Street Art & Urban Creativity | Vol. 12, No. 5s, 2026 | ISSN 2183-9956
Street Art & Urban Creativity. Scientific Journal

Too Concentrated to Fail: A Supply Chain Concentration Risk Framework for Healthcare Digital Infrastructure

SHASHANK BHARADWAJ[1], (SHASHANKBHARADWAJ19@GMAIL.COM)[1]*
RETHISH NAIR RAJENDRAN[2] , (RETHISHRNAIR@GMAIL.COM)[2]
[1]INDEPENDENT RESEARCHER, ASTA CRS INC., UNITED STATES.
[2]TECHNICAL DELIVERY MANAGER, UNISYS CORPORATION, UNITED STATES

Abstract

The February 2024 ransomware attack on Change Healthcare halted claims, eligibility, and pharmacy transactions across much of the United States health system. Ninety four percent of surveyed hospitals reported financial impact, 74 percent reported direct effects on patient care, and recovery ran for months. The event was both a cybersecurity failure and a concentration failure, and concentration is what turned one firm's outage into a sector-wide stoppage. Healthcare had routed a critical share of national administrative transaction flow through a single intermediary, a single-sourcing decision that supply chain management has treated as a first-order risk for decades, without importing the risk management apparatus that discipline built around it. This paper develops that diagnosis into a usable instrument. It maps the intermediary classes through which healthcare digital dependency concentrates, covering claims clearinghouses, pharmacy switches, eligibility hubs, cloud-hosted EHR platforms, and payment processors, and shows why the dependency satisfies the tight-coupling conditions under which local failures propagate systemically. It then constructs the Healthcare Digital Infrastructure Concentration (HDIC) framework: a dependency graph of an organization's critical transaction flows, four measurable metrics adapted from supply chain and utility reliability practice (external market concentration via the Herfindahl-Hirschman index, internal dependency share, non-substitutability expressed as time to switch against survivable downtime, and blast radius across critical functions), a normalized counterparty concentration score with explicit calculation and banding rules, and an N-1 contingency test imported from transmission planning. Governance integration maps the metrics into COSO-aligned risk registers, procurement controls, and board reporting, and an assessment protocol with a defined sampling frame and data governance plan is specified for empirical application. Policy analysis compares healthcare, which has no sector-wide oversight regime for its critical digital intermediaries, with the European Union's Digital Operational Resilience Act, and weighs designation, disclosure, and market-structure options against the efficiency benefits concentration also delivers.

Keywords: concentration risk, healthcare IT, supply chain resilience, clearinghouses, enterprise risk management, N-1 contingency, Herfindahl-Hirschman index, operational resilience, Change Healthcare, third-party risk, cloud computing, information governance.

1. Introduction

On February 21, 2024, a ransomware attack forced Change Healthcare, a UnitedHealth Group subsidiary operating the largest medical claims clearinghouse in the United States, to disconnect its systems. The disruption did not stay inside one company. Hospitals lost the ability to submit claims, verify coverage, and receive remittances, and pharmacies lost the ability to adjudicate prescriptions at the counter. Within three weeks the American Hospital Association reported that 94 percent of surveyed hospitals were experiencing financial impact, 74 percent reported direct effects on patient care, and nearly 60 percent were losing at least one million dollars in revenue per day [1], [2]. Establishing the breach scope took far longer than restoring the transactions. Congressional testimony in May 2024 put the exposure at roughly a third of Americans [3]; the count filed with the HHS Office for Civil Rights on July 31, 2025 was approximately 192.7 million individuals, the largest healthcare breach on the federal record [4].

Federal regulators opened investigations while emergency advance payment programs improvised a bridge for providers whose cash flow had stopped [4].

The instinctive framing of this event is cybersecurity, and the cybersecurity lessons are real. The more consequential diagnosis came from an unexpected quarter. The Treasury Department's Office of Financial Research analyzed the incident in the vocabulary of financial system concentration and observed that the damage radius was set less by the attack than by how much of the sector's transaction flow ran through a single node [5]. That is a supply chain diagnosis. Any manufacturing supply chain professional recognizes the structure on sight: a single-sourced critical input, no qualified second source, long requalification time, no buffer inventory. Supply chain management has treated that configuration as a first-order strategic risk for decades and has built a mature apparatus around it, spanning dual sourcing, supplier concentration metrics, time-to-recovery analysis, and contingency planning [6], [7], [8], [9]. Healthcare administration digitized itself into the same configuration and imported none of the apparatus.

This paper builds the missing apparatus. The argument proceeds in four movements. First, it maps where concentration actually lives in healthcare digital infrastructure, which is harder than it sounds, because dependency runs through intermediaries that most provider organizations never contract with directly and often cannot name. Second, it constructs the Healthcare Digital Infrastructure Concentration (HDIC) framework: a dependency graph of an organization's critical transaction flows and four measurable metrics adapted from supply chain practice and utility reliability planning, aggregated into a counterparty concentration score and an N-1 contingency test. Third, it integrates the metrics into enterprise risk governance and specifies an assessment protocol for multi-organization empirical application. Fourth, it examines the policy layer, where healthcare has no oversight regime for the concentrated intermediaries its operations depend on, and evaluates the options without discounting the efficiency arguments that favor concentration. The contribution is a translation with engineering content: concepts that supply chain management and grid reliability have already operationalized, refitted to the transaction anatomy of American healthcare.

2. LITERATURE SURVEY

2.1 Supply Chain Disruption and Concentration Theory

Both the theory and the toolkit this paper adapts come from the supply chain literature. Sheffi framed disruption risk as a function of vulnerability and redundancy, and argued that single-sourced dependencies convert local failures into enterprise failures [6]. Christopher and Peck treated resilience as a design property built through visibility, velocity, and redundancy rather than retrofitted after an incident [7]. Tomlin compared mitigation against contingency strategies and identified the conditions under which dual sourcing dominates buffer holding [8]. That result carries over to transaction intermediaries with one qualification that matters for the rest of this paper. Inventory has no transactional equivalent, so operational continuity depends on routing redundancy. Financial and procedural buffers do exist, in the form of cash reserves, manual and batch workflows, and documented downtime procedures, but they sustain the organization rather than the transaction, and Section III treats the two categories separately for that reason. Craighead and colleagues showed that disruption severity follows from network density and node criticality rather than event magnitude alone [9]. Simchi-Levi and colleagues supplied the measurement turn with the risk exposure index, which assesses each supplier node by time to recovery against performance impact; that central move, measuring a node by how long the buyer can survive without it, is the one this paper adapts [10].

2.2 Concentration Oversight in Other Critical Sectors

Other sectors that discovered concentrated third-party dependency built oversight regimes around it. The European Union's Digital Operational Resilience Act, Regulation (EU) 2022/2554, is the most developed: it subjects ICT providers designated as critical to the financial sector to direct supervisory oversight, mandates exit strategies and substitutability analysis in contracts, and requires financial entities to maintain registers of their ICT dependencies [11]. The Financial Stability Board issued a parallel toolkit for third-party risk across member jurisdictions [12]. Electric power planning offers an older and sharper instrument: North American Electric Reliability Corporation transmission planning standards require the grid to withstand the loss of any single element, the N-1 criterion, as a binding design constraint rather than an aspiration [13]. United States federal guidance on cyber supply chain risk, NIST SP 800-161, supplies vocabulary for supplier dependency analysis but binds no one in healthcare [14]. Antitrust practice contributes the measurement standard: the Herfindahl-Hirschman index has served as the agencies' concentration metric for decades, and the 2023 merger guidelines treat a post-merger market above 1,800 as highly concentrated while adding a separate presumption for any firm holding more than 30 percent share [15].

2.3 Healthcare Cyber Risk Literature and Policy

Healthcare research has documented the threat side thoroughly. Neprash and colleagues counted 374 ransomware attacks on United States healthcare delivery organizations from 2016 through 2021, exposing the protected health information of nearly 42 million patients, with attack frequency more than

doubling over the period [16]. Later economic analysis linked a hospital ransomware database to Medicare claims and found that attacks reduce hospital volume by 17 to 24 percent in the initial week, and that in-hospital mortality among patients already admitted when an attack begins rises by 34 to 38 percent [17]. Policy has responded on the defense axis, with voluntary cybersecurity performance goals issued in January 2024 [18] and the 405(d) program's recognized practices for organizational cyber hygiene [19]. What none of this addresses is structure. The literature measures attacks on organizations and the policy hardens organizations; neither governs the concentration of the intermediaries between them, which is where the 2024 event showed the systemic exposure actually lives [5]. Perrow's normal accident theory explains why the omission matters: tightly coupled systems with concentrated nodes do not fail gracefully, and coupling in healthcare transaction processing is tight by design, since care delivery, pharmacy dispensing, and provider cash flow all wait synchronously on transaction responses [20].

2.4 The Gap

Table I positions the contribution. Supply chain theory holds the concepts but has not been applied to healthcare digital intermediaries. Sector regimes such as DORA and the NERC standards hold the governance instruments but stop at their sector boundaries. Healthcare cyber research and policy address organizational defense rather than structural concentration. The HDIC framework occupies the empty cell: healthcare-specific, structurally focused, and measurable.

TABLE I. Gap Analysis: Existing Bodies of Work Versus the Requirements of Healthcare Concentration Governance

Body of Work	Concentration Metrics	Contingency Testing	Healthcare Specific	Binding Governance Path
Supply chain resilience [6]-[10]	Yes (REI, sourcing)	Yes (TTR/TTS)	No	Corporate practice
DORA and FSB third-party regimes [11], [12]	Partial (registers)	Yes (exit plans)	No (finance)	Yes (finance only)
NERC planning standards [13]	No	Native (N-1)	No (power)	Yes (power only)
NIST C-SCRM guidance [14]	Partial	Partial	No	Voluntary
Healthcare cyber research and policy [16]-[19]	No	No	Yes	Voluntary
HDIC (this work)	Yes (four metrics)	Yes (N-1 adapted)	Yes	Proposed (Sec. VI)

3. THE ANATOMY OF CONCENTRATION IN HEALTHCARE DIGITAL INFRASTRUCTURE

Concentration in healthcare digital infrastructure hides in the middle of the transaction path. A provider organization sees its EHR vendor, its practice management system, and its bank; it frequently does not see, and has never contracted with, the intermediaries those vendors route through. Five classes carry the load. Claims clearinghouses aggregate, validate, and route the HIPAA-mandated transaction sets between hundreds of thousands of providers and thousands of payers. Change Healthcare processed approximately 15 billion healthcare transactions annually before the 2024 incident, spanning eligibility verification and authorization, prescription transmission, claims transmittal, and payment, and its systems touched one in three patient records [21]; the House Committee on Energy and Commerce put its position at nearly 40 percent of United States medical claims [22]. Neither figure is an audited market share, and Section VI returns to why none exists. Pharmacy switches perform the equivalent function for prescription claims at the point of sale, where adjudication is synchronous and an outage stops dispensing in real time. Eligibility and prior authorization hubs concentrate coverage verification. Cloud platforms concentrate the EHR itself as hosted deployment displaces on-premises installation; class-level hosting shares for healthcare are not published, but the infrastructure layer beneath them is measurably concentrated, with Synergy Research Group placed the three largest providers at 63 percent of worldwide enterprise cloud infrastructure spending in the third quarter of 2025, at shares of 29, 20, and 13 percent [23]. Those three contribute roughly 1,410 index points before the remaining 37 percent is counted, putting the market within reach of the highly concentrated threshold on the leaders alone [15]. Payment and remittance processors concentrate the final financial leg. The classes differ in tempo, since a pharmacy switch outage bites in minutes and a remittance outage in weeks, but they share the property that matters: many organizations, few nodes, synchronous dependence.

TABLE II. Intermediary Classes in Healthcare Digital Infrastructure and Their Failure Characteristics

Intermediary Class	Concentrated Function	Failure Tempo	Coupling	Available Slack
Claims clearinghouses	Claims, eligibility, remittance routing (HIPAA transaction sets)	Days to weeks	Buffered by cash	Dual routing; cash reserves
Pharmacy switches	Real-time prescription adjudication	Minutes	Synchronous	Fallback adjudication procedures
Eligibility and prior authorization hubs	Coverage verification, authorization	Hours	Synchronous	Secondary hub; manual verification
Cloud EHR hosting	Clinical system availability	Minutes to hours	Synchronous	Multi-region design; downtime procedures
Payment processors	Remittance and payment posting	Weeks	Buffered by cash	Manual posting; reserves

Two features make this concentration more dangerous than its industrial analogues. The first is invisibility. In a manufacturing supply chain the buyer chose its single source and can name it; in healthcare, dependency is frequently inherited through vendor subcontracting, so an organization's true counterparty exposure is a fact about its vendors' vendors that no internal system records. During the 2024 incident, many organizations discovered their Change Healthcare dependency only when transactions stopped [4], [5]. The second is coupling. Perrow's conditions for systemic failure, tight coupling and limited slack, hold throughout: dispensing waits on adjudication, admission workflows wait on eligibility, payroll waits on remittance [20]. Slack here has to be read carefully. Routing redundancy is the only slack that keeps transactions flowing; cash reserves, manual and batch workflows, and downtime procedures keep the organization solvent while the transactions are stopped, which is a different and more limited protection [8]. Both must be arranged before the failure and neither can be improvised during it. The framework in Section IV makes both measurable: invisibility yields to dependency mapping, coupling to explicit survivable downtime arithmetic. The 2024 timeline is itself the coupling evidence, with disconnection on February 21, pharmacy disruption within hours, cash flow distress within days, federal payment flexibilities within weeks, and restoration stretched over months, a sequence that tracks each function's buffer exactly as the survivable downtime concept predicts [1], [4], [5]. A sector that wants a different sequence next time has to change the structure that produced this one, and structure is what the framework measures.

4. The HDIC Framework

4.1. The Dependency Graph

The unit of analysis is the organization's critical transaction flow graph. Nodes are the organization's critical business functions, namely claims submission, eligibility verification, pharmacy adjudication, remittance posting, and clinical system availability, together with the external providers those functions traverse, including inherited providers reached through vendor subcontracts. Edges carry two attributes: the share of the function's volume routed through the provider, and whether the dependence is synchronous, meaning the function stops when the provider stops, or buffered, meaning it degrades over a stated interval. Constructing the graph is the first deliverable of an HDIC assessment. The authors' working expectation, formed in Medicaid claims operations, clearinghouse integration, and enterprise delivery engagements across public sector and commercial settings, is that this step yields the largest single increment of new information, because inherited dependencies surface only when someone traces each mandated transaction set from end to end. That expectation is offered as a proposition rather than a finding: the protocol in Section V measures it directly, as the invisibility rate, so it can be confirmed or refuted. The tracing method is prosaic. For each critical function, follow one production transaction from origination to final response, recording every system and organization it traverses, then confirm the route against the vendor contract chain and the remittance trail. Sources include EDI trading partner configurations, vendor subcontractor disclosures, and the operational knowledge of revenue cycle and pharmacy staff, who almost always know the real topology better than the architecture diagrams do.

4.2 Four Metrics

Metric 1, external market concentration (EMC), measures the sector-level structure of each intermediary class with the Herfindahl-Hirschman index computed over transaction volume shares,

$$EMC = HHI = 10{,}000 \times \sum_{i=1}^{n} s_i^2, \quad s_i \in [0, 1] \qquad (1)$$

where n is the number of providers in the class and si is provider i's share of class transaction volume expressed as a fraction between 0 and 1, so a monopolist returns 10,000 and the result is directly comparable with antitrust practice. The thresholds supply the interpretation: a class above 1,800 is highly concentrated, and a single provider holding more than 30 percent share triggers a separate presumption in its own right [15]. The clearinghouse class cannot be scored exactly, because transaction shares are not published, but it can be bounded. Taking the leading firm at the reported range of one third to 40 percent of claims [21], [22] and dividing the remainder among ten equally sized competitors gives a class index between roughly 1,540 and 1,960. The upper end clears the highly concentrated threshold, and the leading firm's share clears the 30 percent presumption at either end. A tail of fewer and larger competitors would raise both figures; a longer tail of smaller ones would lower them. This is an order-of-magnitude bound with its assumptions stated, not a measurement, and the disclosure gap that forces the assumption is the subject of Section VI. EMC is a sector fact an organization cannot change, but it prices the environment: high EMC means qualified alternatives are few and switching queues will be long on the day everyone switches at once.

Metric 2, internal dependency share (IDS), is organization-specific. For each provider and critical function, it is the share of that function's volume routed through the provider, taken from the dependency graph. Where a provider serves several functions, the provider-level value is the criticality-weighted mean over the functions materially dependent on it, with materiality set by a governance threshold on routed share, fixed at 0.05 in the worked example below.

Metric 3, non-substitutability (NSUB), adapts the Simchi-Levi time logic [10]. For each provider and function, the organization estimates the time to switch (TTS) to a qualified alternative under incident conditions and compares it with the survivable downtime (SDT), the interval that function can be unavailable before harm becomes material: days of cash for remittance, hours for eligibility, minutes for pharmacy adjudication. NSUB is the largest value of min(1, TTS / SDT) across the functions the provider serves, so the binding function governs, and the ratio saturates at 1 because a switch that cannot complete inside survivable downtime is, for scoring purposes, no switch at all. The metric is named for what it measures: a rising value means escape is harder, and calling this quantity substitutability inverted its own sense.

Metric 4, blast radius (BR), captures Craighead's insight that severity is a network property [9]. It is the sum of the criticality weights of the functions materially dependent on the provider, reported on a 0 to 5 scale. Weights are assigned on patient-safety and financial-consequence grounds and held constant across providers so scores remain comparable. The worked example uses clinical system availability at 5.0, pharmacy adjudication at 2.0, and claims submission, eligibility verification, and remittance posting at 1.5 each: clinical availability alone saturates the scale because its loss stops care delivery, pharmacy adjudication ranks next because its loss stops dispensing within minutes, and the administrative functions carry equal, lower weight because their failure degrades cash and workflow before it reaches patients.

4.3 Composite Score, Banding, and the N-1 Test

The metrics aggregate into a counterparty concentration score for each provider j,

$$CCS(j) = 100 \times IDS(j) \times [\, BR(j) / BR_{max} \,] \times NSUB(j) \qquad (2)$$

with BRmax fixed at 5. Each factor is bounded by 1, so the product is bounded by 1 and the leading constant places CCS on a 0 to 100 scale without further normalization. The multiplicative form is deliberate. Quantitative risk definition since Kaplan and Garrick treats risk as scenarios carrying a likelihood and a consequence [24]. CCS holds the scenario fixed at complete provider loss and models the consequence, itself the product of how much flow the provider carries, how far the loss reaches, and whether the organization can escape inside its own tolerance. A provider that fails any one of those conditions is not exposed, which is a multiplicative statement rather than an additive one.

Score movement needs an interpretation rule, or the number becomes decoration. Four points govern its use. CCS is ordinal rather than cardinal: 80 denotes greater exposure than 40, not twice the exposure, and cross-organization comparison is valid only where the same criticality weights and survivable downtime values apply, which the Section V protocol standardizes. Scores are banded for action at 0 to 24 monitored, 25 to 49 material, 50 to 74 high, and 75 to 100 critical; the material threshold marks where a provider carrying half a function set's volume across half the criticality-weighted radius with no viable switch would score. Uncertainty sits almost entirely in TTS, so TTS is recorded as a low and high bound and CCS is reported as the resulting interval. That interval is informative in itself: where TTS greatly exceeds SDT the ratio saturates across the range and the interval collapses to a point, so the score is most robust where exposure is worst, while providers whose switching time sits near their survivable downtime carry wide intervals and warrant measurement before management. Action is triggered by a band crossing, by an interval whose upper bound crosses a boundary, by a failed N-1 test at any score, or by a year-over-year rise above ten points. No validation dataset yet exists; the Section V protocol would produce one, and until it does the bands are governance conventions rather than calibrated risk levels.

The organization-level instrument is the N-1 contingency test, adapted from transmission planning [13]. For each provider scoring in the material band or above, the organization must demonstrate, by

documented plan and periodic exercise rather than assertion, that every critical function survives complete loss of that provider for its stated survivable downtime. N-1 converts concentration from a score into a pass or fail engineering question, which is what made it effective in power systems: it is difficult to argue with a contingency you cannot demonstrate. Table III presents a worked example for a hypothetical mid-size health system, constructed to show the arithmetic and asserting nothing about any actual organization.

TABLE III. Illustrative HDIC Assessment for a Hypothetical Mid-Size Health System (Constructed Example, Not Empirical Data)

Provider (Class)	Functions Touched	IDS	BR (0-5)	TTS vs SDT (NSUB)	CCS (interval)	Band and N-1 Result
Clearinghouse P1	Claims, eligibility, remittance	0.85	4.5	45 d vs 21 d (1.00)	76.5 (76.5 to 76.5)	Critical. Fail: no dual routing
Pharmacy switch P2	Adjudication	1.00	2.0	30 d vs 0.5 d (1.00)	40.0 (40.0 to 40.0)	Material. Fail: no fallback
Cloud host P3 (EHR)	Clinical availability	1.00	5.0	90 d vs 3 d (1.00)	100.0 (100.0 to 100.0)	Critical. Fail: single region
Eligibility hub P4	Eligibility (secondary)	0.15	1.5	10 d vs 21 d (0.476)	2.1 (1.5 to 3.0)	Monitored. Pass
Payment processor P5	Remittance posting	0.60	1.5	14 d vs 30 d (0.467)	8.4 (6.0 to 12.6)	Monitored. Pass: manual fallback

4.4 Reading the Illustrative Assessment

The calculation is stated in full for P1 so that every other row can be checked against it. P1 carries 85 percent of the volume of the three administrative functions it serves, giving IDS = 0.85. Those functions carry criticality weights of 1.5 each, so BR = 4.5 and the normalized radius is 4.5 / 5 = 0.90. Time to switch under incident conditions is 45 days against a survivable downtime of 21 days, a ratio of 2.14 that saturates at NSUB = 1.00. Applying (2), CCS = 100 × 0.85 × 0.90 × 1.00 = 76.5. Because the low and high TTS bounds of 30 and 60 days both exceed survivable downtime, the ratio saturates across the range and the interval collapses to a point. The same arithmetic gives 40.0 for P2, 100.0 for P3, 2.1 for P4, and 8.4 for P5.

The example rewards a close read because its arithmetic reproduces recognizable failure patterns. P1 reaches the critical band not because any input is extreme but because the multiplication compounds heavy dependency share, wide radius, and a switching time double the survivable downtime. P2 shows why dependency share alone misleads: dependency is total, yet a narrow radius holds the score to the middle of the material band while the N-1 test fails outright, because minutes of survivable downtime forgive nothing. That case is the reason for keeping N-1 separate from the score. P3 illustrates the cloud exposure flagged in Section III, where total dependency, maximum radius, and a rehosting time no tolerance absorbs produce the ceiling score and the finding that matters most: single-region deployment of the clinical system is the largest concentration exposure on the books. P4 and P5 show the framework declining to cry wolf. Material dependencies with workable substitutes or adequate buffers score low and pass, and both carry visible intervals because their switching times sit inside survivable downtime. Keeping low scores low is what keeps the instrument credible with the executives who must act on the high ones.

5. GOVERNANCE INTEGRATION AND ASSESSMENT PROTOCOL

5.1 *A. Enterprise Risk Management Integration*

The framework is designed to land inside existing enterprise risk machinery rather than beside it. Each provider scoring in the material band or above becomes a risk register entry in the COSO or ISO 31000 aligned structure the organization already maintains. The event is provider loss, the likelihood axis draws on the sector threat record [16], [17], the impact axis is computed from IDS, BR, and daily exposure rather than estimated by workshop consensus, and the control set is enumerable: dual routing, exit and portability clauses, downtime procedures, and cash reserves sized as daily receivables exposure multiplied by time to switch [25], [26]. Board reporting gets what third-party risk almost never has, which is a number that moves. Quarterly reporting of the leading CCS values with their intervals, the N-1 status, and the trend line converts vendor dependency from an annual questionnaire ritual into a managed exposure. The Govern function introduced in version 2.0 of the NIST Cybersecurity Framework asks organizations to establish, communicate, and monitor cybersecurity supply chain risk management outcomes, and CCS with its N-1 result supplies a measurable outcome where the framework otherwise leaves the organization to invent one [27]. Procurement gains leverage at contract time: an HDIC review

before intermediary selection prices concentration into the decision, weighing a cheaper single provider against the CCS it creates, exactly as sourcing strategy weighs unit cost against supply risk in manufacturing [8].

State Medicaid agencies deserve separate mention because they occupy both sides of the exposure. As payers, their managed care organizations and fiscal agents route enrollment, encounter, and payment transactions through the same concentrated intermediaries as commercial flows, so a clearinghouse stoppage propagates into capitation reconciliation, provider payment timeliness rules, and federal reporting obligations at once. As stewards, they hold contractual leverage individual providers lack: managed care contracts and fiscal agent procurements can require HDIC-style dependency disclosure, N-1 exercise evidence, and dual routing capability as conditions of award, converting the framework from voluntary practice into a procurement-enforced standard across a state's program. The 2024 incident forced states into emergency payment flexibilities improvised under pressure [4]; a dependency graph maintained in advance is the difference between improvising and executing.

5.2 Assessment Protocol for Empirical Application

The empirical program is specified so that it can be executed, criticized, and replicated. The sampling frame is United States provider organizations that submit HIPAA X12 transaction sets under their own tax identification number, and the unit of analysis is the organization rather than the transaction or the vendor. Recruitment proceeds through state hospital associations and revenue cycle professional chapters, with purposive stratification on the two dimensions theory predicts will drive variation: organizational scale, across critical access, independent community, and integrated delivery network strata, and routing arrangement, across direct clearinghouse contracts and vendor-mediated routing. The first wave targets a minimum of five organizations with at least one from each scale stratum. Five sites is a feasibility and instrument-validation sample, and its purpose should not be overstated. It can establish that the tracing procedure runs inside a live revenue cycle operation, expose ambiguities in the metric definitions, and bound the invisibility rate to an order of magnitude. It cannot support sector-level inference. Estimating population distributions of CCS requires a probability sample from a national frame such as the CMS Provider of Services file, a second-stage design contingent on the first wave showing that the instrument works.

Three phases follow. The dependency census has participating organizations construct their transaction flow graphs while the protocol records how many material dependencies were unknown before tracing; that quantity, the invisibility rate, is a publishable measurement and the direct test of the proposition stated in Section IV. Scoring computes CCS for every material provider, with time-to-switch estimates grounded in documented switching experience where it exists, including the 2024 incident record, in preference to vendor assertion. The N-1 exercise is a tabletop loss of each organization's highest-scoring provider, with time to restore each critical function measured against its stated survivable downtime. Pre-registered outcomes are the distribution of CCS values across organizations and classes, the invisibility rate, the N-1 pass rate, and the gap between assumed and exercised recovery times.

Data governance requires more care than an earlier framing of this protocol allowed. No proprietary market-share data is needed from intermediaries, which is why the study can proceed without their cooperation, but the protocol is not free of confidential data. It requires organization-internal routing shares, contract chains, transaction configurations, recovery times, and daily receivables exposure, all commercially sensitive and, taken together, a map of where an attacker would strike. That material stays inside the organization. Sites score their own dependency graphs using the supplied instruments, and only derived, de-identified values leave the site under a data use agreement: CCS values and bands, invisibility counts, N-1 outcomes, and recovery-time gaps. Sites are reported under codes, intermediaries by class rather than firm, and cells representing fewer than three sites are suppressed; no dependency graph, routing topology, or contract term is published. Transmitted data are held encrypted with access limited to the named research team, and the study will be submitted for institutional review board determination rather than assuming an exemption.

5.3 Procurement and Contract Controls

Because concentration is created at contract time, the controls with the most leverage are contractual, and the supply chain playbook translates clause by clause. Dual routing capability: the contract permits, and the implementation maintains, live routing of a stated minimum share of transactions through a second qualified intermediary, the transaction equivalent of dual sourcing, which keeps the alternative qualified and the switching path warm [8]. Exit and portability: enrollment data, payer connection configurations, and transaction mappings are exportable in documented formats within stated timelines, which shortens TTS and lowers CCS through the term the organization can actually influence. Continuity obligations: the provider commits to recovery time objectives per transaction set, evidenced by exercise rather than attestation, mirroring the exit plan requirements DORA imposes on financial entities [11]. Transparency: the provider discloses its own material subcontractors, attacking the invisibility problem one tier down. Financial protections: service credits are sized against the buyer's measured daily exposure rather than nominal fees. None of these clauses is exotic, and all are routinely absent from healthcare

intermediary agreements, largely because no one at the table was carrying a number for what the dependency was worth.

6. POLICY DISCUSSION

Above the organizational layer sits a policy question the 2024 incident posed and nothing has since answered. Healthcare has no sector-wide oversight regime for the concentrated digital intermediaries its operations traverse, and the gap stands out against the two comparators used throughout this paper. In finance, the Digital Operational Resilience Act empowers the European Supervisory Authorities to designate critical ICT third parties and oversee them directly [11]. In bulk power, transmission planning standards make survival of any single element loss a binding design requirement [13]. Healthcare has voluntary organizational hygiene goals [18], aimed at a different layer; hardening every hospital does nothing about a single node carrying the sector's transactions. The claim is bounded to those comparators and to the instrument at issue, namely designation and oversight of concentrated intermediaries. Healthcare entities are not unregulated in third-party risk generally, since HIPAA business associate requirements, state contracting terms, and ordinary vendor risk management all apply. None of them measures or limits how much of a sector runs through one firm. Table IV maps the translation.

TABLE IV. Translating Established Concentration Governance Instruments to Healthcare Digital Infrastructure

Instrument (Source Regime)	Source Sector Mechanism	Healthcare Translation
Critical provider designation (DORA) [11]	European Supervisory Authorities designate and directly oversee critical ICT third parties	HHS designation of systemically important health IT entities by transaction share
Dependency registers (DORA) [11]	Financial entities maintain ICT dependency registers	Provider organizations maintain HDIC dependency graphs, with disclosure to regulators
N-1 planning criterion (NERC) [13]	Grid must withstand loss of any single element under a binding standard	Designated entities and large providers demonstrate exercised single-provider loss survival
Concentration measurement (DOJ and FTC) [15]	HHI and share thresholds trigger merger scrutiny	Published class-level transaction share data; EMC monitoring of intermediary markets
Third-party toolkit (FSB) [12]	Cross-jurisdiction supervisory toolkit for third-party risk	HHS and state Medicaid agency joint assessment protocol for shared intermediaries

Three policy instruments merit evaluation. Designation would identify systemically important health IT entities by transaction share thresholds computed with the EMC arithmetic of Section IV, attaching resilience obligations to the designated entities themselves in the form of exercised contingency plans, exit support duties, and incident transparency, which is the DORA model translated. Disclosure would oblige intermediaries to publish transaction share data so that EMC becomes computable by anyone and organizations can price the concentration they are buying into; today those shares are estimated from incident reporting and congressional oversight because no one is required to state them, which is why Section IV could offer only a bounded range [5], [22]. Market-structure remedies, meaning merger review attention to transaction infrastructure and portability mandates that shorten TTS sector-wide, work slowly but attack the root.

Honesty requires the other side of the ledger. Concentration exists because it delivers real efficiencies: one integration instead of dozens, standardized transaction handling, scale economics smaller intermediaries cannot match, and a single node that is, in ordinary times, genuinely reliable. A policy that fragmented transaction infrastructure without portability standards could raise costs and error rates while merely distributing, rather than reducing, aggregate risk. The framework's contribution to the policy debate is therefore not a verdict but a measuring instrument. Designation thresholds, disclosure content, and the cost-benefit arithmetic of structural remedies all require the quantities HDIC defines and the Section V protocol would collect: class-level HHI, dependency shares, and time-to-switch distributions. Regulating what has never been measured is how sectors end up with rules that miss. Measurement first is the sequence the other sectors followed, and it is available to healthcare now.

7. LIMITATIONS AND FUTURE WORK

Limitations and Future WorkSix limitations bound the claims. The framework is proposed and illustrated rather than field-applied; the Section V protocol exists precisely because the invisibility rate,

the CCS distributions, and the N-1 pass rates are empirical unknowns until organizations measure them. Time-to-switch estimates are the softest input, resting on switching experiences that are rare, poorly documented, and incident-dependent; the protocol's grounding requirement and the reported CCS interval mitigate this without eliminating it. The score is uncalibrated, in that no dataset yet links CCS values to realized loss, so the bands in Section IV are governance conventions rather than validated risk levels, and establishing that link is the natural second study. Sector EMC computation is constrained by the disclosure gap Section VI describes, so early applications will carry share estimates with stated uncertainty rather than measurements. The multiplicative form encodes judgments, specifically that dependency, radius, and non-substitutability compound rather than add, which are defensible from the disruption literature [9], [10] and from the structure of quantitative risk definition [24] but deserve sensitivity analysis against additive and maximum-of-terms alternatives. Finally, the framework addresses availability concentration; data concentration, meaning the privacy exposure of a substantial share of a nation's records transiting one node [4], shares the structure but needs its own instrument. Future work includes the cross-sector comparison the framework makes possible, since computing healthcare intermediary EMC against the financial and power sector baselines their regulators publish would locate healthcare's structural risk on a common scale for the first time, and an extension to cyber insurance, where underwriters price organizational controls but not counterparty concentration and CCS distributions would supply the exposure variable the 2024 loss experience showed was missing.

8. CONCLUSION

The Change Healthcare incident is usually filed under cybersecurity, and the filing misses the finding. Ransomware of this kind is ordinary in American healthcare, with 374 attacks recorded against United States healthcare delivery organizations between 2016 and 2021 at a frequency that more than doubled across the period [16]. What made this one a sector-wide stoppage was where it landed: on a node through which the sector had, invisibly and without governance, single-sourced its administrative circulation. Supply chain management would never have let that configuration pass unexamined, and grid planning would have failed it on inspection. This paper translated those disciplines into a healthcare instrument. A dependency graph makes inherited exposure visible, four metrics make it comparable, a composite score with stated calculation and banding rules makes it reportable, an N-1 test makes it falsifiable, governance integration makes it manageable, and a policy analysis makes the sector-level options concrete without wishing the efficiency case for concentration away. The premise is modest and the implication is not: what healthcare cannot name and cannot measure, it will govern only after the next stoppage. The measuring can start before.

References

1. American Hospital Association, "AHA survey: Change Healthcare cyberattack significantly disrupts patient care, hospitals' finances," Washington, DC, USA, Mar. 2024.
2. Reuters, "Almost all US hospitals took financial hit from Change hack, AHA says," Apr. 30, 2024.
3. U.S. Senate Committee on Finance, "Hacking America's Health Care: Assessing the Change Healthcare Cyber Attack and What's Next," hearing with testimony of A. Witty, Washington, DC, USA, May 1, 2024.
4. U.S. Department of Health and Human Services, Office for Civil Rights, "Change Healthcare cybersecurity incident frequently asked questions," Washington, DC, USA, updated Jul. 31, 2025. [Online]. Available: https://www.hhs.gov/hipaa/for-professionals/special-topics/change-healthcare-cybersecurity-incident-frequently-asked-questions/index.html
5. A. Fliegelman, "The cyberattack on Change Healthcare: Lessons for financial stability," Office of Financial Research Brief 24-05, U.S. Department of the Treasury, Washington, DC, USA, Nov. 2024.
6. Y. Sheffi, The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage. Cambridge, MA, USA: MIT Press, 2005.
7. M. Christopher and H. Peck, "Building the resilient supply chain," International Journal of Logistics Management, vol. 15, no. 2, pp. 1-14, 2004, doi: 10.1108/09574090410700275.
8. B. Tomlin, "On the value of mitigation and contingency strategies for managing supply chain disruption risks," Management Science, vol. 52, no. 5, pp. 639-657, 2006, doi: 10.1287/mnsc.1060.0515.
9. C. W. Craighead, J. Blackhurst, M. J. Rungtusanatham, and R. B. Handfield, "The severity of supply chain disruptions: Design characteristics and mitigation capabilities," Decision Sciences, vol. 38, no. 1, pp. 131-156, 2007, doi: 10.1111/j.1540-5915.2007.00151.x.
10. D. Simchi-Levi, W. Schmidt, and Y. Wei, "From superstorms to factory fires: Managing unpredictable supply-chain disruptions," Harvard Business Review, vol. 92, no. 1-2, pp. 96-101,

2014.
11. European Parliament and Council, "Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA)," Official Journal of the European Union, L 333, pp. 1-79, Dec. 27, 2022.
12. Financial Stability Board, "Enhancing third-party risk management and oversight: A toolkit for financial institutions and financial authorities," Basel, Switzerland, Dec. 2023.
13. North American Electric Reliability Corporation, "Standard TPL-001-5: Transmission system planning performance requirements," Atlanta, GA, USA, 2020.
14. National Institute of Standards and Technology, "Cybersecurity supply chain risk management practices for systems and organizations," NIST SP 800-161r1, Gaithersburg, MD, USA, May 2022, doi: 10.6028/NIST.SP.800-161r1.
15. U.S. Department of Justice and Federal Trade Commission, "Merger Guidelines," Washington, DC, USA, Dec. 18, 2023.
16. H. T. Neprash, C. C. McGlave, D. A. Cross, B. A. Virnig, M. A. Puskarich, J. D. Huling, A. Z. Rozenshtein, and S. S. Nikpay, "Trends in ransomware attacks on US hospitals, clinics, and other health care delivery organizations, 2016-2021," JAMA Health Forum, vol. 3, no. 12, art. e224873, Dec. 2022, doi: 10.1001/jamahealthforum.2022.4873.
17. H. Neprash, C. McGlave, and S. Nikpay, "Hacked to pieces? The effects of ransomware attacks on hospitals and patients," American Economic Journal: Economic Policy, vol. 18, no. 1, pp. 256-281, Feb. 2026, doi: 10.1257/pol.20240594.
18. U.S. Department of Health and Human Services, "Healthcare and public health sector-specific cybersecurity performance goals," Washington, DC, USA, Jan. 2024.
19. U.S. Department of Health and Human Services, 405(d) Program, "Health industry cybersecurity practices: Managing threats and protecting patients," Washington, DC, USA, 2023 ed.
20. C. Perrow, Normal Accidents: Living with High-Risk Technologies. Princeton, NJ, USA: Princeton University Press, 1999.
21. American Hospital Association, "Change Healthcare cyberattack underscores urgent need to strengthen cyber preparedness for individual health care organizations and as a field," Washington, DC, USA, 2024. [Online]. Available: https://www.aha.org/change-healthcare-cyberattack-underscores-urgent-need-strengthen-cyber-preparedness-individual-health-care-organizations-and
22. U.S. House of Representatives, Committee on Energy and Commerce, "What we learned: Change Healthcare cyber attack," Washington, DC, USA, May 3, 2024. [Online]. Available: https://energycommerce.house.gov/posts/what-we-learned-change-healthcare-cyber-attack
23. Synergy Research Group, "Cloud market share trends: Big three together hold 63% while Oracle and the neoclouds inch higher," Reno, NV, USA, Nov. 19, 2025.
24. S. Kaplan and B. J. Garrick, "On the quantitative definition of risk," Risk Analysis, vol. 1, no. 1, pp. 11-27, 1981, doi: 10.1111/j.1539-6924.1981.tb01350.x.
25. Committee of Sponsoring Organizations of the Treadway Commission, Enterprise Risk Management: Integrating with Strategy and Performance. New York, NY, USA: COSO, 2017.
26. International Organization for Standardization, "Risk management: Guidelines," ISO 31000:2018, Geneva, Switzerland, 2018.
27. National Institute of Standards and Technology, "The NIST Cybersecurity Framework (CSF) 2.0," NIST CSWP 29, Gaithersburg, MD, USA, Feb. 2024, doi: 10.6028/NIST.CSWP.29.