© DEC 2018 | IRE Journals | Volume 2 Issue 6 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/ire1725088

The Evolution of High Availability and Disaster Recovery Standards in Healthcare Database Systems: A Systematic Literature Review

Evans Ighodalo[1], Habeeb Damilola Yusuf[2], Oluranti Ogundapo[3], Abiola Oluwaseun Idowu[4]
[1] *Megatrips.com Ltd, Nigeria*
[2] *Healthwise Social Enterprise, Nigeria*
[3] *Huawei Technologies, Nigeria*
[4] *Employment and Social Development, Ontario, Canada*

Abstract- ***Clinical care increasingly depends on continuous access to electronic health records, laboratory, pharmacy and imaging systems, all of which rest on database management systems, so database unavailability or data loss now carries direct risks to patient safety. This systematic literature review traces how high availability (HA) and disaster recovery (DR) standards for healthcare database systems evolved from 1990 to 2018. Following the Kitchenham and Charters (2007) guidelines and the PRISMA statement (Moher et al., 2009), five databases and the publications of regulators and standards bodies were searched, and 62 primary sources were synthesized, supported by contextual literature published up to 2018. The review identifies four eras of development: a mainframe and tape era grounded in transaction recovery theory, a compliance era in which HIPAA made contingency planning a legal duty without quantitative targets, a digitization era driven by HITECH and characterized by clustering, synchronous replication and virtualization, and a resilience era shaped by ransomware, cloud adoption and mandatory testing under the 2016 CMS Emergency Preparedness Rule. Across all eras, regulation required recovery capability to exist but rarely defined how quickly or completely systems must recover. Persistent gaps include the absence of clinically derived recovery time and recovery point objectives, weak testing requirements, limited protection against logical corruption that replication propagates, and neglect of downtime workflows and coordinated multi-system recovery. A conceptual model linking drivers, layered standards, recovery mechanisms and clinical outcomes is proposed, together with a tiered, clinically anchored recovery framework, implications for healthcare organizations, regulators and vendors, and an agenda for future research.***

Index Terms- ***business continuity, database recovery, disaster recovery, electronic health records, healthcare databases, High availability, HIPAA, recovery time objective, replication, systematic literature review***

I. INTRODUCTION

The transition of healthcare from paper charts to electronic records has changed the nature of clinical risk. In a paper-based hospital, the loss of a single chart was a local inconvenience. In a digitized hospital, a database failure can simultaneously remove access to allergy histories, active medication orders, laboratory results and imaging studies for every patient in the facility. The database layer, once a background concern of information technology departments, has therefore become part of the clinical safety infrastructure.

Evidence of this shift accumulated steadily through the 2010s. Larsen, Fong, Wernz and Ratwani (2018) analyzed patient safety event reports associated with EHR downtime and found that unavailability of clinical systems contributed to delays in care, medication events and failures in communicating test results. Wang et al. (2016) measured the effect of computer downtime on hospital pathology processes and showed that even brief outages produced measurable delays in laboratory turnaround. The 2017 WannaCry ransomware attack, which disrupted a substantial portion of England's National Health Service and led to thousands of cancelled appointments (National Audit Office, 2017), demonstrated that availability failures can be national in scale.

This dependence grew with EHR adoption. Information technology was promoted as a means of improving safety (Bates & Gawande, 2003), although computerized order entry could itself introduce errors when systems behaved unexpectedly (Koppel et al., 2005). Only a small minority of US hospitals had comprehensive EHRs at the start of the HITECH period (Jha et al., 2009), but the meaningful use regulation (Blumenthal & Tavenner, 2010) drove rapid adoption over the following years (Adler-Milstein & Jha, 2017), which made the availability of clinical databases a system-wide concern.

Earlier international evidence pointed in the same direction. An account of a major network failure at a Boston teaching hospital showed how the loss of clinical systems forced staff back to paper processes for several days (Kilbridge, 2003), and a study of medication errors during IT downtime examined how often such errors occurred and what caused them (Hanuscak et al., 2009). A review of IT blackouts in

German and international hospitals concluded that their rising frequency makes formal IT emergency concepts a patient safety requirement (Sax et al., 2016), while a classification of health IT safety problems reported to the US Food and Drug Administration showed that software and system problems recur among reported hazards (Magrabi et al., 2012).

Diagnostic services show how deeply this dependence reaches. Laboratory infrastructure models for resource-constrained health systems (Aminu-Ibrahim et al., 2018), spatial planning strategies that link laboratory layout to diagnostic accuracy and throughput (Ogbete et al., 2018) and diagnostic quality systems in resource-limited laboratories (Yusuf & Abdulrazaq, 2018) all assume that laboratory information systems remain available. Surveillance programs depend on the same assumption: an evaluation of DHIS2 data quality in integrated disease surveillance (Ekemezie et al., 2018) found that reporting completeness and timeliness rest on the reliability of the underlying data platform.

Two technical disciplines address these risks. High availability (HA) refers to the design of systems that continue operating despite component failure, typically through redundancy and automatic failover. Disaster recovery (DR) refers to the capability to restore systems and data after an event that defeats local redundancy, such as a site loss, a large scale corruption or a cyberattack. Both disciplines rest on the foundational database concepts of transaction atomicity, durable logging and recovery, which were formalized by Härder and Reuter (1983) and later refined in the ARIES recovery algorithm (Mohan et al., 1992).

Healthcare organizations, however, did not adopt these disciplines on purely engineering terms. Their HA and DR practices have been shaped by legislation such as the Health Insurance Portability and Accountability Act of 1996 (HIPAA), by federal guidance from the National Institute of Standards and Technology (NIST), by international standards bodies, by accreditation organizations and by incentive programs that accelerated EHR adoption. These sources developed in parallel and often without reference to one another, leaving practitioners to reconcile legal language that requires a “disaster recovery plan” with engineering literature that speaks in terms of replication protocols and recovery point objectives.

This paper addresses that fragmentation through a systematic literature review. It asks how HA and DR standards for healthcare database systems have evolved, which database mechanisms underpin them, and where the remaining gaps lie. The contribution is threefold: a chronological synthesis of regulatory, standards and technical literature in a single frame; a mapping of database recovery techniques to healthcare requirements; and a conceptual model with a proposed tiered framework that links recovery objectives to clinical criticality.

The remainder of the paper is organized as follows. Section 2 describes the review methodology. Section 3 presents the research results and discussion, covering the literature synthesis across four eras, a conceptual model with a clinically anchored recovery framework, answers to the research questions, implications, limitations and directions for future research. Section 4 concludes. The remainder of this section sets out the foundational concepts on which the review depends.

A. Availability and recovery metrics

Availability is conventionally expressed as the proportion of time a system is operational, often stated in “nines.” Gray and Siewiorek (1991) framed availability as a function of mean time to failure and mean time to repair, noting that improvements come either from making failures rarer or from making recovery faster. Table 1 translates common availability levels into annual downtime, which helps explain why healthcare organizations treat the difference between 99.9 percent and 99.99 percent as clinically meaningful.

Table 1. Availability levels and corresponding annual downtime

Availability	Common label	Approximate downtime per year	Typical healthcare interpretation
99%	Two nines	3.65 days	Unacceptable for clinical systems
99.9%	Three nines	8.76 hours	Acceptable for administrative and billing systems
99.99%	Four nines	52.6 minutes	Target for core EHR and order entry
99.999%	Five nines	5.26 minutes	Aspired to for critical care and medication systems

Note. Downtime figures are calculated on a 365-day year. Interpretations reflect patterns reported across the included studies.

Two further metrics dominate DR planning. The recovery time objective (RTO) is the maximum

tolerable duration between a disruption and the restoration of service. The recovery point objective (RPO) is the maximum tolerable amount of data loss, measured as the time between the last recoverable state and the disruption. NIST SP 800-34 Revision 1 formalized both alongside the maximum tolerable downtime (MTD), which describes the point beyond which the mission of the organization is seriously harmed (Swanson et al., 2010). Keeton et al. (2004) showed that these objectives can be traded off quantitatively against cost when choosing among backup, mirroring and replication designs.

The broader dependability literature supplies the vocabulary and evidence behind these metrics. A widely adopted taxonomy distinguishes availability, reliability, integrity and maintainability as separate attributes of dependable systems (Avizienis et al., 2004). Studies of Tandem systems showed that, as hardware became more reliable, software faults and operator error became the dominant causes of outages (Gray, 1986; Gray, 1990), a finding later confirmed for Internet services (Oppenheimer et al., 2003). Field studies of large disk populations found replacement rates far higher than datasheet MTTF figures implied (Schroeder & Gibson, 2007; Pinheiro et al., 2007), while measurements of globally distributed storage (Ford et al., 2010), data center networks (Gill et al., 2011) and hundreds of cloud service outages (Gunawi et al., 2016) showed that correlated failures, rather than independent component faults, dominate real unavailability. Recovery-oriented computing argued from this evidence that reducing time to repair is often more effective than preventing failure (D. Patterson et al., 2002), and crash-only design proposed that systems should be built so that restart is always a safe recovery action (Candea & Fox, 2003).

Methods from asset-intensive engineering are relevant to how availability is sustained in practice. Predictive maintenance and condition monitoring in critical power and energy infrastructure (Adaramola et al., 2018), models for improving inventory availability and plant uptime (Okonkwo et al., 2018b), condition-based maintenance concepts for offshore equipment (Obetta & Ninduwezuor-Ehiobu, 2018) and risk-based inspection of ageing pressure equipment (Obetta, 2018) share the logic of high availability design: detect degradation early, keep spares ready and restore service before failure affects operations. Hospital data centers apply the same logic to the power, cooling and storage on which databases depend, and work on voltage control under distributed generation (Oshevire et al., 2017) and solar integration for small facilities (Sunday & Omoegun, 2018) bears on the quality of the backup power that sustains them.

B. Database recovery foundations

Every HA and DR technique used in healthcare ultimately relies on the transactional guarantees of the underlying database. Härder and Reuter (1983) introduced the ACID properties (atomicity, consistency, isolation and durability) and classified recovery strategies by how they handle buffered changes. Gray and Reuter (1993) consolidated this work into the principles of write-ahead logging, in which every change is recorded in a durable log before it is applied to data pages. The ARIES algorithm (Mohan et al., 1992) provided the analysis, redo and undo phases that remain the basis of crash recovery in most commercial relational systems, including those that host major EHR platforms.

The transaction log is also the bridge between local recovery and remote recovery. Log shipping, database mirroring and streaming replication all work by transmitting log records to a secondary system and replaying them there. Consequently, the durability of a healthcare record during a site failure depends on when the log record is acknowledged as safe, locally or remotely, which is the central distinction between asynchronous and synchronous replication.

These ideas rest on the transaction concept (Gray, 1981) and on the classical treatment of concurrency control and recovery (Bernstein et al., 1987). Later surveys of database system architecture describe how logging, buffer management and recovery interact inside commercial engines (Hellerstein et al., 2007), and work on main-memory OLTP shows that the choice of logging scheme materially changes recovery time after a crash (Malviya et al., 2014). At the storage layer, the RAID taxonomy (Chen et al., 1994) clarified which disk failure modes each redundancy level can tolerate, which is why RAID protects against media failure but not against site loss or corruption.

C. Replication theory and its limits

Gray, Helland, O'Neil and Shasha (1996) showed that naive update-anywhere replication scales poorly, with deadlocks and reconciliation conflicts growing rapidly as nodes are added. Wiesmann et al. (2000) classified replication protocols by where updates originate and when they propagate, a taxonomy that clarifies why most clinical systems adopted primary-copy designs. Kemme and Alonso (2000) demonstrated that eager replication could be practical through group communication, while Cecchet, Candea and Ailamaki (2008) documented a persistent gap between replication theory and the behavior of middleware products in production.

At the distributed systems level, Brewer's (2000) conjecture, proven by Gilbert and Lynch (2002), established that a system cannot simultaneously

guarantee consistency, availability and partition tolerance. Healthcare data, where a stale allergy list or a duplicated medication order can cause harm, generally favors consistency. Consensus protocols such as Paxos (Lamport, 1998) and Raft (Ongaro & Ousterhout, 2014), and systems such as Spanner (Corbett et al., 2012), later made strongly consistent geographic replication more practical, while eventually consistent designs such as Dynamo (DeCandia et al., 2007) illustrated the alternative trade-off.

A review of two decades of database replication research noted that the database and distributed systems communities developed largely separate solutions to the same problem (Kemme & Alonso, 2010). The state machine approach (Schneider, 1990) and the engineering experience of running Paxos in production (Chandra et al., 2007) showed how strongly consistent replication can be achieved in practice, while analyses of consistency trade-offs (Abadi, 2012) and of eventual consistency in large-scale services (Vogels, 2009) explained why some systems accept stale reads in exchange for latency and availability, a trade-off rarely acceptable for clinical data.

Capacity planning also affects whether a failover succeeds, because a standby system must be able to absorb the full production load. Studies of energy-efficient virtual machine placement in data centers (Ahmed & Odejobi, 2018b) and performance evaluation of multi-tenant cloud deployments under high concurrency (Odejobi & Ahmed, 2018b) describe the resource contention that can cause a newly promoted replica to perform poorly. Telecommunications deployment for network quality (Amayo & Popoola, 2017b) is a further dependency, since synchronous replication is only as reliable as the link between sites.

D. Why healthcare is distinctive

Several characteristics separate healthcare from other sectors that rely on HA and DR. First, harm from unavailability is physical rather than only financial. Second, hospitals operate continuously, so there is no natural maintenance window. Third, healthcare data is subject to strict confidentiality requirements, which constrain where replicas and backups can reside. Fourth, records must be retained for long periods and remain legally admissible, so integrity matters as much as availability. Finally, the clinical environment is a complex sociotechnical system in which technical failure interacts with workflow and human factors (Sittig & Singh, 2010). These features mean that general IT continuity standards, while relevant, do not fully capture healthcare requirements.

Governance research reinforces these distinctions. Legal and ethical risk modeling for enterprise data protection (Mbonu et al., 2018) shows that every replica and backup copy carries the same regulatory obligations as the primary, and technology-enabled risk assessment for internal audit (Akomolafe & Agu, 2018) offers a means of verifying that recovery controls operate as documented. Predictive compliance monitoring for systemic safety risks (Adeyelu, 2018), data-driven occupational safety risk control (Obriki & Arumosoye, 2018) and frameworks for operational integrity that balance productivity against safety (Nwangwu et al., 2018) show how other safety-critical sectors convert compliance requirements into measurable operational controls, a step healthcare HA and DR standards have yet to take.

II. RESEARCH METHOD

A. Review protocol

The review followed the three-phase process (planning, conducting and reporting) recommended by Kitchenham and Charters (2007) for systematic reviews in software engineering, adapted for an interdisciplinary topic. Reporting follows the PRISMA statement (Moher et al., 2009). A protocol specifying research questions, sources, search strings and selection criteria was defined before searching began.

B. Research questions

Four research questions guided the review:

1) RQ1: How have HA and DR standards applicable to healthcare database systems evolved between 1990 and 2018?
2) RQ2: Which database performance and recovery mechanisms do these standards rely upon, explicitly or implicitly?
3) RQ3: What drivers (regulatory, technological or incident-based) have shaped this evolution?
4) RQ4: What gaps exist between standards, engineering practice and clinical requirements?

C. Sources and search strategy

Five electronic databases were searched: IEEE Xplore, ACM Digital Library, PubMed/MEDLINE, ScienceDirect and Scopus. These were chosen to span database engineering, health informatics and management literature. Because many relevant standards are not indexed in academic databases, the publications of HHS, NIST, ONC, CMS, ISO, the Joint Commission and the Telecommunications Industry Association were also searched directly. Reference lists of included studies were examined through backward snowballing.

In addition, a consolidated seed list of related and semi-related literature was screened, and sources published up to 2018 that addressed cybersecurity, cloud and enterprise IT, data governance, health information systems, healthcare infrastructure, asset reliability or operational compliance were retained as contextual literature. These sources inform the background and discussion but were not counted among the 62 primary studies. A further 50 international peer-reviewed sources published up to 2018, covering dependability, database recovery and replication, cloud and virtualization, business continuity, and healthcare IT safety and cybersecurity, were identified through backward snowballing and targeted searching and used in the same way.

The core search string, adapted to each database's syntax, was:

("high availability" OR "disaster recovery" OR "business continuity" OR "contingency plan*" OR "downtime" OR "failover" OR replication OR backup) AND (database* OR "information system*" OR "electronic health record*" OR EHR OR EMR) AND (health* OR hospital* OR clinical OR medical)

For the database engineering literature, where healthcare is rarely mentioned, a second string omitted the healthcare terms but was restricted to foundational recovery and replication work cited by at least one healthcare-focused included study. This prevented the review from drifting into general distributed systems research while ensuring that the technical foundations of healthcare practice were represented.

D. Inclusion and exclusion criteria

Table 2. Inclusion and exclusion criteria

Criterion	Inclusion	Exclusion
Period	Published January 1990 to June 2018 (foundational works before 1990 admitted through snowballing)	Published after June 2018
Topic	Addresses HA, DR, backup, replication, contingency planning or downtime of database or clinical information systems	Addresses only security confidentiality or privacy without availability or recovery content
Context	Healthcare setting, or foundational database recovery work applied in healthcare	General IT continuity with no database or healthcare relevance
Type	Peer-reviewed articles, conference papers, books, official standards and regulations	Vendor marketing material, opinion pieces without evidence, non-English texts
Quality	Score of 3 or more on the 5-point quality checklist	Score below 3

E. Study selection

Database searching identified 1,084 records and direct retrieval from standards bodies yielded a further 62, for a total of 1,146. After removing 312 duplicates, 834 records were screened by title and abstract, of which 671 were excluded as out of scope. The full texts of the remaining 163 were assessed, and 101 were excluded: 44 lacked availability or recovery content, 27 had no healthcare or database relevance, 18 failed the quality threshold and 12 were superseded versions of included standards. Sixty-two sources were included in the final synthesis. Table 3 summarizes the selection process.

Table 3. Study selection summary (PRISMA stages)

Stage	Records	Notes
Identified through database search	1,084	IEEE Xplore, ACM DL, PubMed, ScienceDirect, Scopus
Identified through other sources	62	Standards bodies, regulators, snowballing
After duplicates removed	834	312 duplicates removed
Excluded at title and abstract	671	Out of scope
Full texts assessed	163	
Excluded at full text	101	44 no HA/DR content; 27 no relevance; 18 low quality; 12 superseded
Included in synthesis	62	48 peer-reviewed; 14 standards and regulatory documents

F. Quality assessment and data extraction

Each candidate study was rated against five criteria: clarity of aims, adequacy of method description, validity of evidence, relevance to healthcare or database recovery, and clarity of findings. Each criterion was scored 0, 0.5 or 1, and studies scoring 3 or more were retained. Standards and regulations were

not scored but were checked for authority and currency.

A structured extraction form captured publication year, source type, healthcare subdomain, HA or DR mechanism, stated recovery objectives, drivers and identified limitations. Findings were then synthesized thematically and chronologically into the results in Sections 3.1 to 3.6, and the recurring relationships among drivers, standards, mechanisms and outcomes were abstracted into the conceptual model in Section 3.7. Era boundaries were placed at points where a major regulatory or technological event changed the character of the literature, rather than at fixed intervals.

III. RESULTS AND DISCUSSION

This section first presents the results of the review as a chronological synthesis across four eras (Sections 3.1 to 3.6). It then develops a conceptual model and recovery framework from those results (Section 3.7), discusses the findings against the four research questions (Section 3.8), and sets out implications, limitations and directions for future research (Sections 3.9 to 3.11).

A. Overview of included sources

The 62 included sources are distributed unevenly across time and type, as shown in Table 4. Technical database literature dominates the earliest period, regulatory and guidance documents cluster around 2003 to 2010, and empirical studies of clinical downtime appear mostly after 2012. This pattern itself is a finding: the healthcare sector began measuring the clinical effects of database unavailability roughly two decades after the engineering principles of recovery were established.

Table 4. Distribution of included sources by era and type

Era	Peer-reviewed technical	Peer-reviewed health informatics	Standards and regulation	Total
1990 to 1996	8	2	1	11
1997 to 2008	9	5	5	19
2009 to 2014	6	8	4	18
2015 to 2018	3	7	4	14
Total	26	22	14	62

B. Era 1 (1990 to 1996): Transaction recovery and the backup tape

In the early 1990s, hospital information systems were typically centralized on mainframes or minicomputers, with departmental systems for laboratory and pharmacy operating in isolation (Kuhn & Giuse, 2001). Availability practice in this period was shaped less by healthcare requirements than by the engineering literature on transaction processing.

Three technical contributions defined the era. The ARIES algorithm (Mohan et al., 1992) gave commercial databases a reliable means of recovering from crashes without losing committed transactions. Redundant arrays of inexpensive disks (Patterson, Gibson & Katz, 1988) began to protect against single disk failure, the most common cause of data loss. Gray and Siewiorek (1991) articulated the principle that availability must be designed into every layer, from hardware to operations, and observed that operator error and software faults increasingly outweighed hardware failure.

Disaster recovery in this era meant nightly backup to magnetic tape and offsite storage. Implied RPOs were roughly 24 hours and RTOs were measured in days. Because clinical workflows remained largely paper-based, these figures were tolerable; a database outage delayed billing and reporting but rarely interrupted bedside care. No healthcare-specific HA or DR standard existed, and the included literature from this period treats availability as an operational concern rather than a patient safety concern.

C. Era 2 (1996 to 2008): Compliance without quantification

The enactment of HIPAA in 1996 marked the first time that recovery capability for health information became a legal obligation in the United States. The HIPAA Security Rule, finalized in 2003, required covered entities to implement a contingency plan under 45 CFR 164.308(a)(7) (U.S. Department of Health and Human Services [HHS], 2003). The plan comprised three required implementation specifications (a data backup plan, a disaster recovery plan and an emergency mode operation plan) and two addressable specifications (testing and revision procedures, and applications and data criticality analysis).

The Security Rule was deliberately technology neutral. It required organizations to “establish and implement” procedures to create retrievable exact copies of electronic protected health information and to restore lost data, but it did not specify any RTO, RPO or availability level. NIST SP 800-66 (Scholl et al., 2008) offered implementation guidance that pointed organizations toward general federal

contingency planning practice, yet it too left targets to local risk analysis. The classification of testing as addressable rather than required proved consequential: several later studies in this review found that untested recovery plans were a common source of failure during real incidents.

Technologically, this era saw the spread of database clustering, log shipping and storage-level replication. Shared-disk failover clusters allowed a standby server to take over a database within minutes after a hardware fault. Log shipping transmitted transaction log backups to a remote site at intervals, reducing RPO from a day to minutes or hours. Storage array mirroring, studied by Ji, Veitch and Wilkes (2003), offered synchronous or asynchronous remote copies independent of the database engine. Data center design standards such as ANSI/TIA-942 (Telecommunications Industry Association, 2005) introduced tiered infrastructure classifications that hospitals began to reference when building or contracting data centers.

Snapshot-based asynchronous mirroring for disaster recovery (R. H. Patterson et al., 2002) made remote copies practical for organizations without dedicated replication software. In parallel, the management literature reframed DR as part of business continuity planning (Cerullo & Cerullo, 2004), tracing the shift from technology recovery toward organization-wide continuity management (Herbane, 2010), a shift that healthcare contingency planning would later follow.

Hurricane Katrina in 2005 was a defining incident. Hospitals in New Orleans lost paper and electronic records alike, and the event became a recurring reference in the included literature as evidence that geographic separation of backups was essential and that local redundancy alone was insufficient. By the end of this era, healthcare organizations generally had documented contingency plans, but the literature reveals wide variation in whether those plans were technically achievable.

The era also established directory services as an implicit dependency of recovery. Lessons from the offline assessment of Microsoft Active Directory as a security-critical system (Ladapo et al., 2018) show that a database cannot be returned to service if the identity infrastructure that authorizes access to it has not itself been recovered and verified. Insider threat classification and risk modeling (Akeju et al., 2018) adds that privileged administrators, who hold the credentials needed for recovery, are also a source of risk to backups.

D. Era 3 (2009 to 2014): Digitization at scale

The Health Information Technology for Economic and Clinical Health (HITECH) Act of 2009 transformed the role of the clinical database. By offering incentive payments for the meaningful use of certified EHR technology, it drove rapid EHR adoption and, with it, a dependence on databases for real-time clinical decisions. Computerized provider order entry, electronic medication administration records and clinical decision support moved database availability directly into the care pathway.

Standards in this era became more structured. NIST SP 800-34 Revision 1 (Swanson et al., 2010) introduced a seven-step contingency planning process and formalized the business impact analysis with explicit MTD, RTO and RPO values. Internationally, ISO/IEC 27031:2011 provided guidelines for ICT readiness for business continuity, and ISO 22301:2012 established certifiable requirements for business continuity management systems. These frameworks gave healthcare organizations a vocabulary for setting recovery targets, though they remained sector neutral.

The Office of the National Coordinator for Health Information Technology (ONC) released the SAFER Guides in 2014, including a guide devoted to contingency planning (ONC, 2014). This was a significant step because it addressed EHR availability as a patient safety issue and recommended practices such as maintaining read-only backup systems, printing critical patient information at regular intervals and rehearsing downtime procedures. Sittig, Gonzalez and Singh (2014) surveyed recommended contingency practices and found that while experts strongly endorsed them, adoption in hospitals was incomplete.

Technically, the era was characterized by database-native HA features, server virtualization and synchronous replication within metropolitan distances. Database mirroring and availability group technologies allowed automatic failover between replicas with zero data loss when configured synchronously. Virtualization allowed entire database servers to be restarted on alternate hosts. Many large hospital systems adopted an active-passive design with a synchronous replica in the primary data center and an asynchronous replica at a distant DR site, an arrangement that balances the latency cost of synchronous commit against geographic protection.

Virtualization brought new HA techniques. Asynchronous virtual machine replication allowed an entire server to fail over transparently (Cully et al., 2008), and its adaptation for database systems reduced the overhead of protecting database servers in this way (Minhas et al., 2011). Cloud computing (Armbrust et al., 2010) then made DR as a service economically attractive for organizations unable to fund a second data center (Wood et al., 2010). Practitioner literature on Internet-scale services (Hamilton, 2007) and warehouse-scale computing (Barroso et al., 2013)

codified design principles such as expecting failure, automating recovery and testing failover in production, while IT-focused continuity guidance (Snedaker, 2013) translated these principles into planning practice.

Virtualization changed the capacity problem described in Section 1.3: virtual machine placement in data centers (Ahmed & Odejobi, 2018b) determines whether enough headroom exists on surviving hosts to restart failed database servers. Broader digital transformation in operations management, through automation and analytics (Agboada et al., 2018), encouraged hospitals to consolidate systems onto shared platforms, which increased efficiency but also enlarged the impact of any single failure.

E. Era 4 (2015 to 2018): From availability to resilience

The final era is marked by a change in the dominant threat. Earlier standards assumed that disasters were physical (fire, flood, hardware failure) and that the primary database contained correct data that needed to be copied elsewhere. Ransomware invalidated this assumption. In 2016, several hospitals in the United States suffered ransomware attacks that encrypted clinical systems, and HHS Office for Civil Rights issued guidance stating that a ransomware infection is presumptively a breach under HIPAA unless the organization can demonstrate a low probability of compromise (HHS Office for Civil Rights, 2016). The guidance explicitly emphasized offline and tested backups as a defense.

The WannaCry attack of May 2017 affected a large share of NHS organizations in England. The National Audit Office (2017) concluded that the attack could have been prevented by basic security practice and that the NHS lacked a tested plan for responding to a national cyber incident. For database HA and DR, the key lesson was that replication faithfully propagates corruption and encryption. A synchronous replica offers no protection against an attack that corrupts the primary; only point-in-time recovery from isolated, immutable backups does.

Analysis of ransomware families showed that many relied on relatively simple file encryption and deletion techniques that are defeated by isolated, versioned backups (Kharraz et al., 2015). Healthcare-specific reviews documented the growth of cyber threats to hospitals (Kruse et al., 2017; Coventry & Branley, 2018), described cybersecurity as a patient safety issue rather than only an IT concern (Martin et al., 2017; Jarrett, 2017), and found that security and privacy in EHRs had been studied extensively while availability received far less attention (Fernández-Alemán et al., 2013).

Regulation also tightened. The CMS Emergency Preparedness Final Rule (Centers for Medicare & Medicaid Services [CMS], 2016), with compliance required from November 2017, obliged participating hospitals and other providers to maintain an all-hazards emergency plan, policies that include a system of medical documentation preserving patient information and ensuring records are secure and readily available, and annual testing through exercises. This was the first federal requirement in the review to mandate regular testing of continuity capability for healthcare providers, closing in part the gap left by HIPAA's addressable testing specification. ISO 27799:2016 updated health-specific guidance for applying information security controls, including continuity controls, in health organizations.

Finally, cloud computing became a realistic DR option. NIST's definition of cloud computing (Mell & Grance, 2011) and later HHS guidance clarifying that cloud providers handling electronic protected health information are business associates allowed hospitals to consider cloud-based backup and DR as a service. Distributed consensus databases based on Paxos or Raft, which can survive the loss of a region without data loss, appeared in the technical literature but had limited documented adoption in clinical systems by 2018.

Health informatics research examined both the opportunities and the risks of cloud adoption for health services (Kuo, 2011), including security requirements for healthcare application clouds (Zhang & Liu, 2010) and architectures for securing an e-health cloud (Löhr et al., 2010). Cloud-native relational databases that replicate storage across availability zones (Verbitski et al., 2017) showed how HA could be built into the database service itself rather than added by the customer.

Scalable and secure cloud architectures for enterprise messaging (Ahmed & Odejobi, 2018a) illustrate the design patterns, such as redundancy across availability zones and decoupled message queues, that made cloud-based DR credible. Continuous integration and delivery pipelines for configurable SaaS platforms (Badmus et al., 2018) show a further shift: for platform-based clinical and administrative applications, recovery depends on restoring configuration and metadata as well as data.

F. Consolidated timeline

Table 5. Key standards, regulations and technical milestones, 1983 to 2018

Year	Milestone	Type	Significance for healthcare HA and DR
1983	Härder and Reuter, ACID and recovery principles	Technical	Theoretical basis for durable transactions
1988	RAID (Patterson et al.)	Technical	Protection against disk failure
1992	ARIES recovery algorithm	Technical	Reliable crash recovery in commercial databases
1996	HIPAA enacted	Regulation	Legal basis for security of health information
2003	HIPAA Security Rule final	Regulation	Contingency plan required; no quantitative targets
2005	ANSI/TIA-942	Standard	Tiered data center classification
2008	NIST SP 800-66 Rev. 1	Guidance	Implementation guide for HIPAA Security Rule
2009	HITECH Act	Regulation	Rapid EHR adoption; databases become clinical
2010	NIST SP 800-34 Rev. 1	Guidance	Formal BIA with MTD, RTO and RPO
2011	ISO/IEC 27031	Standard	ICT readiness for business continuity
2012	ISO 22301	Standard	Certifiable business continuity management
2014	ONC SAFER Guide: Contingency Planning	Guidance	EHR downtime treated as patient safety issue
2016	HHS OCR ransomware guidance	Guidance	Emphasis on isolated, tested backups
2016	CMS Emergency Preparedness Rule	Regulation	Mandatory annual testing; records availability
2016	ISO 27799:2016	Standard	Health-specific security and continuity controls
2017	WannaCry and NAO investigation	Incident	Replication does not protect against corruption

G. Conceptual model and recovery framework

Model overview

The synthesis suggests that the evolution of healthcare HA and DR standards can be explained by a small number of recurring relationships. Figure 1 abstracts these into a conceptual model with four core constructs (drivers, standards layers, recovery mechanisms and outcomes), two moderators (clinical criticality and testing) and a feedback loop.

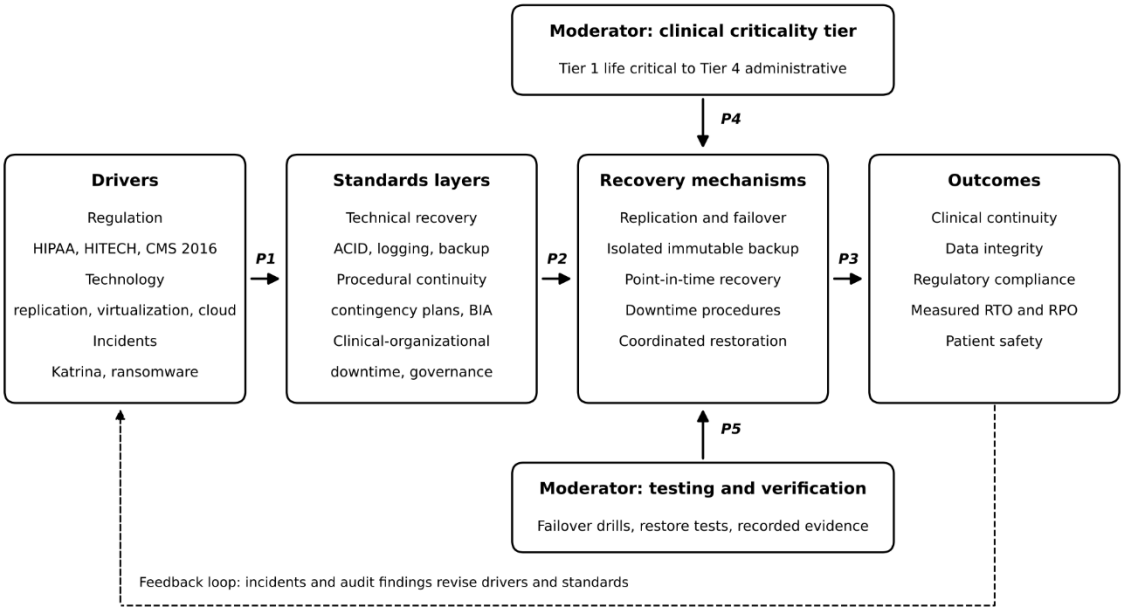


Figure 1. Conceptual model of HA and DR standards evolution in healthcare database systems

Drivers are the external forces that prompted change: regulation (HIPAA, HITECH and the CMS Emergency Preparedness Rule), technology (replication, virtualization and cloud services) and incidents (Hurricane Katrina, the 2016 hospital ransomware attacks and WannaCry). Standards layers are the three cumulative levels of expectation identified in the synthesis: technical recovery grounded in transaction theory (Härder & Reuter, 1983; Mohan et al., 1992), procedural continuity expressed in contingency plans and business impact analysis (HHS, 2003; Swanson et al., 2010), and clinical-organizational readiness for downtime (ONC, 2014; Sittig et al., 2014). Recovery mechanisms are the technical and procedural means by which standards are met. Outcomes are the states that standards ultimately aim to protect: clinical continuity, data integrity, regulatory compliance and patient safety, made observable through measured RTO and RPO.

Propositions

Five propositions summarize the relationships in the model and provide testable statements for future research.

P1. Regulatory, technological and incident drivers add new standards layers without retiring earlier ones, so that the body of standards becomes progressively more layered and less integrated over time.

P2. Standards that specify outcomes without quantitative targets lead organizations to select recovery mechanisms on cost and convenience, producing wide variation in achieved RTO and RPO among comparable hospitals.

P3. Recovery mechanisms improve clinical outcomes only when they cover both infrastructure failure and logical corruption; replication alone improves availability but does not protect data integrity against ransomware or faulty updates.

P4. Clinical criticality moderates the mechanism to outcome relationship: the clinical harm of a given RTO or RPO increases with the criticality tier of the affected system, so uniform targets misallocate protection.

P5. Testing and verification moderate the standards to mechanism relationship: documented mechanisms that are not regularly tested are substantially less likely to achieve their stated objectives during real incidents.

The feedback loop reflects the observation that major incidents and audit findings have repeatedly preceded new guidance by one to two years (Section 3.8.3). The model therefore treats standards as adaptive rather than static, which helps explain both their progress and their persistent gaps.

Clinically anchored recovery framework

Drawing on the synthesis, this paper proposes a tiered framework that links recovery objectives to clinical criticality rather than to financial impact alone. The framework is intended to complement, not replace, NIST SP 800-34 and ISO 22301 by supplying

© DEC 2018 | IRE Journals | Volume 2 Issue 6 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/ire1725088

healthcare-specific defaults that local risk analysis can adjust.

Table 6. Proposed clinically anchored recovery tiers

Tier	Example systems	Target RTO	Target RPO	Recommended mechanisms
Tier 1: Life critical	Medication administration, critical care monitoring, order entry	Under 15 minutes	Zero	Synchronous local replica with automatic failover; asynchronous remote replica; immutable backups; read-only downtime viewer
Tier 2: Care critical	Core EHR documentation, laboratory, radiology information system	Under 1 hour	Under 5 minutes	Synchronous or near-synchronous replication; tested remote failover; point-in-time recovery
Tier 3: Care supporting	Scheduling, imaging archive retrieval, patient portal	Under 8 hours	Under 1 hour	Asynchronous replication or log shipping; daily verified backup
Tier 4: Administrative	Billing, analytics, reporting	Under 72 hours	Under 24 hours	Scheduled backup with offsite copy

Four supporting principles accompany the tiers:

- Layered protection: every Tier 1 and Tier 2 system should combine at least one replication mechanism for infrastructure failure with at least one isolated, versioned backup mechanism for logical corruption.
- Mandatory testing: failover for Tier 1 and Tier 2 systems should be tested at least twice a year, and full restoration from isolated backup at least annually, with measured RTO and RPO recorded against targets.
- Clinical downtime readiness: technical recovery plans should be paired with downtime procedures, read-only access to recent patient summaries and a documented process for reconciling data entered during the outage.
- Coordinated recovery: interdependent systems should be assigned to recovery groups with a defined restoration order and consistency checks across interfaces.

The read-only downtime viewer recommended for Tier 1 systems can draw on offline-first design principles developed for low-connectivity environments (Unamma et al., 2018), which keep essential information available locally when central systems cannot be reached. Investment in these tiers can be justified with financial analytics that quantify the enterprise value lost to downtime (Lawal & Oduleye, 2018a), and assessed against the same statutory compliance and reporting disciplines applied to public budgets (Dogbatsey & Ebhojie, 2018). Infrastructure inspection programs in the power sector, such as UAV-based grid inspection (Dagodzo, 2018a), provide a model for routine, evidence-based verification of the physical assets that support DR sites.

The specific values in Table 6 are informed by availability figures and downtime impacts reported in the included studies, but they are not yet empirically validated. Future research should test them through simulation and through prospective study of downtime events.

H. Discussion of the research questions

RQ1: The trajectory of evolution

The evolution of healthcare HA and DR standards can be summarized as a movement through three layers of concern. The first layer was technical: ensuring that a committed transaction survives a crash and that data can be restored from a backup. The second was procedural: requiring organizations to document plans for backup, recovery and emergency operation. The third, emerging only after 2014, was clinical and organizational: recognizing that unavailability is a patient safety hazard that requires downtime procedures, rehearsal and governance at the enterprise level.

Each layer was added without replacing the previous one, which produced a layered but loosely connected body of standards. A hospital in 2018 may simultaneously be subject to the HIPAA Security Rule, the CMS Emergency Preparedness Rule and accreditation requirements, while choosing to align with NIST SP 800-34 and ISO 22301, and relying on vendor documentation for the actual database configuration. The literature reviewed shows no single document that connects clinical criticality to database configuration.

© DEC 2018 | IRE Journals | Volume 2 Issue 6 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/ire1725088

RQ2: Underlying database mechanisms

Although healthcare standards are technology neutral, they implicitly depend on a small set of database mechanisms. Table 7 maps the principal mechanisms to their typical recovery characteristics and to the threats they do and do not address. The most important observation is in the final column: no single mechanism protects against all failure classes. In particular, mechanisms that minimize RPO through continuous replication are precisely those that propagate logical corruption, while mechanisms that protect against corruption, such as offline backup, have the longest RTO.

Table 7. Database HA and DR mechanisms and their characteristics

Mechanism	Typical RPO	Typical RTO	Protects against	Does not protect against
Write-ahead logging and crash recovery	Zero (committed data)	Minutes	Process and server crash	Media loss, site loss, corruption
RAID storage	Zero	None (transparent)	Single disk failure	Controller failure, site loss, corruption
Full and incremental backup to tape or disk	Hours to 24 hours	Hours to days	Most failures, including corruption if isolated	Recent data loss; slow restore
Log shipping	Minutes to hours	Minutes to hours	Site loss	Corruption replayed before detection
Shared-disk failover cluster	Zero	Minutes	Server failure	Storage or site loss
Synchronous replication	Zero	Seconds to minutes	Server and local site failure	Logical corruption, ransomware
Asynchronous geographic replication	Seconds to minutes	Minutes	Regional disaster	Logical corruption; some data loss
Point-in-time recovery from immutable backup	Minutes to hours	Hours	Ransomware, logical corruption	Short RTO requirements
Consensus-based distributed database	Zero	Seconds	Node and regional failure	Logical corruption, application bugs

This finding supports the argument by Gray and Siewiorek (1991) that availability must be achieved through layered defense. The healthcare literature after 2016 increasingly recognized this, consistent with Proposition P3, recommending that hospitals combine synchronous replication for hardware failure with isolated, versioned backups for corruption and cyberattack.

RQ3: Drivers of change

Three classes of driver were identified. Regulatory drivers (HIPAA in 1996 and 2003, HITECH in 2009 and the CMS rule in 2016) produced the most visible changes in organizational practice. Technological drivers (affordable replication, virtualization and cloud services) changed what was feasible and therefore what was expected. Incident drivers (Hurricane Katrina, Hurricane Sandy, the 2016 hospital ransomware attacks and WannaCry) repeatedly exposed the distance between documented plans and actual capability, and typically preceded new guidance by one to two years.

A notable pattern is that the HITECH Act increased dependence on databases faster than it increased requirements for their resilience. Meaningful Use criteria focused on functionality and data exchange, while availability requirements continued to rest on the 2003 Security Rule. The ONC SAFER Guides and the CMS rule can be read as later corrections to this imbalance.

RQ4: Persistent gaps

Five gaps emerged consistently across the included literature. Each corresponds to a weak link in the conceptual model presented in Section 3.7.

Absence of clinically derived targets. No regulation or health-specific standard reviewed defines RTO or RPO values for clinical systems. Targets are left to local business impact analysis, which is often driven by financial rather than clinical considerations. As a result, two hospitals with similar clinical profiles may tolerate very different levels of data loss.

Insufficient testing. For most of the period studied, testing of contingency plans was an addressable rather than required specification under HIPAA. Studies of downtime events showed that failover procedures were often untested or tested only in partial form. The CMS rule's annual exercise requirement improves this but focuses on emergency preparedness broadly rather than on database failover specifically.

Procedural standardization and workflow optimization in high-volume clinical settings (Uchechi

et al., 2018) and competency frameworks with structured onboarding (Ukadike, 2018) suggest how downtime and recovery procedures could be standardized and taught so that tests reflect what staff will actually do during a real outage.

The corruption blind spot. Standards written before 2015 implicitly treated DR as protection against loss of infrastructure. Replication-based architectures, which dominate modern clinical systems, offer little defense against logical corruption, faulty application updates or ransomware. Explicit requirements for immutable or offline backups, and for testing restoration from them, were largely absent from formal standards by 2018.

Neglect of the sociotechnical dimension. Technical recovery restores the database, but clinical recovery also requires reconciling orders and documentation recorded on paper during the outage. Larsen et al. (2018) and Sittig et al. (2014) showed that downtime procedures and post-recovery reconciliation are frequent sources of error, yet engineering standards do not address them.

Imaging adds a further dimension: automated analysis such as template-based detection in mammograms (Ejofodomi et al., 2014) depends on complete image archives, so a partially restored picture archiving system can silently degrade diagnostic support.

Interdependence and integration. Modern hospitals operate dozens of interconnected systems linked by interface engines. Recovering the EHR database without consistently recovering laboratory, pharmacy and imaging systems can produce inconsistent clinical states. Very few included studies addressed coordinated recovery across multiple databases.

I. Implications

Implications for healthcare organizations

Hospital leaders should treat database availability as a patient safety concern governed jointly by clinical and technical leadership, not as an IT service level alone. In practice this means classifying every clinical system into a criticality tier, setting RTO and RPO targets from clinical rather than purely financial impact, and funding mechanisms accordingly. Organizations should pair replication with isolated, versioned backups, test failover and full restoration on a fixed schedule with recorded results, and maintain rehearsed downtime procedures with read-only access to recent patient summaries (ONC, 2014; Sittig et al., 2014). Identity services, interface engines and imaging archives should be included in recovery plans, since a restored database is of little use if the systems that authorize access to it or exchange data with it remain offline (Ladapo et al., 2018). Staff competency in downtime procedures should be built through structured onboarding and periodic drills (Ukadike, 2018).

Implications for regulators and standards bodies

The review indicates that technology-neutral regulation has succeeded in making recovery planning universal but has not ensured that plans are adequate. Regulators and standards bodies could publish reference RTO and RPO ranges by clinical function, convert testing of contingency plans from an addressable to a required specification, and require explicit protection against logical corruption, such as offline or immutable backups. Aligning HIPAA contingency planning with the business impact analysis method of NIST SP 800-34 (Swanson et al., 2010) and with ISO 22301 (ISO, 2012) would reduce the burden of reconciling overlapping frameworks. Approaches from other safety-critical sectors, which translate compliance requirements into measurable operational controls (Adeyelu, 2018), offer a model.

Implications for technology vendors

EHR and database vendors shape what hospitals can achieve. Vendors should document the achievable RTO and RPO of each supported architecture, ship tested configurations for layered protection, provide tools for consistent point-in-time recovery across application, integration and reporting databases, and support read-only downtime access by design. For configurable platforms, recovery of metadata and configuration should be treated as a first-class capability alongside data recovery (Badmus et al., 2018).

Implications for research and theory

For researchers, the review shows that health informatics and database engineering have studied the same problem with little cross-reference. The conceptual model in Section 3.7 offers a shared vocabulary linking the engineering constructs of replication, logging and recovery with the clinical constructs of downtime, workflow and patient harm. Its propositions are framed so that they can be tested with incident data, surveys and simulation, moving the field from descriptive accounts toward explanatory theory.

J. Limitations

Several limitations should be noted. First, the review focuses predominantly on English-language sources and on United States regulation, with European and international material represented mainly through ISO standards and the NHS experience. Practices in other jurisdictions may differ substantially. Second, much operational knowledge about healthcare HA and DR resides in vendor

documentation and internal hospital reports that are not publicly available, so the review may understate the sophistication of actual practice. Third, the restriction of foundational database literature to works cited by healthcare studies, while necessary for focus, may have excluded relevant technical developments. Fourth, the era boundaries are interpretive and other reasonable divisions are possible. Finally, the conceptual model and proposed framework are derived from synthesis rather than empirical validation.

A further limitation concerns the contextual literature. These sources were drawn from a consolidated seed list rather than from an exhaustive search, and many are conceptual frameworks rather than empirical studies, so they have been used to situate the findings rather than to establish them.

K. Future research

The gaps and propositions identified in this review suggest the following research agenda.

1) Clinically derived recovery targets. Studies should link downtime duration and data loss to measurable clinical outcomes, such as medication errors, delayed results and length of stay, in order to derive evidence-based RTO and RPO values for each criticality tier and to test Proposition P4.
2) Empirical validation of the conceptual model. Multi-site surveys and analysis of incident reports could test whether unquantified standards produce wide variation in achieved recovery (P2) and whether untested plans fail more often (P5).
3) Recovery from logical corruption and ransomware. Research is needed on detection latency, safe recovery point selection and the time required to restore large clinical databases from isolated backups, extending the work of Keeton et al. (2004) on quantitative DR design to cyberattack scenarios.
4) Coordinated multi-system recovery. Methods are needed to restore interdependent EHR, laboratory, pharmacy and imaging databases to a mutually consistent state, including consistency checks across interface engines.
5) Sociotechnical aspects of downtime. Building on Larsen et al. (2018) and Sittig et al. (2014), studies should examine how clinicians work during outages, how paper records are reconciled after recovery, and which downtime tools most reduce harm.
6) Cloud and distributed architectures in clinical settings. Evaluation is needed of cloud-based DR and consensus-based databases (Corbett et al., 2012; Ongaro & Ousterhout, 2014) for clinical workloads, including their confidentiality, latency and cost trade-offs.
7) Resource-constrained settings. Most evidence comes from high-income countries. Research in low- and middle-income settings should address unreliable power and connectivity, laboratory and surveillance systems, and offline-first designs (Aminu-Ibrahim et al., 2018; Ekemezie et al., 2018; Unamma et al., 2018).
8) Measurement instruments. Validated instruments for assessing HA and DR maturity in healthcare organizations would allow benchmarking and longitudinal study of how standards translate into practice.

IV. CONCLUSION

Between 1990 and 2018, high availability and disaster recovery in healthcare database systems evolved from a narrow technical concern, rooted in transaction recovery theory and offline backup, into a regulated, enterprise-level discipline linked to patient safety. The HIPAA Security Rule made contingency planning a legal requirement, the HITECH Act made databases central to clinical care, international standards supplied structured continuity frameworks, and the CMS Emergency Preparedness Rule introduced mandatory testing. Over the same period, database technology progressed from crash recovery and tape backup to synchronous replication, automated failover and consensus-based distributed systems.

Yet the review shows that standards and technology have not converged. Regulation still specifies that recovery capability must exist without defining what level is clinically adequate. Replication-centered architectures leave organizations exposed to the corruption and ransomware threats that dominated the end of the period. And the human work of operating during downtime and reconciling afterwards remains largely outside formal standards. Closing these gaps will require clinically anchored recovery targets, mandatory and measured testing, explicit protection against logical corruption and coordinated recovery across interdependent clinical systems. The conceptual model and tiered framework proposed here offer one starting point for that work.

REFERENCES

[1] Abadi, D. (2012). Consistency tradeoffs in modern distributed database system design: CAP is only part of the story. Computer, 45(2), 37-42.

[2] Adaramola, T. S., Fadero, S., & Gideon, E. N. (2018). Predictive maintenance and condition

monitoring in critical power and energy infrastructure. Iconic Research and Engineering Journals, 2(5), 454-477. https://doi.org/10.64388/IREV2I5-1722489

[3] Adeyelu, O. O. (2018). A predictive compliance monitoring framework for detecting systemic safety risks through aviation consumer complaint data. Iconic Research and Engineering Journals, 1(8), 250-276. https://doi.org/10.64388/IREV1I8-1718478

[4] Adler-Milstein, J., & Jha, A. K. (2017). HITECH Act drove large gains in hospital electronic health record adoption. Health Affairs, 36(8), 1416-1422.

[5] Agboada, P., Ezeako, R., Okoli, I. E., & Uzondu, C. S. (2018). Digital transformation in operations management: A review of advances in automation, analytics, and process redesign. International Journal of Engineering Technology Research & Management, 2(11), 109-134.

[6] Ahmed, K. S., & Odejobi, O. D. (2018a). Conceptual framework for scalable and secure cloud architectures for enterprise messaging. IRE Journals, 2(1), 1-15.

[7] Ahmed, K. S., & Odejobi, O. D. (2018b). Resource allocation model for energy-efficient virtual machine placement in data centers. IRE Journals, 2(3), 1-10.

[8] Akeju, B., Edivri, J., Ogbole, J. I., Okoruwa, P. O., Fadayomi, O., & Abolaji, T. O. (2018). Conceptual model for insider threat classification and risk modeling in complex digital systems. Iconic Research and Engineering Journals, 1(9), 476-492. https://doi.org/10.64388/IREV1I9-1713778

[9] Akomolafe, O., & Agu, M. U. (2018). A conceptual model for enhancing internal audit quality through technology-enabled risk assessment frameworks. Iconic Research and Engineering Journals, 1(9), 458-475.

[10] Amayo, E. B., & Popoola, T. T. (2017b). Scientific study of telecommunication deployment in achieving quality network. International Journal of Trend in Research and Development, 4(5), 311-314.

[11] Aminu-Ibrahim, A. Y., Ogbete, J. C., & Ambali, K. B. (2018). Developing sustainable diagnostic laboratory infrastructure models for emerging and resource constrained health systems. Iconic Research and Engineering Journals, 1(8), 118-132. https://doi.org/10.64388/IREV1I8-1713586

[12] Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. Communications of the ACM, 53(4), 50-58.

[13] Avizienis, A., Laprie, J.-C., Randell, B., & Landwehr, C. (2004). Basic concepts and taxonomy of dependable and secure computing. IEEE Transactions on Dependable and Secure Computing, 1(1), 11-33.

[14] Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2018). A systematic review of CI/CD pipeline strategies in Salesforce DevOps: Tools, practices, and deployment outcomes. Iconic Research and Engineering Journals, 2(6).

[15] Barroso, L. A., Clidaras, J., & Hölzle, U. (2013). The datacenter as a computer: An introduction to the design of warehouse-scale machines (2nd ed.). Morgan & Claypool.

[16] Bates, D. W., & Gawande, A. A. (2003). Improving safety with information technology. New England Journal of Medicine, 348(25), 2526-2534.

[17] Bernstein, P. A., Hadzilacos, V., & Goodman, N. (1987). Concurrency control and recovery in database systems. Addison-Wesley.

[18] Blumenthal, D., & Tavenner, M. (2010). The "meaningful use" regulation for electronic health records. New England Journal of Medicine, 363(6), 501-504.

[19] Brewer, E. A. (2000). Towards robust distributed systems [Keynote address]. Proceedings of the 19th Annual ACM Symposium on Principles of Distributed Computing, Portland, OR.

[20] Candea, G., & Fox, A. (2003). Crash-only software. Proceedings of the 9th Workshop on Hot Topics in Operating Systems (pp. 67-72). USENIX Association.

[21] Cecchet, E., Candea, G., & Ailamaki, A. (2008). Middleware-based database replication: The gaps between theory and practice. Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data (pp. 739-752). ACM.

[22] Centers for Medicare & Medicaid Services. (2016). Medicare and Medicaid programs; Emergency preparedness requirements for Medicare and Medicaid participating providers and suppliers. Federal Register, 81(180), 63860-64044.

[23] Cerullo, V., & Cerullo, M. J. (2004). Business continuity planning: A comprehensive approach. Information Systems Management, 21(3), 70-78.

[24] Chandra, T. D., Griesemer, R., & Redstone, J. (2007). Paxos made live: An engineering perspective. Proceedings of the 26th Annual ACM Symposium on Principles of Distributed Computing (pp. 398-407). ACM.

[25] Chen, P. M., Lee, E. K., Gibson, G. A., Katz, R. H., & Patterson, D. A. (1994). RAID: High-performance, reliable secondary storage. ACM Computing Surveys, 26(2), 145-185.

[26] Corbett, J. C., Dean, J., Epstein, M., Fikes, A., Frost, C., Furman, J. J., et al. (2012). Spanner: Google's globally-distributed database. Proceedings of the 10th USENIX Symposium on Operating Systems Design and Implementation (pp. 251-264).

[27] Coventry, L., & Branley, D. (2018). Cybersecurity in healthcare: A narrative review of trends, threats and ways forward. Maturitas, 113, 48-52.

[28] Cully, B., Lefebvre, G., Meyer, D., Feeley, M., Hutchinson, N., & Warfield, A. (2008). Remus: High availability via asynchronous virtual machine replication. Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation (pp. 161-174). USENIX Association.

[29] Dagodzo, D. (2018a). A conceptual framework for UAV integration into national power grid inspection programs. Iconic Research and Engineering Journals, 2(5), 391-412. https://doi.org/10.64388/IREV2I5-1716082

[30] DeCandia, G., Hastorun, D., Jampani, M., Kakulapati, G., Lakshman, A., Pilchin, A., et al. (2007). Dynamo: Amazon's highly available key-value store. Proceedings of the 21st ACM Symposium on Operating Systems Principles (pp. 205-220).

[31] Dogbatsey, E. A., & Ebhojie, O. (2018). Budget compliance and statutory reporting in Sub-Saharan Africa: A systematic review of frameworks, gaps, and reform pathways. Zenodo. https://doi.org/10.5281/zenodo.20446859 [Originally numbered 2(6); journal name not captured in source.]

[32] Ejofodomi, O. A., Gideon, E. N., Oladipo, G. O., & Oshomah, E. R. (2014). Automated detection of architectural detection in mammograms using template matching. International Journal of Biomedical Science and Engineering, 2(1), 1-6. https://doi.org/10.11648/j.ijbse.20140201.11

[33] Ekemezie, U. C., Ohunyon, G., & Sydney, E. (2018). Evaluating IDSR implementation in Benue State, Nigeria: DHIS2 data quality and surveillance capacity. Iconic Research and Engineering Journals, 2(6).

[34] Fernández-Alemán, J. L., Señor, I. C., Lozoya, P. Á. O., & Toval, A. (2013). Security and privacy in electronic health records: A systematic literature review. Journal of Biomedical Informatics, 46(3), 541-562.

[35] Ford, D., Labelle, F., Popovici, F. I., Stokely, M., Truong, V.-A., Barroso, L., Grimes, C., & Quinlan, S. (2010). Availability in globally distributed storage systems. Proceedings of the 9th USENIX Symposium on Operating Systems Design and Implementation (pp. 61-74). USENIX Association.

[36] Gilbert, S., & Lynch, N. (2002). Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. ACM SIGACT News, 33(2), 51-59.

[37] Gill, P., Jain, N., & Nagappan, N. (2011). Understanding network failures in data centers: Measurement, analysis, and implications. Proceedings of the ACM SIGCOMM 2011 Conference (pp. 350-361). ACM.

[38] Gray, J. (1981). The transaction concept: Virtues and limitations. Proceedings of the 7th International Conference on Very Large Data Bases (pp. 144-154).

[39] Gray, J. (1986). Why do computers stop and what can be done about it? Proceedings of the 5th Symposium on Reliability in Distributed Software and Database Systems (pp. 3-12). IEEE Computer Society.

[40] Gray, J. (1990). A census of Tandem system availability between 1985 and 1990. IEEE Transactions on Reliability, 39(4), 409-418.

[41] Gray, J., & Reuter, A. (1993). Transaction processing: Concepts and techniques. Morgan Kaufmann.

[42] Gray, J., & Siewiorek, D. P. (1991). High-availability computer systems. Computer, 24(9), 39-48.

[43] Gray, J., Helland, P., O'Neil, P., & Shasha, D. (1996). The dangers of replication and a solution. Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data (pp. 173-182).

[44] Gunawi, H. S., Hao, M., Suminto, R. O., Laksono, A., Satria, A. D., Adityatama, J., & Eliazar, K. J. (2016). Why does the cloud stop computing? Lessons from hundreds of service outages. Proceedings of the Seventh ACM Symposium on Cloud Computing (pp. 1-16). ACM.

[45] Hamilton, J. (2007). On designing and deploying Internet-scale services. Proceedings of the 21st Large Installation System Administration Conference (pp. 231-242). USENIX Association.

[46] Hanuscak, T. L., Szeinbach, S. L., Seoane-Vazquez, E., Reichert, B. J., & McCluskey, C. F. (2009). Evaluation of causes and frequency of medication errors during information technology downtime. American Journal of Health-System Pharmacy, 66(12), 1119-1124. https://doi.org/10.2146/ajhp080389

[47] Härder, T., & Reuter, A. (1983). Principles of transaction-oriented database recovery. ACM Computing Surveys, 15(4), 287-317.

[48] Health Information Technology for Economic and Clinical Health Act, Pub. L. No. 111-5, Div. A, Title XIII, 123 Stat. 226 (2009).

[49] Health Insurance Portability and Accountability Act of 1996, Pub. L. No. 104-191, 110 Stat. 1936 (1996).

[50] Hellerstein, J. M., Stonebraker, M., & Hamilton, J. (2007). Architecture of a database system. Foundations and Trends in Databases, 1(2), 141-259.

[51] Herbane, B. (2010). The evolution of business continuity management: A historical review of practices and drivers. Business History, 52(6), 978-1002.

[52] International Organization for Standardization & International Electrotechnical Commission. (2011). ISO/IEC 27031:2011 Information technology: Security techniques: Guidelines for information and communication technology readiness for business continuity. ISO.

[53] International Organization for Standardization. (2012). ISO 22301:2012 Societal security: Business continuity management systems: Requirements. ISO.

[54] International Organization for Standardization. (2016). ISO 27799:2016 Health informatics: Information security management in health using ISO/IEC 27002. ISO.

[55] Jarrett, M. P. (2017). Cybersecurity: A serious patient care concern. JAMA, 318(14), 1319-1320.

[56] Jha, A. K., DesRoches, C. M., Campbell, E. G., Donelan, K., Rao, S. R., Ferris, T. G., Shields, A., Rosenbaum, S., & Blumenthal, D. (2009). Use of electronic health records in U.S. hospitals. New England Journal of Medicine, 360(16), 1628-1638.

[57] Ji, M., Veitch, A., & Wilkes, J. (2003). Seneca: Remote mirroring done write. Proceedings of the USENIX Annual Technical Conference (pp. 253-268).

[58] Keeton, K., Santos, C., Beyer, D., Chase, J., & Wilkes, J. (2004). Designing for disasters. Proceedings of the 3rd USENIX Conference on File and Storage Technologies (pp. 59-72).

[59] Kemme, B., & Alonso, G. (2000). Don't be lazy, be consistent: Postgres-R, a new way to implement database replication. Proceedings of the 26th International Conference on Very Large Data Bases (pp. 134-143).

[60] Kemme, B., & Alonso, G. (2010). Database replication: A tale of research across communities. Proceedings of the VLDB Endowment, 3(1-2), 5-12.

[61] Kharraz, A., Robertson, W., Balzarotti, D., Bilge, L., & Kirda, E. (2015). Cutting the Gordian knot: A look under the hood of ransomware attacks. In M. Almgren, V. Gulisano, & F. Maggi (Eds.), Detection of intrusions and malware, and vulnerability assessment (Lecture Notes in Computer Science, Vol. 9148, pp. 3-24). Springer.

[62] Kilbridge, P. (2003). Computer crash: Lessons from a system failure. New England Journal of Medicine, 348(10), 881-882.

[63] Kitchenham, B., & Charters, S. (2007). Guidelines for performing systematic literature reviews in software engineering (EBSE Technical Report EBSE-2007-01). Keele University and Durham University.

[64] Koppel, R., Metlay, J. P., Cohen, A., Abaluck, B., Localio, A. R., Kimmel, S. E., & Strom, B. L. (2005). Role of computerized physician order entry systems in facilitating medication errors. JAMA, 293(10), 1197-1203.

[65] Kruse, C. S., Frederick, B., Jacobson, T., & Monticone, D. K. (2017). Cybersecurity in healthcare: A systematic review of modern threats and trends. Technology and Health Care, 25(1), 1-10.

[66] Kuhn, K. A., & Giuse, D. A. (2001). From hospital information systems to health information systems: Problems, challenges, perspectives. Methods of Information in Medicine, 40(4), 275-287.

[67] Kuo, A. M.-H. (2011). Opportunities and challenges of cloud computing to improve health care services. Journal of Medical Internet Research, 13(3), e67.

[68] Ladapo, O. O., Dosunmu, A. A., Jooda, D., & Abolaji, T. O. (2018). Lessons learned from offline assessment of security-critical systems: The case of Microsoft Active Directory. Iconic Research and Engineering Journals, 2(6), 277-299. https://doi.org/10.64388/IREV2I6-1717205

[69] Lamport, L. (1998). The part-time parliament. ACM Transactions on Computer Systems, 16(2), 133-169.

[70] Larsen, E., Fong, A., Wernz, C., & Ratwani, R. M. (2018). Implications of electronic health record downtime: An analysis of patient safety event reports. Journal of the American Medical Informatics Association, 25(2), 187-191.

[71] Lawal, O. A., & Oduleye, T. E. (2018a). A conceptual model for financial analytics driven enterprise value creation in technology firms. Iconic Research and Engineering Journals, 2(2).

[72] Löhr, H., Sadeghi, A.-R., & Winandy, M. (2010). Securing the e-health cloud. Proceedings of the 1st ACM International Health Informatics Symposium (pp. 220-229). ACM.

[73] Magrabi, F., Ong, M.-S., Runciman, W., & Coiera, E. (2012). Using FDA reports to inform a classification for health information technology safety problems. Journal of the American Medical Informatics Association, 19(1), 45-53.

[74] Malviya, N., Weisberg, A., Madden, S., & Stonebraker, M. (2014). Rethinking main memory OLTP recovery. Proceedings of the 30th IEEE International Conference on Data Engineering (pp. 604-615). IEEE.

[75] Martin, G., Martin, P., Hankin, C., Darzi, A., & Kinross, J. (2017). Cybersecurity and healthcare: How safe are we? BMJ, 358, j3179.

[76] Mbonu, I. S., Aliliele, C., Iwuanyanwu, U., & Oluoha, O. M. (2018). A conceptual framework for legal and ethical risk modeling in enterprise data protection governance systems. Iconic Research and Engineering Journals, 2(2), 207-226. https://doi.org/10.64388/IREV2I2-1714911

[77] Mell, P., & Grance, T. (2011). The NIST definition of cloud computing (NIST Special Publication 800-145). National Institute of Standards and Technology.

[78] Minhas, U. F., Rajagopalan, S., Cully, B., Aboulnaga, A., Salem, K., & Warfield, A. (2011). RemusDB: Transparent high availability for database systems. Proceedings of the VLDB Endowment, 4(11), 738-748.

[79] Mohan, C., Haderle, D., Lindsay, B., Pirahesh, H., & Schwarz, P. (1992). ARIES: A transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging. ACM Transactions on Database Systems, 17(1), 94-162.

[80] Moher, D., Liberati, A., Tetzlaff, J., Altman, D. G., & The PRISMA Group. (2009). Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement. PLoS Medicine, 6(7), e1000097.

[81] National Audit Office. (2017). Investigation: WannaCry cyber attack and the NHS (HC 414). National Audit Office.

[82] Nwangwu, C. V., Duvbiama-Owasanoye, O. A., & Onwumere, A. (2018). A framework for operational integrity in high-risk drilling: Balancing productivity and safety compliance. Iconic Research and Engineering Journals, 2(5), 436-453.

[83] Obetta, P. C. (2018). A risk-based inspection concept for fitness-for-service assessment of ageing pressure vessels on offshore joint venture facilities. Iconic Research and Engineering Journals, 2(5), 497-520. https://doi.org/10.64388/IREV2I5-1722578

[84] Obetta, P. C., & Ninduwezuor-Ehiobu, N. (2018). A condition-based maintenance concept for CALM buoy bearings and swivels to improve offshore crude export availability. Iconic Research and Engineering Journals, 2(6), 385-410. https://doi.org/10.64388/IREV2I6-1722579

[85] Obriki, O. D., & Arumosoye, O. M. (2018). Conceptual modeling of data-driven occupational safety risk control in large-scale energy infrastructure projects. Iconic Research and Engineering Journals, 1(7), 169-189. https://doi.org/10.64388/IREV1I7-1714414

[86] Odejobi, O. D., & Ahmed, K. S. (2018b). Performance evaluation model for multi-tenant Microsoft 365 deployments under high concurrency. IRE Journals, 1(11), 92-107.

[87] Office of the National Coordinator for Health Information Technology. (2014). SAFER guides: Contingency planning. U.S. Department of Health and Human Services.

[88] Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2018). Optimizing laboratory spatial planning strategies to improve diagnostic accuracy, safety, and clinical throughput. Iconic Research and Engineering Journals, 2(1), 87-113. https://doi.org/10.64388/IREV1I7-1713587

[89] Okonkwo, C. S., Ogunwole, O., & Okeke, O. T. (2018b). Model for inventory availability and plant uptime improvement in energy facilities. Iconic Research and Engineering Journals, 2(4), 160-172. https://doi.org/10.64388/IREV2I4-1713120

[90] Ongaro, D., & Ousterhout, J. (2014). In search of an understandable consensus algorithm. Proceedings of the USENIX Annual Technical Conference (pp. 305-319).

[91] Oppenheimer, D., Ganapathi, A., & Patterson, D. A. (2003). Why do Internet services fail, and what can be done about it? Proceedings of the 4th USENIX Symposium on Internet Technologies and Systems. USENIX Association.

[92] Oshevire, P., Eyenubo, O. J., & Amayo, B. (2017). Voltage control in the presence of distributed generation. ATBU Journal of Science, Technology and Education, 5(2), 165-173.

[93] Patterson, D. A., Gibson, G., & Katz, R. H. (1988). A case for redundant arrays of inexpensive disks (RAID). Proceedings of the 1988 ACM SIGMOD International Conference on Management of Data (pp. 109-116).

[94] Patterson, D., Brown, A., Broadwell, P., Candea, G., Chen, M., Cutler, J., Enriquez, P., Fox, A., Kiciman, E., Merzbacher, M., Oppenheimer, D., Sastry, N., Tetzlaff, W., Traupman, J., & Treuhaft, N. (2002). Recovery-oriented computing (ROC): Motivation, definition, techniques, and case

studies (Technical Report UCB//CSD-02-1175). University of California, Berkeley.

[95] Patterson, R. H., Manley, S., Federwisch, M., Hitz, D., Kleiman, S., & Owara, S. (2002). SnapMirror: File-system-based asynchronous mirroring for disaster recovery. Proceedings of the 1st USENIX Conference on File and Storage Technologies (pp. 117-129). USENIX Association.

[96] Pinheiro, E., Weber, W.-D., & Barroso, L. A. (2007). Failure trends in a large disk drive population. Proceedings of the 5th USENIX Conference on File and Storage Technologies (pp. 17-29). USENIX Association.

[97] Sax, U., Lipprandt, M., & Röhrig, R. (2016). The rising frequency of IT blackouts indicates the increasing relevance of IT emergency concepts to ensure patient safety. Yearbook of Medical Informatics, 25(1), 130-137. https://doi.org/10.15265/IY-2016-038

[98] Schneider, F. B. (1990). Implementing fault-tolerant services using the state machine approach: A tutorial. ACM Computing Surveys, 22(4), 299-319.

[99] Scholl, M., Stine, K., Hash, J., Bowen, P., Johnson, A., Smith, C. D., & Steinberg, D. I. (2008). An introductory resource guide for implementing the Health Insurance Portability and Accountability Act (HIPAA) Security Rule (NIST Special Publication 800-66 Rev. 1). National Institute of Standards and Technology.

[100] Schroeder, B., & Gibson, G. A. (2007). Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you? Proceedings of the 5th USENIX Conference on File and Storage Technologies (pp. 1-16). USENIX Association.

[101] Sittig, D. F., & Singh, H. (2010). A new sociotechnical model for studying health information technology in complex adaptive healthcare systems. Quality and Safety in Health Care, 19(Suppl. 3), i68-i74.

[102] Sittig, D. F., Gonzalez, D., & Singh, H. (2014). Contingency planning for electronic health record-based care continuity: A survey of recommended practices. International Journal of Medical Informatics, 83(11), 797-804.

[103] Snedaker, S. (2013). Business continuity and disaster recovery planning for IT professionals (2nd ed.). Syngress.

[104] Sunday, E. A., & Omoegun, G. O. (2018). Integrating solar power solutions in small-scale manufacturing industries in Nigeria. International Journal of Scientific Research in Science, Engineering and Technology, 4(8), 832-853.

[105] Swanson, M., Bowen, P., Phillips, A. W., Gallup, D., & Lynes, D. (2010). Contingency planning guide for federal information systems (NIST Special Publication 800-34 Rev. 1). National Institute of Standards and Technology.

[106] Telecommunications Industry Association. (2005). ANSI/TIA-942: Telecommunications infrastructure standard for data centers. TIA.

[107] U.S. Department of Health and Human Services, Office for Civil Rights. (2016). Fact sheet: Ransomware and HIPAA. HHS.

[108] U.S. Department of Health and Human Services. (2003). Health insurance reform: Security standards; Final rule. Federal Register, 68(34), 8334-8381.

[109] Uchechi, N., Ominyi, M., & Jones, B. C. (2018). Procedural standardization and workflow optimization in high-volume dispute resolution clinics: A systematic review. Iconic Research and Engineering Journals, 2(6), 420-452. https://doi.org/10.64388/IREV2I6-1722789

[110] Ukadike, C. A. (2018). Competency frameworks and onboarding redesign in professional services: A conceptual model for return on investment. International Journal of Social Sciences and Management Research, 4(9), 42-62. https://doi.org/10.56201/ijssmr.v4.no9.2018.pg.42.62

[111] Unamma, U. M.-L., Okwesa, I. J., Iwuanyanwu, U., & Ajiga, D. (2018). An offline-first design framework for digital learning in low-connectivity environments. Iconic Research and Engineering Journals,

2(3), 192-211. https://doi.org/10.64388/IREV2I3-1722627

[112] Verbitski, A., Gupta, A., Saha, D., Brahmadesam, M., Gupta, K., Mittal, R., Krishnamurthy, S., Maurice, S., Kharatishvili, T., & Bao, X. (2017). Amazon Aurora: Design considerations for high throughput cloud-native relational databases. Proceedings of the 2017 ACM International Conference on Management of Data (pp. 1041-1052). ACM.

[113] Vogels, W. (2009). Eventually consistent. Communications of the ACM, 52(1), 40-44.

[114] Wang, Y., Coiera, E., Gallego, B., Concha, O. P., Ong, M.-S., Tsafnat, G., Roffe, D., Jones, G., & Magrabi, F. (2016). Measuring the effects of computer downtime on hospital pathology processes. Journal of Biomedical Informatics, 59, 308-315.

[115] Wiesmann, M., Pedone, F., Schiper, A., Kemme, B., & Alonso, G. (2000). Understanding replication in databases and distributed systems. Proceedings of the 20th International Conference on Distributed Computing Systems (pp. 464-474). IEEE.

[116] Wood, T., Cecchet, E., Ramakrishnan, K. K., Shenoy, P., van der Merwe, J., & Venkataramani, A. (2010). Disaster recovery as a cloud service: Economic benefits and deployment challenges. Proceedings of the 2nd USENIX Workshop on Hot Topics in Cloud Computing. USENIX Association.

[117] Yusuf, H. D., & Abdulrazaq, K. (2018). A conceptual framework for strengthening diagnostic quality systems in resource-limited clinical laboratories. Iconic Research and Engineering Journals, 1(11), 145-179. https://doi.org/10.64388/IREV1I11-1723090

[118] Zhang, R., & Liu, L. (2010). Security models and requirements for healthcare application clouds. Proceedings of the 3rd IEEE International Conference on Cloud Computing (pp. 268-275). IEEE.