# 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*

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.
DOI: 10.5220/0012092600003555
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<br>"Layer 2.5"<br>Layer 2 | Leased network services. |
| 2 | Campus Area Network | Layer 3<br>Layer 2 | Networks fully operated by us. |
| 1 | Physical | Layer 1<br>"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.