(Reimagining) Resilience Goals for the Internet

Cecilia Testart (Georgia Institute of Technology)
Volker Stocker (Weizenbaum Institute & TU Berlin)[1]
David Clark (MIT)
William Lehr (MIT)

*** *This version: 30 August 2024* ***

Table of Contents

[1] Volker Stocker acknowledges funding by the Federal Ministry of Education and Research of Germany (BMBF) under grant no. 16DII131 (Weizenbaum-Institut für die vernetzte Gesellschaft—Das Deutsche Internet-Institut). Moreover, he gratefully acknowledges funding from the Georgia Institute of Technology for a research visit during which parts of this paper were developed.

Abstract

The Internet has become a critical basic infrastructure for society and the economy. As such, the resilience of the Internet is essential for the functioning and resilience of many sectors of our digital economies. In our increasingly digital society and economy, Internet is a key element of cyber-resilience. Understanding how to assess and promote Internet resilience are important policy challenges. In this paper, we discuss key aspects of the Internet design that differentiate it from other critical infrastructure and that are essential in considering its resiliency. We also explore how to achieve resiliency and provide a series of illustrative examples that reveal the challenges of achieving and measuring Internet resiliency, as well as mechanisms that contribute to it. Finally, we offer our thoughts on the policy challenges and approaches for promoting Internet resiliency.

1. Introduction

The Internet has become a critical basic infrastructure for society and the economy. As such, the resilience of the Internet is essential for the functioning and resilience of many sectors of our digital economies—during normal and abnormal times.[2] This was underscored during the COVID-19 pandemic when the forced shift of social and economic activity online resulted in a deeper integration of connected ICTs into our lives (Feldmann et al., 2021; Stocker et al., 2023).

The word *resilience* is often proposed as an aspiration for critical infrastructure, but without a careful discussion of what that word might actually mean, it is difficult to consider how we might either measure resilience or improve it.

As a starting point, the word resilience captures the idea that a system would continue to function during incidents like natural disasters, cyber-attacks, and other disruptive events due to human or system errors. We contrast resilience with other possible aspirations for critical infrastructure. Resilience implies that adverse events may degrade the system to some extent, but it can continue to provide service and recover effectively. In contrast, one can try to design a system so that it is immune to the consequences of adverse events—a *resistant* rather than a resilient system. For example, we think of bridges as resistant to failure. We expect them to perform as normal, even under adverse circumstances. When an event occurs that is outside the design specification (e.g.,

[2] There is no single definition of what constitutes "the Internet" (see Lehr, Clark et al, 2019a). A commonly cited definition is "a network of networks," which is accurate but not very informative. For our purposes herein, the Internet is the network of publicly routable IP addresses (IPv4 and IPv6) that support end-to-end data packet transport using the Internet suite of protocols (i.e., IP, TCP, UDP, etc.). The cited paper explores this definitional issue more thoroughly. The more limited definition we adopt herein is to help distinguish the Internet from the private networks and overlay networks that are part of the Internet ecosystem more broadly construed.

a collision with a container ship), the bridge may suffer catastrophic failure. *Resistant* systems are immune to adverse events up to a point, and brittle beyond that point.

A 2013 Policy Directive on Critical Infrastructure defined resilience as follows:

> *"The term 'resilience' means the ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions. Resilience includes the ability to withstand and recover from deliberate attacks, accidents, or naturally occurring threats or incidents."*[3]

This definition applies to all critical infrastructures, including digital infrastructures. At a high level, policy-making can concern itself with all of critical infrastructure. More specifically, policy-making can concern itself with cyber-resilience, which encompasses all of the digital realm. In this paper, we focus specifically on the resilience of the Internet as an infrastructure that underpins much of the digital realm, and many critical services that may not initially seem like cyber-infrastructure.

All basic infrastructures share the feature that they are used by virtually all sectors and economic activities, and their availability, accessibility, and performance are largely taken for granted until something goes wrong. A fundamental role for government and public policy is to make sure that citizens and businesses have the basic infrastructure they need to sustain economic growth and meet welfare goals. That is a core mission for economic and industrial policy. Moreover, in the U.S., Europe, and other liberal economies, the pursuit of infrastructure policies should promote, or at least, minimally interfere with the operation of competitive markets. However, the technologies, business/market models, and regulatory policies that ensure basic infrastructure needs are met vary significantly across infrastructures, both non-digital and digital.

In Section 2, we discuss what is distinct about the Internet as infrastructure—how its features, both technical and organizational, shape the challenge of Internet resilience. In Section 3, we return to the challenge of defining and conceptualizing resilience, and explore in more depth how one can achieve resilience, using the specifics of the Internet to illustrate the challenge. In Section 4, we explore a series of examples that illustrate the challenges and mechanisms that contribute to

---

[3] This definition is from a 2013 U.S. Presidential Policy Directive on Critical Infrastructure Security and Resilience (The White House, 2013). In citing this definition, it is important to remember that the directive was not directed solely at the Internet, but was intended to address all critical infrastructures. The Policy Directive explains that the "*Nation's critical infrastructure is diverse and complex. It includes distributed networks, varied organizational structures and operating models (including multinational ownership), interdependent functions and systems in both the physical space and cyberspace, and governance constructs that involve multi-level authorities, responsibilities, and regulations."*

Internet resiliency. In this context, we further explore the relationship between Internet resilience and overall cyber resiliency.

Based on the lessons learned from previous sections, in Section 5, we offer our thoughts on the policy challenges and approaches for promoting Internet resiliency and how these differ from those of ensuring cyber resiliency more generally. The essence of those recommendations is the need to sustain what is uniquely characteristic of the Internet and avoid treating its technical or economic resiliency challenges using methods and approaches that are often appropriate above and below the waist. Section 6 provides a summary of the main points and thoughts about future directions.

2. What is Different About the Internet?

There are three key aspects of the Internet that shape considerations of resilience: generality, layering, and decentralization. These aspects set apart the Internet from other critical infrastructure. This section discusses how those aspects impact how we think about Internet resilience.

2.1. Generality

The Internet's generality is inherently different than other custom-built, single-purpose infrastructures; the Internet serves many purposes and supports many applications that have different requirements.

Users may think of the Internet in terms of those applications—they associate the term Internet with their experience using it—email, social media, streaming content, and the like. But technically, the Internet is a data transport layer on top of which these services are built (e.g., Claffy and Clark, 2014).

The Internet has resilience capabilities as a result of its basic design principles. Data on the Internet is broken into units called packets, and the design of the Internet was based on the assumption that occasionally packets would be lost in transit. Rather than design the packet carriage layer to be *resistant* to failure (assume a packet is never lost, but causes a major disruption to the higher layer service if this assumption was violated), the design is *resilient* to packet loss. The Transmission Control Protocol (TCP) detects lost (or out of order) packets, arranges for lost packets to be resent, puts them in order, and delivers a reliable data stream to the higher layer service. Similarly, circuits and packet-switching hardware (routers) may fail, and the routing protocols of the Internet implement dynamic adaptations to find and exploit functioning paths.

The Internet does not deliver a fixed service with set parameters. Some parts of the Internet have a higher capacity than others; some parts of the Internet manifest longer end-to-end delays than

others. The applications that use this data carriage infrastructure are expected to adapt to these variations as they encounter them.

For example, video streaming services today are designed to adapt to changing transmission capacity (whether due to congestion from excess demand or rerouting over lower capacity links due to failures) by reducing the resolution of the encoded video. The quality of the picture may be reduced, but the video is still delivered. This suggests a basic design approach specific to the Internet: *availability* is more important than *performance*. Providing degraded performance during adversity is a practical response, since applications may be able to adapt and continue to provide an "Internet experience" in a degraded but adequate way that is "better than nothing". This approach has always been one of the basic design principles shaping the Internet.

However, this design approach makes it very difficult to talk about the *quantification* of resilience. The degree of degradation of service will be a response to some adverse event, and so what we need if we are to quantify resilience is some function that relates the degree of degradation to some measure of the degree of adversity. Further, any consideration of a function that relates the degree of service degradation to the nature of the adverse event must include economic considerations of costs and benefits since, technically, there may be no relevance to exceeding minimum resilience,[4] but since providing resilience is costly, too much is suboptimal. However, we do not have a way to quantify adverse events. We return to the challenge of measuring and quantification of resilience in Section 5.

2.2. Layering

As mentioned earlier, the resilience of the Internet only matters because there are critical services running on it. This is the *layered* aspect of the Internet. What the users actually care about is the availability of the higher-level services—and their experience in using these services. Resilience can be designed into the system at every layer. The designers of the physical infrastructure over which the Internet runs (fibers, copper cables, wireless links, satellites, and so on) can try to make them resilient (or resistant). The packet carriage layer can then take steps to provide resilience in the face of failures at the physical layer. And finally, the applications themselves can take steps to be resilient to problems at the packet carriage layer.

The layered approach to resilience can be very powerful, but can also be deceptive, if a designer at one layer fails to understand what the actual characteristics (including failure characteristics)

[4] The US President's Council of Advisors on Science and Technology (PCAST, 2024) has issued a report to the President titled "Strategy for Cyber-Physical Resilience: Fortifying Our Critical Infrastructure for a Digital World" in which they elaborate on performance goals for critical infrastructures and suggest *"to create an integrated set of Critical Infrastructure Performance Goals that define minimum viable delivery objectives for services that are integral to our daily lives."*

are of the layer below. In Section 4 we discuss adverse events where the layered approach made it challenging to recognize the Internet infrastructure vulnerability to a single failure.

2.3. Decentralization

The third key feature of the Internet that shapes any consideration of resilience is its decentralized character. The Internet today is made up of about 75k independent regions that are called Autonomous Systems (ASs) precisely because they can make autonomous decisions about the degree of redundancy they engineer into their part of the infrastructure, what operational and management processes they put in place to deal with unexpected events, and so on. This compounds the complexity of reasoning about resilience—in the case of resilience to hostile attacks, for example, a cyber-attack from country A on a location (e.g., a choke point like an IXP or DNS) in country B is easily feasible, as the distance is only a few milliseconds. Moreover, a local incident can create global harm. In this context, dependencies are crucial considerations.

3. How to achieve resilience

We identify three elements of design for resilience: preparation, limiting and recovering, and learning. [5]

3.1. Preparation

To the extent that designers can identify specific failures, attacks, and the like, they can try to make the system resistant to them, without trying for the larger goal of overall resistance. However, the goal of resilience requires that the designers identify mechanisms that can cope with adverse events without knowing exactly what they will be. A key design approach is *generality* of the mechanisms for recovery and adaptation.

From an infrastructure perspective, there is a need to build redundancy and robustness to prevent single points of failure, reduce vulnerability, and enhance the ability to bounce back and adapt. However, enhancing Internet resilience through redundancy comes at a cost. In a perfect world absent of natural disasters, malicious attacks, and other disruptions like system or human errors, redundancy seems like a wasteful overprovisioning of resources. However, the world as we see it

[5] The following draws and expands on the insights from the literature on organizational resilience and the distinction into three resilience stages (anticipation, coping, and adaptation) as presented in Duchek (2020), who synthesizes insights from the relevant literature into a capability-based conceptualization of organizational resilience. A similar approach that is more granular and focused on regional infrastructure resilience is presented in CISA (2021).

is not perfect. In fact, the threat landscape is broad and continues to evolve.[6] More natural disasters can be expected due to climate change and cyber attacks have emerged as a global problem as more social and economic activity has moved online, geopolitical tensions have grown, and cyber security has emerged as a national priority. As the threats increase, the need for redundancy increases. Redundancy strategies can help to maintain system or service availability at basic functionality or service levels.

Redundancy can be achieved by appropriately provisioning (excess) capacity. One example where this was shown was during COVID-19, where headroom in broadband and server capacity helped to cope with the unexpected surge in demand for online services from home. At the same time, interconnection strategies (more routing diversity) and capacity had to be upgraded (e.g., based on automation at IXPs) (e.g., Stocker et al., 2023; Feldmann et al., 2021). Another example is to have redundant routing paths that can be used in case of (local) failure/outage. While backup routines and backup power supply can help in events like natural disasters (e.g., storms or earthquakes) or malicious attacks on power supply systems, update routines are essential to make software systems less vulnerable and more robust (over their lifecycle).

Moreover, in line with the principle of "availability beats performance" for basic Internet resilience during emergencies, diversification emerges as a key strategy to ensure Internet resilience. Diversity facilitates adaptation via fallback and backup options in case of emergency and makes use of the fact that alternative resources and systems have non-aligned failure modes and/or are deployed in different geographical/physical or contractual contexts. For example, there are a range of broadband access technologies provided by diverse service providers that can support basic connectivity, ranging from wireless to wired, terrestrial to non-terrestrial. Each of these access networks is physically separated. A local outage does not necessarily impair other technologies and services. While hurricanes or earthquakes may impact all terrestrial broadband technologies similarly, potentially leading to the destruction of facilities that need time to be rebuilt, non-terrestrial technologies (e.g., LEO satellites) may come to the rescue as these are not impacted in the same way. Moreover, means to deliver temporary emergency connectivity to regions struck by natural disaster may enhance Internet resilience.[7] While all these examples point to technological forms of redundancy, there is also contractual diversity. End-users (i.e., individuals and companies) may multihome and engage in contractual relationships with multiple firms either horizontally (e.g., interconnections) or vertically (e.g., as complementors on cloud platforms via multi-cloud strategies). Such diversification reduces the dependency on a single company and the availability and performance (incl. security) of their resources and services.

---

[6] Note also that growing numbers and shares of objects are augmented with digital elements, thereby expanding attack surface and new security and resilience concerns.

[7] See, for example, the "cell on wheels" solutions to offer short-term mobile cells to offer emergency connectivity for disaster-struck regions (https://cellsitesolutions.com/products/cell-on-wheels).

From a measurement and monitoring perspective, it is critical to set up measurement and monitoring systems to establish baseline operation benchmarks. Based on the aggregation and analysis of measurement data (e.g., in network operation centers), it is possible to identify anomalous and susceptible behavior and activity. Moreover, monitoring of networks is a precondition for the ability to dynamically adapt system loads in response to current network resource conditions and dynamic rate adaptation as used by CDNs for delivering video streaming content.

As well, it is critical to hire and train humans appropriately so that they are best able to develop and implement operational best practices (i.e., pre-specified situation-specific incident response plans) and contingency plans. That being said, making use of automated response plans (e.g., based on robots at IXPs that facilitate quick and cost-efficient upgrades to interconnection capacity as well as self-patching networks) become increasingly feasible and can be used as part of a comprehensive resilience strategy.

3.2. Limiting and recovering

When an adverse event does occur, resiliency implies two objectives: *limiting* the resulting degradation and *recovering* from the degradation as rapidly as possible. Abstractly, this can be achieved by coping quickly and effectively based on the use (i.e., the implementation, combination, or recombination) of available knowledge, "tools, processes, and practices" (Polk, 2018, p.3).

More specifically, from an infrastructure perspective, strategies that containerize and modularize infrastructure as much as possible can help to contain the scope of the damage. System, hardware, network, and software design flexibility and decisions may help reduce the duration and scope of harms/events. The ability for fast recovery (e.g., load balancing via redirection of loads to redundant backup resources/systems/links) or the ability to dynamically adapt system loads in response to current network conditions (e.g., monitoring of networks and dynamic rate adaptation as used by CDNs for delivering video streaming content), or prioritization in case of emergency-induced scarcity (e.g., CNN turned off certain services to provide more server capacity to others when 9/11 happened) may also present options to mitigate the impact of an incident. Moreover, redundancy and (physical or virtual) isolation strategies can help to keep the problem 'local' and enhance the agility, flexibility, and recoverability of a system to reduce adaptation time.

From an operational perspective, the agility to quickly assess the situation and implement and/or develop incident response and contingency strategies and plans to quickly contain and minimize harm (e.g., locally and temporally) is critical. The quicker the response is implemented, the shorter the duration of the harm. From a measurement and monitoring perspective, real-time monitoring and documentation of incidents is critical to provide the basis for learning and adaptation. In this

process, correct fault attribution (e.g., with regard to the reason/actor, component(s) affected, location, time/duration, etc.) emerges as a critical capability.[8]

Finally, it is important to prepare human capital for crisis and incident response. This ties in with the points made above, as it is crucial that the staff is capable of implementing situation-specific response strategies.

3.3. Learning

A final (and critical) aspect of resilience is learning from and adapting in response to incidents. Through this, systems get stronger. Experience can help improve robustness and redundancy, enhance the capacity and ability to mitigate situation-specific harms, and develop best response strategies and contingency plans.

From an infrastructure perspective, events can help improve the mapping of previously unidentified dependencies in the system—these can be physical or abstract and be horizontal (e.g., via interconnections) or vertical (e.g., between a platform and a complement(or)). In the same vein, new vulnerabilities can be identified, and the degree to which local incidents can spread and have wider, perhaps even systemic effects, can be recognized. The centrality of, and the number of services depending on, a single platform may not be initially appreciated but need to be recognized. Similarly, communication hubs like IXPs may connect hundreds of networks and may have a wide and systemic effect.

From both a measurement and monitoring perspective and an operational perspective, learning may require that situation-specific responses (best practices; *operational*) and contingency plans (*strategic* preparedness to deal with unknown and unknowable incidents) are developed or updated, given the information that has been acquired. Here, a critical distinction must be made. On the one hand, testing to explore preventive and response strategies (e.g., via ‘red teaming’ or similar, simulation-based approaches and digital twinning, etc.) may be used to stress systems, detect vulnerabilities, and enhance the measuring and monitoring system. In addition to such proactive strategies to gather critical insights for learning and adapting systems to make them more resilient, other strategies are more reactive and responses to real-world events. It is a critical capability to learn from incidents and 'almost or near incidents' to prevent similar future incidents and adjust the infrastructure accordingly. Again, this emphasizes the role of the measuring and monitoring system and the necessity to continually update and enhance these systems.

Lastly, it is important that knowledge and information are appropriately shared within and across actors and organizations. Not only is this critical to creating awareness, but it also systematically

[8] See also the discussion on fault attribution in Lehr et al. (2011, Section 4).

enhances resilience. However, the topic of incident reporting is contentious and has been debated for a long time. Beyond incentive problems for different actors to collect relevant data and truthfully share valid insights, aspects related to security and privacy need to be considered.[9]

4. Characterizing and Measuring Internet Resilience

Given it is hard to evaluate and improve without metrics, it would be nice if there were some sort of overall 'figure of merit' for resilience—perhaps a score from 1 to 10 for a region of the Internet. Sadly, we do not believe that such a simple idea can make sense. Resilience is a measure of how well a system responds to adverse events, and adverse events differ in their character, scope, and impact. So, resilience can only be assessed in the context of a particular class of adverse events. To our knowledge there has not been a systematic attempt to construct a taxonomy of adverse event types[10], in particular as they relate to the Internet.

In the following, we look at several specific examples of adverse events to better understand the landscape of resilience. We also include illustrative examples that reveal the challenges of measuring Internet resilience.

4.1. Examples of diversity of failure modes

4.1.1. Link failures

A well-understood class of failure is the failure of a network link—the circuit connecting two packet switches. Routing protocols, both those that are used internally by an ISP and also the global routing protocol Border Gateway Protocol (BGP) that connects ISPs together, are designed to detect link failures and reconstruct viable routes. In this context, one measure of resilience is how quickly the routing protocol can detect the failure and converge on a consistent view of a new routing table. Convergence time is a design metric for routing protocols and an ongoing topic of research.

To understand the actual resilience of an ISP to a link failure, one would have to start with a map of the topology of the ISP, simulate a link failure, and then simulate the execution of the routing protocol to see what paths it finds, and how quickly. Unfortunately, that analysis is not possible with publicly available data. While an ISP does have knowledge of its internal topology, link-level

[9] See, for example, the discussion on measurement challenges and considerations with regard to (strategic) incentives for sharing and data management in Frias et al. (2023).

[10] We are aware of recent attempts in this direction. For example, the NIS Cooperation Group (2024) provided an overview of cybersecurity and resilience risks, threats, dependencies, and vulnerabilities.

public data is far from complete and thus is not enough to have informative simulations and assess the resilience of an ISP to a link failure.

Even though there is more data at the level of connections among ISPs (inter-AS routing using BGP) and a third party can identify the inter-ISP links currently being used (based on analysis of the BGP routing data that is captured by projects such as RIPE and RouteViews), the data still lacks links that are relevant to assess resiliency to link failure. Indeed, back-up links and links used by local traffic are not usually revealed, though they might become visible in the case of a failure. However, there is no guarantee that eventually that data would reveal all the paths that might be activated if a current path fails.

4.1.2. Cable cuts

Evaluating ISP resilience to events causing link failures is an even bigger challenge, as an analysis at the link layer is not enough to assess resilience to cable cuts. As mentioned in Section 2, the layered design of the Internet may hide characteristics of the layer below that impact the resiliency of the layer above. In particular, links may share the same physical infrastructure, and the cut of a cable breaks many links bundled in that cable at the same time. Indeed, cable cuts may rise to the level of geopolitics when it involves critical connections such as undersea cables.

Unfortunately, simulating what would happen with a given cable cut is not trivial as the relevant data is hard to gather. In principle, one could learn which Internet links are in a given bundle, but cables may be installed and operated by private entities, and what they carry is not disclosed. Cables may carry traffic for many networks, not just the public Internet. The topology of critical links such as undersea cables and the resilience of the Internet to potential cable cuts is a topic of current research.

Box 1: Howard Street Tunnel Fire

One way that network operators attempt to improve resiliency is to utilize redundant technologies (such as fiber optic links) that are physically disjoint. The hope is that a single failure will not disable redundant options.[11] ISPs use techniques such as "fiber swapping" with other ISPs to get circuits for their network with independent failure modes.

However, the Howard Street tunnel fire is an example of a substantial outage where this assumption of independence failed. The Howard Street tunnel is a railway tunnel under Howard Street in Baltimore. On July 18, 2001 a freight train derailed in the tunnel, causing a massive underground fire that took days to control. There were fiber optic cables in the tunnel that were destroyed by the fire. ISPs that thought they had independent fiber paths discovered that they

[11] This presumes disjoint failure modes for the redundant options.

followed the same physical path through that tunnel. One of the challenges of a highly layered architecture such as the Internet is that it can be difficult to make the connections across the layers from a fiber circuit to the bundle of fibers that include that fiber to a physical path (a conduit, tunnel, and so on) in which that bundle is installed. This fire caused multiple outages that can be attributed to an incorrect analysis of common mode failures.[12]

4.1.3. Power failures

Power failures are an interesting sort of adverse event, because every power failure potentially has a different scope and duration. A power failure impact on Internet resilience has two sides. One is the "external" impact: assuming everything within the bounds of the power failure goes down, how is the rest of the Internet affected. The other is the "internal" impact: to what extent the region affected by the power failure continues to have Internet service.

Critical infrastructure is expected to have some level of resilience to power failures that can be quantified. As an example, in the era when the landline telephone system was considered critical infrastructure, the phone companies were expected to engineer a high degree of "internal" resilience to power failures. In what has been called the POTS (plain old telephone service) era, the phone in the home was powered by electric current delivered over the phone wires from the telephone central office, which typically had battery backup, which was in turn backed up by generators. So, the measure of resilience was the degree to which this backup was in place and how much diesel fuel was stored at the central office.

The resilience of the Internet to power outages depends on the physical infrastructure supporting connectivity. Each connectivity technology has different power needs and access to backup power. The infrastructure of the cable system, in contrast to the telephone system, has smaller switching points out in the field which need to be powered. In the early days of cable deployment, these smaller locations did not have backup power.[13] That has changed but it is unclear to what extent steps are being taken to keep the cable infrastructure operational during a power failure. The same question can be asked about backup power for cell towers.

Going up the layers, services, and users can make their connection to the Internet more resilient to power failures by depending on multiple types of connections. Indeed, many residential users can access the Internet via fixed broadband and mobile Internet from phones.

---

[12] In the Howard Tunnel, having redundant cables in the Tunnel did not prevent outages since all of those cables were vulnerable to a common failure mode.

[13] One anecdote reported that the reason was the thieves kept breaking into these small, unmanned facilities and stealing the generators.

4.1.4. Monoculture failures

A monoculture failure occurs when there are many copies of the same hardware/software component in the system, and some common event causes them all to fail at the same time. A recent example of a monoculture failure was the flawed Crowdstrike software released by Microsoft,[14] which caused widespread unrelated outages. It is not clear that the Internet itself is as much at risk; normally individual ISPs test new releases before installing them in operational networks, and while one ISP might suffer an event, it would be noticed before other ISPs installed it.

This type of failure can escalate to the level of national security concerns when some foreign vendor might install a back door in their software that might, for example, cause damage to the hardware if triggered. Then, at a time of crisis, the provider triggers this event, causing widespread and long-lasting harm. The consequence of this attack to the vendor would be catastrophic; they would probably never sell another item. For this reason, the probability of this sort of malicious monoculture failure may be very low. Nonetheless, it is a cause of alarm to national security planners, who have suggested that ISPs build their networks using technology from more than one vendor. This mitigation, however, is very costly to the ISP, which must train teams to manage and oversee the two products. Here again, it is very difficult for a third party to evaluate ISP risks to monoculture failures.

4.1.5. Natural and geo-political events

There are at least two sorts of resilience that arise in the context of complex natural and geo-political events impacting Internet service: (i) the ability of the affected region to have Internet service, and (ii) the ability of non-affected regions to continue having connectivity despite the on-going event. In these cases, the resilience analysis depends on developing an external (third party) map of topology and link capacity, including different connectivity technology, and then speculating on how routing protocols are going to respond to natural events or actions by various actors.

Box 2: Hurricane Katrina

The Howard Street tunnel example (see Box 1) has shown that certain events can cause physically disjoint and redundant technologies to fail. In August 2005, Katrina hit the Gulf Coast, causing widespread and a wide range of disruptions to and failures of critical infrastructures and services—including the electricity supply and communications infrastructure (wired and wireless)—and

[14] Bloomberg, *CrowdStrike and the Global IT Outage, Explained*, https://www.bloomberg.com/news/articles/2024-07-19/crowdstrike-microsoft-it-outage-what-caused-it-what-comes-next

affecting large numbers of Internet users (see, e.g., Comfort and Haake, 2006). Physical destruction through wind and flooding caused substantial challenges to restoring Internet resilience.

Among the practical responses to quickly restore connectivity in impacted areas and limit the harm of the event were collaborative efforts among rival companies (e.g., via spectrum sharing and roaming) and between companies and the government and other local public entities (e.g., Abernathy, 2005), and also the provision of physical, on-demand connectivity solutions such as Cell on Wheels to temporarily replace or complement local mobile access infrastructure.[15] Providers typically have crisis response teams, and governments have response plans for critical infrastructures. Different natural disasters affect the Internet in different ways. However, different actors can learn lessons based on ex post evaluations to enhance resilience and response capabilities.

War and censorship by authoritative regimes are two examples of geo-political events that impact Internet service. In the case of censorship, the ability of the users in a country to resist the efforts of their own country to cut them off from the global Internet is an active topic of research. In addition, countries may assess their ability to resist efforts by another country to cut it off from the Internet. The Internet routing system provides a certain level of resilience to censorship.[16] In the U.S., the State Department is interested in these questions and has funded research on the topics. As another example, there has been a great deal of attention to how the connectivity of Russia and Ukraine has changed since the start of that conflict and how it relates to kinetic activity.

We note though that geo-political events can have different scope and different objectives, and it is yet unclear how to generalize their impact on Internet services to the more general question of resilience.

4.1.6. Operator error

While malice and natural disasters get a lot of policy attention, a common cause of failure are errors by network operators, who release a buggy update, misconfigure a routing table, and so on.[17] These failures can take many forms and have a wide range of cascading consequences and they may impact unknown system dependencies. Furthermore, given the distributed nature of the

[15] See "Carriers, contractors assessing communication restoration needs of pummeled Gulf Coast," Wireless Estimator, August 31, 2005, available at https://wirelessestimator.com/content/articles/?pagename=Hurricane%20Katrina.

[16] There is an old saying (attributed to John Gilmore) that "The Net interprets censorship as damage and routes around it."

[17] According to NIS Cooperation Group (2024), *"In general, the major incidents reported (about 160 major incidents each year, from across the EU) fall into four main categories: System failures, typically software or hardware failures (about 60% of reported incidents); Human errors (about 20% of reported incidents); Natural phenomena (about 10% of the reported incidents); Malicious actions (about 10% of the reported incidents)."*

Internet, operators mistakes of one network can impact other networks, as traffic pattern shift or network services (e.g. DNS) depend on infrastructure of other networks (see the box describing a recent incident).

One approach to mitigation is for operators to have a realistic test network into which proposed operator actions can be released. Such an approach is costly and time-consuming. However, bypassing such a pre-testing strategy and introducing solutions directly into the operational network may introduce errors that may be very hard to undo, especially if the update blocks further operator access to address faults.

An ISP's resilience to operator mistakes, whether internal or external, depends on operators' practices, some of them unique to the operation and business model of a network (e.g. many operators have developed in-house automation solutions to reduce operator errors). In addition, the Internet Engineering Task Force (IETF) publishes Best Current Practices documents that include operational practices to reduce misconfigurations and the impact of mistakes. However, even for well-supported best practices, third party verification of whether a network is following best practice recommendations is difficult. Sometimes the lack of good operational practices becomes visible through incidents, but there is usually only anecdotal evidence for the largest and most visible events.

Box 3: Cloudflare DNS resolver outage

One critical aspect to evaluate the resilience of a system is to understand the dependencies. Unfortunately, there are dependencies between the core Internet protocols that are not well understood and which cannot be easily fixed. Cloudflare's 1.1.1.1 public Domain Name System (DNS) resolver outage caused users in 70 countries to be unable to reach provides an example of dependencies between routing and DNS.

On June 27, 2024, a small network in Brazil started advertising in the routing system that it hosted the address space used by the Cloudflare DNS resolver. Most networks would not accept a route for Cloudflare's IP address blocks as Cloudflare has asserted it is the network hosting those resources in many routing databases and security systems network operators use. However, the advertisement from the Brazilian operator had specific characteristics usually used in advertisements to signal the request to blackhole traffic from Denial of Service (DoS) attacks. As such, the wrong route was accepted by many networks, including a large Tier-1 ISP, and the traffic to Cloudflare's DNS resolver was black holed in those networks, making the resolver unavailable to many users.

4.2. Responses to increase resilience

The adverse events we discuss here can to some extent be anticipated. Network operators can prepare for the anticipated events and plan what to do in those situations to minimize the impact on their customers.[18] Services running on top of the Internet can decide to increase Internet resilience by for instance using multiple providers ("multihome"), multiple types of connectivity technologies, or use overlay networks. If users or services have the knowledge and capability to invoke their own "resilience layer" in the form of an overlay network, they may be able to compensate for a lack of resilience at the lower layers, including deliberate attempts to degrade resilience by acts of censorship, politically-motivated cable cuts, and the like.

Overlay networks are a tool that can be used by different actors to try to mitigate the consequence of an adverse event. The general design of an overlay network is a group of computers on the Internet that will receive and forward packets among themselves. Between each of these computers, the path of the data is determined by the normal routing protocols, but when using an overlay network, by picking which intermediate relay computers to utilize, users can somewhat control which parts of the Internet are being exploited as the traffic is being forwarded. Overlay networks illustrate the point that what really matters when we talk about resilience is the resilience of the service (or app) the user is invoking, not just the lower-layer resilience.

Another, more complex challenge is to understand and characterize the degree to which there is resilience in the Internet to low-probability but potential high-impact events—the so-called "black swan" events. A good example of such an event was the attack of 9/11, which, while not primarily targeting the Internet, caused a wide range of failures and disruptions.

What we saw in that case is that network engineers took extraordinary steps to maintain connectivity. They changed their routing policies to carry traffic for their competitors. They actually installed new physical links in multi-tenant colocation centers to create new paths. There is no way to perform measurement to assess the extent to which operators will take such steps to maintain resilience. The only way to assess this sort of "resilience in depth" would be to interview executives from the larger ISPs to understand how they think about their obligation to ensure a resilient Internet.

4.3. Summary

One could look at additional types of adverse events and consider what resilience would mean in that context, but these examples should be sufficient to demonstrate that there can be no single metric of resilience. Resilience can only be assessed and (potentially) measured in the context of

[18] Preparation can take many forms. Provisioning for redundant capacity, operator training, and forensic analysis of past outages to learn are all useful strategies.

a particular sort of adverse event. Further, gathering the data to make an assessment (even through simulation) of actual resilience may be difficult.

5. Policy Lessons and Paths Forward

Internet resiliency is a crucial component for ensuring the resilience of our increasingly digital economies and societies. It is unsurprising that policy efforts to address resilience in the context of critical infrastructures generally and with a special focus on the digital sphere have grown in recent years in the EU and the US. Arguably, the increasing integration of 'networkable' ICT components into virtually all social and economic contexts has expanded the dependency of online access and thus the scope of cyber risks, threats, and vulnerabilities.

The COVID-19 pandemic prompted general recognition of the Internet as basic and critical infrastructure, critical to sustaining economic and social activity when the pandemic disrupted normal modes of operation, while also highlighting inequities in access and the viability of substituting online for face-to-face interactions that varied across jobs and demographics. This appreciation has led to a range of policies aimed at governing and ensuring the availability/existence of basic Internet infrastructure. In the EU and the US, ambitious connectivity goals were established and large-scale investment programs in Internet access infrastructure were launched to achieve those. Examples in the US are the Infrastructure Investment and Jobs Act[19] and the BEAD program.[20] In the EU, digital infrastructure investment by Member States was supported by the Recovery and Resilience Facility (RRF).[21] At the EU level, the EU's Digital Decade Policy Programme aims, among other things, to promote "secure and sustainable digital infrastructures" and has established ambitious EU-wide targets for fixed and mobile connectivity as well as for edge and cloud infrastructure.[22]

Distinct from this set of industrial infrastructure policies are those that focus on the security and resilience of (physical and/or digital) critical infrastructures. Many policy-related documents have offered taxonomies of problems/threats and recommendations or best practices for enhancing resilience.[23] As concerns related to cyber securityand cyber resilience have gained importance in

---

[19] See https://www.congress.gov/bill/117th-congress/house-bill/3684

[20] See https://broadbandusa.ntia.doc.gov/funding-programs/broadband-equity-access-and-deployment-bead-program

[21] See https://commission.europa.eu/business-economy-euro/economic-recovery/recovery-and-resilience-facility_en

[22] See https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/europes-digital-decade-digital-targets-2030_en

[23] For example, the 2013 Presidential Policy Directive on Critical Infrastructure Security and Resilience (The White House, 2013) formulated three strategic imperatives and defined the roles and responsibilities of different stakeholders (including governmental entities). A very comprehensive document that focuses on many aspects we have covered in this paper but covers a broader range of infrastructure contexts with a more regional focus was published in 2021 by

recent years, [24] a growing awareness and consensus of the critical role of cyber security and cyber resilience has produced a wealth of policy-related documents outlining policy objectives and recommended actions. Most of these documents, however, do not focus specifically on Internet resilience[25] or how it can be characterized, measured, and best promoted.

We argue that Internet resilience is a key component of overall cyber resilience that requires a distinct policy posture and treatment given the intrinsic characteristics of how the Internet works.[26]

The cyber-resiliency of critical infrastructures is ultimately a challenge of ensuring that the services citizens and businesses depend on are available and perform as expected during normal times, and in abnormal times (in the face of natural or manmade disasters or threats), are either still able to function or quickly restored. For most people, the resiliency of infrastructures is

---

the Cybersecurity and Infrastructure Security Agency (CISA, 2021). PCAST (2024) focuses on cyber-physical resilience, emphasizing the growing fusion of physical and cyber (or virtual) infrastructures and domains. The White House (2024a) offers comprehensive insights into the security and resilience of critical infrastructures in the US. The Memorandum offers policy principles and objectives and also explains the roles and responsibilities of different stakeholders. The White House published its *Report On The Cybersecurity Posture Of The United States* in 2024. The report emphasized a shift in posture from being reactive to being more proactive, particularly since *"a reactive posture cannot keep pace with fast-evolving cyber threats and a dynamic technology landscape, and that aspiring just to manage the worst effects of cyber incidents is no longer sufficient to ensure our national security, economic prosperity, and democratic values"* (The White House, 2024b, p. ii). While the focus of the report differs from ours, it offers an overview of related US policy actions. In the EU, *Council Directive 2008/114/EC of 8 December 2008 on the identification and designation of European critical infrastructures and the assessment of the need to improve their protection* presented an early effort to consider relevant policies. A recent example from the EU is *Directive (EU) 2022/2557 of the European Parliament and of the Council of 14 December 2022 on the resilience of critical entities*.

[24] See Footnote 23. Moreover, the EU has established a revised cyber security strategy in 2020 (see https://digital-strategy.ec.europa.eu/en/policies/cybersecurity-strategy) and has established several cyber security policies (https://digital-strategy.ec.europa.eu/en/policies/cybersecurity-policies). For example, in 2019, the Cyber Security Act (*Regulation (EU) 2019/881*) entered into force. *Directive (EU) 2022/2555* (NIS2 Directive) updates the 2016 EU cyber security law (NIS1) and establishes "*legal measures to boost the overall level of cybersecurity in the EU"* (see https://digital-strategy.ec.europa.eu/en/policies/nis2-directive#:~:text=The%20Directive%20on%20measures%20for,them%20to%20be%20appropriately%20equipped). Moreover, the proposed EU Cyber Solidarity Act that is aimed to establish a cyber security alert, emergency, and incident review mechanisms *"to reinforce the EU's solidarity and coordinated actions to detect, prepare and effectively respond to growing cybersecurity threats and incidents"* (EC, 2024). The EU's proposed Cyber Resilience Act (*Proposal for a Regulation of the European Parliament and of the Council on horizontal cybersecurity requirements for products with digital elements and amending Regulation (EU) 2019/1020, COM/2022/454 final*) which has been approved by the European Parliament, is focused on cybersecurity requirements for networkable products and services. An overview of the EU activity is provided in NIS Cooperation Group (2024, pp. 5-8).

[25] One report from 2018 (Polk, 2018) explicitly focused on Internet resilience, it did so by more narrowly focusing on a subset of cyber-related threats. We detected the largest overlap with our paper with the NIS Cooperation Group (2024) publication, which focuses on the "*Cybersecurity and resiliency of Europe's communications infrastructures and networks."* The document offers an overview of resilience risks, threats, dependencies, and vulnerabilities, as well as of the EU's major policy actions and tools to promote cyber security and resilience—it also offers recommendations.

[26] For example, PCAST (2024) states: "*Cyber-physical resilience, based on a marrying of cybersecurity, resilience, reliability, and recoverability in information systems, critical infrastructure, and operational technology, is vital to our societal functioning."*

something that is taken for granted and the concern of other parties – only noticed when the desired service is disrupted. The services that depend on the Internet and other digital networks and infrastructures are multifaceted, and the Internet's contribution to overall cyber-resiliency is complex. Furthermore, the Internet depends on the ICT infrastructure.

At a fundamental level, our paper emphasizes the necessity as well as the complexity of characterizing and measuring Internet resilience. Any policy assessment of Internet resiliency should account for the jointly, and at times, separately evolving layers from critical infrastructure services to the networked fabric of ICT resources that enable those services. In addition, the assessment should evaluate the technical and economic intra- and cross-layer dependencies This is critical for understanding the role of Internet resilience for overall resilience, resistance, and customer experience, as well as for anticipating and evaluating related risk and harm scenarios. Furthermore, the evolving role of the public Internet and the changing boundaries between public and private Internet must be factored into any assessment of dependencies and resilience.[27] However, the essence of the Internet has not changed. Whereas the technologies, industry structure, and economic relevance of the Internet has changed significantly since the 1970s, certain core features that have contributed to its emergence as critical infrastructure remain relevant (see Lehr, Clark et al., 2019a).

Policymakers need to understand what the Internet's basic design features and the related complexity mean for Internet resilience. The decentralized nature of the Internet is core to Internet resilience. In contrast to many other critical infrastructures like roads, bridges, and electricity, the Internet is characterized by a diversity of economic and technical controls. These controls are decentralized and distributed among actors with imperfectly aligned incentives. Efforts to enhance Internet resilience cannot be based on a globally centralized approach to control. Instead, policies should preserve the decentralized and distributed control of the Internet. Other critical components at layers below or above may depend on (more) specialized treatment, which in many cases may require closed (or at least less open) systems and central control.[28]

The layered architecture of the Internet and the services running on top of it also impact Internet resilience. On the one hand, the Internet's layered architecture enables resiliency to be composed up the layers. For example, if an end-user has a critical need to support some activity, they can take actions at their layer that attempt to insulate them from failures of resilience at lower layers.

---

[27] See Lehr, Clark et al. (2019b) for how Internet fits into ecosystem of digital platforms.

[28] For example, much of the physical infrastructure may be controllable by design at a sovereign level (national), while the Internet fundamentally is global (international). This is important since much of the digital infrastructure is primarily physical (FTTH, wireless base stations, satellite networks). However, this may change as increasingly much of that is software-based. As a consequence, the boundary between physical and digital and between hardware (physical) and software is blurred, which may have strong implications for control. In terms of specialized treatment, different applications have different security and QoS requirements.

On the other hand, distributed control and various cross-layer and inter-organizational dependencies may lead to complex coordination problems that may impede Internet resilience.

Although the Internet's decentralized and global character and its resistance to a single metric or standard for Internet resilience complicates sovereign or regional policies to promote Internet resilience, that does not mean policy cannot help. National digital infrastructure reliability and resiliency policies come in multiple forms and how these are promoted can help advance (or hamper) Internet resiliency. For example, national universal service policies may promote diverse access options for households[29] (which helps ensure that there are widely available redundant access links for Internet service. Competition policies can seek to ensure that bottlenecks are not monopolized, and where they arise that essential resources are shared.[30] Similarly, national policies that enhance the resiliency of the physical networks that support the Internet can contribute to regional Internet resiliency. Regionally, the Internet may sustain services with a higher expected level of performance and be free from outages than may be feasible globally. For example, wealthy nations may invest in higher-quality and more expensive infrastructure such as Fiber to the Home (FTTH) that is not affordable in other countries, and yet the Internet's rate-adaptive capabilities can still support global communications in the face of local outages and regional disparities in performance. However, efforts to define an aggressive minimal threshold for Internet service is not helpful since it would foreclose technologies, services, and networking options that might otherwise be available to support diverse Internet routing. In telephony networks, operators talked about achieving "5 9's" reliability (i.e., 99.999%) availability as a design goal for core network elements like tandem switches. In the Internet, it would be meaningless to set such a standard since its resiliency comes from the absence of such critical nodes.

Those regional disparities in performance may also be due to disparities in the extent to which ISPs have adopted best practices as identified by the IETF and other Internet consensus bodies. As a counterpoint to options for government regulation, we suggest that voluntary industry standards ought to play a crucial role in maintaining and enhancing Internet resilience. Government policies can also assist in addressing such problems by using their power to nudge markets toward accelerated adoption of those best practices. For example, governments may use their monopsony power to require support for more resilient routing protocols or adoption of best-practice disclosure for entities doing business with government networks. In such ways, local governments might nudge regions of the Internet toward greater resiliency. Better public information about outages can assist in forensic analysis and learning to help make the Internet more resilient over time. For complex systems like the Internet, for which the total state space of possible failure modes cannot be known or anticipated, continuous learning from past problems can help build Internet resiliency.

---

[29] An example of such a policy would seek to ensure all consumers have access to both fixed/wired and mobile wireless broadband access networks.

[30] See Lehr & Sicker (2019).

Lastly, even though resilience needs to be evaluated from a technical and economic perspective, good technical measurements are necessary before the economic consequences can be measured. However, even there, we see fundamental measurement challenges, as well as incentive alignment and other (e.g., security) issues related to information exchange between and coordinated actions across multiple actors (see Sections 3 and 4). Economic measurement issues arise even in the first stage since the design of technical metrics will impact what economic quantums are observable, and economic incentive alignment considerations arise at all stages since any measurements that matter for business and policy are inherently strategic. Given the relevance of social aspects in preparing, limiting, and recovering, and learning to achieve resilience, Internet resilience can be conceptualized as a socio-technical problem, including organization management. Interviewing key actors to measure the level of planning, steps they have taken towards resilience, and other organizational aspects would provide a better understanding and crucial insight into Internet resilience. Thus, assessing Internet resiliency is fundamentally a joint technical and economic challenge.

6. Conclusions & Future Directions

Ensuring the resiliency of critical infrastructures has long been a focus of governments, and with the transition to an increasingly digital society and economy, those include networked computing and communication networks. The Internet is a key component of that fabric of infrastructures, and is a key element of cyber-resilience. Understanding how to assess and promote Internet resilience are important policy challenges.

In this paper, we explain how some of the key features that differentiate the Internet from other digital networks are essential in considering its resiliency. Those include that the Internet is a general purpose network, which is decentralized in its economic control, operational management, and technical design and implementation. The Internet is a global network of networks that includes both wired and wireless technologies operated by a wide diversity of entities with diverse economic interests and resources, spanning the globe. This generality and decentralized techno-economic structure enhances the resilience of the Internet, reducing single points of failure.

The layered nature of the Internet architecture and the simplicity of its core transport protocols facilitate the Internet's generality and decentralized organization. This contributes to the Internet's flexibility and adaptability and ability to sustain its core functionality – enabling end-to-end packet transport, where other digital networking infrastructures cannot. These attributes facilitate the ability of the Internet to recover rapidly from outages of many types and be resistant to network-wide outages.

At the same time, however, the Internet's key features make it difficult to establish and guarantee minimum quality-of-service performance guarantees since those needed to satisfy resource intensive applications would be too expensive for many applications that also rely on the Internet. Additionally, the diversity, distributed and decentralized economic interests that support the Internet make it challenging to coordinate responses to outages

Although no centralized regulatory entity or sovereign nature can control the Internet, there is significant scope for policymakers to promote policies that will contribute to local and regional resiliency, and ultimately, to the resiliency of the Internet globally. Policies to promote more diverse, secure, affordable, and competitive digital network infrastructure (cable and wireless networking) that support the Internet contribute to its regional and local resiliency. Governments can also help promote the adoption of best-practices and facilitate coordination within the Internet through transparency regulations, active measurement initiatives, and use of its monopsony power.

Although there is no single best metric to assess Internet resiliency, building a strong basis for identifying, tracking, and learning from the many sources of threats and actual outages impacting Internet performance continues to prove critical. Furthermore, complementing that information with insights from interviews of key actors about the level of planning, steps taken towards resilience, and other organizational aspects would provide crucial insight into the state of Internet resilience.

References

Abernathy, K. (2005). "Statement of Commissioner Kathleen Q. Abernathy," FCC Commissioner, available at https://docs.fcc.gov/public/attachments/DOC-261097A1.pdf

CISA (2021). Methodology for Assessing Regional Resilience: Lessons Learned from the Regional Resiliency Assessment. https://www.cisa.gov/sites/default/files/publications/DIS_DHS_Methodology_Report_ISD%2520EAD%2520Signed_with%2520alt-text_0.pdf (Accessed August 26, 2024)

Claffy, K. C., & Clark, D. (2014). Platform models for sustainable Internet regulation. *Journal of Information Policy*, *4*, 463-488.

Comfort, L. K., & Haase, T. W. (2006). Communication, coherence, and collective action: The impact of Hurricane Katrina on communications infrastructure. *Public Works management & policy*, *10*(4), 328-343.

Duchek, S. (2020). Organizational resilience: a capability-based conceptualization. *Business research*, *13*(1), 215-246.

EC (2024). EU Cyber Solidarity Act Factsheet. https://digital-strategy.ec.europa.eu/en/library/eu-cyber-solidarity-act-factsheet

Feldmann, A., Gasser, O., Lichtblau, F., Pujol, E., Poese, I., Dietzel, C., ... & Smaragdakis, G. (2020, October). The lockdown effect: Implications of the COVID-19 pandemic on internet traffic. In *Proceedings of the ACM internet measurement conference* (pp. 1-18).

Frias, Z., W. Lehr, V. Stocker, and L. Mendo (2023), "Measuring NextGen Mobile Broadband: Challenges and Research Agenda for Policymaking", 32nd European International Telecommunications Society Conference (EuroITS2023), June 19-20, Madrid, available at https://www.econstor.eu/handle/10419/277959

Lehr, W., Heikkinen, M., Clark, D. D., & Bauer, S. (2011, September). Assessing broadband reliability: Measurement and policy challenges. In *Research Conference on Communication, Information and Internet Policy (TPRC)*. https://ssrn.com/abstract=1979746

Lehr, W., Clark, D., Bauer, S., Berger, A., & Richter, P. (2019a). Whither the public Internet?. *Journal of Information Policy*, *9*, 1-42.

Lehr, W., Clark, D. D., Bauer, S., & Claffy, K. C. (2019b). Regulation when platforms are layered. Available at SSRN: https://ssrn.com/abstract=3427499 or http://dx.doi.org/10.2139/ssrn.3427499

Lehr, W. and D. Sicker (2019), "Telecom Déjà vu: a model for sharing in the broadband Internet," *Information & Communications Technology Law*, 1-36. doi:10.1080/13600834.2019.1653546.

NIS Cooperation Group (2024). *Cybersecurity and resiliency of Europe's communications infrastructures and networks*. https://digital-strategy.ec.europa.eu/en/library/report-cybersecurity-and-resiliency-eu-communications-infrastructures-and-networks

Polk, W. (2018), A Report to the President on Enhancing the Resilience of the Internet and Communications Ecosystem Against Botnets and Other Automated, Distributed Threats, Other, National Institute of Standards and Technology, Gaithersburg, MD, [online], https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=925900 (Accessed August 26, 2024)

PCAST (2024). US President's Council of "Strategy for Cyber-Physical Resilience: Fortifying Our Critical Infrastructure for a Digital World" https://www.whitehouse.gov/wp-content/uploads/2024/02/PCAST_Cyber-Physical-Resilience-Report_Feb2024.pdf

Stocker, V., Lehr, W., & Smaragdakis, G. (2023). COVID-19 and the Internet: Lessons learned. In *Beyond the Pandemic? Exploring the Impact of COVID-19 on Telecommunications and the Internet* (pp. 17-69). Emerald Publishing Limited.

The White House (2013). Presidential Policy Directive -- Critical Infrastructure Security and Resilience. PRESIDENTIAL POLICY DIRECTIVE/PPD-21. https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil

The White House (2024a). National Security Memorandum on Critical Infrastructure Security and Resilience, https://www.whitehouse.gov/briefing-room/presidential-actions/2024/04/30/national-security-memorandum-on-critical-infrastructure-security-and-resilience/

The White House (2024b). 2024 Report On The Cybersecurity Posture Of The United States. https://www.whitehouse.gov/wp-content/uploads/2024/05/2024-Report-on-the-Cybersecurity-Posture-of-the-United-States.pdf