Resilient Networks for Critical Services

Jan Marius Evang

Faculty for Technology, Art and Design
OsloMet – Oslo Metropolitan University
Norway

Spring 2023

OsloMet Avhandling 2024 no 1

ISSN 2535-471X
ISBN 978-82-8364-581-1

OsloMet – Oslo Metropolitan University

University Library

Skriftserien

St. Olavs plass 4,

0130 Oslo,

Telephone (47) 64 84 90 00

Mailing address:

Postbox 4, St. Olavs plass

0130 Oslo

Printed at Byråservice.

Printed on Scandia 2000 white, 80 grams on material pages/200 grams on the cover.

Acknowledgements

I would like to express my sincere thanks to SimulaMet and CRNA for giving me the opportunity to pursue this PhD research. Additionally, I am deeply grateful to Media Network Services for their generous support, including sharing data, providing computing power, and offering network capacity for conducting measurements, all of which greatly facilitated my research endeavors.

I am indebted to my esteemed supervisors, Haakon Bryhni, Foivos Michelinakis, and Olav Lysne, for their unwavering support, insightful discussions, and invaluable guidance throughout the duration of this study.

I would also like to extend my heartfelt appreciation to my colleagues at the Center for Resilient Networks and Applications, particularly my co-authors Azza H. Ahmed, Ahmed Elmokashfi, Ioana Livadariu, and Taric Cicic. Their collaboration, expertise, and input significantly contributed to the quality and breadth of the research.

Scott and Rebekah Allendevaux imparted their extensive knowledge of ISO27001 and risk management, providing me with invaluable insights and encouragement. Without their guidance, I may not have been able to carry out this work with the same level of thoroughness.

Last but certainly not least, I would like to express my deepest appreciation to my wife, Kaja, and my daughters, Sara, Ruth, and Rakel, for their never-ending patience, understanding, and support. Tragically, after the submission of this thesis, my beloved wife Kaja passed away, leaving a profound void in our lives. Her unwavering encouragement and love sustained me throughout this journey, and her absence is deeply felt. I am forever grateful for her immeasurable impact on my life and work.

Preface

The research leading up to this thesis was performed at Simula Metropolitan Center for Digital Engineering (SimulaMet), under the supervision of Haakon Bryhni, Foivos Mihaelinakis, and Olav Lysne. The work is funded by Oslo Metropolitan University and Simula Metropolitan Center for Digital Enginering.

The research is presented in seven peer-reviewed articles. At the time of thesis submission, Papers I was published in the Proceedings of the Applied Networking Research Workshop at IETF-114 (ACM/IRTF), Paper II was published in the Proceedings of the 3rd International Workshop on Information Management (AEIS/IEEE), and Paper V was published in the Proceedings of the 20th International Conference on Security and Cryptography (INSTICC/SciTePress). Papers III and IV have been accepted for publication in the conference proceedings of the 31st International conference on Software, Telecommunications and Computer Networks in September 2023 (IEEE), while Papers VI is under consideration for the ACM hotnets workshop, and Paper VII is under review for publication in IEEE Communications Magazine.

Jan Marius Evang
August 2023
Oslo, Norway

A revised version of Paper VI was accepted for publication at the 9th International Conference on Complexity, Future Information Systems and Risk (COMPLEXIS 2024), and a revised version of Paper VII was accepted for publication at the 4th International Conference on Computer and Communications Engineering (CCCE 2024). Language wash by OpenAI.

Abstract

Network services play a pivotal role in today's society, serving the needs of businesses, governments and for individuals in their daily life. While we often take the seamless functionality of the Internet for granted, its growing use by Critical Services underscores the escalating importance of comprehending both resilience and security challenges.

The interconnected networks that make up the Internet are operated by various actors such as enterprises, governmental agencies, and content delivery networks (CDNs), in addition to global and local Internet Service Providers (ISPs). Maintaining these network services is a complex task. Numerous components could potentially disrupt the service, and it is essential for network service operators to understand the risks associated with each component.

Network security encompasses three fundamental objectives: Confidentiality, Integrity, and Availability [1]. Confidentiality and integrity are often addressed together, due to shared common attack vectors and mitigation solutions. Ensuring availability, however, presents a distinctive challenge. The primary focus of availability is to guarantee that the network service remains operational and usable. Although breaches in confidentiality and integrity can have indirect effects on availability, the nature of risk mitigation strategies differs significantly. In this setting, resilience and redundancy are central concepts.

Together, the papers in this thesis analyse the complete risk landscape applicable to delivering a resilient network for critical services. A majority of the research is performed on the Media Network Services (MNS) global video conferencing network, chosen for its relevance to risk management, and the applicability of results to other network operators.

Papers I and VI use 18 months of measurement data to analyse the root causes of network outages, revealing that the most important outages stem from leased Internet links, physical faults, and human errors. In contrast, relatively few are attributed to local network faults or malicious attacks. This insight into the root causes serves as a

foundational understanding for subsequent analyses. Paper II presents 5 years of risk registry data highlighting the role of management standards like ISO27001 in risk reduction, showcasing their efficacy in fostering a robust risk management framework across various organizational levels. Paper III delves into the intricate domain of Internet risks, demonstrating effective mitigation strategies to enhance network resilience against outages, packet loss and high latency originating from the Internet. Paper IV establishes a co-variation between organisations' security implementations and adherence to two security standards, Mutually Agreed Norms for Routing Security (MANRS) and ISO27001. By verifying Resource Public Key Infrastructure (RPKI) participation, IP spoofer protection, and Internet risk scores for organizations adhering to MANRS and/or ISO27001, we demonstrate that a security-aware company culture is connected to better security practices.

Recent paradigm-shifting incidents like COVID-19 and the Russian incursion into Ukraine demonstrate the importance of considering governance risks. Paper VII extends the scope to encompass national governance risks, specifically the high dependency of national web services on foreign micro services and cloud services, highlighting the imperative of considering broader contextual factors.

Drawing from the collective insights of these papers, combining the theoretical analyses with experiments on an operational network and real-life experiences, Paper V emerges as a synthesis, proposing an innovative cohesive 10-layer model that pragmatically organizes identified risks. This model stands as a testament to the integration of empirical findings into a practical framework, and the results can be generalized to a range of different networks. By utilizing the 10-layer model, network operators will reduce their availability risk and deliver a higher quality service to their customers.

Sammendrag

Nettverkstjenester har en viktig rolle i dagens samfunn. Bedrifter, myndigheter og enkeltpersoner er avhengige av Internett i sitt daglige virke.

Nettverkene som til sammen utgjør Internett, driftes av forskjellige aktører som bedrifter, offentlige etater og innholdsnettverk (CDN-er), i tillegg til globale og lokale internettleverandører (ISP-er). Drift av disse nettverkstjenestene er en kompleks oppgave. Tallrike komponenter kan potensielt forstyrre tjenesten, og det er viktig for nettverkstjenesteoperatører å forstå risikoen knyttet til hver komponent.

Nettverkssikkerhet omfatter tre grunnleggende mål: konfidensialitet, integritet og tilgjengelighet [1]. Konfidensialitet og integritet blir ofte behandlet sammen, da de har felles angrepsvektorer og løsninger. Å sikre tilgjengelighet er en separat utfordring. Hovedfokuset for tilgjengelighet er å garantere at nettverkstjenesten er operativ og kan brukes. Selv om brudd på konfidensialitet og integritet kan ha indirekte påvirkning på tilgjengelighet, er typen risikoreduserende strategier forskjellig, og redundans er sentralt.

Sammen analyserer artiklene i denne avhandlingen det komplette risikolandskapet som skal til for å levere et stabilt nettverk for kritiske tjenester. Store deler av forskningen er utført på Media Network Services' (MNS) globale videokonferansenettverk. Dette nettverket er valgt på grunn av relevansen for risikostyring, og anvendeligheten av resultater for andre nettverksoperatører.

Paper I og VI bruker 18 måneders måledata for å analysere de grunnleggende årsakene til nettverksavbrudd, og avslører at de viktigste bruddene stammer fra leide linjer, fysiske feil og menneskelige feil. Derimot tilskrives relativt få nettverksproblemer til lokale nettverksfeil eller ondsinnede angrep. Denne innsikten i de grunnleggende årsakene fungerer som en basis for etterfølgende analyser. Paper II presenterer 5 år med risikoregisterdata for å fremheve rollen til standarder som ISO27001 for reduksjon av risiko og viser effekten av å implementere et robust rammeverk for risikostyring på tvers av ulike

organisasjonsnivåer. Paper III ser på internettrisiko, og demonstrerer effektive avbøtende strategier som forbedrer nettverkets motstandskraft mot avbrudd, pakketap og høy latenstid der årsakene skyldes Internett. Paper IV viser en samvariasjon mellom organisasjoners sikkerhetsimplementeringer og overholdelse av to sikkerhetsstandarder, Mutually Agreed Norms for Routing Security (MANRS) og ISO27001. Ved å verifisere RPKI-deltakelse (Resource Public Key Infrastructure), IP-spooferbeskyttelse og Internett risk score for organisasjoner som følger MANRS og/eller ISO27001, demonstrerer vi at en sikkerhetsbevisst bedriftskultur er koblet til bedre sikkerhetspraksis.

Nylige hendelser som COVID-19 og den russiske inntrengningen i Ukraina viser viktigheten av også å vurdere governance-risiko. Paper VII utvider forskningen til å omfatte nasjonale styringsrisikoer, spesielt den store avhengigheten til nasjonale webtjenester av utenlandske mikrotjenester og skytjenester, og understreker nødvendigheten av å vurdere bredere kontekstuelle faktorer.

Med utgangspunkt i den kollektive innsikten fra disse artiklene, ved å kombinere de teoretiske analysene med eksperimenter på et operativt nettverk og erfaringer fra det virkelige livet, fremstår Paper V som en syntese, og foreslår en innovativ 10-lags modell som pragmatisk organiserer identifiserte risikofaktorer. Denne modellen integrerer empiriske funn i et praktisk rammeverk, og resultatene kan generaliseres til en rekke ulike nettverk. Ved å bruke 10-lagsmodellen vil nettverksoperatører redusere sin tilgjengelighetsrisiko og levere tjenester av høyere kvalitet til sine kunder.

Contents

List of articles

[Paper I] Jan Marius Evang et al. “Crosslayer Network Outage Classification Using Machine Learning”. In: *Proceedings of the Applied Networking Research Workshop*. ANRW ’22. Philadelphia, PA, USA: Association for Computing Machinery, 2022, pp. 1–7. ISBN: 9781450394444. DOI: `10.1145/3547115.3547193`. URL: `https://doi.org/10.1145/3547115.3547193`.

[Paper II] Jan Marius Evang. “ISO27001 as a Tool for Availability Management”. In: *Proceedings of the International Workshop on Information Management*. WSIM ’22. London, UK, 2022, pp. 82–85. DOI: `10.1109/AEIS59450.2022.00018`. URL: `https://doi.org/10.1109/AEIS59450.2022.00018`.

[Paper III] Jan Marius Evang and Tarik Cicic. “Evolved Cold-Potato routing experiences”. In: *The 31st International Conference on Software,Telecommunications and Computer Networks (SoftCOM 2023)*.

[Paper IV] Jan Marius Evang and Ioana Livadariu. “How Large Is the Gap? Exploring MANRS and ISO27001 Security Management”. In: *The 31st International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2023)*.

[Paper V] Jan Marius. Evang. “A 10-Layer Model for Service Availability Risk Management”. In: *Proceedings of the 20th International Conference on Security and Cryptography - SECRYPT*. INSTICC. SciTePress, 2023, pp. 716–723. ISBN: 978-989-758-666-8. DOI: `10.5220/0012092600003555`. URL: `https://doi.org/10.5220/0012092600003555`.

[Paper VI] Jan Marius Evang. “Outage risk priorities – It’s not the malicious attacks that take down your service”. In: *IEEE Hotnets 2023*. 2023, Submitted for evaluation.

[Paper VII] Jan Marius Evang and Haakon Bryhni. “National ICT Resilience: An analysis of Norway’s cyber infrastructure preparedness”. In: *IEEE Communications Magazine*. 2023, Submitted for evaluation.

Chapter 1

Introduction

The Internet has evolved into an intricate network of networks that spans across the globe and beyond [2]. It plays a role in the daily lives of the majority of the world's population and is critical to the functioning of governments worldwide. Despite its complexity, the foundational technical layer of the Internet, known as Layer 3 connectivity in the ISO/OSI model [3], is surprisingly simple and predominantly facilitated by the Border Gateway Protocol (BGP) [4]. BGP handles interactions among over 50,000 autonomous systems (ASes) and more than 900,000 network prefixes, forming the backbone of Internet communication.

The Internet is not without imperfections, and the challenge of maintaining a stable and functional Internet is of utmost importance. This task becomes even more important when considering critical services that heavily rely on Internet connectivity. While the interaction between networks may appear straightforward, the vast diversity of systems involved and the loosely connected nature of the Internet create an extensive array of potential points of failure.

In response to these challenges, this thesis aims to thoroughly investigate the potential failure points within a service network and propose a comprehensive framework to address the associated risks. Additionally, specific problem cases will be examined to explore effective methods for reducing risks in those particular scenarios. By investigating availability risk management, this research seeks to contribute practical solutions to enhance the resilience and continuity of critical services in the face of all types of risks.

1.1 Critical Services

The US Cybersecurity & Infrastructure Security Agency CISA) has created a list of functions that are regarded as National Critical Functions (NCFs), defined as “The functions of government and the private sector so vital to the United States that their disruption, corruption, or dysfunction would have a debilitating effect on security, national economic security, national public health or safety, or any combination thereof”. Similarly the Norwegian Directorate for Civil Protection (DSB) has compiled a list of Vital Functions in Society. These lists are compared in Table 1.1.

CISA	DSB
Internet and broadcast radio	Electronic communications networks and services
	ICT security
Positioning	
Electricity	Power supply
Transportation/Fuel	Transport
Governance and Elections	Governance and crisis management
Education and Training	
Law enforcement	Law and order
Waste management	
Economic services	Financial services
Medical services	Health and care
Food/Water	Security of supply
	Water and sanitation
Housing	
Production	
Defence	Defence
	Emergency services
	Nature and the environment
	Satellite-based services

Table 1.1: Critical services as defined by the US CISA and the Norwegian DSB

Recent events have illuminated the criticality of Internet connectivity and its profound impact on various aspects of society. During the COVID-19 lockdown in 2020, Internet access became even more essential for a wide range of activities, including work, education, communication, and access to essential services. It played a critical role in enabling remote work and study, facilitating virtual social interactions and providing access to important information and services. The increased reliance on the Internet during this period highlighted its fundamental value in maintaining societal functions and continuity [5].

During the 2022 invasion of Ukraine by the Russian army, the significance of Internet connectivity again came into focus [6, Blog I]. The Internet played a crucial role in this conflict, with the Ukrainian government attempting to block Russia from the Internet, and Russian interests attempting cyber attacks on Ukrainian (and other) targets. Moreover, the Ukrainian network infrastructure sustained physical damage due to the conflict. Notably, Low Earth Orbit (LEO) satellite systems [7] were deployed to ensure resilient connectivity for critical services in the midst of such challenges.

These real-world examples underscore the reality that the Internet itself has evolved into a critical infrastructure that demands robust measures for its security and protection. The crises experienced during the pandemic and the conflict in Ukraine have emphasized the paramount importance of resilience.

Lessons learned from such events also highlight the crucial role of highly reliable national emergency communications networks based on wireless technologies like ad-hoc networks, dedicated emergency networks like TETRA [8, 9], 4G/5G cellular networks, or Low Earth Orbit satellite systems [10]. These platforms are essential for disaster reduction, enabling seamless communication among emergency response teams and supporting the continued operation of critical services [11]. A wide choice of technologies are available for critical services [12], but a global trend for critical services is the increased use of high volume commercial technologies like cellular infrastructure and terminals instead of dedicated emergency networks for critical services, due to higher capacity and lower prices [13]. All the mentioned wireless access technologies rely on a stable backbone network to provide required cloud services for the network itself and the critical applications. Thus, the methods developed in this thesis are applicable for any choice of access network.

Understanding and addressing the risks associated with Internet connectivity and network availability are vital steps to safeguarding critical services and ensuring the resilience of our societies. Each critical service has its own relevant requirements to the underlying network service, such as availability, Time To Recover (TTR), Quality of Service (QoS), and correctness of information. The subsequent chapters of this thesis propose strategies and frameworks to enhance the availability, security, and resilience of critical services, enabling more resilient networks in the face of adversities.

1.2 Availability Risks

In [14] the authors perform a high-level survey of literature which shows that system availability has attracted marginal attention by researchers compared to confidentiality and integrity, even though availability is an important part of risk modeling and represents a Key Performance Indicator (KPI) of equipment and service providers. The Oxford Dictionary defines resilience as “the capacity to withstand or to recover quickly from difficulties”, which is closely linked to availability that we define in terms of the two variables MTBF (Mean Time Between Failures) and MTTR (Mean Time To Recover), or alternatively through measuring or predicting the uptime and downtime percentages.

The Internet today comprises numerous interconnected systems that must collaborate to ensure its continuous operation. Along one axis, there is a large number of independent organizations, organized as Autonomous Systems (ASes), that need to cooperate. Along another axis, multiple layers of network protocols and services combine to form the Internet service. As an illustration of the protocol complexity, The Junos® OS Standards Reference is 93 pages long and is just a list of all standards that Juniper’s network routers and switches substantially support [15].

A failure in any of the layers or any of the organizations involved has the potential to cause availability problems for users. Thus the central focus of this thesis is to investigate the most efficient points at which outage risks can be mitigated. As an illustration of the challenge of uptime, the Norwegian Emergency Network (Nødnett) has a yearly availability goal of 99.95%, which has never been archived since the Norwegian Directorate for Civil Protection (DSB) took over the responsibility in 2017. (Source: DSB annual reports 2017-2022 [16])

In intricate risk scenarios, risks can interdependently emerge from multiple components. Compound risks often demand various calculation methodologies. For instance, when a service relies on two distinct components and the failure of either component halts the service, risk estimation involves using the formula $Uptime_{total} = Uptime_1 \times Uptime_2$. This means that a service with two components, each having 99% uptime, will have a total uptime of 98%. Conversely, in scenarios where two independent redundant components are at play and only one is needed for operation, the corresponding calculation takes the form $Downtime_{total} = Downtime_1 \times Downtime_2$. In this case, a service with two components, each having 99% uptime (1% downtime), will achieve a significantly higher uptime of 99.99%.

To design resilient networks, availability risk management is crucial. While it is very difficult to create a network service that is guaranteed to never see any outage, efforts can be directed towards minimizing the occurrence of outages and expediting the recovery, thereby enhancing resilience and availability.

1.3 The Proposed 10-layer Model for Risk Management

As shown in Chapter 2, none of the established security frameworks give a comprehensive system for managing availability risks. To fill this gap, we propose a new 10-layer model for availability risk. The model is briefly described here, and thoroughly treated in [Paper V]. Table 1.2 provides a visual representation of the proposed new framework for availability risk management, demonstrating how the papers in this thesis relate to its various layers. The papers collectively contribute to a comprehensive understanding of network availability and its associated risk factors, paving the way for practical solutions and improvements in availability risk management practices for critical services.

For the Governance Layer, in addition to traditional research papers, the results of specific small research projects were published in blog posts. These blog posts were chosen as the appropriate medium to address topics of public interest related to Internet availability and resilience. The aim was to increase awareness and understanding of the potential risks associated with Internet dependencies on organizations and locations outside of Norway [Blog I, Blog II].

The classic ISO/OSI reference model [3] is illustrated in Figure 1.1. This model was created in the 1980s [17] as a tool to keep track of all required functionality to design a computer network.

The ISO/OSI reference model breaks down digital communication into seven abstraction layers (Figure 1.1). Layer 1 manages the physical transmission of bits, Layer 2 encompasses local network protocols, Layer 3 facilitates data movement across different networks, Layer 4 supports multiple services over the same network, Layer 5 manages end-to-end communication, Layer 6 enables communication between heterogeneous systems, and Layer 7 contains application-specific protocols. In modern Internet usage, Layers 1-2 are typically handled by the connected device, Layer 3 is managed by the Internet Service Provider (ISP), and Layers 4-7 are the responsibility of the operating system and software.

While the 7-layer ISO/OSI model remains a valuable educational tool, it falls short in accurately depicting the complexity of today's Internet. The rapid evolution of Internet technologies, services, and the growing interconnectivity of various systems have introduced new challenges that the traditional model does not fully address. In response, this thesis introduces the extended layering model (Figure 1.1) extensively discussed in [Paper V], with a specific focus on network availability risk management, adaptable for other contexts involving availability risk.

Within this extended model, all availability risks are systematically cataloged and organized into distinct layers, each corresponding to specific risk factors. The Physical Layer pertains to anything hardware-related, the Local Network Layer encompasses networks fully under the organization's control, the Wide Area Network Layer pertains to network services obtained from sub-contractors, and the Internet Layer comprises other external networks. Additionally, further layers contain additional risks that impact service delivery, including Cloud, Applications, and Services, while the higher layers address indirect yet critical risks originating from Organizations, People, and Governance. Detailed descriptions of each layer can be found in Section 3.3.

Table 1.2 aligns the papers in this thesis with the corresponding layers, collectively contributing to a holistic understanding of availability risk management.

By employing this extended layering model, the thesis offers a structured approach to analyze and manage risks associated with Internet connectivity, bolstering resilience and availability of critical services.

1.4 Research Questions

This thesis aims to address the following research questions:

RQ1: What are the main availability risks to critical services?

RQ2: What is the best way to organize availability risk management?

RQ3: How can the most important availability risks be mitigated?

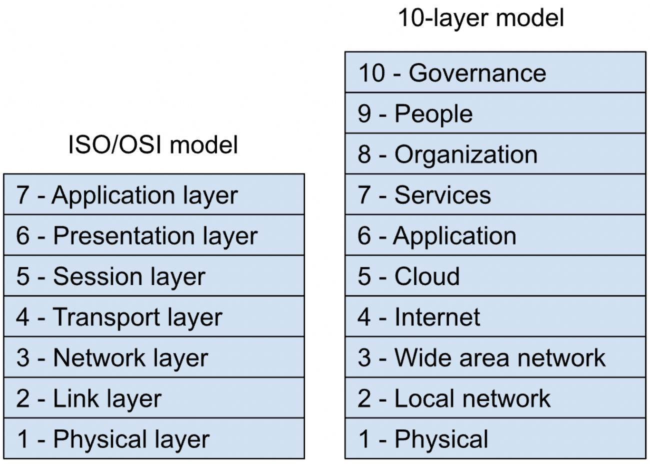

ISO/OSI model

7 - Application layer
6 - Presentation layer
5 - Session layer
4 - Transport layer
3 - Network layer
2 - Link layer
1 - Physical layer

10-layer model

10 - Governance
9 - People
8 - Organization
7 - Services
6 - Application
5 - Cloud
4 - Internet
3 - Wide area network
2 - Local network
1 - Physical

Figure 1.1: Comparison of the 7-layer and 10-layer models.

	I	II	III	IV	V	VI	VII
10 - Governance		x					x
9 - People		x				x	
8 - Organization		x		x		x	x
7 - Services		x					x
6 - Application					x		
5 - Cloud			x		x		x
4 - Internet			x	x	x		
3 - Wide area network	x		x			x	
2 - Local network	x						
1 - Physical	x					x	

Table 1.2: Layers of network connectivity, related to the papers of this thesis.

By investigating these research questions, we can gain insights into the primary risks faced by critical services, explore suitable organizational frameworks for managing availability risks, and identify effective resilience strategies to mitigate the most significant risks. The findings will contribute to the development of a comprehensive understanding of availability risk management and inform the design of practical solutions to enhance the reliability and continuity of critical services.

1.5 Thesis Outline

The thesis encompasses a comprehensive exploration of availability risk management and resilience in the context of critical services. It consists of several distinct sections that collectively contribute to the research objectives. The outline of the thesis is as follows:

- Introduction
  - Provides an overview of the Internet's significance, its underlying simplicity and complexity, challenges in maintaining its reliability, and the objectives of the research.
  - Sets the stage for the subsequent chapters by establishing the research context and objectives.
- Literature Review
  - Explores recent research and standardization efforts in the field of critical services, resilience and availability risk management.
  - Identifies gaps in existing literature and research, highlighting the need for the current study.
- Theoretical Framework
  - Describes in detail the theoretical framework developed for understanding and managing availability risks in critical services.
  - Explores the various layers and components of the proposed framework, highlighting their interrelationships and roles in risk mitigation and resilience.
  - Discusses the theoretical underpinnings of the framework, drawing from relevant theories and concepts in the field.
- Ethical considerations
  - Assesses ethical implications of the research.
  - Describes actions taken to limit the ethical problems.
- Paper presentation
  - Summarizes each paper in the thesis.
  - Highlights the lessons learned from each paper.
- Methodology
  - Presents the research methodology employed to investigate availability risks and validate the proposed framework.

  - Describes the data collection methods, measurement techniques, and analysis procedures utilized.
  - Discusses the rationale behind the chosen methodology and its suitability for addressing the research objectives.

- Experimental Validation
  - Presents the results of the experimental validation conducted to assess the effectiveness of the proposed framework.
  - Provides detailed analyses of real-world case studies to illustrate the practical application of the framework.
  - Discusses the findings, highlighting key insights, successes, and limitations encountered during the validation process.

- Discussion and Conclusion
  - Summarizes the key findings, contributions, and insights gained from the research.
  - Engages in an in-depth discussion of the research findings, contextualizing them within the broader field of availability risk management.
  - Explores the implications of the research for practitioners, policymakers, and researchers.
  - Reflects on the research journey and its significance in advancing the field of resilience and availability risk management.
  - Identifies areas for further investigation and potential improvements to the framework.

- Appendix
  - Contains the full versions of the research papers related to the thesis.

Chapter 2

Literature Review

The field of risk management is heavily influenced by diverse national and international organizations, including the International Organization for Standardization (ISO, an international Non-Governmental Organization), the Information Systems Audit and Control Association (ISACA, a professional membership organization), Payment Card Industry Security Standards Council (PCI-SSC), Intel (a public company), the Center for Internet Security (CIS, an independent, nonprofit organization), Secure Controls Framework Council LLC (US-based organization, sponsored by the consultancy industry), as well as governmental institutions such as British Standards Institution (BSI, the UK national standards body), The European Union Agency for Cybersecurity (ENISA), The council of the European Union, National Institute of Standards and Technology (NIST, US Department of Commerce), the Federal Risk and Authorization Management Program (FedRAMP, a US government program), the Health Insurance Portability and Accountability Act (HIPAA, US Federal Law), Defense Federal Acquisition Regulation Supplement (DFARS, US defense), and FIRST/International Telecom Union (ITU, United Nations).

The standards and frameworks developed by these organizations are often based on laws and regulations in their respective jurisdictions, industry best practices, and expert recommendations. Although the research behind these standards may not be published according to peer-reviewed standards, the standards provide valuable guidelines for risk management.

Choosing a risk assessment model without proper analysis can lead to the implementation of security controls in the wrong areas, resulting in a waste of resources and leaving an organization vulnerable to unforeseen threats. E. G. Amoroso [18] outlines the principles of a good risk classification, which include non-overlapping classes of risks, exhaustiveness, unambiguity, acceptance of established terminology, and usefulness.

A widely used method for risk assessment is to follow the principles outlined in ISO31000 [19] and ISO31010 [20]. These standards offer comprehensive guidelines and methodologies for effective risk management. However, in the context of intricate systems like critical service networks, a more tailored approach to risk discovery is essential. The sheer volume of risks associated with such systems can make it challenging to address all of them collectively, potentially leading to overlooked risks.

While several systems, such as Bayesian Attack Graphs [21, 22, 23], have been proposed for managing discovered risks, they are not specifically tailored towards availability risks and lack a comprehensive risk discovery model.

Models like Joshi and Singh's risk management framework [24] identify phases in the risk assessment process, such as identifying weaknesses, creating a remediation plan, and managing risk to improve security over time. This framework also lacks a model for risk discovery specific for the security context of resilient networks.

The framework called OCTAVE (Operationally Critical Threat, Asset and Vulnerability Evaluation), developed by CERT, provides a model for risk-based information security strategic assessment and planning. While OCTAVE considers assets to include people, hardware, software, information, and systems, its complexity and lack of quantitative risk modeling are significant drawbacks.

The FAIR Risk assessment method (Factor Analysis of Information Risk) from 2005 [25] is complementary to standards like ISO27001, and is useful for breaking down the risk of larger projects onto smaller manageable elements and quantifying those risks. However, FAIR does not cover first time risk assessment or a holistic risk assessment.

The ADMIT methodology [26] captures five major classifiers to characterize the nature of attacks. These are classification by Attack vector, classification by Defense, classification by Method, classification by Impact and classification by attack Target. ADMIT is focused on malicious attacks, and while the methods are useful, they are not applicable to all availability risks since the method's primary risk discovery step only focuses on security scanning and technical asset evaluation.

The NIST RMF (National Institute of Standards and Technology Risk Management Framework) completely ignores the discovery step. It does have a "Categorize" step, but this step assumes that a risk discovery process has already been performed.

The TARA (Threat Agent Risk Assessment) framework is developed by Intel to help organizations manage risks by extracting possible information about security attacks [27]. It primarily focuses on external malicious threats and represents only a fraction of the overall risk landscape.

The proposed OSSF (Online Services Security Framework) framework [28] was developed using the Design Science Research methodology (DSR) [29]. This method defines a process of six activities which result in an artifact, it addresses problem identification but lacks a comprehensive risk discovery process.

The CORAS (Consultative Objective Risk Analysis System) method [30], developed by the Department of Informatics at the University of Oslo, consists of eight steps aimed at identifying threats, risks, and selecting appropriate security measures. However, risk discovery is described as a brainstorming session, and a structured framework for efficient brainstorming of risks for complex network services is missing.

While industries and public sectors are placing a strong emphasis on integrating risk management practices, the corresponding academic research has yet to catch up. Our literature review reveals that most contemporary research is centered on management of availability risks in the domains of power systems and healthcare services. Similarly, significant attention is given to addressing confidentiality and integrity risks within network infrastructures. However, the realm of network availability risk management still remains relatively uncharted in academic literature, even though pertinent references within this field are identifiable, as shown below.

For an overarching but accessible discussion and comparative analysis of risk management standards and guidelines, consult [31].

In the context of evaluating Return on Security Investments (RoSI) and assessing the effects of enhancements, [32] presents a technique for quantifying resilience, [33] explores methods of cost-effectiveness for minimizing SLA breach compensations, and [34] introduces a novel approach for reporting system availability. The application of genetic algorithms in AI to optimize risk mitigation strategies is discussed in [35], while [36] examines the economic viability of sharing backup infrastructure. However, these studies do not cover risk mitigation prioritization methods or analyze the real-world impact on customers. Notably, studies such as [37], [38], and [39] concentrate on risk discovery in physical, technical/human/legal, and human layers respectively, but lack a holistic perspective.

In light of these research gaps and limitations, this thesis aims to contribute to the field of availability risk management by developing a comprehensive theoretical framework that specifically focuses on network availability risks. The proposed 10-layer model aims to provide a holistic view of risk management, addressing the interconnectedness and interdependencies across different layers and factors that influence availability risks. By incorporating a comprehensive risk discovery process and providing a structured approach to organizing risk management efforts, the proposed model seeks to enhance the effectiveness of availability risk mitigation strategies for critical services.

The thesis also aims to fill the gap in the literature by presenting real-world case studies and experimental validation to assess the effectiveness of the proposed framework. By exploring specific problem cases and providing practical examples of risk mitigation, this research seeks to offer valuable insights and guidance to practitioners, policymakers, and researchers in the field of availability risk management.

Through a systematic literature review and critical analysis of existing models, this thesis endeavors to advance the understanding and practice of availability risk management, contributing to the development of more robust and resilient critical services in the face of increasing cyber threats and disruptions.

Chapter 3

Theoretical Framework

3.1 Risk Methodology

Risk management can be described simply as the discovery of what can go wrong and the creation of actions to reduce the likelihood and impact of faults. The risk management methodology in this thesis encompasses a systematic approach to risk management, focusing on the discovery of potential vulnerabilities and the development of actions to mitigate risks effectively. This methodology draws heavily from the methods used for achieving ISO27001 certification and for maintaining the certification through yearly audits, making it a robust foundation for addressing availability risks.

ISO27001 is an internationally recognized standard for Information Security Management Systems (ISMS). The process of attaining an ISO27001 certificate begins with the establishment of an ISMS, which comprises a collection of policies and procedures that govern how the organization operates securely. Also defined is a system of internal audits to improve the organization over time. Central to this process is the need to define Key Performance Indicators (KPIs), i.e. quantified measures that enable organizations to assess their performance, determine the effectiveness of their risk management strategies, and identify areas for improvement. Examples of KPIs for a network operator may be economic like turnover or revenue, they may be capacity related such as total network traffic or link utilization, or as used in this work, measures of network quality. In [Paper II], the KPIs used were risk scores and the number of relevant customer support cases. These risk scores were determined through interviews with relevant stakeholders and reflected the likelihood and impact of identified risks. In [Paper VI], risk importance is quantified, based on previous frequency of outages and previous customer impact of the same outages.

The risk discovery process is a crucial aspect of the risk methodology. It involves identifying and assessing potential risks to the organization's availability. This process begins by categorizing risks into groups, such as risks related to sensitive information leaks or risks arising from the physical environment.

Interviews are then conducted with the identified stakeholders, and a risk registry is updated. The risk registry contains an entry for each identified risk, with an indication whether it is related to the risk objectives of confidentiality, integrity or availability, a likelihood value (0-5) and an impact value (0-5). For the yearly management review and ISO27001 audit, the total risk score for the organization is calculated as the sum of all the risk scores.

The risk scoring process allows for a quantitative assessment of risks. By calculating the sum of all risk scores, the overall risk profile of the organization can be evaluated. However, it is essential to note that the risk value on its own does not provide meaningful insights. Instead, it serves as a basis for comparison and facilitates tracking the organization's risk trends over time.

One of the strengths of this risk methodology is its iterative nature. Risk management is an ongoing process that requires continuous monitoring, evaluation, and improvement. Yearly management reviews and ISO27001 audits provide opportunities to review the ISMS and risk management strategies, enabling organizations to refine their approach and enhance resilience against availability risks.

By following the risk methodology, organizations can systematically identify and address potential availability risks, continually improve their risk management practices, and work towards achieving greater reliability and continuity in critical services.

3.2 Resilience

Security has many facets, and for this thesis, the focus is on resilience, which as mentioned earlier relates to the risk objective of availability. Within a network operator setting, there are numerous availability risks that the Network Operations Center (NOC) must address continuously to enhance network quality and minimize service disruptions.

Within the 10-layer model proposed in this thesis, various availability risks may cause service disruptions, and extensive research and product development has provided a large number of protocols and methods to reduce the impact these failures have on the service delivery.

Service disruptions caused by failures and undercapacity may cause packet loss, jitter or total outage for a longer or shorter time, in turn causing customer complaints. Useful quality indicators are packet loss percentage and frequency, as well as outage duration and frequency. If mitigations are successful, these indicators will show improvements, leading to a more resilient network, and eventually a reduction in customer complaints.

In addition to resilience, other related terms like durability (the continual existence of an artefact, whether available or not), reliability (the availability and correctness of data), and fault tolerance (the ability of recovering after a failure) contribute to the overall understanding of network robustness. These concepts collectively form the foundation for designing effective risk management strategies within the 10-layer framework.

As we proceed with the theoretical framework, we will explore the application of these resilience principles across the different layers of the network infrastructure. Understanding how resilience considerations extend throughout the 10-layer model is crucial for implementing effective risk mitigation measures.

In the subsequent chapters, we will delve deeper into the theoretical underpinnings of availability risk management, focusing on how resilience is integrated into each layer of the 10-layer model. By applying theoretical insights and principles, we aim to develop comprehensive frameworks that enhance the resilience of critical services and ensure their reliable operation, even in the face of adversities.

3.3 Cross-layer RISK Description

In this section, we present a comprehensive analysis of the 10-layer model proposed in [Paper V] for availability risk management. This model offers a practical and structured approach to organizing availability risk enabling organizations to effectively address and mitigate potential disruptions. It acknowledges that many risks that have a root cause in one layer may be mitigated using methods from other layers, highlighting the interconnected nature of availability risk management.

The 10-layer model represents a significant advancement in understanding and managing availability risks across different aspects of critical services. By delving into each layer, from the technical infrastructure to the human and governance elements, organizations can enhance their overall resilience and availability.

The interplay between these layers underscores the importance of a holistic approach to risk management, where comprehensive frameworks and strategies are essential to tackle the diverse challenges posed by the evolving Internet landscape. Throughout this section, we will explore each layer of the proposed model in detail, outlining its specific risk factors and contributions to overall availability risk management. By understanding the nuances of each layer and their interactions, network operators and organizations can design tailored risk mitigation approaches that address vulnerabilities and enhance the reliability of their critical services.

As the Internet continues to evolve and face new challenges, research and practical implementations should continue to explore these layers, refine frameworks, and identify effective strategies for mitigating availability risks. This ongoing effort will ensure that critical services remain robust and resilient in the face of ever-changing threats and disruptions.

3.3.1 Physical Layer

Layer 1 in the OSI model (physical layer) consists of the cables and the corresponding protocols that run on these cables, as well as the wireless equivalents. In today's Internet, the physical risks also include risks towards building infrastructure and other physical infrastructure like radio towers and satellites. Both the confidentiality/integrity risks of a physical break-in to the premises for the purpose of stealing/altering information, and the availability risks posed by natural disasters, human error or purposeful destruction must be considered. While it is crucial to protect the physical layer as it forms the foundation of all systems, it is important to acknowledge that it is not infallible, as seen in [Paper VI] and [Blog I]. Thus, higher layers must anticipate faults and have appropriate measures in place to effectively handle confidentiality, integrity, and availability risks.

In the studies presented in [Paper I] and [Paper II], comprehensive analyses of the physical layer and network infrastructure's risks are conducted. These analyses consider risks from confidentiality, integrity, and availability perspectives, shedding light on the necessary measures to safeguard critical services in these areas. The insights gained from these studies inform the development of robust strategies to bolster the physical layer's resilience and ensure the continuity of critical services.

Cryptographic techniques, such as encryption and digital signatures, play a crucial role in addressing confidentiality and integrity risks during data transmission and at rest. For availability, various mitigations can be employed, such as Redundant Array of Inexpensive Disks (RAID) configurations for reducing the impact of disk failures and implementing redundant components in servers. Additionally, protective measures at the building level can safeguard against external threats like flooding, extreme weather events, or power outages. In some cases, even the trust in equipment manufacturers may need to be scrutinized, as highlighted in [40].

3.3.2 Local Network Layer

The Local Network Layer encompasses a wide array of technologies used to establish connections within a local area, such as cabled copper or fiber Ethernet [41], and various wireless options, including Bluetooth, Wi-Fi, fixed wireless connections, cellular (2G, 3G, 4G, 5G), and satellite communication. In our layered model, the local network also includes local routers, and switches that are administered in the same domain.

Local network protocols integrate security features to ensure data protection and access control. These measures include Wi-Fi encryption, password-based authentication, and International Mobile Equipment Identity (IMEI) authentication for cellular networks. Additionally, Forward Error Correction (FEC), and Wi-Fi´s data recovery capabilities help enhance data reliability at this level.

Critical network services demand high levels of availability. Therefore, redundancy is widely implemented within the Local Network Layer. For vital links, switches, and routers, High Availability (HA) setups are adopted, involving the deployment of redundant devices. Despite encryption being added at this layer, it is not entirely trusted, making end-to-end encryption at higher layers highly recommended for enhanced security.

Common availability protocols at the OSI Layer 2 (link layer) include Spanning Tree (STP), Link Aggregation (LACP), and various fabric systems. Interior Gateway protocols (IGP) such as Open Shortest Path First (OSPF) or Intermediate System to Intermediate

System (IS-IS) are commonly used on OSI layer 3 (network layer) to mitigate faults in the lower layer services.

Cellular services provide resilience through features like cellular handover and Emergency Service, enabling any cellular phone to initiate emergency calls on any network, even without a relevant subscription.

As seen in [Paper VI], the Local Network Layer holds significant importance in risk reduction efforts. It forms a vital part of the 10-layer model for availability risk management. Understanding and securing the Local Network Layer are crucial steps towards enhancing overall network resilience and ensuring uninterrupted service delivery.

3.3.3 Wide Area Network Layer

The Wide Area Network (WAN) represents a critical layer in the 10-layer model for availability risk management, covering leased network links. While OSI Layer 3 (network layer) protocols like IPv4 and IPv6 are used for both local and wide area network communications, the WAN introduces distinct challenges due to its leased physical infrastructure and the involvement of service providers. Unlike local networks, where the network operator has full ownership and management responsibility for all components, including physical cables and devices, WANs rely on leased infrastructure and lower OSI layers (1-3) from external providers. This introduces a higher risk of eavesdropping and reduced visibility into availability risks.

The availability risks in WANs can arise from various sources, including physical faults, network faults like congestion and capacity limitations, as well as protocol errors and human error. These risks can lead to disruptions in service delivery and can significantly impact critical services.

As illustrated in [Paper I], the monitored WAN links experienced around 700,000 incidents of small or large packet loss over a two-year period. This highlights the need for robust design and mitigation strategies to handle such incidents effectively. For critical services, it becomes imperative to implement multiple WAN links from diverse providers to ensure redundancy and avoid shared risks that could impact all connections simultaneously.

Monitoring and assessing the security of WAN links is of paramount importance. While these links may be monitored by various authorities, it is also essential to acknowledge the possibility of illegal actors monitoring them as well, posing additional security challenges.

The Wide Area Network Layer plays a crucial role in the availability and resilience of critical services. Proper management of availability risks in this layer involves selecting reliable providers, implementing redundancy, and employing robust security measures to protect sensitive data and ensure uninterrupted service delivery. As with other layers in the 10-layer model, addressing the risks in the WAN Layer requires a holistic approach, considering its interconnections with other layers to achieve comprehensive availability risk management.

3.3.4 Internet Layer

The Internet Layer, as defined in our 10-layer model, plays a critical role in the interconnection of different networks, forming the backbone of the internet. Located at OSI Layer 3 (Network or Internet Layer), the focus primarily centers around the Border Gateway Protocol (BGP) and its utilization by different organizations, known as Autonomous Systems (ASes). BGP allows ASes to announce to their neighbors which IP addresses (prefixes) they can reach, leading to the generation of a full routing table that is propagated throughout the Internet.

However, the BGP protocol relies on trust, and its security shortcomings have been known since the late 1990s, as mentioned in RFC1771 [42]. The trust-based nature of BGP introduces risks, since the protocol will enable organizations to announce networks they do not have the rights to.

To ensure the resilience and security of critical services, reducing the reliance on trust becomes a crucial objective. [Paper IV] addresses significant BGP security shortcomings and explores the MANRS framework (Mutually Agreed Norms for Routing Security) as a means to reduce risks associated with these shortcomings. Implementing the MANRS framework in conjunction with ISO27001 demonstrates how the risks imposed by BGP can be mitigated effectively. The OECD (Organization for Economic Cooperation and Development) also emphasizes the importance of addressing BGP-related issues in ensuring a secure and resilient Internet infrastructure [43]. BGP-related risks primarily fall under availability, but there have been instances of eavesdropping and integrity breaches, as discussed in [Paper IV].

Vulnerability scanning, a common practice to simulate outside or inside attacks, serves as a mitigation measure at the Internet Layer, helping operators discover potential risks and improve the security posture of their networks.

A case study in [Paper III] illustrates how BGP shortcomings were mitigated by implementing geographic (GeoIP) based routing with latency and packet loss considerations. This approach resulted in improved availability, reduced latency, and lower packet loss for a video conferencing network service, while also reducing risk by circumventing network errors outside the operator's AS.

Addressing BGP vulnerabilities and optimizing routing protocols are crucial steps toward enhancing the resilience, availability, and security of critical services in the ever-evolving Internet landscape. A comprehensive approach combining security best practices, policy frameworks, and technical improvements will contribute to a more resilient and reliable Internet infrastructure for critical applications.

3.3.5 Application Layer

The Application Layer represents a crucial aspect of overall system security, although it has not been the primary focus of this thesis. [Paper II], however, does provide an analysis of risks in the software development process, encompassing both confidentiality and availability risks.

For in-house developed applications, a well-formed development policy must be in place, addressing all potential risks to ensure the resilience and security of the application. Similarly, any third-party applications should be thoroughly evaluated to identify potential risks and assess their potential impacts on critical services.

At the Application Layer, many high-impact risks originating from lower layers can be effectively mitigated. For instance, applications can switch to alternative services if the main service is unavailable, enhancing availability and reducing the risk of service disruptions. A notable example is seen in peer-to-peer file-sharing networks, where peers frequently join and leave the network. Applications in such networks need to be resilient, as they may have to try several different peers before finding one that can be used for file transfers. Additionally, these applications must gracefully handle scenarios where peers leave the network mid-transfer.

While application security remains beyond the scope of this thesis, it is important to acknowledge its significance in enhancing the overall resilience and availability of critical services. Proper development practices, rigorous risk assessments, and robust application design contribute to a more secure and dependable application layer, complementing the efforts made at lower layers to ensure the continuity of critical services.

3.3.6 Services Layer

The Services Layer represents the most visible aspect of a system from the customer's perspective. Service availability directly impacts customer satisfaction, and metrics like SLA (Service Level Agreement) and "uptime" are used to evaluate service quality at this layer.

The service layer is dependent on all the underlying network and hardware layers, as well as the software applications, and services like DNS. Ensuring the resilience and availability of services requires consideration of all these interconnected components.

Our previous work resulting in a US patent [44] focused on providing a highly resilient video service. The implementation utilized a TURN (Traversal Using Relays around NAT) service, which demonstrated excellent availability without a single incident from 2008 to 2023. Many typical methods for improving service availability involve load balancers, anycast, and DNS-assisted techniques. While effective, these methods can introduce new, albeit lower risk, single points of failure. The implementation in [44] effectively mitigated these risks, showcasing exceptional availability.

The global DNS service (information available at `https://root-servers.org/`) is another example of a highly resilient system. It relies on DNS root servers distributed across various locations worldwide. A single root server's availability is sufficient to provide DNS service to a caching local DNS server. These root servers are operated by different organizations, running on diverse software, operating systems, and hardware platforms to add further resilience. However, the DNS system remains susceptible to attacks related to confidentiality and integrity. Operators can potentially log DNS requests or intercept and modify answers, leading to possible service censorship or compromise.

Ensuring the availability and resilience of services requires a comprehensive understanding of the interconnected components at the Service Layer, along with the strategic implementation of robust techniques that address various risks effectively.

3.3.7 Organizations Layer

While the physical layer forms the foundation of the network, the organization is at the top of everything. It plays a pivotal role in conducting risk assessments, formulating and enforcing policies, procuring and configuring equipment, allocating financial resources, appointing and supporting personnel, complying with laws and regulations, and establishing contracts.

The organization's culture and priorities significantly impact aspects like security, confidentiality, integrity, and availability. If these factors are not given due importance, the service's overall performance and resilience may suffer. The adoption of formal frameworks like ISO27001, as demonstrated in [Paper II] and [Paper IV], can prove beneficial in ensuring proper and effective organizational operation.

Risks at this layer can also arise from external entities. For example, employee organizations might go on strike if work conditions are deemed unacceptable, posing risks to operations. Trade organizations may impose regulations, and non-compliance with these regulations can result in operational risk. Additionally, there may be risks originating from customers demanding adherence to certain standards, such as ISO27001, or from suppliers, introducing supply chain risk.

Being a member of Internet associations and network operator groups can generally contribute to reducing risks as well.

Supply-chain problems also fall under the scope of risks at the Organizations Layer. Supply chain risks can be significant, affecting the availability and security of critical components or services. It is crucial for organizations to identify and address such risks to ensure smooth operations.

Loss of data and configuration errors can also be classified within the Organizations Layer. These risks may arise due to inadequate policies, insufficient employee training, or inadequate security measures, all of which fall under the organization's responsibilities.

In conclusion, the Organizations Layer plays a central role in coordinating and overseeing all aspects of risk management, ensuring that the organization prioritizes security, confidentiality, integrity, and availability in its operations and interactions with external entities.

3.3.8 People Layer

The People Layer is an essential aspect of availability risk management, as ultimately, people are involved in all aspects of the organization's functioning. From designing and manufacturing equipment to implementing services and networks, establishing policies, and upholding security, people play a central role. However, people are also susceptible to making mistakes, and in some cases, they may act maliciously, posing significant risks to the organization's resilience.

In the context of the People Layer, the work presented in [Paper II] demonstrates how the ISO27001 framework can be used to analyze direct risks related to personnel. This includes evaluating risks in the process of employing new people or subcontractors, assessing the risk of key personnel leaving the company, and constructing policies for various aspects of the business and service delivery that rely on people's actions.

Addressing internal threats posed by employees is a significant challenge in cyber security and resilience. Implementing compartmentalization and separation of duties are efficient methods for mitigating these risks, ensuring that no single individual has access to everything and that critical work requires mutual verification. A positive work environment and prioritizing employee satisfaction are also crucial factors in reducing stress levels, minimizing turnover, and lowering the likelihood of mistakes or malicious activities.

By focusing on the people aspect, organizations can create a culture of security awareness, responsibility, and diligence, leading to increased resilience and a reduced overall risk to service availability.

3.3.9 Governance Layer

One layer that is frequently overlooked when assessing risk is the Governance Layer. Changes in laws and regulations can introduce additional work or even force the shutdown of certain services. Organizations must also consider the risk of being impacted by embargoes, as demonstrated by the events involving companies conducting business in or with Russia in 2022 (refer to [Blog I] for more information). National sovreginity and Information and Communication Technology (ICT) resilience is assessed in [Paper VII].

The Governance Layer also encompasses organizations responsible for overseeing critical systems on the Internet, such as IANA (Internet Assigned Numbers Authority), ICANN (Internet Corporation for Assigned Names and Numbers), and the Regional Internet Registries. These organizations work for the benefit of the Internet as a whole, but they are subject to local laws, and may be subject to “cease-and-desist” court orders, affecting large parts of the Internet. A recent example comes from the African Network Information Center (AfriNIC), the institution responsible for all IP addresses in Africa, which faced controversies lately, including a court order that temporarily froze all their assets, prompting many organizations to re-evaluate their operational risks (see also [Blog II]).

Additionally, risks associated with IP address exhaustion fall under the Governance Layer. As the pool of available IPv4 addresses depletes, organizations may face challenges in obtaining new IP addresses, impacting network operations, scalability, and service availability. Understanding the risks arising from IP address exhaustion can help organizations plan for the transition to IPv6 and adopt strategies to mitigate the potential disruptions.

At this layer, company governance also plays a role. Decisions bade by top management can have significant implications for risk management. In cases of high-impact security incidents, obtaining top management approval may be necessary to take actions such as shutting down a service to protect sensitive data. Considering governance as part of risk assessment ensures that the broader organizational and legal contexts are taken into account, providing a comprehensive understanding of risk exposure.

To build resilience in the Governance Layer, organizations must stay vigilant about changes in laws and regulations, engage in proactive risk assessments, and formulate contingency plans to address potential disruptions. Compliance with industry standards and best practices, such as ISO27001, can provide a framework for effective governance and risk management, enabling organizations to navigate potential legal and regulatory challenges more effectively. By acknowledging and addressing risks at the Governance Layer, organizations can reinforce the foundation of their risk management strategies and enhance their overall resilience in the face of uncertainty.

3.3.10 Summary of the 10-Layer Model

The proposed 10-layer model for availability risk management offers a comprehensive approach to addressing risks in critical services. Each layer in the model represents a crucial aspect that contributes to the overall resilience and availability of the services. By understanding and managing risks at each layer, organizations can enhance their ability to withstand challenges and disruptions.

At the Physical Layer, risks related to hardware and infrastructure are considered, acknowledging the foundation upon which all systems are built. The Local Network layer encompasses the local network protocols and wireless technologies, highlighting the importance of securing these connections that form the backbone of internal communications.

Moving to the Wide Area Networks Layer, the risks associated with leased infrastructure and network faults are recognized, emphasizing the need for redundancy and diverse network paths. The Internet layer, based on the BGP protocol, highlights the critical role of trust and security in interconnecting different networks to form the Internet.

The Application Layer, although not the primary focus of this thesis, acknowledges the importance of software development and third-party application evaluations in mitigating risks. The Services layer, visible to customers, focuses on service availability and various techniques such as load balancing and DNS-assisted methods to ensure continuity.

The Organization Layer plays a central role in risk assessments, policy formulation, and resource allocation. It also considers external entities, such as trade organizations and suppliers, that can introduce operational risks.

The People Layer recognizes the human element and the impact of personnel on all aspects of service delivery. It emphasizes the importance of proper hiring practices, training, and fostering a positive work environment to reduce both unintentional and malicious risks.

Finally, the Governance Layer encompasses legal and regulatory aspects, which can introduce additional work or even force changes in service delivery.

In conclusion, this 10-layer model provides a comprehensive understanding of availability risk management in critical services. By addressing risks at each layer and recognizing their interconnectedness, organizations can enhance their overall resilience and ensure the continuous delivery of critical services. The research in this thesis aims to contribute valuable insights and practical solutions to enhance the reliability and continuity of critical services in the ever-evolving landscape of the Internet and network environments.

Chapter 4

Methodology

We choose to analyze a global service provider (Media Network Services, MNS), a network designed to provide a high quality, resilient wide area network service for demanding video conferencing customers. The reason for this choice is that the security and resilience challenges and the practical solutions investigated in relation to this service provider will be directly applicable to other networks as this service network is implemented using best practice of connectivity, network elements, and protocols. Due to the focus on a demanding service in the design of the service provider (latency critical video conferencing), the findings using data from this network will be particularly relevant for networks providing critical services. Executing our research on this network facilitated the formulation of experiments and analyses aimed at addressing the research questions.

The focal point on a single network operator naturally narrows the immediate scope of the research to networks of this specific nature. However, this specialization does not constrain the applicability of the research outcomes. To further validate the robustness and universality of the derived models, the model was successfully tested on the research network for the Center for Resilient Networks and Applications (CRNA).

Each paper encompassed within this thesis adopts distinct methodologies tailored to its specific research objectives. The insights gleaned from the exhaustive analysis of the video conferencing network played a pivotal role in the conception of the innovative 10-layer model proposed herein. The inherent versatility of this model lends itself directly to other networks underpinning critical services, thus extending the potential impact of the research beyond its immediate context.

4.1 Data Collection

Each paper in this thesis adopts specific data collection methods tailored to its research objectives. The following is a summary of the data collection methods used in each paper.

For [Paper I], 18 months of network quality data was collected between the networked sites to understand the nature and frequency of network outages. Data from BFD (Binary Forwarding Detection) was collected directly from the routers, and UDP quality data was collected from data-plane measurements. Data was generated by setting up virtual machines (VMs) at each site, and recording packet loss of UDP traffic transmitted between the VMs. Other types of data was also collected (like optical signal strength measurements), but only used in the manual classification process preceding the training stage of the Machine Learning (ML) system. In addition, access to customer support cases and to configuration changelogs was given. This (confidential) data was not used directly, but rather used to determine the impact and root cause of the outages observed in the network quality data.

For risk management in [Paper II] and [Paper V], interviews with stakeholders were performed. These interviews were performed in the ISO27001 context, with the purpose of improving Information Security in the organization, and the aggregated and anonymized results were made available to this research.

The research in [Paper III] is based on 15 years of experience with the MNS network, and the measurement data was retrieved from the working 4-layer heuristics live system. This data was based on commercially available GeoIP data in addition to ICMP ping measurements sent from probes in the global network (See also Section 5).

For [Paper IV], data was collected and compared from various sources, including publicly available databases like MANRS (Mutually Agreed Norms for Routing Security) , PeeringDB, Internet Exchange member lists, the Border Gateway Protocol (BGP) routing table and Routing Public Key Infrastructure (RPKI) validation data, and from Cooperative Association for Internet Data Analysis (CAIDA) data sets. In addition, web-search was used to determine the test objects' ISO27001 certification status. This method of determining the ISO27001 certification status is not a very precise method, and could potentially introduce bias to the data, but no other method was found to get this data, and the method does give a rough estimate of status that can be used to draw high level conclusions.

[Paper V] is a position paper proposing a new model for risk discovery in an availability setting, as described in Section 3.3. Information from various sources including research papers, teaching material, news articles, and blog posts is used to design the proposed classification model. Results were validated by stakeholder interviews with MNS employees, during the preparation for the yearly ISO27001 audits in 2022 and 2023. In addition, the methods were used when deploying the new CRNA network, as presented in Section 7.1.

[Paper VI] is based on the data collected for [Paper I] without new data collection.

In [Paper VII], service location data was collected using the Domain Name System (DNS) and the Maxmind GeoIP database. Data about application dependencies was collected using wireshark to monitor network traffic during application or service usage.

Overall, the data collection methods used are diverse and appropriate for each paper's research focus. Together, they demonstrate a comprehensive approach to gathering data to address availability risk management in the specific context of the small/medium network operator setting.

4.2 Risk Management Methodology: Bridging Theory and Practicality

Within this thesis, risk management is a central theme, forming a crucial bridge between theoretical frameworks and practical implementation of resilient networks. As outlined in Section 2, a substantial body of actors has contributed to the development of various risk management techniques. In our experience, practical risk evaluation for network- and service-providers often relies on commercial standards like ISO27001, NIST800-53, and SOC2/SOC3.

These commercial standards have gained significant recognition in the industry, due to their comprehensive framework for information management and their coverage of relevant topics (see [Paper V]). Many of these standards encompass formal audits leading to certifications, and a large number of consultancy firms specialise in helping companies achieve and sustain these certifications. Both the certification standards and other non-certification standards play a crucial role in demonstrating an organization's commitment to Information Security when interacting with customers and governing bodies.

For the projects undertaken in this research, the methodology of risk management significantly aligns with the tenets of Information Security Management Systems (ISMS) outlined in the ISO27001 standard. Building upon this foundation, the model presented in [Paper V] offers a valuable extension, furnishing a tailored approach to availability risk management.

Incorporating the Common Vulnerability Scoring System (CVSS) [45] within the scope of this research provides a robust risk metric. While initially developed for quantifying risk associated with software vulnerabilities, it is here employed innovatively to construct a Vulnerability Score (VS) dedicated to availability risk management, as detailed in [Paper IV].

Furthermore, [Paper VI] delves into the intricate domain of risk prioritization. It raises pertinent questions about the focus of risk mitigation efforts and suggests utilizing quantified impact and Risk Scores (RS). In a similar vein, the concept of quantified Return on Security Investments (RoSI) is explored in [Paper IV] and [Paper VI]. These explorations underscore the integration of quantitative elements into the risk management framework, enhancing its efficacy and aiding strategic decision-making.

4.3 Artificial Intelligence

Artificial Intelligence (AI), particularly Machine Learning (ML), has gained prominence across technological domains, from self-driving cars to applications like ChatGPT [46]. ML was introduced by in 1959 by Arthur Samuel [47], and has recently gained attraction in a lot of different fields, facilitated by accessible computational resources. In [Paper I] we show one application where ML is trained on packet loss data supervised by root cause analyses, and use the trained model to suggest root causes for unknown outages.

In this thesis, the well established Support Vector Machine model (SVM, from 1992) has been utilized for ML. The computational power required for Machine Learning can be heavy, and the main advantage of SVM is that it transforms the ML problem into a linear separation problem as illustrated in figure 4.1, which requires less computation to solve. SVM is also particularly well suited to classify new data that has not previously been observed.

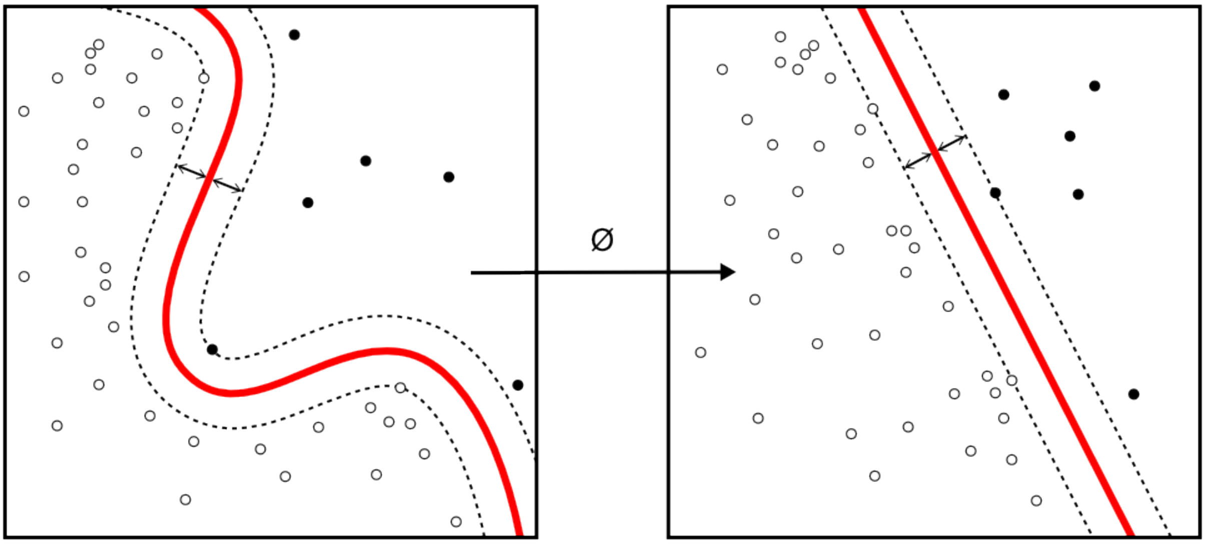

Figure 4.1: Illustration of SVM transformation
Source: `https://upload.wikimedia.org/wikipedia/commons/1/1b/Kernel_Machine.png`

Our results indicate that AI will be an essential tool in achieving network resilience and security. In our research we have demonstrated the particular case of AI's superior ability to quickly find patterns in a large dataset of network quality measurements.

Chapter 5

Ethical Considerations for Internet Measurements

For any research, investigating the ethical aspects is important. Is my experiment appropriate? Is it necessary? Harmful? Can the same results be acquired using data that has already been collected for other purposes? Using Internet measurements to investigate how Critical Services can be more resilient is in my opinion both appropriate and necessary due to the importance of this infrastructure for people, companies and governments. As discussed below, the measurements performed in this research are not harmful. There are other data sets available, but public data sets do not contain "ground truth" like customer complaints and high precision measurements like link status and optical strength that I have used in this analysis, so my work would not have been possible using public available data collected for other purposes. However, ethical questions related to privacy, data collection process, third-party impact of the measurements, and data quality are important.

While the social and medical sciences have the Belmont Report [48] as a standard for ethical research, no such universal recommendations exist in the field of Internet measurements. Some initial guidelines have been published through the preferred channels for the Internet community, as Request For Comments (RFC) documents. RFC1087 [49] "Ethics and the Internet" and RFC1262 [50] "Guidelines for Internet Measure Activities" are fairly dated and rather broad, essentially prescribing "Do the Right Thing" [51].

Contributions like [52, 53, 51] have sought to create a more usable framework for ethics considerations, which I will draw upon to underpin the ethical analysis of my PhD thesis.

The ethical considerations related to my PhD research can be categorized into the following areas:

- Handling of sensitive data, including privacy concerns.
- Usage of data collected by others, including mixing of roles.
- Causing third-party impact from active or passive measurements.
- Data quality.

5.1 Handling of Sensitive Data

For passive measurements, the most important ethical concerns are the legal and privacy implications. In some cases, measurement logs might contain sensitive information and/or personally identifiable information. Hence, a vital step involves identifying and classifying the collected data. Best practices involve swiftly reducing data confidentiality and identifiability through techniques like anonymization or aggregation. When anonymizing data, there are also multiple classes of sensitivity to be aware of. Even anonymized data may be personally identifiable under some circumstances, or may identify a group, for instance if IP addresses are anonymized by prefix only.

A common criterion for data collection is that the data subjects should give "Informed Consent" to the collection of data. When collecting Internet traffic data, this is often not possible, and in such cases, extra care should be exercised. Information about customers is regarded as sensitive data and has been anonymized.

Data that is not Personally Identifiable (PI) may still be confidential, for instance when related to trade secrets or proprietary information.

Storage of collected data must likewise be handled according to the data classification, with classified information being stored only if strictly necessary, and while taking appropriate protection measures. The General Data Protection Regulation (GDPR) [54], the most important European law to ensure correct handling of data, does not apply since this work only processes anonymized and aggregated data.

To ensure proper handling of research data, care was taken to only store data on internal servers, even when the data was judged to be not sensitive.

5.2 Usage of Data Collected by Others

Parts of the data used in this thesis was sourced from logs provided by the Network Operations Center (NOC) at Media Network Services AS. An ethical analysis must always be performed to determine whether data that was collected for other purposes may be used for research. The data usage policy has been defined in an agreement between the data provider and the researcher. In this particular case, the third party logs were collected for the explicit purpose of improving the service delivered to customers, and thus the present research is well within this purpose.

An ethical question is whether to publish results that reflect badly on a company that has cooperated by sharing data for research purposes, and whether to anonymize or not such results. In the papers leading up to this thesis, the source was anonymized, due to the principle of general double-blind evaluation, while in the thesis introduction we have chosen to identify the source since this information would be possible to deduce anyway, from the participants' background.

When using data collected by a third party, it is often good practice to identify and acknowledge the third party and show how the data was attained. It is important to acknowledge any selection bias and any unknown factors in the data collection.

In the context of this thesis, a unique ethical dilemma arises due to the intertwining of personal roles. I hold the dual positions of both researcher and Information Security Officer within the organization that provides the data for this study. This intricate dual role necessitates stringent measures to ensure the replicability of experiments, enabling other researchers to conduct analogous studies. Furthermore, it demands heightened vigilance to maintain the company's established security standards while facilitating the sharing of pertinent research data. To address these concerns, a meticulous protocol has been implemented. All processing and aggregation of sensitive information occur within the data producer's infrastructure, under my capacity as an employee. Subsequently, only anonymized and aggregated data is transmitted to the research computing systems for thorough analysis and potential publication of results. This rigorous approach underscores the commitment to ethical and professional standards while facilitating valuable research insights based on measurements from a live network.

5.3 Potential Third-party Impact

Active measurements introduce the risk of causing disturbance to the system under test. In these cases, a risk/benefit analysis must be performed to gauge whether the potential impact on the service is justified by the possible research outcome. Any measurement systems placed into the core network of a network operator must be done in such a way that it causes as little impact as possible on the live network. A good example of this is the measurements performed in [Paper I] where the only change to the production platform was to enable the transmission of SNMP traps when BFD events occurred, and then off-load the processing and refinement to non-production servers. The generation of SNMP traps is a standard feature in the core routers and mechanisms are already in place to suppress such traps during high CPU usage periods.

A common active measurement method is ICMP ping transmission, which is the same protocol that may be used to perform ICMP flood DDoS attacks. Because of this concern, measures were taken to limit the number of ping packets sent per destination, and to monitor the system to make sure these limits are observed.

It is worth noting that any device connected to the Internet should be able to withstand a certain level of pings, considering the constant stream of such scans reaching all Internet IP addresses, commonly referred to as the “Internet Background Radiation”. Previous research [55] demonstrated that Windows and MacOS webservers could handle ping rates of around 500Mbps before experiencing any service degradation. This value is six orders of magnitude higher than the active measurements conducted in [Paper III]. It is important to differentiate these research measurements from typical Distributed Denial-of-Service (DDoS) attacks, which are often designed to exploit specific vulnerabilities and trigger service disruptions rather than simply inundating the target with a massive ping flood.

In line with ethical research principles, recipients of the testing probes should get information to opt-out or report any unwanted actions. Each IP address used in this research was registered with an “abuse” RIPE contact email address for this purpose.

By carefully managing the number and rate of pings, and by considering the broader context of Internet traffic and server capabilities, the research team ensured responsible and valid measurements while minimizing any potential impact on the devices and networks involved in the study.

Chapter 6

Research Papers Overview

This chapter provides an overview of the seven research papers written as part of this thesis, highlighting their key contributions and findings. The papers collectively explore various aspects of availability risk management, ranging from outage classification and risk assessment frameworks to the analysis of Internet risks and national Information and Communication Technology (ICT) resilience. Through these research endeavors, valuable insights have been gained, contributing to the understanding and enhancement of availability risk management practices in the context of network and service providers.

Each paper offers a unique perspective, addressing specific challenges related to availability risk. The following sections summarize the main objectives, methodologies, and outcomes of each paper, highlighting their significance and how they contribute to the overall body of knowledge in the field. By examining these papers collectively, we gain a comprehensive understanding of the research landscape and the advancements made in availability risk management.

The following sections briefly introduces each of the seven papers.

6.1 Paper I: Crosslayer Outage Classification Using Machine Learning

In this paper, an analysis of outage data from the MNS network is conducted to investigate outage classification techniques. The network quality was continuously monitored for a duration of 18 months, resulting in an extensive dataset of packet loss across all network layers.

To gain further insights, the researchers also analyzed a comprehensive database of 2855 customer complaints. By using information from all layers of the network, the root causes of the outages were identified.

Over the course of the 18-month period, 717352 packet loss events were recorded through active UDP (User Datagram Protocol) measurements across both point-to-point Layer 2 links and Layer 3 routed IP paths. Additionally, passive Layer 2 outage data was collected from BFD (Bidirectional Forwarding Detection) events.

A Machine Learning (ML) model was trained using events that had known root causes and resulted in customer support cases. This trained model was then used to extrapolate and provide indications for the remaining incidents.

The study's results demonstrated the effectiveness of a two-stage approach. In the first stage, only the Layer 2 BFD data was used, enabling accurate identification of almost all Layer 2 outages, including the specific problem type and location. In the second stage, the multi-layer packet loss data was employed to classify the remaining outages.

The proposed method achieved remarkable success, with an f1-score of 0.92 for the majority of outage categories, making it a very useful tool for a resilient network operator.

The results from the research showed significant risk from the physical, local and wide area network layers. In addition many problems originated in maintenance activities, i.e. human errors either in the planning phase or in the execution phase of these maintenances.

6.2 Paper II: ISO27001 as a Tool for Availability Management

This paper discusses the risk management framework implemented at MNS and its effectiveness in ensuring availability management. The company employed stakeholder interviews to identify all potential risks to its operations. Each identified risk was assessed based on its likelihood and potential impact, each assigned a score ranging from 0 to 5.

Furthermore, the paper presents a comprehensive analysis of risk management scores from MNS over a period of 5 years, which were generated as part of the annual ISO27001 audits. These scores encompassed both previously identified risks and newly identified ones.

The findings of the study demonstrate a significant reduction in risk scores following the implementation of the ISO27001 standard. Moreover, the risk scores continued to decrease over the subsequent years, indicating the ongoing effectiveness of the risk management framework in improving availability.

This research shows that the implementation of information management systems like ISO27001 significantly improves resilience risk for network operators, focused on risk management in the Services, Organizations, People, and Governance layers.

6.3 Paper III: Evolved Cold-Potato Routing Experiences

This paper builds upon a previous work from 2013 [56] by the same authors, which described the implementation of a GeoIP-based global routing method and evaluated its advantages, highlighting a considerable reduction in packet loss at the expense of a moderate increase in latency. The method involved manipulating BGP announcements using the “local preference” metric.

While the previous implementation was deployed in production, it revealed areas for improvement as indicated by ongoing customer cases. This paper focuses on further enhancing the routing method, introducing the “Four layer heuristics” approach.

The first heuristic remains GeoIP-based, as outlined in the original paper, but with significant performance enhancements achieved by using a memory-cache in front of the MySQL database. The second heuristic incorporates an offline latency measurement process, periodically updating the GeoIP database with latency measurements for 85.5% of all Internet prefixes. This addresses issues related to imperfect GeoIP data and the mismatch between Internet topology and geography.

Heuristic three involves online probes that continuously measure packet loss and latency to traffic destinations observed in the actual network traffic, allowing for real-time routing adjustments. This heuristic accommodates dynamic network changes occurring within shorter timeframes compared to heuristic two and helps avoid temporary congestion points on the Internet.

The fourth and final heuristic involves manual corrections. In some cases, there may be a desire to prioritize paths based on political considerations, such as keeping traffic within a specific country. Additionally, certain prefixes may be challenging to measure and require manual routing adjustments.

Collectively, the four layer heuristics significantly improve network quality, resulting in a notable decrease in customer cases despite a substantial increase in network traffic. These enhancements demonstrate the effectiveness of the evolved “cold-potato” routing method to improve network resilience in the Wide Area, Internet, and Cloud layers.

6.4 Paper IV: How Large Is the Gap? Exploring MANRS and ISO27001 Security Management

The fourth paper focuses on Internet risks, particularly the potential impact of malicious or accidental routing changes by third parties on network availability. It uses this context to compare the ISO27001 standard and the MANRS (Mutually Agreed Norms for Routing Security) Internet initiative in terms of their approaches to risk management.

To assess the availability risks stemming from BGP (Border Gateway Protocol) Internet events, the paper proposes the use of methods developed for the Common Vulnerability Scoring System (CVSS) to conduct Vulnerability Scoring. This allows for the quantification of risks associated with these events.

Additionally, the paper introduces the Return on Security Investments (RoSI) analysis, which highlights the benefits and cost-effectiveness of implementing the ISO27001 standard versus adopting the relatively inexpensive actions outlined by the MANRS initiative. By employing RoSI analysis, the paper illustrates the synergies between these two approaches and their respective impacts on mitigating availability risks.

Through this comparative analysis, the paper sheds light on the strengths and potential gaps of the ISO27001 standard and the MANRS initiative in organizations addressing Internet risks, particularly in relation to network availability.

6.5 Paper V: A 10-Layer Model for Service Availability Risk Management

The fifth paper builds upon the lessons learned from the previous articles, emphasizing the importance of a clear framework for effective risk management, particularly in relation to availability risk. It highlights the challenges faced when attempting to manage risks without a structured approach, despite employing methods commonly used in the ISO27001 and NIST 800-53 information security standards.

In this position paper, a new 10-layer model is proposed by adapting and extending the ISO/OSI 7-layer model. The key objective is to group risks based on responsibilities and common risk factors rather than focusing solely on technical implementation. This comprehensive model addresses potential risks that are often overlooked, such as those stemming from Cloud, People, Organizations, and Governance.

Furthermore, the paper explores the relationship between the different layers and their ability to mitigate risks across layers. By analyzing outages caused by risks from each layer, it highlights the interconnected nature of these layers and the significance of their interactions in effectively managing and mitigating risks.

Through the introduction of this 10-layer model, the paper provides a valuable framework for service availability risk management, emphasizing the need to consider a wide range of risks and their relationships across multiple layers. Some layers described in this paper were further elaborated from (or not covered by) previous papers, namely Cloud, Applications, and Services.

6.6 Paper VI: Outage Risk Priorities - It's not the malicious attacks that take down your service

The sixth paper builds upon the insights gained from previous studies to analyze risks and outages in the context of risk management. Using the same data material as [Paper I], this research aims to identify the types of outages that have the greatest impact on service delivery and determine the most common types of outages.

By leveraging this data, the paper introduces an Impact Score calculation method that enables the evaluation of mitigation actions in terms of achieving the best Return on Security Investments (RoSI). Contrary to common perception, the findings reveal that malicious attacks rarely cause service availability issues. Instead, the study highlights that instabilities in network links and equipment maintenance or failure have a far more significant impact on service availability, showing high impact from Physical, Wide Area Networks, Organization, and People risks.

While the occurrence of malicious activity resulting in service availability issues, particularly in the form of Distributed Denial of Service (DDoS) attacks, is relatively rare, it is noteworthy that such incidents can have a disproportionately high Impact Score.

In summary, this paper provides valuable insights into outage risk priorities, challenging the notion that malicious attacks are the primary cause of service disruptions. By focusing on addressing network link instabilities and equipment maintenance/failure, organizations can effectively allocate their resources and investments to achieve optimal service availability and enhance their Return on Security Investments (RoSI).

6.7 Paper VII: National ICT Resilience: An analysis of Norway's cyber infrastructure preparedness

The seventh paper in this thesis addresses a critical gap in the research, focusing on the risks associated with Governance. The paper uses Norway as an example to assess the extent of individual countries' self-sufficiency in terms of Internet connectivity. While our two blog posts [Blog I, Blog II] have already highlighted the impact of events in Ukraine in 2022, and the controversies surrounding AfriNIC, this paper aims to provide a comprehensive assessment of dependencies on external organizations and locations. Due to a lack of data from the actual services, their public facing web pages are instead evaluated. The study analyzes 207 stately Norwegian web pages, considering three key vectors: web hosting location, mail exchanger domain, and name server domain.

The research findings reveal that nearly all of the analyzed web pages exhibit some level of international dependency. This highlights the potential risks associated with relying on external organizations and locations for Internet services. The study further extends its analysis to include top web pages from the .no, .se, and .dk domains, using the Top 1 million web pages list as a reference.

The results indicate similar trends across all three countries, with a significant number of websites displaying strong international dependencies. These findings underscore the high level of risk associated with such dependencies, emphasizing the importance of assessing and managing these risks to ensure national ICT resilience.

By demonstrating the extent of international dependencies in the context of Internet connectivity, this paper provides insights to the potential vulnerabilities and risks that nations face. The findings serve as a call to action for policymakers and stakeholders to prioritize strategies that enhance national ICT resilience and reduce dependence on external entities, highlighting risks in the Cloud, Services, Organizations, and Governance layers.

6.8 Summary

The papers that form the core of this thesis together encompass resilience and security risks at all levels of network operations, from the low level outages to national governance risks. Together, they provide a thorough theoretical and practical background leading up to the proposed 10-layer framework for Resilient Networks. We demonstrate that such frameworks are well suited to reduce risk, we show what types of risks are present at all levels, and show the efficiency of mitigating actions for some selected risks. Finally, we show that the model works in a real network setting, as further elaborated in Section 7.1.

Chapter 7

Experimental Validation of the 10-layer Model

The theoretical framework presented in [Paper V] showed a practical application by using the 10-layer model when performing risk analysis in the same setting as the research leading up to the model. However, to reinforce its applicability across diverse settings, it was beneficial to subject the model to validation in a different environment. This imperative to establish a stronger link between theory and real-world implementation led to the initiation of a second validation phase.

This drive for a more comprehensive validation prompted the evaluation of the various methodologies proposed in the research papers within an operational scenario. This assessment occurred during the establishment of the new CRNA (Center for Resilient Networks and Applications) research network. The CRNA network was designed as a a small-scale network, covering (parts of) a metropolitan area in Oslo, implemented as a separate Autonomous System (AS). This network consists of routers, firewalls and switches, as well as several servers operated through the OpenStack Cloud Computing Management.

Through this practical application, insights gleaned from the array of research papers played a pivotal role in shaping the design and deployment of the CRNA network. This iterative approach fostered a more nuanced understanding of the challenges of translating theoretical frameworks into tangible operational systems. This synthesis not only provided real-world validation for the proposed methodologies but also underscored their practical utility and adaptability across different operational contexts.

7.1 Outage Classification Using Machine Learning

The CRNA network contains only one WAN link. This link is comprised of two sections, one section is a campus-area dark fiber owned by the Oslo Metropolitan University, and the second section is a rented metropolitan area dark fiber with a distance of around 1 km. From [Paper I], we learned that BFD trap data was an efficient tool for outage classification. Therefore, BFD is configured with SNMP traps sent to a collection server. In addition, optical data and root cause data are collected according to the method from [Paper I].

7.2 ISO27001 as a Tool for Availability Management

The research in [Paper II] and [Paper IV] show that implementation of the methods from the ISO27001 standard are well suited for availability risk reduction for a network operator.

Therefore, an Information Security Management System (ISMS) for the new research network is being developed according to ISO27001 standards. The ISMS contains:

- Policies and procedures
- Configuration management systems
- Monitoring systems
- A risk analysis and management process

The first part of the ISMS contains policies and procedures for users and administrators of the research network. The most important policies concern configuration management and acceptable use. This network will be used by various permanent and temporary employees and also by guests. The various projects may have different security requirements, and the policy serves as a training tool for what is acceptable use and who to contact for various issues (from generic support to security breaches). For administrators, a configuration change audit policy is important to guard against human errors. Furthermore, the suggested policy demands a yearly security audit to make sure all software is up to date, and the state of the firewalls and user accounts is correct. The security audit report is presented to senior management to show that the engineering team is confident in the network operations, and to illuminate any known issues.

A risk analysis was performed at two points in time, before the implementation of the new research network, and after the implementation, similar to that described in [Paper II]. The results show a large reduction in risk in nearly all areas. The total risk score was reduced by 50%, indicating a a significant improvement to the service risk. The long-term effects remain to be seen.

7.3 Internet Routing

[Paper IV] shows how Internet Routing can be optimized in a Wide Area Network. The research network iBGP was set up with two edge routers with one IP transit each, to the same provider, connected by iBGP. This network is not suitable for implementation of the 4-layer heuristics described in [Paper IV]. These methods are more suitable for networks that span a larger geographic area or that have multiple independent BGP exits. These methods may, however, be deployed in a future research project simulating a customer AS.

7.4 Use of MANRS and ISO27001 recommendations

According to [Paper IV], there are significant synergies to implementing MANRS and ISO27001 at the same time, and minimal extra cost. Hence, the necessary actions for MANRS participation were implemented for the new network. The MANRS implementation consisted of:

- IP address verification
- IRR and PeeringDB information
- RPKI signatures
- Reverse path IP filters

IP address verification was done by checking the RIPE (Réseaux IP Européens) entries for own and customer ASes. All these addresses belong to CRNA or closely related entities, and were easily verified, and the policy was enforced by BGP filters.

Contact information and other information in the Internet Routing Registry (IRR, RIPE) was found to be outdated, and was therefore updated. The network was not registered in PeeringDB, so a registration was added.

For RPKI, a Certificate Authority (CA) was created within RIPE's portal, along with a corresponding Route Origin Authorisation (ROA) for the prefixes used in the CRNA research network. At the time of writing, the optional action of Route Validation (ROV) was not implemented in the network, but the routers do support this, and a virtual machine for route validation will be established.

Reverse path IP filters were implemented in the firewall, to make sure only approved source IP Addresses are used inside the research network. One small DMZ (De-Militarized Zone) network is exempted from this policy, to facilitate research into using spoofed source IP addresses where our researchers are dependent on own public IP addresses.

7.5 The 10-layer Model for Service Availability Risk Management

During the discovery phase of the ISMS risk management process, the 10-layer method from [Paper V] was used. 55 risks were identified, many of them with high risk values.

To ensure no risks were ignored, a brief risk discovery process was also conducted using the ISO27001 controls, and the result was that no further availability risks were discovered.

The identified risks were assessed and suitable actions were taken to reduce the risks to an acceptable level.

7.6 Availability Risk Prioritization

During the risk discovery phase for the CRNA network, historical data was not available. In light of this, the impact scores established in [Paper VI] provided a foundational basis for subsequent risk prioritization. Standard industry best practice firewalls were deployed, and special attention was directed towards limiting the impact of faults on the identified prioritized risk areas of wide area links, equipment failure and optical link vulnerabilities. Furthermore, human errors emerged as an identified risk demanding heightened precedence. To effectively attenuate this risk, supplementary procedures for peer review of configuration changes was implemented and comprehensive training resources for users of the network were created.

7.7 National ICT Resilience

The CRNA network has been analysed for dependencies on international services. No email service is operated by CRNA. The IP addresses are provided through a Norwegian LIR, and all research data is stored on servers operated by Norwegian entities SimulaMet, Simula, and OsloMet. However, the domain crnalab.net is handled by US company Cloudflare. All addresses are documented so the network can be operated using IP addresses only in case of a failure of the domain service. All hardware is produced abroad, so supply-chain risk has been considered and central components have been duplicated. Documentation of the CRNA network is stored in Google's cloud service, but backup copies are stored locally. VPN (Virtual Private Networks) and https access are provided using certificates signed by international Certificate Authorities (CA), and might be unavailable in case of a compromised CA. In summary, the CRNA network is sufficiently well designed to run and provide most services independent of any external dependency.

7.8 Results

Although the revamped CRNA network is still relatively new and hasn't undergone an extended operational phase, the implemented measures are functioning well. The network management team appears well-prepared, fostering secure and efficient network operations. This attests to the practicality of the 10-layer model and validates the relevance of findings from the seven research papers.

In parallel, these actions reinforce the significance of the proposed 10-layer model, aligning with insights from the research papers. Together, they affirm the model's utility and its applicability, giving the CRNA network a high level of resilience and security, and contribute to the overall objectives of this thesis.

Chapter 8

Discussion and Conclusion

The most important contribution from this thesis is the 10-layer model for availability risk management [Paper V]. The comprehensive model is field-proven in a global network, practical, and will help any operator provide resilient networks for critical services, both during an initial risk discovery undertaking and as a continuous process of improving resilience.

The research leading up to the 10-layer model was performed to address specific resilience risks:

An efficient and practical method for outage classification was developed and presented in [Paper I], speeding up network troubleshooting and recovery time.

The effects of the security standards ISO27001 and MANRS were examined in [Paper II] and [Paper IV], showing significant resilience risk improvements from their deployment.

A field-proven system for cold-potato optimized BGP routing was created in [Paper III], improving network latency, packet loss, and outage resilience.

Further, a quantified method for risk prioritization was laid out in [Paper VI], showing which root causes lead to actual outages affecting customers, and guiding the focus for efficient security investments.

Finally in [Paper VII], national ICT risk was investigated, showing a strong dependence on international actors for national web pages, and suggesting guidelines for secure and resilient networks for critical services.

In this thesis, we have successfully addressed the research questions posed at the beginning of our study, providing a deeper understanding of availability risk management for critical services. Additionally, our research has yielded valuable insights and made significant contributions to the field.

RQ1: What are the main availability risks to critical services?

Through the analysis of outage data, customer complaints, and network monitoring, we have identified the primary availability risks faced by critical services. Papers I, IV, V, and VI have revealed that network instabilities, equipment maintenance or failure, and human (maintenance) errors are key factors contributing to service outages. These findings emphasize the need for proactive measures to address these risks and ensure the reliability and continuity of critical services. Moreover, Paper VII highlights the risks associated with dependencies on external organizations and locations, further underscoring the importance of assessing and managing these dependencies for ensuring national ICT resilience.

RQ2: What is the best way to organize availability risk management?

Our research has demonstrated the value of adopting established frameworks and standards for availability risk management. Papers II, IV, and V have highlighted the benefits of using ISO27001 as a tool for information security management, providing a comprehensive framework for identifying, assessing, and mitigating risks. The 10-layer model proposed in Paper V offers a structured approach to organizing availability risk management efforts, grouping risks based on responsibilities and commonalities rather than technical implementation. It provides a holistic view of risk management and facilitates effective coordination among different layers and stakeholders. These findings suggest that a combination of a formal framework and a well-defined organizational model can enhance the effectiveness and efficiency of availability risk management.

RQ3: How can the most important availability risks be mitigated?

Our research presented in Papers I, III, IV, and VI has demonstrated the effectiveness of employing machine learning algorithms, vulnerability scoring methods, and heuristic-based routing approaches to improve network resilience and reduce service outages. These findings highlight the importance of leveraging advanced technologies and methodologies to enhance risk mitigation efforts. Additionally, Paper VII emphasizes the importance of assessing and managing dependencies on external organizations and locations, suggesting that service localization can be a key strategy to reduce risks associated with such dependencies. It is important to acknowledge the limitations of our project, particularly regarding the data used for analysis. The reliance on data from the MNS network may introduce biases and reduce the generalizability of our findings to other networks. However, the use of a data from a consistent infrastructure across the research papers enables comparability and relevance, especially for networks similar to MNS, a global network for critical services. The findings were sucessfully tested on the CRNA research network, and future research could consider incorporating data from a broader range of networks to validate and extend our findings.

Throughout the research period, the shape of our study evolved based on the discoveries and insights gained. The initial work on outage classification in Paper I revealed unforeseen patterns, leading to further exploration in Papers II-IV. These papers laid the foundation for the proposed 10-layer model in Paper V, which provided a comprehensive framework for availability risk management. During my work with the 10-layer model it became clear that Governance and Risk prioritization are important aspects of resilient networks, leading to further research in that area in Papers VI and VII.

In conclusion, this thesis has contributed to the understanding of resilience management through a comprehensive exploration of various layers, risks, and mitigation strategies. The findings and frameworks developed in the seven research papers have practical implications and can serve as a resource for organizations seeking to enhance the resilience and availability of their critical services. By combining theoretical insights with real-world data and practical applications, our research aims to drive advancements in the field and support the continuous improvement of availability risk management practices. The findings should also serve as a call to action for policymakers and stakeholders to prioritize strategies that enhance national ICT resilience and reduce dependencies on external entities.

Throughout the research, numerous potential avenues for further exploration have emerged, signifying the dynamic nature of availability risk management. Notably, the proposed 10-layer model exhibits its efficacy in facilitating network risk management. However, the prospect of adapting similar models to cater to the distinct risk landscapes of various industries holds significant promise. Tailored frameworks could equip industries beyond networking with a structured approach to tackle their unique risk challenges systematically.

The rapid proliferation of transformative technologies such as Artificial Intelligence (AI), Internet of Things (IoT), and blockchains introduces intriguing dimensions of both risk and risk reduction. Further research into AI has the potential to optimize connections, predict faults, and enhance overall network resilience. However, the adoption of AI also introduces potential risks, particularly when dealing with non-explainable AI decision-making processes.

Similarly, IoT projects, while promising to reduce physical risk by gathering data from sensors, require deeper investigation into their implications on confidentiality, integrity, and availability. The use of blockchains can be explored as a tool for enhancing traceability, such as tracking configuration changes, thereby minimizing the impact of human errors.

Investigating the impact of these technologies on availability risk and resilience in various domains represents a captivating avenue for future inquiry.

The domains of Cloud services and Governance risk present ripe opportunities for more extensive exploration. A deeper investigation into the unique risks and mitigation strategies inherent in cloud-based critical services could yield valuable insights. As Cloud services continue to transform the IT landscape, understanding their influence on availability risk becomes essential.

Moreover, delving into Governance risk opens doors to cross-country comparisons, unveiling the subtleties of distinct risk mitigation strategies adopted by nations with varying regulatory frameworks. These cross-comparisons offer the potential to illuminate diverse approaches to managing availability risk in critical services on a global scale.

Another pressing area for future research lies in evaluating supply chain risks. In an interconnected world, the supply chain encompasses various components, from hardware suppliers, via software libraries to third-party services. Understanding how vulnerabilities within these chains impact the availability and resilience of critical services is paramount. Investigating strategies to assess, mitigate, and manage these risks can significantly contribute to enhancing overall network security and reliability.

By delving into these avenues, future research can contribute to the continual evolution of availability risk management strategies, ensuring the resilience and security of critical services across various sectors and technological advancements.

Bibliography

[1] James P Anderson. *Computer security technology planning study.* Tech. rep. ANDERSON (JAMES P) and CO FORT WASHINGTON PA FORT WASHINGTON, 1972.

[2] Rute Sofia et al. “IEEE Access Special Section: Internet of Space: Networking Architectures and Protocols to Support Space-Based Internet Services”. In: *IEEE Access* 10 (Jan. 2022), pp. 92706–92709. DOI: 10.1109/ACCESS.2022.3202342.

[3] H. Zimmermann. “OSI Reference Model-the ISO model of architecture for open systems interconnection”. In: *IEEE Trans. Communication (USA)* COM-28.4 (Apr. 1980). IRIA/Lab., Rocquencourt, France, pp. 425–432.

[4] Yakov Rekhter, Susan Hares, and Tony Li. *A Border Gateway Protocol 4 (BGP-4).* RFC 4271. Jan. 2006. DOI: 10.17487/RFC4271. URL: https://www.rfc-editor.org/info/rfc4271.

[5] Colleen McClain et al. “The internet and the pandemic”. In: (2021).

[6] Reethika Ramesh et al. “Network Responses to Russia’s Invasion of Ukraine in 2022: A Cautionary Tale for Internet Freedom”. In: *32nd USENIX Security Symposium (USENIX Security 23).* Anaheim, CA: USENIX Association, Aug. 2023, pp. 2581–2598. ISBN: 978-1-939133-37-3. URL: https://www.usenix.org/conference/usenixsecurity23/presentation/ramesh-network-responses.

[Blog I] Jan Marius Evang. *What is happening to the Internet in Ukraine and Russia?* CRNA’s Blog, https://crna.substack.com/p/what-is-happening-to-the-internet.

[7] Amritha Jayanti. *Starlink and the Russia-Ukraine War: A Case of Commercial Technology and Public Purpose?* https://www.belfercenter.org/publication/starlink-and-russia-ukraine-war-case-commercial-technology-and-public-purpose. 2023.

[8] European Telecommunications Standards Institute (ETSI). *ETSI EN 300 392-2 - Terrestrial Trunked Radio (TETRA).* https://www.etsi.org/deliver/etsi_en/300300_300399/30039202/03.02.01_60/en_30039202v030201p.pdf. 2007.

[9] Athina Lazakidou, A. Ioannou, and Fotis Kitsios. "Use of TETRA Networks in Crisis Situations for Health Information Transfer Strategies". In: *International Journal of Reliable and Quality E-Healthcare* 3 (Dec. 2014), pp. 1–8. DOI: 10.4018/ijrqeh.2014010101.

[10] Inigo Del Portillo, Bruce G Cameron, and Edward F Crawley. "A technical comparison of three low earth orbit satellite constellation systems to provide global broadband". In: *Acta astronautica* 159 (2019), pp. 123–135.

[11] Qian Wang et al. "An Overview of Emergency Communication Networks". In: *Remote Sensing* 15.6 (2023). ISSN: 2072-4292. DOI: 10.3390/rs15061595. URL: https://www.mdpi.com/2072-4292/15/6/1595.

[12] Sanjoy Debnath et al. "A comprehensive survey of emergency communication network and management". In: *Wireless Personal Communications* 124.2 (2022), pp. 1375–1421.

[13] Abhaykumar Kumbhar et al. "A Survey on Legacy and Emerging Technologies for Public Safety Communications". In: *IEEE Communications Surveys & Tutorials* 19.1 (2017), pp. 97–124. DOI: 10.1109/COMST.2016.2612223.

[14] Theodore Tryfonas, Dimitris Gritzalis, and Spyros Kokolakis. "A Qualitative Approach to Information Availability". In: *Information Security for Global Information Infrastructures*. Ed. by Sihan Qing and Jan H. P. Eloff. Boston, MA: Springer US, 2000, pp. 37–47. ISBN: 978-0-387-35515-3.

[15] Juniper Networks, Inc. *Junos OS(R) Standards Reference.* juniper.net, https://www.juniper.net/documentation/us/en/software/junos/standards/standards.pdf. 2023.

[16] Norwegian Directorate for Civil Protection (DSB). *DSB annual reports 2017–2022*. https://www.dsb.no/rapporter-og-evalueringer. 2017–2022.

[Paper V] Jan Marius. Evang. "A 10-Layer Model for Service Availability Risk Management". In: *Proceedings of the 20th International Conference on Security and Cryptography - SECRYPT*. INSTICC. SciTePress, 2023, pp. 716–723. ISBN: 978-989-758-666-8. DOI: 10.5220/0012092600003555. URL: https://doi.org/10.5220/0012092600003555.

[Blog II] Jan Marius Evang. *Africa and the Internet*. CRNA's Blog, https://crna.substack.com/p/africa-and-the-internet.

[17] ANDREW L. RUSSELL. *OSI: THE INTERNET THAT WASN'T*. spectrum.ieee.org, https://spectrum.ieee.org/osi-the-internet-that-wasnt. July 2013.

[18] John Howard and Thomas Longstaff. "A Common Language for Computer Security Incidents". In: (Feb. 1970). DOI: 10.2172/751004.

[19] *ISO 31000:2018 Risk management – Guidelines*. 2018. URL: https://www.iso.org/standard/65694.html.

[20] *ISO 31010:2019 Risk management – Risk assessment techniques*. 2019. URL: https://www.iso.org/standard/72140.html.

[21] Steven Noel et al. "Measuring Security Risk of Networks Using Attack Graphs." In: *IJNGC* 1 (Jan. 2010).

[22] Anoop Singhal and Xinming Ou. "Security Risk Analysis of Enterprise Networks Using Probabilistic Attack Graphs". In: *Network Security Metrics*. Cham: Springer International Publishing, 2017, pp. 53–73. ISBN: 978-3-319-66505-4. DOI: 10.1007/978-3-319-66505-4_3. URL: https://doi.org/10.1007/978-3-319-66505-4_3.

[23] Masoud Khosravi-Farmad and Abbas Bafghi. "Bayesian Decision Network-Based Security Risk Management Framework". In: *Journal of Network and Systems Management* 28 (Oct. 2020). DOI: 10.1007/s10922-020-09558-5.

[24] Chanchala Joshi and Umesh Kumar Singh. “Information security risks management framework – A step towards mitigating security risks in university network”. In: *Journal of Information Security and Applications* 35 (2017), pp. 128–137. ISSN: 2214-2126. DOI: https://doi.org/10.1016/j.jisa.2017.06.006. URL: https://www.sciencedirect.com/science/article/pii/S2214212616301806.

[25] Bill Dixon. *Understanding the FAIR Risk Assessment.* certconf.org, https://www.certconf.org/presentations/2009/files/TA-2.pdf. 2009.

[26] Chanchala Joshi and Umesh Kumar Singh. “ADMIT- A Five Dimensional Approach towards Standardization of Network and Computer Attack Taxonomies”. In: *International Journal of Computer Applications* 100 (2014), pp. 30–36.

[27] Matt Rosenquist. *Prioritizing Information Security Risks with Threat Agent Risk Assessment.* intel.com, https://download.intel.com/it/pdf/Prioritizing_Info_Security_Risks_with_TARA.pdf.

[28] Jan Meszaros and Alena Buchalcevova. “Introducing OSSF: A framework for online service cybersecurity risk management”. In: *Computers & Security* 65 (2017), pp. 300–313. ISSN: 0167-4048. DOI: https://doi.org/10.1016/j.cose.2016.12.008. URL: https://www.sciencedirect.com/science/article/pii/S0167404816301791.

[29] Ken Peffers et al. “A design science research methodology for information systems research”. In: *Journal of Management Information Systems* 24 (Jan. 2007), pp. 45–77.

[30] Mass Lund, Bjørnar Solhaug, and Ketil Stølen. *Model-Driven Risk Analysis - The CORAS Approach.* Jan. 2011. ISBN: 978-3-642-12322-1.

[31] Alexandra Horobet and Lucian Belascu. “The Standardization of Risk Management Practices at the International Level”. In: *Ovidius University Annals, Economic Sciences Series* XV (Jan. 2015), pp. 216–220.

[32] N. Miller et al. “Resilience analysis and quantification for critical infrastructures”. In: Jan. 2020, pp. 364–384. ISBN: 9781680836868. DOI: 10.1561/9781680836875.ch20.

[33] R. Clemente et al. “Risk management in availability SLA”. In: *DRCN 2005). Proceedings.5th International Workshop on Design of Reliable Communication Networks, 2005.* 2005, p. 8. DOI: 10.1109/DRCN.2005.1563900.

[34] David Tipper. “Reconsidering Network Availability and Time”. In: *2019 15th International Conference on the Design of Reliable Communication Networks (DRCN).* 2019, pp. 114–121. DOI: 10.1109/DRCN.2019.8713690.

[35] Piotr Chołda, Piotr Guzik, and Krzysztof Rusek. “Risk mitigation in resilient networks”. In: *2014 6th International Workshop on Reliable Networks Design and Modeling (RNDM).* 2014, pp. 23–30. DOI: 10.1109/RNDM.2014.7014927.

[36] Krzysztof Rusek et al. “RiskNet: Neural Risk Assessment in Networks of Unreliable Resources”. In: *Journal of Network and Systems Management* 31.3 (2023), p. 64.

[37] Alexander Marder et al. “Access Denied: Assessing Physical Risks to Internet Access Networks”. In: *32nd USENIX Security Symposium (USENIX Security 23).* Anaheim, CA: USENIX Association, Aug. 2023, pp. 6877–6892. ISBN: 978-1-939133-37-3. URL: https://www.usenix.org/conference/usenixsecurity23/presentation/marder.

[38] Kentrell Owens et al. “Electronic Monitoring Smartphone Apps: An Analysis of Risks from Technical, Human-Centered, and Legal Perspectives”. In: *31st USENIX Security Symposium (USENIX Security 22).* Boston, MA: USENIX Association, Aug. 2022, pp. 4077–4094. ISBN: 978-1-939133-31-1. URL: https://www.usenix.org/conference/usenixsecurity22/presentation/owens.

[39] Julia Slupska et al. “They Look at Vulnerability and Use That to Abuse You: Participatory Threat Modelling with Migrant Domestic Workers”. In: *31st USENIX Security Symposium (USENIX Security 22).* Boston, MA: USENIX Association, Aug. 2022, pp. 323–340. ISBN: 978-1-939133-31-1. URL: https://www.usenix.org/conference/usenixsecurity22/presentation/slupska-vulnerability.

[Paper II] Jan Marius Evang. “ISO27001 as a Tool for Availability Management”. In: *Proceedings of the International Workshop on Information Management.* WSIM ’22. London, UK, 2022, pp. 82–85. DOI: 10.1109/AEIS59450.2022.00018. URL: https://doi.org/10.1109/AEIS59450.2022.00018.

[Paper VI] Jan Marius Evang. “Outage risk priorities – It’s not the malicious attacks that take down your service”. In: *IEEE Hotnets 2023.* 2023, Submitted for evaluation.

[Paper I] Jan Marius Evang et al. “Crosslayer Network Outage Classification Using Machine Learning”. In: *Proceedings of the Applied Networking Research Workshop.* ANRW ’22. Philadelphia, PA, USA: Association for Computing Machinery, 2022, pp. 1–7. ISBN: 9781450394444. DOI: 10.1145/3547115.3547193. URL: https://doi.org/10.1145/3547115.3547193.

[40] Olav Lysne. *The Huawei and Snowden Questions.* Jan. 2018. ISBN: 978-3-319-74949-5. DOI: 10.1007/978-3-319-74950-1.

[41] IEEE. “IEEE Standard for Ethernet”. In: *IEEE Std 802.3-2022 (Revision of IEEE Std 802.3-2018)* (2022), pp. 1–7025. DOI: 10.1109/IEEESTD.2022.9844436.

[42] Y. Rekhter and T. Li. *A Border Gateway Protocol 4 (BGP-4).* RFC 1771. Mar. 1995. DOI: 10.17487/RFC1771. URL: https://www.rfc-editor.org/info/rfc1771.

[Paper IV] Jan Marius Evang and Ioana Livadariu. “How Large Is the Gap? Exploring MANRS and ISO27001 Security Management”. In: *The 31st International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2023)*, to appear.

[43] OECD. *Routing security.* www.oecd-ilibrary.org, https://www.oecd-ilibrary.org/content/paper/40be69c8-en. 2022. DOI: https://doi.org/https://doi.org/10.1787/40be69c8-en.

[Paper III] Jan Marius Evang and Tarik Cicic. “Evolved Cold-Potato routing experiences”. In: *The 31st International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2023)*, to appear.

[44] Tarik Cicic et al. “Network Routing”. U.S. pat. 9426055. Aug. 23, 2016.

[Paper VII] Jan Marius Evang and Haakon Bryhni. “National ICT Resilience: An analysis of Norway’s cyber infrastructure preparedness”. In: *IEEE Communications Magazine*. 2023, Submitted for evaluation.

[45] Forum of Incident Response and Security Teams. *CVSS - Common Vulnerability Scoring System*. `https://www.first.org/cvss/`. 2023.

[46] Tom B Brown et al. “Language Models are Few-Shot Learners”. In: *arXiv preprint arXiv:2005.14165* (2020).

[47] A. L. Samuel. “Some Studies in Machine Learning Using the Game of Checkers”. In: *IBM Journal of Research and Development* 3.3 (1959), pp. 210–229. DOI: `10.1147/rd.33.0210`.

[48] National Research. “The Belmont Report. Ethical principles and guidelines for the protection of human subjects of research”. In: 81 (June 2014), pp. 4–13.

[49] Defense Advanced Research Projects Agency Internet Activities Board. *Ethics and the Internet*. RFC 1087. Jan. 1989. DOI: `10.17487/RFC1087`. URL: `https://www.rfc-editor.org/info/rfc1087`.

[50] V.G. Cerf. *Guidelines for Internet Measurement Activities*. RFC 1262. Oct. 1991. DOI: `10.17487/RFC1262`. URL: `https://www.rfc-editor.org/info/rfc1262`.

[51] Christos Papadopoulos and John Heidemann. “Towards best practices for active network measurement”. In: *Proceedings of the CAIDA AIMS Workshop*. 2009.

[52] Jedidiah Crandall, Masashi Crete-Nishihata, and Jeffrey Knockel. “Forgive Us our SYNs”. In: Aug. 2015, pp. 3–3. DOI: `10.1145/2793013.2793021`.

[53] Jeroen Van der Ham and Roland Rijswijk-Deij. “Ethics and Internet Measurements”. In: *Journal of Cyber Security and Mobility* 5 (Jan. 2017), pp. 287–308. DOI: `10.13052/jcsm2245-1439.543`.

[54] European Parliament, Council of the European Union. *Regulation (EU) 2016/679 - General Data Protection Regulation*. `https://eur-lex.europa.eu/eli/reg/2016/679/oj`. 2016.

[55] Ganesh Reddy and Sanjeev Kumar. “Do ICMP Security Attacks Have Same Impact on Servers?” In: *Journal of Information Security* 10 (July 2017), pp. 274–283. DOI: 10.4236/jis.2017.83018.

[56] Ahmed Elmokashfi et al. “Geography Matters: Building an Efficient Transport Network for a Better Video Conferencing Experience”. In: *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies*. CoNEXT ’13. Santa Barbara, California, USA: Association for Computing Machinery, 2013, pp. 369–380. ISBN: 9781450321013. DOI: 10.1145/2535372.2535395. URL: https://doi.org/10.1145/2535372.2535395.

Glossary

ADMIT Attack vector, Defense, Method, Impact, Target. 18

AfriNIC African Network Informaition Center. 32

AI Artificial Intelligence. 38, 62

AS Autonomous System. 7, 10, 27, 53, 55

BFD Bidirectional Forwarding Detection. 46, 54

BGP Border Gateway Protocol. 7, 27, 36, 55

BSI British Standards Institute. 17

CA Certificate Authority. 56, 57

CAIDA Center for Applied Internet Data Analysis. 36

CERT the United States Computer Emergency Readiness Team. 18

CIS Center for Internet Security. 17

CISA (US) Cybersecurity Infrastructure Security Agency. 8

CORAS Consultative Objective Risk Analysis System. 19

COVID-19 Corona Virus Disease 2019. 8

CRNA Center for Resilient Networks and Applications. 35, 37, 53

CVSS Common Vulnerability Scoring System. 38

DDoS Distributed DoS. 44

DFARS Defense Federal Acquisition Regulation Supplement. 17

DMZ De-Militarized Zone. 56

DNS Domain Name System. 29

DSB Direktoratet for Sivil Beredskap. 8, 10

DSR Design Science Research methodology. 19

ENISA European Union Agency for Cybersecurity. 17

FAIR Factor Analysis of Information Risk. 18

FEC Forward Error Correction. 25

FedRAMP Federal Risk and Authorization Management Program. 17

GDPR General Data Protection Regulation. 42

HA High Availability. 25

HIPAA Health Insurance Portability and Accountability Act. 17

IANA Internet Assigned Names Authority. 32

iBGP Internal BGP. 55

ICANN Internet Corporation for Assigned Names and Numbers. 32

ICMP Internet Control Messaging Protocol. 36

ICT Information and Communication Technology. 32

IGP Interior Gateway Protocol. 25

IMEI International Mobile Equipment Identity. 25

IoT Internet of Things. 62

IP Internet Protocol. 32, 42, 55

IPv4 Internet Protocol version 4. 32

IPv6 Internet Protocol version 6. 32

IRR Internet Routing Registry. 55

IS-IS Intermediate System to Intermediate System. 26

ISACA Information Systems Audit and Control Association. 17

ISMS Information Security Management System. 21, 54, 56

ISO International Standardisation Organization. 11, 37

ISP Internet Service Provider. 12

ITU International Telecom Union. 17

KPI Key Performance Indicator. 10, 21

LACP Link Aggregation Control Protocol. 25

MANRS Mutually Agreed Norms for Routing Security. 27, 36, 55

ML Machine Learning. 38

MNS Media Network Services. 35

MTBF Mean Time Between Failures. 10

MTTR Mean Time To Recover. 10

NCF National Critical Function. 8

NIST National Institute of Standards and Technology (USA). 17, 37

NOC Network Operations Centre. 43

OCTAVE Operationally Critical Threat, Asset and Vulnerability Evaluation. 18

OECD Organization for Economic Cooperation and Development. 27

OSI Open Systems Integration. 11

OSPF Open Shortest Path First. 25

OSSF Online Services Security Framework. 19

PCI-SSC Payment Card Industry Security Standards Council. 17

PI Personally Identifiable. 42

QoS Quality of Service. 9

RAID Redundant Array of Inexpensive Disks. 25

RFC Request For Comments. 41

RIPE Reseaux IP Europeens. 56

ROA Route Origin Authorization. 56

RoSI Return on Security Investments. 38

RPKI Resource Public Key Infrastructure. 36, 55

RS Risk Score. 38

SNMP Simple Network Management Protocol. 44, 54

SOC2 Systems and Organization Controls 2. 37

STP Spanning Tree Protocol. 25

SVM Support-Vector Model. 38

TARA Threat Agent Risk Assessment. 19

TTR Time To Recover. 9

VPN Virtual Private Network. 57

VS Vulnerability Score. 38

WAN Wide Area Network. 26, 54

Appendix A

Paper I

Jan Marius Evang et al. “Crosslayer Network Outage Classification Using Machine Learning”. In: **Proceedings of the Applied Networking Research Workshop.** ANRW ’22. Philadelphia, PA, USA: Association for Computing Machinery, 2022, pp. 1–7. isbn: 9781450394444. doi: 10.1145/3547115.3547193. url: `https://doi.org/10.1145/3547115.3547193`.

Status: Published

Crosslayer Network Outage Classification Using Machine Learning

Jan Marius Evang
marius@simula.no
SimulaMet, OsloMet
Oslo, Norway

Azza H. Ahmed
azza@simula.no
SimulaMet, OsloMet
Oslo, Norway

Ahmed Elmokashfi
ahmed@simula.no
SimulaMet
Oslo, Norway

Haakon Bryhni
haakonbryhni@simula.no
SimulaMet
Oslo, Norway

ABSTRACT

Network failures are common, difficult to troubleshoot, and small operators with limited resources need better tools for troubleshooting. In this paper, we analyse two years of outages from a small global network for high-quality services. Then, we develop a machine learning model for outage classification that can be set up with little effort and low risk. We use passive Bidirectional Forwarding Detection (BFD) data to classify Layer2 problems and add active packet loss data to classify other problems. The Layer2 problems were classified with a 99% accuracy and the other problems with 40%–100% accuracy. This is a significant improvement when we observe that only 35% of the customer cases we studied received any Reason for Outage (RFO) response from the Customer Support Centre.

CCS CONCEPTS

• **Networks → Public Internet**; **Network measurement**; • **Computing methodologies → Supervised learning**.

ACM Reference Format:
Jan Marius Evang, Azza H. Ahmed, Ahmed Elmokashfi, and Haakon Bryhni. 2022. Crosslayer Network Outage Classification Using Machine Learning. In *Applied Networking Research Workshop (ANRW '22), July 25–29, 2022, PHILADELPHIA, PA, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3547115.3547193

1 INTRODUCTION

Today's Internet comprises a group of small, medium, large and extra large networks as far as geographic presence and traffic volume are concerned. The end-to-end network service is produced following a three-layer model that is similar to the lower levels of the OSI reference model [1].

Table 1: Manually classified causes.

Class	Cause	Count
MultiLoss	Multiple Layer2 providers	870
CogentLoss	Cogent's network	556
Customer	Customer's equipment	411
TeliaLoss	Telia's network	316
Layer3	Layer3 only	222
InternMaint	Internal maintenance	114
Optic	3dB Optical change	74
ProvMaint	Provider maintenance	70
EquinixLoss	Equinix Cloud Fabric	58
SubseaCable	Subsea cable outages	42
EquipFail	Equipment failure	40
FiberCut	Fiber cut in provider network	39
Layer1	Leased Layer1 lines	18
Metro	Metropolitan area links	18
DoS	Denial of Service attacks	4

A few large providers sell Layer2 capacity based on the global mesh of Layer1 optical fibres, which are used by Layer3 providers to compose end to end services.

This layered architecture is exposed to various types of faults, such as physical fiber faults, equipment faults, planned maintenance and malicious attacks. Our data shows that the Layer1/Layer2 service has a high number of faults (see Table 1). Smaller networks that lease Layer1/Layer2 services need to quickly attribute such faults and report them to the respective providers. This is important for two reasons. First, it can help shorten the resolution time. Second, faults must be reported during the incident to be acknowledged according to the Service Level Agreements (SLAs).

Unlike large networks with sizable organization and abundant resources, small and medium network operators have a much smaller Network Operations Centre (NOC) with limited resources and staff. A typical small-medium NOC either operates a single enterprise network or is a speciality Internet Service Provider (ISP) providing a service to select customers in a narrow business area or in a geographic area.

Smaller NOCs often have a small but highly demanding customer base, for instance their co-workers in an enterprise, people in their

ANRW '22, July 25–29, 2022, PHILADELPHIA, PA, USA
© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-9444-4/22/07…$15.00
https://doi.org/10.1145/3547115.3547193

own geographic area or specialized service providers. This makes detecting and isolating faults very important yet a demanding task.

The NOC usually has automatic network monitoring systems in operation, but they can suffer from large numbers of both false positives (alerts without a real fault) and false negatives (faults that do not generate an alarm). This often causes true positives to be overlooked [2]. In an outage event where one component in the network has failed, causing interruption to network traffic, an overwhelming amount of log messages and alerts will be arriving from different monitoring systems. This makes the NOC waste time and effort to find the real cause. In other cases, a problem may not be noticed until customers complain. Customer Support (CS), may not have enough information to respond to a customer case because the NOC is busy troubleshooting. Alarm Consolidation systems exist but they suffer from high complexity [3], narrow field [4] or high compute requirements [5]. In this work, we tackle these problems by developing a generic model to assist NOCs and CSes. We leverage supervised learning to assist in classifying different outages. For classification, we use the Support-Vector Machine model (SVM) [6]. Our system is two-stage. In the first, it discriminates Layer1/Layer2 problems from Layer3 ones. Here, we identify a set of easily to collect metrics that can help achieving this in an efficient manner. In the second, it classifies Layer3 problems based on their root causes.

The research in [7] claims that supervised learning for fault classification is often suffering from low quality of training data, but in our research we have access to precise outage data, including root cause data.

Our system requires minimal changes to the network, and has a minimal impact on networking equipment and computing power. We also demonstrate that our proposed system is implementable and can be used to assist an existing provider efficiently.

With the system developed here, the NOC will speed up troubleshooting, quickly create trouble tickets with the providers, and the CS will improve customer satisfaction by giving informed feedback to all customer support cases. Compared to similar systems such as [5], investments in time and equipment are small, changes to configuration is minimal, and causes are successfully predicted with an f1-score of 0.99 for Layer2 cases and f1-score of 0.66 for other cases (see Section 4.1). Without the tool, only 35% of the cases received any outage report from CS.

2 RELATED WORK

Various works have used machine learning and other statistical methods for attributing faults for specific network protocols, however, there is still lack of work that leverages logs from different layers, and predict causes across network layers.

Existing research such as [3] implements a complex system of user defined scenarios, while they do not require detailed knowledge of the underlying system, they cannot detect problems outside manually defined failure scenarios. Our labeled data and two-stage approach makes classification of known faults across all layers possible and efficient, and the feedback loop handles new fault classes. Moreover, several projects [4, 8–12] examine how very detailed measurements of optical signal strength can be used to gain knowledge about the underlying Layer1 links. However, these methods require measurement of q-factor [13] telemetry [14] which is unavailable to higher layers providers.

The research in [15] also uses customer tickets for anomaly detection, but focuses only on Layer1 and last mile. The authors in [16] analyse Layers 2-7, while we analyse backbone Layers 1-3, and a common "No issue" class for any problems in other layers or outside the backbone network. Unlike [17], our system does not need any knowledge about the underlying network, only the manual feedback needs this.

Some commercial service providers have implemented systems for anomaly detection in system logs, for instance [18]. These systems have the advantage that they analyse the existing logs, and therefore are easy to start using, however, there is a high risk of exposing confidential information to a third party. In our system, only the feedback loop will have any confidentiality risk.

Finally, the authors in [5] analyse traffic by using a distributed Apache Storm [19] system in combination with data obtained from the Netflow [20] protocol. This puts extra stress on the networking equipment [21] and demands much more storage and CPU power, making it undesirable unless Netflow is already used for other purposes. BFD, on the other hand, is usually implemented in hardware.

The objective in this paper is to fill the gap and use simple data logs from various layers together with customer support data to classify outages in a fast and easily-implementable low-impact solution.

3 METHODOLOGY

3.1 Description of system

Our system consists of a data collection unit, a classification model (See Section 3.4), and alert and feedback units as shown in Figure 1.

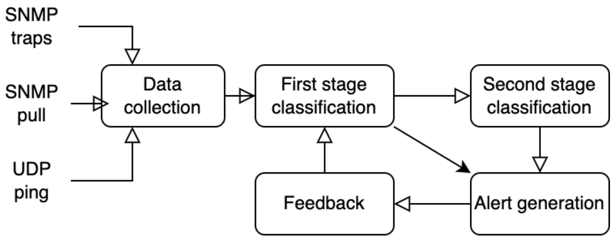

Figure 1: Our proposed outage classification system.

We collected the measurements from a global network covering 12 cities around the world, which we depict in Figure 2. The network uses three different Layer2 service providers to interconnect its points of presence (PoPs). The first is Telia VPLS, which is full-mesh Layer2 switched network based on VPLS/ELAN [22] over their global backbone network. The VPLS service supports Q-in-Q switching [23], so individual point-to-point VLANs [24] are configured, with each VLAN having member ports from only two cities. The second is Cogent L2C, which is a point-to-point MPLS [25] based service where multiple point-to-point Layer2 links are provided over the same physical interface. The third is Equinix Cloud Exchange Fabric (ECXF), which is a service of multiple point-to-point Layer2 links over the same physical interface.

The network is set up with IS-IS + BFD [26, 27] as Interior Gateway Protocol (IGP). The IGP makes sure that in case of issues on one link or device, customer traffic is automatically re-routed to an alternative path.

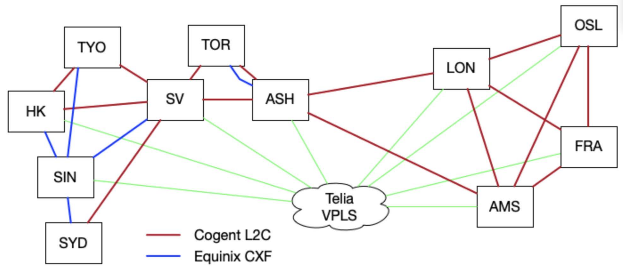

Figure 2: Network design.

3.2 Data description

In this work, we collected data from the monitored ISP over two years (2019-08-08 to 2021-10-01). Below, we list these measurements alongside their description.

Optical signal strength measurements. Every 30 minutes, optical received signal strength for the local link to the provider was read via SNMP [27]. There were 12 total outages on 5 interfaces, 88 drops in optical strength of more than 3dB on 14 interfaces and 53 increases of more than 3dB on 10 interfaces.

Interface error counters. Every 10 minutes, interface error counters for all devices were logged by SNMP polling. No interface errors were recorded for Layer2 switch ports, because any such errors would have been revealed and corrected during pre-production testing.

Buffer overflows/Tail-drops. Every 10 minutes, the buffer overflow/tail-drop counters were logged by SNMP polling. Only two interfaces showed tail drops, altogether 40 incidents. The NOC had especially amended this risk by over-provisioned the network to handle network traffic peaks without packet loss.

Layer2 packet loss data. In each of the 11 cities shown in Figure 2, 2 probe Virtual Machines (VMs) were set up. Our probe software is based on OpenNetNorad [28], which we rewrote in C to improve performance and reduce CPU consumption. This "pinger" transmits 100 UDP 64 byte packets every 0.5 seconds and waits for responses, and the "ponger" immediately returns any received packets to the sending IP address. The number of lost packets is then recorded. Packets are transmitted from 100 different UDP ports to detect any issues related to link aggregation or Equal Cost Multipath (ECMP) within the provider network. In addition to the probes measuring point-to-point (P2P) loss over the Layer2 links, a full mesh of probes (FM) were set up to measure the Layer3 service. One or more lost UDP packets in a 0.5sec interval generates one loss report. There were 196 million loss reports for 36 different pairs of probe VMs. These were pre-processed to 717352 unique events (see Section 3.3).

BFD traps. For each point-to-point link or VLAN, BFD (Bidirectional Forwarding Detection) [29] is configured to send one packet every 100ms. If 3 packets in a row are lost, the link is declared down and an SNMP trap message is sent to a collector. SNMP trap data is passively collected and stored in a database for later processing. The IS-IS protocol also receives BFD events and takes care of re-routing traffic.

Software crash logs. There were 62 instances of software crash/core-dumps incidents on the routers and switches. Most of these did not cause any interruption to network traffic since the Forwarding Engines were still operational.

Configuration change logs. Configuration change logs indicate which piece of equipment was configured and when. Also a textual description of the work was performed.

Customer complaints data. The customers' systems have strict network requirements for latency, packet loss and jitter. Customer cases were raised upon any violation of these requirements. The data was anonymized and made available for this work. During the period, there were 19399 customer cases, of which 8120 were related to the network. The complaints were reduced to 2855 unique cases on 21 different paths.

Customer service response data. For each customer case, CS analysed logs and provided a Reason For Outage (RFO) if possible. Out of 2855 cases, 1014 (35%) received RFO from CS, 109 of these were "no issue found".

Manual analysis of customer reports. We looked at all available data for each customer reported case and determined the reason for the incident. In most cases the cause was in a Layer2 provider's network. For other cases the cause could be determined more precisely from CS responses. The results are presented in Table 1.

In some cases, there were losses in multiple providers at the same time, which may be caused by either an (undetected) failure in the monitored network, a larger failure that impacted multiple providers, short traffic peaks that caused packet loss and therefore triggered a re-routing to another provider and subsequent loss there, or could be just a coincidence.

Multiple customer complaints received within a 5-minute interval were counted as one case. Still, a single root cause could cause multiple cases over a longer time. Some incidents were caused by planned or unplanned maintenance. These were recorded as cases, if they caused customer complaints even when the customer had been informed ahead of time.

The 713857 events that did not correspond to customer cases were not manually analysed.

3.3 Data preprocessing

The data used for the Machine Learning algorithm was BFD SNMP events (BFD), point-to-point UDP pings (P2P) and full-mesh UDP pings (FM). The other data was used only in the manual classification process of all the cases. The result of the manual classification was used to train the supervised machine learning system.

Due to small delays in detection and collection of test data, the resolution of the timestamps had to be reduced to match events from different sources. Each measuring point was added as a separate feature, with an aggregation of the number of such events per minute. One minute aggregation was chosen as a trade-off between fast detection and data size. For the BFD and P2P data, the measuring points were each link, for the FM data, the measuring points were the unique pairs of PoPs. This resulted in a dataset of 717352 unique events and 2855 unique cases. The features were 47 BFD, 32 P2P and 125 FM. The classes with < 4 cases were omitted.

3.4 Model description

We tested both Multilayer Perception neural networks (MLP) and SVM. SVM had both shortest processing time and highest classification accuracy, and is used in this paper. The data was split 75:25 into a training dataset and a testing dataset, and we tuned the hyperparameters using grid search. The optimal kernel was the Radial Basis Function (RBF) kernel with $C = 150$ and $\gamma = 7.5 \times 10^{-5}$. SVM is in general resistant to overfitting and we verified this by ShuffleSplit [30] and saw that the f1-score remained the same.

The first stage classification used only BFD data for classifying the largest and most precisely defined classes, i.e. the Layer2 provider cases. The output was five classes. One per each Layer2 provider, A fourth class that involve cases where more than one Layer2 provider, and one "Layer3" class for cases which were not caused by Layer2 events. A large number of events were processed in the first stage, but since fewer features were used, processing requirements were greatly reduced. The second stage classification used BFD, P2P and FM data for the Layer3 class to give an indication of the root cause. Since a much smaller subset of events was processed in this stage, the addition of more features did not lead to a large increase in processing power requirement. See also Section 4.4.

The feedback loop is used by NOC/CS when a prediction has failed, to manually correct the case label in the data and re-train the model.

After training the two machine learning models on the case data, the trained models were applied to all events, to see what knowledge could be gained.

4 PERFORMANCE EVALUATION

4.1 Evaluation metrics

We used the precision, recall and f1-score to assess our classifier.

For each class, the precision is the number of correctly predicted cases divided by the total predictions in that class. Recall is the number of correctly predicted cases divided by the number of true cases in that class. F1-score is the harmonic mean of precision and recall [31].

To visually evaluate the output of the classification process, we plot the Confusion Matrices. These show how well the model was able to assign a correct "predicted label" to each class of "true labels". The diagonals of the matrices show the correct predictions.

4.2 Accuracy and Feature importance

We performed the first classification stage initially by including all features, which resulted in a precision of 0.89, a recall of 0.89 and an f1-score of 0.92 (see the confusion matrix is in Figure 3a).

Using only BFD features showed much better scores for Layer2 cases, but did (as expected) not distinguish between Layer3 and Customer issues as seen in the confusion matrix in Figure 3b and scores in Table 2. Total f1-score was now 0.99 with a combined Customer+Layer3 class . Further, repeating the first stage while including only FM and only P2P gave poor results with f1-score 0.35 for FM and and f1-score of 0.26 for P2P (see Figures 3c and 3d).

Table 2: First stage evaluation, based on BFD.

class	precision	recall	f1-score
CogentLoss	1.00	0.99	0.99
TeliaLoss	1.00	1.00	1.00
MultiLoss	0.99	0.99	0.99
EquinixLoss	0.88	1.00	0.94
Customer	0.65	1.00	0.79
Layer3	0.00	0.00	0.00

Table 3: Second stage prediction scores (Based on BFD+P2P+FM)

class	precision	recall	f1-score
InternMaint	0.65	0.72	0.68
Optic	0.75	0.43	0.55
ProvMaint	0.33	0.55	0.41
SubseaCable	0.92	0.92	0.92
EquipFail	0.40	0.40	0.40
FiberCut	0.75	0.64	0.69
Layer1	0.83	1.00	0.91
Metro	1.00	0.43	0.60
DoS	1.00	1.00	1.00

The BFD analysis contained only 4 misclassifications: 2 MultiLoss events classified as EquinixLoss were caused by two unrelated coinciding loss events where the EquinixLoss event affected multiple Equinix links, and 2 CogentLoss events classified as MultiLoss were multiple coinciding Cogent events. The analysis including all features added the capability of distinguishing between Layer3 loss and Customer loss, at the expense of requiring more computing time and adding more "noise" to the various Layer2-classifications. Still, we see a relatively small number of misclassifications (14 misclassified and 429 correctly classified Layer2 events).

For the second stage, the events that were identified by the first stage classification were removed, and a new supervised classification was attempted for the remaining events. After hyperparameter tuning, this classification showed an f1-score of 0.66. The size of the dataset in this analysis is only 437 cases with 204 features, and the results were not as good as for the first stage, but a reasonable suggestion for a root cause might still provide valuable input to the NOC's troubleshooting process. Figure 4 and Table 3 show the confusion matrix and classification score for stage 2, respectively. We can clearly see that determining the exact root cause can be hard for a few types of failures. For instance, ProvMaint and InternalMaint events may cause a wide variety of different error symptoms, that may be indistinguishable from the other classes. Interestingly, subsea cable cuts (f1-score 0.92) and fiber cuts (f1-score 0.69) had relatively good classification scores, even though these were thought to be difficult to distinguish. A point for future study might be to understand why.

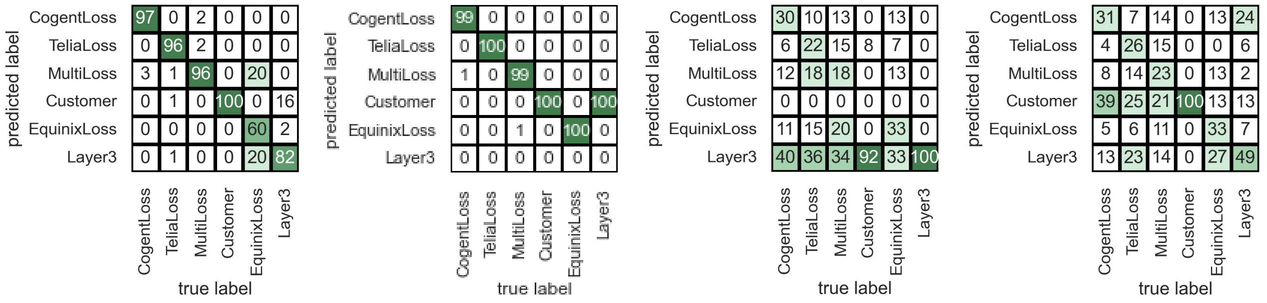

(a) All features **(b) BFD features only** **(c) FM features only** **(d) P2P features only**

Figure 3: First stage classifications

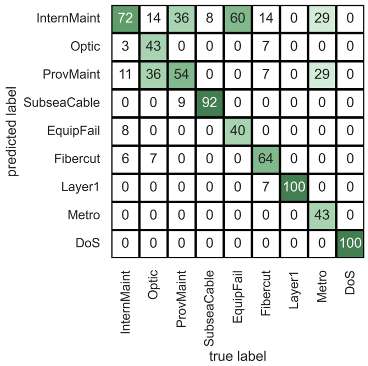

Figure 4: Confusion matrix for second stage classification

Table 4: First stage data of the Layer2 cases and predictions

class	support cases	extrapolated cases
CogentLoss	556 (19.4%)	49997 (7.0%)
TeliaLoss	316 (11.1%)	64859 (9.1%)
MultiLoss	870 (30.5%)	47227 (6.6%)
EquinixLoss	58 (2.0%)	15346 (2.1%)
Customer+Layer3	633 (22.2%)	536428 (75.1%)
Other	14.8%	

4.3 Extrapolation

Using the first stage model, BFD-trained on the cases with well known cause and symptoms, we ran a prediction on all the events where we did not get any customer complaints, to get an idea of how common the various types of problems are in these events. The very high f1-score of the model fitted on the complaint-data means that the predictions on the non-complaint-data will be highly relevant for our research. However, selection bias in that some hidden class of outages never leads to complaints might reduce the accuracy of the extrapolation.

For the first stage model, the results can be seen in Table 4. The most interesting observation is that the "Customer+Layer3" classification is much more common than in the cases where the customers filed complaints. (75.1% of the events, versus 22.2% of the cases). This means that the test network does a good job of hiding Layer3 problems from customers, and Layer2 problems are more likely to cause customer complaints, but still only 0.4% of all events caused customer cases.

Table 5: Second stage data for the Layer3 cases and predictions

class	cases	predictions
InternMaint	114 (27.1%)	185169 (34.5%)
Optic	74 (17.6%)	216591 (40.4%)
ProvMaint	70 (16.7%)	26883 (5.0%)
SubseaCable	42 (10.0%)	19522 (3.6%)
EquipFail	40 (9.5%)	18308 (3.4%)
FiberCut	39 (9.3%)	58099 (10.8%)
Layer1	18 (4.3%)	10450 (1.9%)
Metro	19 (4.5%)	1329 (0.2%)
DoS	4 (1.0%)	77(0.01%)

The "MultiLoss" class is only 6.6% of the events in the non-complaint dataset, vs 30.5% of the complaint-cases. This indicates that the network is better at hiding Layer2 problems in a single provider, and problems affecting multiple providers are more likely to generate customer complaints.

Further, we used the model fitted on the Second stage data from the cases, and made a prediction using only the "Customer+Layer3" class from the first stage non-complaint events. Applying the second stage model to the non-complaint data gives an indication that Internal maintenance and Optic events are less likely to cause customer complaints than the other classes, but the size of the dataset and the lower accuracy of the model makes these results much less certain. See Table 5.

4.4 Processing performance

BFD is implemented in hardware on our routers and do not put any load on the routers' CPU. To compare, Netflow would cause 15%-20% CPU impact according to [5], which matches our own experience. SNMP traps produced by the routers using the lowest priority processes, and all data is transmitted blindly using UDP, also reducing processing. Our ML processing on an M1 Pro 10 core

CPU took <1sec. The amount of stored data for the ML system is low. For each BFD trap we store timestamp+link-id and for each UDP measurement we store timestamp, source/destination address and loss percentage.

5 DISCUSSION

Our analysis of two years of outage data shows that a two-stage classification system is well suited to classify network outages, providing the NOC with useful predictions on where to start troubleshooting, and providing CS with RFO for all cases with a much better success rate than the observed 35% of CS responses during the period of the study. BFD data exhibits their high importance in the classification. In contrast, although the active P2P and FM raw data provides very precise measurements, they are not highly contributing to discriminating features in the classification model.

One important shortcoming is that we do not have latency measurements. But as our analysis reveals, BFD SNMP traps are very good indicator of problem types and location, so latency changes would probably not have a great impact on this result. Moreover, in this work, customer complaints are the only source for determining whether a packet loss event is regarded as an outage. Only the cases that are received as customer complaints are analysed in detail. This means that some outages may be overlooked if the customer did not complain, and some complaints may be groundless (i.e caused by other factors than the test network). A customer complaint is only counted as a network outage if the timestamp is reported as within 60 seconds of an internal packet loss or BFD trap event.

There are many features that show some correlation, which might disturb the machine learning classification model since one event is likely to affect multiple features. But since the features have a large geographic spread, and since there are many features, a certain degree of correlation should not cause problems for our analysis.

Model Drift (MD) is another consideration. During the 2 years of data collection, there were continuous changes to both the network topology, the routing protocols and the customer's monitoring system. MD may have degraded our analysis, in that patterns for the various classes of events change over time. However, this will also reflect more accurately a real-life situation. The results prove that our first stage analysis was not significantly affected by MD. In the future, we plan to gain insight into how our model may degrade, for instance by temporal cross validation, and how to rectify it through a system for retraining while running in the production. A future improvement, especially for the second stage, would be to also report the second ranked classification for an outage.

Another very important practical consideration is the difference in complexity of gathering the data for the first stage and the second stage. The passive BFD data used for the first stage is very easy to collect. Most networks already use the BFD protocol as a part of the IGP protocol, but very few actually gather the SNMP trap data from BFD. The changes and risk to the network will be small, the BFD events are already being detected, so the only change is to generate SNMP trap messages and set up one central location to store these (optionally a second location for redundancy.)

The active P2P and FM raw data are very accurate and provides very precise measurements, but showed less precision in case classification, and the system used to gather this data is much more expensive in management and computing power.

6 CONCLUSION

We have developed a system that Network Operations Centres and Support Centres for smaller operators can use in a failure situation. Using minimal resources, we passively collect BFD data and classify the Layer2 events to an f1-score of 0.99. By adding a second stage with active monitoring to collect UDP ping data we predict other types of root cases with a 0.66 f1-score. Our analysis interestingly shows that BFD features, which are the easiest to collect, give the best results for outage classification.

REFERENCES

[1] J. Day and H. Zimmermann, "The osi reference model," *Proceedings of the IEEE*, vol. 71, no. 12, pp. 1334–1340, 1983.

[2] B. AlAhmadi, L. Axon, and I. Martinovic, "99% false positives: A qualitative study of soc analysts' perspectives on security alarms," *USENIX Association*, 2021. [Online]. Available: https://ora.ox.ac.uk/objects/uuid:0be05f6b-7470-4210-acb6-2018d5dc6ca0

[3] K. Appleby, G. S. Goldszmidt, and M. Steinder, "Yemanja—a layered fault localization system for multi-domain computing utilities," *Journal of Network and Systems Management*, vol. 10, pp. 171–194, 2004.

[4] T. Christopoulos, O. Tsilipakos, G. Sinatkas, and E. E. Kriezis, "On the calculation of the quality factor in contemporary photonic resonant structures," *Opt. Express*, vol. 27, no. 10, pp. 14 505–14 522, May 2019. [Online]. Available: http://opg.optica.org/oe/abstract.cfm?URI=oe-27-10-14505

[5] Y. Du, J. Liu, F. Liu, and L. Chen, "A real-time anomalies detection system based on streaming technology," in *2014 Sixth International Conference on Intelligent Human-Machine Systems and Cybernetics*, vol. 2, 2014, pp. 275–279.

[6] R. Soentpiet *et al.*, *Advances in kernel methods: support vector learning.* MIT press, 1999.

[7] S. Ayoubi, N. Limam, M. A. Salahuddin, N. Shahriar, R. Boutaba, F. Estrada-Solano, and O. M. Caicedo, "Machine learning for cognitive network management," *IEEE Communications Magazine*, vol. 56, no. 1, pp. 158–165, 2018.

[8] C. Natalino, A. di Giglio, M. Schiano, and M. Furdek, "Root cause analysis for autonomous optical networks: A physical layer security use case," in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–4.

[9] L. Shu, Z. Yu, Z. Wan, J. Zhang, S. Hu, and K. Xu, "Low-complexity dual-stage soft failure detection by exploiting digital spectrum information," in *45th European Conference on Optical Communication (ECOC 2019)*, 2019, pp. 1–4.

[10] C. Delezoide, P. Ramantanis, L. Gifre, F. Boitier, and P. Layec, "Field trial of failure localization in a backbone optical network," in *2021 European Conference on Optical Communication (ECOC)*, 2021, pp. 1–4.

[11] Ujjwal, J. Thangaraj, and A. A. Dias Barreto, "Accurate qot estimation for the optimized design of optical transport network based on advanced deep learning model," *Optical Fiber Technology*, vol. 70, p. 102895, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1068520022000785

[12] C. Miao, M. Chen, A. Gupta, Z. Meng, L. Ye, J. Xiao, J. Chen, Z. He, X. Luo, J. Wang, and H. Yu, "Detecting ephemeral optical events with OpTel," in *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22).* Renton, WA: USENIX Association, Apr. 2022, pp. 339–353. [Online]. Available: https://www.usenix.org/conference/nsdi22/presentation/miao

[13] "Recommendation O.201: Q-factor test equipment to estimate the transmission performance of optical channels," International Organization for Standardization, Geneva, CH, Standard, 2003.

[14] [Online]. Available: https://www.juniper.net/documentation/us/en/software/junos/interfaces-telemetry/index.html

[15] J. Hu, Z. Zhou, and X. Yang, "Characterizing Physical-Layer transmission errors in cable broadband networks," in *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22).* Renton, WA: USENIX Association, Apr. 2022, pp. 845–859. [Online]. Available: https://www.usenix.org/conference/nsdi22/presentation/hu

[16] J. Iurman, F. Brockners, and B. Donnet, "Towards cross-layer telemetry," in *Proceedings of the Applied Networking Research Workshop*, ser. ANRW '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 15–21. [Online]. Available: https://doi.org/10.1145/3472305.3472313

[17] R. R. Kompella, J. Yates, A. Greenberg, and A. C. Snoeren, "IP fault localization via risk modeling," in *2nd Symposium on Networked Systems*

Design & Implementation (NSDI 05). Boston, MA: USENIX Association, May 2005. [Online]. Available: https://www.usenix.org/conference/nsdi-05/ip-fault-localization-risk-modeling

[18] [Online]. Available: zerbium.com

[19] The Apache Software Foundation, "Apache storm," https://storm.apache.org/.

[20] E. B. Claise, "Cisco systems netflow services export version 9," Internet Requests for Comments, RFC Editor, RFC 3954, 8 2004, http://www.rfc-editor.org/rfc/rfc3954.txt. [Online]. Available: http://www.rfc-editor.org/rfc/rfc3954.txt

[21] "Netflow services," p. 74, 2003.

[22] Metro Ethernet Forum, "Ethernet services definitions - phase 2," 4 2008.

[23] "Provider bridges, ieee std. 802.1ad," 2005.

[24] "Bridges and bridged networks, ieee std. 802.1q-2018," 2016.

[25] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol label switching architecture," Internet Requests for Comments, RFC Editor, RFC 3031, 1 2001, http://www.rfc-editor.org/rfc/rfc3031.txt. [Online]. Available: http://www.rfc-editor.org/rfc/rfc3031.txt

[26] International Organization for Standardization, *Information technology — Telecommunications and information exchange between systems — Intermediate System to Intermediate System intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service*, ISO/IEC10589:2002 ed. Vernier, Geneva, Switzerland: International Organization for Standardization, 2015. [Online]. Available: https://www.iso.org/standard/30932.html

[27] D. Katz and D. Ward, "Bidirectional forwarding detection (bfd) for ipv4 and ipv6 (single hop)," Internet Requests for Comments, RFC Editor, RFC 5881, 6 2010.

[28] Facebook Inc, "Opennetnorad," https://github.com/fbsamples/OpenNetNorad, 2017.

[29] R. Presuhn, "Version 2 of the protocol operations for the simple network management protocol (snmp)," Internet Requests for Comments, RFC Editor, STD 62, 12 2002, http://www.rfc-editor.org/rfc/rfc3416.txt. [Online]. Available: http://www.rfc-editor.org/rfc/rfc3416.txt

[30] Q.-S. Xu and Y.-Z. Liang, "Monte carlo cross validation," *Chemometrics and Intelligent Laboratory Systems*, vol. 56, no. 1, pp. 1–11, 2001. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0169743900001222

[31] N. Chinchor, "MUC-4 Evaluation Metrics," in *Proceedings of the 4th Conference on Message Understanding*, ser. MUC4 '92. USA: Association for Computational Linguistics, 1992, p. 22–29. [Online]. Available: https://doi.org/10.3115/1072064.1072067

Appendix B

Paper II

J. M. Evang, “ISO27001 as a Tool for Availability Management”. ,In: **Proceedings of the International Workshop on Information management**. WSIM ’22. London, UK, 2022, pp. 82–85. doi: 10.1109/AEIS59450.2022.00018. url: `https://doi.org/10.1109/AEIS59450.2022.00018`.

Status: Published

[Article not attached due to copyright]

Appendix C

Paper III

Jan Marius Evang and Tarik Cicic. “Evolved Cold-Potato routing experiences”. In: **The 31st International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2023)**.

DOI: https://doi.org/10.23919/SoftCOM58365.2023.10271604

[Article not attached due to copyright]

Appendix D

Paper IV

Jan Marius Evang and Ioana Livadariu. “How Large Is the Gap? Exploring MANRS and ISO27001 Security Management”. In: **The 31st International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2023)**, to appear.

DOI: https://doi.org/10.23919/SoftCOM58365.2023.10271691

[Article not attached due to copyright]

Appendix E

Paper V

Jan Marius. Evang. “A 10-Layer Model for Service Availability Risk Management”. In: **Proceedings of the 20th International Conference on Security and Cryptography - SECRYPT**. INSTICC. SciTePress, 2023, pp. 716–723. isbn: 978-989-758-666-8. doi: 10.5220/0012092600003555. url: `https://doi.org/10.5220/0012092600003555`.

Status: Published

A 10-Layer Model for Service Availability Risk Management

Jan Marius Evang[1,2]
[1] *Oslo Metropolitan University, Oslo, Norway*
[2] *Simula Metropolitan Center for Digital Engineering, Oslo, Norway*
marius@simula.no

Keywords: Risk Assessment, Availability Management.

Abstract: Effective management of service availability risk is a critical aspect of Network Operations Centers (NOCs) as network uptime is a key performance indicator. However, commonly used risk classification systems such as ISO27001:2013, NIST CSF, and NIST 800-53 often do not prioritize network availability, resulting in the potential oversight of certain risks and ambiguous classifications. This paper presents a comprehensive examination of network availability risk and proposes a 10-layer model that aligns closely with the operational framework of NOCs. The 10-layer model encompasses hardware risk, risks across various network layers, as well as external risks such as cloud, human errors, and political governance. By adopting this model, critical risks are less likely to be overlooked, and the NOC's risk management process is streamlined. The paper outlines each layer of the model, provides illustrative examples of related risks and outages, and presents the successful evaluation of the model on two real-life networks, where all risks were identified and appropriately classified.

1 INTRODUCTION

From the advent of computer networks, disruptions to the network service have been a persistent challenge. A Network Operations Centre (NOC)'s most important goal has been to make these disruptions invisible to the end users, since they can lead to lost productivity, revenue, and erode customer trust. At all times, businesses have performed some form of risk management, whether formally or informally, and countless books have been written on the subject, to the point where an official standard was created with the 1st Edition of ISO31000 (ISO, 2018) in 2009.

A "top down" approach to risk *identification* is to conduct interviews with key stakeholders, based on one of the common security standard frameworks's classification system. This approach may be confusing and not optimal for a NOC team. Sometimes these categories are very generic, for instance the ISO27001:2013 (ISO, 2022a) standard has chapters like "Cryptography" and "Communications Security", and NIST CSF (Barrett, 2018) has "Protective Technology", while the updated ISO27001:2022 has only four themes of ”People”, ”Organizations”, ”Technology” and ”Physical”[1]. Furthermore, one network availability risk often spans multiple categories, for instance NIST800-53's (NIST, 2022) controls[2] of "Audit", "Security Assessment", "Contingency Planning", "Incident Response", "Media Protection", "Planning", "Performance Measurement", "System and Communication Protection", "System Integrity" and "Supplier Risk" have significant overlaps. We experimentally verify these in Section 3.

To address these challenges, this paper proposes a novel framework for the discovery and classification of availability risk in network services. Our model is based on the ISO/OSI 7-layer reference model (OSI model) (Zimmermann, 1980), which has proved to be a very suitable tool for dividing network functions into manageable compartments (See Figure 1). The OSI model is not perfect, but is used in some form in network courses, research and standardisation processes. The layers of the OSI model are well defined, common network protocols map reasonably well to the layers and the model is universally recognized in the networking business.

However, when it comes to network availability risk management, a different separation of layers is

[1] ISO27001:2022 may be an improvement over ISO27001:2013, but detailed risk discovery data based on this model was not available to us at the time of writing.

[2] Actions, devices, procedures, techniques, or other measures that reduce the vulnerability of an information system.

Evang, J.
A 10-Layer Model for Service Availability Risk Management.
In *Proceedings of the 20th International Conference on Security and Cryptography (SECRYPT 2023)*, pages 716-723
ISBN: 978-989-758-666-8; ISSN: 2184-7711

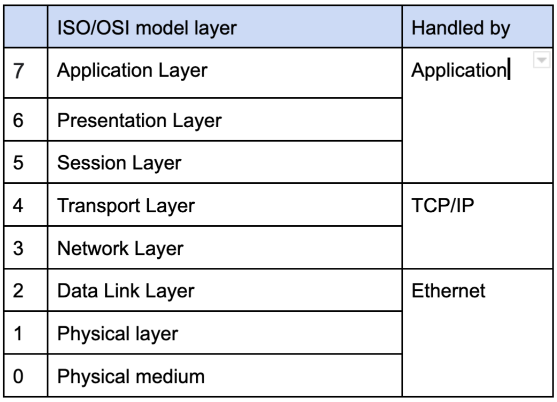

	ISO/OSI model layer	Handled by
7	Application Layer	Application
6	Presentation Layer	
5	Session Layer	
4	Transport Layer	TCP/IP
3	Network Layer	
2	Data Link Layer	Ethernet
1	Physical layer	
0	Physical medium	

Figure 1: The ISO/OSI model, as used in a typical network service.

suggested in this paper. Some risks lie outside the OSI model layers, and we slightly modify the layer division to better match the risks that a NOC needs to manage. Although the idea of additional layers beyond the 7-layer OSI model is not new, as seen in previous works like (Taylor and Wexler, 2003; Kachold, 2009), a comprehensive description of all the layers has not been published until now. In this paper, we use named layers to describe the new proposed layers, while numbered layers refer to the layers of the OSI model, to avoid confusion.

Information security is often classified into three main objects: confidentiality, integrity, and availability (Anderson, 1972). While confidentiality and integrity are typically addressed together, availability is often handled separately by a Network Operations Centre (NOC). This chapter focuses on the topic of availability and its importance for all types of NOCs, whether in-house or outsourced. In today's interconnected world, organizations heavily rely on information availability across various layers, encompassing customer interactions and service delivery. However, due to the multitude of risks involved, identifying and managing these risks can be challenging. To facilitate the risk identification process, common approaches involve grouping risks into manageable areas and analyzing them individually to gain a comprehensive overview. This paper aims to categorize and discuss risk topics associated with operating a network service, highlighting examples of availability breaches at each layer. Mitigation strategies within the same layer or across different layers are also presented. Please note that the references cited mostly refer to media coverage of outages, as detailed research on such incidents is seldom available, and the provided content may include speculations.

Risk is defined as the impact of uncertainty on objectives, and it is typically expressed in terms of the likelihood of an event occurring and its consequences or impact, which can be qualitative or quantitative. Numerically, we define the risk level as the product of likelihood and impact. The impact can be measured in various ways, such as packet loss, total downtime, or financial loss.

Every layer within the model poses its own set of risks, necessitating a holistic approach where the NOC considers all layers, quantifies associated risks, and determines appropriate mitigation actions. This comprehensive perspective is crucial for effective risk management and ensuring the availability of network services.

2 SECURITY LAYERS

The security topics in our proposed 10-layer model are defined with the service layer in the middle, where the total availability (uptime) is measured. Below the service layer, we have layers whose risks are predictable and directly affect service delivery, and where industry standards have emerged to handle these risks. Above the service layer, we find topics that indirectly and less predictably contribute to the availability risk of the NOC, like risks associated with human errors, company culture and legal responsibilities.

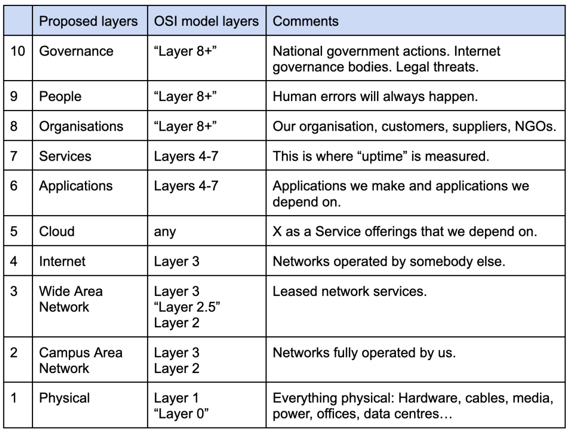

	Proposed layers	OSI model layers	Comments
10	Governance	"Layer 8+"	National government actions. Internet governance bodies. Legal threats.
9	People	"Layer 8+"	Human errors will always happen.
8	Organisations	"Layer 8+"	Our organisation, customers, suppliers, NGOs.
7	Services	Layers 4-7	This is where "uptime" is measured.
6	Applications	Layers 4-7	Applications we make and applications we depend on.
5	Cloud	any	X as a Service offerings that we depend on.
4	Internet	Layer 3	Networks operated by somebody else.
3	Wide Area Network	Layer 3 "Layer 2.5" Layer 2	Leased network services.
2	Campus Area Network	Layer 3 Layer 2	Networks fully operated by us.
1	Physical	Layer 1 "Layer 0"	Everything physical: Hardware, cables, media, power, offices, data centres…

Figure 2: The proposed 10-layer model for network service risk assessment.

2.1 Physical Layer

This category encompasses risks associated with physical hardware, including cables, networking equipment, server equipment, workstations, phones, and IoT devices. Outages at the physical layer can be caused by equipment defects, broken cables, planned maintenance activities, power failures, and physical security breaches. Controls for managing these outages can be found in ISO27002 Clause 11 (Physical & Environmental Security) (ISO, 2022b) and NIST800-53's PE controls.

Physical layer outages often have longer durations and may require on-site technician visits, resulting in extended Time To Recover (TTR). Therefore, it is crucial to mitigate these risks proactively. Duplication and clustering of networking and server hardware, along with redundant components such as power supplies and hard disks with automatic failover, can be implemented. Critical network links may require duplicate network cables and the use of network protocols to maintain service availability during Physical Layer failures. One particularly severe physical layer failure is a fire in a server room triggering a fire suppression system, potentially causing permanent equipment failure. Mitigating such an outage involves distributing the service across multiple geographic locations to ensure service continuity.

Selecting high-quality hardware and having hardware service agreements can enhance the likelihood of maintaining reliable physical layer operations. For low TTR requirements, keeping spare parts in-house can be considered, based on a Return on Investment assessment.

Risk discovery at the Physical Layer is relatively straightforward, as every physical asset can fail and should be included in the risk registry. Evaluating the likelihood of failures and implementing measures to reduce their impact are essential.

Examples of outages include the Jan 2020 earthquake in Puerto Rico (Santiago et al., 2020), which caused prolonged power outages and network faults, leading to significant internet disruptions. However, communications were still upheld through the resilient cellular network during these events (NETBLOCKS, 2020). As another example, multiple subsea cables following the same paths in the Suez Canal have posed increased risks of shared-fate problems, resulting in several outages (Burgess, 2022).

2.2 Local Network Layer

The local network refers to the network infrastructure within a building or campus, where the NOC owns and manages the hardware and cabling. This layer includes networks such as server-room networks, building cabling, office-space networks, as well as wireless networks like WiFi, cellular, and IoT.

Risks at the Local Network Layer primarily stem from firmware or configuration errors in network equipment, along with capacity issues like full disks, out-of-memory situations, and network capacity limitations. Monitoring and proactive planning are key measures to mitigate these risks. Additionally, this layer plays a crucial role in mitigating most of the risks originating from the Physical Layer by implementing local (network) protocols like RAID (Patterson et al., 1988), LACP (C/LM - LAN/MAN Standards Committee, 2000), VRRP (Hinden, 2004), High Availability protocols, and Interior Gateway Protocols (IGP) such as IS-IS (ISO, 2002) and OSPF (Moy, 1998).

Examples of outages include one of GitHub's major outages in December 2012, which occurred due to the failure of multi-chassis link aggregation protocols at the local network layer when a switch experienced partial malfunctioning (Imbriaco, 2012). Another significant outage took place in February 2020, where the RIPE RPKI repository experienced a three-day outage caused by a full disk quota, leading to the invalidation of all RIPE RPKI routes (Trenaman, 2020).

2.3 Wide Area Network Layer

The wide area network (WAN) encompasses networks that are logically part of the NOC's operations but physically leased from network service providers. These networks can include optical fibers, Layer 1 wavelengths, Layer 2/2.5 MPLS-like services (Viswanathan et al., 2001), or overlay networks like SD-WAN over a Layer 3 service. WANs typically span metropolitan, national, or international areas, and may also include in-building or space-based networks. Additionally, Layer 1/2 interconnections with remote customers and suppliers of network services are considered within this layer.

Wide area networks often experience full or partial outages, as documented in (Evang et al., 2022). These outages can have various root causes, including physical layer or local network layer events, congestion, or issues from other layers. However, the common symptoms are outages or packet loss. Mitigation strategies for network outages in WANs often involve duplicate links, redundancy protocols, MPLS, VXLAN (Mahalingam et al., 2014), BFD (Katz and Ward, 2010), and IGP protocols such as IS-IS and OSPF. However, the time taken for failover (TTR) is usually longer due to the distances involved, which may cause delays in protocol updates. Capacity risks are more significant in wide area networks since services are typically purchased based on capacity, and service providers may drop packets if the agreed traffic rate is exceeded. Mitigating this risk requires careful consideration, including over-purchasing of capacity, planning for backup links, assessing shared-fate risks of links, and potentially engaging multiple providers to safeguard against total provider failure.

Example of outage: In June 2022, simultaneous outages occurred in two major subsea cable systems, leading to congestion and packet loss for numerous wide area networks traversing the Suez Canal (Belson, 2022).

2.4 Internet Layer

The internet layer focuses on the risks associated with connectivity to external networks that are beyond the direct control of the NOC, where they best-case have a contractual agreement, and worst-case have no control whatsoever.

The predominant protocol at this layer is BGP (Rekhter et al., 2006), which encompasses IP transit, Internet Exchanges, private peering, and BGP customers. While BGP effectively navigates the intricate Internet landscape, it suffers from security limitations (Freedman et al., 2019). The protocol relies on trust and does not verify the validity of exchanged data, leading to significant confidentiality and availability risks as highlighted in the OECD Routing Security paper of 2022 (OECD, 2022). Efforts are underway to address these systemic flaws, with promising technologies like RPKI (Bush and Austein, 2013) employing cryptographic signatures to mitigate origin hijacking risks. Other initiatives such as BGPsec (Lepinski and Sriram, 2017) and SCION (Rustignoli and de Kater, 2022) tackle BGP path hijacking risks but encounter their own challenges (Durand, 2020).

Examples of outages: In February 2008, a Pakistani network operator mistakenly announced YouTube's IP addresses via BGP, resulting in a two-hour global service blackhole (Hunter, 2008). These announcements, intended for internal use only, were leaked to their upstream provider and subsequently propagated throughout the entire internet.

In June 2015, Telecom Malaysia leaked 179,000 prefixes to Level3, causing a significant volume of traffic to traverse Telecom Malaysia's backbone, leading to network overload, severe packet loss, and internet slowdown worldwide (Toonk, 2015).

The deficiencies in BGP have also been exploited maliciously. In August 2020, AS209243 announced the IP addresses of a critical smart contract user interface for the Celer Bridge cryptocurrency exchange. The attacker obtained authorized HTTPS certificates and reportedly stole a total of USD 234,866.65 worth of various cryptocurrencies (The SlowMist Security Team, 2022; Kacherginsky, 2022).

2.5 Cloud Layer

Today, numerous services are delivered through various cloud providers, ranging from on-premises solutions to Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). The level of risk varies depending on the extent of responsibility transferred from the NOC to the dedicated teams of the service providers. However, it's important to assess the Return on Security Investments (RoSI) considering the costs involved. ISO27017 (ISO, 2015) provides a specific code of practice for securing Cloud Services. Cloud-related risks also extend to supporting services such as email systems, documentation systems, and customer management systems.

During an outage at a major cloud provider, the impact can be severe, leaving the NOC with little to do but wait. To mitigate cloud risks, systems can be distributed across multiple cloud providers and failover protocols can be implemented.

Examples of outages: In December 2021, Amazon Web Services (AWS) experienced a significant outage in their IaaS service, causing disruptions to numerous dependent services (Goovaerts, 2021; AWS, 2021).

In October 2022, the Cloudflare Content Delivery Network (CDN) cloud service suffered an outage due to a software bug, resulting in a failure rate of around 5% for over six hours (Graham-Cumming, 2022).

2.6 Applications Layer

Application risks arise from both internally developed applications and those developed by third parties. To mitigate risks associated with third-party applications, thorough sandbox testing and duplication strategies are employed for critical services.

Ensuring well-written applications with minimal software errors and effective error handling is crucial for reducing availability risks. While confidentiality risks are beyond the scope of this document, it's worth noting that breaches in confidentiality can also impact availability. ISO27002's Clause 14 provides recommended controls for secure application development and service protection.

Application-based redundancy can be implemented to safeguard the service from significant outages at lower layers. In such cases, if the primary backend service fails, the application can utilize a secondary backend service.

An example of an application causing availability issues is the Facebook outage in 2021, which resulted from a software bug and potentially led to significant revenue losses in the tens of millions (Integrated Human Factors, 2022).

2.7 Services Layer

The services provided by the organization are what customers ultimately experience. These services depend on all underlying layers and may also depend on purchased services. Mitigation measures are implemented at lower layers to minimize service outages.

Customer contracts often include Service Level Agreements (SLAs) that define expected availability. If the sold service has a better SLA than the purchased service, risk mitigation is necessary. SLA levels can vary widely, ranging from 99% to 99.999% uptime per year. SLAs are addressed in ISO27002's Clause 18.

While planned maintenance is typically exempt from SLA contracts, it still impacts availability and requires mitigation. Risks may also arise from failures of subcontracted supporting services, such as payment services. Using redundant services can reduce risk but increases costs.

The root DNS service exemplifies a highly critical service with a resilient design. It is distributed across independent servers, avoiding dependency on any single entity. Even during heavy DDoS attacks (ICANN, 2007), the DNS service remained robust and did not significantly disrupt internet traffic.

An undisclosed root cause led to the September 2022 Zoom outage, causing the Video Conferencing service to be unavailable and resulting in numerous failed video meetings (Goyal, 2022; Silberling, 2022; Zoom, 2022).

2.8 Organizations Layer

The quality of service delivery relies heavily on the organization itself. A positive company culture, strong policies, and employees who adhere to those policies can significantly reduce human errors.

Implementing a robust Information Security Management System (ISMS) with comprehensive risk policies and effective mitigation measures is essential. Considering the culture, policies, and certifications of providers and peers is also important, as customers may require adherence to standards like ISO27001 or NIST800-53.

Furthermore, organizations may have dependencies on overarching entities such as trade unions, employer organizations, industry associations, and Regional Internet Registries and network operators' groups.

Examples of outages include a 10-day IT outage in July 2022 at the UK's largest hospital, attributed to a lack of attention to IT security in the company culture (Thimbleby, 2022). Another instance was nationwide internet shutdowns in Lebanon in 2022 due to a strike by employees of the state-owned telco, Ogero (Barton, 2022).

2.9 People Layer

Human errors are inevitable, and a NOC must take measures to protect the service against common mistakes. Implementing effective procedures and reducing stress can help mitigate this risk. It is also important to address the risk of disloyal employees through compartmentalization, need-based access rights, and a strong Human Resources team.

Other people-related risks include the impact of sick leave and employee departures, which can lead to knowledge loss and potential exposure to competitors or attackers. Documentation plays a crucial role in mitigating these risks, ensuring that no individual possesses irreplaceable knowledge within the company.

Numerous significant outages in the internet world have been caused by human errors that went undetected by control systems. Examples include the June 2022 Cloudflare outage (Belson, 2022), the October 2021 outage affecting Facebook, WhatsApp, and Instagram (Integrated Human Factors, 2022), and the February 2017 AWS outage (AWS, 2017).

2.10 Governance Layer

The risk of breaking local regulations or national laws is most often associated with Confidentiality and Integrity, but the punishments may be severe and even cause availability outages, for instance if a court orders the temporary or permanent shutdown of a service. The financial impact of a breach of contract or breach of regulations, or even a customer boycott must also be considered, as this may lead to cost cuts, including cut of security measures.

Example of outages: When the Russian army entered Ukraine, western countries deployed sanctions towards Russian entities. On the 3. March 2022, Cogent terminated services to Russian organisations with 24 hours notice and stated they would turn off all co-located equipment and prepare it to be picked up. Lumen at the same time disconnected all their hardware in Russia (Madory, 2022).

2.11 Governance Layer

Governance risks are often underestimated in risk evaluations. These risks can arise from national governments, central internet governance bodies like ICANN and RIR, and centralized services such as IRR, RPKI, and Root DNS. Critical services must be prepared to withstand potential outages of these governance services.

Static risks in the Governance Layer exist during implementation, while dynamic risks involve changes in laws and regulations. Other risks include IPv4 address exhaustion, legal actions such as "cease and desist" letters, and being blocked by governmental filters or embargoes.

Failure to comply with local regulations or national laws on confidentiality and integrity, may lead to severe punishments, which again might impact availability. Breaches can lead to legal orders for temporary or permanent service shutdowns, financial penalties, and customer boycotts, potentially necessitating cost cuts and reduced security measures.

Example of outage: In March 2022, following the Russian army's entry into Ukraine, Western countries imposed sanctions on Russian entities. Cogent terminated services to Russian organizations with 24 hours' notice, while Lumen disconnected their hardware in Russia, causing service disruptions (Madory, 2022)

3 MODEL VERIFICATION

The efficiency of the 10-layer model was verified for two different networks.

3.1 Risk Registry Analysis of Exiting Network

To test the new 10-layer model, we were allowed access to the risk registry from a global network provider, and mapped all the risks that were identified during their ISO27001:2013 risk discovery process into the proposed model as well as into the ISO27001:2013 and NIST800-53 models for comparison. The risks are anonymized, but the statistics may be published.

We see that for ISO27001, each risk maps to on average 8.9 controls (median 8), and for NIST800-53, each risk maps to an average of 4.8 controls (median 5). In the 10-layer model, however, only three risks map to two layers, while all other risks maps to a single layer. For ISO27001 and the 10-layer model, all risks were covered, but for NIST800-53, eight risks were not discovered by any of the sections. The types of missed risks were Governance risks and risks to non-production equipment like lab equipment and equipment during transport.

3.2 Risk Discovery Process for a New Network Service Provider

Our second verification project uses the new 10-layer model to discover risks associated with the implementation of a new small research network for a local research organization. The network spans a metropolitan area, with two sites and two separate IP transit sessions.

The risks for this network was discovered by interviewing the NOC for the new research network, using the 10-layer model as basis. After this risk discovery process, the ISO27001 and NIST800-53 frameworks were briefly consulted to discover any risks that were un-noticed by the 10-layer procedure.

The result of the risk discovery was 55 risk points across all 10 layers, out of which 48 were assigned a mitigation plan.

The second risk discovery process, using the ISO27001 and NIST800-53 frameworks did not reveal any new risk points, and the interviewees (subjectively) found this process more confusing and less straightforward than the process based on the 10-layer model. When asked to elaborate, the subjects stated that the risk areas were not well defined when applied to Network Availability and the 10-layer model was easier to follow.

4 DISCUSSION

The certification market has grown into a multi-billion dollar industry, with standards like ISO27001, NIST800-53, and SOC2 gaining significant momentum. However, we believe that the inherent classification in these standards may not be well-suited for effectively managing network and service availability risks. Relying solely on these standards for risk discovery can lead to confusion, oversights, and unnecessary work, resulting in incomplete risk management and employee frustration.

While none of these standards provide a mandatory risk discovery interview template, we propose our 10-layer model as a suitable foundation for conducting such interviews in alignment with any security standard. This model is familiar to the Network Operations Center (NOC) and encompasses all relevant risks, making it easy to understand and facilitating classification. By using this model, the NOC can gain confidence in their ability to handle all risks effectively.

It's important to note that mitigating every single risk may not be necessary, but being aware of all risks and making informed management decisions about whether to accept or mitigate them is crucial. By confidently producing a comprehensive risk management report using this model, a NOC manager can instill trust in top management, reassuring them that the network and/or service is in capable hands.

In conclusion, while existing certification standards have their merits, our proposed 10-layer model offers a practical and comprehensive approach to risk discovery and management. It empowers the NOC with a familiar framework, facilitates risk classification, and ultimately contributes to a more confident and capable handling of network and service risks.

REFERENCES

Anderson, J. P. (1972). Computer security technology planning study. Technical report, ANDERSON (JAMES P) AND CO FORT WASHINGTON PA FORT WASHINGTON.

AWS (2017). Summary of the amazon S3 service disruption in the northern virginia (us-east-1) region. aws.amazon.com, https://aws.amazon.com/message/41926/?ascsubtag=[]vx[p]14556677[t]w[r]google.com[d]D.

AWS (2021). Summary of the AWS service event in the northern virginia (US-EAST-1) region. aws.amazon.com, https://aws.amazon.com/message/12721/.

Barrett, M. (2018). Framework for improving critical infrastructure cybersecurity version 1.1.

Barton, J. (2022). Networks down in lebanon as ogero workers strike. developingtelecoms.com, https://developingtelecoms.com/telecom-business/operator-news/13926-networks-down-in-lebanon-as-ogero-workers-strike.html.

Belson, D. (2022). AAE-1 & SMW5 cable cuts impact millions of users across multiple countries. blog.cloudflare.com, https://blog.cloudflare.com/aae-1-smw5-cable-cuts/.

Burgess, M. (2022). The Most Vulnerable Place on the Internet. www.wired.com, https://www.wired.com/story/submarine-internet-cables-egypt/.

Bush, R. and Austein, R. (2013). The Resource Public Key Infrastructure (RPKI) to Router Protocol. RFC 6810.

C/LM - LAN/MAN Standards Committee (2000). IEEE standard for information technology - local and metropolitan area networks - part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications-aggregation of multiple link segments. *IEEE Std 802.3ad-2000*, pages 1–184.

Durand, A. (2020). Resource public key infrastructure (RPKI) technical analysis.

Evang, J. M., Ahmed, A. H., Elmokashfi, A., and Bryhni, H. (2022). Crosslayer network outage classification using machine learning. In *Proceedings of the Workshop on Applied Networking Research*, ANRW '22, New York, NY, USA. Association for Computing Machinery.

Freedman, D., Foust, B., Greene, B., Maddison, B., Robachevsky, A., Snijders, J., and Steffann, S. (2019). Mutually agreed norms for routing security (MANRS) implementation guide.

Goovaerts, D. (2021). Extended AWS outage disrupts services across the globe. www.fiercetelecom.com, https://www.fiercetelecom.com/cloud/extended-aws-outage-disrupts-services-across-globe.

Goyal, R. (2022). Zscaler digital experience detects outage. www.zscaler.com, https://www.zscaler.com/blogs/product-insights/zoom-outage-detected-zscaler-digital-experience-zdx.

Graham-Cumming, J. (2022). Partial cloudflare outage on october 25, 2022. blog.cloudflare.com, https://blog.cloudflare.com/partial-cloudflare-outage-on-october-25-2022/.

Hinden, B. (2004). Virtual Router Redundancy Protocol (VRRP). RFC 3768.

Hunter, P. (2008). Pakistan youtube block exposes fundamental internet security weakness: Concern that pakistani action affected youtube access elsewhere in world. *Computer Fraud & Security*, 2008(4):10–11.

ICANN (2007). Factsheet root server attack on 6 february 2007. www.icann.org, https://www.icann.org/en/system/files/files/factsheet-dns-attack-08mar07-en.pdf.

Imbriaco, M. (2012). Downtime last Saturday. github.blog, https://github.blog/2012-12-26-downtime-last-saturday/.

Integrated Human Factors (2022). Facebook & instagram outage likely caused by human error. www.ihf.co.uk, https://www.ihf.co.uk/facebook-instagram-outage-by-human-error/.

ISO (2002). *ISO/IEC 10589:2002 Information technology — Telecommunications and information exchange between systems — Intermediate System to Intermediate System intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service*. International Organization for Standardization, Geneva, Switzerland.

ISO (2015). *ISO/IEC 27017:2015 Information technology — Security techniques — Code of practice for information security controls based on ISO/IEC 27002 for cloud services*. International Organization for Standardization Geneva, Switzerland.

ISO (2018). *ISO 31000:2018(en) Risk management — Guidelines*. International Organization for Standardization, Geneva, Switzerland.

ISO (2022a). *ISO/IEC 27001:2022(en) Information security, cybersecurity and privacy protection — Information security management systems — Require-*

ments. International Organization for Standardization, Geneva, Switzerland.

ISO (2022b). *ISO/IEC 27002:2022, Information security, cybersecurity and privacy protection — Information security controls*. International Organization for Standardization, Vernier, Geneva, Switzerland, ISO/IEC 27002:2022 edition.

Kacherginsky, P. (2022). Celer bridge incident analysis. www.coinbase.com, https://www.coinbase.com/blog/celer-bridge-incident-analysis.

Kachold, L. (2009). Layer 8 linux security: OPSEC for linux common users, developers and systems administrators. linuxgazette.net, https://linuxgazette.net/164/kachold.html.

Katz, D. and Ward, D. (2010). Bidirectional Forwarding Detection (BFD). RFC 5880.

Lepinski, M. and Sriram, K. (2017). BGPsec Protocol Specification. RFC 8205.

Madory, D. (2022). Cogent and lumen curtail operations in russia. www.kentik.com, https://www.kentik.com/blog/cogent-disconnects-from-russia/.

Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and Wright, C. (2014). Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks. RFC 7348.

Moy, J. (1998). OSPF Version 2. RFC 2328.

NETBLOCKS (2020). Mobile internet provides lifeline after earthquake knocks out Puerto Rico infrastructure. netblocks.org, https://netblocks.org/reports/puerto-rico-earthquake-internet-outage-dAmqEDA9.

NIST (2022). Security and privacy controls for federal information systems and organizations. Technical Report NIST Special Publication 800-53, National Institute of Standards and Technology, U.S. Department of Commerce, Washington, D.C.

OECD (2022). *Routing security*. Number 330. The Organization for Economic Cooperation and Development, OECD Digital Economy Papers.

Patterson, D. A., Gibson, G., and Katz, R. H. (1988). A case for redundant arrays of inexpensive disks (RAID). *SIGMOD Rec.*, 17(3):109–116.

Rekhter, Y., Hares, S., and Li, T. (2006). A Border Gateway Protocol 4 (BGP-4). RFC 4271.

Rustignoli, N. and de Kater, C. (2022). SCION Components Analysis. Internet-Draft draft-rustignoli-panrg-scion-components-01, Internet Engineering Task Force. Work in Progress.

Santiago, R., de Onís, C. M., and Lloréns, H. (2020). Powering life in puerto rico. *NACLA Report on the Americas*, 52(2):178–185.

Silberling, A. (2022). Zoom is down in a major outage. www.techcrunch.com, https://techcrunch.com/2022/09/15/zoom-is-experiencing-a-major-outage/.

Taylor, S. and Wexler, J. (2003). Mailbag: OSI layer 8 - money and politics. www.networkworld.com, https://www.networkworld.com/article/2339786/mailbag--osi-layer-8---money-and-politics.html.

The SlowMist Security Team (2022). Truth behind the Celer Network cBridge cross-chain bridge incident: BGP hijacking. medium.com, https://medium.com/coinmonks/truth-behind-the-celer-network-cbridge-cross-chain-bridge-incident-bgp-hijacking-52556227e940.

Thimbleby, H. (2022). Failing IT infrastructure is undermining safe healthcare in the NHS. www.bmj.com, https://www.bmj.com/content/379/bmj-2022-073166/rr.

Toonk, A. (2015). Massive route leak causes internet slowdown. www.bgpmon.net, https://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/.

Trenaman, N. (2020). Downtime last Saturday. www.ripe.net, https://www.ripe.net/ripe/mail/archives/routing-wg/2020-February/004015.html.

Viswanathan, A., Rosen, E. C., and Callon, R. (2001). Multiprotocol Label Switching Architecture. RFC 3031.

Zimmermann, H. (1980). OSI reference model-the ISO model of architecture for open systems interconnection. *IEEE Trans. Communication (USA)*, COM-28(4):425–432. IRIA/Lab., Rocquencourt, France.

Zoom (2022). Issues starting and joining meetings incident report for zoom. status.zoom.us, https://status.zoom.us/incidents/k7fm2j5q8lx1.

Appendix F

Paper VI

Jan Marius Evang. “Outage risk priorities – It’s not the malicious attacks that take down your service”. In: **IEEE Hotnets 2023**. 2023, Submitted for evaluation.

Status: Accepted for evaluation

[Article not attached due to copyright]

Appendix G

Paper VII

Jan Marius Evang and Haakon Bryhni. “National ICT Resilience: An analysis of Norway’s cyber infrastructure preparedness”. In: **IEEE Communications Magazine**. 2023, Submitted for evaluation.

Status: Accepted for evaluation

[Article not attached due to copyright]