NIIN

National Industry Innovation Network

Digital Resilience in a New Age

Boosting the digital resilience of Australia's critical infrastructure



Frank den Hartog | UNIVERSITY OF CANBERRA
Mizanur Rahman | UNIVERSITY OF CANBERRA
Daniella Gullotta | UNIVERSITY OF CANBERRA
Simone Chitsinde | UNIVERSITY OF CANBERRA

Research Development + Innovation Program

Innovation Central Network

Skills + Talent Development

Cyber + Health Alliances

Executive Summary

Australia's critical infrastructure highly depends on interconnected digital systems, including information technology (IT), operational technology (OT), and communications networks. While these systems can fail, possibly leading to large-scale disruptions of critical services, they can also play an important role in providing resilience against such disruptions. This whitepaper examines how Australian critical infrastructure operators understand, assess, and strengthen digital resilience, defined as the ability of systems and networks to respond to, recover from, and adapt to disruption while maintaining essential functions and acceptable performance. Its aim is to provide solution architects with best practices and recommendations that can improve the resilience of the critical services their organisations provide to their customers.

Through literature-based case study analysis and in-depth interviews with critical infrastructure operators across multiple sectors, our findings indicate that digital resilience is now regarded as a core, all-of-business, organisational capability rather than a purely technical concern. While redundancy remains an important design principle, the most resilient organisations are not necessarily those with the greatest level of technical duplication, but those with the strongest capabilities for detection, isolation, recovery, adaptation, and operational decision-making.

A significant finding is the critical role of maintaining a minimum level of observability and situational awareness also after the event. Across multiple incidents, recovery was delayed not by the initial fault itself but by an inability to rapidly identify and understand the failure and observe the recovery progress. Monitoring and observability infrastructure should be significantly more robust and resilient against threats than the production systems.

Critical infrastructure increasingly relies on complex interdependent ecosystems of cloud services, on-prem systems, networks, and systems for authentication, localisation, timing, etc. Many of these dependencies are not known or understood and can lead to cascading failures. Digital resilience will increasingly depend on an organisation's ability to understand complex dependencies and maintain visibility across interconnected systems.

We also found substantial variation in resilience practices, without a common approach across sectors. Many organisations rely on pragmatically and historically grown practices rather than industry best practices, due to a lack of forums where engineering practitioners can effectively share common insights.

©University of Canberra. 2026

We, thus, landed on the following recommendations for the industry:

Recommendations

1. **Focus investment on monitoring and observability infrastructure to obtain more and better data, and to be the most robust, independent, and resilient of the entire estate.**

2. **Improve change management and migration strategies and practices, as ever more complex digital ecosystems require more rigorous approaches to obtain full control during change.**

3. **Continuously analyse hidden dependencies and single points of failure.**

4. **Drive greater collaboration between the engineering practitioners across different critical infrastructure operations to grow the maturity and adoption of best practice frameworks, common metric sets and widely accepted maturity models.**

5. **Leverage common metrics, maturity models and practice to continuously validate the actual level of resilience in critical infrastructure operations.**



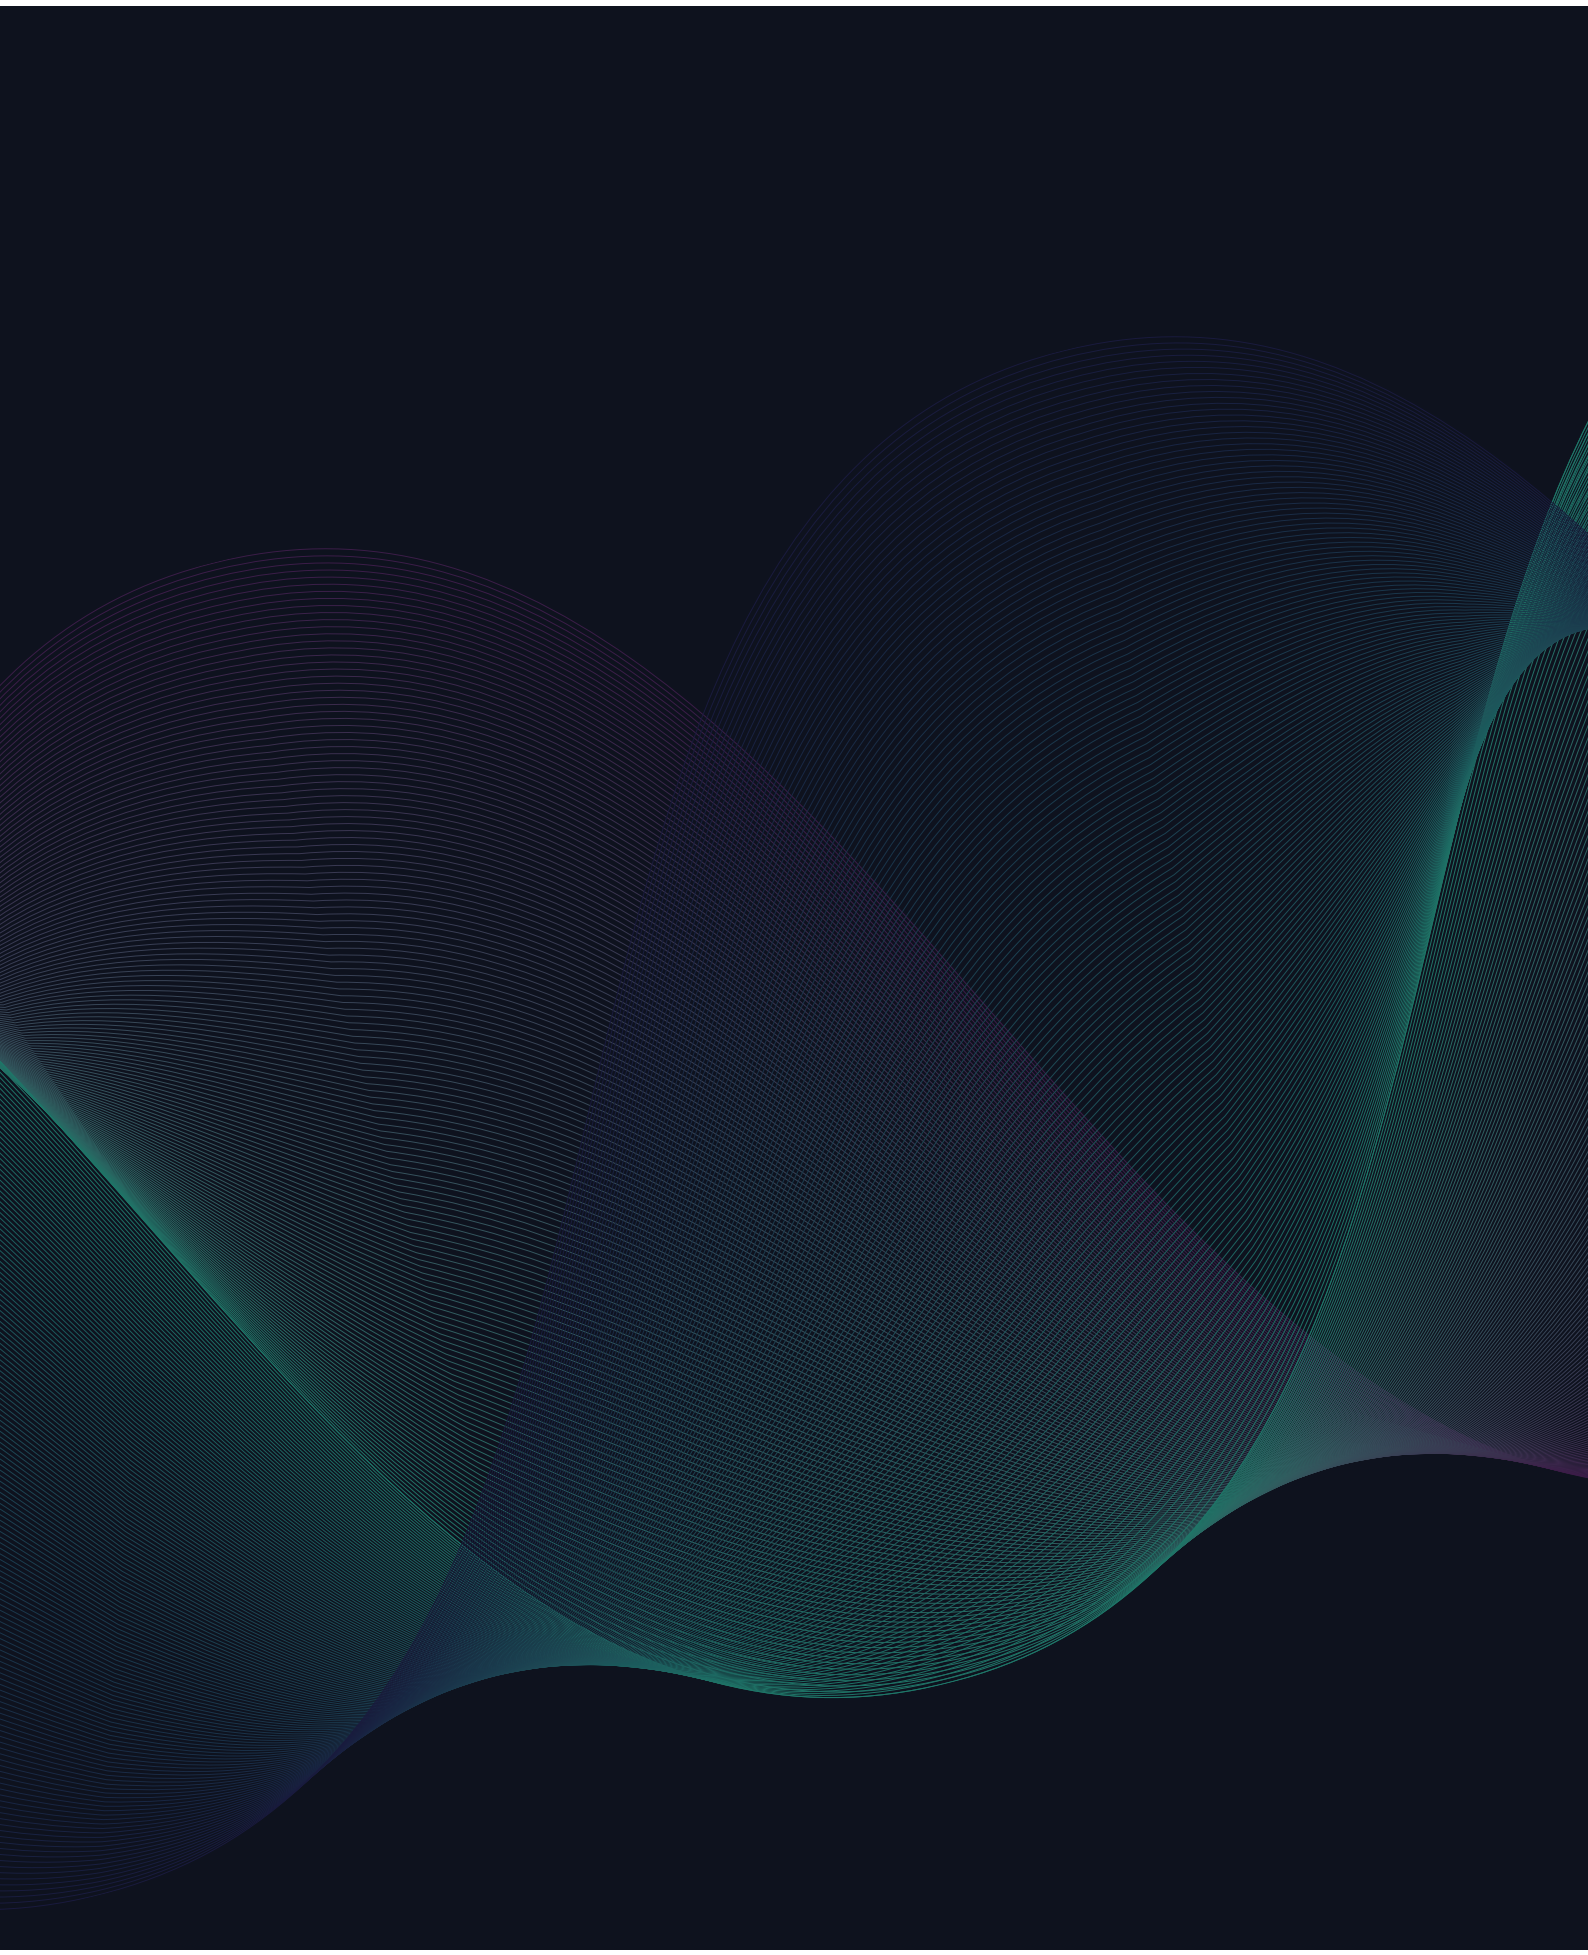


Table of Contents

Executive Summary 1

Table of Contents 3

1. Introduction 5

2. Resilience and digitisation of critical infrastructure 7

2.1 The digitisation of our critical infrastructure 7

2.2 How to define resilience 8

2.3 Relevant standards, frameworks, and guidelines 9

3. Relevant use cases from literature 11

3.1 Methodology 11

3.2 Results 11

4. Interviews with critical infrastructure operators 16

4.1 Methodology 16

4.2 Understanding resilience 16

4.3 Impact quantification 18

4.4 Network and system dependencies 19

4.5 Previous incidents 22

4.6 Resilience standards and framework 24

4.7 Testing and training 24

4.8 Cyber threats, cascading failures, and physical threats 26

5. Analysis 29

5.1 Migration strategies and change management 29

5.2 Lack of joint approaches 29

5.3 Deeply integrated production and recovery architectures 30

5.4 Single points of failure and cascading failure 32

6. Conclusions, future work, and recommendations 34

6.1 Migration strategies and change management 34

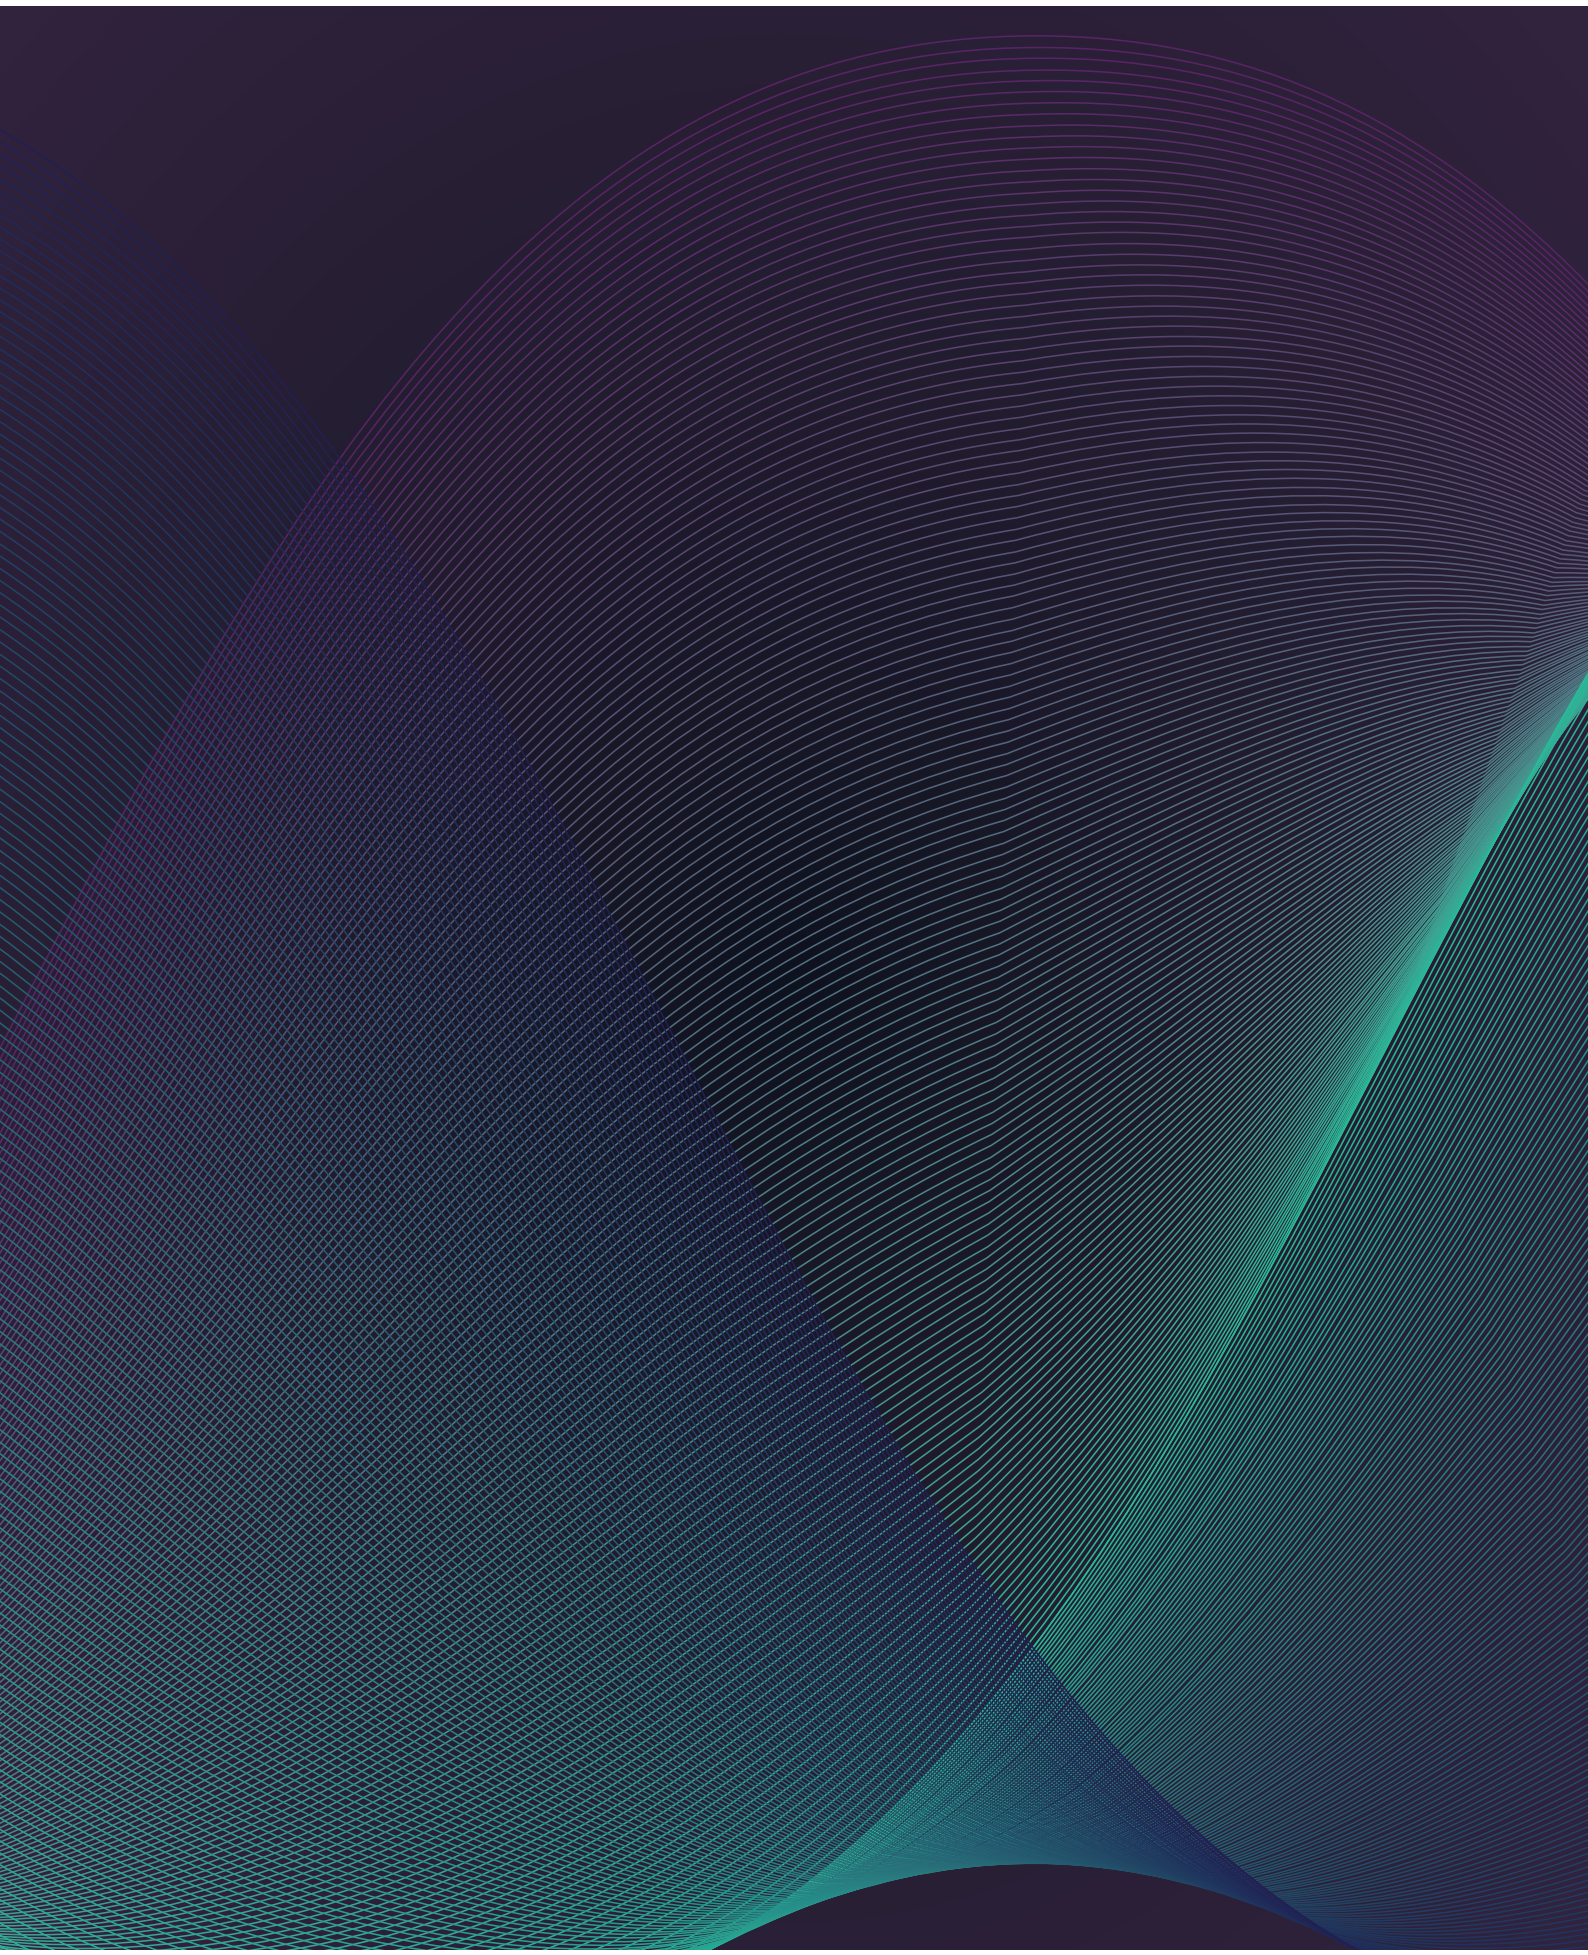
6.2 Lack of joint approaches 35

6.3 Deeply integrated production and recovery architectures 37

Acknowledgements 38

Contact 38

References 39



1. Introduction

Like critical infrastructure elsewhere in the world, Australia's critical infrastructure is continuously exposed to external threats, be it physical such as weather events or digital such as cyber security incidents. Regardless, two main questions arise: how do our critical infrastructure operators protect their systems against such events, and how do they get these systems back up running again when protection has failed? The latter is the focus of this whitepaper, particularly in the context of the digital systems that our critical infrastructure so much depends on today: IT (Information Technology) systems, OT (Operational Technology) systems, and digital communication networks. We call this digital resilience. It is the capacity of a system or network to respond to, recover from, and adapt to disruptions while maintaining its essential functions and acceptable performance.

The digital systems underpinning critical infrastructure are becoming increasingly complex as AI, autonomous operations, cyber defence, and continuous optimisation become integral to operations. Likewise, managing their resilience becomes more intricate, and the risk of malinvestment in digital resilience has become ever more real. For instance, a common misconception equates redundancy with resilience, often resulting in costly yet less effective strategies that fail in real-world conditions. Contemporary incidents such as the CrowdStrike-related events in 2024 highlight systemic gaps where resilience strategies collapsed despite redundant systems, exposing dependence on a single vendor's update process, lack of observability, and insufficient operational readiness. These failures underscore the urgent need for an approach that prioritises visibility, intelligence, and proactive updates to change, migration, and recovery procedures, to ensure resilience is both measurable and sustainable.

The aim of this whitepaper is to provide an analysis of decisions currently being made in digital resilience across critical infrastructure IT and OT environments as well as digital communication networks. We carried out this research by comparing different interpretations of resilience relevant for critical infrastructure, analysing common principles and mistakes in resilience strategies across sectors, examining strategically selected use cases from different industries where resilience strategies failed or were successful, and engaging industry experts through in-depth interviews. The intended audience of this paper are IT and OT solution architects working with critical infrastructure operators, who are tasked with improving the digital resilience of their estate and the critical services they provide to their customers.

The paper is structured as follows. The next chapter discusses definitions and standards related to critical infrastructure and digital resilience. This is then followed by an extensive use case analysis based on academic and grey literature. Next, we present the results of our interviews with twelve experts from different critical infrastructure sectors, followed by an extensive analysis of our results. The paper concludes with our main findings and a set of recommendations.


Data Sector : 8183 04588 79
ANALYSIS DATA NODE

2. Resilience and digitisation of critical infrastructure

2.1 The digitisation of our critical infrastructure

The Australian Government's 2023 Critical Infrastructure Resilience Strategy (Cyber and Infrastructure Security Centre, 2023) defines critical infrastructure as *those physical facilities, systems, assets, supply chains, information technologies and communication networks which, if destroyed, degraded, compromised or rendered unavailable for an extended period, would significantly impact the social or economic wellbeing of Australia as a nation, or its states or territories, or affect Australia's ability to conduct national defence and ensure national security.*

The Security of Critical Infrastructure Act 2018 (Australian Government, 2026a), or SOCI Act in short, mentions eleven business sectors that the government has recognised as sectors in which critical infrastructure assets are deployed. They include communications, financial services and markets, data storage or processing, defence industry, higher education and research, energy, food and grocery, health care and medical, space technology, transport, and water and sewerage.

It is interesting to note that the SOCI Act only came into effect eight years ago, while every society had to defend its critical infrastructure already for centuries. The driver behind the need for the SOCI Act becomes apparent with the observation that much of the impetus behind this act relates to cyber security. Over the past decades, our critical infrastructure has rapidly become digitised. Digitisation means that computers, software and digital communication networks have entered deep into the fabric of the infrastructure and, directly or indirectly, have connected many critical infrastructures to the Internet and with each other. Over the course of two decades, digital assets across the entire economy have more than doubled, as firms invested not just in IT but also in digitising their physical assets (Gandhi et al., 2016).

In this paper, criticality is defined as per Cyber and Infrastructure Security Centre (2023), which focuses on the impact that the unavailability of critical services has on the social or economic wellbeing of the whole nation. But many conclusions from our research may also be applicable to organisations, whether public or private, that operate infrastructure that may not be critical for the nation, but whose resilience is nevertheless critical to their own organisations and customers.

It is also worth highlighting that enforcement of the SOCI Act will strengthen this year (Critical Infrastructure Security Centre, 2026), and that its obligations can flow down through supply chains. Operators with direct obligations under SOCI can be expected to push aspects of their resilience requirements onto vendors, suppliers and partners who play a role in the resilient operation of the infrastructure.

The bottom line is that the availability of our critical services has become strongly dependent on the uptime and performance of the underlying infrastructure, which is more than ever depending on IT and connectivity.

It, thus, is exposed to all the typical threats and weaknesses that IT and connectivity suffer from, including lack of maintenance, mistakes in architectures, software bugs, hidden dependencies, cyber attacks, and physical threats such as fibre cuts.

2.2 How to define resilience

The concept of resilience has evolved substantially over time. Nevertheless, there is no consistent understanding of what resilience means and how to measure it across industries.

The term resilience originates from the Latin *resalire*, meaning "to spring back," and has long been used across ecology, psychology, and engineering to describe how a system responds to disturbance. In engineering, computing, and networking, this idea was later extended to address systems that must operate in dynamic, uncertain, and sometimes adversarial environments. As a result, resilience is now understood as more than merely "bouncing back" after failure. In systems and networks, it refers to the broader ability to continue functioning through disruption, recover in a controlled way, and adjust to changing conditions over time (Khaloopour et al., 2024).

Rather than focusing only on correct operation under stated conditions or predefined fault assumptions, resilience is concerned with how systems behave when they face disruption, uncertainty, evolving threats, or conditions that may exceed their original design assumptions. This is schematically depicted in Figure 1, which is inspired by the work described in (Kott et al., 2024) about the cyber resilience of Ukrainian critical infrastructure. The figure shows the performance of a system (or a system of systems, over time). The system withstands events successfully, until at some stage it does not. The performance of the system then degrades and will come back. Resilience is the net effectiveness of all the measures put in place to minimise the area of degradation. As "performance" is measured differently between different infrastructures, also resilience will be defined and measured differently.

That said, there is a clear trend in the literature to define resilience more broadly and, eventually, to capture every measure to protect availability without exception (see, for instance, (Australian Government, 2026b). In this whitepaper, however, we scope the term down to the relatively narrow definition of *the capacity of a system or network to respond to, recover from, and adapt to disruptions while maintaining its essential functions and acceptable performance.*

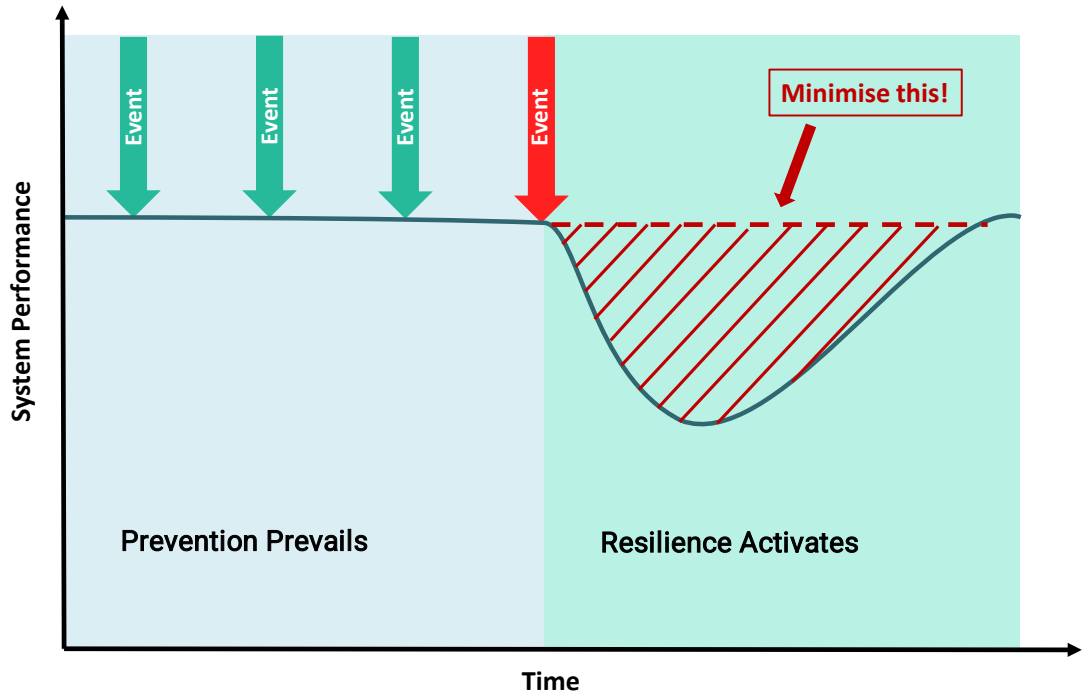


Figure 1: Resilience are all the measures put in place to minimise system performance degradation over time after an impactful event such as a flood or a successful cyber attack.

2.3 Relevant standards, frameworks, and guidelines

In practice, resilience in systems and networks is usually not implemented through one single framework.

Most industries use a layered combination of frameworks that cover strategy, governance, risk, business continuity, cyber security, systems engineering, sector-specific operations, and regulatory assurance. For systems and network resilience, the most relevant frameworks are the ones that address system design, cyber-resilient engineering, ICT recovery readiness, industrial control protection, continuous monitoring, and secure interconnection between networks.

They include various National Institute of Standards and Technology (NIST) cybersecurity standards such as NIST SP 800-160 Volume 2 Rev. 1: Developing Cyber-Resilient Systems (Ross et al., 2021), and NIST Cybersecurity Framework (CSF) 2.0 (National Institute of Standards and Technology, 2024), as well as NIST SP 800-37 Rev. 2: Risk Management Framework (RMF) (Joint Task Force, 2018) which is more generally focussed on risk management.

ISO/IEC 27031:2025 (International Organization for Standardization, & International Electrotechnical Commission, 2025) is arguably the most relevant ISO standard for Information and Communication Technology (ICT) and network resilience. It focuses on ICT readiness for business continuity, helping organisations ensure that ICT systems can prevent, respond to, and recover from ICT-related disruptions affecting critical operations. Even though it connects to continuity, its actual focus is technical readiness of ICT infrastructure, services, and dependencies, which makes it highly relevant to systems and network resilience.

In Australia, these international standards are supported by a range of Australian frameworks. For this whitepaper the most relevant ones besides the earlier-mentioned SOCI Act are the Information Security Manual (ISM) (Australian Signals Directorate, 2026), the Australian Cyber Security Centre (ACSC) Principles of OT Cyber Security (Australian Signals Directorate, 2024), and the Critical Infrastructure (CI) Fortify guidance documents (Australian Signals Directorate [ASD], 2025).

ISM is the broadest and strongest technical cyber-resilience framework of this set. It is a framework organisations can apply, using their risk management framework, to protect IT and OT systems from cyber threats. It follows a full lifecycle approach, including system definition, control selection, implementation, assessment, authorisation, and monitoring, which makes it directly relevant to secure design, operation, and continuous resilience of systems and networks.

The ACSC Principles of OT Cyber Security is explicitly aimed at helping organisations design, implement, and manage OT so it remains safe, secure, and able to support business continuity for critical services, whereas CI Fortify (ASD, 2025) is a clear guidance for severe-disruption resilience in critical infrastructure. It emphasises two hard resilience goals: being able to isolate vital OT and enabling systems while maintaining critical services and being able to rapidly rebuild those systems if compromise cannot be confidently removed.

There are also standards specifically written for individual pieces of non-IT infrastructure that refer to resilience in one or the other way. They include, for instance, Australian Energy Sector Cyber Security Framework (AESCSF) (Australian Energy Market Operator, 2023) for energy, Australian Standard (AS) 7770 (Rail Industry Safety and Standards Board, 2018) for rail, and ISA/IEC 62443 (International Electrotechnical Commission, n.d.) for OT. They are out of scope for this paper, as we here focus on commonalities in digital resilience between sectors.



3. Relevant use cases from literature

3.1 Methodology

There are many examples documented of critical infrastructure failing because of issues with their IT and data communication infrastructure.

In this chapter we discuss the ones of the past ten years that tell us the most about the resilience measures in place before the incident, the cause of the incident (failure type), the magnitude of the impact and the duration of the service degradation, the resilience measures that were in place before the incident, and the measure taken during and after the incident.

We have also tried to cover a wide range of critical infrastructure sectors, including telecommunications, energy, transport and logistics, healthcare, government and digital infrastructure. We also looked at their growing interdependence, meaning that disruption in one sector can propagate into others (Gong et al., 2023). The selected cases cover network configuration changes, routing failures, software defects, destructive cyberattacks, ransomware, endpoint failures, power loss and physical infrastructure disruption.

In line with the previous chapter, resilience is assessed here as a time-dependent capability rather than by impact magnitude alone. For each incident, the analysis distinguishes the depth of performance loss, the duration of degraded service, the time to partial and near-full restoration, and any residual recovery period.

3.2 Results

Figure 2 provides a summary of our use case analysis.

The selected use cases are described in terms of the type of failure that started the chain of events, the magnitude of the impact, the recovery duration (where we distinguish initial recovery from full recovery), the resilience mitigations that were in place before the incident, and the mitigations that were put in after.

The analysis shows that around half of the cases were triggered by some failure in change or configuration management, be it network configuration change, routing configuration, timing issues, or other configuration failures. Changes to shared infrastructure can affect many services simultaneously, but their resilience implications differ greatly when recovery time is considered. For example, Telstra's July 2026 outage affected about 45% of calls and data sessions at peak, yet most services were restored the same morning (Brady, 2026). Rogers' July 2022 outage lasted up to 26 hours for some services, with root-cause identification delayed because the failed core also supported network management and staff communications (Xona Partners Inc., 2024). The latter highlights that it is not just connectivity to the infrastructure that is important, but also to the people who support, manage and restore the operation of the infrastructure. Cloudflare's 2022 routing-policy incident affected major data centres handling roughly half of the total number of requests but lasted only 75 minutes (Strickx & Hartman, 2022). AT&T's February 2024 configuration failure affected more than 125 million devices and persisted for at least 12 hours, because mass device re-registration overloaded control-plane systems even after the erroneous change was reversed (Federal Communications Commission, 2024). These cases demonstrate that resilience depends not only on the scale of disruption but also on fault isolation, rollback capability, management-plane independence, workforce enablement, and restoration capacity.

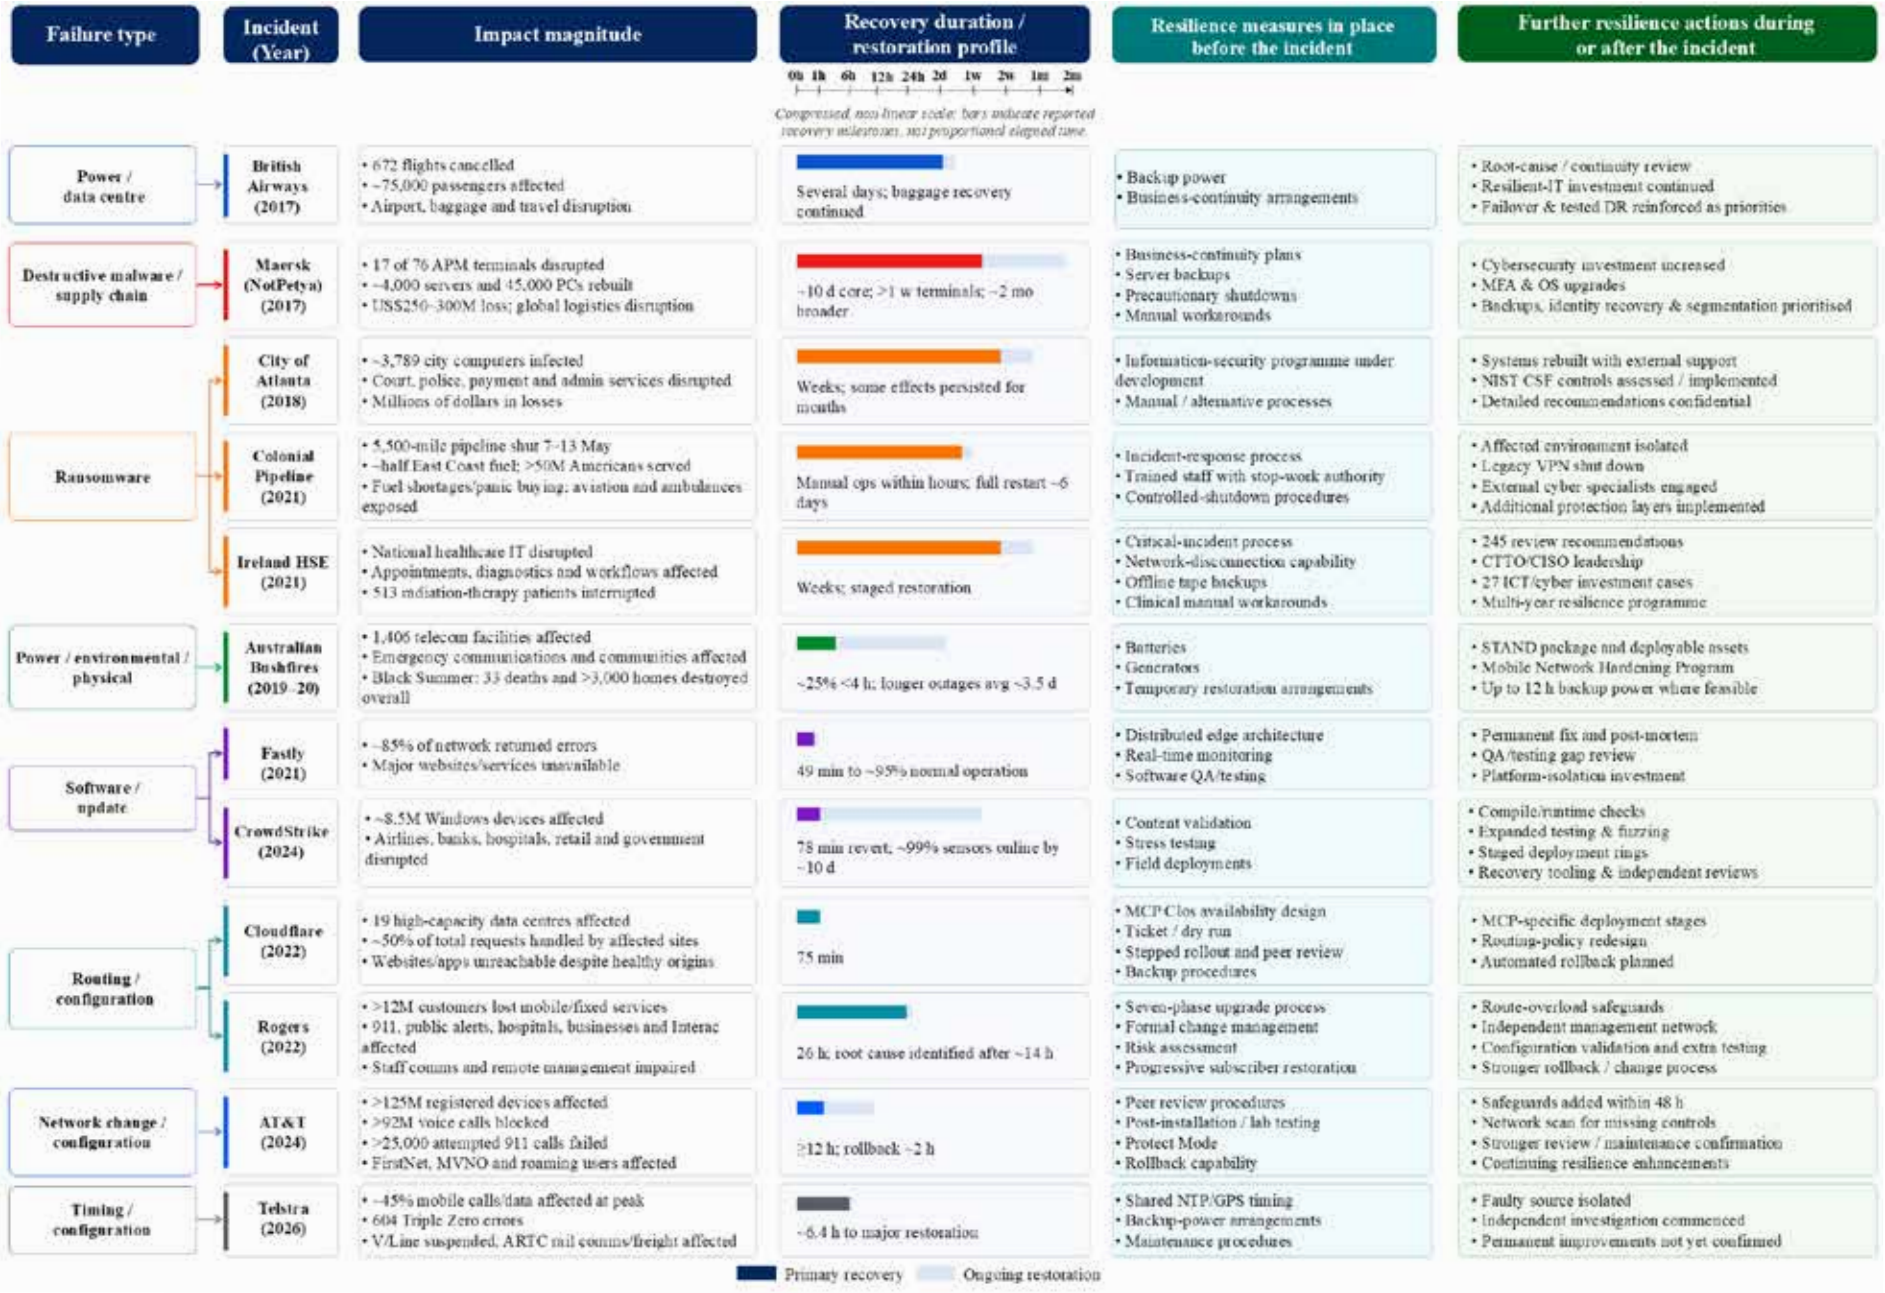


Figure 2: Failure types, impact magnitude, recovery duration and resilience measures across selected critical-infrastructure incidents. None of the incidents led to fatalities.

COLONIAL PIPELINE, UNITED STATES, 2021 | ENERGY: RANSOMWARE AND IT-OT DEPENDENCY

Colonial Pipeline supplies nearly half of the U.S. East Coast's fuel and serves over 50 million people. DarkSide ransomware compromised Colonial Pipeline's enterprise IT environment, leading the company to shut pipeline operations, causing fuel shortages.

Pre-incident resilience controls or measures: Colonial had a company-wide incident-response process, trained employees with stop-work authority, and controlled-shutdown procedures for situations in which system safety was uncertain. This enabled rapid shutdown and containment. Some local lines were returned to manual operations within hours. The entire system was restarted and product delivery commenced to all markets after approximately six days.

Further resilience actions during or after the incident: After identifying the likely entry point Colonial implemented additional layers of protection across the enterprise and engaged external cyber security experts to support remediation and strengthen its cyber security defences (Blount, 2021).



CrowdStrike's July 2024 incident illustrates the difference between removing a fault and restoring operations: defective Falcon content was withdrawn within 78 minutes, yet many Windows systems required manual intervention because they could no longer boot normally. Although many organisations recovered within hours, CrowdStrike reported that about 99% of Windows sensors were online only after ten days (CrowdStrike, 2024; Weston, 2024). This highlights that common software dependencies can produce rapid global failures, while recovery duration depends on whether remediation can be automated or requires endpoint-by-endpoint intervention.

Destructive cyberattacks stand out in Figure 2 as the cases with the longest recovery time. This is likely related to the fact that recovery from e.g. ransomware attacks may require rebuilding complete new trusted infrastructure rather than restarting services. The NotPetya attack spread through Maersk's interconnected Windows environment, disrupting booking, communications, terminals and logistics operations (Androjna et al., 2020) and causing losses estimated at US$250–300 million (A.P. Moller-Maersk, 2018). Maersk reportedly rebuilt approximately 4,000 servers and 45,000 computers within ten days, but broader software restoration continued for almost two months (Greenberg, 2018). Recovery depended on a surviving offline domain controller, demonstrating that resilience requires not only backups but also ensuring that identity, management and recovery assets do not share the same failure domain. Following the attack, Maersk strengthened controls such as multi-factor authentication and operating-system upgrades.

Ransomware incidents are a separate category of cyber attacks, as they demonstrate the complexity of the dependence of critical services on enterprise IT systems. Colonial Pipeline shut down operations following a ransomware attack in May 2021 and resumed full product delivery after roughly six days (U.S. Department of Energy, 2021). Although the shutdown was precautionary, it revealed that physical services could not continue normally when billing, monitoring and coordination systems were unavailable or untrusted (Janvrin et al., 2026). Ireland's Health Service Executive experienced disruption for several weeks, with clinical and administrative functions recovering at different rates (Flavin et al., 2022). Atlanta's municipal systems also recovered unevenly over weeks and, in some cases, months after the SamSam attack (Department of Justice, 2018). These incidents show that ransomware resilience depends on trusted recovery environments, segmentation, local service fallback and the ability to restore shared identity and application services in a controlled sequence.

Power, environmental and physical failures can affect both network and application availability. During the 2019–20 Australian bushfires, over a thousand telecommunications facilities were impacted (Australian Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts [ADITRDCSA], 2020). While about one-quarter recovered within four hours (ADITRDCSA, 2020), longer outages took several days because restoration was constrained by battery and generator endurance, fuel supply and physical access (Marshall et al., 2023). British Airways' 2017 data-centre power failure, compounded by an uncontrolled restoration of power that damaged servers, affected approximately 75,000 passengers and disrupted baggage operations beyond restoration of the flight schedule (Press Association, 2017). The case demonstrates that redundancy alone does not ensure resilience if recovery actions create common-mode failures, and reinforces the importance of independent power domains, controlled restart procedures, application and data failover, and regularly tested disaster-recovery arrangements.

Overall, the cases show that resilience depends on both impact magnitude and recovery time. Recovery ranged from less than 75 minutes for Fastly, to nearly two months for Maersk's broader restoration. But, most interestingly, most mitigations put in place after a major outage have focused on increasing initial robustness, and only few have significantly improved resilience.

Apparently, the initial response from the industry is still to focus investment on prevention, rather than on resilience.



CROWDSTRIKE, GLOBAL, 2024 | CYBERSECURITY: FAULTY SOFTWARE UPDATE

A defective Falcon content update caused Windows systems to crash, affecting an estimated 8.5 million devices (Weston, 2024). The update was reverted 78 minutes after release but affected endpoints did not recover automatically. Many systems returned within hours, but 99% of them were back online only after several days.

Pre-incident resilience controls or measures: CrowdStrike used content validation, stress testing and field deployments before release. They did not detect the parameter mismatch.

Further resilience actions during or after the incident: CrowdStrike introduced compile-time validation, runtime bounds checks, expanded testing and fuzzing, staged deployment rings, customer control over deployment timing, additional recovery tooling and independent reviews of its code, quality and release processes (CrowdStrike, 2024).

4. Interviews with critical infrastructure operators

4.1 Methodology

We have conducted 12 in-depth interviews with operators of Australian critical infrastructure in the following sectors:

- communications
- financial services and markets (here abbreviated as "finance")
- data storage or processing (here abbreviated as "data")
- energy
- space technology (here abbreviated as "space")
- transport
- water and sewerage (here abbreviated as "water")
- mining
- non-corporate government entities (here abbreviated as "government")

The top seven align with seven sectors as defined by the SOCI Act. We have added the mining sector and the non-corporate government entities' own critical infrastructure. Although government and mining sectors are not mentioned in the SOCI Act as critical infrastructure sectors, they often operate infrastructure that nevertheless must adhere to the SOCI obligations. We have omitted the health care sector, the education sector, defence industry, and the food and grocery sector for various reasons, but mostly because these sectors have such unique approaches in which they handle their infrastructure that it is nearly impossible to apply a standard list of questions and compare answers.

The interviews were conducted in mid-2026 via Microsoft Teams. The research team opted for an interview design with pre-defined questions, exploring the participant's understanding of resilience, network and system dependencies, previous incidents, adopted resilience measures, threat perceptions and future improvement areas. This design ensured that responses across different sectors could easily be compared whilst allowing the participants and interviewers to engage in the natural flow of conversation and explore answers more deeply.

The interview participants held an executive or senior level role where they are directly responsible for operational and business continuity, cyber resilience and/or service delivery. Roles include Chief Information Officer, System Architect, Head of Operations, and similar.

Chief Information Security Officers (CISOs) were not included, as we wanted to probe how cyber resilience was positioned against other forms of digital resilience at higher executive levels.

As we have anonymised and aggregated the results, the interviewees are only mentioned by the name of their sector. For the sectors Energy, Transport, and Mining we interviewed two executives from two different organisations. They are identified below as "Energy 1", "Energy 2", "Transport 1", etc.

4.2 Understanding of resilience

We asked the interviewees how they understood the concept of resilience for their organisation and the sector.

We were particularly interested in how they compare the importance of resilience of their IT and communication estate against other organisational risks. Across the nine sectors we found that the most critical infrastructure operators placed digital resilience on par with or higher than other organisational risks. When asked about how mature their organisation is today in relation to digital resilience on a scale of 1 to 5, we found a large spread of self-rated organisational resiliency maturity levels ranging from mid (2-3) to high (4-5). When asked later in the interview where they would like their organisation to be, participants are working towards a higher maturity level. A few

participants explicitly mentioned that obtaining a 5 would be great but specified that ideal maturity level of 4 was a more realistic target either due to sizing resilience investment to their risk profile or due to budget constraints.

The results for current and target organisation maturity for digital resilience are graphically shown in Figure 3. It should be noted that although 12 (interviewees) is a relatively small number, it is large enough to obtain statistically relevant results: the standard deviation in the results (indicated by the dark blue shade) is significantly smaller than the mean. Figure 4 shows how important the interviewees rate digital resilience compared to other organisational risk factors.

Most interviewees rate digital resilience as their main concern.

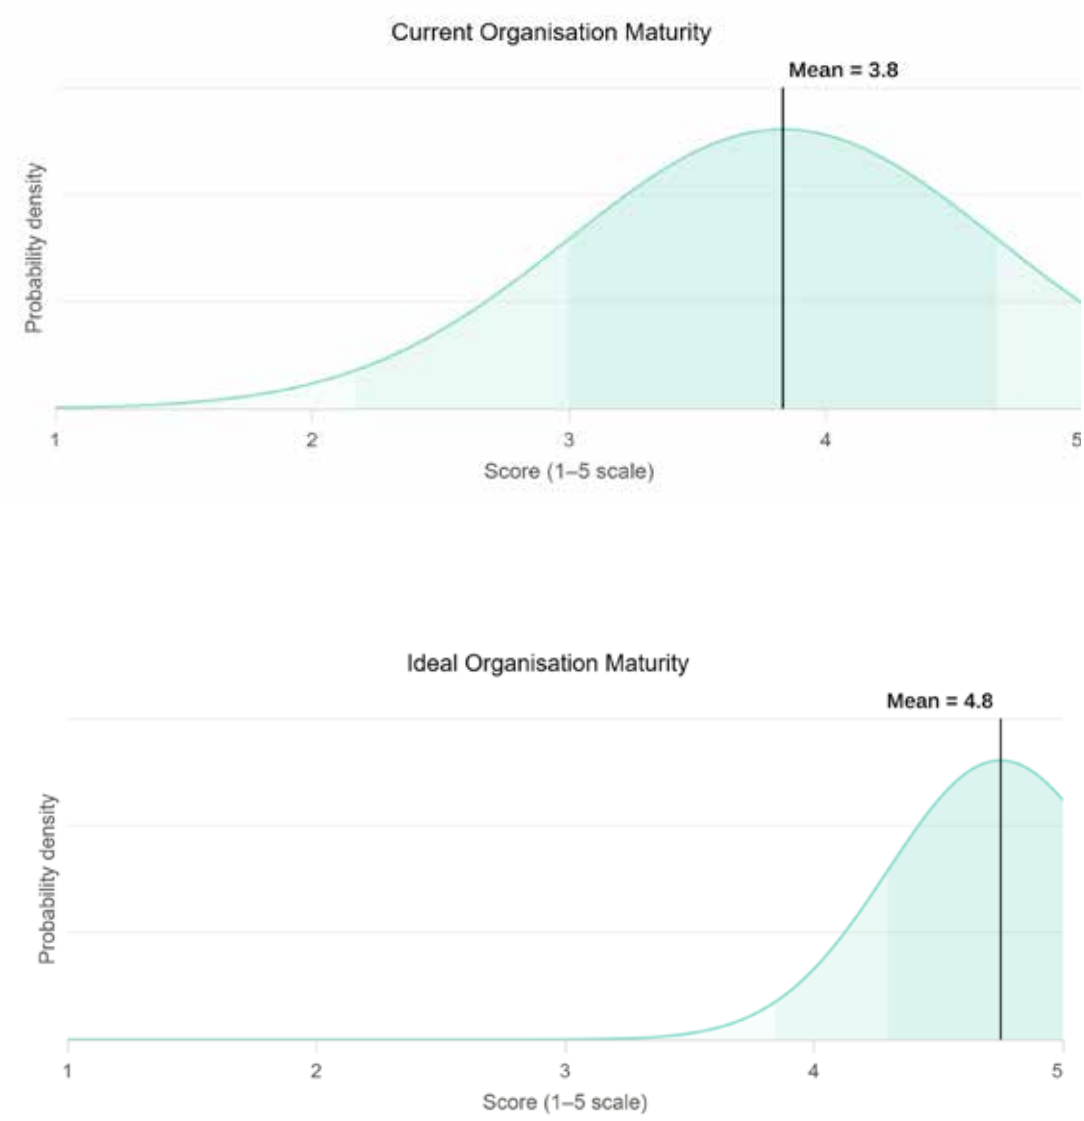


Figure 3: Top: current organisational maturity for digital resilience according to the interviewees. Bottom: target organisational maturity for digital resilience.

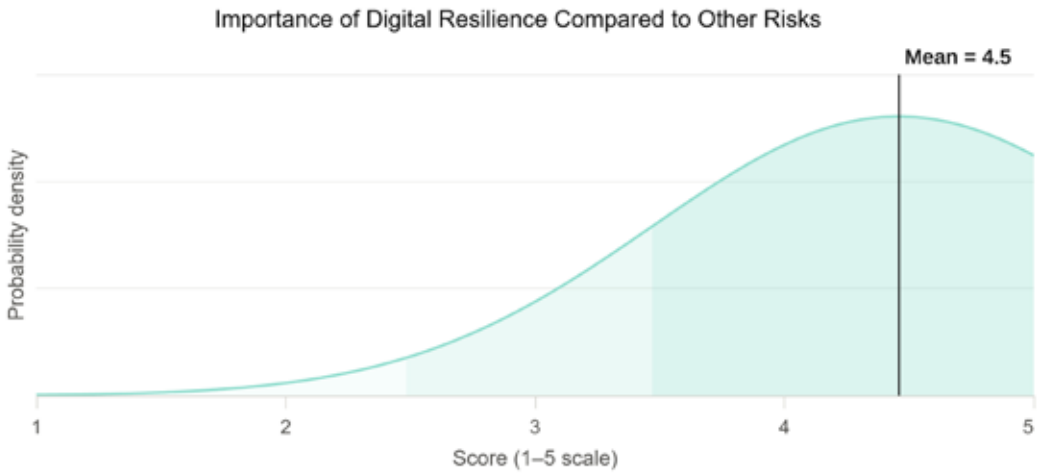


Figure 4: Importance of digital resilience compared to other organisational risks in critical infrastructure.

4.3 Impact quantification

When asked to quantify resilience in terms of the key business impact that needs to be minimised, respondents indicated that network connectivity interruption has the potential to completely halt or severely degrade core business operations.

Participants described impacts including production stoppages, transport gate closures, loss of customer services, and disruption to financial transactions and contractual obligations impacts (particularly for data centre and communications industry). Many interviews also highlighted safety and public risk consequences, particularly where communications support emergency services (triple zero), public safety, or operational field activities. Regulatory and compliance consequences were also identified, including licence obligations, regulatory scrutiny and contractual breaches.

A recurring cross-sector theme was the loss of operational control and visibility, with several organisations stating they would become effectively "blind" without connectivity, losing the ability to safely monitor, manage, dispatch or control critical assets and services.

These impacts frequently translated into production and revenue loss through reduced output, lost sales, contractual penalties and customer impacts.

Respondents reported significant variation in downtime tolerance, reflecting differences in service criticality. Some organisations identified recovery objectives or maximum tolerable disruption periods of one hour for their most critical services, while others reported tolerances extending to four, eight, twelve hours, or even several days. The lowest tolerance was consistently associated with emergency communications, public safety functions, Supervisory Control and Data Acquisition (SCADA) environments, process control systems, trading services and other operational technology platforms, while data centre environments described expectations of near-continuous availability.

Many organisations have established formal approaches to impact quantification, including Business Impact Assessments (BIA), Maximum Tolerable Period of Disruption (MTPD), Recovery Time and Recovery Point Objectives (RTO/RPO), service criticality assessments, major incident frameworks, risk matrices and service-level based methodologies. Financial impact assessment was a common component of these approaches, with disruption often measured through lost production, revenue impacts, financial penalties, customer outcomes and broader business disruption costs. While some participants identified a capability for

their organisation to fallback to manual operations as a resilience measure, these workarounds were generally described as degraded modes of operation that reduce efficiency, visibility and control rather than eliminating the underlying dependency on resilient network connectivity.

4.4 Network and system dependencies

Table 1 shows which types of connectivity the resilience of critical infrastructure mostly depends on.

Fixed Network Infrastructure includes (dark) fibre, WAN, and other fixed-line services and copper/legacy access where relevant. Private Mobile Networks includes private Long-Term-Evolution (LTE), private 5G, industrial wireless and organisation-owned wireless networks.

Public Mobile Networks includes mobile carrier services, public cellular voice/data, cellular failover and mobile backup connectivity. Radio and Operational Communications includes Digital Mobile Radio (DMR) and operational radio networks for safety-critical field communications. Satellite Connectivity includes Starlink and other Low Earth Orbit (LEO) and Geostationary Earth Orbit (GEO) satellite services, satellite phones, satellite backhaul and satellite resilience links. We reiterate that the Interview IDs are not sectors, but IDs for individuals. The results provided in this and the following tables are not representative for the whole sector.

Table 1 shows that most operators combine fixed network connections with a variety of wireless services. Regarding fixed networks, the specifics differ between operators. Some use dark fibre, sometimes redundant but not always, and some rent capacity from a telecommunication service provider.

Table 1: The primary network connections explored by participants.

Interview ID	Fixed Networks	Private Mobile Networks	Public Mobile Networks	Radio Comms	Satellite
Communications	✓		✓		✓
Transport 1	✓	✓			
Energy 1	✓			✓	✓
Mining 1	✓	✓			✓
Mining 2	✓	✓	✓	✓	✓
Data	✓				
Energy 2	✓	✓		✓	✓
Government	✓				
Water	✓		✓		✓
Finance	✓				
Space	✓				
Transport 2	✓		✓		✓

Recent incidents that drew significant attention from the media related to the reliance of critical services on public mobile networks. We, therefore, asked the interviewees to rate how much their reliance on public mobile networks and satellite networks would compare to their reliance on fixed networks. The results are shown in Figure 5. The spread in answers is quite large: in most cases, our critical services are not that reliant on mobile and satellite networks. The largest surprise was that some interviewees requested to revisit their answers after they realised that their dependency of Global Positioning System (GPS) also qualified as dependence on satellite systems. One interviewee worded this as follows, when asked about dependence on satellite coverage compared to fixed network: "Probably 2... but if you include GPS... then we're at a 5."

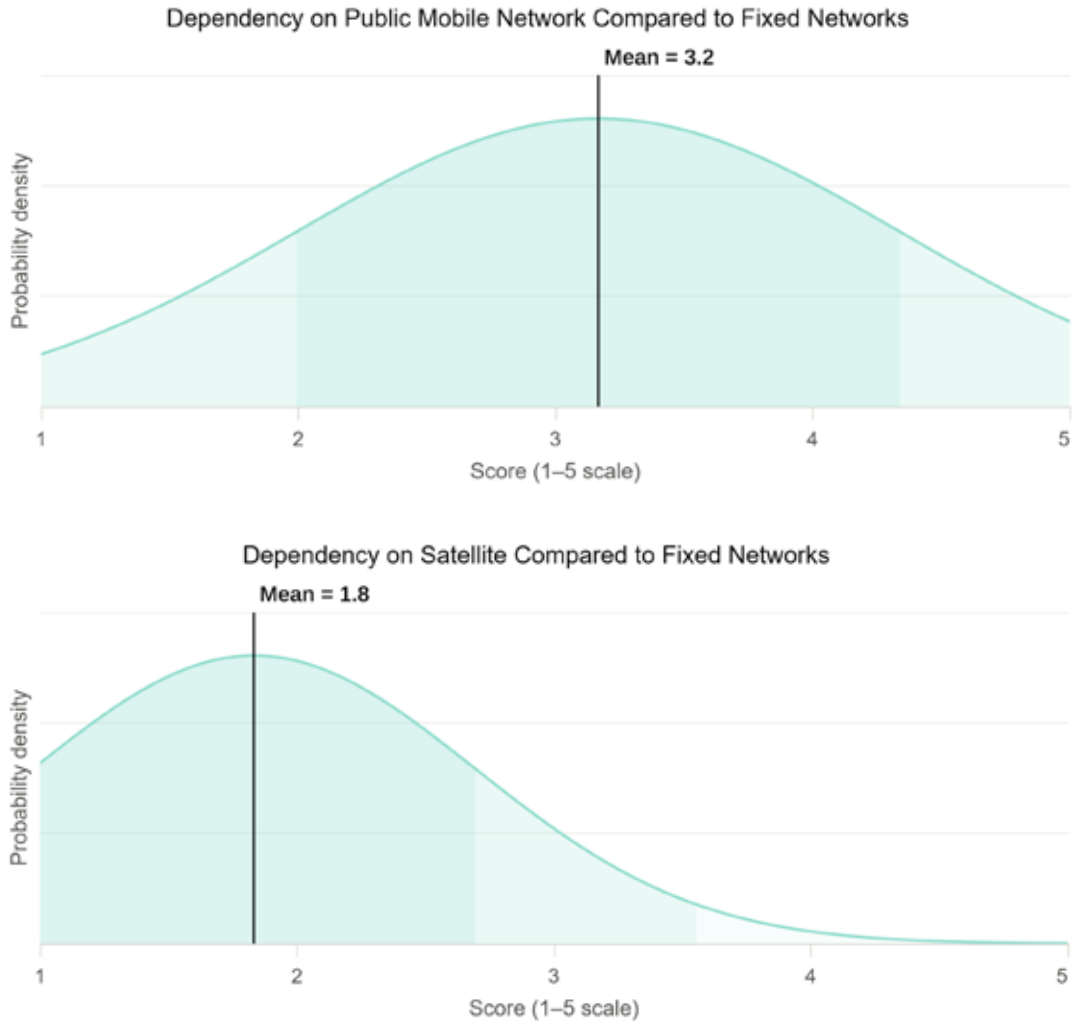


Figure 5: Top: Reliance on public mobile networks compare to fixed networks. Bottom: Reliance on satellite networks compared to fixed networks.

Table 2 shows the self-reported reliance on cloud and Software-as-a-Service (SaaS) products for critical operations. On-Premises means traditional on-prem systems, private cloud, organisation-operated data centres and internally managed infrastructure. Public Cloud are hyperscaler-hosted services, typically from Amazon Web Services, Microsoft, and Google. SaaS includes third-party managed applications such as SAP SaaS, Customer Relationship Management (CRM) systems, Microsoft 365, Oracle Cloud applications, specific Application Programming Interfaces (APIs) in the cloud, etc. The table shows a very mixed picture.

Every operator relies on cloud services in one way or the other, but their implementation varies greatly.

Table 2: Self-reported reliance on cloud and SaaS products for critical operations.

Interview ID	On-Prem	Public Cloud	SaaS
Communications	✓		
Transport 1	✓		✓
Energy 1	✓		
Mining 1	✓ (Critical Services)	✓	✓
Mining 2	✓ (Critical Services)	✓	✓
Data	—		✓
Energy 2	✓	✓	
Government	✓	✓	✓
Water	✓	✓	✓
Finance	✓	✓	✓
Space		✓	
Transport 2		✓	✓

Table 3: Self-reported single points of failure across all interviews.

Interview ID	Any known points of failure?	Details	Importance to resolve (1-5)
Communications	Yes	Access network	Not stated
Transport 1	Yes	Insider threat	Not stated
Energy 1	Yes (assumed)	Not known	1
Mining 1	Yes	GPS	4.5
Mining 2	Yes	GPS	4
Data	No	N/A	N/A
Energy 2	Yes	DMR (fallbacks are less reliable)	5
Government	Yes	Not stated	Not stated
Water	Yes	Telemetry network	1
Finance	Yes	Identity provider	Not stated
Space	No	N/A	N/A
Transport 2	Yes	Not stated	Not stated

We subsequently asked if the interviewee's system had a known single point of failure and if that needed urgent solving. The answers are provided in Table 3. Participants agreed that operating with single points of failure means to operate with elevated risk. Obviously, these answers only include the known single points of failure.

Actually, various interviewees only recognised a single point of failure (e.g. GPS, or Domain Name System (DNS)) after us asking explicitly.

One participant recalled that previous incidents, such as CrowdStrike, had exposed hidden single points of failure in their operations, which were then prioritised and fixed.

4.5 Previous incidents

This section provides an overview of incidents described by the participants that happened in their own infrastructure.

They are a mix of positive and negative experiences, in the sense of "our resilience measures worked well and as intended" (positive) and "we lacked the appropriate resilience measures to limit the impact to an acceptable level"(negative). The incidents are summarised below on a per sector basis.

DATA STORAGE & PROCESSING

The CrowdStrike outage impacted Microsoft-dependent systems across the organisation such as workstations. However, customer-facing services continued operating and critical facilities remained online. The organisation activated crisis-management procedures, relied on alternative communications including WhatsApp, deployed additional operational personnel and maintained service continuity while remediation activities were undertaken. Service recovery took approximately six hours. Customer services were reportedly unaffected. The incident subsequently drove greater technology diversity, including support for both macOS and Microsoft environments, and a move toward making operational facilities more autonomous and resilient.

ENERGY

A recent mobile network outage resulted in the loss of 4G connectivity across remote operational areas, affecting 15 field assets. Operators temporarily lost visibility of asset telemetry and operational data; however, assets were engineered to fail safely and automatically shut down where required. Staff reverted to DMR as a backup communication channel and operations continued without major production or safety impacts. The incident reinforced the importance of layered communications architecture incorporating LTE, DMR and satellite technologies.

GOVERNMENT

A physical accidental incident involving flooding from an internal bathroom disabled a Uninterruptible Power Supply (UPS) supporting critical infrastructure. Despite the loss of supporting equipment, existing redundancy in other branches functioned as designed and prevented any material impact on business operations.

FINANCE

During a network outage, recently integrated branch locations lost their primary connectivity path. Staff were able to continue operating through alternative internet and mobile connectivity options, including tethered devices, enabling customer services to continue with limited disruption. The lesson was to not simply have multiple links, but to ensure genuine diversity of connectivity methods and delivery paths.

SPACE

A cloud-zone failure affected the control and management plane of a critical system. The organisation identified the root cause and used the event to drive architectural change. The outcome was a transition toward multi-zone and multi-cloud infrastructure to reduce single point of failure risk and improve resilience against future cloud-service failures.

COMMUNICATIONS

In one example, customers were unable to register on the mobile network, or place calls, creating nationwide service disruption. This led to brand damage, significant financial cost and regulatory exposure. Root cause analysis identified software faults that prevented available capacity from being used effectively. The incident led to substantial changes in resilience practices, including N+2 capacity planning, resilience verification exercises and proactive failover testing where network components are deliberately taken out of service to validate recovery processes.

TRANSPORT

A fire at an external telecommunications exchange resulted in a six-hour communications outage. Without network connectivity, operational processes were significantly disrupted, with gate operations halted and the movement of trucks and trains unable to proceed normally. Logistics operations effectively stopped until communications were restored. The outage directly affected logistics throughput and vessel servicing activities until communications were restored later that morning.

MINING 1

A major technology failure involving storage infrastructure and a power disruption corrupted both sides of a redundant environment, creating a loss of visibility and control across operational systems. Personnel were unable to safely operate parts of the processing environment and workers were evacuated from an underground operation. Critical services were restored in approximately six hours, while full recovery took around two weeks. The outage resulted in significant financial impacts in addition to operational disruption and safety risk and prompted substantial changes to business continuity planning, operational procedures and recovery preparedness.

MINING 2

The event discussed originated from a failed network change where a cable was connected into the wrong port, creating a network loop that propagated throughout the environment. The resulting outage disrupted mine-site operations for approximately four hours and affected rail, port, processing and supporting technology services. It took approximately one hour to identify the fault, after which the cable was removed, network devices were rebooted and services progressively restored. The incident drove tighter change-management controls, greater contractor oversight and improved operational governance.

WATER

Multiple resilience events were described, including a backhoe striking a fibre link, firewall-instability issues, architectural weaknesses within core networking infrastructure and the wider impacts of the CrowdStrike event. The incidents affected connectivity to water-treatment environments, SCADA traffic and operational technology systems, creating periods where operators lost visibility of critical services. In response, the organisation introduced greater carrier diversity, cellular backup capability, distributed core-network architecture and firewall-design improvements.

TRANSPORT 2

The CrowdStrike event created widespread operational disruption across industry, requiring coordination with customers, government and logistics partners to prioritise critical freight and maintain deliveries. In response, the organisation relied on business continuity arrangements and has since increased its focus on cross-sector cyber exercises, coordination planning, governance processes and resilience testing across interconnected supply-chain partners.

4.6 Resilience standards and frameworks

This section describes which standards and frameworks the interviewees use when designing and developing resilience measures and procedures.

The results of this question are summarised in Table 4. Here, "NIST" refers to any NIST or a NIST-aligned framework. "ISO" refers to ISO standards, ISO certification, or specific ISO standards such as ISO/IEC 27031 but also ISO 22301, or ISO 27001. "ITIL" refers to Information Technology Infrastructure Library (ITIL) and vendor standards, or other industry frameworks that are not ISO/NIST. "Regulatory Frameworks" refers to compliance obligations, sector-specific regulation, the SOCI Act, critical infrastructure risk management, or similar regulatory requirements. Interviewees also refer to jurisdiction-specific "Jurisdictional Frameworks", such as state/territory-based cyber/security frameworks, and to "Internal Frameworks", i.e. the organisation has created or maintains its own framework, standards, or methodology (e.g. business continuity plans) while referencing established frameworks such as ISO or NIST.

The results are somewhat surprising: there is not a clear "winner" when it comes to standards and frameworks.

There seems to be no apparent logic in which is applied when.

4.7 Testing and training

This section describes the answers of the interviewees on their testing and training practices.

The results are summarised in Table 5. Failover Testing means validating traffic/services switch to alternate paths and backups can be restored. Recovery Testing means validating disaster recovery and service restoration. Tabletop Exercises means discussion-based scenario testing, and Live Simulations means controlled outages or real-world drills. standards and frameworks. There

Table 4: Frameworks underpinning resilience management as reported by participants.

Interview ID	NIST	ISO	ITIL	Regulatory Frameworks	Jurisdictional Frameworks	Internal Frameworks
Communications		✓				✓
Transport 1	✓	✓		✓		
Energy 1						
Mining 1	✓					
Mining 2		✓				
Data	✓	✓		✓		✓
Energy 2	✓	✓				✓
Government						
Water		✓		✓	✓	
Finance			✓	✓		
Space		✓				
Transport 2				✓		✓

seems to be no apparent logic in which is applied when.

Participants reported role-based training, operational playbooks, major incident management procedures, and regular resilience exercises to prepare staff for operational disruptions. In addition to scheduled failover and recovery testing, some participants noted that full failover or 'chaos monkey' testing is commonly undertaken as part of major capital projects or significant changes.

Across the interviews, the frequency of these tests ranged from monthly to annually with one participant distinguishing frequency based on local and regional testing activities. One participant shared that their organisation rotates simulation scenarios annually, such as focusing on cyber incidents one year and third-party or supply chain failures the next. Participants also described a range of testing approaches, from tabletop exercises and disaster recovery simulations to software and infrastructure testing. One participant called out that they have failover and backup options, but they are not tested. In the event of an incident or outage, participants shared a range of alternative communication methods to support staff coordination, including Short Message Service (SMS), digital radio, and messaging platforms such as WhatsApp.

Overall, the findings across the interviews suggest that resilience depends on both technical recovery capabilities and the ability of staff to operate effectively under degraded conditions.

Table 5: Testing activities and frequencies reported by participants.

Interview ID	Failover & Backup Testing	Recovery Testing	Tabletop Simulation	Real-world Simulations
Communications	✓ (annually)	✓ (annually)	✓	Not stated
Transport 1	✓ (annually)	✓ (annually)	✓	✓ (annually)
Energy 1	✓ (monthly)	Not stated	✓	Not stated
Mining 1	✓ (annually)	Not stated	Not started	✓ (annually)
Mining 2	Not stated	✓ (annually)	✓ (annually)	✓ (annually)
Data	✓ (annually)	✓ (annually)	✓ (quarterly)	✓ (annually)
Energy 2	Not stated	✓ (annually)	✓ (as needed)	Not stated
Government	✓ (annually)	Not started	Not stated	Not stated
Water	✓ (annually)	✓ (annually)	Not stated	Not stated
Finance	✓ (annually)	✓ (annually)	Not stated	Not stated
Space	✓ (monthly local / annual regional)	✓ (monthly local / annual regional)	Not stated	Not stated
Transport 2	None	Not stated	✓ (annually)	✓ (as needed)

4.8 Cyber threats, cascading failures, and physical threats

Interview participants were asked to individually rate how concerned (with 5 being very concerned) they were about resilience issues after cyber incidents, physical incidents (such as fibre cuts, equipment damage and natural disasters) and cascading failures (such as the CrowdStrike incident from 2024).

The results are summarised in Figure 6.

From a cross-sector perspective we found a strong agreement across our participants that cyber threats followed by cascading failures are perceived as greater risks to network operation than physical threats. Participants expressed confidence in their organisation's ability to remain operational in the event of physical incidents such as fibre cuts, storms, fires and natural disasters. These participants often referenced existing resilience measures, including redundancy, alternate connectivity options, established recovery processes, and prior experience managing physical disruptions. Physical incidents are generally viewed as known and anticipated operational challenges. In contrast, increased levels of concern regarding cyber security and cascading failure stem from a lack of confidence in the organisation's ability to recover and maintain service continuity during and after such events.

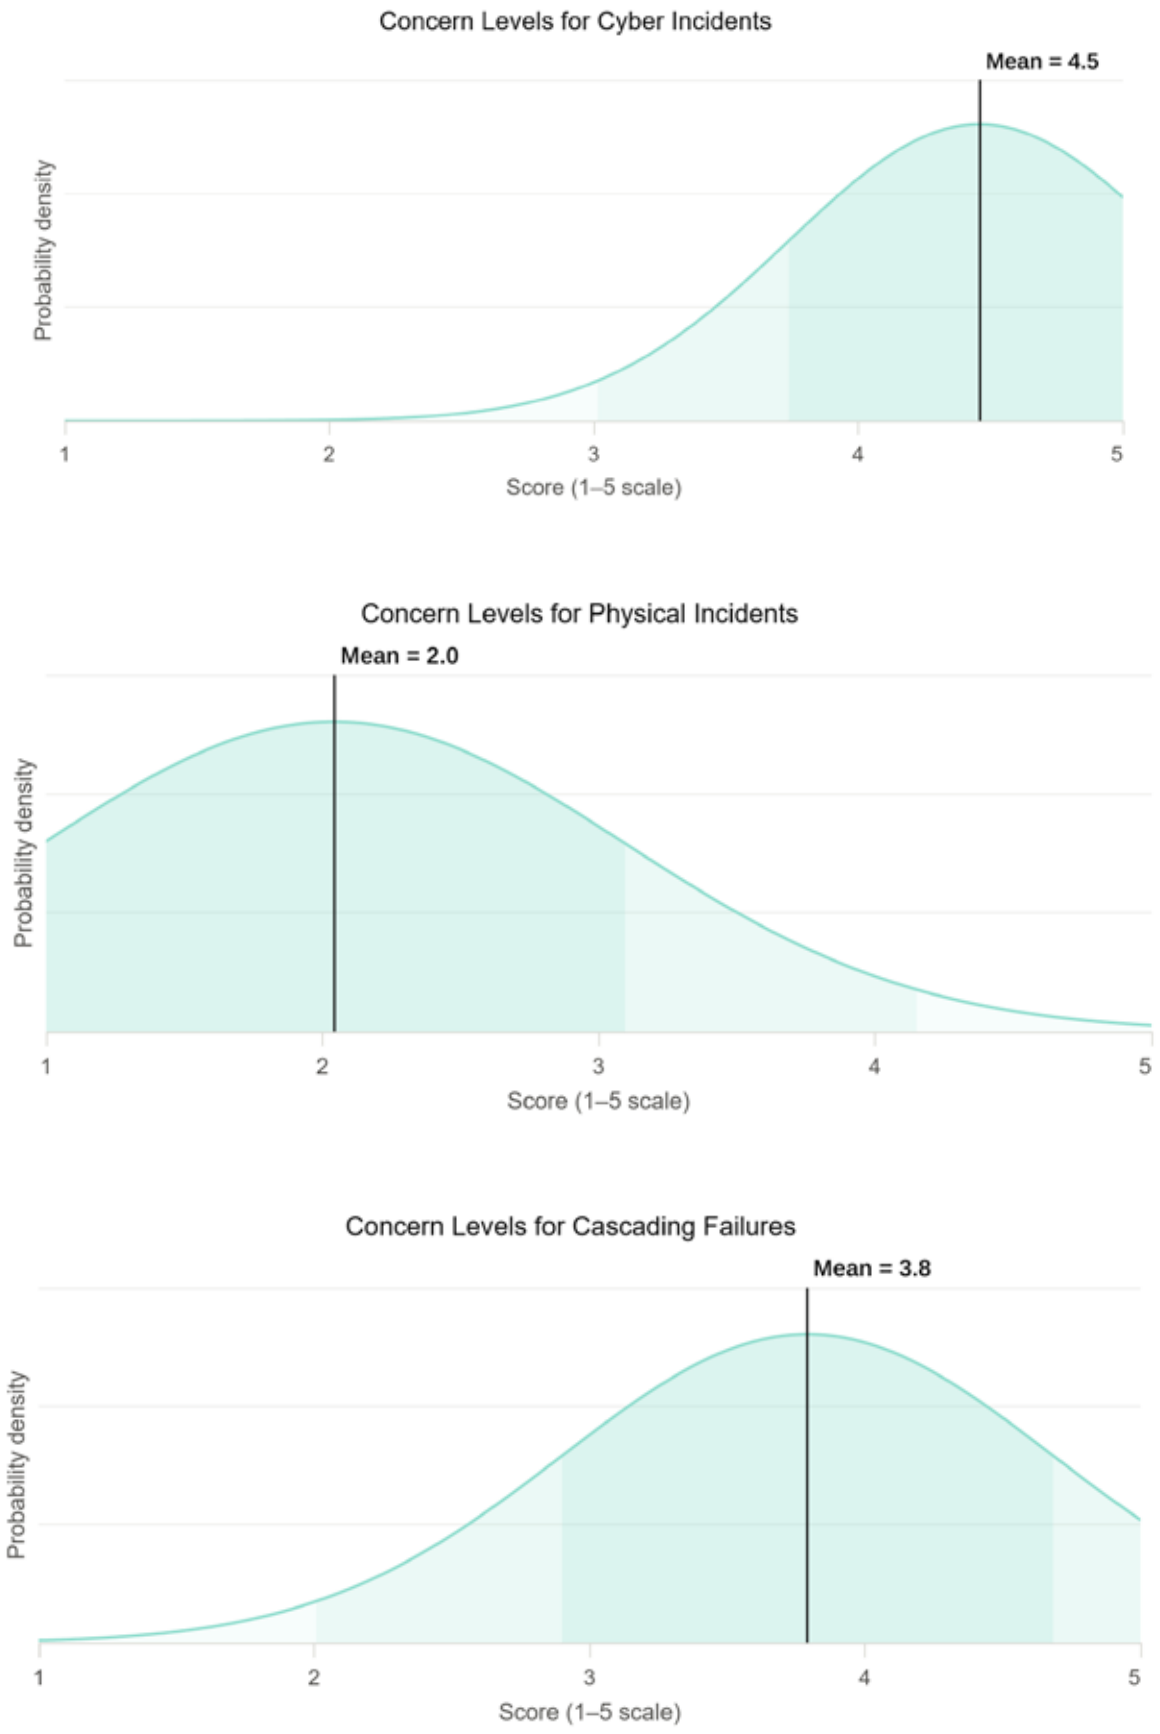


Figure 6: Rate of concern regarding resilience to physical threats (top), cyber threats (middle) and cascading failure (bottom).



5. Analysis

5.1 Migration strategies and change management

Most failures start with changes and updates being made to the existing infrastructure.

This was the case with six out of the twelve use cases described in Figure 2, and of those, five were the most recent use cases overall. It suggests that improvements need to be made in change management procedures, despite change management being quite established as an expertise with well thought through protocols.

However, with critical infrastructure systems becoming more complex than ever, we believe that more effort needs to be spent on maintaining full control. Particularly the extended interconnectedness of systems in multiple managed domains often leads to new dependencies that are not known or well understood (see also our discussion below on cascading failures). Change management and migration procedures and protocols need to be continuously reviewed and updated to capture those new dependencies.

At the same time, a continuous effort should be made in making systems less complex rather than more, for instance by regularly reviewing the architecture of the (systems of) systems, reviewing how dual vendor policies increase resilience but also complexity, and accelerating the migration of legacy systems. In other work, we already emphasised the need to replace legacy systems (particularly the ones that are end-of-life) because of increased cyber security risks fuelled by AI (den Hartog et al., 2026). This would also be an opportunity to declutter the digital estate and reduce the number of hidden dependencies.

5.2 Lack of joint approaches

One of our key findings is that most operators of critical infrastructure take a relatively isolated approach to improving resilience.

There is some awareness of standards, which then are applied as seen fit, but none of the interviewees reported active participation in setting and improving those standards or, on a lower level, in regular structured formal or informal discussions and knowledge sharing on digital resilience.

Instead, most operators rely heavily on their historic approach to resilience, the belief that they are all "unique", that there is little common best practice, and that improving resilience is largely organisation specific. However, we found several common themes in the literature and amongst operators that would justify best-practice sharing and joint approaches. They include developing a common understanding of dependency on GPS and on public mobile networks, of the changing dynamics around data centres, and of the interdependency between critical infrastructures. The sectors should also look at the applicability of international standards to Australian critical infrastructure, defining a common framework of resilience metrics and maturity model that are appropriate to the Australian critical infrastructure context. Finally, considering resilience as a part of the public good of critical infrastructure operations may require greater collaboration between otherwise competing service providers in the operators' supply chains, requiring discussions around antitrust, confidentiality, and shared accountability.

The government has already created forums such as the Trusted Information Sharing Network (TISN) to connect the leaders of many organisations with obligations under SOCI, to help coordinate and develop responses to build cyber resilience. However, we are not aware of forums that look cross-sector and at digital resilience beyond cyber and connect those at the practitioner level. The aim of such forums should be to

develop common practices, share insights and knowledge and, on occasion, provide feedback to regulators and others where, in practice, there are unintended counterproductive consequences of recommendations and directives. Organisations such as CI-ISAC were not mentioned by the interviewees, and although not limited to a single sector, CI-ISAC appears to focus on the health-care sector in its external communications, and does not discuss resilience (CI-ISAC, 2026). Since 2022, the financial sector is running resilience exercises as per the Cyber Operational Resilience Intelligence-led Exercises (CORIE) Framework (Council of Financial Regulators, 2022). While CORIE is also restricted to cyber resilience, and uses a much wider definition for resilience, it mentions "Gold Team Exercises" in its final chapter, which could possibly be applied to other sectors as well.

We, therefore, strongly recommend the institution of forums where organisations can structurally collaborate on the level of engineering practitioners with other critical infrastructure operators within the own sector as well as with other sectors on determining best practices in digital resilience, moving away from the current pragmatic, isolated, history-based, piecemeal approaches. To avoid the policy vs practice gap (den Hartog et al., 2026), a positive success factor is the inclusion of practitioners in these discussions, and the recognition that a balance needs to be found between antitrust regulation and the need to increase resilience across Australia's critical infrastructure estate.

5.3 Deeply integrated production and recovery architectures

A recurring theme in the literature as well as the interviews was that optimal resilience requires a minimal functional recovery capability, including monitoring and observability, and recovery functionality such as automatic failovers and backup activation.

However, production and recovery systems are often deeply integrated, with monitoring functionality etc. often implemented in-band. Thus, when disaster happens, it often debilitates the recovery capability in its stride. We recommend that future investments should prioritise making the most fundamental part of the monitoring and observability infrastructure significantly more robust and resilient against threats, preferably being independent of the entire estate.

This recommendation is well aligned with the ASD CI Fortify guidelines (ASD, 2025). While CI Fortify largely focusses on isolating OT from IT in its advice for isolating vital systems, we add that these guidelines may also inform the architecture of the most fundamental monitoring and observability and other recovery infrastructure.

We also recommend that the overall level of observability is increased in terms of the ability to obtain more, and more fine-grained data and information about events and their effect on performance. Successful recovery is partly a matter of being able to quickly understand and isolate the root cause of the event, and having access to high-quality data from before and during the event has often been a game changer in previous incidents.



5.4 Single points of failure and cascading failure

Critical infrastructure has become heavily dependent not only on their own IT and network infrastructure but on complete supply chains.

Banking, telecommunications, healthcare, transport, and emergency services now operate atop layered software, cloud and network ecosystems built and maintained across hundreds of independent providers. This includes the interdependencies between critical infrastructure sectors, i.e. communications on power, mining on transport, etc. Recent incidents have demonstrated how fragile this arrangement is. Particularly the earlier discussed CrowdStrike 2024 use case is a typical example of this, but also the more recent Telstra use case shows elements of this fragility: these incidents are known as cascading failures. Cascading failures are events in which the failure of one component in an interdependent system, triggers further failures in connected components, generating failure that progressively spreads via positive feedback, resulting in partial or total system collapse (Ulrich, 2017; Chamberlin, 2022).

Cascading failure is well studied in the field of physical infrastructure, where research has produced topology-based models of how localised faults propagate into cascading failures across coupled networks (Buldyrev et al., 2010; Valdez et al., 2020). Such fault-propagating properties are much harder to detect in modern large interconnected dynamic IT supply chains, due to their global, complex, and highly dynamic nature.

Nowadays, large interconnected dynamic IT supply chains have a dependency structure that accumulates incrementally through reuse of components, integration with other services and third-party adoption rather than through centralised design; they are organically grown systems-of systems. For example, a typical application nowadays may depend on dozens of cloud services and software packages, with most of its dependencies or connection inherited indirectly, rather than chosen explicitly, becoming hidden dependencies unbeknownst to system operators (Zimmermann et al., 2019; Decan et al., 2019; McGuire, 2025).

Three elements dictate how cascading failures operate in organically grown systems-of-systems. First: hidden dependencies create blind spots and pathways through connections that are inherited indirectly, for example, via shared cloud services or software packages. They map out where a system is brittle and vulnerable to the second condition: single points of failure. Often, these connections are deeply embedded within large interdependent systems that no one system operator can fully map or have total awareness of these dependencies (Zimmermann et al., 2019; Schueller & Wachs, 2024). Third: positive feedback loops are the primary mechanism that drive and accelerate a cascading failure through the paths of hidden dependencies once a single point of failure is triggered (Ulrich, 2017; Chamberlin, 2022). When a component degrades or fails, its load or function redistributes to connected nodes. This redistribution overloads the remaining system, causing subsequent failures to cascade.

The resilience of a system therefore depends on both the extent of that degradation and how long it takes for lost capability to be recovered (Brtis et al., 2026). The consequences of a cascading failure are, therefore, not captured by the scale of the initial loss alone, and Figure 1 as an illustration of the definition of resilience is then not applicable anymore. Figure 7 provides a better representation: the initial degrading event produces a chain reaction of dependent events leading to further degradation and a longer time to recover.

These structural dependencies also affect the resilience of the wider IT supply chain. Particularly in large interconnected dynamic IT supply chains, the resilience of an individual system is partly shaped by dependencies beyond the system boundary visible to its operator. Accordingly, where those dependencies are inherited, hidden, or poorly understood, the effect of an upstream failure may be difficult to anticipate before a cascade occurs. Assessing resilience only within the boundary of an individual system can therefore overlook hidden dependencies that materially influence how that system responds to failure.

This also has a consequence for what often is regarded as the first phase in responding to a failure: falling back to some fail-safe default. It is cited as an effective action to prevent further cascading failure (Segovia-Ferreira, 2024), but if the dependencies are not known, the fallback may very well fail.

This issue has come up several times in the use cases research as well as the structured interviews. But the current literature landscape lacks a unified methodology to proactively model these elements in large interconnected dynamic IT supply chains. Common methods explored include topology-based network models to statically model dependency structures to retroactively piece together cascading failures. Other methods include anomaly detection models which achieve strong results in single-operator environments, but are designed to trigger only after anomalous behaviour, such as a cascade, has commenced, rather than predicting or detecting it prior (Soldani & Brogi, 2022). This is a knowledge gap that needs to be addressed urgently.

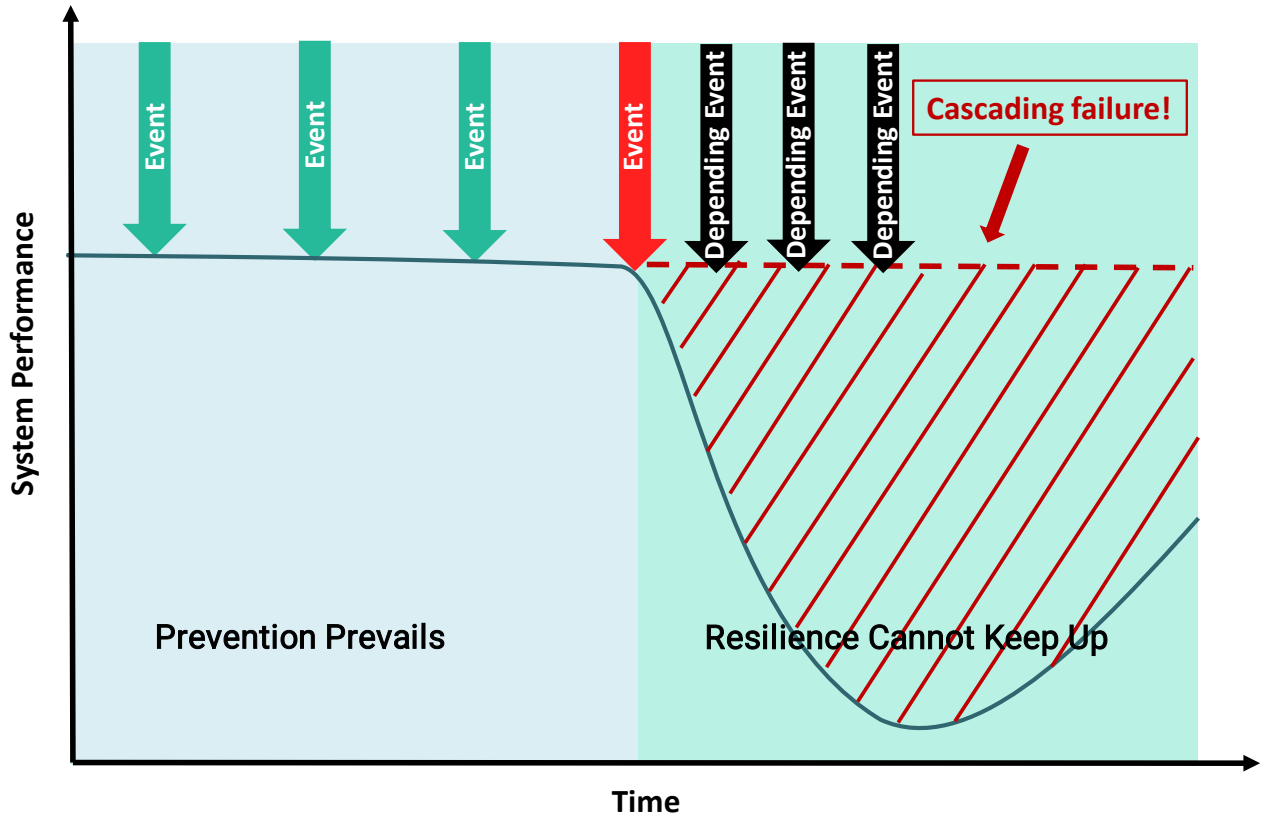


Figure 7: Traditional resilience measures often cannot keep up with the chain reaction of events triggered in a cascading failure incident, creating a much larger impact and over a much longer period than anticipated.

6. Conclusions, future work, and recommendations

6.1 Conclusions

In this whitepaper we examined how Australia's critical infrastructure sectors understand, implement, and assess digital resilience in an environment characterised by increasing digitisation, interdependence, and exposure to cyber and operational risks.

Through a combination of literature-based case study analysis and twelve in-depth interviews with critical infrastructure operators across multiple sectors, we have shown that digital resilience is now regarded as a core, all-of-business, organisational capability rather than a purely technical concern. Interview participants consistently ranked digital resilience among their highest strategic priorities, reflecting the extent to which modern critical services depend on interconnected IT (in the cloud as well as on-prem), OT, and communication systems.

While redundancy remains an important design principle, historical incidents repeatedly show that duplicate systems alone do not guarantee continuity of operations. Failures at various organisations illustrate that loss of service is often driven by shared dependencies, software supply chains, management-plane failures, configuration errors, or recovery constraints rather than the absence of backup infrastructure. The most resilient organisations were not necessarily those with the greatest level of technical duplication, but those with the strongest capabilities for detection, isolation, recovery, adaptation, and operational decision-making. Apparently, redundancy and diversity are becoming increasingly complex and multifaceted and must be understood in detail if they are to be relied on as a tool to increase resilience.

A recurring theme throughout both the case studies and interviews is the growing importance of hidden dependencies and cascading failures. Critical infrastructure operators increasingly depend on complex ecosystems of cloud services, on-prem systems, networks, systems for authentication, localisation, and timing, and on third-party suppliers. Many of these dependencies are indirect and only become visible during major incidents. Our findings suggest that traditional resilience approaches, which focus primarily on individual systems or organisational boundaries, are then no longer sufficient. Resilience must increasingly be understood and managed at the level of interconnected systems-of-systems.

The interviews also revealed substantial variation in resilience practices, standards adoption, organisational maturity, and consistency in approaches to measuring resilience. While most organisations employ combinations of ISO, NIST, regulatory, and internally developed frameworks, there is no common approach across sectors. Similarly, testing practices vary considerably, ranging from frequent failover exercises and live simulations to environments where backups or recovery arrangements are only partially validated. This inconsistency suggests that many organisations rely on pragmatic and historically grown practices rather than industry best practices.

A significant finding of this study is the critical role of maintaining a minimum level of observability and situational awareness also after the event. Across multiple incidents, recovery was delayed not by the initial fault itself but by an inability to rapidly identify and understand the (evolving) failure and observe the recovery progress.

Future investments should focus on making monitoring and observability infrastructure significantly more robust and resilient against threats, for instance by isolating them from the production systems where possible.

Digital resilience should increasingly be determined by an organisation's ability to understand complex dependencies, manage change safely, maintain visibility across interconnected systems, and recover based on industry best practices when disruption occurs. Future resilience strategies should therefore move beyond infrastructure duplication alone and focus on intelligence, observability, adaptation, and cross-sector awareness. As Australia's critical infrastructure continues to digitise, these capabilities will be essential to maintaining the availability, safety, and trustworthiness of the services on which our society depends.

6.2 Future work

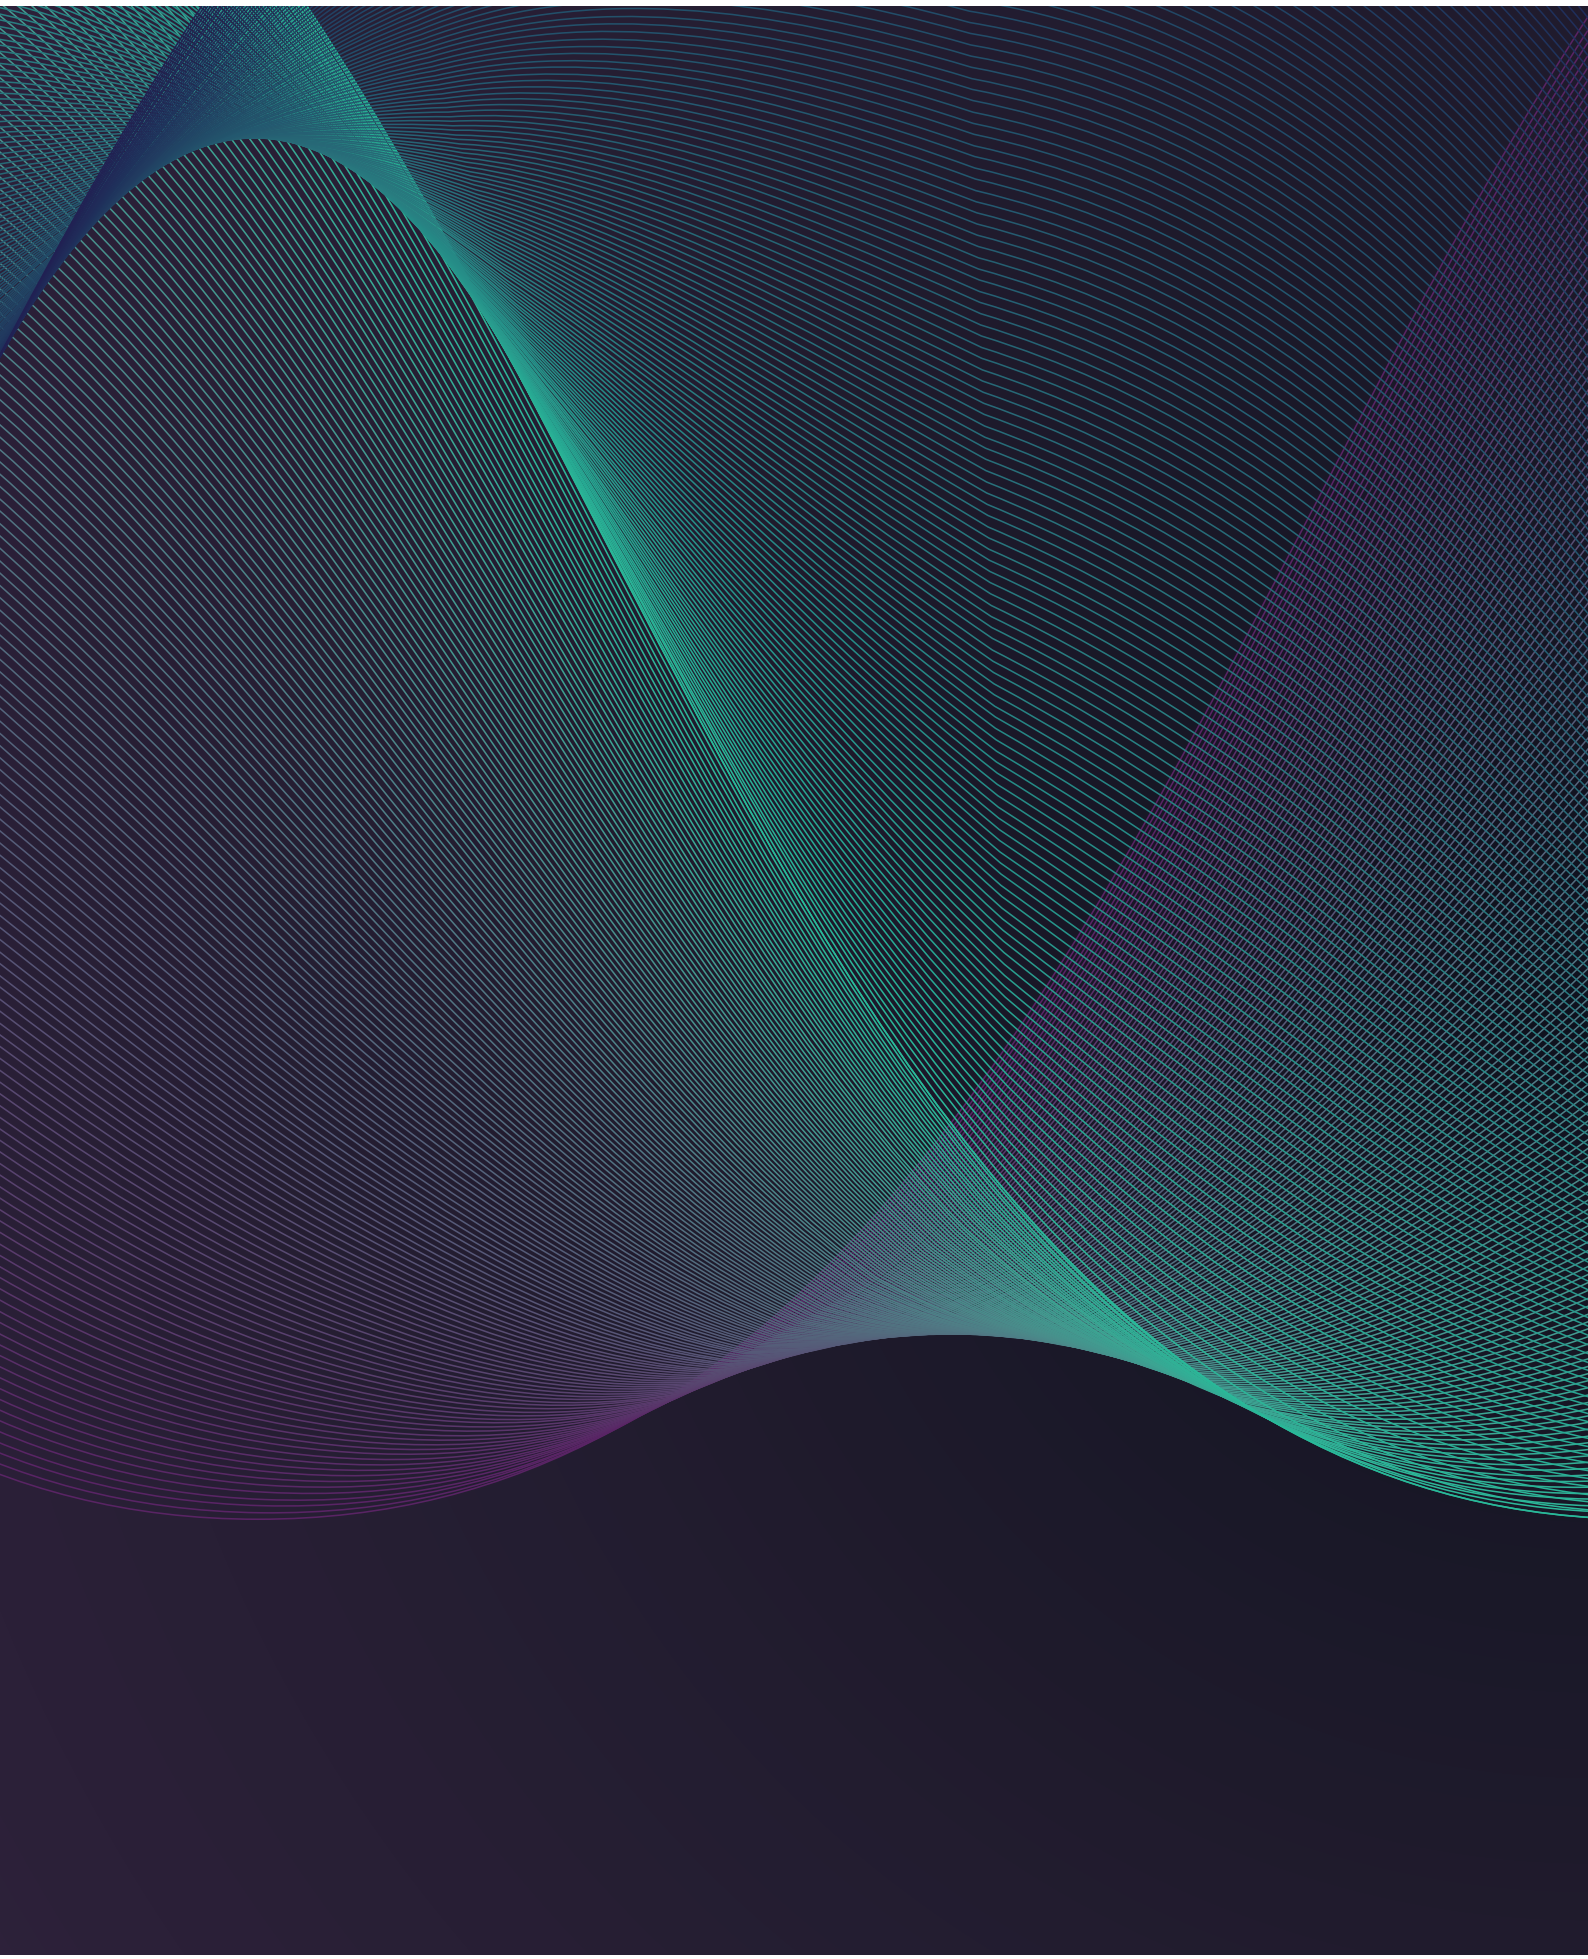
As this study has analysed the current literature as well as the experience of practitioners in the field, there has been little discussion on how resilience operations may have to change due to the introduction of AI: AI in the IT of critical infrastructure operators as well as AI used by future malicious cyber security actors.

AI will likely introduce new failure domains including trustworthiness, explainability, model drift, data poisoning, autonomous decision making, workforce displacement and concentration risk. Many of the current approaches to resilience will also apply to a world with AI. But new approaches may need to be developed too, which may also be based on AI. For instance, AI could enhance network telemetry, assurance, and dependency mapping. This is for further study.

The reported inconsistency in resilience implementations, standards adoption, approaches to measuring resilience, and consequent testing practices can partly be attributed to a lack of understanding in how business outcomes (the impact of failure on customers) and technical measures (e.g., infrastructure uptime) exactly relate to each other. This can lead to a growing gap between provider metrics and critical service resilience. There is a clear need to investigate and develop stronger metrics that go across the technology stack and directly relate to business impact. For instance, there needs to be industry-wide alignment on key resilience and recovery definitions, including RTO, RPO, and tolerance levels. Greater clarity in these areas will help reduce ambiguity, strengthen accountability, and minimise conflict regarding resilience expectations and obligations.

Recent incidents related to the reliance of critical services on public mobile networks drew significant attention from the media. However, our results suggest that such reliance is only moderate. This discrepancy can be explained as follows. First, we asked the interviewees to make a comparative assessment, i.e., to look at their complete digital estate and to compare how much they rely on mobile networks compared to other infrastructure, particularly fixed networks. Second, public mobile networks have historically been designed to provide best-effort services at mass-market scale. Yet we observe organisations increasingly use mobile connectivity to support critical operations, including, incident response and recovery. Further work is needed to better understand how such operational dependencies align with the design characteristics of public mobile networks.

Several case studies show that restoring infrastructure is often easier than coordinating the workforce that handles impacted systems, customers, media, regulators, etc. while the services are technically restored. Resilience is increasingly a people problem as much as it is a technology problem. The role of workforce readiness and organisational response capability requires further exploration. One of the outcomes could be an accelerated reliance on public mobile networks, not only for critical services but also for workforce enablement.



6.3 Recommendations

From our work we can distil the following recommendations for the critical infrastructure industry to

01

RECOMMENDATION 1

Focus investment on monitoring and observability infrastructure in two distinctly different ways:

a. obtain more, and more fine-grained data and information about events and their effect on performance, as well as on dependencies between functions in systems, systems of systems, and automated supply chains, and

b. invest in this very observability and monitoring infrastructure to be the most robust, independent, and resilient of the entire estate.

02

RECOMMENDATION 2

Improve change management and migration strategies and practices.

Most failures in resilience start with changes and updates being made to the existing infrastructure. With critical infrastructure systems having become more complex than ever, more effort needs to be spent on obtaining full control and understanding of how to make changes and how to migrate to new systems. Ideally, a continuous effort should be made in making the system less complex rather than more.

03

RECOMMENDATION 3

Continuously analyse hidden dependencies and single points of failure.

Resilience efforts will be seriously hampered when events trigger each other, and particularly when the chain of events hits a single point of failure. More time and effort need to be spent on continuously analysing the system (of systems) on hidden dependencies and single points of failure, before they become apparent through a disastrous event.

04

RECOMMENDATION 4

Drive greater collaboration between the engineering practitioners across different critical infrastructure operations to grow the maturity and adoption of best practice frameworks, common metric sets and widely accepted maturity models. Move away from the current pragmatic, isolated, history-based, piece-meal approaches.

05

RECOMMENDATION 5

Leverage common metrics, maturity models and practice to **continuously validate the actual level of resilience in critical infrastructure operations.**

Acknowledgements

The research leading to this whitepaper was supported by Telstra Limited and by Cisco Systems Inc. It has been greatly strengthened by the insights, expertise and generous support of the many contributors who engaged with us throughout its preparation. We are deeply grateful to our twelve interviewees as well as experts from Telstra and Cisco, who provided thoughtful feedback, shared valuable operational perspectives, and challenged our assumptions in ways that helped sharpen and refine our analysis.

While this paper has benefited greatly from the time, expertise, and collaboration of all involved, the views, interpretations and recommendations expressed herein are solely those of the authors. They do not necessarily reflect the positions or perspectives of any organisation or individual who contributed to the review process.

This research was proudly supported by Telstra and Cisco

Contact

For further information regarding this whitepaper, its analysis or recommendations, please contact:

National Industry Innovation Network

National Industry Innovation Network: communications@niin.com.au

Innovation Central Canberra

A collaboration led by cisco

FOUNDATION PARTNERS

cisco | University of Canberra

Innovation Central Canberra: innovationcentral@canberra.edu.au

References

Androjna, A., Brcko, T., Pavić, I., & Greidanus, H. (2020). Assessing cyber challenges of maritime navigation. *Journal of Marine Science and Engineering*, 8(10), 776.

A.P. Moller-Maersk. (2018). 2017 annual report. https://investor.maersk.com/system/files-encrypted/nasdaq_kms/assets/2018/04/25/13-00-21/A.P._Moller_-_Maersk_Annual_Report_2017.pdf

Australian Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts. (2020). *Findings from the 2019–20 bushfire season.* Retrieved August 2, 2026, from https://www.infrastructure.gov.au/media-technology-communications/phone/communications-emergencies/findings-from-2019-20-bushfire-season

Australian Energy Market Operator. (2023). *Australian energy sector cyber security framework (AESCSF): Framework overview*. https://www.aemo.com.au/-/media/files/initiatives/cyber-security/aescsf/2023/the-2023-aescsf-overview.pdf

Australian Government. (2026a). *Security of Critical Infrastructure Act 2018*. https://www.legislation.gov.au/C2018A00029/latest/versions

Australian Government (2026b). *Memorandum of Understanding on Digital Resilience to Support the Security of Australia between The Commonwealth of Australia, as represented by the Department of Home Affairs and Microsoft Corporation (Microsoft).* https://www.homeaffairs.gov.au/cyber-security-subsite/files/memorandum-of-understanding-on-digital-resilience-to-support-the-security-of-australia.pdf

Australian Signals Directorate. (2024). *Principles of operational technology cyber security. Australian Government.* https://www.cyber.gov.au/publication/principles-of-operational-technology-cyber-security

Australian Signals Directorate. (2025). *CI Fortify: Guidance for Australian critical infrastructure service continuity and resilience*. Australian Government. https://www.cyber.gov.au/publication/ci-fortify-guidance-for-australian-critical-infrastructure-service-continuity-and-resilience

Australian Signals Directorate. (2026). *Information security manual (September 2026)*. Australian Government. https://www.cyber.gov.au/business-government/asds-cyber-security-frameworks/ism

Blount, J. (2021, June 8). *Testimony of Joseph Blount, President and Chief Executive Officer, Colonial Pipeline Company.* U.S. Senate Committee on Homeland Security & Governmental Affairs. https://www.hsgac.senate.gov/wp-content/uploads/imo/media/doc/Testimony-Blount-2021-06-08.pdf

Brady, V. (2026, July 17). *Triple Zero Senate Inquiry: Opening statement from Vicki Brady.* Telstra. Retrieved August 2, 2026, from https://www.telstra.com.au/exchange/triple-zero-senate-inquiry--opening-statement-from-vicki-brady

Brtis, J., Cureton, K., Jackson, S., & Taylor, I. (2026). *System resilience.* In: The Guide to the Systems Engineering Body of Knowledge (SEBoK). Retrieved on August 9, 2026, from https://sebokwiki.org/wiki/System_Resilience

Buldyrev, S.V., Parshani, R., Paul, G., Stanley, H.E., and Havlin, S. (2010) .Catastrophic cascade of failures in interdependent networks, *Nature*, 464(7291):1025–1028. https://doi.org/10.1038/nature08932 .

Chamberlin, B.W. (2022). *How Failures Cascade in Software Systems*, Brigham Young University, https://scholarsarchive.byu.edu/etd/9474/ .

CI-ISAC (2026), Resources for Critical Infrastructure Organisations, Retrieved on 9 October 2026 from https://ci-isac.org.au/resources/

Council of Financial Regulators. (2022, July). *Cyber Operational Resilience Intelligence-led Exercises (CORIE)® framework: Program guide for financial institutions (including financial market infrastructure) in Australia* (Version 2.0). https://www.cfr.gov.au/publications/policy-statements-and-other-reports/2022/revised-corie-framework-rollout/pdf/corie-framework.pdf

Critical Infrastructure Security Centre. (2026). *Our regulatory principles and approach*. Australian Government. https://www.cisc.gov.au/legislation-regulation-and-compliance/our-regulatory-principles-and-approach

CrowdStrike. (2024, August 6). *External technical root cause analysis—Channel File 291*. https://www.crowdstrike.com/content/dam/crowdstrike/www/en-us/wp/2024/08/Channel-File-291-Incident-Root-Cause-Analysis-08.06.2024.pdf

Cyber and Infrastructure Security Centre. (2023). *Critical Infrastructure Resilience Strategy.* https://www.cisc.gov.au/resources-subsite/Documents/critical-infrastructure-resilience-strategy-2023.pdf

Decan, A., Mens, T., & Grosjean, P. (2019). An empirical comparison of dependency network evolution in seven software packaging ecosystems. *Empirical Software Engineering*, 24(1), 381-416.

den Hartog, F., Liu, E., Gullotta, D., & Sloan, S. (2026). *Securing the Nation: Accelerating government cyber security in an age of digital disruption.* NIIN. https://doi.org/10.26191/00nq-3j74

Federal Communications Commission. (2024, July 22). February 22, 2024 AT&T Mobility network outage: Report and findings. https://docs.fcc.gov/public/attachments/DOC-404150A1.pdf

Flavin, A., O'Toole, E., Murphy, L., Ryan, R., McClean, B., Faul, C., McGibney, C., Coyne, S., O'Boyle, G., Small, C., Sims, C., Kearney, M., Coffey, M., & O'Donovan, A. (2022). A national cyberattack affecting radiation therapy: The Irish experience. *Advances in Radiation Oncology*, 7(5), 100914.

Gandhi, P., Khanna, S., & Ramaswamy, S. (2016, April 1). Which industries are the most digital (and why)? *Harvard Business Review*, 1, 45-48.

Gong, S., Ye, Y., Gao, X., Chen, L., & Wang, T. (2023). Empirical patterns of interdependencies among critical infrastructures in cascading disasters: Evidence from a comprehensive multi-case analysis. *International Journal of Disaster Risk Reduction*, 95, 103862.

Greenberg, A. (2018, August 22). The untold story of NotPetya, the most devastating cyberattack in history. *Wired.* Retrieved August 4, 2026, from https://www.wired.com/story/notpetya-cyberattack-ukraine-russia-code-crashed-the-world/

International Electrotechnical Commission. (n.d.). IEC 62443 series: *Industrial communication networks—Network and system security*. https://webstore.iec.ch/en/iec-search/result?q=IEC%2062443

International Organization for Standardization, & International Electrotechnical Commission. (2025). *Cybersecurity—Information and communication technology readiness for business continuity* (ISO/IEC Standard No. 27031:2025). https://www.iso.org/standard/27031

Janvrin, D. J., Boss, S. R., & Gray, J. (2026). Colonial Pipeline: Responding to a ransomware attack in the energy sector. *Issues in Accounting Education*, 41(2), 131–145.

Joint Task Force. (2018). Risk management framework for information systems and organizations: A system life cycle approach for security and privacy. *NIST Special Publication 800-37*, Rev. 2. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-37r2

Khaloopour, L., Su, Y., Raskob, F., Meuser, T., Bless, R., Janzen, L., ... & Jamali, V. (2024). Resilience-by-design in 6G networks: Literature review and novel enabling concepts. *IEEE access, 12*, 155666-155695. https://doi.org/10.1109/ACCESS.2024.3480275

Kott, A., Dubynskyi, G. Y., Paziuk, A., Galaitsi, S. E., Trump, B. D., & Linkov, I. (2024). Russian cyber onslaught was blunted by Ukrainian cyber resilience, not merely security. *Computer*, 57(8), 82–89. https://doi.org/10.1109/MC.2024.3404568

Marshall, A., Wilson, C.-A., & Dale, A. (2023). Telecommunications and natural disasters in rural Australia: The role of digital capability in building disaster resilience. *Journal of Rural Studies*, *100,* 102996.

McGuire, M. (2025). T*ransitive dependencies are exponentially increasing open source software risk*, Black Duck website, retrieved on April 24, 2026, from https://www.blackduck.com/blog/managing-transitive-dependencies-open-source-software.html

National Institute of Standards and Technology (2024) The NIST Cybersecurity Framework (CSF) 2.0., *NIST Cybersecurity White Paper* (CSWP) NIST CSWP 29. https://doi.org/10.6028/NIST.CSWP.29

Press Association. (2017, June 1). British Airways IT failure caused by 'uncontrolled return of power.' The Guardian. Retrieved August 4, 2026, from https://www.theguardian.com/business/2017/may/31/ba-it-shutdown-caused-by-uncontrolled-return-of-power-after-outage

Rail Industry Safety and Standards Board. (2018). *Rail cyber security* (AS 7770:2018). https://www.ariso.org.au/products/as-7770-rail-cyber-security/

Rockwell, N. (2021, June 8). *Summary of June 8 outage.* Fastly. Retrieved August 2, 2026, from https://www.fastly.com/blog/summary-of-june-8-outage

Ross, R., Pillitteri, V., Graubart, R., Bodeau, D., & McQuaid, R. (2021). *Developing cyber-resilient systems: A systems security engineering approac*h (NIST Special Publication 800-160, Vol. 2, Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-160v2r1

Schueller, W., & Wachs, J. (2024). Modeling interconnected social and technical risks in open source software ecosystems. *Collective Intelligence*, 3(1). https://doi.org/10.1177/26339137241231912.

Segovia-Ferreira, M., Rubio-Hernan, J., Cavalli, A., & Garcia-Alfaro, J. (2024). A survey on cyber-resilience approaches for cyber-physical systems. *ACM Computing Surveys,* 56(8), 1-37.

Soldani, J., & Brogi, A. (2022). Anomaly detection and failure root cause analysis in (micro) service-based cloud applications: A survey. *ACM Computing Surveys (CSUR)*, 55(3), 1-39. https://doi.org/10.1145/3501297 .

Strickx, T., & Hartman, J. (2022, June 21). *Cloudflare outage on June 21, 2022.* Cloudflare. Retrieved August 2, 2026, from https://blog.cloudflare.com/cloudflare-outage-on-june-21-2022/

Ulrich, M. (2017), *Addressing Cascading Failures*, Google SRE website, retrieved March 15, 2026, from https://sre.google/sre-book/addressing-cascading-failures/

U.S. Department of Energy. (2021, May). *Colonial Pipeline cyber incident*. Retrieved August 2, 2026, from https://www.energy.gov/ceser/colonial-pipeline-cyber-incident

U.S. Department of Justice. (2018, November 28). *Atlanta U.S. Attorney charges Iranian nationals for City of Atlanta ransomware attack.* Retrieved August 2, 2026, from https://www.justice.gov/usao-ndga/pr/atlanta-us-attorney-charges-iranian-nationals-city-atlanta-ransomware-attack

Valdez, L.D., Shekhtman, L., La Rocca, C.E., Zhang, X., Buldyrev, S.V., Trunfio, P.A., Braunstein, L.A., and Havlin, S. (2020), Cascading failures in complex networks, *Journal of Complex Networks*, 8(2), https://doi.org/10.48550/arXiv.2007.11075 .

Weston, D. (2024, July 20). *Helping our customers through the CrowdStrike outage*. Microsoft. Retrieved August 2, 2026, from https://blogs.microsoft.com/blog/2024/07/20/helping-our-customers-through-the-crowdstrike-outage/

Xona Partners Inc. (2024, July). *Assessment of Rogers networks for resiliency and reliability following the 8 July 2022 outage.* Canadian Radio-television and Telecommunications Commission. Retrieved August 2, 2026, from https://crtc.gc.ca/eng/publications/reports/xona2024.htm

Zimmermann, M., Staicu, C. A., Tenny, C., & Pradel, M. (2019). Small world with high risks: A study of security threats in the npm ecosystem. In *28th USENIX Security symposium (USENIX security 19)*. 995-1010.


NIIN
cisco
UNIVERSITY OF CANBERRA