EuroDIG
European Dialogue on Internet Governance

European Dialogue on Internet Governance

50 years of the Internet, how it works and how we can protect it

Internet Society

The opinions, recommendations, conclusions and presentation of facts in this publication reflect the authors' views. EuroDIG Messages that reflect the consensus of the EuroDIG stakeholder community are included alongside the authors' text in this publication.

European Dialogue on Internet Governance

50 years of the Internet, how it works and how we can protect it

Internet Society

About EuroDIG

Launched in 2008, EuroDIG, the European Dialogue on Internet Governance, is a unique annual event that brings together Internet stakeholders from throughout Europe (and beyond), and from across the spectrum of government, industry, civil society, academia and the technical community. Stakeholders and participants work over the course of each year to develop, in a bottom-up fashion, a dynamic agenda that explores the pressing issues surrounding how we develop, use, regulate and govern the Internet. EuroDIG participants come away with broader, more informed perspectives on these issues and new partners in responding to the challenges of the information society.

Contents

Foreword: On the Past, Present and Future of the Internet 7

Introduction 11

How the Internet works: origins of a packet-switched network 13
- The first steps 13
- The foundational Internet concepts 15
- The expansion of the Internet 19

The Internet is a set of practices: the Critical Properties 21
- Critical Property 1: An Accessible Infrastructure with a Common Protocol 22
- Critical Property 2: An Open Architecture of Interoperable and Reusable Building Blocks 24
- Critical Property 3: Decentralized Management and a Single Distributed Routing System 25
- Critical Property 4: Common Global Identifiers 26
- Critical Property 5: A Technology Neutral, General-Purpose Network 28

The Enablers of an Open, Globally Connected, Secure and Trustworthy Internet 31
- Supporting an Open Internet 31
  - Enabler: Easy and unrestricted access 32
  - Enabler: Unrestricted use and deployment of Internet technologies 36
  - Enabler: Collaborative development, management, and governance 37
- Supporting a Globally Connected Internet 39
  - Enabler: Unrestricted reachability 39
  - Enabler: Available capacity 40

Supporting a Secure Internet 41
Enabler: Data confidentiality of information, devices, and applications 42
Enabler: Integrity of information, applications, and services 43
Supporting a Trustworthy Internet 44
Enabler: Reliability, resilience, and availability 44
Enabler: Accountability 46
Enabler: Privacy 47

Compiling all together: the Internet Impact Assessment Toolkit 50

The Internet is under threats: what can the community do 52
Identified types of threats 53
Threat 1: Regulating Business Relationships 53
Threat 2: National Internet Gateways 55
Threat 3: Creating Walled Gardens 56
Threat 4: Regulation of DNS Infrastructure 58
Threat 5: Blocking Security Technologies 60
Threat 6: Digital Sovereignty 63
Threat 7: Internet Shutdowns 66
Threat 8: Limiting Global Access 68

Foreword: On the Past, Present and Future of the Internet

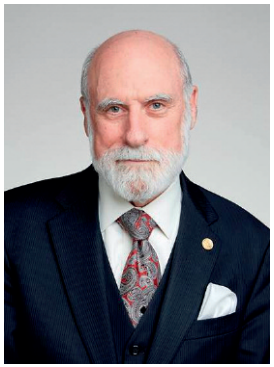

Vinton G. Cerf, Vice president and Chief Internet Evangelist for Google

This year, we celebrated the 50th anniversary of the first publication that described how the Internet could work[1]. Building on the success of the Arpanet project[2], the Internet was made operational on January 1, 1983. In the 1980s, optical fiber networks brought and continues to bring massive communication capacity to the Internet. In the US, in the mid-1980s, additional backbones were built by NSF (NSFNET), NASA (NASA Science Internet) and DOE (ESNET). Internet became publicly available in the US in 1989 with the arrival of UUNET, PSINET and CERFNET and the World Wide Web was announced in December 1991. The MOSAIC browser arrived in 1993, leading to the Netscape Navigator whose developer, Netscape Communications, went public in 1995 triggering the "dot-boom". A "dot-bust" arrived in spring 2000 but the Internet survived and thrived as Internet Service Providers expanded access to the Internet's web servers.

For lack of space, this essay neglects the pre-history of social networks in the form of Bulletin Board Systems in the late 1970s and 1980s, USENET in 1979, Friendster, Myspace, Geocities and systems such as America Online, Prodigy and Compuserve, among others. Pre-WWW Internet search started with the FTP File Search engine, Archie, in 1986. Gopher[3] arrived in 1991 but was rapidly superseded by Web

search engines, beginning with Yahoo! in 1994 and Alta Vista in 1995. Google arrived with its search engine in 1998 and Microsoft launched MSN Search that same year, eventually replacing it and its successors with BING in 2009.

Streaming media starts with a performance by the Severe Tire Damage Band in June 1993 and evolved rapidly, especially with the arrival of YouTube in 2005 and Netflix' introduction of its streaming video services in 2007. A trend was born and reached fruition during the COVID pandemic when so much of everything was done online, real-time, with video conferencing.

The World Summit on the Information Society held its two conferences in 2003 and 2005 and spawned the Internet Governance Forum (IGF) which met first, in Athens, in 2006 and has met yearly since then and beginning in 2008 has grown to include about 75 youth, regional and national IGFs. In 2007, the iPhone triggered a new boom and the Internet and the World Wide Web applications became accessible to anyone with a smart phone and mobile radio (or WiFi) access.

By the mid-2000s and 2010s, a wide range of social media systems were in operation including Twitter (now X), Facebook (now META), YouTube, Baidu, WeChat, Instagram, SnapChat, Telegram, Pinterest, WhatsApp, among many others.

Looking back from 2024, it becomes apparent that the powerful amplifying effects of the Internet and its applications enabled countless new business applications, information services, consumer applications, government service access, educational and entertainment services and myriad other applications. There are literally millions of apps (applications) available for smart phones. What also became very apparent, is that this diverse, global platform was openly accessible to people with harmful or criminal intent. These rotten apples in the Internet barrel have been causing trouble since the Morris Worm in 1988[4]. All of the amplifying power of the Internet is used by bad actors to launch denial of service attacks, malware distribution, misinformation, disinformation, financial fraud, bullying, revenge pornography and a lengthy list of other harms.

These harmful behaviors have triggered government regulatory proposals including the Budapest Convention[5] and the more recent Cybercrime Convention[6]. In Europe, a variety of regulations have been introduced that have had extraterritorial impact including the General Data Protection Regulation (GDPR)[7], the European Digital Markets Act[8] and Digital Services Act[9], EU Network and Information Security Directive (NIS2)[10] among others. Some of these initiatives are intended to be ratified in multilateral settings. Efforts to curtail harm-

ful behavior and to hold parties accountable are increasing. China has a well-developed "Great Firewall" and associated surveillance capacity to control behavior considered unacceptable to the Chinese regime. Other countries shut the Internet down during elections and at will in times of unrest. The attempt to control misbehavior (by varying definitions) is palpable but often misguided. These efforts, while well intended for the most part, are often ineffective and sometimes prevent legitimate usage. Nonetheless, it can be anticipated that governments will continue to try to control what they consider to be harmful behavior.

Looking towards the future, the Internet will demand increasing cyber savvy of its users and the provision of tools to protect privacy and safety online. The threats are exacerbated by the rapid evolution of so-called Artificial Intelligence tools including Machine Learning and Large Language Models (LLMs). These are powerful enablers for the production of text, software, imagery, video and audio. There are concerns about the absorption of intellectual property in the training of LLMs and subsequent production of derivative works without compensation or citation. It seems likely that concepts like an Internet Drivers' License will emerge in the form of training in schools to learn critical thinking skills and to reject various attacks including "fake news", phishing scams, misrepresentation, character assassination and a host of other ills.

All is not lost, however. The same technology that enables harmful behavior also can amplify beneficial efforts across a waterfront of applications from health diagnosis and care to the creation of new industries. The power of multilayer neural networks and large language models already enable speech recognition, production, translation and increasingly effective artificial intelligence agents that can take actions on behalf of users from scheduling airline flights to ordering meals and offering turn-by-turn navigation in city streets. These tools are enhancing collaborations among researchers, helping to analyze massive amounts of data, aiding in the discovery of new medical treatments, managing self-driving vehicles, and generally providing superhuman capabilities especially for large scale, data driven tasks.

What might we expect in the future? Much improved diagnostic tools for personalized medicine, taking into account DNA and medical histories. Improved transportation traffic flow in city streets and in the air. Drone delivery of goods to doorsteps. Development of genetically modified plants well suited to higher temperatures and less water concomitant with global warming. Better weather pre-

dictions including information about flooding and forest fires. Development of non-invasive diagnostics including miniature, steerable, swallowable video “submarines” to navigate the human body. Commercial space exploration and the deployment of an interplanetary Internet to support commercial applications on the Moon, in the asteroid belt and beyond. Optical and neuromuscular implants to restore sight and movement. Sensemaking of non-human signals and perhaps “conversation” between human and non-human species. General Artificial Intelligence remains an aspiration but we will get closer to this as new concepts of cognition are advanced. I remain optimistic about the future of the Internet and the software that animates it. The best is yet to come.

References

1 https://www.cs.princeton.edu/courses/archive/fall06/cos561/papers/cerf74.pdf

2 https://en.wikipedia.org/wiki/ARPANET

3 https://en.wikipedia.org/wiki/Gopher_(protocol)

4 https://en.wikipedia.org/wiki/Morris_worm

5 https://www.coe.int/en/web/cybercrime/the-budapest-convention

6 https://www.unodc.org/unodc/en/frontpage/2024/August/united-nations_-member-states-finalize-a-new-cybercrime-convention.html

7 https://gdpr-info.eu/

8 https://en.wikipedia.org/wiki/Digital_Markets_Act

9 https://en.wikipedia.org/wiki/Digital_Services_Act

10 https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2021)689333

Introduction

by Sandra Hoferichter, EuroDIG and David Frautschy, Internet Society

The Internet has fundamentally changed the world as we know it over the past five decades. What once began as an experimental network is today the backbone of our globalised society, a central platform for communication, innovation and cultural exchange. But this development also comes with enormous challenges: the questions of digital sovereignty, privacy, security and access are more relevant today than ever.

Since its inception, EuroDIG, the "European Dialogue on Internet Governance", has seen itself as an open platform that brings together people and organisations from all over Europe to discuss these crucial issues. In this anniversary year of the Internet, we see it as our responsibility to contribute to the reflection and discussion of the past, present and future of this technology.

With this publication we would like to honour not only the impressive development of the Internet, but also provide food for thought on current threats that are concerning, and how our community can help on the design of its next decades. Our community has learned over the years that the best solutions are when all voices are heard: from civil society, business, governments and the technical community. EuroDIG has led this principle of the multistakeholder approach from the beginning and it remains our foundation.

"50 years of the Internet" is not only an occasion to celebrate, but also an opportunity to pause and reflect on current trends, and also to look forward. For this, the first step is to understand how the Internet works. Only then we will be able to identify ongoing developments – policy proposals and corporate decisions – that pose a threat to the Internet.

This book equips the community with the tools to protect and defend the Internet, advocating for an Internet that is inclusive, safe, and sustainable for all. By recognising the Critical Properties and Enablers that make the Internet exist and thrive, we set up a methodology that can help assess the impacts of policy trends. It starts by grasping the challenge ahead, understanding the information that is required, gathering a team with the necessary skills, and launching a process that may have its difficulties – Internet Society staff are ready to help on that – but that will be rewarding: the community keeping a healthy Internet for future generations. This

book wants to stimulate our community to act on these fronts.

We would like to thank all the authors from the Internet Society who contributed to this publication and to all those who, through their ideas, contributions and discussions, have made EuroDIG what it is today: an actor actively shaping the debate about the Internet of the future.

Let's work together to shape the next steps to keep the Internet open and make it a secure place which everyone can access.

How the Internet works: origins of a packet-switched network

The first steps

Back in in August 1962, J. C. R. Licklider of the Massachusetts Institute of Technology (MIT) envisioned a globally interconnected set of computers through which everyone could quickly access data and programs from any site. In spirit, the concept was very much like the Internet of today. Licklider was the first head of the computer research program at DARPA, the Defense Advanced Research Projects Agency, the research and development agency of the United States Department of Defense responsible for the development of emerging technologies for use by the military.

Leonard Kleinrock at (MIT) published the first book on the subject of packet switching theory in 1964, "Communication Nets: Stochastic Message Flow and Delay". Kleinrock was convinced of the theoretical feasibility of communications using *packets rather than circuits*, which was a major step along the path towards computer networking. The other key step was to make the computers talk together. To explore this, in 1965 the TX-2 computer in Massachusetts was connected to the Q-32 in California with a low speed dial-up telephone line creating the first – however small – wide-area computer network ever built. The result of this experiment was the realization that the time-shared computers could work well together, running programs and retrieving data as necessary on the remote machine, but that *the circuit switched telephone system was totally inadequate for the job*. Kleinrock's conviction of the need for packet switching was confirmed.

In late 1960s, it happened that separate works at MIT (1961 – 1967), at RAND in California (1962 – 1965), and at the National Physical Laboratory (NPL) in London (1964 – 1967) had all proceeded in parallel without any of the researchers knowing about the other work. The word "packet" was adopted from the work at NPL and the proposed line speed to be used in

the ARPANET design, established in DARPA, was upgraded from 2.4 kbps to 50 kbps.

In August 1968, after Roberts and the DARPA funded community had refined the overall structure and specifications for the ARPANET, an RFQ was released by DARPA for the development of one of the key components, the packet switches called Interface Message Processors (IMPs). The RFQ was won in December 1968 by a group headed by Frank Heart at Bolt Beranek and Newman (BBN). As the BBN team worked on the IMP's with Bob Kahn playing a major role in the overall ARPANET architectural design, the network topology and economics were designed and optimized by Roberts working with Howard Frank and his team at Network Analysis Corporation, and the network measurement system was prepared by Kleinrock's team at University of California, Los Angeles (UCLA).

Due to Kleinrock's early development of packet switching theory and his focus on analysis, design and measurement, his Network Measurement Center at UCLA was selected to be the first node on the ARPANET. All this came together in September 1969 when the company BBN installed the first IMP at UCLA and the first host computer was connected. The Stanford Research Institute (SRI) provided a second node and also supported the Network Information Center, including functions such as maintaining tables of host name to address mapping.

One month later, when SRI was connected to the ARPANET, the first host-to-host message was sent from Kleinrock's laboratory to SRI. Two more nodes were added at UC Santa Barbara and University of Utah. Thus, by the end of 1969, four host computers were connected together into the initial ARPANET, and the budding Internet was off the ground. Even at this early stage, it should be noted that the networking research incorporated both work on the underlying network and work on how to utilize the network.

Computers were added quickly to the ARPANET during the following years, and work proceeded on completing a functionally host-to-host protocol and other network software. In December 1970 the Network Working Group (NWG) finished the initial ARPANET host-to-host protocol, called the Network Control Protocol (NCP). As the ARPANET sites completed implementing NCP during the period 1971 – 1972, the network users finally could begin to develop applications.

In October 1972, Kahn organized a large, very successful demonstration of the ARPANET at the International Computer Communication Conference (ICCC). This was the first public demonstration of this new network technology. It was also in 1972 that the initial "hot"

application, electronic mail, was introduced. In March, Ray Tomlinson at BBN wrote the basic email message send and read software, motivated by the need of the ARPANET developers for an easy coordination mechanism. In July, Roberts expanded its utility by writing the first email utility program to list, selectively read, file, forward, and respond to messages. From there email took off as the largest network application for over a decade.

The foundational Internet concepts

The original ARPANET grew into the Internet, which was based on the idea that it would be composed of multiple independent networks of rather *arbitrary* design. The Internet as we now know embodies a key underlying technical idea, namely that of **open architecture** networking. In this approach, the choice of any individual network technology was not dictated by a particular network architecture but rather could be selected freely by a provider and made to interwork with the other networks through a meta-level "Internetworking Architecture". Up until that time there was only one general method for federating networks. This was the traditional circuit switching method where networks would interconnect at the circuit level, passing individual bits on a synchronous basis along a portion of an end-to-end circuit between a pair of end locations. Recall that Kleinrock had shown in 1961 that packet switching was a more efficient switching method. Along with packet switching, special purpose interconnection arrangements between networks were another possibility. While there were other limited ways to interconnect different networks, they required that one be used as a component of the other, rather than acting as a peer of the other in offering end-to-end service.

In an open-architecture network, the individual networks may be separately designed and developed and each may have its own unique interface which it may offer to users and/or other providers. including other Internet providers. Each network can be designed in accordance with the specific environment and user requirements of that network. There are generally no constraints on the types of

network that can be included or on their geographic scope.

The idea of open-architecture networking was first introduced by Kahn shortly after having arrived at DARPA in 1972. This work was originally part of the packet radio program, but subsequently became a separate program in its own right. At the time, the program was called "Internetting". Key to making the packet radio system work was a reliable end-end protocol that could maintain effective communication in the face of jamming and other radio interference, or withstand intermittent blackout such as caused by being in a tunnel or blocked by the local terrain. Kahn first contemplated developing a protocol local only to the packet radio network, since that would avoid having to deal with the multitude of different operating systems, and continuing to use NCP.

However, NCP did not have the ability to address networks (and machines) further downstream than a destination IMP on the ARPANET and thus some change to NCP would also be required. NCP relied on ARPANET to provide end-to-end reliability. If any packets were lost, the protocol (and presumably any applications it supported) would come to a grinding halt. In this model NCP had no end-end host error control, since the ARPANET was to be the only network in existence and it would be so reliable that no error control would be required on the part of the hosts. Thus, Kahn decided to develop a new version of the protocol which could meet the needs of an open-architecture network environment. This protocol would eventually be called the Transmission Control Protocol/Internet Protocol (TCP/IP). While NCP tended to act like a device driver, the new protocol would be more like a communications protocol.

Four ground rules were critical to Kahn's early thinking:

- Each distinct network would have to stand on its own and no internal changes could be required to any such network to connect it to the Internet.
- Communications would be on a best effort basis. If a packet didn't make it to the final destination, it would shortly be retransmitted from the source.
- Black boxes would be used to connect the networks; these would later be called gateways and routers. There would be **no information retained** by the gateways about the individual flows of packets passing through them, thereby keeping them simple and avoiding complicated adaptation and recovery from various failure modes.
- There would be **no global control** at the operations level.

Other key issues that needed to be addressed were:

- Algorithms to prevent lost packets from permanently disabling communications and enabling them to be successfully retransmitted from the source.
- Providing for host-to-host "pipelining" so that multiple packets could be enroute from source to destination at the discretion of the participating hosts, if the intermediate networks allowed it.
- Gateway functions to allow it to forward packets appropriately. This included interpreting Internet Protocol (IP) headers for routing, handling interfaces, breaking packets into smaller pieces if necessary, etc.
- The need for end-to-end checksums, reassembly of packets from fragments and detection of duplicates, if any.
- The need for **global addressing**
- Techniques for host-to-host flow control.
- Interfacing with the various operating systems
- There were also other concerns, such as implementation efficiency, internetwork performance, but these were secondary considerations at first.

Kahn began work on a communications-oriented set of operating system principles while at BBN and documented some of his early thoughts in an internal BBN memorandum entitled "Communications Principles for Operating Systems". At this point he realized it would be necessary to learn the implementation details of each operating system to have a chance to embed any new protocols in an efficient way. Thus, in the spring of 1973, after starting the internetting effort, he asked Vint Cerf to work with him on the detailed design of the protocol. Cerf had been intimately involved in the original NCP design and development and already had the knowledge about interfacing to existing operating systems. So armed with Kahn's architectural approach to the communications side and with Cerf's NCP experience, they teamed up to spell out the details of what became TCP/IP.

The give and take was highly productive and the first written version of the resulting approach was distributed as INWG#39 at a special meeting of the International Network Working Group (INWG) at Sussex University in September 1973. Subsequently a refined version was published in 1974. The INWG was created at the October 1972 International Computer Communications Conference organized by Bob Kahn, et al, and Cerf was invited to chair this group.

Some basic approaches emerged from this collaboration between Kahn and Cerf:

- Communication between two processes would logically consist of a very long stream of bytes (they called them octets). The position of any octet in the stream would be used to identify it.
- Flow control would be done by using sliding windows and acknowledgments (acks). The destination could select when to acknowledge and each ack returned would be cumulative for all packets received to that point.
- It was left open as to exactly how the source and destination would agree on the parameters of the windowing to be used. Defaults were used initially.
- Although Ethernet was under development at Xerox PARC at that time, the proliferation of LANs were not envisioned at the time, much less PCs and workstations. The original model was national level networks like ARPANET of which only a relatively small number were expected to exist. Thus a 32 bit IP address was used of which the first 8 bits signified the network and the remaining 24 bits designated the host on that network. This assumption, that 256 networks would be sufficient for the foreseeable future, was clearly in need of reconsideration when LANs began to appear in the late 1970s.

A major initial motivation for both the ARPANET and the Internet was resource sharing – for example allowing users on the packet radio networks to access the time sharing systems attached to the ARPANET. Connecting the two together was far more economical that duplicating these very expensive computers. However, while file transfer and remote login (Telnet) were very important applications, electronic mail has probably had the most significant impact of the innovations from that era. Email provided a new model of how people could communicate with each other, and changed the nature of collaboration, first in the building of the Internet itself, and later for much of society.

There were other applications proposed in the early days of the Internet, including packet based voice communication (the precursor of Internet telephony), various models of file and disk sharing, and early “worm” programs that showed the concept of agents (and, of course, viruses). A key concept of the Internet is that it was **not designed for just one application, but as a general infrastructure on which new applications could be conceived**, as illustrated later by the emergence of the World Wide Web. It is the general purpose nature of the service provided by TCP and IP that makes this possible.

The expansion of the Internet

The early implementations of TCP were done for large time-sharing systems, but when desktop computers first appeared, it was thought by some that TCP was too big and complex to run on a personal computer. A compact and simple implementation of TCP was designed, and it was fully interoperable with other TCPs, but was tailored to the application suite and performance objectives of the personal computer, and showed that workstations, as well as large time-sharing systems, could be a part of the Internet.

Widespread development of Local Area Networks (LANs), Personal Computers (PCs) and workstations in the 1980s allowed the nascent Internet to flourish. A major shift occurred because of the increase in scale of the Internet and its associated management issues. To make it easy for people to use the network, hosts were assigned names, so that it was not necessary to remember the numeric addresses. Originally, there were a limited number of hosts, so it was feasible to maintain a single table of all the hosts and their associated names and addresses. The shift to having many independently managed networks (e.g., LANs) meant that having a single table of hosts was no longer feasible, and the Domain Name System (DNS) was invented, permitting a scalable distributed mechanism for resolving hierarchical host names (e.g. www.eurodig.org) into an Internet address.

By 1985, Internet was already well established as a technology supporting a broad community of researchers and developers, and was beginning to be used by other communities for daily computer communications. Electronic mail was being used broadly across several communities, often with different systems, but interconnection between different mail systems was demonstrating the utility of broad based electronic communications between people.

As the network grew larger, it became clear that the sometimes ad hoc procedures used to manage the network would not scale. Manual configuration of tables was replaced by distributed automated algorithms, and better tools were devised to isolate faults. In 1987 it became clear that a protocol was needed that would permit the elements of the network, such as the routers, to be remotely managed in a uniform way. Several protocols for this purpose were proposed, including Simple Network Management Protocol (SNMP).

In the next few years, a new phase of commercialization took place. Originally, commercial efforts mainly comprised vendors providing the basic networking products, and service providers offering the connectivity and basic Internet services. The Internet became almost a "commodity" service, and much of the latest attention has been on the use of this global information infrastructure for support of other commercial services. This has been tremendously accelerated by the widespread and rapid adoption of browsers and the World Wide Web technology, allowing users easy access to information linked throughout the globe. Products are available to facilitate the provisioning of that information and many of the latest developments in technology have been aimed at providing increasingly sophisticated information services on top of the basic Internet data communications.

With the arrival of communications in portable form, and specially the development of smartphones, the Internet has become a pervasive communications means. Together with the development of apps – programs or software applications designed to run on mobile devices – the Internet now is the underlaying infrastructure that supports today's email, calendar and contact databases, mobile games, e-commerce, e-learning, e-government, social-media participation, media streaming, and an immense variety of on-the-go activities.

Today, those who have access to the Internet – and it is important to remind that today still 2.6 billion people do not have it – take it for granted that the Internet is here to stay and that it will be fit for purpose. But today, the Internet is in danger. Despite its technological robustness and resilience, it is important to acknowledge that it is politically fragile. Public policy and corporate decisions may splinter it into isolated networks that may not be able to connect or interoperate with one another efficiently.

This book explains how to identify these trends, and how to preserve the properties – set of practices – that make the Internet such a unique invention.

The Internet is a set of practices: the Critical Properties

What makes the Internet 'the Internet'? There have been many kinds of computer networks, but none of them has been embraced by so many people on a global scale and integrated into day-to-day life. What is it about the Internet as a *"network of networks"* that has evolved into an essential global tool, and a whole new space for innovation, growth, and transformation?

The Internet owes its success not only to the technology, but *to the way it operates and evolves*. The Internet provides unprecedented opportunities for advancing social and cultural understanding. The online environment empowers individuals to connect, speak, innovate, share, learn, and organize. There are virtually infinite opportunities in which we can use the Internet as a force for good. To make sure we can keep using it this way, we need to identify, recognize and protect its **critical properties**.

Specific technologies and business models may come and go, but the specific particularities of the way the Internet works – the *Internet Way of Networking* – has been a constant foundation for the success of the Internet from the beginning. For the Internet of the future to be as innovative and sustainable as it has been so far, these critical properties need to guide its evolution.

While the critical properties are foundational pillars, they manifest themselves through the benefits they provide to anyone who uses, builds, develops, and operates various components of the Internet ecosystem. These critical properties have been already hinted at during the previous chapter. While describing the first steps of the Internet, the reader can find them highlighted in bold text, as precursors of the following full list of five critical properties:

- **Critical Property 1:** An Accessible Infrastructure with a Common Protocol that is open and has low barriers to entry.
  Critical Property Benefits: Unrestricted access and common protocols deliver global connectivity and encourage the network to grow. As more and more participants connect, the value of the Internet increases for everyone.
- **Critical Property 2:** Open Architecture of Interoperable and Reusable Building Blocks based on open standards development processes voluntarily adopted by a user community.

Critical Property Benefits: Open architecture creates common interoperable services, which deliver fast and permissionless innovation everywhere. The inclusive standardization process and demand-driven adoption ensures that useful changes are adopted, while unnecessary ones disappear.

- **Critical Property 3:** Decentralized Management and a Single Distributed Routing System which is scalable and agile.
  Critical Property Benefits: Distributed routing delivers a resilient and adaptable network of autonomous networks, allowing for local optimizations while maintaining worldwide connectivity.
- **Critical Property 4:** Common Global Identifiers which are unambiguous and universal.
  Critical Property Benefits: A common identifier set delivers consistent addressability and a coherent view of the entire network, without fragmentation or fractures.
- **Critical Property 5:** A Technology Neutral, General-Purpose Network which is simple and adaptable.
  Critical Property Benefits: Generality delivers flexibility. The Internet continuously serves a diverse and constantly evolving community of users and applications. It does not require significant changes to support this dynamic environment.

Critical Property 1: An Accessible Infrastructure with a Common Protocol

You don't need permission from a central authority to connect to the Internet. You find a point nearby, make arrangements to connect, and you're on the Internet. The network is extended by the many kinds of organizations that connect to it. There is no policy on who can connect or what they should pay; these factors are largely driven by the market, not a centralized authority. Individual nodes connect to the Internet using different physical attachments (e.g., wireless LAN, Ethernet, DSL) and using a variety of underlying networking technologies. However, every hardware connection presents itself, eventually, as a pack-

et-switched interface, and every node has a common, open, network layer protocol available: the Internet Protocol (IP).

This open and accessible infrastructure delivers several key benefits: the first is global connectivity, bringing participants from around the world together and allowing them to reach each other. The second is growth: the network continues to grow because participants find value in connecting, which continues to create even more value for everyone connected. An Internet user trying to use a new application doesn't have to ask questions like "Are they running the same protocol I am?" or "Can I reach their part of the Internet from my part of the Internet?" In fact, most Internet users may not even know to ask these questions, because the Internet's open model means they don't have to think about such things. The network is open to anyone willing to participate, as a consumer, an information provider, an infrastructure builder, or an academic who wants to study how it all fits together. Without a central authority dictating who, how, and where connections are made, the network can grow organically to support the needs of its users. Once a network has surmounted the basic task of connecting to the Internet, they are part of the entire global Internet.

The accessible Internet assumes a market-based approach to growth, which has the effect of disenfranchising those without the means to fund connectivity and services. If you don't have any money to pay for it, there may be no business reason for someone to extend the Internet to your home or business. The Internet is open, but this does not mean that everyone will have access in an organic market. In areas where Internet users have few choices in service providers and connections, the benefits of this critical property may be diminished: Internet users may see a less accessible Internet.

When the property of a common protocol is missing, then users do not experience the full value of the Internet. For example, the Internet is undergoing a transition from IPv4 to IPv6. During that transition period, some users may be "on the Internet," yet unable to connect to some applications because one is on IPv4 and the other on IPv6. The danger of losing connectivity and, therefore, fragmenting the Internet is one of the reasons that the transition has taken so long and been so expensive: no one wants to violate this critical property and isolate themselves from the rest of the network.

Critical Property 2: An Open Architecture of Interoperable and Reusable Building Blocks

The Internet provides well-defined and well-understood services to applications using a simple open architecture. Technology building blocks are assembled in a layered fashion, working together to provide services to applications and end users. Each building block delivers a specific function, like supporting different network types, ensuring reliable transport, enabling security, or providing name resolution. Anyone can add innovation at any point – and Internet users can adopt (or reject) those building blocks that bring value without re-engineering the entire network. When building blocks for new common services are easy to build and install, this speeds deployment and innovation.

This open architecture delivers a key benefit: the common interoperable services and reusable building blocks allow for fast, permissionless innovation everywhere. An application designer does not have to start from first principles and wonder about the architecture and technology of the underlying network. Instead, the Internet's architecture offers a well-understood menu of choices that allow for fast deployment and innovation. Even uncertainties such as whether the underlying network is IPv4 or IPv6 are minimized from the application designer's perspective, because the building blocks responsible for transport functions hide these differences.

The structure of the Internet building blocks tends to push innovation upwards, as developers build on top of what exists, delivering better and more creative services – without requiring changes to the underlying technology.

The process of standardization is open to all interested and informed parties, and the results of this process are deployed on a voluntary basis. Changes are adopted when they serve a purpose, and unnecessary ones die. Even when some of the building blocks are proprietary (the Google Maps API, for example) their definitions are open enough to allow decentralized development and deployment, preventing ossification.

The importance of these open and interoperable building blocks can be seen when we find parts of the network closed. For example, Internet firewalls operate at a level where they "manage" transport layer TCP and UDP con-

nections in between end nodes. These devices have a much more static view of the Internet. This means that even if two end systems agree to run a new transport protocol, it may be difficult to deploy this protocol across the Internet because many Internet firewalls would not have the capability to control it, and thus would block it.

Fast innovation on the Internet is underpinned by an application designer's ability to take advantage of well-defined layered services. This is a great benefit to both the application and its users. For example, the well-known Transport Layer Security (TLS) protocol provides a defined security service to any application, eliminating the need to invent this mechanism from scratch. Experience has shown that trying to reinvent security rather than use standard building blocks, like TLS, often results in security compromises and breaches. Although the Internet is not free of breaches, the ability of security designers to re-use building blocks such as TLS delivers greater security at lower costs.

Critical Property 3: Decentralized Management and a Single Distributed Routing System

As a network of networks, the Internet's infrastructure is based on thousands of independent networks choosing to collaborate and connect together. Each of these networks runs a common, open, protocol (Border Gateway Protocol, BGP) that allows it to exchange routing information with its neighbors. And each of these networks makes independent decisions on how to route traffic to its neighbors, based on its own needs and local requirements. There is no central direction, or a controller dictating how and where connections are made, so the network grows organically, driven by local interests.

The distributed routing system delivers several key benefits: global reach, resilience, and optimized connectivity. Each organization that joins the Internet selects how they connect and how they route their data based on local requirements. They are able to optimize how their Internet connection works to match their needs: price, services available, connection

bandwidth, reliability, or quality, and so on. No central coordination is required because all agreements and policy decisions are between the connecting organization and their neighbors; you don't need to request permission to join the Internet from some central authority. The ability to make independent decisions on a regional, local, or hyperlocal basis allows the Internet to be more agile, scalable, and adaptable to the needs of its users.

The lack of a central routing authority within the Internet, however, does come with disadvantages as well. Without enforcement of a common policy, both human error and deliberate malice can result in interruptions to connectivity and security issues such as spying on Internet traffic or impersonating an organization. By taking a collaborative approach to routing, the Internet relies on peer pressure and community action to resolve issues – and resolution usually occurs very quickly once the community has identified the problem. In the absence of a common distributed routing system, the Internet would lose both agility and scalability. Local decisions and requirements would be impossible to accommodate without updating the central controller. Enforcing centralized routing – or even regional routing – eliminates the ability for end users to choose the best connectivity for their needs, creates scalability problems, brings economic disadvantages, and inevitably degrades the resilience and performance of a network as large as the Internet.

Critical Property 4: Common Global Identifiers

The Internet is an infrastructure that supports complex applications, some of them so large that they spread across continents and have millions of cooperating servers behind them. Internet users see elegant interfaces hiding behind a single name: Google, Facebook, Microsoft, and others. But there's an essential glue that allows every user to connect to the applications they use: IP addresses. Every bit of data flowing between a user's computer and the applications being used is in an IP packet, and every single IP packet has an address that says where it is going. These IP addresses allow any two systems on the Internet

to find each other, without ambiguity. Having common global identifiers delivers a key benefit: consistent addressability. The common identifier space, underneath all the various levels of application, delivers a coherent view of the entire network. From any point on the Internet, a tiny packet of information can be passed from computer to computer, each one examining the same few bits – the address – to clearly identify a destination. When used as designed, the IP address isn't subject to abbreviation or interpretation; IP addresses can't be confused or ambiguous. The common identifier space seems like such a small thing, but the consistency it delivers to the Internet is a critical property.

Closely tied to IP addresses is another group of identifiers: domain names supported by the Internet's Domain Name System (DNS). The DNS has many uses, but the most common is the creation of a consistent mapping between names and IP addresses. The consistency of the DNS is an important part of delivering a predictable and reliable service to every Internet user.

We can see how essential a single common global identifier space is by looking at what happens when this critical property is threatened. The perfect example is the continuing transition from the shorter IPv4 addresses to longer and more plentiful IPv6 addresses. IPv6 addresses are absolutely, increasingly required because there are simply not enough IPv4 addresses to accommodate the growth of the Internet. But with the introduction of IPv6 addresses, there now are two global identifier spaces, and if a device has an address in one space, it may not be able to reach the other. The challenge is that each address family is incompatible with the other, meaning that a device with an IPv4 address cannot exchange data, or 'talk', with an IPv6 device without the need for address translation. This creates fragmentation of the Internet, and the resistance to this fragmentation is one of the reasons that the transition from IPv4 addresses to IPv6 addresses is taking so long.

The common global identifier space of IP addresses means that individual users and network managers all have a single view of the network. Without these common global identifiers, we would have to construct special gateways, install translators, and create mapping tables to keep everything connected. Fracturing other name spaces, such as the DNS, also creates additional costs, overhead, and friction within the network. The utility of the Internet would be reduced, and resources would be wasted. Instead, with common, consistent, and predictable global identifiers, the Internet, a huge 'network of networks', acts as one single connected network.

Critical Property 5: A Technology Neutral, General-Purpose Network

The most popular uses of the Internet have changed dramatically from its first days: remote terminals and file transfer gave way to email and simple collaborative communications systems, which evolved to Web browsing, social networks, and media streaming. This was possible because the Internet was designed as a general-purpose network – not optimized for voice, particular usage patterns, or special traffic characteristics. The Internet is completely agnostic about the type of content that flows through it, guaranteeing neither quality nor connectivity, yet delivering enough of both to be a base layer for information services, commerce, communications, recreation, and more.

The benefit of a general-purpose network is its ability to continuously meet the requirements of a diverse, constantly evolving, environment. With no specific purpose in mind, the network serves data communication needs of billions of people, through an infinite number of applications, all doing different things, all at the same time. The Internet has been adapted for so many uses that it is displacing other types of networks. Dedicated voice telephone lines in the world peaked 15 years ago, replaced in part by Internet telephony. Streaming television and movie services are being delivered over the Internet, partially replacing programming delivered over cable TV and satellite networks. And because the Internet is not attached to any particular data transmission technology, it is able to re-use the cable TV and satellite infrastructure as data communications networks, incorporating them into the Internet as well. The building blocks responsible for services on top of best-effort packet forwarding, like reliable transport or specific applications, reside at the edge nodes of the Internet, and therefore can be rearranged to achieve a desired result without the need for global coordination or fundamental changes to the design of the underlying networks. This architectural approach is often referred to as the end-to-end argument, or principle.

The Internet's general purpose comes with drawbacks: while the Internet can be used for many things, it is not designed to do any particular job especially well. For example, without widespread mechanisms for congestion

control and quality of service, or the ability to centrally manage capacity and scalability of the network, streaming services have had to establish elaborate caching systems to serve their subscribers, i.e. ensuring they can watch high-definition videos or play virtual reality games without endless buffering. But this development also demonstrates the Internet's ability to adjust, adapt and build on top, or amend parts of it.

While the networks constituting the Internet may have been built for specific purposes, the general design was not. Otherwise, the Internet would not have been able to support other types of applications. For example, the first digital telephone networks were optimized for voice, delivering calls with higher quality and greater efficiency than the Internet can. Yet, these networks had to be completely overhauled to deliver a new feature, say a video call, at great expense and considerable difficulty. A general-purpose network may not be perfectly optimized for every new application, but it can support most new applications. A long-lived general-purpose Internet design lets innovators pursue, without permission, their ideas knowing the network's benefits and drawbacks, enabling fast movement forward while in comparison the network changes are small and gradual.

The five properties described here are *Critical* both because they are necessary for the Internet's healthy evolution, and because they convey what makes the Internet unique. They represent the Internet's *optimal state*. By codifying the basics of the ideal Internet model, we have a reference point that helps us tell whether this model is moving away from or towards the best it can possibly be. While the Internet's Critical Properties cannot guarantee the associated benefits, together they form the necessary condition for future evolution in a way most likely to create and disseminate the value that comes from connection.

The critical properties describe the foundation the Internet needs to exist – and illustrate why it is a *set of practices*, rather than a specific technology, and that makes it unique from other networking models. However, to help the Internet thrive we need another set of conditions that enable it to reach its full potential. This potential can be expressed by a set of goals for the Internet. Time and time again, different groups in different parts of the world with different viewpoints keep coming back to a common set of such aspirations:

1. An **Open Internet** that allows everyone to participate with a minimum of barriers, to use it, to innovate, and to grow and sustain the Internet as a force for good.

2. A **Globally Connected Internet** that is inclusive, allowing everyone to interconnect without geographical restrictions and use the full power of the network.
3. A **Secure Internet** that survives attacks, that supports everyone in maintaining integrity and confidentiality of the data. A secure Internet also means that its use does not create insecurity, such as botnets that are used in phishing scams.
4. A **Trustworthy Internet** that people can depend on to be there, so that the Internet can be a base for worldwide services, everything from recreation to commerce to information.

These four goals become guidelines for the journey to a better Internet. They tell us what we want the Internet to be, now, and in the future.

These Internet goals are aspirational statements and because of their broad and abstract nature, it is difficult to use them to analyze how various developments may impact the Internet. To aid this analysis, for each of the Internet goals, we have identified a series of supporting characteristics: things that progress or hold back the Internet's growth and its global goals. Generically, we call these supporting characteristics *"Enablers"*: they advance and enable the targeted goal.

The Enablers of an Open, Globally Connected, Secure and Trustworthy Internet

The purpose of identifying enablers is to simplify the task of analyzing the potential effects of proposed changes or new policy proposals, and in consequence how they may affect the goals. For example, a secure Internet requires that the Internet supports both data confidentiality and data integrity. Each of these is an enabler: data confidentiality supports the goal of a secure Internet, as does data integrity. If either is missing, the security of the Internet is reduced. Since the enablers, not the goals, are the tool for analysis of proposals, they are the focus of this framework.

In this section, we identify enablers that relate to each of the four Internet goals. To help make the meaning of each enabler and how it relates to the attainability of an Internet goal as clear as possible, we have provided examples of different policies or technologies specific to an enabler that either advance or block the goal in the area identified.

It is important to note the enablers are presented in their ideal form. By thinking of them as if they reflect a perfect state, we have a reference point that helps us determine whether a particular development moves the Internet away or towards the identified goals. The enablers can also expose some of the tensions that exist between the goals and make the potential trade-offs clearer. For example, some of the actions may have a positive effect on the security of the Internet, while making it less open at the same time.

Supporting an Open Internet

The Internet is fully open when anyone may create, use, or deploy it according to their own wishes. With a fully open Internet, anyone is free to deploy Internet networks and build

services and applications on the Internet, combine them in novel ways and deploy them without barriers, as long as this is done in interoperable ways. An open Internet is an accessible Internet – it is easy for networks to join, and for users to connect to it and use its services.

Following some enablers of an open Internet are defined. We have also provided some examples of different policies or technologies that either advance or block the goal in the area identified. Note that these examples should be read as illustrations for their effect on the enabler with which they are listed. Some of our examples may have positive effects for one enabler and negative effects for others.

Enabler: Easy and unrestricted access

Description:
It is easy to become part of the Internet, for networks and users alike. That means that for users the Internet is affordable, and Internet services are accessible, and that networks can easily become part of the Internet, without unnecessary regulatory or commercial barriers for both groups.

Questions:
- Does the proposed change create or lower a barrier to entry, such as costs, administrative overhead, or other difficulties?
- Is the effect of the change to restrict who can participate, closing down the Internet?
- Does the proposed change create unnecessary requirements for particular skills, or raise costs?

Example 1:
The Web Content Accessibility Guidelines (WCAG) is a recommendation by the World Wide Web Consortium (W3C) for making Web content more accessible, primarily for people with disabilities. For example, by recommending text alternatives for any non-text content so that it can be changed into other forms, such as speech or large print. In some jurisdictions conformity with these guidelines is required by law to safeguard the rights of people with disabilities.

This is an example of a positive effect on an open Internet for users, strengthening “Easy and unrestricted access” by facilitating the use of the Internet services by everyone.

Example 2:
In some countries Internet access is only available through monopoly service providers. The effect of this is to create higher costs (through lack of competition) as well as limitations on what types of services and connection methods are available. The lack of choice among providers effectively limits overall access.

This is an example of a negative effect on the goal of an open Internet for users and networks by weakening "Easy and unrestricted access" through limitation of connectivity options and higher costs.

EuroDIG discussions: Net Neutrality statement

Analyzing the characteristics that ensure the Internet continues to develop and thrive is an important exercise that also helps to identify policies and business decisions that deviate from the ideal model of the Internet. A key policy that has helped to sustain the Internet is net neutrality. EuroDIG has been at the forefront of discussions about the impact of net neutrality since the annual EuroDIG meeting in Madrid in 2010 and culminating in the Multistakeholder Statement on Net Neutrality, which was drafted by the EuroDIG stakeholder community in Sofia in 2015: https://eurodigwiki.org/wiki/Draft_statement_on_net_neutrality._2.0

Human rights perspective
The network neutrality principle plays an instrumental role in fostering the full enjoyment of Internet users' human rights as well as preserving Internet openness. Network neutrality is the principle according to which Internet traffic shall be treated equally, without discrimination, restriction or interference, independent of the sender, receiver, type, content, device, service or application so that Internet users' freedom is not restricted.

In accordance with the network neutrality principle, Internet access providers shall not restrict Internet users' right to freely seek, impart and receive information

and ideas via the Internet. Accordingly, Internet access providers should manage the transmission of Internet traffic in a non-discriminatory manner and shall provide meaningful, intelligible and transparent information with regard to their Internet access service and traffic management practices, notably with regard to the coexistence of Internet access service and other services. In general, all players in the Internet ecosystem shall comply with privacy and data protection legislation, and act in a respectful way with relevant human rights. Accordingly, any techniques to inspect or analyse Internet traffic shall be in accordance with privacy and data protection legislation.

In order to ensure network neutrality, guaranteeing the full enjoyment of Internet users' human rights, the competent national authority shall be mandated to regularly monitor, conduct independent testing and report on Internet traffic management practices and usage policies.

End-user perspective

Internet access services enable end-users to communicate, to access, deliver and share content and Internet applications. In general, users expect Internet traffic that they send and receive to be conveyed in a manner that is independent of its source, content or destination and in a manner that respects their privacy. The continued success of the Internet as a communications medium and an engine for innovation and growth depends upon the continued enablement of new services and applications and end-user Internet traffic not being blocked or otherwise degraded by Internet service providers or other actors in the Internet ecosystem.

Choice and transparency are at the heart of a user's Internet experience, enabling them to remain in control of their Internet experience, and thereby allowing them to benefit from, and participate in, the open Internet. Internet subscribers may choose to block, prioritise or otherwise modify Internet traffic they send or receive but do not expect to have these choices made for them by third-parties without their consent.

Business perspective

The Internet should be open and accessible to all people in a non-discriminatory fashion. Net neutrality holds that subscribers of Internet access services should be able to access and share content, applications and online services of their choice, when legally entitled to do so. In accordance with such principle, market participants shall refrain from behaving in an anti-competitive way to the detriment of consumers or competition. To safeguard net neutrality it is essential to preserve effective competition among providers of access services to the Internet as well as in all other elements of the Internet ecosystem as well as to ensure transparency to end users by providing clear and meaningful information that facilitates informed customer choices when matching offers with their heterogeneous demands.

A set of general and global principles would best suit an Internet in permanent evolution. It is important to note that a "one size fits all" approach would not best benefit different countries and regions. Open Internet guidelines should promote access and openness, while encouraging network operators and Internet players to innovate and deliver the ample range of services demanded by customers, assure a satisfactory user experience over the Internet and promote the goal of universal internet connectivity. The possibility to provide commercially differentiated offers, including specialised services (see technical definitions section below) whilst at the same time providing strong safeguards preserving open and robust Internet access services, in order to develop yet unforeseen new business models along the digital value chain has to be preserved as a mean to increase customer choice. In order to best serve the interests of the end-users and all Internet players, providers of Internet access services shall not block, throttle or discriminate against specific content, applications or services except as necessary and for as long as necessary for the application of reasonable traffic management. Reasonable Internet traffic management is justified provided it is done in a transparent, non-discriminatory and proportionate manner.

Enabler: Unrestricted use and deployment of Internet technologies

Description:
The Internet's technologies and standards are available for adoption without restriction. This enabler extends to end-points: the technologies used to connect to and use the Internet do not require permission from a third party, operating system (OS) vendor, a network provider, or any other third party. The Internet's infrastructure is available as a resource to anyone who wishes to use it. Existing technologies can be mixed in and used to create new products and services that extend the Internet's capabilities.

Questions:
- Does the proposed change restrict how the Internet's technologies can be used or deployed?
- Is the effect of the change to create an unfair or discriminatory limit?
- Does the proposed change unreasonably limit how end users can manage and control their own devices?

Example 1:
The RSA SecurID authentication system was an early multi-factor authentication system protected by patents and trade secrets. The resulting vendor lock-in and expense was good business for RSA, but also limited the ability of developers to include SecurID technology in Internet applications, depriving users of the opportunity for better security and protection against credential theft.

This is an example of a negative effect on the goal of an open Internet caused by weakening the enabler restricting the use and deployment of technologies.

Example 2:
Google and Oracle engaged in a complicated 9-year battle over the Java Application Programming Interface (API), based on Oracle's ownership of Java and Google's use of the API to make their own written-from-scratch Android operating system compatible with Java applications. Oracle claimed that the API was theirs to own and control as much as the rest of the Java source code, while Google claimed that their use of the API was not subject to Oracle's copyright under the "fair use" doctrine. In the end, the US Supreme Court agreed with Google: Oracle's copyright on Java did not restrict Google's use of the API.

This particular case was largely economic: had Oracle won, Google would have had a huge bill to pay based on the immense popularity of Android. But the larger issue of

whether the use of the API was fair use affects future Internet innovation.

This is an example of a positive effect on the goal of an open Internet that strengthened "Unrestricted use and deployment of Internet technologies", as the US Supreme Court made it clear that some types of innovation and use cannot be restricted.

Enabler: Collaborative development, management, and governance

Description:
The Internet's technologies and standards are developed, managed, and governed in an open and collaborative way. This open collaboration extends to the building and operation of the Internet and services built on top of the Internet. The development and maintenance process is based on transparency and consensus, and has as its goal the optimization of infrastructure and services to the benefit of the users of these technologies.

Questions:
- Does the proposed change limit collaboration during development, operation, and governance?
- Is the goal of the proposed policy a restraint on collaboration?

Example 1:
The Internet's address space is a limited resource that requires careful administration. Rather than centralize decision-making, each of the major Internet regions is responsible for governance of the address space in their region. Local address space policymaking is based on a collaborative process driven by the regional community. Policies that reach consensus are implemented by the Regional Internet Registry. This creates an environment where the Internet itself can be "open" in a context that fits the region.

This is an example of a positive effect on the goal of an open Internet, strengthening "Collaborative development, management, and governance" through community-driven management of Internet resources.

Example 2:
Internet Exchange Points (IXPs) offer community network operators the opportunity to connect and exchange Internet traffic. By bringing together Internet stakeholders such as ISPs, municipal networks, and content delivery networks, local Internet traffic is routed more efficiently, and local users enjoy better

and more resilient access to regional traffic. IXPs often include competitors working together for the benefit of themselves and of the local community. Many IXPs have unrestricted open peering policy leaving the decision of who to interconnect with to individual participants. A multilateral policy of peering with everyone is very common.

IXPs are an example of a positive effect on the goal of an open Internet, strengthening "Collaborative development, management, and governance" through facilitating open access, using non-discriminating policies and building local communities.

EuroDIG discussions: the role of the IXPs

The importance of Internet exchange points (IXPs) in the Internet's infrastructure and Internet governance generally was discussed at EuroDIG 2014 in Berlin. The workshop entitled "*The Role of IXP (Internet Exchange Points) in Internet Governance*" concluded with strong messages which underlined the crucial role that IXPs play in collaborative governance of the open and decentralized global Internet. The messages emphasized, in particular, their importance in providing a solid and reliable infrastructure that backs up the worldwide Internet and serves as a neutral marketplace for interconnection that is open for all Internet users.

Source:
https://eurodigwiki.org/wiki/The_Role_of_IXP_(Internet_Exchange_Points)_in_Internet_governance_%E2%80%93_WS_07_2014

Supporting a Globally Connected Internet

With a truly interconnected Internet, anyone who wants to be part of the Internet can participate and exchange traffic with other participants without restrictions. A globally connected Internet is not just a technical capability, but one in which all barriers to connection are minimized and everyone who wants to use it can get a fast, reliable, and affordable connection to end-points (users, services or resources like storage, computing, sensing, and actuating) no matter where they are located.

Enabler: Unrestricted reachability

Description:
Internet users have access to all resources and technologies made available on the Internet and are able to make resources available themselves. Once a resource has been made available in some way by its owner, there is no blocking of legitimate use and access to that resource by third parties.

Questions:

- Does the proposed change restrict which resources a user can use and access, or restrict the resource the user may contribute to the Internet?
- Is the effect of the change that a third party can block access to significant parts of the resources of the Internet, or creates single points of failure?

Example 1:
The Internet community worked hard to mitigate negative effects of widespread use of the Network Address Translation (NAT) devices between end-user networks and the Internet, which allow for the more efficient use of limited public IPv4 address space. While NAT is harmless for most Internet users, it is extremely disruptive to some protocols, such as Voice over IP (VoIP). To work around the problems created by NAT, the Internet community has defined other protocols, including STUN, TURN, and ICE. For many users, these protocols and technologies enable use of peer-to-peer type communications that would normally be blocked as a side effect of NAT.

This is an example of a positive effect on the goal of a globally connected Internet, strength-

ening "Unrestricted reachability" through removing barriers to end-to-end connectivity.

Example 2:
In some countries, VoIP services are blocked because of government policy maintaining a state-granted monopoly on (and ability to tax) long-distance voice communications. This blockage creates economic inefficiencies, imposes higher costs, and serves to isolate users in that country from other Internet users taking advantage of VoIP technologies and services.

This is an example of a negative effect on the goal of a globally connected Internet weakening "Unrestricted reachability" by blocking services.

Enabler: Available capacity

Description:
The capacity of the Internet is sufficient to meet user demand. No one expects the capacity of the Internet to be infinite, but there is enough connection capacity – ports, bandwidth, services – to meet the demands of the users.

Questions:
- Does the proposed change act to increase the availability of Internet resources, such as bandwidth or other capacity?
- Is the effect of the policy to limit growth and capacity, either directly or indirectly?

Example 1:
Stakeholders in the Democratic Republic of the Congo came together in 2012 to set up an IXP in the capital and primary city of Kinshasa (KINIX). After some years, the community felt the need for a second IXP in the second biggest city Lubumbashi. It was launched in 2019. The community launched a third IXP in Goma called GOMIX in September 2021. The establishment of these IXPs significantly reduced connectivity costs (in 2020 estimated yearly savings per network were as high as $163,000) and increased available communication capacity.

This is an example of a positive effect on the goal of a globally connected Internet by improving "Available capacity" in areas of scarcity.

Example 2:
Networks that choose not to deploy IPv6 but rely on Carrier Grade Network Address Trans-

lation (CGNAT) – essentially a much larger version of the NAT mentioned earlier – will reach a point where their users will starve the available resources. CGNATs necessarily restrict the number of simultaneous connections individual users can use. Without a CGNAT an individual user can use 64 thousand connections simultaneously, while in a worst-case scenario a million people behind a CGNAT only have 16 simultaneous connections available per individual user. Compare that number to the 50 simultaneous connections needed to load a typical website.

This is an example of a negative effect on the goal of a globally connected Internet. The exhaustion of IPv4, combined with the lack of IPv6 deployment, leads to starvation of resources for users, thereby weakening “Available capacity”.

Supporting a Secure Internet

A secure Internet is resistant to attacks on its infrastructure, delivering a robust service to its user community. This means that its protocols and infrastructure, such as routing and DNS, should present a secure base that is resistant to both intentional attacks and accidents. In a secure Internet, data should have its confidentiality, integrity, and availability protected. Ideally, a secure Internet also does not create insecurity, such as botnets that are used in phishing scams. And the services and applications that run over the Internet itself should be secure, to the greatest extent providing defense in depth.

In this context, “secure” complements and relates to “trustworthy.” When evaluating different proposed policies through the lens of the enablers both will often be taken into account.

Enabler: Data confidentiality of information, devices, and applications

Description:
Data confidentiality, usually accomplished with tools such as encryption, allows end users to send sensitive information across the Internet so that eavesdroppers and attackers cannot see the content or know who is communicating. Allowing the transfer of sensitive information helps create a secure Internet. Data confidentiality also extends to data-at-rest in applications and on devices. (N.B., "confidentiality" also contributes to privacy, which is part of a trustworthy Internet)

Questions:
- Does the proposed change strengthen or weaken the ability of users to preserve the confidentiality of their information in transit or at rest?
- If this change is implemented, will the underlying protocols of the Internet provide stronger or weaker confidentiality?

Example 1:
Mauritius proposed in 2023 that all social media traffic be decrypted, inspected, and archived. The rationale was based on the relatively obscure local language and lack of physical presence of major social media organizations, leading to a poor response time to legitimate complaints from Mauritius. If it had not been dismissed, the proposed change would have impacted Internet security for people in Mauritius and their correspondents by dramatically reducing the confidentiality of information exchanged over social networks. Unfortunately, such proposals are becoming common around the world.

This is an example of a negative effect on the goal of a secure Internet by weakening "Data confidentiality of information, devices, and applications" by requesting the removal of protection of information exchanged between correspondents online.

Example 2:
Most websites have added encryption to provide data confidentiality for their users. This added security is based on standardized protocols (TLS/HTTPS) and frameworks to ground the encryption process securely using digital identity certificates and trusted Certification Authorities (CAs). Together, the certificates, CAs, and all the processes and rules that govern this ecosystem are generically referred to as "the WebPKI" (for "web public key infrastructure"). While there are many legitimate criticisms of different aspects of the WebPKI, the overall result is that Internet users can eas-

ily and transparently use encryption to provide greater confidentiality to their Internet activities.

This is an example of a positive effect on the goal of a secure Internet, strengthening "Data confidentiality of information, devices, and applications" through a combination of industry standardization and implementation by application developers.

Enabler: Integrity of information, applications, and services

Description:
The integrity of data sent over the Internet, and stored in applications, is not compromised. That is, information sent over the Internet shouldn't be modified in transit, unless directed by the communicating parties (e.g., a captioning bot may be useful to turn spoken words into text).

Critical underlying Internet services, such as DNS and the routing system, cannot be manipulated or compromised by malicious actors.

Data stored in applications cannot be manipulated or compromised by third parties.

Questions:

- Does the proposed change strengthen or weaken the integrity of data, or the ability of users to verify that data are not corrupted?
- Does the proposed change strengthen or weaken the accuracy and integrity of Internet services, such as DNS?

Example 1:
Resource Public Key Infrastructure (RPKI) is a combination of technology standards and Internet-hosted databases that help to increase the integrity of routing across the Internet. With RPKI, IP address holders can publish information about how their address blocks should be routed; at the same time, network operators such as ISPs can use the information published in the RPKI system to validate routing updates and avoid malicious behavior such as hijacking of IP address space.

This is an example of a positive effect on the goal of a secure Internet by increasing "Integrity of information, applications, and services" as far as the Internet routing system is concerned.

Example 2:
In 2011, S.978, the "PROTECT IP Act" was proposed in the United States Senate. The bill contained a variety of technical requirements

for Internet content filtering, including intercepting and modifying DNS queries and responses. The proposed legislation would also have been inconsistent with an important Internet security technology, DNS Security Extensions (DNSSEC), and would have resulted in a stifling of DNSSEC deployment. S.978 would have accomplished its goals, in part, by attacking the soundness and stability of DNS. (The bill never went to a vote at the Senate).

This is an example of a negative effect on the goal of a secure Internet, by weakening "Integrity of information, applications, and services" for the DNS system.

Supporting a Trustworthy Internet

Unlike security, trustworthiness depends not only on the state of the Internet, but also on the state of people and organizations that use and participate in it. The extent to which the Internet may be considered trustworthy depends upon an informed base of users who have the tools to evaluate trustworthiness, based on their current knowledge of the Internet's vulnerabilities and threats to it.

The concepts of trustworthy Internet and a secure Internet are tightly intertwined: if the Internet is not secure, it cannot possibly be trustworthy. However, a completely secure Internet could still be untrustworthy if it violated user expectations or if some of its participants were not worthy of trust. Trustworthiness is not simply a matter of security.

Enabler: Reliability, resilience, and availability

Description:
The Internet is reliable when technology and processes are in place that permit the delivery of services as promised. If, for example, an Internet service's availability is unpredictable, then users will observe this as unreliable.

This can reduce trust not just in one single service, but in the Internet itself.

Resilience is related to reliability: a resilient Internet maintains an acceptable level of service even in the face of errors, malicious behavior, and other challenges to its normal operations.

Questions:

- Does the proposed change create unpredictable variations in the Internet's reliability or in the reliability of a service or set of services?
- Will users be unable to know, from day to day, whether they can use the Internet and its services?
- Does the proposed change increase or reduce the overall level of the Internet's resilience to malfunction?

Example 1:

Statuspage.io is a commercial product that focuses on showing service status – up, down, or something in between – to the service's users. By creating an entire product around delivering this information, the statuspage.io team has delivered a tool to the Internet community that fills a niche and makes it easy to transparently communicate service status. Because statuspage.io is commercial, it has also inspired competitors – including open source – further making this type of transparency tool widely available to network operators and information service providers.

This is an example of a positive effect on the goal of a trustworthy Internet, by strengthening "Reliability, resilience, and availability" by increasing transparency about service performance.

Example 2:

Deliberate Internet shutdowns are often used at the country level during stressful moments, such as during highly contested elections, or at times of civil unrest. Governments in Belarus, India and Venezuela have all shut down their nations' Internet in recent years. The result of these shutdowns can be unpredictable Internet reliability. The reasoning and duration of these shutdowns is rarely transparent, further reducing trust.

This is an example of a negative effect on the goal of a trustworthy Internet by weakening "Reliability, resilience, and availability" through deliberate interruption of services.

See ISOC Pulse tracker https://pulse.internetsociety.org/shutdowns for more specific examples.

Enabler: Accountability

Description:
Accountability on the Internet gives users the assurance that organizations and institutions they interact with are directly or indirectly acting in a transparent and fair way.

In an accountable Internet, entities, services, and information can be identified and the organizations involved will be held responsible for their actions.

Questions:
- Does the proposed change create non-transparent authorities or hidden actors?
- Is the effect of the change to create unaccountable or anonymous authorities who will affect the trust users have in the Internet?

Example 1:
RFC 7725 defines a way for a website owner to signal to users that information is not available as a consequence of a legal demand. Websites, ISPs, and search engines that are blocking access to information because of some sanction can use the mechanisms from this RFC to be very transparent: "you can't see what you are looking for because, legally, we are forbidden from showing it to you."

This is an example of a positive effect on the goal of a trustworthy Internet by strengthening "Accountability" through increased transparency.

Example 2:
The Global Internet Forum to Counter Terrorism (GIFCT) aims to "prevent terrorists and violent extremists from exploiting digital platforms." As part of this mission, GIFCT maintains a database of images and videos deemed "terrorist content" to be shared among its members. When a video or image that matches one in the database is detected, the database users can act on this information. This database has grown to have more than 300,000 images and videos. Because the database uses one-way hashes to identify elements, the actual videos and images are not stored. This creates a lack of transparency – researchers cannot review the database; regulators cannot audit it – with significant risks for freedom of expression, such as censorship of non-English content. The breadth of the database is hard to control too, since in some cases it is hard to determine "terrorist content" with global applicability.

This is an example of a negative effect on the goal of a trustworthy Internet, by weaken-

ing "Accountability" through designing and widely deploying opaque blocking tools with little oversight and often little recourse for the user.

EuroDIG discussions: accountability of organizations and institutions

Accountability has been an important cross-cutting issue in various EuroDIG discussions. One of the most notable examples of this were the EuroDIG sessions on accountability during the process of ending the US government's oversight of ICANN (known as the IANA stewardship transition) in 2014 – 2016. The session "*The IANA stewardship transition: a test case for Internet governance?*" held during EuroDIG 2015 in Sofia stressed the importance of trust and accountability in structures and organisations, including ICANN, during that landmark process of transition from government oversight to multistakeholder governance.

Source:
https://eurodigwiki.org/wiki/The_IANA_stewardship_transition:_a_test_case_for_Internet_governance%3F_%E2%80%93_PL_04_2015

Enabler: Privacy

Description:
Privacy on the Internet is the ability of individuals and groups to be able to understand and control what information about them is being collected and how, and to control how this is used and shared. Privacy often includes aspects of anonymity, removing linkages between data, devices, and communications sessions and the identities of the people to which they pertain.

Questions:

- Does the proposed change improve, reduce or eliminate users' ability to understand or control how their information is collected, or to control how this information is used and shared?
- Is the effect to provide or eliminate the possibility for a user to act anonymously or pseudonymously?

Example 1:

The United Kingdom's Online Safety Act (2023) obliges service providers to monitor and delete "legal but harmful" content. Campaigners suggest that in practice this will oblige service providers to scan content – including private messaging. The only way to do this would be to break end-to-end encryption, meaning "private" message platforms would be subject to surveillance by non-state and state actors, even without any allegation or suspicion of crime being committed.

This is an example of a negative effect on the goal of a trustworthy Internet as it weakens "Privacy" by creating hidden actors and hidden actions based on pervasive traffic monitoring .

Example 2:

The California Consumer Privacy Act (CCPA) gives consumers more control over the personal information that businesses collect about them. While most of the CCPA is about manually communicated policies and disclosures, the current CCPA also requires that businesses honor opt-out signals from web browsers similar to those of "Do Not Track"; this gives a simple and easily managed way for the end user to signal their privacy preferences. By encouraging standardized and automated privacy controls, such as the Global Privacy Control (GPC), the CCPA helps users better manage their privacy.

This is an example of a positive effect on the goal of a trustworthy Internet, strengthening "Privacy" by providing consumers with more control over their personal information that businesses collect.

Example 3:

The EU GDPR was sweeping legislation that had as one goal the creation of significant privacy enhancements for Internet users. Although implementation of the GDPR requirements in many cases resulted in sub-optimal user experience, the regulation has largely positive effects on trustworthiness by making user privacy and privacy policies are part of the business model and design of current Internet services, which has a long-term effect of increasing trust.

This is an example of a positive effect on the goal of a trustworthy Internet, as it strengthens "Privacy" by making the requirements part of the business model and design of the services.

EuroDIG discussions: privacy

Privacy has been an integral part of EuroDIG discussions since the inception of EuroDIG in 2008 and this is reflected in many of the EuroDIG messages agreed and published following the annual meetings. For example, a session at EuroDIG 2018 in Tbilisi shortly after the EU's General Data Protection Regulation (GDPR) came into effect, highlighted the importance of privacy in relation to trust on the Internet, transparency, and accountability. The messages from that session emphasized that privacy-enhanced technologies should be everyone's right.

Source:
https://eurodigwiki.org/wiki/Privacy_is_everywhere:_how_to_deal_with_emerging_problems%3F_%E2%80%93_WS_02_2018

Compiling all together: the Internet Impact Assessment Toolkit

Up to this point, this book has presented a short overview of how the Internet was created, which are its *Critical Properties* and also the *Enablers* that make it a meaningful infrastructure. All these concepts interact together as components of the Internet Impact Assessment Toolkit (IIAT), designed and made available by the Internet Society to help people around the world assess the implications of changes that may affect the Internet.

The toolkit provides a framework for the analysis of proposed changes in policies, legislation, technologies, applications, business models, and regulations. The goal of this toolkit is to help Internet stakeholders identify the possible effects, both positive and negative, and make more informed decisions.

Using these analytical tools, users may get a clearer view of the potential effects of proposed policies – be it public policies, new laws, or corporate decisions –, and be able to find the arguments to debate, or make alternative suggestions, that can achieve the intended goals without harming the Internet. And based on these analyses, stakeholders are able to clearly explain the pros and cons of a proposal in a balanced way.

The toolkit focuses on the *"Internet Way of Networking"* – as described on page #21, including the five Critical Properties – and the commonly accepted targets for the Internet, i.e. what we call the "Internet goals", that are represented by the Enablers. The basic premise of the toolkit is that the Internet Way of Networking and the four goals for the Internet unlock the broader benefits of the Internet for everyone. Should legislation, or corporate decisions, be proposed that could potentially affect the Internet, the toolkit may be used to build a picture of what the likely effect on the Internet will be.

The target of an IIAT analysis is to determine if the issue at hand impacts the Critical Properties of the Internet Way of Networking, and/or the enablers of the Internet goals. Proposed changes may have mixed effects, positive in some ways and negative in others.

A "two-pass approach" may be used: first to evaluate; and then to analyze in depth.

First evaluate…

- The proposed changes must be evaluated against *each* of the critical properties and the enablers.
- When looking at the goals, use the enablers to deep-dive and better understand the effects. Typically, this evaluation stage will turn up only a few things worth evaluating in depth.

… then analyze

- When an enabler or critical property is impacted, the effect may be rated as positive or negative.
- It is also helpful to divide any impacts into *strong* and *weak* effects, to focus on the most important issues.

The result of this analysis may be an *Internet Impact Brief* (IIB), a document that answers to the question: *"How might this issue impact the Internet, positively or negatively?"* And this document should also provide for the right arguments to engage with policy formulators, to discuss the ammendments that would be needed to align their proposals with an open, globally connected, trustworthy and secure Internet, preserving its core values.

The Internet Society has complemented this toolkit with detailed information on how to conduct the analysis, how to produce a brief[1] and even how to launch an campaign to engage with decision-makers. All this information, and more resources, including published Internet Impact Briefs, are available on their website: https://www.internetsociety.org/resources/doc/2020/internet-impact-assessment-toolkit/introduction/

[1] Infographic "How to do an Internet Impact Brief" – https://www.internetsociety.org/wp-content/uploads/2022/10/How-to-do-an-Internet-Impact-Brief-Infographic.pdf

The Internet is under threats: what can the community do

Any policy or decision that undermines the open, global, interoperable nature of the Internet – the very attributes that have empowered people everywhere to benefit from it – contributes to *fragmenting* the Internet. Instead of the same global, seamless Internet that the world now relies on and benefits from, we could end up with a splintered, degraded version of the Internet that is unrecognizable to the one we have today. One where ordinary people are not able to create, share, and connect freely.

Internet fragmentation is the division or splintering of the unified, open, global Internet into smaller, isolated networks subject to different rules, regulations, and technical standards – which may not be able to interconnect or interoperate seamlessly.

The Internet works well because no single person or entity controls it. Anyone can choose to connect to it, and the network grows and adapts to fit our needs. When all this works correctly, our experience of the Internet should be the same no matter who or where we are, because we are all connecting to the same unified, global, seamless Internet.

But this is changing *quickly*. The Internet is under a variety of threats, and we cannot take the open, global Internet for granted. If we are not careful, we could lose the Internet and all of the benefits it offers people worldwide. As our socio-economic dependence on it grows, decisions made about the Internet could have a profound effect on how it will function for generations to come. Instead of the same global, seamless Internet that we rely on, we could end up with an unrecognizable, splintered version of it – fragments of what it originally was.

Policies and decisions that threaten fragmentation of the Internet can emerge in different forms and can have technical variations. However, such threats can be analyzed to mark out trends and patterns, thereby helping advocates to find out what they have in common. For instance, some policies may give governments more control over what people can access over the Internet, while other threats could break the seamlessness with which networks interconnect with one another.

For our efforts to uphold the open, global Internet to be more effective and efficient, it is useful to categorize fragmentation threats into *'threat categories'*. This exercise can help us understand if a similar policy has been pro-

posed elsewhere in the world, and what existing research and advocacy resources already exist. We present eight types of threats, giving concrete examples of each of them and explaining the negative effects they have on the Internet.

Threats may not be water-tight, and may threaten fragmentation on a number of axes. They may fall into multiple categories. Ultimately, these categories are a way to organize specific cases of fragmentation to aid understanding and advocacy.

Identified types of threats

Threat 1: Regulating Business Relationships

An increasing number of governments are proposing policies to force big Internet-based platforms to pay for telecom infrastructure costs. Or they want to compensate media organizations for making news available on their platforms. They argue that big Internet platforms 'free-ride' on the Internet, and should pay to be able to use the Internet's infrastructure. But these proposals threaten the decentralized nature of the Internet either by fundamentally impacting the way that networks interconnect with one another or by centralizing the management of Internet functions and content.

In an attempt to regulate Internet-based businesses, such proposals will critically undermine the Internet's infrastructure, and have cascading effects on people, businesses, global trade and supply chains, and the seamlessness and resilience of the Internet.

Example:
The EU's "Fair Share" (aka "cost-sharing", "network fees", or "dispute settlement mechanism") model will force large platforms to pay more, and it will fragment the Internet. This proposal is broadly opposed across academia, regulators, small and medium-sized ISPs, IXPs, civil society organizations, consumer advocates, and content providers. In other words, almost everyone, apart from the largest network operators, sees this as a po-

tentially damaging decision. Even smaller ISPs would be at a disadvantage, and anyone who sends a large amount of traffic would find it difficult to compete globally, or even with larger players within Europe.

The public consultation's results were published in October 2023. The 'Fair share' model was clearly rejected by a majority of those who responded. Moreover, most of the European Parliament is also opposed to this. But because the large operators have broad influence, and they've been arguing for some form of this for almost two decades, we can't afford to turn our attention away.

Effect:

This policy would undermine most of the principles of the Internet Way of Networking that make the Internet valuable to all of us. This isn't just a policy that would also contribute to fragmentation – it is, at its core, intentional fragmenting of the Internet, which would serve nobody.

Talking Points:

People who use the Internet already pay for their connectivity, and so do the businesses that provide services on the Internet. That means there is no market failure – everyone is already paying and getting paid.

It sounds like a good way to keep large platforms in check, but only the largest telecom operators actually benefit from these policies. They already have a significant monopoly on traffic termination, and this would give them even more room to exploit that.

Right now, it is enough to be connected to the Internet, and in doing that, you can be part of this global network of networks. These rules break that idea, and fragment the Internet as we know it.

EuroDIG discussions: the "Fair Share" debate

Following renewed calls from telecom operators in Europe to make technology companies pay for network usage, a workshop entitled *"Models to support investment in the network infrastructure in Europe: what is the way forward?"* held at EuroDIG 2023 in Tampere examined the concerns raised in relation to the EU's "Fair Share" proposals. While the session highlighted significant differences in the per-

spectives of various stakeholders regarding this issue, it also conveyed a strong message about the risk of Internet fragmentation that cost-sharing approaches create.

Source:
https://eurodigwiki.org/wiki/Models_to_support_investment_in_the_network_infrastructure_in_Europe:_what_is_the_way_forward%3F_%E2%80%93_WS_01_2023

Threat 2: National Internet Gateways

Internet exchange points are places where networks come together to exchange traffic from different sources. In an open, global Internet, networks are free to interconnect with other networks across geopolitical borders. This is why you can easily use services that are based in countries other than the one you live in.

A national gateway uses these exchanges as digital checkpoints, where a government can block, throttle, or filter that traffic. If traffic has to run through government-mandated gateways, the free flow of information can become difficult, or even impossible, affecting the user experience and undermining the network itself.

Example:
Cambodia's National Internet Gateway. In 2021, Cambodia issued a decree that required all Internet traffic in Cambodia to be rerouted through the National Internet Gateway by February 2022. The plans remain delayed, but if it's implemented, it will not only affect Cambodia, but also the networks that interconnect with Cambodian networks. This could affect traffic within the region. If this plan is implemented, this would mean that networks must connect to government-mandated locations, increasing the technical and financial barriers that network operators face in becoming part of the global Internet. Since no other networks in the country can access the global Internet directly or independently, this would severely impact a network's global reach and limit collaboration between Cambodia and the rest of the world by establishing barriers across the Internet ecosystem. The result is likely to be a significant degradation of network performance and

increase in costs. In all likelihood, these costs will ultimately be passed on to Internet users.

Effect:
This threat undermines three of the five fundamental properties of the global Internet. There's a risk that this policy could spread to other countries, creating an even greater risk to the network and everything that relies on it.

Talking points:
A national Internet gateway is typically a government-mandated gateway through which all Internet traffic is routed. It centralizes control over all local and international Internet traffic – both incoming and outgoing.

With a national Internet gateway in place, networks must connect to government-mandated locations, increasing the technical and financial barriers that network operators face in becoming part of the global Internet.

A national Internet gateway is likely to result in a significant degradation of network performance and increase in costs. In all likelihood, these costs will ultimately be passed on to Internet users.

Threat 3: Creating Walled Gardens

Walled gardens are closed ecosystems created by large companies, where users are confined to a specific set of services, applications, or products offered by those companies. These closed environments limit user choice and interoperability with third-party services. The companies behind these walled gardens control the entire user experience, often prioritizing their own products and services over competitors. This can lead to a lack of innovation, reduced competition, and a stifling of diversity in the digital landscape.

Such walled gardens can lead to monopolistic practices, hindering fair competition and limiting the potential for smaller, innovative companies to thrive. As a result, users may find themselves with limited options, facing higher prices, and experience a lack of privacy and control over their digital experiences.

Example:
Zero-rated content. Some telecom operators and ISPs offer their customers a deal where some types of content are not counted against their data plan. For example, they might offer access to certain apps, streaming services, or messaging platforms, as much as users want, for no extra cost. This is referred to as a zero-

rated Internet service. It's often framed as a benefit for customers, and is part of how they compete with other providers.

While it is advertised as offering efficient costs and improved access, zero-rated services often are the result of commercial alliances between large platforms and telecom operators. The practice raises concerns about net neutrality and competition. It also increases the risk that service providers will disproportionately shape users' preferences, and even their opinions.

With zero-rated services, people sitting next to each other on a train, for example, might have a very different experience. People may not be reading the same news or listening to the same music – or even be able to send a message to one another because of access to only certain platforms. If there is one free social network, one free news provider – and so on – these can become what the Internet is for many users. It's a fragmented experience of the Internet. This practice creates new Internet gatekeepers, who have the power to shape our tastes and preferences by creating a walled garden.

Effect:

Zero-rating violates one of the core elements of net neutrality because all traffic is not treated equally. This is not just the case with technical traffic management, but also applies to the commercial practices of Internet Service Providers (ISPs). It fosters potential market consolidation by already dominant players, and has a negative impact on the open Internet. It undermines the open, global Internet and harms innovation and growth.

Talking points:

If your mobile provider offers you free online services as part of your data plan, it might not affect you directly, but it affects others. For those who cannot afford to pay for extra data, this means a fragmented Internet experience, based on who you are, where you live, and how much you can pay.

The people who need the Internet the most, and benefit most from the opportunities it provides, are the most adversely affected by zero-rating practices.

Zero-rated services harm net neutrality by favoring certain services over others. And this, in turn, increases the risk of the market consolidating in the hands of large telecom providers, ISPs, and large platforms. This distorts the market and fair competition.

EuroDIG discussions: zero rating

The issue of 'zero rating' was discussed in-depth at EuroDIG 2016 in Brussels. A workshop entitled *"Zero rating: what is it?"* considered the definition of zero rating in order to achieve a common understanding of the issue amongst stakeholders. The discussions at a follow-up session on the impact of zero rating illustrated the diversity of views on zero-rated services, with some interventions highlighting how zero rating harmed Internet neutrality, innovation, and connectivity, and fragments the Internet user's experience.

Sources:
https://eurodigwiki.org/wiki/Zero_rating_what_is_it%3F_%E2%80%93_WS_07_2016,
https://eurodigwiki.org/wiki/Impact_of_zero_rating_%E2%80%93_WS_07_follow_up_2016

Threat 4: Regulation of DNS Infrastructure

DNS is a system that translates domain names into numeric IP addresses. When you type a website address or open an app, DNS is like a digital phone book that makes sure your device connects with the information you're looking for. If you have consistent DNS resolution, it means that everyone who uses the Internet sees the same things, wherever they are in the world. An Internet Service Provider (ISP) usually chooses a DNS resolver for its customers, so most users don't need to think about it. And if there's an outage or attack, the traffic can be rerouted, making the system resilient.

Some DNS resolvers filter certain domain names that are known to be malicious, such as links to malware. But when DNS resolvers start to filter specific types of content, this poses a threat to the Internet overall. It means that the resolver is no longer just a translator.

This creates a scenario where what you can see and do online depends on what DNS resolver is being used – something you don't always control.

Example:
Russia's National DNS. In 2019, the Russian government passed a law titled 'Sovereign Internet'. As part of this legislation, a National DNS was created, which replicates the global DNS. The stated reason was that Russia wanted to mitigate the threat of being disconnected from the DNS (specifically something called the global root).

An Autonomous System (AS) is a collection of networks of IP numbers that are managed by a single entity, all following the same set of rules. Everyone who operates an AS in Russia is required to connect to the NDNS and perform their name resolution through it. They are expected to use a local root server, which will give a government-approved backup copy of the root zone. Or, they could use a public National DNS resolver, directly, or through the network's own resolvers. Name resolution is normally done by a global DNS provider. Several companies have already been fined for failure to connect to the Russian National DNS. This approach to domain name resolution fragments the global Internet.

Effects:
Russia's National DNS is based on an approach that fundamentally fragments the global DNS, and, as a result, undermines and fragments the global nature of the Internet itself. It uses what's called an alternative root. Because all DNS requests have to go through this system, it can be used as a tool for censorship and surveillance, violating citizens' privacy and security.

Even though this policy's stated aim is to mitigate the threat of being disconnected from the global DNS, this approach also creates a single point of failure. It could affect the availability, performance, and resilience of the DNS resolution service.

Talking points:
An alternative DNS system threatens to fragment the Internet, creating an alternative name space that could make it difficult to even connect to the global Internet again.

Mandated routing such as National DNS creates a nationwide censorship and surveillance tool. It enables surveillance of Internet traffic, and can cause users to be blocked or redirected to specific content.

An alternative DNS system like this creates a single point of failure, which affects the availability, performance, and resilience of DNS resolution in the country.

It's important not to disconnect any country from the global DNS. Not only does it fragment the Internet, it makes it even more difficult for people who need information the most to get access to it.

EuroDIG discussions: DNS infrastructure and impact of regulation

The stakeholder community discussed in a session at EuroDIG 2022 in Trieste how EU proposed regulations that were under consideration at that time could potentially impact domain name system (DNS) infrastructure. The debate mentioned the possible fragmentation of the DNS as a consequence of some of the EU proposals, including one on regulating root name servers in the EU NIS2 directive. While the language on root zone servers had by that time already been removed from the draft directive, the discussions reiterated the dangers of setting this as a precedent.

Source:
https://eurodigwiki.org/wiki/Digital_sovereignty_impact_on_the_Internet_infrastructure._%E2%80%93_FA_01_Sub_02_2022

Threat 5: Blocking Security Technologies

Encryption keeps our messages and information safe when we send them over the Internet. It scrambles the words and pictures into a format that only the person with the right key can understand. End-to-end encryption is the strongest form of encryption, where only the sender and recipient have access to what has been shared – even the platform does not know the content of the messages.

With the amplification of harms online, governments have been trying to address crime, hate speech, and harmful content online. Gov-

ernments argue that having access to encrypted content is necessary for safety and public security, but weakening or breaking encryption opens the door for bad actors, like hackers or other governments, to exploit the same vulnerabilities. Diluting encryption also puts personal, business, and national security at risk. And with lack of security, it will be impossible to build on users' trust.

Example:

UK Online Safety Act. The UK Online Safety Act is a law that requires platforms to scan content for material that's deemed harmful or exploitative. It also requires them to take action against it. The lawmakers say that their aim is to make the UK the safest place in the world to be online, and the best place to run a digital business. They argue that this law will address a wide range of harms that stand in the way of that, including child sexual abuse material, violence against women and girls, underage access to pornography, and misinformation.

Companies that offer user-to-user communication, allow user-generated content, recommended algorithms, search functions, or services that might be accessed by children would be held liable for content on their platforms.

The Act does not specifically mention encryption in relation to content sent between individual users, but the demands in it would undermine this by default. The providers have to guarantee that, even if the communications are encrypted, either they or law enforcement will be able to get access to it. It also demands that providers are able to reliably verify or estimate a user's age. The law requires that they use 'accredited technologies' to do so, which the UK government itself has acknowledged do not exist.

The Act also applies to any service that can be accessed by UK users, no matter where that company is based, even if it's outside the UK. This means that the Online Safety Act is likely to have an extra-territorial effect, which leaves companies in the position where they have to choose between offering services in the UK, or providing end-to-end encryption.

If companies that offer encryption are forced to cut off UK users, then this would not only cut UK users from truly encrypted services, it also cuts them off from people who use those services outside the country. If these services choose to stay in the UK and comply with the law, people who communicate with those in the UK would also lose their confidentiality.

This makes the UK a less safe place to run a digital business. The Online Safety Act undermines confidence and trust in services developed in the UK. Because anyone implementing those services will have to assume those

products are complying with it – and all that goes with it. The overall prospect is that the UK would become an untrustworthy endpoint in any digital communication, both inside the country, and across borders.

Effects:

The Online Safety Act fails the tests of necessity and proportionality. Its requirement to either undermine or circumvent encryption makes the Internet less secure for everyone, but is unlikely to deliver the intended protection for children and other at-risk individuals.

This Act contains multiple negative impacts on the open, global, and secure Internet. These include constraints on ease of access and permissionless use of endpoint technologies; reduction in reachability, if UK endpoints must be considered untrustworthy; impact on data confidentiality; and an increase in the attack surface available to malicious third parties and hostile governments, with corresponding impact on resilience and cybersecurity.

Talking points:

The Online Safety Act compromises the security of all UK users. This includes children it claims to protect. Even senior national security professionals from UK agencies have said that the benefits of confidential communication far outweigh the risks of harm, but the government has ignored their advice.

Not only does the Online Safety Act affect people in the UK, it would also undermine the security and confidentiality of anyone who communicates with people in the UK. This means it has an extra-territorial effect, imposing the laws of the UK onto people outside of its borders.

Companies offering end-to-end encrypted services have maintained that they will not be compromising on the security of their platforms, and would rather quit the UK market if they are compelled to undermine end-to-end encryption. People in the UK would have access to less secure online services. These less secure services are more vulnerable to cyber attacks, hackers, and bad actors who aim to harm children and vulnerable people – the same ones this law claims to deter.

Rather than make the UK the best place in the world to do business, the Online Safety Act would undermine global trust in UK businesses, since they would have to assume that products and services made there comply with this law. This means that people who value privacy or need encrypted services might be less likely to do business with a UK company.

EuroDIG discussions: security and encryption

The UK's proposals in its Online Safety Bill were discussed in the broader context of user safety and encryption during the Virtual EuroDIG 2021 in a session entitled *"Crypto Wars 3.0 – can privacy, security and encryption co-exist?"*. The discussion highlighted how trust in encrypted communications and the safety of users would be undermined by the access of authorities to encrypted messages. The agreed EuroDIG messages emphasized that so-called "backdoor" access to encrypted communications would not prevent all criminals from encrypting their communication in an unbreakable way, but it did create a risk of generally weakening everybody's security online and possibly make the solution worse than the problem created by a criminal minority. The session also called for avoiding the false framings and false dichotomies around encryption, privacy, and security.

Source:
https://eurodigwiki.org/wiki/Crypto_Wars_3.0_%E2%80%93_can_privacy,_security_and_encryption_co-exist%3F_%E2%80%93_WS_05_2021#Messages

Threat 6: Digital Sovereignty

Digital sovereignty is a broad, umbrella concept that can include governments that wish to control how Internet operations and resources are run; local businesses that decry the dominance of foreign tech platforms; indigenous communities that want to safeguard local knowledge and resources; and individuals who want to assert their autonomy over their interactions with devices, platforms, and how they manage their data.

Policies that have elements of digital sovereignty may adversely affect how the Internet

works and undermine our ability to make use of the Internet.

Example:
Indian CERT Cybersecurity Directions. Every Internet service relies on the correct time to maintain secure, compliant systems, especially where systems and users are spread across broad geographies.

Therefore, everything on the Internet connects to a Network Time Protocol (NTP) server. This is how devices and applications determine and coordinate time across distances, devices, and connections. This is how your phone automatically resets the time when you enter a new time zone. There are around 3,000 publicly available NTP servers around the world. Connecting to multiple NTP servers means more resilience and accuracy. This is considered an industry best practice.

The Indian government mandates that all entities covered under the Indian Computer Emergency Response Team's (CERT-In) Cybersecurity Directions must connect to two government-controlled NTP servers. These are the National Informatics Centre and the National Physical Laboratory.

Even if there's no malicious intent, it's important for time servers to be aligned. For example, if you have time servers that aren't coordinated, and the discrepancy is large enough, you wouldn't know the correct time, so you might not show up for a meeting, or know that you're about to miss your flight. Being able to see the correct time on the user side is important, but it can make things even more complicated on the back end.

Even tiny misalignments can be catastrophic for financial transactions, which rely on time that's accurate to the millisecond, or cybersecurity. Correct time logs are important for spotting and responding to attacks, which means that if a time log is off somewhere, a legitimate interaction or transaction could be treated as malicious. This type of disruption could be difficult to track, and could lead to widespread problems for Internet users and providers.

Even a lag in one of the NTP servers can reverberate across the Internet, and undermine its resilience globally.

Effects:
Internet Society carried out an impact brief, and wrote to CERT-In and the IT ministry. CERT-In should reconsider its one-size-fits-all approach and respect the decentralized nature of the network, and the long-established practice of depending on multiple NTP servers for the time. https://www.internetsociety.org/resources/doc/2022/internet-impact-brief-india-cert-in-cybersecurity-directions-2022/

Talking points:

The Indian government requires that all entities covered under these directions must connect to two government-mandated NTP servers at the National Informatics Centre and the National Physical Laboratory.

There are around 3,000 publicly available NTP servers around the world. Connecting to multiple NTP servers means more resilience and accuracy. This is considered an industry best practice.

CERT-In should reconsider its one-size-fits-all approach and respect the decentralized nature of the network, and the long-established practice of depending on multiple NTP servers for the time.

EuroDIG discussions: digital sovereignty

The EuroDIG community has discussed digital sovereignty at many recent meetings, linking it in particular to the risks of Internet fragmentation. This issue was debated during EuroDIG 2016 in Brussels in a session entitled “Internet fragmentation and digital sovereignty: implications for Europe”. Stakeholders emphasized the need to protect global connectivity on the grounds that the economic interests of nation-states relied on an open Internet. At EuroDIG 2022 in Trieste, the discussions in a session entitled “Digital sovereignty – is Europe going in the right direction to keep the Internet safe and open?” highlighted that regulation relating to digital sovereignty should avoid collateral damage to the services and operators regarding economic costs and availability, and mitigate any possible risk of fragmentation of the critical infrastructure of the global Internet.

Sources:

https://eurodigwiki.org/wiki/Internet_fragmentation_and_digital_sovereignty:_implications_for_Europe_%E2%80%93_PL_04_2016#Messages;

https://eurodigwiki.org/wiki/Digital_sovereignty_%E2%80%93_is_Europe_going_in_the_right_direction_to_keep_Internet_infrastructure_secure_and_open%3F_%E2%80%93_FA_01_Sub_03_2022#Messages

Threat 7: Internet Shutdowns

An Internet shutdown is an intentional disruption of Internet or electronic communications, rendering them inaccessible or effectively unusable, for a specific population or within a location, often to exert control over the flow of information.

Any disruption to the Internet, no matter its duration, has a detrimental effect on the economy and employment. Internet shutdowns disrupt productivity, prevent e-commerce from happening, generate monetary losses in time-sensitive transactions, and increase unemployment. Internet shutdowns erode trust in the Internet as a reliable and neutral space for communication, collaboration, and the exchange of ideas, undermining its potential as an engine for positive social change.

Example:
Internet Shutdowns during Exams. Some governments shut down the Internet during exam time, to prevent cheating. But it is disproportionate, and it cheats everyone. Internet access is an essential part of life for many people around the globe – but for millions of people, things become extremely hard during the school exam period. Every year, governments around the world – including Sudan, Jordan, Iraq, Algeria, Syria, and India – make the draconian decision to implement Internet shutdowns during school or competitive exam periods.

In many countries, general public exams determine much of what a student is able to do later on, such as what university they can go to and which subjects they are qualified to study. In some countries, cheating on these exams can mean jail time, but the pressure is so huge on these young people that some will resort to cheating anyway.

Governments in these countries believe that this temporary shutdown will help prevent students from cheating or sharing exam questions, but it does more than shut millions

of students off from their communication, and even their study material.

Exam-related shutdowns result in public outrage, since they also cause essential social, political, economic, and cultural activities to come to a standstill. They particularly impact women and gender minorities. A shutdown means being unable to communicate with loved ones, access money, or even check a transit schedule. Internet shutdowns are never a proportionate response to anything, no matter how long they last. Even if a shutdown were to prevent exam cheaters from communicating, it also prevents everyone else from using online services. It is not an effective anti-cheating mechanism, and it comes at a cost to all of society.

In addition to the direct and immediate effect of losing access to online services, shutdowns normalize disruptions to access. It makes those disruptions seem, not just normal, but even accepted. This damages the reliability of and trust in the Internet, which is of vital importance for socio-economic development everywhere.

Rather than prevent cheating, Internet shutdowns during exam periods cheat everyone out of their ability to access essential services. Every shutdown also reduces resilience, and damages confidence in the availability of the Internet as a global resource.

Effects:

Internet shutdowns harm societies and economies, and they reduce the reliability of the global Internet infrastructure. We urge governments and decision-makers everywhere to support policies that keep the Internet on and strong, in order to build robust and resilient economies and give people the opportunity to build their own prosperous futures. Internet shutdowns are disproportionate, and there is no evidence to suggest that cutting off access to the Internet during exams actually prevents cheating.

Talking points:

Internet shutdowns are sometimes used as a swift, sweeping response to a non-technical problem, like preventing students from cheating during exam periods. But even if these tactics did prevent cheating – and there is no evidence that they do – cutting off everyone's access to the Internet is an extremely disproportionate response.

Every shutdown not only cuts off people in a region from the Internet, it also cuts the world off from that region. Many web services use backend components based in multiple geographies, so even a company outside the area could be affected by the shutdown. For the Internet to be open, secure, and trustworthy, it needs to be global, and every shutdown makes it less global.

Reliable access to the Internet is especially important for marginalized individuals and groups, and cutting off their access means they are even more constrained. It becomes even more difficult to communicate, travel, or do business safely and securely

Threat 8: Limiting Global Access

In an interconnected world, information and ideas transcend borders. An open and global Internet is an important engine for a globally informed, engaged population. But the Internet is increasingly politicized, especially during geopolitical conflicts and tensions.

Limiting or blocking access to the global Internet is a disproportionate response, and it causes widespread harm. Blocking services and applications limits global access for people who need connectivity to freely communicate, access information, and participate in the global digital space.

Example:
Disconnecting Countries from the Global Internet in times of war. In times of crisis or conflict, access to the open Internet becomes more important than ever. Not just for accurate information, opportunity, and communication, but also for hope and reassurance. Any restrictions imposed on the Internet will only be a burden to people in these countries.

The war in Ukraine spurred a wide range of reactions around the world, focused on ways to discourage Russia's attack. In February 2022, officials in Ukraine asked two bodies, ICANN and RIPE NCC, to take action against Russia. RIPE NCC is an Internet registry responsible for allocation of addresses. ICANN is a non-profit, multistakeholder organization that's responsible for securing some of the Internet's vital operations.

They asked ICANN to revoke the rights of Russia's ccTLDs – the domain name endings .RU, .SU and .рф, and to shut down root servers in Russia. They also asked RIPE NCC to revoke the rights of Russian members to IPv4 and IPv6 addresses.

At the same time, countries were also issuing sanctions against Russia that had the effect of disconnecting parts of the Internet's infrastructure from Russian networks. This had an impact on Internet infrastructure companies and Internet exchange points. ICANN and RIPE NCC refused the request because it conflicted with their policies and their multistakeholder approach.

When a country's networks are disconnected from the Internet, there are a number of unintended consequences – for the people in that country as well as for the Internet more broadly.

It hinders the reliability, resilience, and availability of the Internet, and can disrupt the network's operations beyond borders. This country-level disconnection splinters the Internet along geographical, political, commercial, and technical lines. It also sets a dangerous precedent that undermines trust in the Internet's multistakeholder governance processes and sets a precedent of using the Internet as a geopolitical weapon.

At the societal level, it can undermine the use of the Internet at a time when people need it most, to find accurate information and access means of safety. Rather than another dimension of accountability for nation states, this approach legitimizes a dangerous playbook. Other regimes might use it to control Internet access in future conflicts.

Effects:

We cannot let the Internet become a pawn of geopolitics. Politicizing decisions about the Internet's inner workings sets a dangerous precedent that puts us on the fast track to a 'splinternet' – an Internet artificially carved up along political, economic, and technological boundaries. The effects may be irreversible, opening the door for further restrictions across the globe. Governments, businesses, and organizations worldwide must ensure that the day-to-day technical governance of the Internet is not politicized. It is vital that the management and operations of Internet infrastructure, including the naming, addressing, routing and security systems, remain apolitical. Sanctions should not disrupt access to and use of the Internet. Where needed, sanctions regimes should offer exemptions to ensure continued service of Internet infrastructure.

Talking points:

In times of war and conflict, governments, service providers, and other organizations are increasingly considering actions that could irreversibly damage the global Internet. Even when there are valid reasons for sanctions, we cannot support decisions that would harm the very nature of the network.

We need to prevent geopolitical decisions from fragmenting the Internet. Otherwise, we will increasingly lose the global Internet we rely upon to attempts to carve it up along political, economic, technological boundaries. This will significantly harm our ability to communicate, create, and connect.

Supporting victims of geopolitical conflict means ensuring free and open access to the Internet, which is a critical lifeline for civilians.

EuroDIG discussions: keeping the Internet global in times of war and conflict

The issue of disconnecting countries from the global Internet in times of war was debated in the session entitled *"Shattered Neutrality: Internet at Crossroads of War and Geopolitics"* at EuroDIG 2023 in Tampere. This discussion considered the decision of ICANN not to disconnect Russia from the DNS following its illegal invasion of Ukraine. ICANN's decision was justified on the grounds that no single actor had the power or authority, or the right, to remove people or an entire country from the global Internet. EuroDIG's messages from this session highlighted the fundamental problems caused by actions that interfere with the Internet's operation as a global communications infrastructure, such as revoking the delegation of a top-level domain (TLD) or an IP address. EuroDIG called for the Internet to be maintained as a single, globally interoperable space for communication, free from any disruption caused by geopolitical tensions.

Source:
https://eurodigwiki.org/wiki/Shattered_Neutrality:_Internet_at_Crossroads_of_War_and_Geopolitics_%E2%80%93_TOPIC_01_Sub_01_2023

The Internet community can do a lot to combat these threats. This document has been compiled to explain what the Internet is, what makes it different from other types of networks, which are its Critical Properties and the Enablers that make it thrive, and inform the Internet community how to create an Internet Impact Assessment to determine the nature of the threats.

One important action is to monitor policies and decisions locally, have your eyes open to what is happening and being proposed in your part of the world and flag them online, in the media, or at events like EuroDIG. Spread the

word. Talk to like-minded people about an emerging threat that may need to be monitored and advocated against.

Study in detail the eight types of existing threats described in this document, and the talking points. Analyze any threat you spot using the Internet Society's Internet Impact Assessment Toolkit to produce a brief that can be used as evidence.

Write to your policymakers, engage with them using public consultation processes, or reach out to them online.

Ultimately, every voice matters, and every action counts. Talk and raise awareness about the threats the Internet is facing, because together we must protect it for generations to come.

Acknowledgements

We want to thank the people who worked in the development of the Internet Impact Assessment Toolkit (IIAT): Andrew Sullivan, Joseph Lorenzo Hall, Carl Gahnberg, Natalie Campbell, Katie Watson Jordan, Konstantinos Komaitis, Andrei Robachevsky, and Olaf Kolkman.

This section would not be complete if we missed expressing gratitude to Leslie Daigle, for the work on the *Invariants*, the precursor concept that led to the Critical Properties and the Enablers.

This publication would not have been possible without the participation of Tatiana Tropina, David Frautschy, Nadia Tjahja, Mark Carvell and the EuroDIG team: Sandra Hoferichter and Rainer Rodewald.

Imprint

Published by:
EuroDIG Support Association
Schächlistrasse 19, CH-8953 Dietikon
email: office@eurodig.org
web: www.eurodig.org

Assistant Editor: Mark Carvell

This publication is a joint project of Internet Society and EuroDIG.

Graphic and production: monade · agentur für kommunikation GmbH, Leipzig
2024