ITUPublications

International Telecommunication Union
Standardization Sector

ITU-T Technical Paper

(06/2026)

XSTP.epoch

Global coordination requirements for 2038-class rollover events (including but not limited to 2036, 2038, 2106) – June 2026 edition

Technical Paper ITU-T XSTP.epoch

Global coordination requirements for 2038-class rollover events (including but not limited to 2036, 2038, 2106) – June 2026 edition

Summary

Technical Paper ITU-T XSTP.epoch consolidates global understanding of risks associated with a cluster of structurally distinct timekeeping discontinuities arriving between 2036 and 2038, with the 2106 horizon treated as a long-tail consideration (collectively referred to as 2038-class risks). The Technical Paper documents known and emerging failure patterns across the deployed substrate of ICT systems incorporating 32-bit time representations, and examines systemic challenges in discovery, inventory, in-situ versus simulated testing, coordinated vulnerability management, and cross-sector coordination.

The paper develops the structural limitations of remote detection and laboratory simulation for 2038-class failure characterization, drawing on historical analogues including the 1996 Ariane 5 incident to illustrate why in-situ validation under operational conditions is the binding requirement where simulation-based testing cannot capture environment-specific behaviour. Effective response to 2038-class risks depends on multi-stakeholder coordination across operators, vendors, regulators, certification bodies, and standards organizations whose ordinary remediation pathways are not constituted to address risks that span their jurisdictional boundaries – and on the common terminology, conceptual models, and shared prioritization approaches that such coordination requires.

It establishes shared terminology, a coherent problem framing, and a high-level analytical foundation; identifies areas requiring further study and inter-SDO collaboration; and surfaces operational response categories including harm-minimization strategies for substrate conditions where full remediation is not achievable on the available timeline. Its objective is to support Member States, operators, vendors, standards bodies, and incident response communities in preparing for potential correlated failure pathways before rollover thresholds are reached.

Keywords

2038-class rollover, coherent time, critical infrastructure coordination, epoch rollover, non-intentional digital disruption, NTP era rollover, temporal root of trust, time security, triage-by-failure-tempo.

Note

This is an informative ITU-T publication. Mandatory provisions, such as those found in ITU-T Recommendations, are outside the scope of this publication. This publication should only be referenced bibliographically in ITU-T Recommendations.

Change Log

This document contains the June 2026 edition of the ITU-T Technical Paper XSTP.epoch on "*Global coordination requirements for 2038-class rollover events (including but not limited to 2036, 2038, 2106)*" approved at the ITU-T Study Group 17 meeting held in Geneva, 1-10 June 2026.

This Technical Paper is maintained on an active revision cycle aligned with ITU-T Study Group 17 plenary meetings. Each subsequent revision is anticipated to land in coordination with an SG17 plenary, with interim editorial updates between plenaries as substantive material warrants. Table 1 summarises the approximate remaining update window for each cluster event.

Table 1 – Approximate remaining update window by cluster event

Event	Date	Plenary cycles remaining	Annual quarters remaining
NTP era rollover	7 February 2036	~19	~39
Unix time_t overflow	19 January 2038	~23	~47
GPS week-number rollover	20 November 2038	~25	~50
NOTE – Assuming the SG17 plenary cadence of approximately two plenaries per year is sustained across the 2025-2028 study period and successive study periods, at the time of this revision (June 2026).			

The cluster of structurally distinct timekeeping discontinuities described in clause 6 arrives within a finite, knowable, and rapidly contracting window. The number of revision cycles this Technical Paper will undergo before the events it describes is bounded – and small. Each cycle is an opportunity to expand sectoral coverage, surface additional documented incidents, refine the analytical framework, and integrate the empirical signals that the 2036 NTP era rollover will produce.

NOTE – The work this paper supports is finite in duration, structurally bounded by the events it addresses, and dependent on sustained cross-sector contribution across the remaining update window.

Acknowledgements

This Technical Paper was prepared under the leadership of Mr Gianpaolo Scalone (Vodafone, United Kingdom) with contributions from André Baptista (Ethiack, Portugal), Vint Cerf (Individual, United States), Trey Darley (Proper Tools SRL, Belgium), Brian DeWyngaert Jr. (CISA, United States), Lorenzo Dina (Nobug, Italy), James N. Duncan (Individual, United States), Andres Erni (BEOZ Association, Switzerland), Allan Friedman, PhD (Institute for Security and Technology, United States), Daniel E. Geer, PhD (In-Q-Tel, United States), Chris Gibson (FIRST, United Kingdom), Wes Hardaker (Google, United States), Christian (fukami) Horchert (CrabNebula, Ltd, Malta), Christian O. Hunt (Voidlab STR, United States), Pascal Huppert (University of Twente / University of Münster, Netherlands), Shreyas Konjerla (TU Delft, Netherlands), Georgy Kucherin (Kaspersky, Russian Federation), John Lange (Y2038.com LLC, United States), William J. Malik (Individual, United States), Danny Mayer (Network Time Foundation, United States), Sarah Novotny (Trinity College Dublin, Ireland), Niyo Little Thunder Pearson (Individual, United States), Harlan Stenn (Network Time Foundation, United States), Laurie Tyzenhaus (Software Engineering Institute, United States), Pedro Umbelino (Bitsight Technologies, Portugal), and Bernhard M. Wiedemann (SUSE, Germany).

Names and affiliations are shown as provided by the contributors; listing reflects each contributor's consent to be named and does not imply institutional endorsement.

© ITU 2026

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

Table of Contents

		Page
1	Scope	1
2	References	2
3	Definitions	2
3.1	Terms defined elsewhere	2
3.2	Terms defined in this Technical Paper	2
4	Abbreviations and acronyms	4
5	Reader's guide	6
6	Background and motivation	6
6.1	Foundational time concepts	10
6.2	Failure mode taxonomy	10
6.3	Rollover event timeline	15
6.4	Risk windows spanning canonical dates	16
7	Problem statement	16
7.1	Observed failure modes in real systems	18
7.2	Temporal root of trust dependencies	19
8	Analysis framework	21
8.1	Integrity of time-dependent operations	21
8.2	Resilience and correlated failure pathways	22
8.3	Lifecycle and supply-chain exposure	23
8.4	Observability and detectability limitations	25
8.5	Protocol interoperability dependencies	27
8.6	Critical infrastructure considerations	29
8.7	Coordinated vulnerability management constraints	31
8.8	Implications for incident response and digital forensics	33
8.9	Coherence invariant and the distributed-proof framing	34
9	Reference architecture	35
9.1	Global time dependency stack	35
9.2	Rollover propagation chain	35
9.3	Cloud and platform synchronization dependencies	35
9.4	Protocol layers with embedded timestamp fields	36
10	Cross-sector exposure survey	36
10.1	Telecommunication networks	41
10.2	Financial systems	45
10.3	Biological research and pharmaceuticals	47
10.4	Electric and connected vehicles	48
10.5	Implantable and hospital medical devices	49
10.6	CSIRT/PSIRT/CERT coordination	51

Page

10.7 Emergency services operations 52
10.8 ICS/OT environments 52
10.9 Cloud and ICT infrastructure 52
10.10 GNSS-dependent services 53
10.11 Space infrastructure 53
10.12 Broadcasting and emergency alerting 54
10.13 Government services and public administration 54
10.14 Insurance and reinsurance 55
10.15 Epistemic coordination layer 55
10.16 Buildings and built environment 56
10.17 Manufacturing and industrial production 57
10.18 Logistics and distribution 57
10.19 Mining and extraction 58
10.20 Water and wastewater 58
10.21 Food and farming 58
10.22 Oil and gas 59
10.23 Power plants 59
10.24 Power grid 59
10.25 Vehicles and traffic infrastructure 60
10.26 Railroad 60
10.27 Bridges and tunnels 60
10.28 Aviation 61
10.29 Shipping 61
10.30 Wastes and recycling 61

11 Threat model 62
11.1 Threat categories 62

12 Research landscape and empirical observations 62
12.1 Publicly reported field observations 62
12.2 Academic and conference literature 62

13 Identified inter-SDO gaps 62
13.1 Inter-SDO coverage map 62

14 Gap analysis 63
14.1 Discovery and inventory limitations 63
14.2 Testing and testbed gaps 66
14.3 Cross-border coordination challenges 66
14.4 Supply chain dependencies 66
14.5 Absence of inter-SDO harmonization 66

15 Recommendations 67
15.1 Shared terminology and conceptual framework 67

Page

15.2 Testing guidance 67
15.3 Coordination practices 68
15.4 Inter-SDO liaison and certification considerations 68
15.5 Procurement considerations 69
15.6 Harm minimization considerations 70
15.7 Areas for future study 72
15.8 Success criteria for this Technical Paper 72
Bibliography 74

Introduction

This Technical Paper addresses global coordination requirements related to a 2036-2038 cluster of timestamp rollover events: the network time protocol era rollover (2036), the overflow of 32-bit signed Unix time_t (2038), and the next GPS week-number rollover (2038). The 2106 horizon is treated as a long-tail consideration. These events are deterministic in timeline and heterogeneous in operational effect; their temporal proximity compounds the coordination challenge they collectively present.

The Technical Paper develops the calendar-deadline-specific instance of the non-intentional digital disruption category surfaced in [b-WTISD26]. It introduces the concept of a *temporal root of trust*, the implicit assumption that participating components in a distributed system observe a coherent time and examines how security properties downstream of this assumption degrade together when it is violated, whether through rollover, manipulation, synchronization failure, or implementation divergence. The paper distinguishes 2038-class events from conventional software vulnerabilities by surfacing the application-binary-interface character of the underlying remediation, which differentiates 2038-class response from ordinary patch-cycle response in ways consequential for how preparation work is structured.

The paper consolidates known and emerging failure patterns across the deployed substrate, organising them into a taxonomy of seven failure-mode classes (representation and range, interpretation and semantic mismatch, synchronization and coherence, lifecycle and longevity, testing and observability gaps, compound interactions, and delayed and operationally cadenced manifestation). It surveys cross-sector exposure across thirty sectors organized against the [b-Hayes] infrastructure taxonomy, with developed treatments of telecommunications, financial systems, electric and connected vehicles, and implantable and hospital medical devices, and with scaffolded sector stubs supporting cross-sector contribution in subsequent revisions. The paper examines structural challenges in discovery, asset inventory, in-situ versus simulated testing, and coordinated vulnerability management; it surfaces correlated and silent failure modes, lifecycle and supply-chain inheritance of vulnerabilities, interoperability dependencies that amplify risk at scale, and the *epistemic coordination layer* on which every other sector's remediation work depends including the structural conditions of institutional-memory loss and AI-mediated reproduction of legacy patterns that compound the substrate's remediation profile.

The Technical Paper does not introduce normative requirements or prescribe mitigation measures. It surfaces operational response categories distinct from full remediation, including harm-minimization strategies for substrate conditions where full remediation is not achievable on the available timeline, and a *triage-by-failure-tempo* organising principle for prioritising response across sectors. Its objective is to establish shared terminology, a coherent problem framing, and a high-level analytical foundation to support Member States, operators, vendors, standards bodies, and incident response communities in preparing for rollover-class risks and identifying areas requiring further study or inter-SDO coordination.

Editorial note

This Technical Paper is intended as a living document, maintained across multiple revisions during the period leading up to and beyond the 2038-class rollover events it addresses. The editors approach the document with explicit awareness of several conditions that may bear on its durability and usefulness.

The deployed technical infrastructure that the paper describes will, in many cases, outlive the operational capacity of the organizations and individuals presently responsible for it. The remediation pathways the paper addresses depend on knowledge that is not transferred at the rate at which it is being lost – a structural condition developed in clause 10.15.

The period during which this paper will be read and referenced coincides with a period in which substantial volumes of variable-quality content about 2038-class risks will enter circulation. Some of this content will be produced or mediated by AI systems whose training corpora incorporate prior content in the same domain, producing a recursive dynamic in which content generation feeds training distributions that feed further content generation. The paper's value as an authoritative reference will depend in part on its resistance to dilution by such contamination and on the maintenance of clear provenance for its substantive claims.

This Technical Paper is therefore developed under a stewardship posture: cross-sector contribution from domain experts, transparency of editorial logic, careful distinction between what is established and what remains under development, and acknowledgement that the document's eventual usefulness will depend less on the completeness of any single revision than on the durability of the collaborative process through which subsequent revisions are produced.

Certain sections remain deliberately scaffolded pending future development — most notably the threat model (clause 11) and the research landscape (clause 12), which carry drafting notes and are staged for completion in subsequent revisions.

A further consideration warrants direct acknowledgement. This Technical Paper does not presume the failure scenarios it surveys will materialize. The structural conditions developed across clause 8 and clause 10 indicate that 2038-class events will exercise the deployed substrate in ways the substrate has not previously been tested against; whether that exercise produces systemic coherence events or localized, recoverable failures absorbed by accumulated redundancy is an empirical question this Technical Paper is positioned to inform rather than to resolve. Where subsequent revisions develop empirical evidence that the substrate retains coherence under stress, including evidence from the 2036 NTP era boundary as a leading indicator, the Technical Paper will record that finding as directly as it records the structural conditions that motivated the work.

Technical Paper ITU-T XSTP.epoch

Global coordination requirements for 2038-class rollover events (including but not limited to 2036, 2038, 2106) – June 2026 edition

1 Scope

While this Technical Paper does not prescribe mitigation strategies, it highlights the structural limitations of relying solely on laboratory or simulated testing for 2038-class rollover failures.

Similar structural risks have been documented in other complex engineered systems. One of these is the [b-Ariane 5] launch failure, in which inertial reference code reused from the previous system behaved reliably under simulation but failed catastrophically in operation because the new vehicle's trajectory exceeded the value ranges the legacy code had been validated against. This incident underscores how legacy assumptions can remain hidden until a system enters a new operational envelope – a failure mode analogous to 2038-class rollover risks in modern information and communication technology (ICT) ecosystems. The Technical Paper therefore emphasizes the importance of in-situ validation and cross-disciplinary awareness when considering the assessment and testing of time-related dependencies.

This Technical Paper is intended for Member States, telecommunication operators, service and cloud providers, equipment and software vendors, critical infrastructure operators, national cybersecurity authorities, including CERTs/CSIRTs, and standards development organizations concerned with time-dependent ICT systems. It serves as a shared reference, coordination aid and consolidates existing knowledge on 2038-class timestamp rollover risks and provides common terminology and analytical framing to support cross-sector and cross-border understanding, risk awareness, and long-term planning.

As this Technical Paper is non-normative, it does not define compliance requirements, testing mandates, protocol changes, certification criteria, or prescriptive mitigation actions. It is not intended to replace existing standards or regulatory frameworks, but to provide a neutral foundation to support informed discussion, further study, and potential future work within ITU-T Study Group 17 and other relevant bodies.

Addressing 2038-class timestamp rollover risks requires long-term, cross-sector, and cross-border coordination. Affected systems span telecommunications, cloud services, embedded and industrial platforms, and critical infrastructure, often with extended lifecycles. No single stakeholder has full visibility or ownership of these dependencies.

This Technical Paper does not define governance structures or assign responsibilities. However, it highlights the importance of information sharing, alignment of terminology, and sustained coordination among Member States, operators, vendors, national cybersecurity authorities, and relevant standards development organizations. Existing mechanisms within ITU-T Study Group 17, together with collaboration with external bodies and communities, may provide appropriate forums for such coordination.

Given the multi-year horizon of rollover class risks and the potential for correlated and silent failure modes, coordination activities may need to extend beyond conventional vulnerability disclosure or short-term remediation models. This includes maintaining continuity of awareness across study periods, supporting inter-SDO dialogue, and enabling consistent public risk framing over time.

These considerations are indicative only and are intended to support discussion and further study. They do not constitute prescriptive governance models, operational mandates, or formal coordination requirements.

This Technical Paper is intended to provide a consolidated problem framing and shared understanding of 2038-class timestamp rollover risks. It does not address, define, or imply the following:

– Normative requirements or obligations, including compliance criteria, conformance targets, or mandatory controls for Member States, operators, or vendors.
– Protocol specifications, protocol changes, or implementation mandates, including modifications to existing standards, interfaces, or time representations.
– Prescriptive mitigation strategies or migration plans, such as required remediation approaches, upgrade paths, or timelines for replacement of affected systems.
– Testing, certification, or audit requirements, including pass/fail criteria, readiness scoring, or comparative assessment between organizations or jurisdictions.
– Vulnerability disclosure procedures or timelines, including CVE assignment, severity scoring, or coordinated disclosure mandates.
– Attribution of responsibility or liability for rollover related failures across supply chains, operators, vendors, or other stakeholders.

The absence of these elements is intentional as this Technical Paper does not seek to replace existing standards, regulatory frameworks, or operational guidance. Its role is, rather, to support informed discussion, coordination, and further study by establishing a common reference point for understanding rollover class risks.

2 References

[ITU-T X.1051] Recommendation ITU-T X.1051 (2023), *Information security, cybersecurity and privacy protection – Information security controls based on ISO/IEC 27002 for telecommunications organizations.*

3 Definitions

3.1 Terms defined elsewhere

None.

3.2 Terms defined in this Technical Paper

This Technical Paper defines the following terms:

3.2.1 2038-class rollover event: A failure condition arising from finite binary time representations in computing systems, including but not limited to the Network Time Protocol era rollover on 7 February 2036, the overflow of 32-bit signed Unix time_t on 19 January 2038, the GPS week-number rollover on 20 November 2038, and the overflow of 32-bit unsigned timestamp representations in 2106. The defining characteristic is the interaction between bounded timestamp representations and system lifecycles that exceed their original design assumptions, rather than any specific calendar date. The name derives from the most commonly discussed case – the signed 32-bit Unix time_t value, measured from the 1 January 1970 epoch, which overflows at 03:14:07 UTC on 19 January 2038.

3.2.2 capacity-to-rebuild: The collective technical, institutional, and human capacity available across a substrate or sector to maintain, modify, or reconstruct deployed infrastructure when remediation is required. It is distinct from the operational state of the substrate itself: infrastructure may continue to function while capacity-to-rebuild has atrophied through commercial reorganization, personnel dispersion, documentation decay, or loss of build environments and toolchains. Where it decays below the threshold a coordinated remediation would require, the substrate exhibits the infrastructural vestigialisation condition (clause 3.2.6).

3.2.3 coherence invariant: A property of distributed systems, expressible informally as the condition that for any two participating systems, time-dependent validations produce compatible outcomes under shared protocol assumptions. It holds across the deployed substrate when participating systems share a coherent representation of time across event boundaries, and fails when partial migration, manipulation, or representational divergence produces non-convergent validation states. Its formal characterization is developed in clause 8.9.

3.2.4 coherent time: The property of a distributed system in which participating components share a common, monotonically advancing, and comparable notion of time sufficient for the system's coordination requirements. It does not require absolute accuracy against any external reference, only that participating systems agree on temporal ordering and progression within tolerances sufficient for their operational dependencies. Loss of coherent time is a foundational failure mode capable of undermining security, safety, and coordinated incident response.

3.2.5 epoch rollover: A failure condition in which a system's time representation, encoded as an offset from a defined reference epoch, reaches the upper bound of its representable range and either overflows, wraps around, or transitions to an undefined state. The term refers to the exhaustion of the representable range, not to the reference epoch itself, which remains unchanged. Epoch rollovers are deterministic in timeline but heterogeneous in effect, varying by representation width, signedness, epoch definition, and the behaviour of dependent software.

3.2.6 infrastructural vestigialisation: A structural condition in which deployed technical infrastructure continues to function while the institutional and human capacity required to maintain, modify, or reconstruct it has atrophied below the threshold a coordinated rebuild would require. The term draws on the biological notion of vestigialisation – the persistence of a structure after its underlying function has decayed — and on the systems literature on complexity collapse [b-Tainter]. 2038-class events do not produce this condition but expose its present extent; it interacts with capacity-to-rebuild (clause 3.2.2) and the epistemic coordination conditions of clause 10.15.

3.2.7 monotonicity / monotonic time: The property of a time representation in which successive observed values are non-decreasing – that is, time as observed by a consuming system progresses forward and does not reverse, jump backward, or repeat. It is a structural requirement for log ordering, distributed consensus, replay protection, and audit-trail integrity, and may be lost through clock adjustment (including ordinary NTP correction), representational overflow at rollover boundaries, or a trusted time source being replaced by one carrying an earlier value. Monotonicity at the system level does not guarantee it across systems; the latter depends on coherent time (clause 3.2.4).

3.2.8 temporal root of trust: A foundational assumption, often implicit in distributed system design, that participating components share access to coherent time as a precondition for higher-level security properties. Authentication validity periods, certificate lifetimes, replay protection, log ordering, distributed consensus, and audit-trail integrity all rest on it. When it is degraded – through rollover, manipulation, synchronization failure, or implementation divergence – the security properties anchored to it degrade with it.

3.2.9 time delta discontinuities: Many time calculations compute a delta between two reference points. Errors can arise when one value wraps or overflows, or when hard-coded values are used in the calculation. Such errors may surface long after the triggering event, or be fleeting – the delta incorrect only while the two reference values sit on opposite sides of a wrapping boundary.

3.2.10 time desynchronization: A condition in which two or more participating systems in a distributed environment hold representations of time that have diverged beyond the tolerances required for their coordination. It may arise from clock drift, failure of time-distribution mechanisms, manipulation of time sources, implementation divergence, or representational overflow. It is distinct from absolute inaccuracy: systems sharing an inaccurate time may stay synchronized with one another, while systems sharing accurate time may desynchronize – with consequences ranging from negligible to the immediate failure of systems that depend on coherent time (clause 3.2.4).

3.2.11 time security: The discipline concerned with the integrity, authenticity, and availability of time information as it is produced, distributed, and consumed within distributed systems. It extends beyond time synchronization – which concerns clock alignment – to the security of time-distribution protocols, the resilience of time sources against manipulation, and the implementation correctness of time-handling logic. It is closely tied to the temporal root of trust (clause 3.2.8).

3.2.12 triage-by-failure-tempo: An organising principle for prioritising response to 2038-class events according to the rate at which the human consequences of inadequate remediation manifest after rollover. Failure modes affecting life safety manifest at the shortest tempo and bear the highest priority; those affecting command-and-control, critical-infrastructure substrates, and economic systems manifest at progressively slower tempos, allowing correspondingly greater flexibility in response. It is developed further as an area for study in clause 15.7.

3.2.13 wall clock time: The civil-time representation of a moment as it would be read from a conventionally-displayed clock – calendar date, time-of-day, time-zone offset, and any civil adjustments such as daylight-saving transitions or leap-second insertions. It is the layer through which time is communicated between humans and through which most user-facing systems present temporal information and is distinct from the underlying offset-based representations of epoch rollover (clause 3.2.5) and the 2038-class rollover event (clause 3.2.1).

4 Abbreviations and acronyms

This Technical Paper uses the following abbreviations and acronyms:

ABI	Application Binary Interface
ATG	Automated Tank Gauge
ATSC	Advanced Television Systems Committee
BAS	Building Automation System
BGP	Border Gateway Protocol
BMS	Building Management System
BSS	Business Support System
CAD	Computer-Aided Dispatch
CBDC	Central Bank Digital Currency
CERT	Computer Emergency Response Team
CPE	Customer Premises Equipment
CRA	Cyber Resilience Act
CSIRT	Computer Security Incident Response Team
CVD	Coordinated Vulnerability Disclosure
CVE	Common Vulnerabilities and Exposures
CVSS	Common Vulnerability Scoring System
DAB+	Digital Audio Broadcasting Plus
DCS	Distributed Control System
DNS	Domain Name System
DNSSEC	Domain Name System Security extension
DTCC	Depository Trust and Clearing Corporation

DVB-T2	Digital Video Broadcasting – Second Generation Terrestrial
EAS	Emergency Alert System
ECU	Electronic Control Unit
ERP	Enterprise Resource Planning
ERTMS	European Rail Traffic Management System
FRMCS	Future Railway Mobile Communication System
FIRST	Forum of Incident Response and Security Teams
GMP	Good Manufacturing Practice
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
GSM	Global System for Mobile Communications
HACCP	Hazard Analysis and Critical Control Points
HVAC	Heating, Ventilation, and Air Conditioning
ICS	Industrial Control System
ICT	Information and Communication Technology
IPv4/6	Internet Protocol version 4/6
LIMS	Laboratory Information Management System
M2M	Machine-to-Machine
MES	Manufacturing Execution System
NITZ	Network Identity and Time Zone
NTP	Network Time Protocol
OSS	Operations Support System
OT	Operational Technology
P25	Project 25
PKI	Public Key Infrastructure
PLC	Programmable Logic Controller
POS	Point of Sale
PSAP	Public Safety Answering Point
PSIRT	Product Security Incident Response Team
PTP	Precision Time Protocol
RPKI	Resource Public Key Infrastructure
RTGS	Real-Time Gross Settlement
RTLS	Real-Time Location System
SBOM	Software Bill of Materials
SCADA	Supervisory Control and Data Acquisition
SIG	Special Interest Group
SDO	Standards Development Organization

SNTP	Simple Network Time Protocol
SOA	Start of Authority
SSH	Secure Shell
TBoM	Temporal Bill of Material
TETRA	Terrestrial Trunked Radio
time_t	POSIX time type
TLS	Transport Layer Security
URLLC	Ultra-Reliable Low-Latency Communication
UTC	Coordinated Universal Time
V2X	Vehicle-to-everything
VEX	Vulnerability Exploitability exchange

5 Reader's guide

This Technical Paper is organized in three movements. clause 1 through clause 8 establish the problem space: scope, terminology, observed failure modes, and the analytical framework through which 2038-class rollover risks are examined in the clauses that follow. Clause 9 and clause 10 develop the cross-sector substance: a reference architecture for time-dependency analysis, and an extended survey of sectoral exposure organized against the [b-Hayes] infrastructure taxonomy. Clause 11 through clause 15 address the coordination response: threat models, inter-SDO mapping, gap analysis, recommendations, and references for further study.

Readers pressed for time are encouraged to read clause 10 before forming a view on the urgency or scope of the work this Technical Paper describes.

Clause 10 develops the document's primary empirical contribution – a survey of thirty sectors whose 2038-class exposure has been identified by the Forum of Incident Response and Security Teams (FIRST) Time Security Special Interest Group (SIG) community and SG17 contributors, with developed treatments of telecommunications, financial systems, electric and connected vehicles, and implantable and hospital medical devices. The cross-sector substance bears directly on whether 2038-class risks should be regarded as a sector-specific concern or as a substrate condition affecting the broader ICT landscape.

Readers approaching this document at the institutional or policy level may wish to begin with the Editorial Note above, which describes the document's stewardship posture and its intended evolution across subsequent revisions. Readers concerned with ITU-T process and inter-SDO coordination may wish to begin with clause 13 and clause 15.

NOTE – Passages marked "drafting note:" identify material scaffolded in the present revision and flagged for development in a subsequent one. They are editorial signposts rather than content and are expected to be resolved as the sections they annotate are completed. In most cases they are also invitations for subject matter experts with pertinent experience to join this effort and help us develop these areas more completely.

This Technical Paper is non-normative. The institutional context for that choice is described in clause 1.

6 Background and motivation

This Technical Paper addresses not a single timekeeping event but a cluster of structurally distinct timekeeping discontinuities arriving within a 3-year window. Listed in chronological order, the cluster comprises: the network time protocol era rollover on 7 February 2036; the overflow of 32-bit signed Unix time_t on 19 January 2038; and the next GPS week-number rollover expected on

20 November 2038. Each event has different exposure characteristics, affects different installed bases, and produces different operational consequences. Their temporal proximity is itself a systemic condition: remediation capacity claimed by responses to the earlier events in the cluster is unavailable for responses to the later events.

This Technical Paper develops a category that has recently been identified at the institutional level. The expert report [b-WTISD26], with a foreword by the ITU Secretary-General, identifies non-intentional digital disruption as a distinct risk category that current cybersecurity and resilience frameworks address only partially, and surfaces invisible cross-sector substrate dependencies as the foundational structural condition. The 2038-class rollover cluster addressed in the present Technical Paper is the calendar-deadline-specific instance of that broader category: a stress case whose timing is determined not by adversary intent but by the arithmetic of bounded representations meeting deployed-system lifecycles. The framings are complementary; the present document develops the technical and coordination dimensions of the specific instance against the broader category articulated by [b-WTISD26].

The network time protocol (NTP) era rollover in 2036 warrants particular attention as the structurally most consequential rollover event in the cluster. NTP is the time-distribution substrate [b-Mills] on which many systems depend for coherent time across distributed deployments. A system whose host time representation (time_t) is properly implemented in 64-bit may nonetheless experience operational failure if the NTP infrastructure delivering time to it experiences era-rollover problems. NTP implementations vary in how they handle era boundaries; the variations interact in ways that may not be fully visible until 2036 actual exposure produces clear empirical signals.

The 2036 event also functions as a leading indicator: organizations that fail to detect or remediate 2036 NTP-layer problems are unlikely to have the resources or organizational capacity to remediate their 2038 exposure in the narrow window remaining between the two events.

The 2038 time_t overflow remains the most widely recognized event in the cluster and the event from which the term "2038-class rollover" derives. Many computing and communication systems rely on 32-bit signed time formats that will overflow on 19 January 2038, producing undefined or system-dependent behaviour as some system clocks even go negative in value. While modern systems increasingly use 64-bit time, significant installed bases of embedded devices, industrial control systems, transport systems, and legacy software (frequently load-bearing in critical systems) remain vulnerable. Timestamp width expansion in implementation, while necessary, may not be sufficient[1]: protocol-layer encoding mismatches and interoperability between mixed-width implementations introduce additional correctness considerations addressed in clause 8.5.

The GPS week-number rollover expected on 20 November 2038 will affect global navigation satellite system (GNSS)-derived timing across receiver populations whose firmware lifecycles commonly outlast the platforms incorporating them. Prior GPS week-number rollovers in 1999 and 2019 produced documented operational consequences in aviation and other sectors [b-GPS.gov], as referenced in clause 7.1 and clause 10.28. The 2038 rollover will exercise the same failure pathways across an expanded deployed base.

---

[1] Bit-width expansion alone does not determine the rollover horizon; storage resolution is co-determinant. The 64-bit signed integer maximum at seconds-resolution (POSIX time_t convention) gives a horizon of approximately 292 billion years (see footnote 2). At finer resolutions the horizons shorten by three orders of magnitude per step: millisecond-resolution 64-bit storage (as in Java System.currentTimeMillis(), JavaScript Date.now()) gives approximately 292 million years; microsecond-resolution storage (POSIX struct timeval tv_usec) gives approximately 292,277 years; nanosecond-resolution storage (POSIX struct timespec, Linux CLOCK_REALTIME, Go time.Time internal representation) gives approximately 292.5 years. Nanosecond-resolution 64-bit timestamps initialized against a 1970 epoch are therefore exposed to rollover within the operational lifetimes of long-lived infrastructure deployed today. Bit-width is necessary but not sufficient; resolution choice is co-determinant of remediation adequacy.

Four structural properties of the 2036-2038 cluster warrant explicit acknowledgement at the outset:

1) First, the events are deterministic: their occurrence is not a matter of probability but of calendar arithmetic and regulatory deadline. The 2036 and 2038 rollover events occur when their underlying counters or representations reach their representable bounds. Unlike most failure modes addressed in cybersecurity literature, the question is not whether these events will occur, but what their direct and emergent consequences will be, and how failures propagate across the heterogeneous, tightly-coupled ecosystem that is loosely termed "*critical infrastructure*".

2) Second, the consequences are difficult to discover in advance. Inventory of time-handling implementations across deployed systems is not generally available to operators; supply-chain opacity means manufacturers themselves frequently lack visibility into the time-handling behaviour of subcomponents in their own products; in-situ behaviour under rollover conditions cannot be reliably reproduced through laboratory simulation, as developed in clause 10. The deployed base of legacy 32-bit timestamp implementations across embedded systems, industrial control infrastructure, transport platforms, and long-lifecycle software running in critical sectors is large; the proportion remediated by the dates of the cluster events is not currently knowable with confidence. Early empirical evidence, including the GPS week-number rollover failures of 1999 and 2019, documented telecommunication timing incidents, and observed authentication failures from out-of-range timestamps, confirms that the failure modes described in this Technical Paper are not theoretical projections but active patterns in present-day production systems.

3) Third, the institutional memory of prior rollover-class events shapes the environment in which the present cluster will be addressed in ways the engineering substance of the events themselves does not capture. The Y2K transition is widely remembered as a non-event, and the memory is misleading. Y2K's operational uneventfulness was the product of coordinated global remediation work undertaken sufficiently early to be effective by the time the calendar date arrived; the absence of cascading failures at the moment of transition was the consequence of remediation, not evidence that the remediation was unnecessary. The institutional-memory dynamic this produces – *in which successful preparation is remembered as evidence that preparation was unnecessary* – is a structural condition the present Technical Paper must surface, since it shapes the institutional environment in which the 2036–2038 cluster will be addressed and informs the cadence at which substantive preparation can be expected to receive resourcing across the remaining update window summarized in Table 1.

   The same dynamic operates in the opposite direction as the cluster approaches. Where the memory of past remediation produces under-attention, the arrival of the events themselves – and their treatment in general-interest media – may produce the reverse: a rapid surge of public attention arriving in advance of, or in place of, institutional readiness. Both conditions are failures of calibration rather than of engineering, and both shape the environment in which technical remediation must proceed. An uncoordinated surge of attention that reaches decision-makers without prior preparation constitutes a distinct risk in its own right, capable of distorting remediation priorities and compressing the time available for considered response, whatever the eventual technical trajectory of the underlying exposure.

4) Fourth, the deployed base of 32-bit time-handling implementations does not appear to be contracting at a rate commensurate with the approaching event window – and may be expanding. Component-level sales of 32-bit microcontrollers, embedded processors, and adjacent silicon continue at high and apparently accelerating volume; products incorporating these components are entering current production today and will, in many cases, remain in operational deployment well beyond the 2036–2038 cluster events. In parallel, the rapid growth in AI-assisted code generation introduces a distinct structural concern: large language models trained on historical code corpora – which contain decades of 32-bit time_t patterns

embedded across embedded-systems, industrial control system (ICS), and legacy-application codebases – exhibit a statistical tendency to reproduce those patterns in newly-generated code. The combination of continued 32-bit silicon deployment and AI-mediated reproduction of legacy time-handling patterns in new code is not deterministic but is suggestive: the substrate of 32-bit time representations may be expanding through ordinary commercial activity even as the date of its predictable failure approaches.

Timestamp inconsistencies arising from any of these events propagate through distributed authentication systems, billing and settlement infrastructure, network control and data planes, logging and evidence chains, embedded systems with long replacement cycles, and safety-critical systems in maritime, aviation, railway and industrial domains. Because these systems are intertwined across operator, sector, and jurisdictional boundaries, the resulting risks are inherently global. The structural conditions making coordinated remediation difficult – heterogeneous lifecycles, supply-chain opacity, regulatory fragmentation, the lack of cross-sector visibility into time-handling dependencies, and the strained state of the international coordination mechanisms on which cross-jurisdictional response would depend – are examined in subsequent clauses of this Technical Paper.

In November 2022, the 27th General Conference on Weights and Measures adopted [b-CGPMRES4E], deciding that the permitted difference between UT1 and coordinated universal time (UTC) will be increased in or before 2035, *with implementation details to be finalized at the 28th CGPM in 2026.* The resolution was not unanimous; its adoption involved negotiated compromise between parties favouring earlier action and parties whose national time-distribution infrastructure required additional preparation time. Although leap-second reform is regulatory rather than representational and is therefore distinct from the rollover-class failures examined in this Technical Paper, its implementation period substantially overlaps with the 2036–2038 cluster. The same international coordination mechanisms, institutional processes, and operational communities required for leap-second transition planning will also be engaged in rollover-related remediation and interoperability work; the demonstrated diplomatic strain on timekeeping consensus is itself a structural condition shaping the environment in which the cluster will be addressed.

The 2036-2038 cluster examined throughout this Technical Paper is not an isolated set of events but the leading edge of a longer pattern of rollover-class risks clustering at bit-width boundaries in deployed substrates. Subsequent boundary events – including the 2106 unsigned 32-bit rollover and additional rollovers in legacy systems that adopted bespoke epochs to defer the 2038 event without coordinating those deferrals with other affected systems – bear systematic treatment as part of the longer-horizon pattern. The present Technical Paper focuses on the 2036-2038 cluster as the most operationally consequential occurrence in the present revision window; the structural pattern itself, and the implications of treating individual rollover events as instances of that pattern rather than as discrete events, is anticipated in subsequent revisions.

This Technical Paper addresses the structural patterns of 2038-class timestamp rollover risks across ICT systems, with particular focus on the 2036-2038 cluster of timekeeping discontinuities described in this clause: NTP era rollover, time_t overflow, GPS week-number rollover. The paper documents known and emerging failure classes, examines dependencies of telecommunications, networked systems, and cybersecurity processes on coherent time, and surveys cross-sector exposure to the structural conditions identified in clause 8 and clause 10. It also identifies areas requiring further analysis or future inter-SDO collaboration. The 2106 horizon is treated as a long-tail consideration in clause 10.2.3; it is not a primary focus of the present version.

NOTE – This Technical Paper does not define testing requirements, protocol changes, or specific implementation guidance.

6.1 Foundational time concepts

The brief list of terms and definitions provided in clause 3 is sufficient for this baseline version. Full definitions, including formal mathematical treatment where applicable, are deferred to subsequent versions.

6.2 Failure mode taxonomy

The failure modes associated with 2038-class timestamp rollover risks are heterogeneous in their manifestation but can be organized into a small number of recurring classes. This proposed taxonomy is organized by how time-handling assumptions are violated, rather than where such failures manifest. The classes are not mutually exclusive: in practice, compound failures involving multiple classes are common and frequently more difficult to diagnose than failures of any single class in isolation. Table 2 summarises the failure mode taxonomy used in this Technical Paper.

Table 2 – Failure mode taxonomy

Class	Failure mechanism	Concrete example	Remediation difficulty
Representation and range	Fixed-width counters, truncated timestamps, and temporal-reference-point-relative encodings overflow or wrap when their representable range is exceeded.	[b-Alstom]: In 2017, attempting to set a date beyond 2037 on rolling-stock embedded software revealed that approximately 38 trains across more than one-third of the network could not represent post-2037 dates. The Tribunal Administratif de Paris ruled on 13 November 2025 that [b-Alstom] must remediate the defect; remediation is allowed five years following a twelve-month assessment period.	Deterministic but heterogeneous in remediation cost. May require changes across multiple layers when representations are embedded deep within libraries, firmware, or protocol implementations.
Interpretation and semantic mismatch	Time values remain representable but are interpreted inconsistently across components – differences in epoch definition, time zone handling, leap adjustments, or signedness.	NTPv3 [b-IETF RFC 1305]/NTPv4 [b-IETF RFC 5905] wire-semantic ambiguity in widely deployed time-synchronization infrastructure. As surfaced in the [b-FIRST Time SIG] by maintainers from the NTP development community, the NTPD codebase supports both NTPv3 and NTPv4 packet formats for backward compatibility, but version differentiation in the absence of explicit handshake signals can be non-deterministic – meaning that participating systems may interpret the same wire bytes according to different semantic models, with consequences for synchronization accuracy, era-rollover handling, and security-property assumptions downstream of the NTP layer. Certificate validity timestamps and revocation-list semantics are similarly subject to mismatch between components – differences in how notBefore/notAfter boundaries are interpreted across implementations, or in how revocation freshness is judged, can produce inconsistent trust decisions at scale.	Difficult to attribute to a single faulty element. Often requires coordinated changes across interfaces between components that each behave correctly according to local assumptions.

Table 2 – Failure mode taxonomy

Class	Failure mechanism	Concrete example	Remediation difficulty
		Civil clock-shift events such as daylight-saving transitions are structurally similar in mechanism: the same wire bytes represent legitimate but differently handled time values across systems depending on local-time-conversion logic. Failures of this kind surface not only in code but in coordination-layer artefacts – documented Y2K-era incidents include hemispheric daylight-saving-time differences producing multi-hour discrepancies in follow-the-sun support-engineering handoff schedules, with consequences for active-incident coverage continuity.	
Synchronization and coherence	Assumptions about bounded clock skew, monotonic progression, or synchronization convergence are violated through degraded time sources, network partitioning, or interference with synchronization protocols.	2019 GPS week-number rollover incident: more than a dozen aircrafts, were grounded after flight-management GPS receivers rolled back to dates in 1999 following the rollover. Operators grounded affected aircraft pending GPS firmware remediation, per [b-SIB 19-01R1]. The proximate cause was operator non-application of vendor service information letters issued in March 2019 specifically warning of the rollover – an exemplar of the lifecycle and supply-chain exposure conditions developed in clause 8.3.	Cascading effects in distributed systems are common; failures may degrade availability, safety, or security properties without producing explicit error conditions.
Lifecycle and longevity-induced	Components designed with implicit assumptions about operational duration remain in service far beyond their intended horizon.	Industrial control systems with 10-30 year operational lifetimes frequently incorporate 32-bit timestamp logic originally specified for shorter consumer-software lifecycles. Equipment refresh cycles are often decoupled from firmware/software refresh cycles, resulting in legacy binaries persisting on otherwise modern hardware.	Often discovered late, when remediation options are constrained by operational, regulatory, or safety requirements. Replacement may be infeasible before rollover thresholds are reached.

Table 2 – Failure mode taxonomy

Class	Failure mechanism	Concrete example	Remediation difficulty
Testing and observability gaps	Failure modes are difficult to exercise through conventional testing frameworks due to dependence on extended time horizons, coordinated multi-system behaviour, or interactions with external time sources.	[b-Ariane 5]: Code reused from the [b-Ariane 4] inertial reference system behaved reliably under simulation and pre-deployment testing but failed catastrophically in operation due to previously unexercised boundary conditions in the new operational envelope.	Compliance and functional testing pass while latent vulnerabilities persist. Detection often requires in-situ validation under realistic operational conditions.
Interaction and compound failures	Failure modes from multiple classes interact: representation limits compound with synchronization assumptions; semantic mismatches compound with lifecycle constraints; testing gaps obscure the combined effect of multiple weak assumptions.	The clustering of consequential time-handling events between 2036 and 2038 – including the NTP era rollover (2036), the Unix time_t rollover (2038), and the GPS week-number rollover (2038) – creates conditions under which developer touch-points on time-critical code are concentrated, increasing the risk that adjacent latent bugs are surfaced or introduced during remediation efforts. Transport layer security (TLS) authentication is a near-universal cross-class instance: clock-skew conditions (synchronization/coherence) interact with certificate-validity semantics (interpretation/semantic) to produce trust failures that are often misdiagnosed as networking faults, with the underlying timing dependency invisible to operational telemetry. A more severe variant of the same dynamic appears in distributed consensus systems, where clock divergence produces permanent chain forks; this case is developed in clause 7.	Particularly challenging to analyse; manifestations may be delayed, indirect, or context-dependent. Coordinated remediation across classes is required.

Table 2 – Failure mode taxonomy

Class	Failure mechanism	Concrete example	Remediation difficulty
Delayed and operationally-cadenced manifestation	Failures activated not at the rollover threshold but instead surface at the cadence of the operational logic that exercises the affected code path, rather than at the rollover event itself. A population of latent failures becomes simultaneously eligible to manifest at the rollover threshold, but each individual failure surfaces according to its own system's operational tempo: cron jobs that run hours or more later, monthly batch processes, annual reporting cycles, certificate validity checks at next renewal. A particularly insidious sub-category produces *operating-incorrectly-without-failing* – for example, signed certificates whose validity arithmetic overflows in a direction that extends their validity beyond what was intended rather than expiring them, *producing systems that continue to operate but with security properties materially different from those the operators believe to be in effect*.	Y2K analogue: an annual data-summarization job scheduled for 31 December would not surface its rollover-class failure until its next scheduled execution. The Y2K record contains documented cases of this dynamic – failures that did not manifest until weeks, months, or in some cases more than a year after the calendar transition that activated them. The relevant population for 2038-class risk is broader still, given the longer typical horizon of operational scheduling logic and the prevalence of validity-arithmetic dependencies in deployed cryptographic infrastructure.	Particularly difficult to plan for. The incident-response posture appropriate to the rollover event itself may relax before the bulk of consequences have surfaced, producing a discovery profile that extends over months or years following the event rather than concentrating around it. Risk-management frameworks that treat the rollover as a discrete event miscalibrate the response cadence required.

Compound failures merit emphasis. The classes above are analytically separable but operationally entangled. A representation failure in one component may produce a semantic mismatch at an interface that propagates as a synchronization failure to a downstream system that was itself lifecycle-constrained from receiving updates that would have detected the original representation issue. This entanglement is one of the structural reasons why 2038-class risks do not fit conventional vulnerability frameworks and why remediation timelines must accommodate dependencies across organizational and technical boundaries.

A more severe variant of the same dynamic appears in distributed consensus systems, where clock divergence beyond consensus tolerances produces permanent chain forks with no technical mechanism for resolution; this case is developed in clause 7.

This taxonomy is intended to support analysis rather than to enumerate exhaustively. By organising temporal fragility into recurring classes with concrete instances, it provides a framework for comparing disparate incidents, identifying common root causes, and clarifying where existing mitigation efforts address symptoms rather than underlying assumptions.

6.3 Rollover event timeline

The events surveyed in clause 6 are defined relative to specific temporal reference points – epochs – from which the affected representations measure their offsets. Each epoch and the rollover event it produces are co-defined: a rollover is the moment at which a representation, measured from its reference epoch, reaches the upper bound of its representable range. The timeline in Table 3 records both, organized chronologically.

Table 3 – Rollover event timeline

Date	Event	Type
Monday 1 January 1900 00:00:00 UTC	NTP epoch established & two-digit year encoding epoch established	Reference point
Thursday 1 January 1970 00:00:00 UTC	Unix epoch established	Reference point
Sunday 6 January 1980 00:00:00 UTC	GPS epoch established	Reference point
Saturday 21 August 1999	GPS week-number rollover (era 0 → era 1)	Historical reference; documented operational consequences
Saturday 1 January 2000	Y2K calendar transition (two-digit year representation rollover)	Rollover (mitigated)
Saturday 6 April 2019	GPS week-number rollover (era 1 → era 2)	Rollover (mitigated, partial)
Thursday 7 February 2036 06:28:16 UTC	NTP era-1 rollover (32-bit unsigned seconds counter from 1900 epoch)	Rollover (pending)
Tuesday 19 January 2038 03:14:07 UTC	Unix 32-bit signed time_t rollover (from 1970 epoch)	Rollover (pending)
Saturday 20 November 2038 approximately 23:59:42 UTC (assuming GPS-UTC offset remains +18 seconds; offset is subject to further leap-second insertion or omission)	GPS week-number rollover (era 2 → era 3)	Rollover (pending)

Table 3 – Rollover event timeline

Date	Event	Type
Sunday 7 February 2106 06:28:16 UTC	Unix 32-bit unsigned rollover (from 1970 epoch)	Rollover (long-tail)
Approximately 292 billion years after the Unix epoch[2]	Unix 64-bit signed time_t rollover (from 1970 epoch)	Rollover (long-tail; cosmological timescale)

The reference points themselves remain unchanged across rollover events: a rollover refers to range exhaustion of a representation, not to revision of the underlying epoch from which the representation measures. The post-remediation 64-bit row is included for completeness rather than for coordination purposes – the 64-bit signed time_t horizon exceeds the projected remaining lifetime of the Sun's main sequence by a substantial margin, and is in this sense not a deferred rollover but a structural removal of the rollover from any timescale on which the deployed substrate could plausibly persist. Subsequent revisions of this Technical Paper will extend this chronology with documented incidents demonstrating the failure modes developed in clause 6.2 and additional cross-references to the sectoral treatments developed in clause 10.

6.4 Risk windows spanning canonical dates

As developed in clause 8 and clause 7.2, rollover-class risks span overlapping temporal windows; failures may surface before or after canonical dates depending on validation logic, certificate lifetimes, counter rollovers, and assumptions of monotonic time. Earlier events such as the 2036 NTP era rollover function as indicators of latent vulnerabilities in deployed substrate.

Additionally, for software performing delta calculations between two clock measurements, a problem may surface only when the two measurements bound the rollover date and will self-correct once the latent measurement surpasses the rollover date.

7 Problem statement

Time differs structurally from other resources on which distributed computing systems depend. Memory, storage, and computational capacity are bounded resources whose consumption is local to a system or process; failures in their management produce localized consequences and admit localized remediation.

Time, by contrast, is a globally shared resource: each participating system in a distributed environment depends on a coherent and comparable representation of time across all other participating systems with which it must coordinate.

Time is also mutable in ways memory and storage are not – sources of time can be substituted, manipulated, lost, or made incoherent through the failure of intermediate distribution mechanisms, with the consuming system frequently unable to detect the divergence.

---

[2] No day of the week is given for the 64-bit signed time_t rollover, and not for want of arithmetic. The arithmetic is straightforward; what is not straightforward is the civil-interpretive scaffolding required to project a day-of-the-week label across approximately 292 billion years of calendar drift, leap-year conventions, hypothetical future calendar reforms, and the relativistic effects of changes in Earth's rotation, orbital dynamics, and the SI second definition over geological timescales. Each of these assumptions is doing substantial work, and none of them is stable over the horizon involved. The phenomenon is not unique to 64-bit time_t: civil-time meaning-attachment decays at the rate of institutional-memory persistence, which is shorter than the horizon of physical clock-tick precision by many orders of magnitude. The observation is structurally adjacent to the conditions developed in clause 10.15 and clause 8.3 and is offered here as a reader-orientation note rather than as analytical content.

The combination of global reach and mutability produces a structural condition in which time-related failures propagate across system, organizational, and jurisdictional boundaries faster than coordination mechanisms typically operate. This condition is foundational to the failure modes examined throughout this Technical Paper.

Conventional software vulnerabilities admit a response posture in which discovery, patching, and deployment are operationally distinct phases performed by identifiable parties along established remediation pathways; the response is bounded by the affected component and proceeds at the cadence of ordinary security operations. 2038-class events differ structurally: the conditions surveyed in clause 8 and clause 10 indicate that effective response depends on coordinated action across operators, vendors, regulators, certification bodies, and standards organizations whose ordinary remediation pathways are not constituted to address risks that span their jurisdictional boundaries. The distinction is consequential for how preparation work is structured: response to a bug-class risk can be aggregated post-hoc through vendor-by-vendor patch cycles, *while response to a vulnerability-class risk of the present scope depends on coordination mechanisms operating in parallel with – and in advance of – the ordinary remediation pathways that any single participant would otherwise rely on.*

2038-class remediation in 32-bit systems is not, in the general case, achievable through application-level patching: it requires modification of the underlying time-handling interface between the operating system, the standard library, and the userland binaries compiled against them. *This interface – conventionally described as the application binary interface – is a category of system artifact that ordinary security patching is not constituted to modify.* Standard patch deployment preserves the interface; remediation of 2038-class exposure requires changing it.

The practical consequences are operationally consequential: shipping a newer application binary on top of an unmodified 32-bit interface does not remove the underlying exposure; vendor-applied surface remediations may produce systems whose test behaviour appears clean while the underlying interface remains exposed at rollover; and the deployed-substrate observability constraints developed in clause 8.4 apply with distinctive force, since remote indicators of remediation status correlate imperfectly with the interface-level remediation that would actually close the exposure. *These conditions differentiate 2038-class remediation from ordinary security response in ways meriting explicit acknowledgement when an institutional response posture is being designed.*

A significant portion of the world's digital and telecommunication infrastructure continues to rely on time representations, libraries, userland, protocols, and embedded systems that will be affected by the 2038-class timestamp rollover and related time-handling limits.

These risks are not hypothetical: failures have already been observed in production systems when timestamps exceed expected bounds or when time desynchronization occurs due to protocol errors, system drift, or GNSS anomalies. Citable historical incidents are surveyed in clause 7.1.

Historical precedent for Internet-scale protocol transitions – most notably the prolonged IPv4-to-IPv6 coexistence – suggests that assumptions of clean migration are rarely borne out at deployment scale. Engineering for prolonged coexistence of divergent time semantics, rather than a discrete transition event ("flag day"), may be the more realistic design posture.

The implication is that 2038-class remediation strategies that assume coordinated cutover across deployed systems are likely to encounter the same coexistence dynamics that have characterized every prior large-scale Internet protocol transition.

While the 2038 rollover is receiving increasing attention, earlier rollover events – notably the NTP era wrap scheduled for 7 February 2036 – will serve as a practical early warning. These earlier transitions are structurally similar to the 2038 overflow (counter or epoch boundary crossing) but with a smaller installed base exposed.

Observed failures in 2036 may provide empirical signals of latent vulnerabilities, particularly in environments where 32-bit time representations, legacy protocol glue, or embedded components persist unnoticed. The narrow window between 2036 and 2038 reinforces the need for proactive coordination: systems that fail silently in 2036 may cascade through dependent infrastructures before 2038, when operational impact is likely to be far broader.

Those organizations that hear the 2036 warning but have not remediated their 2038 exposure by that point are unlikely to have the resources available to remediate in the narrow window remaining.

The consequences of 2038-class failures are systemic and cross-cutting. Disruption may affect critical services across energy, telecommunications, financial, and transportation sectors. Trustworthiness of logs and evidence used in national-level investigations may be reduced. Continuity of public-safety systems requiring secure, sequential timestamps may be lost.

Coordination between computer emergency response teams (CERTs), computer security incident response teams (CSIRTs), and product security incident response teams (PSIRTs) responding to incidents triggered by time anomalies may be impaired. Interoperability between jurisdictions with differing patch levels or legacy device exposure may fail.

Operational risk may accumulate in embedded systems that cannot be patched or replaced before 2038.

These failure modes are unlikely to manifest as a single coordinated event; rather, 2038-class failures may occur gradually over the coming decade as systems reach rollover thresholds at different times.

This phased manifestation increases the difficulty of coordinated mitigation and amplifies the need for shared international understanding of the underlying structural conditions.

Selected failure categories observed in production systems are surveyed in clause 7.1; the underlying temporal-trust dependencies they share are developed in clause 7.2.

7.1 Observed failure modes in real systems

Failures observed in production systems to date span multiple sectors and failure-mode classes. The categories below are non-exhaustive:

– **Authentication failures**: Certificate validation errors, replay-protection malfunctions, and token-expiry miscalculations triggered by out-of-range timestamps.

– **Logging and audit failures**: Non-sequential logs, negative timestamps, and gaps in evidence chains that complicate national incident response.

– **Industrial and embedded system faults**: Legacy programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) endpoints, maritime, railway, and aviation systems exhibiting unexpected resets, shutdowns, or undefined behaviour when internal clocks drift or overflow.

– **Telecommunication anomalies**: Billing system faults, lawful intercept issues, session timeout errors, and synchronization drift affecting routing and service continuity.

– **GNSS and time-distribution incidents**: Real-world GNSS week-number rollovers have already caused outages in aviation, maritime navigation, and emergency services.

– **Financial infrastructure**: Distributed consensus systems represent the most severe instance of this coupling. Blockchain networks and smart-contract platforms depend on bounded clock agreement among globally distributed nodes for chain integrity; clock divergence

beyond consensus tolerances produces chain forks with no technical mechanism for resolution, and smart contracts deployed with implicit assumptions about how time will be reported across decades inherit the lifecycle exposure described in clause 6.2 without the operational visibility that traditional financial infrastructure retains.

– **Software library rollover defects**: Widely used libraries (including NTP, C standard libraries [b-GNU C Lib], and OS date-handling code) have demonstrated rollover vulnerabilities when tested against future dates.

Without a common vocabulary, shared risk framing, and coordinated preparation, Member States may face cascading failures in systems that rely on synchronized time to maintain operational safety and security.

7.2 Temporal root of trust dependencies

The security properties of distributed systems rest on a shared assumption – usually unstated, frequently invisible to operators – that participating components observe a coherent time. The term *temporal root of trust* (defined in clause 3.2.8) names this assumption explicitly so that its dependencies can be examined as a class rather than treated as accidents of individual subsystems. Time differs structurally from other resources distributed systems depend on, as developed in clause 7. It is globally shared rather than locally bounded, and mutable in ways memory and storage are not. Security properties that rest on it inherit both characteristics.

The dependency surface is broad and operationally entangled. Transport layer security (TLS) depends on temporal trust at multiple points – establishing certificate validity windows, evaluating revocation freshness, accepting or rejecting cryptographic handshakes – and is consequently the near-universal substrate against which temporal-trust degradation surfaces in practice. Clause 6.2 compound-failures paragraph treats TLS as the canonical cross-class instance; here it serves as the canonical instance of how temporal trust manifests at the cryptographic-protocol layer.

Authentication exchanges more broadly depend on temporal trust to bound the validity of credentials, the freshness of nonces, and the acceptance windows for cryptographic challenges. Certificate-based trust depends on temporal trust to establish when a certificate becomes valid, when it expires, and whether a revocation should be honoured – including in public key infrastructure (PKI) deployments whose validity windows assume sustained access to coherent time across the period of trust. Logging and audit systems depend on temporal trust to establish the ordering of events across distributed sources [b-Lamport-1]; without it, forensic reconstruction is unreliable and evidence-chain integrity is compromised, developed in clause 8.8. Distributed consensus protocols depend on temporal trust to bound clock skew across participating nodes within tolerances their consensus algorithms assume – a dependency developed in clause 7.1, under *financial infrastructure*, for the case of blockchain and smart-contract platforms. Replay protection depends on temporal trust to distinguish a current message from a captured-and-resubmitted copy.

The structural observation is that these are not separate dependencies on time. They are a single dependency – the temporal root of trust – manifesting across different security properties. Degradation of the underlying assumption degrades all of them together, frequently without producing failure modes the affected systems are designed to surface. A clock-skew condition that would individually produce a TLS handshake failure, a logging-sequence ambiguity, a Kerberos ticket rejection, and a consensus tolerance violation is not four failures with a common symptom; *it is one failure with four symptoms*. Diagnosis is consequently difficult: each subsystem reports a different fault, none of which names the underlying temporal condition, and operational telemetry frequently does not surface the time-coordination state on which the subsystems depend.

The temporal root of trust is one instance of a broader category – *substrate-trust commitments* – that distributed systems make to substrates whose correct operation their operators cannot directly verify. Compiler integrity, supply-chain provenance, hardware-implementation correctness, and time coordination share the structural property that the assertion they support flows downstream into every

higher-layer security property the system relies on, while the verification of the assertion itself sits at a layer below conventional analytical instrumentation. The 2038-class events surveyed in this Technical Paper exercise the time-coordination instance of this category at operational scale; the broader category itself warrants explicit acknowledgement in the analytical framing of subsequent revisions.

Machine-to-machine (M2M) operational contexts amplify the consequences. Where human operators interact with authentication or certificate-validation systems, a temporal-trust failure may produce a dialog the operator can override, escalate, or defer until the underlying condition is diagnosed. Autonomous systems – embedded controllers, industrial automation, vehicle-to-infrastructure communication, autonomous trading, machine-to-machine settlement, and the broad class of operational technology described in clause 10.8 have no equivalent fallback. A temporal-trust failure in an autonomous context produces a binary outcome: the system either continues operating on degraded trust assumptions, with consequences proportional to what those assumptions were underwriting, or it stops. There is no intermediate state in which a human is consulted, and there is no operational pathway for graceful degradation that does not require the system to have been designed to recognize the condition in advance. The autonomous-system case is the worst-case manifestation of the broader temporal-trust dependency: failure modes cannot be deferred to human judgment, and remediation pathways cannot be invoked from outside the failed system.

A further dimension bears explicit mention. The temporal root of trust on which the foregoing dependencies rest is itself a shared substrate, and shared substrates with this combination of dependency reach and operational invisibility have historically attracted sustained adversarial attention. Public reporting in October 2025 surfaced nation-state-level accusations concerning multi-year operations targeting national timing infrastructure; the accused party did not publicly respond to the substance of the accusations.[3] The structural inference – that nation-state-level adversarial interest in timing infrastructure is established, sustained, and likely multilateral – follows from the public record without requiring attribution to any specific actor.

The operational consequence is that the temporal substrate may already be subject to persistent access established under prior threat-model assumptions, in ways not surfaceable through detection methodologies that postdate the establishment of that access. The mapping problem is asymmetric: defenders attempting to assess the integrity of their temporal substrate face a verification task whose evidentiary baseline may itself have been compromised before the verification effort began. This condition is not unique to the temporal substrate among substrate layers but is particularly consequential where the substrate's correct operation is presupposed by every higher-layer security property the system relies on. The 2038-class events surveyed in this Technical Paper interact with this condition: each event functions as a forcing function that may surface latent positioning that would otherwise remain dormant, while simultaneously stressing the legitimate operational coordination required to maintain the temporal substrate's integrity through the cluster.

2038-class events surface this class of dependency directly. Each event in the cluster as set forth in clause 6 is a deterministic, calendar-driven stress test of the temporal root of trust across the deployed substrate. Where the temporal root of trust holds – where systems share a coherent time across the event boundary, and where the substrate's integrity has not been previously compromised – dependent security properties continue to function as designed. Where either condition fails, the dependent properties degrade in correlated patterns whose technical disentanglement is, as observed throughout this Technical Paper, structurally difficult.

---

[3] See contemporaneous public reporting [b-Reuters] on accusations relating to operations against national timing-service infrastructure. The present Technical Paper cites the case as an instance of nation-state-level interest in timing infrastructure entering the public record; specific attribution and the merits of the underlying accusations are outside the scope of this document.

8 Analysis framework

This Technical Paper applies an interdisciplinary analysis framework that maps the 2038-class timestamp rollover risks to established cybersecurity principles. The goal is to provide Member States and sector stakeholders with a structured method for assessing, comparing, and prioritizing vulnerabilities related to time representation failures across heterogeneous systems. The analysis framework consists of the following dimensions.

8.1 Integrity of time-dependent operations

Critical system functions across deployed infrastructure routinely assume forward-only monotonic time. The 2038-class rollover events described in clause 6 violate this assumption directly: at the moment of rollover, subsequent to it (and in some cases even prior to it), processes tied to timestamp arithmetic – scheduled tasks, cryptographic key rotations, log sequencing, certificate validation, retry-and-backoff loops – may terminate unexpectedly, enter undefined behaviour, or loop indefinitely depending on how the underlying implementation handles values that exceed representable bounds. The failure modes are not uniform across affected systems: empirical observation indicates that some failures present as silent state corruption without any operator-visible alert, while others cascade through downstream dependent processes whose behaviour becomes unpredictable in ways that are difficult to attribute to the originating temporal-trust failure.

Four categories of operational risk follow directly from these conditions.

1) *Authentication and authorization infrastructure* fails when credentials, tokens, and certificate-validation chains encounter timestamps outside their expected ranges – producing expired-credential rejections for valid users, accepted credentials for expired sessions, or non-deterministic mixtures of both depending on the path through the validation logic.
2) *Logging and forensic capability* degrades when timestamp ordering can no longer be relied on to reconstruct causal sequences of events, with consequences for both operational incident response and legal-evidentiary use of audit records, as treated further in clause 8.8.
3) *Process scheduling* produces anomalous behaviour when task-execution timing depends on time arithmetic that no longer behaves monotonically – tasks may execute out of order, repeat unexpectedly, or deadlock on resource acquisition pathways that assume forward progress.
4) *Distributed-systems consensus mechanisms* – including those underpinning database replication (for an early timestamp-based approach, see [b-IETF RFC 677]), container orchestration, blockchain validation, and the broader landscape of coordination protocols that depend on agreement about "current" time – may produce conflicting views of system state across nodes, with downstream consequences for any application relying on the consensus layer for correctness.

Across the analytical dimensions developed in this clause and the cross-sector exposure survey developed in clause 10, the present Technical Paper distinguishes four levels of evidentiary status for the conditions it surfaces: *observed conditions*, for which empirical evidence is established and citable; *reported conditions*, surfaced through field observations from the [b-FIRST Time SIG] community, industry incident reports, and operator engagements but not yet systematically verified; *suspected conditions*, theoretically plausible on the basis of the structural analysis developed in this Technical Paper but not yet examined empirically; and *unexamined conditions*, representing evidence gaps where the analytical question has not been posed in a form that empirical investigation could address.

The distribution of conditions across these four levels is itself part of the analytical contribution this Technical Paper makes: more conditions presently sit in the suspected and unexamined categories than in the observed category, and the structural conditions developed across clause 8 and clause 10 indicate that closing this distribution – through cross-sector empirical work undertaken across the

remaining update window summarized in Table 1 – is among the most operationally consequential research priorities the present revision can surface.

NOTE – The evidentiary-status vocabulary defined here (observed, reported, suspected, unexamined) is applied selectively in the present revision; consistent labelling across the developed sector treatments in §10 is scheduled for a subsequent revision.

8.2 Resilience and correlated failure pathways

Timestamp rollover risks differ structurally from most failure modes that conventional resilience frameworks are designed to address. Conventional resilience engineering assumes that failures arrive as isolated defects – a single component failing for a single reason at a single moment, with the rest of the system available to absorb the failure through redundancy, failover, or graceful degradation. 2038-class rollover events do not behave this way. They arrive as a simultaneous trigger condition affecting every system that incorporates the relevant time representation, across organizational and jurisdictional boundaries, at the same instant. The result is the structural possibility of globally synchronized failure events – a category of risk that disaster recovery and incident coordination frameworks designed against isolated-defect assumptions are not generally prepared to absorb.

Several structural properties compound this exposure. Legacy 32-bit timestamp handling can persist silently inside modern 64-bit infrastructure through compatibility layers, statically-linked libraries, protocol glue code, and the broader long tail of code paths that remain present even after nominal modernisation; the apparent modernity of a host environment does not by itself attest to the time-handling behaviour of every code path executing within it, *and the assumption that it does produces a false sense of safety that is particularly difficult to detect by remote observation*, as treated further in clause 8.4. Time-dependent crash behaviour in embedded, industrial control, and machine-to-machine contexts propagates non-linearly across dependent nodes: a single subsystem entering an undefined state at rollover can perturb control loops, monitoring pipelines, logging infrastructure, and process-coordination layers in ways whose downstream effects depend on the specific topology of dependencies and are therefore difficult to characterize in advance.

Risk is amplified where temporal misalignment combines with limited observability – a structural condition particularly characteristic of embedded and operational-technology environments where 32-bit application state may not be visible to remote monitoring infrastructure even where the host environment is. The failure may have occurred without producing any externally observable signal, with consequences propagating through the system before the failure itself is detected.

The correlated character of 2038-class failure has a further consequence for the fault-tolerance machinery on which much of the deployed substrate relies. The distributed-systems tradition descending from [b-Lamport-2] achieves tolerance of arbitrary component failure through replication and agreement, but it does so under two structural assumptions: that component failures are independent of one another, and that faulty components constitute a bounded minority of participants – in the canonical formulations, fewer than one third. 2038-class rollover inverts both assumptions simultaneously. Affected components do not fail independently; they fail identically, at the same instant, for the same arithmetic reason. And nothing bounds the faulty set to a minority: where a shared time representation is embedded across a replica set, the rollover produces a correlated majority – or a correlated entirety – of participants emitting the same incorrect values in mutual agreement.

Under these conditions, agreement machinery does not contain the failure; it confers authority upon it. A quorum of replicas that share an unremediated time representation will traverse the rollover boundary in flawless consensus on an incorrect time, and the consensus itself – ordinarily evidence of system health – becomes the vehicle by which the incorrect value acquires the status of validated output [b-Thomas]. Replication, similarly, provides no protection where the replicas share the failure's cause: three time sources that incorporate the same bounded representation are, with respect

to 2038-class exposure, a single time source. Redundancy mitigates independent failure; against deterministic correlated failure it merely reproduces the defect with greater confidence.

The practical implication is that this exposure is a property of the time representation, not of the supply of time: multiplying sources does not touch it where those sources share a bounded representation, and agreement or majority-voting offers little detection or containment value where the participants share the failing representation. What distinguishes this failure class from the arbitrary-failure model it otherwise resembles is its determinism – the same property that, as developed in clause 8.4, makes it tractable.

Recovery from globally co-occurring failures introduces complexity that recovery from locally isolated failures does not. Disaster recovery procedures and incident coordination frameworks are typically structured around the assumption that the resources required to recover one failed system – replacement components, vendor support, regulator engagement, mutual-aid arrangements – remain available because they have not themselves failed. Where the failure mode produces synchronized exposure across the resources that recovery would draw upon, the recovery model itself becomes load bearing in ways it was not designed to be. The same vendor's support engineers may be simultaneously called by every customer; the same mutual-aid network may face requests from every participant at once; the same regulator may receive notifications from every regulated entity in parallel. This is a category of correlated-failure condition that resilience modelling has historically addressed in narrow contexts (pandemic planning, large-scale power outages, certain financial-crisis scenarios) but that 2038-class events introduce into a much broader cross-sector landscape.

These observations suggest that preparedness work for 2038-class events would benefit from correlated-failure modelling in test regimes – including simulation of synchronized rollover under load, distributed-dependency stress testing, and exercise of recovery procedures under conditions in which the recovery resources are themselves understood as potentially exposed. Detailed methodology for such testing is anticipated in subsequent revisions of this Technical Paper and is among the items for which inter-SDO and cross-sector coordination would substantially advance the state of practice.

8.3 Lifecycle and supply-chain exposure

The structural exposure described in the preceding clauses is shaped substantially by the lifecycle and supply-chain conditions under which the affected systems were produced and continue to operate. 2038-class vulnerabilities frequently originate upstream of the systems in which they manifest – embedded in inherited libraries, legacy SDKs, toolchains, or hardware modules authored well before 2038-class risk was widely recognized by their authors as a concern. The lifecycle mismatch is direct: devices and industrial systems designed for operational lifetimes of ten to thirty years, and in some sectors substantially longer, routinely incorporate 32-bit timestamp logic originally specified for consumer-software lifecycles measured in single-digit years. The component was correct for the lifecycle it was specified against; the system has outlasted that lifecycle.

A second category of mismatch arises between equipment refresh cycles and firmware or software refresh cycles, which are typically governed by different operational, contractual, and budgetary structures and which therefore decouple in ways that produce systems where legacy binaries persist on otherwise modern hardware. The canonical example is a 64-bit industrial controller running 32-bit supervisory applications, where the host environment has been refreshed but the application layer has not. The host environment's modernity does not propagate into the application layer's time-handling behaviour, and the disjunction is rarely visible to operators who do not have direct introspection into the application substrate. Within the software stack itself, remediation undertaken at the kernel layer does not automatically propagate to userland binaries, system libraries, and build toolchains. These operate under different update cadences, different ownership structures, and different remediation incentives. *A kernel patched against 2038-class exposure may continue to host userland applications and toolchain-compiled binaries that carry the original exposure unmodified, and the*

operator-visible signals of kernel-layer remediation can mask the absence of remediation at the layers above it.

Responsibility for remediation of components incorporating 2038-class exposure is frequently distributed across multiple parties whose contractual relationship to the deployed system is fragmented: the original equipment manufacturer, the systems integrator who incorporated the component into a larger product, the operator who deployed the product, and in some sectors the maintenance contractor who has assumed long-term operational responsibility. Each of these parties may have visibility into a different layer of the system and may have differing economic incentives to investigate, characterize, or remediate 2038-class exposure in components they did not author and may not fully understand. The result is structural ambiguity about risk ownership: in many deployments, no single party has both the technical capacity and the contractual standing to drive remediation through to completion.

Procurement and component-revision dynamics compound this exposure. A single firmware module, chipset revision, or supplier substitution in the ongoing production of a long-lived product can silently reintroduce 2038-class exposure into deployments that operators believed to be remediated, where downstream verification of time-handling behaviour across supplier substitutions is not part of standard procurement practice. Tracing the timestamp-handling behaviour of a deployed component through a multi-tier supply chain – custom system-on-chip implementations, board-support packages, vendor libraries, third-party protocol-glue code – is operationally difficult in the general case and is frequently impossible where documentation gaps exist or where the relevant authoring teams have dispersed.

Architectural remediation introduces its own constraints. Where 32-bit timestamp logic is embedded at the level of application binary interfaces, library structures, or wire-format encodings shared across many dependent functions, replacement may require not isolated patching but coordinated redesign across thousands of dependent code paths. Such redesigns are operationally expensive, frequently break backwards compatibility in ways that have downstream consequences for other deployed systems, and are accordingly difficult to justify on routine maintenance budgets even where the technical case for remediation is clear.

A structural failure mode that warrants direct naming arises when the cumulative effect of these lifecycle and supply-chain conditions exceeds the capacity-to-rebuild as explained in clause 3.2.2 that remediation would require. The deployed substrate continues to function while the institutional and human capacity to maintain, modify, or reconstruct it atrophies – a condition with close analogues in evolutionary biology, where the term vestigialisation describes structures that persist after their underlying function has decayed, and in the systems literature where Tainter's complexity collapse describes the broader societal pattern. Applied to the present case, infrastructural vestigialisation as described in clause 3.2.6 names the condition in which deployed time-handling infrastructure remains operational while the engineering tradition required to remediate it across the rollover cluster ceases to be available at the scale the cluster would demand. The condition is not hypothetical: the supplier-team-lifecycle observations developed above, the dispersed-authoring-team conditions characteristic of the automotive sector developed in clause 10.4, and the structural production of knowledge loss described in clause 10.15 together describe a substrate that is, in load-bearing portions, already operating in this mode. The 2038-class events do not produce vestigialisation; they expose its present extent.

A category of structural mechanism that may merit consideration in some procurement contexts is the use of independent third-party escrow arrangements for source code and build-system artefacts of systems incorporating digital elements, particularly where such systems are acquired with public funds, are deployed in long-lived critical-infrastructure roles, or carry safety-of-life implications. Source-and-build escrow does not by itself remediate 2038-class exposure in a deployed component, nor does it remove the original vendor's role as the natural party to undertake remediation under most conditions. What it does is preserve the structural possibility of remediation by parties other than the

original vendor in conditions under which the original vendor is no longer in a position to remediate – including vendor dissolution, product discontinuation, loss of institutional memory about a legacy product, or absence of economic incentive to address exposure in a product approaching end-of-life. In such conditions, escrowed source and build artefacts make it operationally feasible for successor vendors, regulators acting through delegated authority, infrastructure operators with sufficient technical capacity, or coordinated multi-party arrangements to undertake remediation that would otherwise be technically impossible. The arrangement interacts substantively with the legal and contractual frameworks surrounding modern products incorporating digital elements, addressed in clause 8.6: where intellectual property protections and anti-tamper provisions would otherwise close remediation pathways for parties other than the original vendor, escrow arrangements can operate as a structural counterweight by establishing legitimate alternative pathways at the moment of acquisition rather than negotiating them after deployment. Detailed treatment of escrow mechanisms, their interaction with procurement frameworks across jurisdictions, and the technical questions of what specifically must be escrowed to preserve remediation capacity (which can include build environments, toolchain dependencies, signing keys, and test infrastructure beyond source code alone) is anticipated in subsequent revisions of this Technical Paper.

Supply-chain assurance frameworks intended to support visibility into deployed components – software bills of materials (SBOMs), vulnerability exploitability exchange (VEX) documents, and the broader category of attestation mechanisms emerging across multiple jurisdictions – do not in their current form provide standardized mechanisms for declaring or querying 2038-class exposure. The absence is not an oversight in any individual framework but reflects the cross-cutting and time-deferred character of 2038-class risk, which does not map cleanly onto the vulnerability-disclosure schemas around which SBOM and VEX practices have been developed. Regulatory and certification frameworks emerging across multiple jurisdictions touch on supply-chain transparency in ways that may interact substantively with 2038-class remediation pathways; comparative analysis of those interactions is anticipated in subsequent revisions of this Technical Paper, as flagged in clause 1.

Considered in combination, these lifecycle and supply-chain conditions suggest that timestamp-rollover evaluation has structural value as an explicit element of procurement criteria, lifecycle governance, and long-tail support agreements. The development of formal declaration mechanisms – including SBOM extensions, VEX-equivalent constructs, and supplier-attestation language adapted to 2038-class exposure – would substantially advance the state of practice and is among the items for which inter-SDO coordination would be productive.

8.4 Observability and detectability limitations

The structural exposure described in the preceding clauses is compounded by the limited tractability of detecting and characterizing 2038-class exposure in deployed systems by means available to parties other than the original vendor. Remote scanning techniques – including the inspection of protocol version strings carried in Secure Shell (SSH) banners, NTP version negotiation, SNMP responses, and the broader category of network-observable metadata – can support probabilistic triage of risk profiles across populations of systems, *but they cannot confirm architectural safety of any individual system*. The distinction matters operationally: remote-scanning techniques are useful for prioritising which systems warrant deeper investigation, but they cannot stand as evidence that a system is not exposed.

The false-positive and false-negative profile of remote scanning is asymmetric in ways that matter. False positives – systems flagged as potentially exposed that turn out, on closer inspection, to be safe – produce wasted investigation effort and are operationally tolerable. False negatives – systems reported as likely safe that in fact retain 32-bit time-handling logic in load-bearing code paths – are operationally dangerous because they produce false confidence that no further investigation is warranted. A host environment reporting a 64-bit operating system may nonetheless run critical 32-bit applications, legacy containers, kernel drivers, or statically linked libraries whose time-handling

behaviour will fail at rollover; remote observation lacks visibility into runtime process architecture and application code paths sufficient to rule out these conditions.

Several structural conditions limit what can be inferred from remote observation. Architecture cannot reliably be inferred from protocol metadata, because the metadata-emitting layer is not the layer at which time-handling behaviour is determined. Package-level patching, where it has been applied, does not in general resolve failures arising from 32-bit timestamp logic embedded in firmware or compiled binaries below the level at which package management operates; backported security patches cannot reach into binary code paths whose application binary interface (ABI) assumptions the patches do not change. Large-scale discovery efforts are further hindered by the opaque supply chains and undocumented components characteristic of operational-technology, embedded-systems, and industrial-controller environments, where proprietary firmware frequently provides limited introspection capability even to operators with privileged access to the host system.

Laboratory simulation, which is the natural alternative when remote observation is insufficient, carries its own limits. Many systems and system-interactions only fail under operational workload conditions that laboratory environments cannot fully replicate, where the failure mode depends on specific sequences of inputs, specific timing of concurrent operations, or specific interactions with adjacent systems whose behaviour the laboratory cannot reproduce. The [b-Ariane 5] incident, in which a software fault that had been laboratory-tested under the conditions appropriate to the predecessor system propagated into operational failure under the different conditions of the new vehicle, is a frequently-cited illustration of how environment mismatch can evade detection during pre-deployment testing. The implication for 2038-class exposure is direct: pre-deployment laboratory testing of timestamp-rollover behaviour, however careful, cannot in the general case substitute for in-situ behaviour under operational workload.

A further category of observability constraint arises from the interaction between time-handling failure modes and time-based trust mechanisms. Audit logs, performance metrics, software-update verification chains, and certificate-validation infrastructure all depend on timestamp integrity for their own operation; where timestamp integrity has degraded, the observability infrastructure that would normally detect the degradation may itself produce false stability signals – appearing to confirm correct operation through mechanisms that themselves depend on the integrity of the same time substrate whose failure they are supposed to detect. Time-based trust mechanisms can mask failures until thresholds are crossed at which the masking becomes externally visible – at which point the failure may have propagated well beyond its origin.

This masking effect can be characterized more precisely in terms of the failure model against which the affected systems were designed. Conventional resilience engineering achieves tractable, affordable fault tolerance by restricting its assumptions about how components fail – most commonly through a fail-stop assumption, under which a failed component ceases to respond but does not respond incorrectly. In [b-Lamport-2], the formulation of what the distributed-systems literature terms the *Byzantine Generals problem* (the problem of achieving correct coordination when components may fail in arbitrary rather than merely silent ways) observed that "it is often assumed that a computer may fail to respond but will never respond incorrectly," and that this assumption is precisely what permits less expensive fault-tolerance designs. Where this assumption cannot be made, the full cost of tolerating arbitrary failure is incurred. 2038-class rollover violates the fail-stop assumption directly: a system that has crossed a rollover boundary does not generally fall silent but continues to operate, emitting confidently incorrect timestamps. Affected systems are thereby relocated from the fail-stop failure class, for which their surrounding resilience and monitoring infrastructure was in many cases designed, into a failure class that infrastructure was not constituted to detect. Observability mechanisms that assume a failed source will stop responding faithfully propagate the incorrect values of a source that responds wrongly.

This relocation is, however, partial in a manner that is consequential for preparedness. Unlike arbitrary failure in the general case, in which a faulty component may behave in any manner

(including adversarially coordinated deception) rollover-induced incorrect responses are deterministic. Every affected system fails at a known date and in a known manner determined by the arithmetic of its time representation. This determinism is simultaneously the source of the danger as well as the basis for tractable response. It is the source of the danger because deterministic incorrect output defeats fail-stop monitoring as surely as arbitrary output would. It is the basis for tractable response because the form, timing, and propagation pathways of the incorrect output are knowable in advance, in a way that genuinely arbitrary failure is not – which is what makes systematic preparation across the remaining update window possible at all. The adversarial dimension developed in clause 8.5 marks the boundary at which this determinism can be deliberately removed: an attacker capable of inducing timestamp manipulation converts a predictable failure into an arbitrary one, reintroducing the full cost that the deterministic character of unmodified rollover would otherwise spare.

These observations suggest that effective discovery of 2038-class exposure across deployed populations of systems requires a multi-layer approach rather than reliance on any single observation technique. Remote heuristics support probabilistic risk classification at population scale; the highest-priority subset identified through such classification warrants targeted in-situ testing; in-situ testing must be conducted under conditions including forward-time-shift simulation under operational workload rather than under laboratory conditions; and the relationship between observability infrastructure and the temporal substrate it depends upon must be characterized explicitly when interpreting results from any layer of the discovery pipeline.

The complexity of discovery scales with the integration and interdependency of the systems being characterized, not with the count of individual devices. The interconnected resilience assumptions across modern infrastructure mean that a single undetected failure can propagate across networks with consequences disproportionate to the count of affected systems. The development of standardized methodologies for progressive risk identification – including the triage thresholds, priority criteria, and minimum in-situ simulation requirements appropriate to different sector contexts – would substantially advance the state of practice and is among the items for which inter-SDO and cross-sector coordination would be productive.

8.5 Protocol interoperability dependencies

Timestamp handling in deployed infrastructure routinely spans multiple protocol layers – application, transport, meta-protocol, and embedded firmware – with the consequence that 2038-class exposure in any single layer can cascade across the others in ways that are difficult to isolate by inspection of any one layer alone. A 32-bit timestamp field in a network-management protocol such as SNMP can propagate upward into observability and service-level-agreement infrastructure that depends on the protocol's outputs, or downward into the operational-logic layer through alarm and reporting interpretation; the exposure manifests not where the 32-bit field is encoded but wherever consuming systems assume the encoded value behaves consistently across the rollover boundary.

Several structural conditions amplify this layered exposure. Time fields encapsulated within binary protocol formats, vendor software-development-kit wrappers, or proprietary protocol bindings are frequently not inspected during conventional integration testing, with the consequence that rollover-boundary behaviour in such fields is non-deterministic in environments that assume forward-only monotonicity. Protocol encodings exist in both fixed-width and variable-length encodings, each of which may be subject to encoding errors upon a rollover event. The breadth of potential encoding errors is wide, ranging from simple rollover to zero, to introducing negative values, and to over-insertion of data in the wireframe beyond what the specifications require (for example, encoding a 32-bit time value in a fixed-width field as 40 bits by accident). Where one subsystem has been remediated to use width-expanded time representations but dependent subsystems retain 32-bit interfaces, the dependent subsystems can produce malformed timestamps or crashes that propagate into the otherwise-compliant remediated system: the patched system remains exposed to failures

originating in its less-remediated neighbours, with the propagation pathway invisible to the testing regime that validated the remediation in isolation.

A particularly consequential category arises from the relationship between implementation-layer width expansion and protocol-layer wire encoding. *Expanding a struct or variable from 32-bit to 64-bit time_t within a system does not, by itself, guarantee correct protocol-layer behaviour where the wire encoding carries no semantic marker distinguishing the legacy representation from its successor.* Without an explicit version or semantic indicator on the wire, different systems interpreting the same bytes according to different time models may reach different conclusions about the same nominal timestamp – a class of silent disagreement that is operationally difficult to detect because both systems are behaving according to their own internally-consistent interpretation. Protocol implementations that convert between time representations – for example between struct tm and epoch time – are further vulnerable to integer-overflow conditions at rollover boundaries in legacy libraries whose conversion routines were not designed against these conditions.

Heterogeneous environments – multi-vendor, multi-version, with components at different patch levels – frequently manage time through aggregation services such as logging hubs, security-information-and-event-management platforms, and time-synchronization infrastructure. A single non-compliant component within such an environment can send corrupted timestamps that disrupt time-series interpretation across the aggregated view, with consequences for any downstream analysis that depends on consistent temporal ordering. The aggregation layer was not generally designed against the assumption that a participating component might emit deterministically wrong timestamps at known future dates; the design assumption was that timestamps would be approximately correct except in pathological cases, and the case described here is structurally more systematic than the pathological framing anticipated.

Backward-compatibility pressure complicates remediation across protocol boundaries. Legacy protocol support is preserved across deployed infrastructure for technical, regulatory, and business-continuity reasons that are individually defensible – including the continued operational presence of devices that cannot be updated, the certification status of installed systems, and the contractual expectations of long-running customer relationships. The cumulative effect is that vulnerable protocol behaviours are retained in environments whose host systems would otherwise be safe – *for example, the continued operational presence of SNMPv1 in environments running modern 64-bit platforms*, where the platform's modernity does not propagate into the protocol layer's time-handling.

Conformance testing for protocol implementations typically validates nominal-case behaviour under current operational conditions, where current operational conditions do not include the rollover boundary or its immediate vicinity. Rollover-boundary behaviour is rarely included in standard conformance certification processes, with the consequence that protocols may be certified as conformant whilst remaining exposed to predictable failures at known future dates. The structural limitation is not a defect in any individual certification regime but reflects the assumption embedded across the broader certification landscape that protocol implementations will be tested against conditions plausibly encountered in operation, rather than against conditions whose timing is known but whose appearance in operational testing requires deliberate construction.

A further category of exposure arises from the interaction between protocol-layer time handling and adversarial inputs. Time-drift attacks and induced desynchronization – as observed in documented GPS, NTP [b-Malhotra], and network identity and time zone (NITZ) spoofing cases – demonstrate that anomalous timestamp input from one protocol source can alter trust anchors across federated protocols, with consequences propagating through public-key infrastructure, consensus mechanisms, blockchain validation chains, and distributed database systems. *The adversarial dimension does not create the underlying exposure but compounds it*: an attacker capable of inducing timestamp manipulation across a protocol boundary gains access to a class of failure modes that the protocols themselves were not designed to resist.

Considered together, these protocol-layer conditions suggest that 2038-class remediation conducted within the boundaries of a single system, however carefully, cannot in the general case guarantee the system's continued correct operation when interacting with peer systems whose remediation status, protocol version, and time-encoding semantics it does not control. The implication is structural: remediation at the layer of any single implementation is necessary but not sufficient, and cross-implementation coordination at the protocol layer is among the items for which inter-SDO engagement would substantially advance the state of practice.

8.6 Critical infrastructure considerations

The structural conditions described in the preceding clauses take their most consequential form in critical-infrastructure environments – energy, telecommunications, transport, aerospace, water, healthcare, and the broader landscape of sectors developed in clause 10 – where the constraints on remediation are particularly stringent and the consequences of non-remediated exposure are particularly severe. System upgrades in these environments are routinely constrained by limited maintenance windows, lengthy certification processes, safety-regulation approval timelines, export-control restrictions shaped by geopolitical conditions, and sovereign-control requirements that vary across jurisdictions. Legacy 32-bit embedded architectures are deployed in operational roles whose physical and contractual lifetimes were specified against the assumption that the deployed hardware would remain in service across the upcoming rollover cluster – installed bases that cannot in the general case be replaced before 2038 without operational disruption disproportionate to the remediation goal.

Failure modes in critical-infrastructure environments do not generally present as isolated asset failures. Rollover-related disruption propagates through supervisory control and data-acquisition systems (SCADA), distributed control systems, logging infrastructure, backup routines, safety-instrumented-system shutdown triggers, and telemetry pathways in ways that produce coordinated rather than isolated degradation. This pattern matches the correlated-failure-pathway framing developed in clause 8.2 and is amplified by the structural conditions specific to critical-infrastructure operation – including the prevalence of long-lived deployed assets, the multi-tier supplier relationships through which time-handling behaviour was originally established, and the operational constraints on inspection and remediation.

Rollover vulnerability persists even in critical-infrastructure environments where modernization has nominally occurred. Where 32-bit runtime components remain present inside 64-bit host environments – through compatibility layers, soft-PLC implementations, programmable-logic-controller drivers, vendor software-development kits, application programming interface bridges, and legacy human-machine-interface software – the host environment's modernity does not propagate into the application substrate where time-handling behaviour is actually determined. Critical-infrastructure operators may believe their environments to be remediated on the basis of host-platform modernization while application-layer exposure persists invisible to the verification regime that informed the belief. The discovery and observability constraints developed in clause 8.4 apply with particular force in these environments, where in-situ testing is frequently constrained by safety, operational-continuity, and certification considerations that limit the operator's ability to validate time-handling behaviour under conditions that would actually reveal the exposure.

Supply-chain remediation cycles in critical infrastructure typically exceed three to five years, and operational lifetimes for critical-infrastructure hardware routinely extend across ten to twenty-five years or longer. The arithmetic is direct: where the remediation cycle for a class of deployed asset exceeds the time remaining to the rollover events described in clause 6, full remediation is not achievable through ordinary operational pathways for that class of asset within the available window. This condition does not apply uniformly across critical-infrastructure sectors or across asset classes within sectors, but it applies to a significant subset of the deployed landscape, with the consequence that partial remediation, harm minimization, and resilience under conditions of incomplete

remediation become operational categories of equal practical importance to full remediation itself, as further developed in clause 15.

In-situ testing in critical-infrastructure environments is frequently restricted or prohibited by safety, operational, and certification constraints – the same constraints that protect the systems against ordinary operational hazards also limit the testing modalities available to validate behaviour under non-routine conditions including rollover. The result is a heightened reliance on laboratory simulation, the limits of which were addressed in clause 8.4 and which the [b-Ariane 5] inquiry findings establish as insufficient for confident pre-deployment validation of behaviour under environment-mismatched conditions. The class of conditions to which 2038-class rollover belongs is precisely the class for which laboratory simulation is structurally inadequate as the sole validation pathway.

Time-dependent behaviour in critical-infrastructure environments operates across safety and performance envelopes whose tolerances are routinely tight. Rollover-induced time shifts or timestamp corruption can affect control-loop behaviour in ways that trigger protective shutdowns, produce overshoot conditions, or desynchronize coordinated operations across distributed sites. Indirect impacts propagate through dependent services, including GPS-derived timing distribution, NTP synchronization infrastructure, secure certificate validation chains, and automated recovery procedures that themselves depend on coherent time for correct operation; failures originating in critical-infrastructure systems can cascade into adjacent sectors with substantially different operational character, and failures originating in adjacent sectors can propagate into critical-infrastructure systems through the same dependency pathways.

A further structural constraint on remediation pathways bears direct surfacing. The legal and contractual frameworks surrounding modern products incorporating digital elements – including supply-chain attestation regimes, intellectual property protections, anti-tamper provisions in product certifications, and the broader contractual posture of software-as-a-service deployments – increasingly limit the technical possibilities for post-deployment inspection, reverse-engineering, and patching of components by parties other than the original vendor. These frameworks are individually defensible on the grounds of supply-chain integrity, intellectual-property protection, and security-against-tampering; their cumulative effect is to narrow the remediation window for 2038-class exposure to interventions the original vendor undertakes, on the original vendor's timeline, with the original vendor's operational capacity.

Where a vendor no longer exists, has discontinued support, has lost institutional memory of a legacy product, or simply lacks economic incentive to remediate a deployed product approaching end-of-life, the legal substrate currently closes pathways that the technical substrate would otherwise leave open.

This consideration interacts with the discovery challenge described in clause 6 and with the supply-chain opacity observations developed in clause 8.3 and clause 10 throughout: not only is it difficult to inventory deployed time-handling implementations and to test their behaviour in situ, but where such implementations are discovered to be at risk, the legal pathways to remediation by parties other than the original vendor are increasingly constrained.

The cross-sector coordination structure within which critical-infrastructure remediation can be advanced is itself the subject of established resilience frameworks across jurisdictions – including general-purpose business-continuity standards, sector-specific cybersecurity-management frameworks, and the broader landscape of regulatory and certification regimes addressing critical-infrastructure protection. Comparative analysis of how 2038-class remediation interacts with these frameworks is anticipated in subsequent revisions of this Technical Paper, as flagged in clause 1 and clause 8.2.

Considered together, these conditions suggest that effective preparedness work for 2038-class events in critical-infrastructure environments would benefit from coordination mechanisms that integrate critical-infrastructure operators, national computer-emergency-response-team coordination,

timing-service providers, and the regulators whose certification regimes shape the remediation pathways available to operators. The development of harmonized approaches addressing both fault-tolerance and controlled degradation across the rollover window – including the harm-minimization considerations developed in clause 15.6 – is among the items for which inter-SDO and cross-sector coordination would substantially advance the state of practice.

8.7 Coordinated vulnerability management constraints

The coordination machinery developed across the past two decades for managing software and hardware vulnerabilities – coordinated vulnerability disclosure (CVD) processes, the common vulnerabilities and exposures (CVE) program, common vulnerability scoring system (CVSS) severity assessment, and the broader institutional ecosystem connecting researchers, vendors, computer security incident response teams (CSIRTs), and regulators – was designed against a particular implicit model of the vulnerabilities it addresses. That model assumes a vulnerability is tied to identifiable affected software versions and a discrete disclosure date; that severity can be characterized through a scoring system reflecting impact and exploitability at the time of disclosure; that remediation is achievable through a patch issued by an identifiable vendor; and that the appropriate window between private disclosure to the vendor and public disclosure to the broader community can be measured in months. 2038-class exposure satisfies few of these assumptions.

The mismatch begins at the level of vulnerability identification. 2038-class exposure is not tied to a specific software version or a discrete disclosure date; it is an architectural condition of the deployed timestamp representation, present from the moment of original implementation, manifesting deterministically at predictable future dates regardless of disclosure activity. There is no moment at which the exposure was "discovered" in the sense that the CVE model contemplates; the underlying limitation has been documented in the technical literature for decades, and the question is not how to alert defenders to its existence but how to coordinate remediation across deployed systems whose operators are aware of the exposure in general but lack visibility into its specific manifestation in their environments. The CVSS scoring model, which characterizes severity in terms of impact and exploitability at the time of disclosure, does not naturally accommodate vulnerabilities whose severity is time-delayed and whose cross-domain systemic effects depend on the configuration of dependencies external to the affected system itself.

The failure-mode profile of 2038-class exposure further complicates the inherited coordination model. Where CVD processes assume vulnerabilities manifest as discrete observable failures that can be characterized and tested, 2038-class failure modes include silent state corruption, cascading logic divergence, and degraded operation that may not produce any externally observable signal at the moment of failure. Systems may already be operating with corrupted temporal-trust assumptions when exposed to manipulated time inputs from upstream protocol sources (NTP, GPS, NITZ), *in which case the attribution question – which component failed, when, and why – becomes operationally difficult to answer in the timeframes that incident-response coordination ordinarily presumes*.

The vendor-visibility-and-ownership assumption underlying CVD processes is, in many critical contexts, simply not satisfied for 2038-class exposure. The affected systems include products whose vendors have ceased operations, products whose original manufacturers have been acquired and rebranded multiple times, products incorporating original-equipment-manufacturer components whose chain of contractual responsibility has fragmented across multiple successor entities, and products embedded inside larger systems where the integration responsibility is held by a party other than the component author. The standard CVD pathway – researcher discovers vulnerability, contacts vendor, vendor develops patch, coordinated disclosure proceeds – presumes a vendor able and willing to engage. Across the deployed landscape relevant to 2038-class exposure, this presumption holds for some subset of affected products and not for others, with no general mechanism for routing remediation responsibility when the presumption fails.

The testing requirements specific to 2038-class exposure introduce a further category of mismatch. Confident validation of remediation, or even of the presence of exposure, requires the ability to advance system clocks across rollover boundaries under operational workload – an operation that may violate safety constraints, warranty terms, regulatory certifications, and operational-continuity expectations in the environments where 2038-class exposure is most consequential. Live testing under operational conditions is in many sectors not available as a routine option; laboratory simulation, the natural alternative, suffers from the environment-mismatch limitations developed in clause 8.4. The result is that confident pre-remediation characterization of exposure across deployed populations requires testing modalities that the existing coordinated-vulnerability ecosystem does not generally provide and that affected operators cannot generally undertake unilaterally.

Standard CVD timelines, which range from approximately 90 to 180 days between private disclosure and public disclosure, are incompatible with the remediation cycles characteristic of 2038-class exposure. Remediation of architectural time-handling exposure may require application-binary-interface changes, structural redesign of time-handling logic, or wholesale component replacement, none of which is generally achievable on patch-cycle timeframes. In regulated and safety-critical sectors, remediation rollouts may extend across multiple years to comply with certification and validation requirements; the disclosure-to-remediation window the conventional CVD model anticipates is, for these environments, structurally shorter than the remediation operation it is intended to coordinate.

The multi-stakeholder engagement required for 2038-class remediation does not map cleanly onto the bilateral researcher-vendor relationship around which CVD coordination is structured. *Effective response requires sustained engagement across sectors developed in clause 10* – information and communications technology, aviation, industrial control, healthcare, financial services, public safety, and the broader landscape of cross-sector dependencies – with attention to interactions across sector boundaries that no single existing coordination mechanism is structurally positioned to address. FIRST and adjacent communities have established relationships across which such coordination can be developed, but no formal international mechanism currently exists for multi-year predictive risk disclosure and remediation monitoring of the kind that 2038-class exposure indicates.

A further consideration arises from the public character of the underlying exposure. 2038-class events do not require sophisticated reverse-engineering or vulnerability discovery to identify: the timestamp formats and their representable bounds are publicly documented and trivially enumerable, *with the consequence that the structural exposure is known not only to defenders preparing remediation but to any party with adversarial intent and access to the same documentation.* The risk that simultaneous discovery and exploitation may occur across multiple actors before official remediation has progressed substantially is itself a structural condition that the existing coordination ecosystem was not designed to address – *there is no equivalent to the "patch first, disclose second" dynamic that conventional CVD uses to manage the disclosure-exploitation race, because there is no patch in the conventional sense and no single moment of disclosure to manage.*

Considered together, these mismatches between 2038-class exposure and the established CVD ecosystem suggest that the work of preparing for the upcoming rollover cluster requires a coordination mechanism with characteristics distinct from those of conventional vulnerability disclosure. Such a mechanism would be oriented around structural exposure with known deterministic timelines rather than around discrete vulnerabilities with disclosure dates; would support staged disclosure consistent with the multi-year remediation cycles that 2038-class exposure indicates; would provide pre-validation testing guidance addressing the testing-modality constraints described above; would maintain coordinated remediation tracking across long timeframes; and would integrate the dependencies on time-synchronization providers, certification authorities, and safety regulators that conventional CVD ecosystems coordinate only incidentally. Development of such a mechanism, and its relationship to existing CVD coordination infrastructure, is among the items for which inter-SDO and cross-sector engagement would substantially advance the state of practice and toward

which the [b-FIRST Time SIG] has been positioned as one of the venues where such coordination work can proceed.

Whether such a coordination mechanism can be developed and sustained depends substantially on the organizational culture in which it operates. The typology articulated by [b-Westrum] – distinguishing pathological, bureaucratic, and *generative* organizational cultures by the way information flows in response to anomaly – suggests that effective preparation for 2038-class events will require generative rather than bureaucratic handling: information moved to the right person in the right form and the right time frame, with cross-departmental and cross-organizational bridging encouraged rather than tolerated, and with response to discovered anomaly oriented toward inquiry rather than scapegoating or procedural justice. The structural conditions described throughout this Technical Paper – silent failure modes, fragmented vendor visibility, multi-year remediation cycles, cross-jurisdictional dependency – do not yield to bureaucratic information handling or punitive punish-the-messenger working culture. The cultural disposition of the institutions through which the response is coordinated is therefore not incidental to the technical work but is a key aspect on which the technical work depends.

8.8 Implications for incident response and digital forensics

The 2038-class timestamp rollover events described elsewhere in this Technical Paper carry consequences specific to incident response and digital forensics that call for explicit treatment. Incident response and forensic investigation depend on time coherence to establish event sequencing, attribute action to actor, and reconstruct the operational state of affected systems across the period of interest. Where the temporal substrate that produces and records the timestamps relied on for these purposes has itself failed, the analytical methods that incident response and forensic investigation conventionally employ are degraded in ways that do not have ready substitutes.

The first-order consequence is loss of timestamp monotonicity in evidentiary records. Log files may contain non-sequential records, negative or implausible timestamps, gaps where audit-trail entries should appear, and overlapping records from systems whose clocks have diverged. Each of these conditions impairs the analytical work that incident response and forensic investigation routinely undertake: reconstructing the timeline of an incident, correlating events across multiple systems, distinguishing legitimate from anomalous activity, and establishing the causal sequence by which an incident developed. The conventional methodologies for these analyses assume temporal coherence in the underlying records; the assumption does not hold under 2038-class rollover conditions, and substitute methodologies have not been systematically developed.

The second-order consequence is evidentiary. Forensic artefacts that rely on time metadata may no longer meet the integrity and reliability expectations of the legal, regulatory, and procedural frameworks under which they are presented. Where a court, regulator, or other competent authority is asked to admit time-stamped evidence, the question of whether the timestamps were produced under conditions of temporal coherence becomes load-bearing for admissibility. *A 2038-class failure affecting the systems that produced an evidentiary record may render the record inadmissible, or admissible only with caveats that limit its analytical weight*. The legal and procedural communities have not yet systematically engaged with the question of how to handle evidence produced under known temporal-substrate failure, and the development of such guidance is itself anticipated work for the period leading up to and beyond the rollover cluster.

The third-order consequence is coordination. Cross-border and multi-jurisdictional incidents – which are increasingly the rule rather than the exception in the digital domain – require correlation of evidence and event records across systems and organizations operated under different timekeeping implementations, regulatory regimes, and operational practices. Where the underlying timestamps are coherent, correlation across boundaries is operationally challenging but analytically tractable. Where the timestamps are incoherent, divergent interpretations of event timing across jurisdictions can delay response, hinder information sharing, reduce confidence in shared situational awareness, and in

extreme cases produce contradictory determinations about what actually occurred. The coordination mechanisms developed among CSIRTs, CERTs, PSIRTs, and other authorities for cross-jurisdictional response presume timestamp coherence as a precondition; their behaviour under conditions where the precondition has failed has not been systematically characterized.

These consequences interact with the constraints on coordinated vulnerability management developed in clause 8.7. Incident response and forensic investigation are downstream consumers of the trust assumptions that conventional vulnerability disclosure depends on; where those assumptions are degraded, the operational practices that follow from them inherit the degradation. A coordinated incident response that cannot rely on consistent timestamping across affected systems faces the same structural mismatch with established coordination frameworks that the vulnerability-management discussion identified at a different layer.

The implications for incident response and digital forensics therefore extend considerably beyond questions of system availability or operational continuity. They affect the trust, accountability, and coordinated-response capabilities on which the broader cybersecurity ecosystem itself depends for its own function. The structural observations made elsewhere in this Technical Paper – concerning the limitations of conventional vulnerability disclosure, discussed in clause 8.7, the discovery and observability constraints on detecting 2038-class exposure, as set forth in clause 8.4, and the cross-sector coordination requirements developed in clause 10 apply with particular force to the incident response and digital forensics communities, whose operational work depends on the temporal substrate whose failure modes this paper surveys.

Detailed treatment of incident-response-specific and forensic-specific methodologies for operating under conditions of temporal-substrate degradation is anticipated in subsequent revisions, in coordination with the FIRST community whose practitioners hold the operational expertise on which such treatment will depend.

Together, these dimensions form the analytical backbone for the Technical Paper, enabling consistent evaluation across devices, architectures, and Member State environments.

8.9 Coherence invariant and the distributed-proof framing

The eight analytical dimensions developed across clause 8.1 through clause 8.8 may be unified through a more formal observation. The Internet's operational coherence rests on what functions as a large distributed proof: the Request for Comments corpus maintained by the IETF, the Recommendations corpus maintained by ITU-T, and the adjacent standards bodies whose specifications interlock to constitute the deployed substrate together describe a system of local semantics whose interaction produces the global property the network's operators rely on. The running network is the constructive witness that this proof holds locally – not perfectly, but sufficiently to sustain global interoperability. This framing is not original to the present Technical Paper; it surfaces in the formal-methods, distributed-systems, and protocol-research literatures across the last several decades. Its applicability to the 2038-class cluster warrants explicit acknowledgement because the cluster is, in this framing, an exercise of the proof at a layer below the explicit semantics of the constituent standards. The forensic dimension of 2038-class exposure is rarely surfaced in the broader discussion of rollover-class risks, and its absence from that discussion is itself a structural condition this Technical Paper must surface: *forensic capacity is, by construction, the capacity that becomes relevant after an incident has occurred, which is precisely the moment at which the costs of its degradation become non-recoverable.*

A coherence invariant may be informally stated as follows: for any two participating systems A and B, time-dependent validations produce compatible outcomes under shared protocol assumptions. The 2038-class events, considered through this framing, are the first sustained operational stress on the conditions under which the coherence invariant has held across the deployed substrate. Partial migration – the operational state in which some participating systems have adopted width-expanded or otherwise remediated time representations while others have not – produces the structurally

distinctive failure mode the present Technical Paper surveys: non-convergent validation states in which the invariant holds for some pairs of participating systems and fails for others, with the failure pattern depending on the specific configuration of remediation across the substrate rather than on any single system's correctness in isolation. The system may remain functional in this state; it becomes structurally harder to reason about, to coordinate across, and to defend. The protocol-layer interoperability observations developed in clause 8.5, the discovery and observability constraints developed in clause 8.4, and the coordinated vulnerability management constraints developed in clause 8.7 are operational manifestations of the same underlying coherence-invariant question, examined at different layers of the analytical framework.

The formal characterization of the coherence invariant – including its precise statement, the protocol-specific conditions under which it holds, the empirical question of whether contemporary network topology retains the fault-tolerance properties established in earlier eras (notably [b-Boehm] distributed-routing simulations, which established that packet-switched networks could absorb substantial node loss while maintaining routing convergence), and the testable question of under what specific conditions partial migration violates the invariant in the deployed substrate – is among the items for which inter-SDO engagement and engagement with the formal-methods and protocol-research communities would substantially advance the state of practice. Detailed treatment, including consideration of whether a formal proof of the invariant's preservation under specified partial-migration conditions is achievable for load-bearing protocol subsets, is anticipated in subsequent versions of this Technical Paper.

9 Reference architecture

High-level description of architectural layers affected by time failures.

NOTE – This clause is staged deliberately. The reference architecture is intended to function as the analytical scaffold on which the cross-sector survey in clause 10 rests – a layered description of the global time-dependency stack, from physical-layer time sources through NTP/precision time protocol (PTP) distribution to application-layer timestamp consumers (clause 9.1); the propagation of rollover conditions across that stack and the layers at which graceful degradation is most achievable (clause 9.2); the synchronization dependencies introduced by cloud and platform architectures (clause 9.3); and the protocol layers carrying embedded timestamp fields (clause 9.4). This clause will be developed in coordination with the FIRST time security SIG's in-progress empirical work on hyperscaler NTP architectures rather than to present an illustrative architecture in advance of that grounding. This choice is consistent with the document's own position, developed in clause 1 and clause 8.4, that simulated and illustrative characterizations of time-dependency behaviour are structurally insufficient substitutes for observation under operational conditions; the reference architecture is held to the same standard the paper applies to the systems it surveys. The structural commitments described across clauses 9.1-9.4 are stable; their development to the analytical density of clause 8 is anticipated in subsequent revisions as the SIG's empirical observations mature.

9.1 Global time dependency stack

NOTE – A figure illustrating the global time dependency stack – from physical layer time sources through NTP/PTP distribution to application-layer timestamp consumers – is in preparation for subsequent versions. The accompanying prose enumerates the layers and their dependencies.

9.2 Rollover propagation chain

NOTE – A figure illustrating rollover propagation across the dependency stack is in preparation. The prose treatment below traces propagation paths and identifies the layers at which graceful degradation is most achievable.

9.3 Cloud and platform synchronization dependencies

NOTE – The cloud and platform synchronization dependencies described here will be developed further in subsequent revisions, in coordination with [b-FIRST Time SIG] empirical observations on hyperscaler NTP architectures.

9.4 Protocol layers with embedded timestamp fields

NOTE – The protocol layers enumerated here are illustrative rather than exhaustive. Subsequent revisions will expand the enumeration with citations to the relevant IETF and ITU-T specifications **(non-normative references to common patterns)**.

10 Cross-sector exposure survey

The structural patterns this Technical Paper describes manifest differently across the sectors of the industrial landscape. This clause presents a survey of sectors whose 2038-class exposure has been identified by the [b-FIRST Time SIG] community and ITU-T Study Group 17, organized against the established taxonomy of industrial infrastructure published by [b-Hayes]. Not all sectors are developed in this revision; the selection reflects a combination of acute exposure characteristics, availability of citable empirical material, and the time available before the present revision deadline. Sectors not developed here are flagged for subsequent revisions; their inclusion in this scaffold reflects a commitment to comprehensive coverage rather than completed treatment.

The [b-Hayes] taxonomy, while not a formal regulatory or engineering standard**,** is used here for two reasons. First, it provides a recognisable cross-sector reference frame familiar to readers from infrastructure, regulatory, and operator backgrounds. Second, its ordering – moving from raw materials through processing, distribution, and movement to disposal – corresponds approximately to a supply-chain dependency graph, which is itself useful for thinking about how 2038-class failures in one sector propagate into others. The [b-Hayes] taxonomy is supplemented here with several additions, reflecting the post-publication emergence of additional cross-cutting infrastructure layers with material time-coordination dependencies. Where this Technical Paper introduces sectors or layers that do not appear in the Hayes taxonomy, the corresponding table rows are flagged accordingly.

Where sectors are flagged as stubs in the table below, the developed treatment is anticipated in subsequent revisions. Engagement from sector practitioners – including the constituencies named in each stub – is welcomed. Where a stub closes with a brief note flagging why engagement from a particular sector is structurally important, the note surfaces the sector-specific consideration; absence of such a note in a stub indicates that the engagement invitation is general rather than tied to a distinctive sector property. Table 4 summarises the sectors considered in the cross-sector exposure survey.

Table 4 – Cross-sector exposure survey

Sector	Primary 2038-class exposure	Distinctive structural conditions	Status in this version
Telecommunication networks	Carrier-grade timing infrastructure, mobile network synchronization, billing and interconnect, lawful intercept timestamping, Operations Support System (OSS)/BSS substrate	Cross-sector substrate dependency; eroded fixed/mobile distinction; long equipment lifecycles; deep multi-tier supplier relationships	Developed: see clause 10.1
Financial systems	Trading, clearing, point-of-sale systems, and settlement infrastructure, central bank digital currency (CBDC) architectures, cryptocurrency/blockchain and smart contract substrates	Order-of-precedence requirements; regulatory clock-synchronization regimes; cross-border settlement dependencies	Developed: see clause 10.2

Table 4 – Cross-sector exposure survey

Sector	Primary 2038-class exposure	Distinctive structural conditions	Status in this version
Biological research and pharmaceuticals†	Clinical trial records and audit trails, biological-research laboratory information systems, pharmaceutical manufacturing computerized-systems validation (PLC, SCADA, manufacturing execution system (MES), LIMS, enterprise resource planning (ERP), historian, database, firmware, middleware), pharmaceutical distribution and cold-chain logistics under GxP regimes	Multi-decade regulatory record-retention obligations (25-year clinical trial master file archiving under [b-EU 536/2014]); strict timestamp-integrity requirements under [b-21 CFR 11] and Good Clinical Practice; deep computerized-systems-validation dependencies across manufacturing; cross-jurisdictional pharmaceutical operations	Stub: see clause 10.3
Electric and connected vehicles	Embedded electronic control units (ECUs) in vehicles, charging infrastructure, in-vehicle GNSS receivers, vehicle-to-grid PKI	ECU proliferation; supply-chain opacity; supplier-team lifecycle mismatch; PKI dependencies in vehicle-to-grid communication	Developed: see clause 10.4
Implantable and hospital medical devices	Implanted medical devices, hospital infrastructure, cross-border interaction with patient-implanted device firmware and clinical equipment compatibility	Patient-individual remediation requirements; heavy regulatory oversight; cross-jurisdictional coordination	Developed: see clause 10.5
CSIRT/PSIRT/CERT coordination	Cross-sector incident correlation, multi-jurisdictional response coordination, shared situational awareness across operator boundaries	Operational coordination layer rather than technical sector; correlation depends on coherent timeline across reporting sources	Stub: see clause 10.6
Emergency services operations	Computer-aided dispatch (CAD) systems, public safety answering point (PSAP) call-handling, project 25 (P25)/ terrestrial trunked radio (TETRA) emergency-radio synchronization, mutual-aid coordination timestamps	Multi-agency cross-jurisdictional coordination; legal evidentiary weight of dispatch records; mixed legacy and modern radio infrastructure	Stub: see clause 10.7
ICS/operational technology (OT) environments	PLCs, SCADA systems, distributed control systems, safety instrumented systems, engineering workstations and historians	Long operational lifecycles; safety-critical certification regimes; deep multi-tier supplier relationships; operational impossibility of arbitrary system updates	Stub: see clause 10.8

Table 4 – Cross-sector exposure survey

Sector	Primary 2038-class exposure	Distinctive structural conditions	Status in this version
Cloud and ICT infrastructure	Cloud platforms, hosting infrastructure, identity and access management, observability platforms, enterprise ICT substrate	Hyperscale provider concentration; multi-tenancy propagation of platform-layer time-handling; enterprise identity infrastructure lifecycles; observability dependencies invisible to application owners	Stub: see clause 10.9
GNSS-dependent services	Receiver-side time handling, cross-sector timing dependency	Substrate dependency for multiple sectors above; documented prior rollover events; consumer-grade receiver firmware lifecycles	Stub: see clause 10.10
Space infrastructure	Satellite onboard systems, ground-segment timing distribution, constellation-operations coordination, satellite earth station synchronization	Multi-decade satellite operational lifetimes; cross-jurisdictional ground segments; dependency on coherent time for constellation coordination; convergence with GNSS-derived timing services	Stub: see clause 10.11 (receiver-side GNSS dependencies addressed separately in clause 10.10)
Broadcasting and emergency alerting	Programme scheduling and ad-insertion synchronization, simulcast coordination, [b-EAS]/ [b-ETSI TS 102 900]/ Cell Broadcast emergency-warning chains, digital video broadcasting – second generation terrestrial (DVB-T2)/ advanced television systems committee (ATSC) 3.0/digital audio broadcasting plus (DAB+) digital broadcast timing	Strict latency requirements for emergency warnings; forensic-review dependence on timestamp integrity; geographically dispersed transmitter networks	See clause 10.12
Government services and public administration	Civil registries (births, marriages, deaths, citizenship), court records, tax administration, cadastral records, elections systems, immigration databases, social-benefits administration	Legal force of timestamped records; multi-decade record retention; downstream propagation into pensions, inheritance, property rights, and identity verification chains	Stub: see clause 10.13
Insurance and reinsurance	Claims-processing timestamps, actuarial time-series data, catastrophe-modelling platforms, reinsurance treaty	Solvency frameworks requiring auditable claims chronology; principal market vehicle for absorbing correlated tail risk across all other sectors; pre-event	Stub: see clause 10.14

Table 4 – Cross-sector exposure survey

Sector	Primary 2038-class exposure	Distinctive structural conditions	Status in this version
	event-occurrence and notification windows	pricing of 2038-class risk constrains post-event capital response	
Epistemic coordination layer	Human knowledge attrition, institutional memory loss, documentation and provenance decay, AI-mediated training-data contamination	Substrate of capacity-to-know underlying every other sector's remediation; recursive AI dynamics novel relative to Y2K; structural production of knowledge loss by commercial reorganization	Stub: see clause 10.15
Buildings and built environment†	Building management system/building automation system (BMS/BAS), access-control and surveillance timestamping, fire-suppression sequencing, elevator dispatch coordination, heating, ventilation, and air conditioning (HVAC) integration with utility demand-response	Substrate category – essentially all other sectors operate within buildings; expanding integration with cloud-based facilities management and broader ICS/OT environments; heterogeneous exposure profile (operationally tolerable to consequential)	Stub: see clause 10.16
Manufacturing and industrial production†	PLC and distributed control system (DCS) sequencing on production floors, MES/ERP timestamping integration, quality-control batch records, pharmaceutical [b-GMP]/[b-GAMP] and food-safety [b-HACCP] timestamps, semiconductor fabrication metrology	Multi-month production cycles in semiconductor fabrication; regulatory consequences of batch-record corruption for safety-critical products (pharmaceuticals, automotive, food); deep multi-tier supplier dependencies	Stub: see clause 10.17

Table 4 – Cross-sector exposure survey

Sector	Primary 2038-class exposure	Distinctive structural conditions	Status in this version
Logistics and distribution†	Supply-chain chain-of-custody timestamps, customs clearance, cold-chain integrity verification, pharmaceutical good distribution practice, warehouse automation real-time location system (RTLS), last-mile delivery coordination, retail-distribution endpoints including automated tank gauges (ATGs) and fuel-station point of sale (POS)	Cross-border operations spanning many jurisdictions and protocol generations; integrating layer between manufacturing, retail, and consumption; cold-chain regulatory consequences for pharmaceuticals, vaccines, and food safety	Stub: see clause 10.18
Mining and extraction†	Embedded control systems in extraction, processing, and ore-transport equipment	Long equipment lifecycles; remote and harsh operating environments; heterogeneous SCADA, PLC, across mixed vendors, generations, and communication protocols	Stub: see clause 10.19
Water and wastewater†	Embedded control in dam and reservoir management, pumping, treatment, and distribution; SCADA timestamping	Critical public health dependency; municipal operator scale; long asset lifetimes	Stub: see clause 10.20
Food and farming†	Precision agriculture timing, GNSS dependence, cold-chain logistics, processing-line automation	GNSS-derived positioning; supply-chain traceability requirements; seasonal operational windows, biological perishability, food safety	Stub: see clause 10.21 (GNSS dependencies addressed in clause 10.10)
Oil and gas†	Pipeline control systems, refinery automation, custody-transfer metering, distribution	Multi-decade infrastructure lifecycles; cross-jurisdictional pipeline networks; safety-critical certification (e.g., [b-IEC 61508], [b-IEC 61511], [b-IEC 62443])	Stub: see clause 10.22
Power plants†	Generation control systems, protection relays, plant historians	Multi-decade lifecycles; safety-critical certification regimes (e.g., [b-IEC 61850], [b-IEC 62443])	Stub: see clause 10.23
Power grid†	Grid synchronization, phasor measurement units, transmission control, billing	Coherent-time dependency at the seconds-to-microseconds scale; cross-jurisdictional interconnection; safety-critical certification	Stub: see clause 10.24

Table 4 – Cross-sector exposure survey

Sector	Primary 2038-class exposure	Distinctive structural conditions	Status in this version
		regimes (e.g., [b-IEC 61850], [b-IEC 62443])	
Vehicles and traffic infrastructure†	Roadside traffic management systems, signal coordination, tolling infrastructure, vehicle-to-everything (V2X) roadside units, congestion management sensors	Municipal and transport-authority operator landscape; long civil-infrastructure lifecycles; integration with vehicle-side systems addressed in clause 10.4	Stub: see clause 10.25
Railroad†	Signalling systems (e.g., ERTMS), positive train control, train-to-ground communications over GSM-R/FRMCS, scheduling systems	Multi-decade infrastructure lifecycles; safety-critical certification; cross-operator interoperability	Stub: see clause 10.26
Bridges and tunnels†	Structural monitoring sensors, traffic management, ventilation control	Long infrastructure lifecycles; instrumentation heterogeneity	Stub: see clause 10.27
Aviation†	Avionics, air traffic management, GNSS dependencies, maintenance and certification records	Multi-decade airframe lifecycles; safety-critical certification; documented prior incidents (GPS week rollover 2019)	Stub: see clause 10.28 (documented incidents referenced in clause 7.1)
Shipping†	Maritime navigation, port operations, container tracking, GNSS dependencies	GNSS-derived positioning; long vessel lifecycles; cross-jurisdictional operations	Stub: see clause 10.29
Wastes and recycling†	Processing-facility automation, materials-tracking systems	Operational-technology heterogeneity; environmental compliance timestamping	Stub: see clause 10.30

† Sectors marked with a dagger were drawn from the infrastructure taxonomy in [b-Hayes]; their inclusion reflects editorial coverage rather than analytical priority. All listed sectors are exposed to 2038-class risk along the failure-mode dimensions described in clause 6.2.

10.1 Telecommunication networks

Telecommunication networks represent a cross-cutting exposure class for 2038-class rollover risks, both as a sector with its own structural conditions and as the substrate on which most other sectors in clause 10 depend for delivery of their services. Telecommunication infrastructure carries the time-coordination dependencies of every connected device, every transaction routed across the network, and every coordinated workflow that spans geographic distance. The structural conditions of the sector – long equipment lifecycles, infrequent software updates, coexistence of multiple generations of technology, deep multi-tier supplier relationships, and the operational dependence of downstream sectors on continuous service – combine to produce remediation considerations distinctive from those of other sectors.

Networks of systems greatly compound the problem through their interdependence and reliance on communication protocols that may not themselves be resilient against 2038-class rollovers. Resilient

systems need not only protect their own internal time storage mechanisms but must expect that external systems may transmit non-conformant temporal values. Communication failures themselves come in multiple forms, such as fixed-width timestamp fields in networking protocol definitions, to variable-length encoding systems that have maximum-length constraints. Communications from remote systems may err because the system is wholly non-conformant, partially non-conformant, because the communication protocol itself is fully or partially susceptible. Group communication systems bring further challenges as some members of the group may be hardened against 2038-class rollovers while others remain unsupported. Resilient systems must guard against sudden incomprehensible values or discontinuities.

A structural observation worth surfacing at the outset: the historical distinction between "fixed" and "mobile" telecommunications has substantially eroded. In the era of virtual circuits, VoIP, software-defined networking, and packet-switched delivery of services that were previously circuit-switched, the population of operational telecommunication engineers retaining hands-on experience with pure circuit-switched infrastructure is approaching zero. The implication for 2038-class remediation is that the deployed telecommunications base is more homogeneously packet-based than the regulatory and operational frameworks that describe it, and that legacy circuit-switched skills required for fallback operation in degraded conditions may not be available at the scale a coordinated rollover event would require.

10.1.1 Carrier-grade timing infrastructure and mobile network synchronization

Telecommunication networks depend on time-synchronization infrastructure at multiple layers. Carrier-grade deployments of the network time protocol, the precision time protocol [b-IEEE 1588], and [b-ITU-T G.8271] synchronization profiles provide the time substrate on which both fixed and mobile services operate. These deployments commonly use GNSS-derived primary references, with multiple layers of holdover and distribution between the reference and the consuming systems.

Mobile network synchronization requirements have tightened progressively across successive technology generations. Modern 5G deployments using time division duplex configurations require time alignment across base stations within tolerances measured in microseconds, with the most demanding ultra-reliable low-latency communications (URLLCs) configurations requiring even tighter alignment. The next-generation digital infrastructure depending on these capabilities – including smart-city deployments, industrial automation, autonomous vehicle infrastructure, and precision agriculture – inherits the time-handling dependencies of the underlying 5G synchronization layer. A 2038-class failure affecting carrier-grade timing infrastructure does not stay confined to telecommunications; it propagates into every sector whose operational architecture assumes continuously available, coherent time delivered through mobile network synchronization.

10.1.2 Legacy and access-layer coexistence

Telecommunication infrastructure is exposed to a range of interconnected risks driven by increasing network complexity, long equipment lifecycles, and dependencies on time-synchronized and software-controlled operations. Legacy access technologies – such as older DSLAM platforms, legacy aggregation nodes, and first-generation customer premises equipment (CPE) deployed at customer premises – often remain in service well beyond their original design horizon and may rely on outdated firmware, embedded software, or unsupported components. These assets frequently coexist with modern virtualized, cloud-native, and automated network functions, creating hidden interoperability and resilience risks. Limited observability into legacy systems, constrained maintenance windows, and supply-chain or vendor support limitations can hinder proactive remediation. As a result, failures or vulnerabilities in legacy elements may propagate across interconnected network layers and may even traverse compliant systems, potentially leading to service degradation, loss of synchronization, billing or authentication anomalies, and cascading impacts on critical services that depend on the telecom network.

10.1.3 Billing, settlement, and cross-operator interconnect

Carrier billing and settlement infrastructure depends on timestamp coherence across operator boundaries. Call detail records, data session records, and equivalent event-generation systems produce time-stamped events that are subsequently aggregated, rated, and settled across the boundaries between operators, between roaming partners, and between carriers and their wholesale customers. Cross-operator settlement frameworks – including international roaming clearing, interconnect billing, and wholesale settlement arrangements – depend on consistent interpretation of timestamps across systems operated by different commercial entities under different regulatory regimes.

A 2038-class failure affecting timestamp interpretation at the boundary between connections and endpoints introduces ambiguity that is structurally similar to the order-of-precedence problem treated in clause 10.2. Unsettled inter-operator obligations are not benign: they propagate into commercial dispute, regulatory escalation, and in extreme cases service disruption (e.g., through de-peering when an operator suspends interconnect with a counterparty whose timestamp integrity cannot be reconciled).

10.1.4 Lawful intercept and regulatory timestamping

Telecommunication operators are subject to regulatory obligations across most jurisdictions requiring time-stamped retention and, where lawfully requested, production of communication metadata. The technical implementations of these obligations vary across jurisdictions but share structural conditions: time-stamped event records are retained for defined periods, may be requested by competent authorities, and are subsequently used in legal and regulatory proceedings where their temporal integrity is consequential.

The cross-border dimension introduces a coordination problem distinctive to telecommunications. A communication session spanning multiple jurisdictions may produce time-stamped records under different timekeeping implementations at the originating, transit, and terminating operators. A 2038-class failure affecting timestamp integrity at any point in this chain may render cross-border lawful-intercept records inconsistent with one another, with consequences ranging from evidentiary admissibility questions to coordination failures between investigating authorities relying on the records to establish event sequencing. The structural observation parallels that made for digital forensics generally in clause 8.8: the integrity of time-stamped evidence is consequential for downstream legal and regulatory processes that depend on it, and the cross-border dimension amplifies the coordination challenge.

10.1.5 OSS/BSS and operational substrate

Operational support systems (OSS) and business support systems (BSS) – the operational software substrate on which carriers manage their networks, customers, and commercial relationships – carry their own 2038-class exposure. OSS platforms handle network management, fault correlation, performance monitoring, and provisioning workflows; BSS platforms handle customer management, billing, and revenue assurance. Both depend on coherent timestamping across distributed systems, frequently incorporating legacy software components whose 2038-class exposure has not been systematically evaluated.

The operational consequence is that even where carrier-grade timing infrastructure is remediated correctly, OSS/BSS-layer 2038-class exposure may still produce operational failures that manifest as service incidents, billing errors, or fault-management blind spots, as explained further in clause 10.1.1. The remediation pathway is constrained by the same factors as broader telecommunication remediation: vendor relationships, certification requirements, and the operational impossibility of taking critical systems offline for extended periods.

10.1.6 Telecommunications as cross-sector substrate

Telecommunication networks are distinctive among the use cases in this Technical Paper because they are simultaneously a sector with their own structural exposure and the substrate on which most other sectors depend for service delivery. A 2038-class failure affecting telecommunication infrastructure propagates into the dependent sectors – financial, healthcare, transportation, energy, public safety, and emergency response – at the speed of network operation, which is to say faster than coordinated response mechanisms typically operate.

The combination of long equipment lifecycles, eroding capacity for circuit-switched fallback, dependence of downstream sectors on continuous synchronization, cross-border interconnect coordination requirements, and the regulatory dimensions of lawful intercept produces a remediation profile that requires coordinated attention across operators, regulators, standards bodies (including ITU-T's own work on time and synchronization in [b-ITU-T G.8271] and related Recommendations), and the broader community of sectors depending on telecommunication carrier infrastructure for their operational continuity.

10.1.7 Routing and naming infrastructure: BGP/RPKI and DNSSEC

The Internet's routing and naming substrates depend on cryptographically signed assertions whose validity is bounded by time intervals embedded in the assertions themselves. Two protocol systems warrant particular attention in this Technical Paper, both because their correct operation is foundational to the operation of essentially every other system addressed in clause 10 and because the time-bounded validity arithmetic on which they depend is structurally exposed to 2038-class events in ways that are not always recognized in their respective communities.

The resource public key infrastructure (RPKI) supports validation of routing announcements in the border gateway protocol (BGP), which carries the routing information on which all inter-network traffic depends. Route origin authorizations (ROAs) and the cryptographic objects supporting their validation carry explicit notBefore and notAfter timestamps that bound their validity interval. The validation logic by which routing systems determine whether an announcement is currently valid performs arithmetic operations on these timestamps relative to the validating system's current time. The structural exposure follows directly: implementations performing this arithmetic with 32-bit signed time representations are subject to the failure modes developed in clause 6.2, and the consequences of incorrect validity assessment in routing infrastructure are operationally severe. A system that incorrectly determines an authorization to be expired may refuse to propagate routes that should be accepted; a system that incorrectly determines an authorization to be not yet valid may produce equivalent refusal. The cross-sector consequences of routing-layer disruption are immediate and broad: routing failures propagate at network operation speed across the dependencies developed throughout this clause.

The domain name system security extensions (DNSSEC) provide origin authentication and integrity protection for domain name system (DNS) responses through a hierarchy of cryptographic signatures. Each signature carries explicit inception and expiration timestamps; validation requires that the validating system's current time fall within the inception-to-expiration interval. The same structural exposure applies: implementations performing this arithmetic with 32-bit signed time representations face the failure modes developed in clause 6.2, with the consequence that DNSSEC validation may fail in ways that cause the affected name resolutions to be rejected as bogus. The downstream effects propagate across every service whose endpoint discovery depends on DNS – which is to say, essentially every service surveyed in clause 10. A further consideration warrants explicit mention: the long-lived nature of certain DNSSEC signing keys and the operational complexity of key rollover under degraded conditions introduces an additional remediation dimension distinct from the protocol-level timestamp arithmetic itself.

The structural property that BGP/RPKI and DNSSEC share – time-bounded validity arithmetic operating as a precondition for the correct operation of the substrate on which essentially all other sectors depend – distinguishes them from sector-specific time-handling concerns. Their treatment in this Technical Paper is anchored in clause 10.1 because the routing and naming substrates are properly understood as part of the telecommunication carrier infrastructure on which downstream sectors depend; the analytical conditions are nonetheless cross-cutting and bear on the conditions developed in clause 10.9, clause 10.6, and the cross-sector substrate framing developed in clause 10.1.6. The relationship between these protocol-layer concerns and the broader temporal root of trust framing developed in clause 7.2 is direct: routing and naming infrastructures are operational instances of substrate-layer time-coordination commitments whose correct operation depends on every participating system sharing a coherent and arithmetic-consistent representation of time.

10.2 Financial systems

Financial market infrastructure represents a particularly acute exposure class for 2038-class rollover risks. Unlike many other sectors where coherent time is a useful operational property, in financial systems coherent time is structurally constitutive of the system's function. Modern trading, clearing, and settlement systems depend on precise, monotonically increasing timestamps to establish order of precedence between transactions. When two transactions affect the same asset and arrive at the same venue, the question of which transaction occurred first is not a matter of operational convenience – it is the basis on which legal title is determined. Loss of timestamp coherence, even briefly, translates directly into ambiguity in property rights.

This dependency is formalized in regulatory frameworks. In the European Union, Article 22c of [b-EU 600/2014], supplemented by [b-EU 2025/1155], requires trading venues and other market participants to synchronize the business clocks they use to record reportable events, with maximum permitted divergence from UTC measured in microseconds for the most time-sensitive activities. In the United States, [b-CAT] imposes equivalent obligations on reportable order events. Asia-Pacific financial jurisdictions maintain analogous regulatory frameworks. These regimes assume that participating systems will continue to produce comparable, synchronized timestamps indefinitely. They do not specify behaviour under timestamp-representation failure.

The chain of consequences from timestamp ambiguity to systemic financial impact is short. Trades that cannot be definitively sequenced cannot be definitively cleared. Clearing failures propagate into settlement failures. Unsettled positions in financial systems are not benign artifacts – they propagate into counterparty exposure, intraday liquidity stress, and ultimately into regulatory and supervisory action. The [b-WTISD26] expert report explicitly identifies satellite-timing failure causing transaction rejection as a representative cascade scenario; the 2038-class transition introduces a deterministic, calendar-driven version of this failure mode that, unlike a satellite-timing event, will affect all systems with vulnerable time-handling implementations simultaneously rather than in localized clusters.

10.2.1 Cross-border settlement infrastructure

Cross-border settlement systems amplify the exposure because mixed-implementation states – some systems 64-bit-aware, some not – produce boundary conditions across jurisdictional lines. SWIFT, which carries the substantial majority of cross-border interbank messaging, depends on timestamp coherence across thousands of participating institutions running heterogeneous infrastructure. The depository trust and clearing corporation (DTCC), which clears and settles a large fraction of US securities transactions, depends similarly on coherent time across its participant base. Real-time gross settlement (RTGS) systems operated by central banks depend on timestamp coherence as a basic precondition for queue management and final settlement.

The structural observation is that no single operator in these networks has visibility into the time-handling implementations of all participating systems. A participant whose internal systems handle 2038-class transitions correctly may nonetheless experience operational failure if a counterparty's systems do not. The failure mode propagates through the dependencies that make cross-border settlement work at all.

10.2.2 Central bank digital currencies

Central bank digital currency (CBDC) implementations are emerging as a forward-facing exposure surface. The [b-DigiEuro] project is the most concretely advanced major CBDC initiative in a jurisdiction whose time-handling regulations are explicit and well-documented. CBDC architecture is being designed now, with time-handling assumptions being established at the foundational layer of systems intended to remain in production for decades. The window for surfacing 2038-class considerations during architectural design – before implementation choices are locked in and propagated across participating institutions – is open at present and will close as deployment progresses.

This is not unique to the digital euro. CBDC projects in production deployment or advanced pilot stages in multiple jurisdictions face similar architectural decisions. The cross-jurisdictional interoperability of future CBDCs will depend on time-handling assumptions made today; mixed-implementation states across jurisdictional CBDC infrastructures would compound the cross-border settlement exposure described above.

10.2.3 Cryptocurrencies and decentralized finance

Cryptocurrencies and decentralized finance protocols inherit time-coordination dependencies despite being designed, in many cases, to minimize centralized trust assumptions. The dependencies operate at several layers.

At the protocol layer, block-timestamp validation rules constrain how far a proposed block's timestamp may deviate from network time. Bitcoin's nTime field is a 32-bit unsigned integer with a rollover horizon in 2106 rather than 2038. However, the protocol's resilience against 2038-class disruptions is not solely a function of its own field widths. Network cohesion depends on continued coordinated behaviour across thousands of independently operated nodes, many of which run on operating systems and software stacks with 2038-class exposure below the application layer. A sufficient fraction of nodes experiencing time-handling failure at the operating-system or middleware layer could produce network partitions or consensus failures that the protocol's own time-handling does not protect against.

Network-cohesion shearing below the application layer is a failure mode distinct from, and not addressed by, the application-layer time-handling design.

Ethereum and other proof-of-stake systems exhibit similar dependencies with different specifics. Smart contracts may include block.timestamp comparisons that propagate timestamp assumptions into application logic, creating exposure surfaces that are not always visible from the protocol design alone. DeFi protocols implementing time-locked contracts, vesting schedules, or oracle-based price feeds may carry 2038-class assumptions inherited from underlying libraries or from interaction with external systems.

The structural observation is that decentralized systems are not insulated from time-coordination problems by their decentralization. The dependencies are inherited from the substrate on which the systems run, and the systems' own design choices about time-handling do not address dependencies below their own design layer.

10.2.4 Financial systems: Temporal ambiguity to settlement risk

Financial systems are distinctive among the use cases in this Technical Paper because temporal ambiguity in financial systems translates immediately into legal and economic ambiguity. A transaction that cannot be definitively sequenced cannot be definitively settled; an unsettled transaction in a regulated financial system is not a deferred operation but an active liability. The 2038-class transition introduces a deterministic, calendar-driven amplification of failure modes that have manifested episodically in documented historical incidents, including but not limited to time-handling errors that contributed to or compounded events such as the [b-Flash Crash] and the [b-Knight Capital] events. These historical incidents are not themselves 2038-class failures, but they demonstrate the operational sensitivity of financial systems to time-handling errors at much smaller scales than a coordinated 2038-class rollover event would produce. These dependencies extend further than is commonly recognized: high-frequency trading systems, exchange matching engines, and cross-border settlement infrastructure all assume coherent, sub-millisecond-accurate time, and the chain of consequences from timestamp ambiguity to transaction repudiation to systemic clearing failure is short.

10.3 Biological research and pharmaceuticals

Biological research, clinical research, and pharmaceutical manufacturing and distribution constitute a category of 2038-class exposure with structural conditions distinctive from the sectors developed elsewhere in this cross-sector exposure survey. The category encompasses clinical trial infrastructure across the trial lifecycle, biological research records and laboratory information systems, pharmaceutical manufacturing operations under [b-GMP] and [b-GAMP], pharmaceutical distribution and cold-chain logistics under [b-GDP], and the broader landscape of regulatory recordkeeping required by national and international pharmaceutical and clinical-trial frameworks.

Distinctive structural conditions in this category include the multi-decade record-retention requirements imposed by clinical-trial regulation – [b-EU 536/2014] requires sponsors and investigators to archive clinical trial master files for at least twenty-five years after the end of the trial, with subject medical files archived in accordance with national law, and analogous record-retention requirements obtain in other jurisdictions; the strict timestamp-integrity obligations imposed by [b-GCP] and by frameworks such as [b-21 CFR 11] in the United States, where electronic-record audit trails carry direct regulatory and evidentiary weight; the deep deployment of computerized systems requiring formal validation across pharmaceutical manufacturing environments – including programmable logic controllers, supervisory control and data acquisition systems, manufacturing execution systems, laboratory information management systems, enterprise resource planning systems, historians, databases, firmware, and middleware – each of which carries potential 2038-class exposure within the validated state of the computerized system; the pharmaceutical-distribution and cold-chain logistics dependencies on accurate timestamping for expiration management, serialization, proof of delivery, temperature-log integrity, and supply-chain traceability under [b-GDP]; and the cross-jurisdictional nature of contemporary pharmaceutical operations, where products manufactured in one jurisdiction are distributed and consumed across many, with regulatory timestamps required to be coherent across the full chain.

The category interacts substantively with several sectors developed elsewhere in this survey: the implantable and hospital medical devices treatment in clause 10.5 addresses the device-side patient-facing dimension of healthcare; the manufacturing and industrial production treatment in clause 10.17 addresses the pharmaceutical [b-GMP] and [b-GAMP] dimensions at the production-floor level; the logistics and distribution treatment in clause 10.18 addresses pharmaceutical cold-chain at the distribution layer. The present clause is intended to consolidate the regulatory-recordkeeping and clinical-research dimensions that the device, manufacturing, and distribution treatments touch on but do not develop as a category.

NOTE – Further development of this clause requires expertise in biological research, clinical research, pharmaceutical manufacturing and distribution, GxP computerized-systems validation, quality assurance, regulatory affairs and cold-chain logistics. Particular considerations include the regulatory consequences of record-integrity failures, long record-retention periods and operations spanning multiple jurisdictions.

10.4 Electric and connected vehicles

The automotive sector exhibits 2038-class exposure patterns that combine the lifecycle mismatches discussed elsewhere in this paper with sector-specific structural conditions that warrant explicit attention. Modern vehicles are computing platforms with operational lifetimes measured in decades, built from components sourced through deep multi-tier supply chains, with embedded software lifecycles that diverge sharply from the platform lifecycles they support. The electrification of the vehicle fleet, the deployment of supporting charging infrastructure, and the integration of vehicles into broader grid and identity systems amplify these patterns rather than resolving them.

10.4.1 ECU proliferation and supply-chain opacity

Contemporary passenger vehicles ship with substantial numbers of embedded computing modules. Industry observers report that vehicles rolling off production lines today contain on the order of 150 to 300 electronic control units (ECUs), each running its own firmware, frequently sourced from different first-tier suppliers, and in many cases incorporating subcomponents and software libraries from second- and third-tier suppliers whose relationships to the vehicle manufacturer are mediated through layers of contractual abstraction.

The structural consequence is that the vehicle manufacturer's visibility into the time-handling implementations within its own platforms is incomplete by design. A given ECU's firmware may have been developed by a supplier-of-suppliers using toolchains and libraries chosen for reasons of cost, availability, and certification that did not include 2038-class considerations during the design phase. The vehicle manufacturer typically integrates the ECU as a black-box component meeting interface specifications; the internal representation of time within the ECU is not surfaced during the integration review.

10.4.2 Supplier-team lifecycle versus vehicle lifecycle

The mismatch between development-team lifecycles and vehicle lifecycles is particularly acute in the automotive sector. Vehicle platforms commonly remain in service for fifteen to twenty years, with major model variants supported for substantial portions of that period. The supplier development teams responsible for ECU firmware, and the build systems they used to produce shipped binaries, frequently have substantially shorter operational lifetimes than the vehicles into which their work is integrated. Development teams reorganize, dissolve, or move to subsequent generations of work within a year or two of a model launch; build systems are decommissioned or migrated to successor toolchains whose binary-reproducibility relationship to the shipped firmware may not be preserved.

The result is that vehicles in active service today contain firmware whose authoring teams no longer exist, built on toolchains that may no longer be operable, sourced from supplier-of-supplier relationships that may no longer be traceable. Remediation of 2038-class exposure in such firmware is not merely a matter of identifying the issue and producing a patch; it may require reconstructing development environments that have not been maintained since the original shipping date, with the cooperation of organizations whose relevant personnel and institutional memory have moved on.

10.4.3 Charging infrastructure dependencies

The transition to electric vehicles introduces additional time-handling dependencies through the charging infrastructure. Standards governing vehicle-to-grid communication, including [b-ISO 15118] and successors, specify public-key-infrastructure-based authentication mechanisms for charging sessions. These mechanisms depend on certificate validity periods, certificate-revocation

checking, and time-based session integrity in ways structurally similar to the PKI dependencies discussed elsewhere in this paper.

Charging infrastructure additionally depends on coherent time for billing, load management, grid-services coordination, and the cross-operator interoperability that allows vehicles to charge across networks operated by different commercial entities. A 2038-class failure in either the vehicle-side or the charger-side time-handling implementations may manifest as charging-session authentication failure, billing dispute, or in degraded cases the inability to charge at all. The exposure surface is heterogeneous: vehicles from one manufacturer attempting to charge at infrastructure operated by another may exhibit different failure modes than same-manufacturer charging sessions, and the boundary between them is precisely where supply-chain opacity is most consequential.

10.4.4 Navigation and telemetry dependencies

Modern vehicles additionally depend on GNSS receivers for navigation, telematics, and increasingly for the precise positioning required by advanced driver-assistance and autonomous-driving functions. The 2038-class exposure of consumer-grade GNSS receivers is treated separately in clause 10.10; the vehicle-specific consequence is that the long platform lifecycles and supply-chain opacity discussed above apply with particular force to the GNSS receivers embedded in vehicles being manufactured and sold today.

10.4.5 Resource-constrained embedded contexts

Vehicle ECUs frequently operate under resource constraints that complicate remediation. Microcontroller selection in automotive contexts is governed by cost, power consumption, thermal envelope, and certification requirements that have historically favoured smaller, less capable processors than would be deployed in general-purpose computing environments. Remediation strategies that assume processor-level changes (32-bit to 64-bit migration), library updates, or expanded memory footprints may be infeasible within the constraints under which a given ECU was specified and certified.

This interacts with the supply-chain opacity issue in compounding ways. Even where vehicle manufacturers identify 2038-class exposure in specific ECUs, the remediation path may require coordination with suppliers whose contracts and operational capacities are not structured for retrofit work on shipped platforms.

10.4.6 Vehicles: Lifecycle, supply chain, and embedded constraint compounding

The automotive sector combines three of the structural conditions this paper identifies as amplifying 2038-class risk: long platform lifecycles, opaque multi-tier supply chains, and embedded contexts with limited remediation flexibility. The current acceleration of fleet electrification and charging-infrastructure deployment is layering additional time-coordination dependencies on top of these existing conditions, at a pace that has outrun the development of coordinated remediation practices within the sector.

Vehicles being manufactured and sold today, with documented service-life expectations extending well past 2038, are entering operational service with embedded time-handling implementations whose continued correct operation past rollover horizons cannot be asserted by the manufacturers selling them. Charging infrastructure being deployed in parallel is depending on time-handling assumptions that have not been systematically evaluated against rollover scenarios. This structural exposure invites coordinated attention from manufacturers, charging-network operators, regulators, and standards bodies concerned with the longevity and interoperability of the deployed vehicle fleet.

10.5 Implantable and hospital medical devices

Medical devices represent a 2038-class exposure class with structural conditions distinct from those of other sectors examined in this Technical Paper. The patient-facing nature of medical device deployment, the regulatory regimes governing their lifecycle, the embedded operational lifetimes of

implanted devices, and the role of patient confidence in clinical decision-making combine to produce remediation considerations that warrant explicit attention.

This clause addresses two related but distinct subcategories: implantable medical devices, where individual patients carry embedded computing for years or decades, and hospital infrastructure, where institutional deployment of medical computing systems creates aggregate exposure across patient populations. The medical-device cybersecurity literature has been active for over a decade and has surfaced multiple categories of risk; the present treatment concerns time-handling exposure specifically, which has received comparatively less attention.

10.5.1 Implantable medical devices

Modern implantable medical devices – pacemakers, implantable cardioverter-defibrillators, neurostimulators, cochlear implants, insulin pumps, and successor categories under active development – contain embedded computing components that operate continuously inside patients for operational lifetimes measured in years or, increasingly, decades. Battery life and physical replacement cycles vary by device class, but a substantial fraction of currently implanted devices will remain in service past 2038. The embedded firmware in these devices typically derives from supply chains and toolchains whose 2038-class exposure was not tested during the original certification process.

Regulatory regimes governing medical device approval and post-market surveillance vary across jurisdictions but share structural properties consequential for 2038-class remediation. The United States Food and Drug Administration regulates devices through premarket approval pathways and post-market cybersecurity guidance updated in 2023. The European Union's Medical Device Regulation establishes equivalent obligations within the EU and EEA. Health Canada, Japan's Pharmaceuticals and Medical Devices Agency, the United Kingdom's Medicines and Healthcare products Regulatory Agency, and analogous bodies in other jurisdictions maintain their own regulatory frameworks. The common structural property is that any modification to a deployed medical device – including firmware updates addressing time-handling exposure – requires regulatory engagement that may extend beyond the timeline available before rollover horizons.

A consequence does not present in other sectors is that the patient is the platform. Remediation of 2038-class exposure in an implanted device cannot be performed without clinical engagement: the patient must present at a clinical facility, the device must be communicated with through manufacturer-specific programming equipment, and the patient's individual clinical status must be considered in any modification to a device on which their life or wellbeing may depend.

10.5.2 Cross-border patient mobility and device labelling

Patient mobility across jurisdictions introduces a coordination problem distinctive to medical devices. A patient implanted with a device in one jurisdiction may travel to another for work, family, or care reasons, and may require interaction with their device – interrogation, reprogramming, battery-status assessment, or rollover-related firmware update – in a jurisdiction other than the one where the device was approved. The regulatory regime governing such interaction is structurally ambiguous: the approving authority is the device's home regulator, but the clinical event takes place under the supervising authority of the patient's current location.

The 2038-class transition introduces a specific manifestation of this general problem. A patient whose device has been updated to address rollover exposure may, on travelling, encounter clinical equipment in another jurisdiction that has not been similarly updated, or that is uncertain about compatibility with the updated firmware. The clinical question may not be substantive – the device may function correctly across the equipment in question – but the question being unanswerable in the moment may be sufficient to alter clinical decision-making. A patient uncertain whether interrogation of their device by available equipment is safe may defer or refuse necessary clinical interaction. The failure mode is not technical; it is the propagation of unresolved uncertainty into clinical hesitation.

A clear labelling scheme indicating 2038-class remediation status of both devices and clinical equipment, with cross-jurisdictional recognition, may be valuable in preserving clinical confidence during the transition period. The technical content of such a scheme – what fields are conveyed, how compatibility is asserted, how cross-border recognition is established – would need to be developed in consultation with relevant medical device regulatory authorities and clinical professional bodies. The structural observation surfaced here is that the patient-confidence dimension of medical device remediation introduces requirements absent in sectors where remediation is invisible to end users.

10.5.3 Hospital infrastructure

Beyond implanted devices, hospital infrastructure contains substantial deployed computing whose 2038-class exposure follows the broader patterns documented in this paper. Patient admission and registration systems, picture archiving and communication systems, electronic health record platforms, infusion pumps, patient monitoring networks, laboratory information systems, and surgical robotics all depend on coherent time for record ordering, audit trail integrity, drug-administration timing, and the coordination of distributed clinical workflows. Foundational systems such as patient admission and registration occupy a particular structural position: their failure can produce hospital-wide operational interruption even when other clinical systems remain technically functional, with consequences disproportionate to the apparent simplicity of the systems themselves. The lifecycle of hospital information systems is typically shorter than that of implanted devices but longer than that of consumer computing infrastructure, with major platforms commonly in service for ten to fifteen years. The 2038-class exposure of foundational hospital-information systems is consequently load-bearing in ways that the enumeration of clinical and diagnostic systems above does not by itself surface.

The remediation pathways available for hospital infrastructure differ from those for implanted devices. Hospital systems can be patched, replaced, and reconfigured on operational timelines that implanted devices cannot match, but the operational constraint of continuous clinical service introduces its own coordination problem: maintenance windows in clinical environments are short, scheduling is constrained by clinical workflow, and the consequences of time-handling failures may manifest as patient-safety events rather than as administrative inconvenience.

10.5.4 Medical devices: Patient confidence and cross-jurisdictional coordination

Medical devices are distinctive among the use cases in this Technical Paper because the consequences of time-handling failure manifest at the boundary between technical infrastructure and individual patient outcomes. The remediation pathway includes regulatory engagement, clinical engagement with individual patients, cross-jurisdictional coordination not required in other sectors, and attention to patient confidence as a factor in clinical decision-making. The combination of long device lifetimes, heavy regulatory oversight, patient-individual remediation requirements, and the propagation of uncertainty into clinical hesitation produces a remediation profile that calls for sustained engagement from medical device manufacturers, clinical professional bodies, regulatory authorities across jurisdictions, and standards bodies concerned with the lifecycle of patient-facing computing systems.

The structural observations made elsewhere in this paper concerning supply-chain opacity, lifecycle mismatch, and remediation complexity apply with particular weight in the medical device sector. The labelling and cross-jurisdictional coordination considerations surfaced above represent specific dimensions of the broader coordination problem this paper addresses, manifesting in a sector where the human consequences of unresolved uncertainty are immediate.

10.6 CSIRT/PSIRT/CERT coordination

Incident response coordination teams – including national computer emergency response teams (CERTs), sectoral computer security incident response teams (CSIRTs), and vendor product security incident response teams (PSIRTs) – are positioned at the operational coordination layer where

2038-class incidents will surface, propagate across sectors, and require multi-stakeholder response. The structural distinctiveness of the CSIRT/PSIRT/CERT exposure class lies not in the technical systems they operate, but in the cross-sector and cross-border coordination function they perform.

National-level CERT/CSIRT coordination depends on shared situational awareness across sectoral boundaries, consistent incident classification across jurisdictions, and the ability to correlate observations from multiple operators against a coherent timeline. A 2038-class event affecting multiple sectors simultaneously – as the analysis throughout this paper suggests is structurally likely – will test these coordination mechanisms in ways for which they have not been explicitly designed.

NOTE – Engagement invited from CSIRT/PSIRT/CERT practitioners – national CERT coordinators, sectoral CSIRT operators, and FIRST-member incident response teams. This clause is structurally important given the operational role these communities will hold during a coordinated 2038-class event.

10.7 Emergency services operations

Police, fire, emergency medical services, and integrated dispatch infrastructure depend on accurate time across multiple coordination layers. Computer-aided dispatch (CAD) systems timestamp call receipt, assignment, response, and clearance – data that is both operationally load-bearing and legally consequential, since dispatch records routinely appear in inquiry proceedings and civil litigation. Emergency-radio networks rely on synchronized time for trunked-channel coordination, talkgroup management, and cross-jurisdictional interoperability during major incidents. Public-safety answering points (PSAPs) processing emergency calls timestamp every interaction in audit logs that survive for years and that propagate into mutual-aid arrangements, insurance claims, and after-action reviews. The 2038-class exposure here is the cross-jurisdictional coordination case: when neighbouring PSAPs, regional fire commands, or mutual-aid partners disagree on time at the protocol level, the operational degradation is silent until the moment of major-incident response, when it is too late to discover.

NOTE – Engagement invited from PSAP operators, computer-aided dispatch system vendors, emergency-radio network operators, and the public-safety communications standards community. This clause is structurally important given the operational role these communities will hold during a coordinated 2038-class event and the legal weight that PSAP and CAD timestamping carries in post-incident review.

10.8 ICS/OT environments

Industrial control systems (ICS) and operational technology (OT) environments exhibit 2038-class exposure with structural conditions distinctive from those of information technology systems. Long operational lifecycles, safety-critical certification regimes, deep multi-tier supplier relationships, and the operational impossibility of arbitrary system updates combine to produce remediation profiles that warrant treatment separately from the broader IT exposure addressed elsewhere in this paper.

ICS/OT environments commonly incorporate programmable logic controllers, supervisory control and data acquisition systems, distributed control systems, safety instrumented systems, and field instrumentation with operational lifetimes commonly measured in decades. The 2038-class exposure manifests across the timestamping behaviour of these systems, the synchronization infrastructure connecting them, and the engineering workstations and historians used to operate them.

NOTE – Engagement invited from process control engineers, ICS-specific incident response practitioners, and safety certification specialists. This clause is structurally important given the safety-critical nature of these environments and the operational constraints on remediation pathways.

10.9 Cloud and ICT infrastructure

Cloud computing platforms, hosting infrastructure, identity and access management services, observability platforms, and the broader ICT substrate on which sector-specific applications are deployed exhibit 2038-class exposure with structural conditions distinctive from the telecommunication infrastructure addressed in clause 10.1. Where telecommunication networks carry

the delivery substrate connecting endpoints, ICT infrastructure carries the compute and platform substrate on which the applications served across those networks actually run.

Distinctive structural conditions include the concentration of cross-sector workloads on a small number of hyperscale cloud providers (which produces correlated exposure across sectors using shared platforms), the multi-tenancy assumptions underlying cloud platform architecture (which propagate platform-layer time-handling decisions across tenant workloads regardless of tenant awareness), the long lifecycles of enterprise identity infrastructure deployed on hybrid cloud architectures, and the increasing dependency of sector-specific applications on cloud-native observability, monitoring, and alerting infrastructure whose 2038-class exposure may not be visible to the application owners depending on it.

The structural observation parallels that made for telecommunications in clause 10.1: ICT infrastructure is simultaneously a sector with its own structural exposure and the substrate on which most other sectors depend.

NOTE – Engagement invited from cloud platform engineers, enterprise identity specialists, observability and platform reliability practitioners, and cloud-sector cybersecurity practitioners. This clause is structurally important given the cross-cutting role of cloud and ICT infrastructure in supporting most other sectors addressed in this paper.

10.10 GNSS-dependent services

Global navigation satellite systems (GNSS) provide a substrate layer of positioning and timing services on which multiple sectors addressed elsewhere in this Technical Paper depend. The receiver-side time-handling exposure of consumer-grade and commercial-grade GNSS modules is distinct from the transmitter-side and ground-segment exposure addressed in clause 10.11, warranting separate treatment.

A structural property distinctive to GNSS within the broader landscape of timing-dependent services warrants explicit acknowledgement: timing is not metadata in GNSS-derived positioning but is the ranging primitive from which position is computed. Rollover or ambiguity in the time substrate propagates directly into position and velocity outputs, with the consequence that the 2038-class exposure of GNSS receiver populations operates at the layer of the service's primary output rather than at a peripheral metadata layer.

Distinctive structural conditions include the 2019 GPS week-number rollover incident affecting aircrafts, as documented in clause 7.1, the deployment scale of consumer-grade GNSS receivers across automotive, telecommunications, financial, and agricultural sectors, and the lifecycle characteristics of GNSS receiver firmware which is commonly deployed and forgotten rather than systematically maintained across the operational lifetimes of the platforms incorporating it. See also clauses 10.1, 10.2, 10.4 and 10.21.

NOTE – Engagement invited from GNSS receiver design engineers, GNSS service operators, and downstream-sector specialists whose operations depend on GNSS-derived timing. This clause is structurally important given the substrate role of GNSS in supporting multiple other sectors addressed in this paper.

10.11 Space infrastructure

Space-based infrastructure has become a cross-cutting substrate for time-coordination dependencies across multiple sectors examined elsewhere in this Technical Paper. Satellite communications systems, satellite earth stations, ground-segment timing distribution, and constellation-operations infrastructure all depend on coherent time for both intra-system coordination and the services they provide to downstream consumers. The transmitter side of GNSS – distinct from the receiver-side dependencies treated in clause 10.10 – is part of this layer, as is the ground segment supporting non-GNSS satellite operations.

The 2038-class exposure of space infrastructure includes onboard satellite systems whose operational lifetimes commonly extend across multiple decades, ground-segment time distribution infrastructure

whose remediation is constrained by cross-jurisdictional operational arrangements, and constellation-coordination systems whose continued correct operation depends on time-handling implementations across both satellite and ground components. Since the Hayes taxonomy was originally articulated, commercial space activity has accelerated markedly, with large constellation operators deploying satellites at unprecedented scale alongside a broader diversification of orbital service providers, substantially expanding the deployed footprint of time-dependent space infrastructure.

A failure mode distinctive to space infrastructure is the potential for time-handling failures to contribute to orbital debris generation. Satellite attitude control, orbit maintenance, conjunction assessment, end-of-life disposal, and constellation-scale deconfliction all depend on coherent time across onboard and ground-segment systems. A 2038-class failure affecting these dependencies may produce missed collision avoidance manoeuvres, uncontrolled disposal, or attitude and orbit control errors whose consequence is the creation of debris in operationally important orbits. Unlike many failure modes treated elsewhere in this Technical Paper, orbital debris is structurally non-remediable: debris created in operational orbits is not recoverable and may render orbital regions unusable for time horizons measured in decades. The Kessler-syndrome dynamics documented in the [b-WTISD26] expert report – wherein collisions between orbital objects generate additional debris that increases collision probability for remaining objects – describe a cascade mechanism that 2038-class failures could contribute to – or initiate.

Space infrastructure exposure to 2038-class events shares structural conditions with its exposure to space-weather events of the kind documented in the [b-WTISD26] expert report: substrate failures cascading through invisible dependencies into sectors whose operators may not have visibility into the failure source. The structural pattern is consistent with that addressed throughout this Technical Paper; the specific stressors differ, but the cascade dynamics and remediation challenges are structurally similar.

NOTE – Detailed treatment of space infrastructure dependencies, including specific operator-side and ground-segment exposure classes, to be developed in subsequent versions of this Technical Paper.

10.12 Broadcasting and emergency alerting

Broadcast media infrastructure – radio, television, streaming distribution, and content delivery networks – depends on accurate time across multiple operational layers, from programme scheduling and ad-insertion synchronization to simulcast coordination across geographically separated transmitter sites. Modern digital broadcasting standards including DVB-T2, ATSC 3.0, and DAB+ embed precise timing requirements that propagate from broadcast operations centres through terrestrial transmitter networks. The emergency-alerting subset deserves particular attention: [b-EAS], [b-ETSI TS 102 900] and equivalent national systems, and cell broadcast Service used for public warning all carry strict latency and timestamp requirements that, when violated, undermine both operational effectiveness and forensic reconstruction of warning chains after major incidents. The 2038-class exposure is dual: scheduled-broadcast disruption is a visible operational failure that broadcasters would detect and remediate; alert-timestamp integrity is the more consequential exposure because emergency-warning chains are subject to legal review and depend on cross-system time consistency that is rarely tested under stress.

NOTE – Engagement invited from broadcast operations engineers, DVB and ATSC standards bodies, emergency alert system (EAS) and cell broadcast operators, and national public-warning authorities. This clause is structurally important given the public-safety consequences of warning-chain timestamp degradation and the legal review to which emergency-alerting infrastructure is routinely subjected.

10.13 Government services and public administration

Government information systems carry some of the most consequential time-dependencies in civil infrastructure. Civil registries – recording births, marriages, deaths, citizenship, and residency – produce timestamped records that establish legal identity and underpin essentially every downstream

civic interaction; corruption of these timestamps cascades into pensions, inheritance, immigration, and voting eligibility. Court records and judicial filings depend on precise timestamping for procedural validity, statutes of limitation, and chain-of-custody determinations. Tax administration systems use timestamps to establish filing compliance, assessment periods, and audit windows. Cadastral records – the registries of land ownership, parcels, boundaries, and encumbrances – represent a particularly load-bearing case: the legal validity of property rights depends on the integrity of timestamped registry entries that may extend back decades and that, if corrupted, become extraordinarily difficult to reconstruct. Elections systems, immigration databases, and social-benefits administration similarly carry timestamps with downstream legal force. The 2038-class exposure across government services is that timestamp corruption in any single registry can propagate into legal proceedings, benefit eligibility determinations, and identity-verification chains years or decades after the originating error.

NOTE – Engagement invited from civil-registry administrators, court records technologists, tax administration system operators, cadastral and land-registry authorities, elections systems engineers, and the broader e-government standards community. This clause is structurally important given the legal force that government timestamps carry across the full range of downstream civic and economic interactions.

10.14 Insurance and reinsurance

The insurance and reinsurance industries occupy a distinctive position with respect to 2038-class risk because their core function – pricing and capital-allocating against tail risk – depends on accurate time across both operational and modeling layers. Operationally, claims-processing systems timestamp incidents, notifications, payments, and reserves; corruption of these timestamps creates regulatory exposure under solvency frameworks that require auditable claims chronology. Actuarial systems and catastrophe-modeling platforms rely on time-series data spanning decades and on real-time market and meteorological feeds for pricing and capital management. Reinsurance treaties – through which primary insurers transfer correlated tail risk to global reinsurers – carry detailed temporal specifications for event-occurrence, notification, and aggregation periods, and disputes over treaty interpretation routinely turn on timestamp evidence. The 2038-class exposure here is structurally important to the broader analysis of this Technical Paper: insurance and reinsurance markets are the principal vehicle through which the financial consequences of correlated infrastructure failures are absorbed across the economy, and if reinsurance models do not adequately price 2038-class risk before the event window opens, the post-event capital response across all other sectors will be constrained by the limits of insurance-industry solvency.

NOTE – Engagement invited from underwriters, claims-systems technologists, actuarial and catastrophe-modeling specialists, reinsurance treaty technical staff, and the broader insurance-supervisory community. This clause is structurally important given the role insurance markets play in absorbing the financial consequences of correlated infrastructure failures and the dependence of reinsurance pricing models on accurate forward-looking risk assessment.

10.15 Epistemic coordination layer

The technical exposure classes addressed elsewhere in this Technical Paper depend on a substrate that is rarely named explicitly in infrastructure taxonomies: the epistemic coordination layer comprising the human knowledge, institutional memory, build provenance, documentation chains, training data, and shared technical understanding without which the underlying technical infrastructure cannot be maintained, remediated, or correctly described.

This layer is not a sector in the conventional sense; it is the substrate of capacity-to-know on which every other sector's remediation work depends. Its 2038-class exposure is distinct from, and compounding with, the technical exposures addressed throughout this paper.

Four structural conditions warrant explicit identification:

1) The cohort of engineers who designed and built the deployed 32-bit infrastructure is reaching retirement during precisely the period in which their knowledge is most needed for

remediation. Systems whose internal time-handling assumptions are not fully documented, build environments operable by a small number of individuals, and undocumented operational dependencies maintained through institutional habit all become structurally fragile as the individuals carrying that knowledge exit the workforce. The remediation pathway for many 2038-class exposures depends on knowledge that is not being transferred at the rate at which it is being lost.

2) Mergers, acquisitions, divestitures, outsourcing arrangements, and supplier consolidations produce situations in which the organization currently responsible for a deployed system has no continuity with the organization that built it. Build systems decommissioned during integration, documentation lost in migrations, personnel reassigned or made redundant, and contractual relationships substituted for in-house engineering all produce structural loss of institutional memory. The systems remain in service; the capacity to understand and modify them does not. This pattern is not incidental to commercial activity but structurally produced by it, and the rate of loss has accelerated as outsourcing and consolidation have intensified across the sectors addressed in this paper.

3) Technical documentation and provenance chains decay across the lifecycles of deployed systems. Documentation written under one set of assumptions becomes increasingly difficult to interpret as those assumptions cease to be common knowledge. Architectural records describing systems that have been refactored multiple times no longer match the systems they purport to describe. The chain of design intent connecting current system behaviour to original assumptions becomes impossible to reconstruct for long-lived systems whose documentation has not been actively maintained as a first-class artefact.

4) Structurally novel relative to prior epoch transitions such as Y2K – the increasing role of AI-mediated knowledge production introduces a recursive risk to the epistemic substrate itself. As public attention to 2038-class risks grows in the coming decade, substantial volumes of variable-quality content will be produced about the subject. Some of this content will reflect careful technical analysis; much will reflect speculation, commercial framing, panic-driven simplification, or factual error introduced through retransmission. AI systems trained on this content will incorporate the resulting mixture without quality weighting, producing downstream outputs that reflect the contaminated training distribution rather than the underlying technical reality. The recursive dynamic – content production feeding training data feeding output production feeding further content production – operates precisely during the period in which authoritative technical guidance is most needed and threatens to degrade the very knowledge substrate on which informed remediation depends. The cumulative effect of these four conditions is a structural condition for which the term infrastructural vestigialisation, developed in clause 8.3, names the failure mode: deployed substrate continuing to function while the capacity-to-rebuild atrophies below the threshold a coordinated rebuild would require.

NOTE – Engagement invited from domain experts in knowledge management, institutional memory and engineering documentation, technical-debt accounting, training data integrity, and the broader study of epistemic infrastructure. This clause is structurally important given the cross-cutting role of this substrate in supporting the remediation pathways for every other sector addressed in this paper.

10.16 Buildings and built environment

Building management system (BMS) and building automation system (BAS) coordinate heating, ventilation, air-conditioning, lighting, elevators, access control, fire-suppression, and surveillance across modern commercial, institutional, and increasingly residential buildings. These systems depend on accurate time at multiple operational layers: access-control logs and surveillance video timestamps carry both security and legal weight, fire-suppression sequencing relies on synchronized time across detector and actuator networks, elevator dispatch algorithms coordinate across multiple cars and floors, and HVAC scheduling integrates with utility demand-response programmes that

require synchronized operation across thousands of buildings simultaneously. The integration of BMS infrastructure with cloud-based facilities-management platforms and with broader ICS/OT environments has expanded the attack surface significantly over the past decade, as addressed in clause 10.8. The 2038-class exposure is heterogeneous: some failures (HVAC scheduling drift) are operationally tolerable, while others (fire-suppression sequencing, surveillance-video timestamp integrity in security incidents) are consequential. The buildings sector is also notable as a substrate category – essentially all other sectors operate within buildings, and BMS failures can propagate into the operational continuity of every tenant.

NOTE – Engagement invited from BMS and BAS engineers, facilities-management technologists, fire-and-life-safety systems integrators, access-control and physical-security vendors, and the broader smart-buildings standards community (including BACnet, KNX, and OPC UA stakeholders). This clause is structurally important given the substrate role buildings play across essentially every other sector and the legal and life-safety consequences of access-control, surveillance, and fire-suppression timestamp failures.

10.17 Manufacturing and industrial production

Industrial production – automotive assembly, semiconductor fabrication, pharmaceutical manufacturing, food and beverage processing, chemical production, and the broader landscape of factories that produce the physical goods supporting all other sectors – depends on time-coordinated operation at multiple layers. Programmable logic controllers (PLCs) and distributed control systems (DCS) on the production floor require synchronized time for sequencing operations, coordinating multi-station assembly lines, and producing the timestamped event records that feed into manufacturing execution systems (MES) and enterprise resource planning (ERP) integrations. Quality-control timestamps establish batch traceability obligations under regulatory frameworks including [b-ISO 9001], automotive [b-IATF 16949], pharmaceutical [b-GMP] and [b-GAMP], and food-safety [b-HACCP] requirements; pharmaceutical batch records in particular face strict regulatory scrutiny in which timestamp integrity is foundational to product release and to post-market surveillance, and cold-chain logistics for temperature-sensitive biologics carry analogous regulatory weight. Semiconductor fabrication processes, which can run for months across hundreds of process steps, are particularly time-sensitive at the wafer-level metrology layer. The 2038-class exposure across manufacturing is both operational (production interruption from time-coordination failures) and regulatory (batch-record integrity for safety-critical products).

NOTE – Engagement invited from PLC and DCS vendors, MES and ERP technologists, quality-systems engineers, pharmaceutical and semiconductor manufacturing specialists, and the broader industrial automation standards community (including [b-IEC 62264] and [b-IEC 62443] stakeholders). This clause is structurally important given the regulatory consequences of batch-record timestamp corruption for safety-critical products and the role of industrial production as a substrate dependency for essentially every other sector's physical-goods supply.

10.18 Logistics and distribution

Logistics and distribution networks – postal services, parcel carriers, warehouse automation systems, last-mile delivery platforms, customs and border-control processing, and retail-distribution endpoints including fuel-station automatic tank gauges (ATGs) [b-ICSA-25-261-07] and [b-ICSA-25-296-03] and point-of-sale terminals – depend on synchronized time across cross-border operations spanning many jurisdictions and protocol generations. Supply-chain timestamps establish chain-of-custody, customs-clearance windows, cold-chain integrity verification under [b-GDP], particularly for pharmaceuticals, vaccines, and food-safety, expiration-date and shelf-life controls, and the legal records on which insurance and liability determinations rest. Warehouse automation increasingly relies on real-time-location systems (RTLS) and coordinated robotic operations that require precise time synchronization. The retail-distribution tail – point-of-sale terminals at fuel stations, payment-card authorization chains, and the automatic tank gauges that monitor underground fuel storage – represents a particularly heterogeneous deployment surface, with mixed legacy and modern equipment operating across decades of vintage. The 2038-class exposure across logistics is

structurally significant because logistics is the integrating layer between manufacturing, regulated distribution, retail, and consumption; cascading time-coordination failures here propagate into every downstream sector and may undermine [b-GDP] evidence, expiry-date handling, recall execution, and downstream quality or liability determinations.

NOTE – Engagement invited from postal and parcel-network operators, warehouse-automation vendors, customs and border-control technologists, cold-chain logistics specialists and [b-GDP] specialists, pharmaceutical and vaccine distribution operators, retail point-of-sale system operators, and fuel-distribution infrastructure engineers. This clause is structurally important given the role logistics plays as the integrating layer between manufacturing, retail, and consumption across cross-jurisdictional supply chains.

10.19 Mining and extraction

Mining and extraction operations incorporate embedded control systems across extraction equipment, processing facilities, ore-transport infrastructure, and safety monitoring systems. Long equipment lifecycles, remote and harsh operating environments, and heterogeneous SCADA and PLC deployments across mixed vendors, generations, and communication protocols produce 2038-class exposure conditions consistent with the broader ICS/OT pattern addressed in clause 10.8, with additional sector-specific considerations.

Distinctive structural conditions include the operational remoteness of many mining and extraction sites, which complicates remediation logistics; the safety-critical role of timestamping in fatigue analysis, maintenance scheduling, and incident reconstruction; and the increasingly automated nature of extraction operations, where time-coordination failures may produce safety and environmental consequences directly.

NOTE – Engagement invited from mine automation engineers, extraction industry safety specialists, and operators of remote SCADA deployments.

10.20 Water and wastewater

Water and wastewater infrastructure represents critical public-health infrastructure with 2038-class exposure across dam and reservoir management, pumping stations, treatment facilities, distribution networks, and the SCADA systems that coordinate them. The sector is distinguished by municipal operator scale, public-health criticality, and asset lifetimes commonly measured in decades.

Distinctive structural conditions include the public-health consequences of operational failure (which constrain maintenance windows and elevate the consequences of unintended downtime), the heterogeneous operator landscape (with municipal utilities operating at substantially smaller scale and resource depth than transmission-grid operators or major telecommunication carriers), and the integration of cybersecurity considerations into a sector whose traditional engineering practice has been physical-systems-focused.

NOTE – Engagement invited from municipal utility operators, water-sector ICS specialists, and public-health authorities with infrastructure oversight responsibilities. This clause is structurally important given the public-health implications of water-sector service continuity.

10.21 Food and farming

Food production and agricultural systems exhibit 2038-class exposure across precision agriculture timing systems, GNSS-dependent positioning for autonomous and semi-autonomous equipment, cold-chain logistics monitoring, and processing-line automation. The sector is distinguished by seasonal operational windows, biological perishability, food safety, supply-chain traceability requirements, and the increasing penetration of digital systems into operations historically governed by manual practice.

Distinctive structural conditions include the dependency of modern precision agriculture on GNSS-derived positioning and timing, addressed in clause 10.10, the food-safety implications of

cold-chain timestamping failures, and the seasonal nature of many agricultural operations, which produces operational windows in which remediation is feasible and other windows in which it is not.

NOTE – Engagement invited from precision agriculture engineers, food-safety regulators, and cold-chain logistics operators.

10.22 Oil and gas

Oil and gas infrastructure incorporates pipeline control systems, refinery automation, custody-transfer metering, and the SCADA infrastructure coordinating extraction, transport, and refining operations. The sector is distinguished by multi-decade infrastructure lifecycles, cross-jurisdictional pipeline networks, and safety-critical certification regimes governing equipment modifications.

Distinctive structural conditions include the cross-border nature of major pipeline systems (where coordination across operators and jurisdictions is a precondition for safe operation), custody-transfer measurement (where timestamping integrity has direct commercial and contractual significance), and the safety-critical role of timing in process control systems whose failure modes can produce environmental and public-safety consequences.

NOTE – Engagement invited from pipeline operations engineers, custody-transfer measurement specialists, and refinery automation practitioners.

10.23 Power plants

Electrical power generation facilities incorporate generation control systems, protection relays, plant historians, and the operational technology infrastructure coordinating safe and efficient generation across diverse fuel sources and generation technologies. The sector is distinguished by multi-decade lifecycles, safety-critical certification regimes (notably [b-IEC 61850] and [b-IEC 62443]), and the cross-cutting role of generation in supporting the broader electrical grid addressed in clause 10.24.

Distinctive structural conditions include the safety-critical role of timestamping in protection relay coordination, the integration of generation control systems with the wider grid synchronization infrastructure, and the operational consequences of unexpected generation-side failures for grid stability and downstream consumer service.

NOTE – Engagement invited from generation control engineers, protection relay specialists, and operators of large-scale generation facilities.

10.24 Power grid

Electrical transmission and distribution grids depend on coherent time at multiple scales – from the seconds-to-minutes timescale of dispatching and market operations to the microseconds-to-milliseconds timescale of phasor measurement, protection coordination, and synchrophasor-based wide-area monitoring. The sector is distinguished by cross-jurisdictional interconnection, the safety-critical and economic significance of timing precision, and the operational impossibility of taking transmission infrastructure offline for arbitrary remediation.

Distinctive structural conditions include the dependency of modern grid operations on GNSS-derived timing for synchrophasor measurement (with downstream propagation through wide-area monitoring and protection systems), the cross-border nature of major interconnections (which produces coordination requirements analogous to those addressed for financial systems in clause 10.2 and telecommunications in clause 10.1, and the increasing dependency on coherent time as renewable generation and distributed energy resources expand the operational complexity of grid coordination.

NOTE – Engagement invited from transmission system operators, distribution utility engineers, synchrophasor specialists, and grid-cybersecurity practitioners. This clause is structurally important given the cross-cutting role of grid infrastructure in supporting most other sectors addressed in this paper.

10.25 Vehicles and traffic infrastructure

Vehicles and traffic infrastructure constitute the civil-infrastructure layer of road transportation: traffic management systems, signal coordination networks, tolling infrastructure, roadside V2X units, congestion management sensors, and the increasingly digitalized supervisory systems that coordinate vehicle movement across municipal and regional road networks. The 2038-class exposure of this layer is distinct from the vehicle-side exposure addressed in clause 10.4, though the two are operationally entangled.

Distinctive structural conditions include the municipal and transport-authority operator landscape (which operates at substantially smaller scale and resource depth than the vehicle-manufacturer supply chains addressed in clause 10.4, the long lifecycles of deployed civil infrastructure (where traffic-signal controllers, tolling gantries, and roadside sensors commonly remain in service for fifteen to twenty-five years), the increasing dependency on roadside V2X infrastructure as autonomous and semi-autonomous driving deployments mature, and the integration of legacy electromechanical traffic-control systems with modern digital supervisory platforms. The boundary between vehicle-side and infrastructure-side systems – where vehicles communicate with roadside infrastructure for navigation, safety, tolling, and traffic coordination – is precisely where 2038-class exposure across the two domains compounds. The integration of traffic-management systems into broader smart-city platforms further amplifies the cross-sector consequences of 2038-class failures in this layer.

NOTE – Engagement invited from municipal transport engineers, intelligent transportation systems specialists, V2X infrastructure operators, and tolling-system operators. This clause is structurally important given the integration of civil road infrastructure with the vehicle-side systems addressed in clause 10.4 and the cross-sector consequences of traffic-coordination failures.

10.26 Railroad

Railway infrastructure incorporates signalling systems (e.g., ERTMS), positive train control, train-to-ground communications over GSM-R/FRMCS, scheduling and dispatching systems, and the operational technology infrastructure coordinating safe operations across multi-decade asset lifetimes. The sector is distinguished by safety-critical certification regimes, cross-operator interoperability requirements (particularly in jurisdictions with multiple operators sharing infrastructure), and the integration of legacy and modern signalling systems within shared operational environments.

Distinctive structural conditions include the safety-critical role of timestamping in train control and signalling coordination, the cross-jurisdictional nature of major rail networks (where rolling stock and operational coordination cross borders), and the long lifecycle of signalling infrastructure whose modification requires extensive certification and operational planning.

NOTE – Engagement invited from signalling engineers, positive train control specialists, and railway cybersecurity practitioners.

10.27 Bridges and tunnels

Major bridges and tunnels increasingly incorporate digital infrastructure including structural monitoring sensors, traffic management systems, ventilation control, and emergency response coordination. The sector is distinguished by long infrastructure lifecycles (frequently exceeding the lifecycles of the digital systems instrumenting them), instrumentation heterogeneity across operators and jurisdictions, and the safety-critical role of monitoring systems in structures whose failure consequences are catastrophic.

Distinctive structural conditions include the integration of digital monitoring and control systems with infrastructure designed under earlier engineering paradigms, the long-term archival significance of monitoring data for structural integrity assessment, and the operational dependency of major

transport corridors on the continued correct operation of tunnel ventilation and traffic management systems.

NOTE – Engagement invited from structural monitoring specialists, tunnel operations engineers, and transport-infrastructure cybersecurity practitioners.

10.28 Aviation

Aviation exhibits 2038-class exposure across avionics, air traffic management, GNSS-dependent navigation, and the maintenance and certification records governing airworthiness across multi-decade airframe lifecycles. The sector is distinguished by safety-critical certification regimes that constrain remediation timing, cross-jurisdictional operational coordination (involving ICAO, regional air navigation service providers, and individual operators), and documented prior incidents establishing empirical reality – notably the 2019 GPS week-number rollover referenced in clause 7.1.

Distinctive structural conditions include the long operational lifetimes of airframes and avionics (frequently exceeding the lifecycles of the supporting ground infrastructure), the cross-jurisdictional nature of international air operations, and the safety-critical certification overhead that constrains the remediation pathways available to operators and manufacturers.

NOTE – Engagement invited from avionics engineers, air traffic management specialists, aviation cybersecurity practitioners, and aviation regulatory authorities. This clause is structurally important given the documented prior incidents in the sector and the safety-critical consequences of timing failures in operational aviation.

10.29 Shipping

Maritime shipping incorporates navigation systems, port operations, container tracking, and the cross-jurisdictional operational coordination characteristic of an inherently international sector. GNSS-derived positioning is foundational to modern maritime operations, with downstream dependencies addressed in clause 10.10, and long vessel lifecycles produce structural exposure conditions analogous to aviation above.

Distinctive structural conditions include the cross-jurisdictional nature of international shipping (where vessels operate under one flag state's regulations while trading into other jurisdictions' ports), the long operational lifetimes of commercial vessels, and the increasing dependency of port operations on digital infrastructure whose 2038-class exposure may not have been systematically evaluated.

NOTE – Engagement invited from vessel operations specialists, port automation engineers, and maritime cybersecurity practitioners.

10.30 Wastes and recycling

Waste management and recycling operations incorporate processing-facility automation, materials-tracking systems, and the operational technology coordinating collection, sorting, processing, and disposal across municipal and commercial waste streams. The sector is distinguished by operational-technology heterogeneity, environmental compliance timestamping requirements, and the integration of digital tracking into operations historically governed by manual practice.

Distinctive structural conditions include the regulatory significance of timestamping in environmental compliance reporting (where audit trails are consequential for regulatory enforcement), the heterogeneous operator landscape (spanning municipal authorities, commercial waste operators, and recycling processors), and the increasing automation of sorting and processing facilities whose 2038-class exposure may compound with the lifecycle mismatches affecting other ICS/OT environments.

NOTE – Engagement invited from municipal waste operators, recycling facility automation engineers, and environmental compliance specialists.

11 Threat model

NOTE – The threat model presented here outlines the principal threat categories relevant to 2038-class risk. The sub-clauses below will be developed in subsequent versions in coordination with [b-FIRST Time SIG] workstreams. The treatment will be calibrated to inform defensive engagement without providing material uplift to adversarial planning; where this calibration is in tension with the document's broader commitment to structural transparency, the calibration is resolved in favour of defensive utility. The conditions under which this calibration trade-off becomes operationally consequential, and the principles by which it is resolved in specific cases, will be developed alongside the threat-category treatments.

11.1 Threat categories

11.1.1 Time desynchronization scenarios

NOTE – To be developed in subsequent versions.

11.1.2 GPS/NTP/NITZ spoofing

NOTE – To be developed in subsequent versions.

11.1.3 Rollover-induced failures

NOTE – To be developed in subsequent versions.

11.1.4 Protocol chain reactions

NOTE – To be developed in subsequent versions.

11.1.5 Silent failure modes

NOTE – To be developed in subsequent versions.

12 Research landscape and empirical observations

12.1 Publicly reported field observations

NOTE – To be developed in subsequent versions.

12.2 Academic and conference literature

Two recent empirical studies have characterized the deployed NTP ecosystem in ways relevant to 2038-class risk assessment. [b-Huppert] characterized seven major time service providers across 9,000+ vantage points in 170 countries, surfacing significant divergence between published protocol guidance and observed production deployments – including providers serving substantial client populations from single NTP servers in violation of [b-IETF RFC 8633] best current practices, providers running NTPv3 on the wire, and providers exhibiting clock errors on the server side that cannot be explained by network latency alone. [b-Konjerla] evaluated eight NTP/simple network time protocol (SNTP) clients under both normal operation and time-shift attacks, demonstrating that the major operating-system default clients are vulnerable to time-shift attacks achievable through DNS hijack or NTP Pool compromise [b-Moura] – without requiring an active man-in-the-middle position. The empirical picture both studies surface is consistent: the deployed NTP ecosystem diverges from the canonical reference implementation in ways that compound 2038-class exposure across the timing substrate.

13 Identified inter-SDO gaps

13.1 Inter-SDO coverage map

Table 5 provides an indicative overview of inter-SDO coverage related to time and rollover risks.

Table 5 – Indicative inter-SDO coverage related to time and rollover risks

Organization	What is covered	What is not covered
ITU-T	Protocols, architectures, security frameworks	No consolidated, cross-SG treatment of rollover-class time failures
IETF	Internet protocols (e.g., time distribution, security protocols)	No system-wide lifecycle or cross-sector rollover coordination
ISO/IEC JTC 1	Information security, software engineering standards	No deterministic, timeline-based treatment of epoch rollover risks
IEEE	Hardware, timing technologies, local protocols	No end-to-end assessment of rollover propagation across systems
FIRST/CERT communities	Incident response and vulnerability coordination	No long-horizon handling model for deterministic rollover events

14 Gap analysis

The rapid growth in AI-assisted code generation introduces a distinct structural concern: large language models trained on historical code corpora – which contain decades of 32-bit time_t patterns embedded across embedded-systems, ICS, and legacy-application codebases – exhibit a statistical tendency to reproduce those patterns in newly-generated code. The dynamic warrants particular attention for a structural reason that extends beyond the 2036-2038 cluster itself: training corpora reflect the historical state of the code that has been written, not the operational state of the substrate the code runs on. *A model whose training corpus was assembled before the rollover events does not become aware of them simply because the calendar has advanced past them; the same statistical patterns that produced 32-bit time_t idioms in code written today will continue to produce them in code generated after the rollover events have occurred, until and unless the training corpus is updated to reflect the post-rollover state of practice.* This decoupling between corpus state and substrate state is a structural condition of AI-mediated code production, not a transient feature of the period leading up to the cluster events. The combination of continued 32-bit silicon deployment, AI-mediated reproduction of legacy time-handling patterns in new code, and the corpus-substrate decoupling that prolongs this dynamic past the rollover events themselves is not deterministic, but is suggestive: the substrate of 32-bit time representations may be expanding through ordinary commercial activity not only as the date of its predictable failure approaches, but plausibly continuing to expand after that date has passed.

The indicative inter-SDO mapping in clause 13.1 illustrates that existing standards and coordination mechanisms address aspects of time handling, but do not provide a consolidated treatment of rollover-class risks.

14.1 Discovery and inventory limitations

A structural gap warranting explicit acknowledgement is the absence of authoritative, comprehensive, and globally coordinated inventory of 32-bit time-handling implementations across the deployed substrate. The conditions developed in clause 8.4 and clause 10.15 together indicate that no such inventory presently exists, that the conditions under which one could be assembled are constrained in ways that ordinary asset-inventory methods do not encounter, and that the inventory gap is not merely a precursor to remediation work but is itself a load-bearing condition shaping the remediation landscape.

The methods through which 2038-class exposure may be characterized in the absence of authoritative inventory fall into several categories.

Server-side fingerprinting of internet-exposed services – including network time protocol daemons, secure shell implementations, web server software, simple network management protocol responses, and protocol-specific banners – supports identification of architecture bitness, software version, and operating system family for systems whose service banners are externally observable; such methods do not characterize non-exposed systems and are subject to the visibility constraints developed in clause 8.4.

Client-side observation through endpoint metrics, browser telemetry, and application-layer behaviour supports inference about devices that initiate connections without exposing services, including the substantial population of consumer and industrial devices whose internet-facing surface is limited.

Software-corpus analysis through public repositories and software bill-of-materials artefacts supports identification of 32-bit time-handling patterns at the source level, where access to such artefacts is available. Firmware and hardware analysis through device specifications, technical documentation, and reverse engineering supports characterization of embedded systems whose time-handling implementations are not reflected in higher-layer signals.

Each method category produces distinctive strengths and limitations. Server-side methods support order-of-magnitude population estimates at internet scale but do not characterize the substantially larger population of non-exposed systems. Client-side methods complement server-side characterization but encounter privacy, contractual, and ethical constraints that bound their applicability. Software-corpus methods produce high-fidelity characterization of specific implementations but cannot infer deployment scale from source-level presence. Firmware and hardware methods support authoritative characterization of individual device populations but do not scale to internet-wide assessment. None of these methods, individually or in combination, produces a complete picture. This is the context within which the indicators developed in clause 14.1.1 and the measurement-methodology considerations addressed there are best understood.

The in-situ testing required to authoritatively characterize time-handling behaviour under rollover conditions, developed in clause 8.4, encounters operational, contractual, safety, and ethical constraints that bound its applicability across the deployed substrate. The historical record provides one citable example of the consequences of inadequate in-situ characterization: the loss of the [b-Ariane 5] inaugural flight in June 1996 was attributed in significant part to the re-use of inertial reference system software validated against the [b-Ariane 4] operational envelope, without revalidation against the new vehicle's actual flight trajectory. An arithmetic overflow occurred approximately 37 seconds after the start of the main engine ignition sequence (approximately 30 seconds after lift-off), producing complete loss of guidance and attitude information; vehicle disintegration followed at approximately H0+39 seconds, with total loss at approximately H0+40 seconds. The case is referenced here not as analogue to 2038-class events specifically, *but as illustration of the principle that simulation and laboratory validation, however rigorous, do not substitute for testing under operational conditions when the operational conditions differ from those designed against.* Detailed treatment of in-situ testing methodology, including the conditions under which probabilistic remote assessment can effectively support prioritization of authoritative testing, is anticipated in subsequent revisions of this Technical Paper.

14.1.1 Metrics and indicators

This Technical Paper does not define readiness metrics, maturity models, or assessment requirements. The structural conditions *developed across* clause 8 *and* clause 10 suggest that conventional metrics, which presume bounded and inventory-visible exposure, do not map cleanly onto 2038-class risk; comparative scoring across organizations or Member States on the basis of such metrics would be premature and could obscure the cross-sector substrate conditions that the present paper surfaces.

A set of descriptive indicators may nonetheless support situational awareness and prioritization across stakeholders preparing for the rollover window. Such indicators are diagnostic rather than evaluative: they characterize the shape of an organization's exposure rather than rate its preparedness. Examples include:

– **Inventory completeness**: The extent to which time representations within deployed systems are documented, attributable to known authoring teams, and traceable through the supply chain. The lifecycle conditions developed in clause 8.3 suggest that inventory gaps are themselves the load-bearing signal, not merely a precursor to remediation work.

– **Legacy-substrate persistence**: The presence and operational role of 32-bit time dependencies in software, firmware, network protocols, and the cross-sector substrate layers identified in clause 10. Substrate-layer persistence carries different implications than application-layer persistence; the indicator is informative only when paired with the layer at which it is observed.

– **Observability under forward-time conditions**: The ability of an organization to identify, test, and characterize time-dependent behaviour under operational workload when system clocks are advanced past rollover boundaries. The discovery and detectability limitations developed in clause 8.4 establish that this capacity is itself frequently the binding constraint.

– **Incident-derived signals**: Timestamp anomalies, silent failure modes, and interoperability disruptions surfaced through field observations, incident reports, and the coordinated-vulnerability disclosure channels developed in clause 8.7. These signals are particularly valuable because they reflect actual deployed behaviour rather than declarative attestation.

Measurement of 2038-class exposure warrants methodological treatment distinct from that applied to conventional vulnerability classes. The interface-level nature of the underlying exposure, developed in clause 7, produces a distinctive measurement profile: false-positive rates for remote indicators of exposure are lower than is conventional in vulnerability assessment, since the interface-level remediation required cannot be silently backported through ordinary package management; false-negative rates, conversely, are higher, since architectures running unmodified 32-bit application binaries on otherwise remediated host environments remain exposed while presenting remediation-consistent surface signals. Probabilistic and heuristic methods – drawing on community installed-base datasets, fingerprint distributions, ecosystem reporting, and version-architecture correlation – support directional and relative-exposure assessment at scales that authoritative inventory cannot reach. Such methods produce strengths and limitations distinct from conventional methods: they support order-of-magnitude estimates of where exposure resides and how rapidly the substrate is shifting, but do not support exact counts or individual-system attribution. *The combination of low false-positive and high false-negative rates suggests that probabilistic remote assessment is most effectively positioned as a triage instrument supporting prioritization of authoritative in-situ testing, rather than as a substitute for it.* The measurement methodology appropriate to 2038-class exposure assessment is itself an emerging discipline; its development is anticipated in subsequent revisions of this Technical Paper and across the cross-sector practitioner community whose engagement the present paper invites.

The indicators above are descriptive: they support shared understanding of where exposure resides and what remediation pathways remain available. They are not designed to support comparative assessment between organizations or jurisdictions, and this Technical Paper notes explicitly that any attempt to convert them into compliance criteria or scoring rubrics would misuse the present framing. Their value is in informing coordination, further study, and longer-term risk management discussion among stakeholders who already share the cross-sector concerns this Technical Paper develops.

14.2 Testing and testbed gaps

The absence of shared, accessible, and authoritative testbed infrastructure for 2038-class exposure assessment is a structural condition this Technical Paper surfaces. The observability and detectability constraints developed in clause 8.4 indicate that authoritative characterization of time-handling behaviour under rollover conditions cannot be reliably reproduced through laboratory simulation alone; in-situ testing of deployed systems is frequently the binding requirement, and the operational, contractual, and safety constraints surrounding such testing are themselves substantial. Coordinated cross-sector testbed development – including shared methodology, reproducibility conventions, and the calibration of probabilistic remote assessment methods against in-situ findings developed in clause 14.1.1 – is among the items for which inter-SDO engagement and engagement with the empirical research community would substantially advance the state of practice. Detailed treatment is anticipated in subsequent revisions of this Technical Paper.

14.3 Cross-border coordination challenges

2038-class events do not respect jurisdictional boundaries: the deployed substrate examined throughout this Technical Paper crosses operator, sector, and Member State boundaries in ways that ordinary national-level remediation pathways are not constituted to address. The cross-jurisdictional conditions surveyed in clause 10 – particularly those affecting telecommunications (clause 10.1), GNSS-dependent services (clause 10.10), cloud and ICT infrastructure (clause 10.9), and aviation (clause 10.28) – indicate that coordination mechanisms operating below the international level will encounter structural gaps that cannot be closed unilaterally. The legal, regulatory, and contractual heterogeneity across jurisdictions, developed in clause 8.6, introduces additional constraints on coordinated response. Detailed treatment of cross-border coordination pathways, including the role of ITU-T and adjacent multilateral bodies, is anticipated in subsequent revisions of this Technical Paper.

14.4 Supply chain dependencies

The lifecycle and supply-chain conditions developed in clause 8.3 indicate that 2038-class exposure is frequently inherited rather than authored: components incorporating time-handling implementations specified for short consumer-software lifecycles persist in deployed systems whose operational lifetimes extend across and beyond the rollover cluster, and responsibility for remediation is distributed across upstream and downstream parties whose contractual relationships fragment in ways that obscure risk ownership. This condition is not specific to any sector but manifests with distinctive force in sectors with long deployment lifecycles, particularly those addressed in clause 10.5, clause 10.8, and clause 10.17. The supply-chain dimension of 2038-class exposure remains substantially under-characterized at the cross-sector level, and systematic treatment – including the role of source-and-build escrow arrangements, supplier attestation regimes, and the temporal bill of materials concept introduced in clause 15.5 – is anticipated in subsequent revisions of this Technical Paper.

14.5 Absence of inter-SDO harmonization

The inter-SDO landscape relevant to 2038-class risk surveyed in clause 13 indicates that no single standards-developing organization owns the analytical or coordination scope this Technical Paper addresses. Coverage is distributed across ITU-T, IETF, ISO/IEC, IEEE (including its rollover-preparedness project [b-IEEE P4150]), CEN/CENELEC, 3GPP, and adjacent bodies whose remit, cadence, and process conventions differ in ways that produce gaps at the intersections between their respective scopes. The harmonisation challenge is not merely one of coordination between bodies that share scope; it includes the more difficult case of identifying which body should hold scope over substrate-layer concerns that fall between traditional protocol-layer and application-layer boundaries. The certification-pathway analysis developed in clause 15.4 provides one operational frame for engagement with this gap. Detailed treatment of harmonisation pathways, including specific liaison

candidates and the protocols of engagement, is anticipated in subsequent revisions of this Technical Paper.

15 Recommendations

This clause outlines high-level guidance and suggested directions for future work.

15.1 Shared terminology and conceptual framework

Effective coordination across the Member States, operators, vendors, and standards bodies preparing for the 2036-2038 rollover cluster depends on a shared vocabulary for the phenomena under discussion. The terms developed in clause 3 including *coherent time*, *temporal root of trust*, *time desynchronization*, *monotonic time*, and *2038-class rollover event* – are offered for adoption beyond the present Technical Paper. Common terminology reduces the coordination overhead of cross-sector and cross-jurisdictional engagement, supports consistent classification of incident reports and field observations, and helps prevent the dilution of authoritative framing by variable-quality content circulating in the period leading up to the rollover events, a structural concern developed in clause 10.15.

The conceptual framework consolidated across clause 6, clause 7, and clause 8 – the cluster framing, the failure-mode taxonomy, the temporal-root-of-trust formulation, and the eight analytical dimensions of clause 8 – is similarly offered as a shared reference rather than as a normative imposition. Adoption is encouraged where it supports cross-sector coherence; adaptation to sector-specific context is expected. This technical paper notes that the durability of this framework as a coordination aid will depend on sustained engagement across future editions of this Technical Paper, and that contribution to the conceptual development is welcomed.

Continued refinement of the terminology and conceptual framework across the revision cycle is anticipated, with particular attention to terms surfaced through cross-sector contribution from sectors not yet developed in clause 10.

15.2 Testing guidance

The discovery and detectability limitations developed in clause 8.4, together with the Ariane 5 inquiry findings discussed in clause 6.2, establish that laboratory simulation cannot in the general case substitute for in-situ behaviour under operational workload. The class of conditions to which 2038-class rollover belongs is precisely the class for which environment-mismatched pre-deployment testing has historically proven insufficient. Testing guidance for 2038-class exposure must accordingly accommodate the structural limitations of conventional testing modalities rather than presume their sufficiency.

Several considerations follow. Testing under forward-time-shifted conditions should be conducted on operationally representative systems under realistic load wherever this is feasible, with laboratory replication serving as a complementary rather than substitute pathway. Validation should span both host-environment and application-layer behaviour, since the host environment's modernity does not propagate into the application substrate where time-handling behaviour is determined, as developed in clause 8.3 and clause 8.6. Test horizons should extend across the full 2036-2038 cluster rather than focusing narrowly on the 2038 boundary, since exposure at earlier boundaries – particularly the 2036 NTP era rollover – functions as a leading indicator of broader substrate fragility. Test regimes should incorporate time-manipulation scenarios including backward jumps and rapid forward shifts, both to characterize behaviour under adversarial input, as developed in clause 8.5, and to surface assumptions about monotonic progression that may not hold under operational conditions.

Where remote or probabilistic scanning techniques are employed across deployed populations, the guidance developed in clause 8.4 applies directly: such techniques support population-scale triage *but cannot stand as evidence of architectural safety in any individual system*, and the false-negative

profile in particular warrants explicit acknowledgement when scanning results inform remediation prioritization. Interdependencies, assumptions, and substrate conditions surfaced during testing warrant explicit documentation, ideally in a form that supports integration with broader ICT risk registries and that preserves the institutional memory on which subsequent revision cycles will depend.

The development of standardized methodologies for progressive risk identification and in-situ validation – including triage thresholds, priority criteria, and minimum simulation requirements appropriate to different sector contexts – is among the items for which inter-SDO and cross-sector coordination would substantially advance the state of practice.

15.3 Coordination practices

The structural conditions developed across clause 8 and clause 10 indicate that effective preparation for the 2036-2038 rollover cluster requires coordination practices distinct in several respects from those developed for conventional vulnerability response. The constraints on coordinated vulnerability management developed in clause 8.7 – including the multi-year remediation cycles that 2038-class exposure indicates, the absence of a discrete disclosure moment, and the cross-sector engagement requirements that conventional CVD frameworks coordinate only incidentally – apply with particular weight to the coordination practices appropriate for the period leading up to and beyond the rollover cluster.

Several considerations warrant explicit acknowledgement. Maintenance and remediation windows in critical-infrastructure environments are routinely constrained by safety, operational-continuity, and certification considerations, as developed in clause 8.6; coordination practices that presume the availability of arbitrary maintenance windows are not generally tractable in these environments, and the development of harmonized scheduling approaches across sectors and jurisdictions would reduce the coordination overhead of remediation campaigns whose timeline is bounded by the calendar dates of the cluster events. Cross-sector information sharing – including the sharing of observed incidents, characterized exposure patterns, and remediation outcomes – supports the development of empirical signal across the deployed substrate and is one of the principal mechanisms through which the 2036 NTP era rollover can function as a leading indicator for the 2038 events, as developed in clause 6. Coordinated monitoring across operators and jurisdictions during the rollover events themselves – particularly the 2036 NTP era boundary, the 2038 time_t boundary, and the 2038 GPS week-number boundary – would support real-time identification of cascading effects and reduce the latency of cross-jurisdictional response.

The Forum of Incident Response and Security Teams (FIRST) and its Time Security Special Interest Group have established relationships across which such coordination can proceed; the operational role of national CERTs and CSIRTs, sectoral coordination bodies, regulators with critical-infrastructure oversight responsibilities, and the timing-service provider community is developed throughout clause 10. Continuity of awareness and coordination across SG17 study periods is itself a structural condition of effective preparation, given the multi-year horizon of the work and the rate at which institutional memory is lost across the commercial reorganizations affecting the substrate, as explained in clause 10.15.

15.4 Inter-SDO liaison and certification considerations

Liaison will be conducted through existing ITU-T SG17 channels and supplemented by informal coordination through FIRST membership networks. The relationship between this Technical Paper and adjacent standards work in IETF, IEC, ITU-R and CEN/CENELEC is acknowledged in clause 13 and clause 14.5.

A structural observation worth surfacing for further study: certification and compliance pathways may represent a higher-leverage mechanism for driving uniform 2036/2038-class remediation behaviour across the deployed substrate than protocol-level standardization alone. The argument is operational rather than theoretical. Vendor incentives to address rollover-class exposure are constrained by the absence of explicit testable conformance criteria in certification frameworks that vendors are already subject to. Whether 2038-class rollover behaviour might be incorporated into existing certification pathways (without changing the underlying protocols) is identified here as an area for further study and inter-SDO liaison rather than a recommended measure: such pathways already operate with the cadence, audit infrastructure, and enforcement mechanisms capable of driving change at scale, but this Technical Paper defines no conformance criteria. Prior discussions in the NTP community indicate that this framing has been considered since approximately 2013; the present Technical Paper surfaces it as a coordination opportunity rather than a prescriptive recommendation.

The interaction between this observation and recent EU regulatory developments – particularly [b-EU CRA] and adjacent frameworks – warrants explicit treatment in subsequent revisions. The [b-EU CRA] product-cybersecurity conformance regime is a candidate vehicle for embedding 2038-class testing requirements in the certification flow without requiring new normative protocol work. Analogous certification regimes in other jurisdictions (FDA pre-market submission for medical devices addressed in clause 10.5; ICS/OT certification regimes addressed in clause 10.8; aviation safety certification frameworks addressed in clause 10.28) face structurally similar opportunities.

A further observation, noted for future study rather than current development: *the integrity-checking function of audit, oversight, and governance mechanisms depends not only on their substantive design but on their cadence relative to the operational cadence of the work they are meant to govern.* Certification regimes whose audit cycles operate at substantially slower cadence than the remediation cycles of the deployed substrate they certify face a structural challenge that warrants explicit analysis, particularly under conditions where remediation timing is bounded by external events such as the 2036-2038 cluster. The general principle – that cadence-matching between operational work and integrity-checking mechanisms is a structural design property distinct from the substantive design of either – applies beyond the certification case and bears systematic treatment in subsequent revisions.

NOTE – Sector-specific certification-pathway analysis is anticipated in subsequent revisions. The structural observation surfaced here is intended to support inter-SDO coordination discussions rather than prescribe specific certification-criteria language.

15.5 Procurement considerations

Procurement decisions taken in the period leading up to and beyond the 2036-2038 rollover cluster have structural consequences for the remediation pathways available to the operators and end-users of the procured systems. The lifecycle and supply-chain conditions developed in clause 8.3 and the legal-and-contractual constraints developed in clause 8.6 together indicate that the procurement moment is frequently the last point at which several remediation pathways remain open by default; once a deployment is established under standard procurement terms, the contractual and legal substrate may narrow remediation to interventions the original vendor undertakes on its own timeline and with its own operational capacity.

Several considerations may accordingly warrant attention during procurement of systems incorporating digital elements with operational lifetimes extending across the rollover cluster.

1) **Awareness of time-handling implementations and their representable bounds**, including the underlying timestamp widths, epoch definitions, and signedness assumptions of the procured components – supports informed assessment of exposure that procurement specifications have not historically required suppliers to disclose.

2) **Supplier attestation regarding time-handling behaviour across the cluster events**, where it can be obtained, supports later remediation planning even where it does not by itself remediate exposure. One candidate mechanism for operationalising such attestation in structured form, offered here as a future-revision direction rather than a present-revision recommendation, is what has been proposed in adjacent practitioner discussions under the working name **temporal bill of materials** (TBoMs): a structured disclosure of the time-handling characteristics of components incorporated into a procured system, analogous in spirit to the *software bill of materials* (SBOM) artefact familiar from cybersecurity supply-chain practice. The analogy is suggestive rather than exact – the time-handling dimension of components is presently characterized at a lower fidelity than the dependency-and-version information SBOM artefacts typically carry, and the standardization pathways through which a TBoM artefact could be developed remain to be elaborated. The concept is surfaced here for future-revision consideration; the supply-chain dimension referenced in clause 14.4 is the broader analytical context within which structured-disclosure mechanisms of this kind warrant systematic treatment.

3) Where systems are acquired with public funds, deployed in long-lived critical-infrastructure roles, or carry safety-of-life implications, the **source-and-build escrow arrangements** discussed in clause 8.3 may merit consideration as a structural counterweight to vendor-dependency conditions that the contractual substrate would otherwise produce.

4) **Maintenance and support agreements whose terms extend across the rollover cluster**, with explicit treatment of time-handling remediation as a covered category of work, preserve operational pathways that ordinary maintenance contracts do not generally anticipate.

These considerations are indicative only and do not constitute procurement mandates, do not specify contractual language, and do not propose conformance criteria; they reflect the structural observation that procurement is one of the points at which the remediation landscape of a deployed system is most amenable to shaping, *and at which the cost of preserving future optionality is substantially lower than the cost of recovering it after deployment*. The procurement moment is one instance of a broader category – work performed before the precipitating event in a coordinated infrastructure-risk scenario – that exhibits structurally different cost, optionality, and coordination properties than equivalent work performed after the event. Detailed treatment of this category, and of the analytical implications of structuring preparedness work in advance of the conditions it is meant to address, is anticipated in subsequent revisions of this Technical Paper. Future revisions are expected to include jurisdiction-specific procurement framework analysis and the interaction with certification regimes addressed in clause 15.4.

15.6 Harm minimization considerations

The non-normative posture of this Technical Paper does not extend to claiming that the conditions it describes are without available response. A category of action distinct from full remediation – namely harm minimization – remains available across the deployed infrastructure landscape and warrants explicit acknowledgement in this revision even where its detailed treatment is anticipated in future revision cycles.

The observation is structural rather than tactical. Even where full remediation of 2038-class exposure across a deployed system is not achievable on the available timeline, a range of harm-minimization strategies remain operationally meaningful: deliberate de-digitalization of subsystems whose digital implementation no longer adds operational value commensurate with its 2038-class exposure; redesign toward hybrid analogue/digital architectures in which time-coordination dependencies are decoupled from operationally critical control pathways; introduction of human-legible, human-updatable, and human-operable override mechanisms at the level where time-coordination failures would otherwise produce silent or unrecoverable system states; staged degradation modes that preserve essential functionality under conditions of time-coordination failure rather than producing binary success/failure outcomes.

These approaches draw on a substantial design tradition in resilience engineering and on the broader category of loose-coupling in complex socio-technical systems ([b-Perrow]), and are present in mature form in some safety-critical industries (aviation flight controls, nuclear plant operations, certain classes of medical devices) but absent or under-developed across substantial portions of the deployed infrastructure landscape. Harm minimization is structurally distinctive in conditions of infrastructural vestigialisation, as developed in clause 8.3; the response set available to systems whose capacity-to-rebuild has atrophied is narrower than the response set assumed by remediation frameworks that presume an ordinary engineering substrate, and the harm-minimization pathways accordingly carry disproportionate weight in vestigialised portions of the deployed landscape.

Whether such approaches can be sustained across the rollover window depends substantially on the organizational culture in which they are developed and deployed: Westrum (2004) notes that bureaucratic responses to anomaly characteristically favour suppression, encapsulation, or local fixes over inquiry and system-wide remediation, while generative responses surface anomalies for cross-sector inquiry. Harm minimization as a category of action remains structurally available; whether it is operationally available in any given institutional context depends on the culture of the organizations through which it would have to be coordinated.

A category of successful engineering precedent for rollover-aware design warrants brief acknowledgement alongside the harm-minimization framing developed above. Several deployed protocols incorporate rollover-aware semantics by construction: SNMP Counter32 values are defined as differential rather than absolute, with the rollover boundary explicitly handled by the consuming application's delta calculation; the serial-number arithmetic specified in [b-IETF RFC 1982] provides defined comparison semantics across the modular boundary for DNS start of authority (SOA) records and analogous use cases; and adjacent protocol families have developed similar conventions where rollover-aware semantics were anticipated at design time. These precedents demonstrate that rollover-aware design is operationally feasible at the protocol-specification layer where the failure mode is anticipated and the boundary semantics are defined deliberately. A detailed survey of rollover-aware engineering precedents across the protocol landscape is anticipated in subsequent revisions; the survey will support the harm-minimization considerations developed above by establishing that the design pattern has been successfully instantiated in deployed practice.

Even where the trajectory of a system toward a known future failure is not alterable, the consequences of that failure are not predetermined, and the period between present awareness and the failure event represents a substantial harm-minimization window. Foreclosing the harm-minimization question through either inattention or despair removes from the available response set a category of intervention that does not require full remediation and that may, in many sectors, be the most operationally tractable form of preparation.

Engagement invited from resilience engineering, safety-critical systems design, and human-factors practitioners – including those working on hybrid analogue/digital architectures, human-machine interface design for failure modes, and the broader design tradition surrounding loose coupling and graceful degradation. This clause is structurally important given the operational tractability of harm-minimization approaches relative to full remediation across substantial portions of the deployed infrastructure landscape, and the significant design-tradition base from which such approaches can be developed.

15.7 Areas for future study

– Design considerations for successor time representations that preserve operational properties (total ordering, compact representation, low cognitive overhead) while providing semantic distinguishability from legacy encodings. Active discussions are underway in adjacent standards communities; coordination with these efforts is anticipated.

– The interaction between automated code generation tools and legacy time-handling patterns in training corpora warrants systematic study, particularly with respect to ecosystem-level propagation effects.

– The IPv4-to-IPv6 transition bears systematic comparison with 2038-class remediation as a precedent for the partial-migration dynamics that arise when a hard-coded width in a load-bearing protocol becomes inadequate to operational needs and the transition to a wider replacement proceeds asymmetrically across the deployed substrate. The IPv4-to-IPv6 case offers an empirical record – the only one currently available – of how operators, vendors, and standards bodies have navigated such a transition at planetary scale, including the coordination structures that emerged, the duration of the partial-migration window, the categories of failure observed during it, and the conditions under which remediation incentives proved sufficient or insufficient to drive sustained progress. Systematic treatment of this precedent, including identification of which elements transfer to 2038-class remediation and which do not, is anticipated in subsequent revisions.

– The recommendations developed in the clauses below are organized primarily by coordination mechanism (terminology, testing, coordination practice, inter-SDO liaison, procurement, harm minimization) rather than by failure-mode urgency. A complementary organising principle is *triage-by-failure-tempo* – which takes into primary consideration the rate at which the human consequences of inadequate remediation manifest after rollover. Failure modes affecting life safety, for example in the cases of medical devices, emergency communications, or aviation – see clause 10.5, clause 10.7, and clause 10.28 – manifest consequences at the shortest tempo and accordingly bear the highest prioritization. Failure modes affecting critical-infrastructure substrates (power, water, transport, telecommunications, addressed in clause 10.1, clause 10.20, clause 10.24 and adjacent clauses) manifest at a slower tempo but with broader cross-sector propagation. Failure modes affecting economic and financial systems – including settlement, billing, insurance, as addressed in clause 10.14 and clause 10.18– manifest at the longest tempo and therefore admit greater operational flexibility in response. *Response capacity, where it is bounded, may be most effectively allocated along the tempo axis rather than across the coordination axis.*

A further consideration warrants surfacing alongside triage-by-failure-tempo: the failure-tempo axis interacts with the delayed-manifestation conditions developed in clause 6.2. Response posture calibrated to the rollover event itself may relax before the bulk of consequences have surfaced, particularly for failure modes whose manifestation cadence is governed by operational scheduling logic – annual reporting cycles, certificate renewal intervals, monthly batch processing, and similar. *Response cadence appropriate to 2038-class events accordingly extends over months to years following the threshold, not concentrated around it.*

15.8 Success criteria for this Technical Paper

The success of this Technical Paper may be assessed through qualitative indicators reflecting its usefulness as a shared reference, rather than through compliance or adoption metrics.

Indicative signs of success include:

– Adoption as a reference document by Member States, operators, standards bodies, and coordination communities when discussing or assessing 2038-class timestamp rollover risks.

– Consistent use of common terminology and concepts introduced in this paper across related discussions, studies, and coordination activities within ITU-T Study Group 17 and beyond.
– Contribution of structured inputs to future analytical or normative work, where appropriate, without prejudging outcomes, scope, or the need for such work.

These criteria are indicative only. They are intended to support reflection on the impact and relevance of the Technical Paper over time, rather than to define targets, performance measures, or obligations for any stakeholder.

Bibliography

[b-ITU-T G.8271] Recommendation ITU-T G.8271/Y.1366 (2020), *Time and phase synchronization aspects of telecommunication networks.*

[b-21 CFR 11] U.S. Food and Drug Administration. (1997). *Electronic records; electronic signatures*. 21 C.F.R. Part 11.

[b-Ariane 4] European Space Agency. (2003), *Ariane 4: The launcher that gave Europe its independence in space* (ESA Briefing Paper). ESA Communications.

[b-Ariane 5] Lions, J.-L. (1996), *Ariane 5 Flight 501 failure: Report by the Inquiry Board*. European Space Agency.

[b-Boehm] Boehm, S. and Baran, P. (1964), *On distributed communications: II. Digital simulation of hot-potato routing in a broadband distributed communications network*. RAND Corporation Memorandum RM-3103-PR.

[b-CAT] U.S. Securities & Exchange Commission (SEC) (2012), *Consolidated audit trail*. 17 C.F.R. § 242.613.

[b-CGPMRES4E] General Conference on Weights and Measures (CGPM) (2022), *On the use and future development of UTC,* Resolution 4 of the 27th CGPM, International Bureau of Weights and Measures (BIPM).

[b-DigiEuro] European Central Bank, (2025). *Eurosystem moving to next phase of digital euro project* (Press Release). <https://www.ecb.europa.eu/press/pr/date/2025/html/ecb.pr251030~8c5b5beef0.en.html>

[b-EAS] Federal Communications Commission (2022), *Emergency Alert System (EAS)*.

[b-ETSI TS 102 900] ETSI TS 102 900 (2023), *Emergency Communications (EMTEL); European Public Warning System (EU-ALERT) using the Cell Broadcast Service*.

[b-EU 536/2014] Regulation (EU) No 536/2014, *Medicinal products for human use, and repealing Directive 2001/20/EC*. Official Journal of the European Union, L 158, 1-76.

[b-EU 600/2014] Regulation (EU) No 600/2014, *Markets in financial instruments and amending Regulation (EU) No 648/2012*. Official Journal of the European Union, L 173, 84-148.

[b-EU 2025/1155] Commission Delegated Regulation (EU) 2025/1155, *Supplementing Regulation (EU) No 600/2014 of the European Parliament and of the Council with regard to regulatory technical standards specifying the input and output data of consolidated tapes, the synchronisation of business clocks and the revenue redistribution by the consolidated tape provider for shares and ETFs, and repealing Commission Delegated Regulation (EU) 2017/574*. Official Journal of the European Union.

[b-EU CRA] Regulation (EU) 2024/2847, *horizontal cybersecurity requirements for products with digital elements and amending Regulations (EU) No 168/2013, (EU) 2019/1020 and Directive (EU) 2020/1828 (Cyber Resilience Act)*, Official Journal of the European Union, L series, 2024/2847.

[b-Flash Crash] U.S. Commodity Futures Trading Commission (CFTC) and U.S. Securities & Exchange Commission (SEC) (2010), *Findings regarding the market events of May 6, 2010*.

[b-FIRST Time SIG] FIRST Time Security Special Interest Group. <https://www.first.org/global/sigs/time/>

[b-GAMP] ISPE GAMP 5 (2022), *A risk-based approach to compliant GxP computerized systems (Second Edition)*.

[b-GCP] International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use (ICH) (2025), *Guideline for good clinical practice*, ICH Harmonised Guideline E6(R3).

[b-GDP] WHO TRS 1025 (2020), *Annex 7: Good storage and distribution practices for medical products*.

[b-GMP] WHO TRS 986 (2014), *Annex 2: WHO good manufacturing practices for pharmaceutical products: Main principles Good Manufacturing Practice (GMP)*.

[b-GPS.gov] U. S. Government (2019), *GPS week number rollover*. <https://www.gps.gov/news/gps-week-number-rollover>

[b-GNU C Lib] *Y2038 Proofness and the time_t Transition* (2019) <https://sourceware.org/glibc/wiki/Y2038ProofnessDesign> (accessed January 2026).

[b-HACCP] WHO and FAO CXC 1-1969 (2023), *General principles of food hygiene*.

[b-Hayes] Hayes, B. (2014), *Infrastructure: A field guide to the industrial landscape*.

[b-Huppert] Huppert, P., Moura, G.C.M., Davids, M., de Boer, P.-T., Holz, R., Fernhout, R., Smaragdakis, G. and Hesselman, C. (2025), *BigTime: Characterizing large time service providers*. Technical Report. Arnhem: SIDN Labs.

[b-IATF 16949] IATF 16949:2016, *Quality management system requirements for automotive production and relevant service parts organizations (1st ed.)*.

[b-ICSA-25-261-07] ICSA-25-261-07 (2025), *Dover fueling solutions progauge MagLink LX4 devices* (CVE-2025-55068). <https://www.cisa.gov/news-events/ics-advisories/icsa-25-261-07>

[b-ICSA-25-296-03] ICSA-25-296-03 (2025), *Veeder-Root TLS4B automatic tank gauge system* (CVE-2025-55067). <https://www.cisa.gov/news-events/ics-advisories/icsa-25-296-03>

[b-IEC 61508] IEC 61508-series, *Functional safety of electrical/electronic/programmable electronic safety-related systems*.

[b-IEC 61511] IEC 61511-series, *Functional safety – Safety instrumented systems for the process industry sector*.

[b-IEC 61850] IEC 61850-series, *Communication networks and systems for power utility automation*.

[b-IEC 62264] IEC 62264-series, *Enterprise-control system integration*.

[b-IEC 62443] IEC 62443-series, *Security for industrial automation and control systems*.

[b-IEEE 1588] IEEE 1588 (2019), *IEEE Standard for a precision clock synchronization protocol for networked measurement and control systems*.

[b-IEEE P4150] IEEE P4150, *Recommended Practice for Assessing and Documenting Preparedness for Time Rollover Events in Years 2036, 2038, and 2106. IEEE Standards Association project*; PAR approved 4 June 2026. <https://standards.ieee.org/ieee/4150/12603/>

[b-IETF RFC 677] IETF RFC 677 (1975), *Maintenance of duplicate databases*.

[b-IETF RFC 1305] IETF RFC 1305 (1992), *Network time protocol (version 3) specification, implementation and analysis*.

[b-IETF RFC 5905] IETF RFC 5905 (2010), *Network time protocol (version 4) protocol and algorithms specification*.

[b-IETF RFC 1982] IETF RFC 1982 (1996), *Serial Number Arithmetic*.

[b-IETF RFC 8633] IETF RFC 8633 (2019), *Network time protocol best current practices*.

[b-ISO 9001] ISO 9001:2015, *Quality management systems*.

[b-ISO 15118] ISO 15118-series, *Road vehicles – Vehicle to grid communication interface*.

[b-Alstom] Alstom 2038-class remediation order (2025). French press coverage: *Le Figaro, Le Parisien, L'Informé, Presse-citron*.

[b-Knight Capital] U.S. Securities & Exchange Commission (SEC) (2013), *In the matter of Knight Capital Americas LLC* (Administrative Proceeding File No. 3-15570; Release No. 34-70694).

[b-Konjerla] Konjerla, S., Moura, G.C.M., Smaragdakis, G. and Dittrich, T. (2025), *Are NTP clients always right?* Technical Report. Arnhem: SIDN Labs.

[b-Lamport-1] Lamport, L. (1978), *Time, clocks, and the ordering of events in a distributed system*. Communications of the ACM, 21(7), 558-565.

[b-Lamport-2] Lamport, L., Shostak, R., and Pease, M. (1982), *The Byzantine generals problem*. ACM Transactions on Programming Languages and Systems, 4(3), 382-401.

[b-Malhotra] Malhotra, A., Cohen, I. E., Brakke, E., and Goldberg, S. (2016), *Attacking the network time protocol*. In Proceedings of the Network and Distributed System Security (NDSS) Symposium 2016.

[b-Mills] Mills, D. L. (2011), *Computer network time synchronization: The network time protocol on Earth and in space*. 2nd ed. CRC Press.

[b-Moura] Moura, G. C. M., et al. (2024), Deep dive into NTP Pool's popularity and mapping. Proceedings of the ACM on Measurement and Analysis of Computing Systems (SIGMETRICS). <https://doi.org/10.1145/3639041>

[b-Perrow] Perrow, C. (1984), *Normal accidents: Living with high-risk technologies*. Basic Books; updated edition Princeton University Press, 1999.

[b-Reuters] Lee L. (2025), *China accuses US of cyber breaches at national time centre, Reuters*. <https://www.reuters.com/world/china/china-accuses-us-cyber-breaches-national-time-centre-2025-10-19/>

[b-SIB 19-01R1] EASA SIB 2019-01R1 (2019), *GPS week rollover*.

[b-Tainter] Tainter, J. A. (1988), *The Collapse of Complex Societies*. Cambridge University Press.

[b-Thomas] Thomas, R. H. (1979), *A Majority Consensus Approach to Concurrency Control for Multiple Copy Databases*. ACM Transactions on Database Systems, 4(2), pp. 180-209. <https://doi.org/10.1145/320071.320076>

[b-Westrum] Westrum, R. (2004), *A typology of organisational cultures*. Quality and Safety in Health Care, 13(Suppl II): ii22–ii27.

[b-WTISD26] ITU, UNDRR and Sciences Po (2026), *When digital systems fail – An expert report on the hidden risks of our digital world*.