIBILITY LAYERS

Layers (inner to outer):

Public (dashed)

Customer

Carrier Internal

Government

KEY INSIGHT



The operational consequence is that an organization may know more than it can safely say, while still needing to preserve a defensible record for the future.

Incident Overview

Salt Typhoon represents the clearest example in this research set of constraint-driven regulatory concurrency.

At a high level, the incident involved:

- Suspected compromise of major telecommunications infrastructure linked to Chinese nation-state threat actors
- Unauthorized access to sensitive telecom environments
- Theft of call-record-related data
- Compromise of limited private communications
- Potential access to information tied to lawful surveillance processes

Unlike the other cases in this research, Salt Typhoon was not primarily a classic consumer breach narrative. It was a national-security-sensitive infrastructure compromise affecting a highly regulated sector.

That changes the reporting problem fundamentally. In most incidents, the challenge is how to disclose quickly enough. In Salt Typhoon, the challenge is how to satisfy reporting obligations when full disclosure may be operationally or legally constrained.

Why This Case Is Different

Salt Typhoon is not the largest incident in this set by obligation count, the clearest by timeline, or the easiest to model quantitatively. Therefore, it introduces a fifth concurrency archetype – constraint-driven concurrency – where reporting complexity is created by limitations on what can be said, when it can be said, and to whom it can be said.

These limitations arise from national security concerns, active law-enforcement operations, intelligence sensitivities, telecommunications-sector rules, and the possibility that public transparency itself may create additional risk.

That means the reporting challenge is one of controlled execution under restricted visibility.

Timeline (Publicly Visible Layer)

Because much of the operational detail is not public, the timeline must be reconstructed from government and carrier disclosures.

DATE	EVENT	REPORTING SIGNIFICANCE
Oct 25, 2024	FBI/CISA public statement	Public confirmation of investigation; affected companies reportedly notified
Nov 13, 2024	Additional FBI/CISA details	More specific acknowledgement of data compromise
Dec 2024	Government hardening/guidance activity	Indicates sustained sector-level coordination
Jan 10, 2025	Major carrier states incident contained	Public carrier-level operational milestone
Apr 24, 2025	FBI requests additional tips/public assistance	Ongoing national-security posture

The public record shows a meaningful lag between government acknowledgement, technical containment guidance, and carrier-level public completion messaging.

That lag illustrates the central problem of this case: **external observers cannot easily determine when critical regulatory triggers were met.**

Regulatory Mapping (Full View)

Salt Typhoon sits at the intersection of several obligation layers:

REGIME	TRIGGER CONDITION	CLOCK / STANDARD	APPLIES
FCC / CPNI Rules	Breach involving protected telecom information	Reasonable determination; prompt law enforcement notice	Yes
Law Enforcement Notification	Telecom-sector breach / national security compromise	Immediate / urgent coordination	Yes
Public Carrier Disclosure	Materiality or public disclosure judgment	Context-specific	Likely for some entities
Customer Notification	Depending on impact and sensitivity	Context-dependent	Selective / partial
National Security Constraints	Intelligence / operational risk	Overrides public transparency norms	Yes

In most breaches, reporting complexity comes from **more clocks**. In Salt Typhoon, complexity comes from clocks that may start based on nonpublic thresholds, obligations satisfied through nonpublic channels, and public disclosures that may be intentionally incomplete.

That makes this case difficult to count, but highly important to understand.

Obligation Count Derivation

This is the most conservative modeled case in the research, because public evidence is incomplete by design. We therefore do not attempt a speculative high-end count. Instead, we model a minimum defensible public-obligation set:

1. Carrier reporting into federal law-enforcement channels
2. FCC/telecom privacy-related notification obligations
3. Carrier-level customer or targeted-notification obligations
4. Public disclosure / containment messaging
5. Government-level public statements and guidance coordination

Across the visible record, a conservative modeled total is 10 concurrent obligations. This should be read as not a total lifecycle count, and not a full internal reporting count. It is instead a marker of the publicly defensible minimum complexity.

Compared with Change Healthcare or Snowflake, Salt Typhoon produces a smaller modeled obligation count for three reasons:

1. The public record is intentionally incomplete
2. Some reporting likely occurred through closed government channels
3. National security constraints suppressed broader public disclosure

A lower visible obligation count does not imply a lower operational burden. In fact, the opposite may be true.

Disclosure Artifacts

Although the public artifact set is smaller, the coordination burden is unusually high. Likely artifact categories include:

1. Carrier-to-government incident notifications
2. Law-enforcement coordination records
3. FCC-related reporting documentation
4. Public carrier statements
5. Targeted customer or stakeholder notices
6. Internal executive and legal decision records
7. Technical guidance and remediation communications

What makes these artifacts difficult is not quantity alone, but their sensitivity and segmentation, with some for government only, some for carrier leadership only, some for affected customers only, and some unsuitable for broad public release.

That segmentation increases the risk of fragmentation and inconsistency.

Fact Stability Under Secrecy Constraints

Most incidents in this research suffer from unstable facts because investigation takes time. Salt Typhoon suffers from unstable facts because they may be known internally but cannot be broadly disclosed externally.

That creates a different kind of instability. External stakeholders see partial information, public narratives lag private understanding, and disclosure is shaped by what can be safely released, not merely what is known.

This matters because reporting systems are usually built on the hidden assumption that once facts are known, disclosure can proceed. In Salt Typhoon, that assumption breaks.

The operational consequence is that an organization may know more than it can say, while still needing to coordinate with regulators, reassure leadership, communicate with selected stakeholders, and preserve future defensibility. That is a uniquely difficult reporting environment.

Judgment & Decision Pressure

Salt Typhoon concentrates pressure into a narrower set of high-consequence decisions.

DECISION	CONSIDERATIONS
When is "reasonable determination" reached?	This question is foundational in telecom/privacy reporting, but the operative date may never be public.
What can be disclosed without compromising the investigation?	Public transparency may be constrained by ongoing counterintelligence activity, law-enforcement sensitivity, or broader national-security concerns.
Who should be notified—and how broadly?	The difference between targeted customer notification, broader public notice, and regulator-only reporting becomes unusually consequential.
How should the public record be sequenced?	Organizations may need to notify some parties immediately, delay broad disclosure, and issue carefully bounded public statements later.

The key risk is that the central failure mode is not merely late disclosure, it is misalignment between private knowledge, regulatory reporting, and public narrative.

Dependency and Coordination Model

Salt Typhoon is not highly dependent in the same sense as Change Healthcare. There is no single upstream vendor whose fact pattern for which all others await. But there is still a meaningful coordination structure because carriers depend on government threat information, government depends on carrier telemetry, and public messaging depends on both.

That creates a triangular coordination problem across regulated companies, law enforcement/national security authorities, and affected stakeholders.

We classify the dependency level of this case as medium, not because of downstream consumer-notice multiplication, but because disclosure legitimacy depends on cross-institution coordination.

Concurrency Analysis

Concurrency Type: Constraint-Driven

 DIMENSION	 OBSERVATION
Time	Moderate to extended public disclosure arc
Organizations	Limited visible set, but sector-wide sensitivity
Fact Stability	Internally uneven; externally constrained
Disclosure Freedom	Low
Coordination Burden	High

This case follows a different logic than the other four because of its maximal constraint on communicability.

In operational terms, the incident response team must manage obligations that may be satisfied privately, statements that cannot be complete, and a public record that may always appear partial.

Salt Typhoon shows that concurrency can be created by limits on disclosure, not just the expansion of disclosure. It demonstrates four important points:

- 1. Fewer visible obligations can still create severe reporting complexity** – This is not a simple "more obligations equals more risk" story.
- 2. National security can break normal assumptions about transparency** – Public disclosure may lag not because the organization is unprepared, but because disclosure itself is constrained.
- 3. Reporting systems must handle asymmetric visibility** – Different stakeholders may have access to different parts of the truth at the same time.
- 4. Defensibility matters even when public detail is limited** – Months later, organizations may still need to explain what they knew, when they knew it, and why broader disclosure was limited.

Key Finding (Case-Level)

The Salt Typhoon incident triggered a **comparatively modest number of publicly visible reporting obligations** but did so under significant national-security and law enforcement constraints—demonstrating a form of regulatory concurrency in which the central challenge is not volume, but the controlled coordination of incomplete public disclosure.

At this point, the pattern across all five cases is clear. Regulatory concurrency is not one thing. It can't be reduced to a single metric. It is a family of operational failure modes that express themselves differently depending on time, scale, dependency, fact stability, and disclosure freedom.

The next section synthesizes these five cases into a single analytical model and shows why they point to the same conclusion: **modern incident reporting has outgrown manual coordination.**

CROSS-INCIDENT ANALYSIS: THE FIVE MODELS OF REGULATORY CONCURRENCY

From Five Incidents to One Underlying Pattern

At first glance, the five cases in this research appear materially different:

- Change Healthcare is a healthcare mega-breach with prolonged downstream reporting.
- Snowflake is a multi-victim credential-compromise campaign.
- Salesforce is a supply chain/identity compromise with cascading effects.
- 700Credit is a financial-services-adjacent breach with clock-start ambiguity.
- Salt Typhoon is a telecom/national-security case constrained by secrecy.

They differ by sector, attack type, victim structure, and regulatory regime. And yet, they all show that the reporting problem grows faster than the organization's ability to coordinate it manually.

Regulatory concurrency is not a niche issue, but a generalized operational condition that emerges whenever four things occur at once:

1. Multiple obligations exist
2. The obligations do not start at the same time
3. Facts are unstable
4. Different stakeholders require different outputs

Every case in this research satisfies those conditions. What varies is how the problem scales.

The Five Models of Regulatory Concurrency

The five incidents reveal five distinct but related concurrency archetypes.

5 Models of Concurrency


 MODEL	 FEATURE	 INCIDENT
 Vertical	Obligations over time	Change Healthcare
 Horizontal	Obligations across organizations	Snowflake
 Cascading	One incident creates others	Salesforce
 Definition-driven	Ambiguity drives concurrency	700Credit
 Constraint-driven	External limits shape reporting	Salt Typhoon

THE 3-DIMENSION DIAGNOSTIC				
Model	⦿ High	⦿ High	Low	-
Time	⦿ Medium	⦿ Medium	Low	-
Scale	⦿ High	⦿ Medium	Low	-
Fact Stability	⦿ High	Low	-	Low

Regulatory concurrency is not five different problems.
It is a distributed system problem.

The Unifying Framework: Concurrency Across Three Dimensions

Although the five models differ, they can all be understood through three core dimensions.



The future of incident reporting is not better coordination. It is better systems.

- 1. Time** – How long does the organization remain under active reporting pressure?
 - Time matters because reporting burden is not a single event. It is a sustained state of coordination.
- 2. Organizational Scale** – How many organizations must independently execute reporting workflows?
 - Scale matters because each organization introduces separate legal interpretation, separate timeline decisions, separate communications, and separate risk.
- 3. Fact Stability** – How stable are the facts when reporting obligations begin?
 - Fact stability is the hidden driver in all five cases. It determines whether narratives remain consistent, whether decisions can be reconstructed, and whether early disclosures become liabilities later.

	TIME	SCALE	FACT STABILITY
Change Healthcare	Very long duration	One core organization, many dependent entities	Very low early stability
Snowflake	Shorter per entity, but repeated across many entities	Many primary reporting organizations	Inconsistent stability across victims
Salesforce	Moderate duration with mid-investigation expansion	Many primary and secondary reporting organizations	Decreasing stability as scope expands
700Credit	Moderate duration, high timing sensitivity	Vendor plus downstream dealer ecosystem	Timeline instability more than scope instability
Salt Typhoon	Duration shaped by controlled disclosure	Limited visible organizations, sector-sensitive coordination	Internal facts may exceed publicly communicable facts

A More Useful Mental Model: Concurrency as State Management

Traditional compliance thinking treats incident reporting as a list of obligations, identifying the rule, determining applicability, and submitting the notice.

That framing is increasingly obsolete. A more accurate mental model is incident reporting as state management across a distributed system.

In this model, each regulator is a node with its own timing expectations, each stakeholder requires a different output, facts are a changing system state, and teams must keep outputs synchronized as state changes.

This is why manual methods fail. Humans are being asked to perform the equivalent of distributed state synchronization using email, spreadsheets, meetings, and narrative memory.

That is not a scalable operating model.

Why Obligation Count Alone Is Not Enough

The obligation counts in this research are useful because they give a tangible measure of burden.

CASE	OBLIGATION COUNT
Change Healthcare	107 (or more)
Snowflake	209 (lower bound)
Salesforce	300+ (lower bound)
700Credit	104
Salt Typhoon	~10 (visible minimum)

Raw counts, however, do not fully capture risk. A 104-obligation case can be more operationally dangerous than a higher-volume case if the clock starts are disputed, downstream responsibilities are unclear, or the public narrative cannot match private knowledge.

This is why the research does not argue that the biggest number is always the hardest case. Instead, it argues that organizations need to understand what kind of concurrency they are facing and which dimension is most likely to break first.

What Breaks First in Each Model

Each concurrency archetype has a characteristic failure mode.

 CONCURRENCY MODEL	 PRIMARY FAILURE MODE
 Vertical	Version drift over time
 Horizontal	Fragmentation across organizations
 Cascading	Scope explosion during response
 Definition-driven	Clock-start misalignment
 Constraint-driven	Narrative asymmetry under restricted disclosure

This is an important practical insight.

If organizations know which model they are in, they can predict what decisions will become bottlenecks, which artifacts will become inconsistent, and where defensibility will break down later.

Across all five models, manual coordination cannot reliably maintain a synchronized reporting state across time, teams, jurisdictions, and outputs. This is the true common denominator.

The issue is structural, i.e. too many moving parts, too many outputs, too many time-sensitive dependencies, and too much instability. When this happens, organizations risk generating a record that cannot later be defended.

Taken together, the five cases show that regulatory concurrency is not industry-specific, attack-type-specific, not reducible to one regulation, and not solvable through incremental process tuning.

It is a systems problem characterized by overlapping clocks, unstable facts, distributed responsibility, and escalating coordination burden.

That is why the central conclusion of this research is that modern incident reporting has outgrown the coordination capacity of manual workflows.

If that is true—and the evidence from each case suggests it is—then the next question is:

What breaks inside current operating models, and what must change to make reporting defensible at scale?

That is the focus of the next section.

WHAT THIS MEANS FOR ORGANIZATIONS: WHERE CURRENT WORKFLOWS FAIL

The Central Operational Reality

Most organizations do not experience reporting failure as a single dramatic event. Instead, failure appears in more subtle—and often more dangerous—forms:

- A filing that was technically on time but based on unstable assumptions
- A notice letter that differs from a regulator submission
- A timeline that cannot be reconstructed six months later
- A business unit that believed legal was handling the issue while legal believed the vendor was still validating facts
- A downstream partner that files based on stale or incomplete inputs
- A board briefing that cannot be reconciled cleanly with the public record

These are the predictable byproducts of a reporting environment in which obligations multiply, facts evolve, and responsibility is distributed.

The problem is that even strong teams are often operating within a workflow model built for a simpler era. Most current reporting workflows are still built around a set of assumptions that no longer reflect reality.

Assumption 1: The organization can establish the facts before major disclosures begin

That assumption may have been directionally true when reporting clocks were longer, obligations were fewer, and incidents were more contained. It is no longer true in many modern cases.

As Change Healthcare demonstrates, the first major disclosure can occur within a day, while meaningful fact stabilization takes months. In the case of Snowflake, some organizations will begin disclosure while other affected organizations are still determining whether they were compromised at all. And as Salesforce shows, the set of relevant facts may continue expanding during response.

In other words, reporting now begins before incident understanding is complete.

Assumption 2: Reporting follows a mostly linear sequence

The implicit model in many organizations still looks like this:

1. Detect the incident
2. Investigate
3. Determine scope
4. Notify relevant parties
5. Close out reporting

That sequence is increasingly obsolete. The workflow is now parallel, recursive, and state-dependent.

In practice, organizations now investigate while drafting disclosures; notify while scope is still changing; coordinate with regulators while customer communications are still being revised; and document decisions after the fact because the reporting process outpaced the governance process.

Assumption 3: A central team can manually manage the reporting state

Many organizations still rely on incident counsel, privacy counsel, security leadership, communications, and executive stakeholders to maintain shared awareness through meetings, email chains, spreadsheet trackers, and ad hoc review loops.

That model fails when the number of deadlines increases, the number of outputs multiplies, or the truth changes faster than human coordination can propagate updates. Teams can work around the clock and still fail to preserve a synchronized reporting state.

The Five Common Failure Points

Across the cases analyzed, five recurring organizational failure points emerge.

System Failure Diagnostics
WHERE THINGS BREAK



FAILURE BY CONCURRENCY TYPE

TYPE	Vertical	Horizontal	Cascading	Definition	Constraint
FAILURE NODE	Version drift	Responsibility gaps	Artifact sprawl	Clock misalignment	Decision loss

1. Clock-Start Misalignment

This occurs when different teams are operating from different assumptions about when the reporting clock began.

Examples include security treats awareness as preliminary, legal treats discovery as the operative trigger, the business assumes the timer starts when scope is confirmed, and a regulator later adopts a different interpretation altogether.

This is especially acute in 700Credit, Salt Typhoon, and any incident involving sector-specific thresholds like "reasonable determination."

The immediate risk is timing confusion. The longer-term risk is retrospective indefensibility.

2. Version Drift Across Disclosures

Version drift occurs when different outputs reflect different fact states.

Examples include an SEC filing that references preliminary scope, a regulator notice that uses updated counts, a board summary that uses a working estimate, a customer letter that uses narrower language for legal reasons, and a downstream partner communication that reflects yet another version.

This is almost inevitable when facts are changing, documents are being drafted in parallel, and there is no enforced system of record.

Change Healthcare is the clearest illustration of this pressure over time, but every case in the research shows some version of it.

3. Responsibility Gaps at Organizational Boundaries

Concurrency becomes especially fragile where one party depends on another.

Examples include covered entities depending on a business associate, customers depending on a vendor, a telecom carrier depending on government input, and dealers depending on a service provider.

These boundaries create a persistent reporting risk because each party assumes another party is closer to the facts—or closer to the obligation.

That creates delayed escalation, duplicated work, conflicting disclosures, or unowned tasks. The more dependencies exist, the less likely manual coordination is to hold.

4. Inability to Reconstruct Decision Logic Later

This is one of the most underestimated failure modes.

During an incident, teams often know that a decision was made, that the decision seemed reasonable at the time, and that several people were aligned.

Months later, when challenged by a regulator, a plaintiff, a board committee, or internal audit, the organization often struggles to answer the following questions:

- What evidence supported the decision?
- What facts were known at the time?
- Who approved it?
- What alternatives were considered?
- Why was one timeline used instead of another?

This is especially dangerous in definition-driven and fact-instability-driven cases, because the organization may later appear inconsistent even when it was acting in good faith.

5. Artifact Proliferation Without Control

Once concurrency rises, organizations are no longer managing "the disclosure." They are managing numerous artifacts, including filings, legal work products, customer communications, executive briefings, and internal decision records.

Each of these is a separate artifact, with different recipients, different timing, and different legal consequences. Without disciplined orchestration, artifacts proliferate faster than they can be synchronized. As a result, organizations end up with inconsistent data counts, conflicting dates, stale statements, and a fragmented evidentiary record.

What Breaks First by Concurrency Type

The most useful executive takeaway is that not every incident breaks the same part of the workflow first. Different concurrency models have different failure signatures.

CONCURRENCY TYPE	WHAT BREAKS FIRST
Vertical Concurrency (Change Healthcare)	• Version control over time • Downstream synchronization • Confidence in long-duration reporting consistency
Horizontal Concurrency (Snowflake)	• Coordination across entities • Comparability of reporting decisions • Enterprise visibility into ecosystem-wide exposure
Cascading Concurrency (Salesforce)	• Incident scoping • Responsibility boundaries • The assumption that the relevant system boundary is stable
Definition-Driven Concurrency (700Credit)	• Timeline interpretation • Legal alignment • Confidence in deadline compliance
Constraint-Driven Concurrency (Salt Typhoon)	• Narrative consistency across restricted audiences • Alignment between private and public reporting • Defensibility of selective disclosure

This is important because organizations often over-invest in the wrong controls. For example, a team may focus on template notice letters when the real risk is unresolved clock-start logic; or build a strong public relations review process, when the real failure point is dependency coordination across partners.

Why This Becomes a Leadership Problem, Not Just a Legal One

At lower levels of complexity, reporting can be treated as a legal, privacy, or compliance process. At higher levels of concurrency, that framing breaks because the reporting challenge now affects decision rights, operating tempo, executive confidence, board oversight, and enterprise risk posture.

Boards and executive teams are increasingly exposed to operating model questions such as:

- Are we confident we know which clocks are running?
- Are we relying on assumptions or validated facts?
- Are all outward-facing statements internally consistent?
- Can we explain our decisions six months from now?
- Do we know which partner or vendor dependencies could delay our obligations?

This is where regulatory concurrency becomes a governance issue.

The Hidden Cost: Defensibility Erosion

The immediate cost of concurrency is visible, resulting from things like more counsel time, coordination, and document production. The deeper hidden cost comes from defensibility erosion.

This occurs when the organization's record becomes harder to defend over time because facts changed but documentation did not keep pace, rationales were discussed but not logged, dates were used inconsistently, or internal and external narratives diverged. That is the long-tail risk in nearly every case studied.

A company may survive the initial reporting wave and still face later exposure because the record it created under pressure cannot withstand retrospective scrutiny.

In that sense, concurrency creates a delayed liability because the organization is not judged only on what happened, but on how coherently it documented and communicated what it believed was happening.

What This Means in Practical Terms

If an organization's current reporting process is built primarily on email, spreadsheets, legal trackers, and manually synchronized narrative drafts, then the organization should assume it is vulnerable in at least one of the following ways:

- It will struggle to keep all active clocks visible in one place.
- It will have difficulty ensuring all disclosures use the same fact state.
- It will be exposed to partner/vendor dependency delays.
- It will struggle to preserve a defensible record of why key calls were made.
- It will not know which concurrency model it is in until something breaks.

This is the practical consequence of the research.

The issue is not whether organizations have an enterprise process that can hold reporting state under concurrency. Many organizations today cannot.

The cyber reporting problem many organizations believe they have “covered” is no longer the one they are facing. They need an operating model capable of managing overlapping clocks, preserving fact consistency, coordinating across dependencies, and logging decision logic in real time. This is a different category of response capability.

If current workflows fail because they were designed for a simpler reporting environment, then the next question is what kind of operating model is now required?

That is the focus of the next section.

THE OPERATING MODEL SHIFT: FROM MANUAL COORDINATION TO SYSTEM-DRIVEN REPORTING

The Core Shift

The central conclusion of this research is that **the underlying operating model for incident reporting is no longer fit for purpose.**

Historically, reporting could be managed as a legal workflow driven by counsel and supported by ad hoc coordination across teams. That model assumed a limited number of obligations, relatively stable facts, and a mostly linear sequence of decisions.

None of those assumptions reliably hold today. What has changed is not just scale, but the structure of the problem. Modern incident reporting requires organizations to manage multiple clocks, multiple narratives, evolving facts, and distributed responsibility, all at once.

In essence, it is a system orchestration problem. A more accurate way to understand modern reporting is to treat it as a distributed system that must maintain a consistent state across time, teams, and outputs.

In this system, facts are the system state. Regulators and stakeholders are nodes. Deadlines are time constraints on state updates. Disclosures are outputs derived from state. Decisions are state transitions.

The goal is to ensure that every output reflects a coherent, defensible version of the system state at a given moment in time. That is a fundamentally different requirement.

Why Incremental Improvements Are Not Enough

Many organizations respond to increased complexity by adding more checklists, engaging more outside counsel, building more templates, or increasing meeting cadence.

These actions can help at the margins, but they do not address the structural problem. They do not solve for state synchronization, version control across parallel outputs, real-time propagation of fact updates, or consistent decision logging.


In fact, as concurrency increases, these approaches often increase coordination overhead, slow decision-making, and amplify the risk of inconsistency.

The result is a system that becomes more complex without becoming more reliable.

The Capabilities Required for Modern Reporting

To operate effectively under regulatory concurrency, organizations need a set of capabilities that do not exist in traditional workflows. These capabilities fall into six categories.

Operating Model Shift & Capabilities



1. Real-Time Regulatory Obligation Mapping

Organizations must be able to identify all applicable regulatory regimes at the moment of detection, determine which obligations are triggered, and track each obligation as an active clock.

This requires a dynamic mapping of jurisdiction, data type, industry, and incident characteristics. Without this, teams are always discovering obligations late or operating with incomplete visibility.

2. Centralized "Clock Ledger"

Every active reporting obligation must be visible in one place, tied to a clearly defined trigger, associated with a specific deadline, and continuously updated.

Critically, the system must record why a clock started when it did, not just when it is due. This directly addresses the failure mode seen in 700Credit, which was ambiguity around incident identification, breach determination, and notification milestones, and in Salt Typhoon, which was ambiguity around reasonable determination.

3. Real-Time Evidence Capture

When decisions are made, organizations must capture what facts were known, what evidence supported those facts, and how confident the team was in those conclusions.

This is essential because the organization will later be judged based on what it knew at the time—not what it knows now. Without this capability, decisions become difficult to reconstruct and defensibility erodes.

4. Fact Consistency Enforcement

All outward-facing artifacts must draw from a single, continuously updated source of truth. This requires structured data models for incident facts, version control across changes, and propagation mechanisms that ensure updates flow to all active disclosures.

Without this, version drift becomes inevitable and inconsistencies appear across filings, notices, and communications.

5. Workflow Orchestration Across Stakeholders

Reporting is no longer confined to a single team. It involves security, legal, privacy, communications, executives, and often external partners.

Organizations need defined workflows for each obligation, clear ownership for each step, and visibility into progress across all active tasks.

This is particularly critical in Change Healthcare (downstream entities), Snowflake (multiple independent organizations), and Salesforce (multi-system cascades).

6. Decision Rationale Logging

Every critical judgment call must be explicitly recorded, tied to the evidence available at the time, and associated with the individuals who made or approved the decision.

This includes decisions such as materiality determination, scope inclusion, clock-start interpretation, and disclosure timing.

This is the only way to ensure that the organization can defend its decisions months or years later.

What This Looks Like in Practice

When these capabilities are in place, the reporting process changes in several important ways:

1. **From fragmented to unified** – Instead of multiple spreadsheets, disconnected email threads, and siloed team knowledge, the organization operates from a shared, structured system of record.
2. **From reactive to orchestrated** – Instead of responding to deadlines as they emerge, the organization anticipates obligations, tracks them continuously, and manages them proactively.
3. **From narrative-driven to state-driven** – Instead of drafting disclosures independently, the organization generates outputs from a controlled fact base, ensuring consistency across all artifacts.
4. **From implicit to explicit decision-making** – Instead of relying on institutional memory, the organization logs decisions in real time, creating a defensible record.

Why Automation Is Not Optional

At lower levels of complexity, these capabilities can be approximated manually. At the scale demonstrated in this research, they cannot.

Consider:

- 100+ concurrent obligations
- 10+ artifact types
- Multiple teams
- Evolving facts
- Multiple jurisdictions

Even with highly skilled teams, there is no reliable way to maintain synchronized state across all dimensions without system support. This is a question of feasibility.

Organizations that continue to rely on manual coordination will face increasing exposure to missed or misinterpreted deadlines, inconsistent disclosures, regulatory scrutiny, and post-incident defensibility challenges.

Organizations that adopt system-driven reporting will gain visibility into all active obligations, confidence in fact consistency, the ability to explain decisions under scrutiny, and a more controlled, predictable reporting process.

The shift required is categorical:

- From coordination to orchestration
- From documents to systems
- From memory to recorded state

This is the only model that aligns with the reality demonstrated across all five cases.

CONCLUSION

The five incidents analyzed in this paper are not anomalies. They are early, highly visible examples of a broader shift that is already underway. Reporting obligations are increasing, timelines are compressing, expectations for consistency are rising, and incidents are becoming more interconnected.

Taken together, these trends produce a new baseline operating condition, which is that **regulatory concurrency is now a permanent feature of incident response.**

What the Evidence Shows

Across five distinct cases, spanning different industries, attack types, and regulatory regimes, the same pattern emerges:

- **Change Healthcare** shows how reporting obligations expand over time while facts remain unstable.
- **Snowflake** shows how obligations multiply across organizations from a shared root cause.
- **Salesforce** shows how one incident can generate additional incidents and reporting layers.
- **700Credit** shows how ambiguity in definitions can create risk independent of scale.
- **Salt Typhoon** shows how external constraints can reshape reporting even when obligation counts are lower.

These are different expressions of the same underlying problem.

The most important insight from this research is not the number of obligations, but the structure of the problem. Obligations do not start at the same time, facts are not stable when reporting begins, responsibility is distributed across teams and entities, and outputs must remain consistent across all of it.

That combination creates a condition in which incident reporting behaves like a distributed system under load. And like any such system, it fails when state is not synchronized, updates are not propagated consistently, or decision logic is not recorded.

Why This Matters for Leadership

At the executive level, this shift changes the nature of cyber risk. The risk is no longer limited to whether an incident occurs or how quickly it is contained. It now includes whether the organization can meet overlapping regulatory obligations, whether its disclosures remain consistent as facts evolve, and whether its decisions can be defended months or years later.

This moves incident reporting from a compliance function to a governance and operating model issue.

CISOs, leadership teams, and boards must now ask:

- Do we have visibility into all active reporting obligations?
- Can we ensure consistency across all disclosures?
- Can we explain our decisions under scrutiny?
- Can our current processes hold up under concurrent pressure?

If the answer to any of these is uncertain, the organization is exposed.

The implication is that they need to operate differently. Processes built for fewer obligations, slower timelines, and more stable facts will not scale to 100+ concurrent reporting requirements, evolving incident scope, and distributed responsibility.

The gap cannot be closed by more meetings, more emails, or more templates. It requires a shift to system-driven reporting workflows that can maintain state, enforce consistency, and preserve defensibility in real time.

Regulatory concurrency is a structural transformation in how incident response must be executed.

Organizations that recognize this—and adapt their operating models accordingly—will be able to respond with confidence and defensibility.

Organizations that do not will increasingly struggle to keep pace with the demands placed on them, even when they are acting in good faith.

The future of incident reporting is not better coordination. It is better systems.

**Want to learn more about BreachRx?
Contact us today.**

Learn more



About Us

BreachRx's Rex Platform uses agentic AI to orchestrate enterprise-wide incident response, enabling organizations to manage high-volume attacks, reduce chaos, and make fast, coordinated, defensible decisions across teams.



SOC 2 Type II

Contact Us

info@breachrx.com
linkedin.com/company/breachrx
breachrx.com