# Outage Risks: It is not the Malicious Attacks that Take Down Your Service

Jan Marius Evang[1,2] and Alojz Gomola[1]
[1] *Simula Metropolitan Center for Digital Engineering, Norway*
[2] *Oslo Metropolitan University, Norway*

Keywords: Risk Management, Networks, Resilience.

Abstract: Network operators often prioritize risks to optimize resource allocation. This paper introduces an analysis model for prioritizing network outage risks, a practice common among large operators but underrepresented in research. We propose metrics such as Risk Value and clarify their definitions. The study reveals insights into the impact of incidents, classifying them based on customer support cases. Through the application of the presented method to a global network operator, we demonstrate the generation of unexpected insights and outcomes: Short outages are very frequent and regularly cause customer complaints. We examined both safety and security related incidents in relationship to customer report events, with surprising turn out for malicious attacks.

## 1 INTRODUCTION

In the ever-evolving landscape of network operations, ensuring the uninterrupted delivery of services is paramount. Network operators face the perpetual challenge of mitigating outage risks while optimizing the allocation of finite financial and personnel resources. While industry practices often involve prioritizing these risks, the scholarly exploration of methodologies and models in this domain remains surprisingly limited.

This paper aims to address this gap by introducing an advanced analysis model tailored specifically for prioritizing network outage risks. This practice, while commonplace among large-scale network operators, has yet to receive the depth of attention it deserves within the academic research community.

At the core of our contribution lies the introduction of new metrics, notably the Risk Value and Impact Score. Figure 2 illustrates the connection between the classic risk matrix and the new metrics. Note that the highest Impact Score will be attained where few incidents lead to many cases, and the highest Risk Value are for root causes with high Impact Score and many cases. These metrics provide a robust and standardized framework for assessing and prioritizing outage risks, enhancing the overall resilience of network infrastructures.

By using our proposed analysis model, we go on to a comprehensive study into the impact of incidents, leveraging the unique perspective of customer support cases. The empirical application of our model to a global network operator not only validates its efficacy but also yields unexpected and insightful outcomes.

Our findings challenge conventional assumptions, revealing that short outages, despite their brevity, are both pervasive and a consistent source of customer complaints. On the contrary, our study finds that malicious attacks rarely causes outages, yet have a high impact when they do occur.

This research not only seeks to refine and enhance the practical methodologies employed by network operators but also contributes to the theoretical foundations of network risk management. By bridging the gap between industry practices and academic discourse, our analysis model and associated metrics offer a framework for addressing the complex landscape of network outage risks.

## 2 PREVIOUS WORK

In 2018, Aceto et al. conducted a survey study on Internet outages (Aceto et al., 2018). Although the results of the study may be somewhat outdated, the theoretical descriptions and analysis presented remain relevant. It is worth noting, however, that the focus of the current paper is specifically on outages that im-

Evang, J. and Gomola, A.
Outage Risks: It is not the Malicious Attacks that Take Down Your Service.
DOI: 10.5220/0012505400003708

In *Proceedings of the 9th International Conference on Complexity, Future Information Systems and Risk (COMPLEXIS 2024)*, pages 77-82
ISBN: 978-989-758-698-9; ISSN: 2184-5034

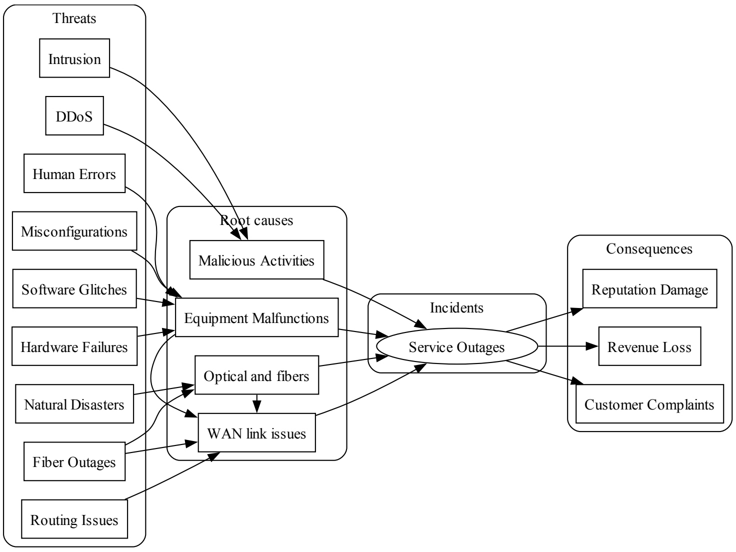

Figure 1: A bowtie diagram illustrating the identified network outage risks.

pacted an Internet Service Provider's (ISP) internal global network, rather than the Internet as a whole.

Another significant contribution to the field is the analysis conducted by Wang et al. (Wang and Franke, 2020) on the economic aspects associated with network outages in enterprises. Their study primarily examined two key aspects: the cost implications of long-term outages and the role of insurance premiums in mitigating the financial impact. The research provided valuable insights into the economic considerations that decision-makers should take into account when managing network outage risks.

Another study by the same authors (Franke and Buschle, 2016) explored the investment decisions of decision-makers regarding security measures. The research revealed a general tendency among decision-makers to misallocate security investments. This finding underscores the importance of raising awareness among decision-makers about security prioritization to improve Return on Security Investments (RoSI).

Another notable analysis in (Govindan et al., 2016) focused on 100 high-impact outages in the Google Network Infrastructure. The paper provides useful insights into the network components causing outages, indicates that most outages stem from maintenance activities, and offers important advice on risk reduction. Notably, the paper does not mention malicious activities as potential causes for outages and lacks a comparison to low-impact incidents.

Acklyn et al. in (Murray and Rawat, 2021) developed a model for network hazard flow and malware detection based on a Long-Short Term Memory (LSTM) algorithm. While extensive, the focus is on malware and not on network outages as such.

Despite these valuable contributions, there remains a research gap in guiding decision-makers to prioritize mitigation efforts and investments for network outages, specifically within the context of an ISP's internal global network. While previous studies have emphasized economic considerations and misallocation of investments, there is a clear need for practical frameworks or guidelines tailored to this specific domain.

Therefore, the current research aims to address this research gap by proposing a practical method that enables decision-makers in an operational ISP environment to effectively prioritize their mitigation efforts and allocate resources based on the specific risks and characteristics of their internal global network.

# 3 TYPES OF NETWORK OUTAGES

The main areas of network outage risks are summarized in Table 1 (Evang, 2023), providing insight into the vulnerabilities that can disrupt a typical large-scale network. Physical layer outages are often triggered by natural disasters and weather-related events, while local network layer outages primarily result from maintenance activities (Franken et al., 2022). Wide area Layer 2 network services are susceptible to frequent packet loss incidents, and applications face the risks of hacking and Denial of Service (DoS) attacks. The Internet layer encompasses risks stemming from both mistakes and malicious attacks, including cloud service risks and challenges in service delivery. Cyber attacks and hacking attempts usually receive the highest focus in the Internet layer, alongside instances of human error and organizational policy flaws. Regrettably, governance risks are frequently overlooked, warranting attention in comprehensive risk assessments.

Table 1: The 10-layer model for network service risk assessment.

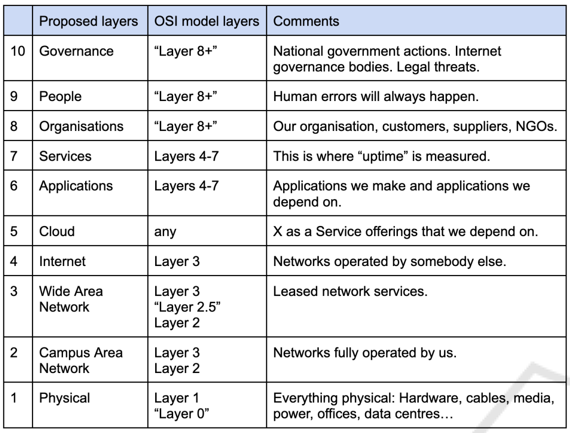

| | Proposed layers | OSI model layers | Comments |
|---|---|---|---|
| 10 | Governance | "Layer 8+" | National government actions. Internet governance bodies. Legal threats. |
| 9 | People | "Layer 8+" | Human errors will always happen. |
| 8 | Organisations | "Layer 8+" | Our organisation, customers, 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… |

In Section 6, we provide a comprehensive analysis showcasing the frequency of different types of network outage incidents encountered by a global network provider. This examination sheds light on the relative occurrence of various incidents, enabling a better understanding of the prevailing risks in network operations.

## 4 FACTORS TO CONSIDER WHEN PRIORITIZING AND EVALUATING OUTAGE RISKS

Risk assessment (Rausand and Haugen, 2020) commonly involves assessing two key variables: likelihood and impact. Likelihood refers to the probability of an incident occurring, while impact puts a numerical value to the consequences of such an incident. The risk value is calculated by multiplying impact by likelihood, and the risk management policy determines acceptable and actionable risk values. Thus, a high likelihood risk may be tolerable if the impact is low, and conversely, the risk of a high-impact fault may be acceptable if the likelihood is low.

When managing critical infrastructure, the repercussions of an outage can be significant, underscoring the importance of measures to diminish the probability and alleviate the potential consequences. Evaluating the impact of an incident is frequently more straightforward than gauging the likelihood. For example, a non-redundant fiber failure would lead to a service disruption, with potentially severe consequences for critical services. Impact mitigation is usually implemented by redundant services and fast failover. In certain scenarios, a solitary fiber cut in a central location may have a cascading effect, affecting multiple services such as both internet and cellphone connectivity, amplifying the overall impact.

Estimating likelihood can be more challenging, and it is often based on subjective expert opinions. Mean Time Between Failures (MTBF) estimates provided by equipment manufacturers can be useful for physical equipment. Alternatively, analyses like the one described in (Evang et al., 2022) can be conducted. The investigation in that paper involved analyzing a database of customer complaints related to network services. The root causes of these complaints were identified, and a machine learning system was developed to accurately classify new outages. The machine learning model was trained on cases with known root causes and customer complaints. Subsequently, the model was applied to all outage data, uncovering some interesting insights, as described in Section 6.

When critical infrastructure is at stake, the impact of an outage could be substantial, necessitating proactive measures to minimize the likelihood of incidents and, if feasible, reduce their potential consequences. Outages can lead to various impacts, including revenue loss, reputational damage, and diminished customer satisfaction. In the case of critical services, outages may even result in harm to individuals' health or loss of life. Outages that exceed agreed-upon Service Level Agreements may constitute a breach of contract and, in certain cases, a violation of laws and regulations governing critical service providers, potentially leading to legal repercussions.

## 5 BEST PRACTICES FOR MANAGING NETWORK OUTAGE RISKS

Implementing effective risk management practices is essential for mitigating network outage risks. Various frameworks, such as those found in the ISO27001 Information Security Management Systems (International Organization for Standardization, 2022) and ISO31000 Risk Management standards (ISO 31000:2018(en), 2018), offer valuable guidance. When preparing for ISO27001 certification, organizations undergo a comprehensive risk evaluation process. All risks within the company are assessed, and mitigation strategies are determined for those

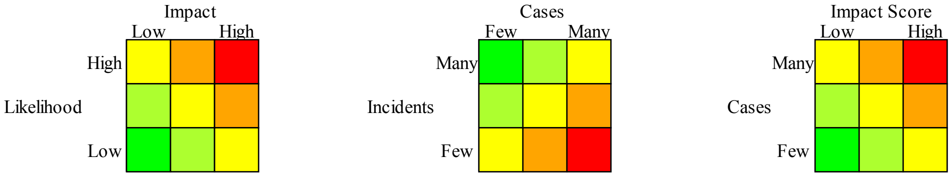

(a) Classic risk matrix (b) Impact Score matrix (c) Risk Value matrix

Figure 2: Comparison of matrices. Red indicates areas that require special attention.

deemed significant. Even insignificant risks are registered, and appropriate mitigation actions are identified. Company policies are scrutinized to ensure they adequately address the identified risks and mitigations, and updates are made as necessary. Regular re-evaluation of the risk registry allows for the discovery of new risks and assessment of the effectiveness of implemented mitigations.

For all identified risks, conducting an explicit or implicit Return on Security Investments (RoSI) evaluation is crucial. This evaluation involves estimating the monetary impact of an outage by analyzing the financial losses incurred during similar past incidents. Additionally, the frequency of such outages is estimated. The cost of implementing mitigation actions is then evaluated and compared to the potential incident cost over a comparable time period. Any mitigation measure where the impact cost exceeds the cost of the mitigation will be a worthwhile security investment.

By employing these practices, organizations can make informed decisions regarding risk mitigation, ensuring that resources are allocated effectively to reduce the likelihood and impact of network outages.

## 6 EXPERIMENTAL VALIDATION

In this section, we conduct an exploration of our proposed analysis model's effectiveness in prioritizing network outage risks. The methodology involves the application of our model to a global network operator, allowing us to draw actionable insights and validate the robustness of our approach.

The dataset used for this analysis is referenced in (Evang et al., 2022). Over an 18-month period, active and passive network measurements revealed a total of 700,000 network **incidents**, while 2,855 customer **cases** were reported in relation to network incidents during the same time frame. Careful post-mortem of all the customer cases allowed for retrospective determination of the root causes. Notably, it was found that only around 35% of the customer cases received a correct root cause response at the time of reporting. In this paper, we will use the **cases** as an indication of the **impact** of an incident. An incident without a customer-case is regarded as low-impact, while an incident resulting in a customer case is regarded as high-impact.

An intriguing aspect of the study was the application of the trained machine learning model to the outages that did not prompt any customer complaints. This provided an estimate of the number of outages of each type that actually resulted in customer complaints, serving as an indicator of the impact level associated with each type of outage. The method we employ involves comparing the percentage of each type in the customer case dataset to the percentage of each type predicted in the non-complaint outages, leading us to define an impact score (IS) using Equation 1.

$$IS = \frac{\text{cases\%}}{\text{incident\%}} \tag{1}$$

With the corresponding Risk Value from Equation 2.

$$RV = \text{cases\%} * \text{IS} \tag{2}$$

Table 2 provides an overview of the root causes of cases and incidents, presenting the corresponding Impact Score and Risk Values. The analysis revealed that the predominant cause of network incidents was attributed to packet loss and outages in the leased Layer 2 WAN links, accounting for 85% of cases and 25% of incidents. While the precise origins of these outages were not conclusively identified, potential factors include physical fiber outages, equipment failures, and maintenance and human errors.

The second most prominent cause of outages was attributed to equipment maintenance or equipment failure, accounting for 53% of the non-WAN incidents and 43% of the related customer cases.

The third most significant cause was linked to failures in optical links, subsea cables, and metro-connects. (46% of non-WAN cases and 57% of non-wan incidents)

A particularly surprising result emerged from the data, revealing the near absence of malicious attacks. Only four customer complaints (1%) were attributed to such attacks, and the estimated number of incidents was only 0.01%.

Evaluation of the impact demonstrated that Layer2 WAN outages and incidents involving subsea cables, metro-connects, and Layer1 infrastructure had a substantial impact. Additionally, although the number of DoS attacks was relatively low, they exhibited a high impact.

Table 2: Variables used in impact and likelihood calculation (non-wan percentages).

| Type of outage | cases | incidents | IS | RV |
|---|---|---|---|---|
| WAN link issues | 85% | 25% | 3.4 | 2.9 |
| Equipment | 53% | 43% | 1.2 | 0.64 |
| Optical and fibers | 46% | 57% | 0.8 | 0.37 |
| Malicious attacks | 1% | 0.01% | 100 | 1 |

Table 2 serves as a valuable reference point for the operator, prevalence of short outages prompts a reconsideration of their significance in outage management strategies. Simultaneously, the infrequency yet high impact of malicious attacks underscores the need for targeted security measures.

## 7 CONCLUSION

In this study, our main objective was to investigate and analyze network outage risks and their impact on a global Internet Service Provider (ISP). Through analysis of passive and active outage measurement data and examination of customer cases, we gained valuable insights into the causes and consequences of network outages.

Our investigation identified packet loss and outages in leased Layer 2 WAN links as the primary contributors to network incidents. While the definitive causes of these outages were not ascertained, factors such as physical fiber outages, equipment failures, and maintenance/human error are likely contributors. Equipment maintenance and failures emerged as significant causes of outages, representing a substantial portion of incidents.

The relatively low number of cases (2855) compared to the total incidents (700,000) is a result of the implementation of fast failover mechanisms and a deliberate focus on achieving “fail open” risk reduction strategies.

Consistent with observations in (Govindan et al., 2016), malicious attacks were nearly absent from the data, with only a minimal number of customer complaints attributed to such attacks. This suggests that the existing security measures implemented by the ISP have proven effective in mitigating this specific risk.

Our impact evaluation highlighted the significant consequences of Layer2 WAN outages and optical failures. Although the number of Denial-of-Service (DoS) attacks was relatively low, they exhibited a high impact when they did affect the service.

The findings of this study have important implications for network operators and service providers. By understanding the key causes of outages and their impact, operators can prioritize their resources and efforts to effectively mitigate risks and minimize disruptions. Additionally, the near absence of malicious attacks emphasizes the importance of maintaining robust security measures to prevent potential future threats.

It is important to acknowledge the limitations of this study. Our analysis focused on one specific ISP, and the findings may not be generalizable to other network operators. Furthermore, the underlying causes of certain outages were not definitively identified, warranting further investigation.

To further advance research in this area, future studies could explore the specific mechanisms and root causes of different types of outages, allowing for more targeted risk mitigation strategies. Additionally, examining the effectiveness of various security measures and their impact on reducing the likelihood and impact of outages would provide valuable insights for network operators.

In conclusion, this study has shed light on the risks and impacts associated with network outages for a global ISP. We show that the most important focus area is the physical layer, in making sure that outages of cables and equipment are handled. Outages caused by malicious attacks have a high impact, but do not significantly contribute to the number of outages.

By leveraging this knowledge, risk management can be performed continuously at an operational stage. Impact Score can be easily calculated, and the number of cases can be reported. This way network operators can ensure the continuity of their services, minimize disruptions to customers, and maintain a secure and reliable network infrastructure. Ultimately, this research contributes to the broader understanding of network outage risks and supports efforts to enhance network security and reliability in an increasingly interconnected world.

## ACKNOWLEDGEMENTS

Language improvements by (OpenAI, 2022).

## REFERENCES

Aceto, G., Botta, A., Marchetta, P., Persico, V., and Pescap, A. (2018). A comprehensive survey on internet outages. *Journal of Network and Computer Applications*, 113:36–63.

Evang, J. M. (2023). A 10-layer model for service availability risk management. In *Proceedings of the 20th International Conference on Security and Cryptography - SECRYPT*, page to appear. INSTICC, SciTePress.

Evang, J. M., Ahmed, A. H., Elmokashfi, A., and Bryhni, H. (2022). Crosslayer network outage classification using machine learning. In *Proceedings of the Applied Networking Research Workshop*, ANRW '22, page 17, Philadelphia, PA, USA. Association for Computing Machinery.

Franke, U. and Buschle, M. (2016). Experimental evidence on decision-making in availability service level agreements. *IEEE Transactions on Network and Service Management*, 13(1):58–70.

Franken, J., Reinhold, T., Reichert, L., and Reuter, C. (2022). The digital divide in state vulnerability to submarine communications cable failure. *International Journal of Critical Infrastructure Protection*, 38:100522.

Govindan, R., Minei, I., Kallahalla, M., Koley, B., and Vahdat, A. (2016). Evolve or die: High-availability design principles drawn from googles network infrastructure. In *Proceedings of the 2016 ACM SIGCOMM Conference*, SIGCOMM '16, page 5872, New York, NY, USA. Association for Computing Machinery.

International Organization for Standardization (2022). *ISO27001, Information security, cybersecurity and privacy protection Information security management systems Requirements*. International Organization for Standardization, Vernier, Geneva, Switzerland, ISO/IEC 27001:2022(en) edition.

ISO 31000:2018(en) (2018). *Risk management Guidelines*. ISO, Geneva, Switzerland.

Murray, A. and Rawat, D. B. (2021). Network hazard flow for multi-tiered discriminator analysis enhancement with systems- theoretic process analysis. In *2021 IEEE Global Humanitarian Technology Conference (GHTC)*, pages 55–61.

OpenAI (2022). ChatGPT. https://openai.com/chatgpt. Accessed: 12 February 2024.

Rausand, M. and Haugen, S. (2020). *Front Matter*, pages i–xx. John Wiley & Sons, Ltd.

Wang, S. S. and Franke, U. (2020). Enterprise it service downtime cost and risk transfer in a supply chain. *Operations Management Research*, 13(1-2):94–108.