

A Responsible Internet to Increase Trust in the Digital World

Cristian Hesselman^{1,2}, Paola Grosso³, Ralph Holz², Fernando Kuipers⁴, Janet Hui Xue⁵, Mattijs Jonker², Joeri de Ruiter¹, Anna Sperotto², Roland van Rijswijk-Deij^{2,6}, Giovane C. M. Moura^{1,4}, Aiko Pras², and Cees de Laat³

¹ SIDN Labs, the Netherlands

² University of Twente, the Netherlands

³ University of Amsterdam, the Netherlands

⁴ Delft University of Technology, the Netherlands

⁵ Wolfson College, Oxford University, UK

⁶ NLnet Labs, the Netherlands

cristian.hesselman@sidn.nl

Abstract. Policy makers in regions such as Europe are increasingly concerned about the trustworthiness and sovereignty of the foundations of their digital economy, because it often depends on systems operated or manufactured elsewhere. To help curb this problem, we propose the novel notion of a responsible Internet, which provides higher degrees of trust and sovereignty for critical service providers (e.g., power grids) and all kinds of other users by improving the transparency, accountability, and controllability of the Internet at the network-level. A responsible Internet accomplishes this through two new distributed and decentralized systems. The first is the Network Inspection Plane (NIP), which enables users to request measurement-based descriptions of the chains of network operators (e.g., ISPs and DNS and cloud providers) that handle their data flows or could potentially handle them, including the relationships between them and the properties of these operators. The second is the Network Control Plane (NCP), which allows users to specify how they expect the Internet infrastructure to handle their data (e.g., in terms of the security attributes that they expect chains of network operators to have) based on the insights they gained from the NIP. We discuss research directions and starting points to realize a responsible Internet by combining three currently largely disjoint research areas: large-scale measurements (for the NIP), open source-based programmable networks (for the NCP), and policy making (POL) based on the NIP and driving the NCP. We believe that a responsible Internet is the next stage in the evolution of the Internet and that the concept is useful for clean slate Internet systems as well.

Keywords: trust, digital sovereignty, responsible Internet, cybersecurity, transparency, accountability, controllability

1 Introduction

The Internet has evolved from a local network for a small group of experts in the early 1970s to a global, continuously evolving infrastructure that supports a wide range of services and products that almost all businesses, governments, and citizens depend on today, even more so after the 2020 Covid-19 outbreak.

However, policy makers in regions such as Europe are increasingly concerned about the trustworthiness and sovereignty of the foundations of their digital economy [ENISA19] [EIT20] [GAIAX20a], because it often depends on systems manufactured or operated elsewhere. For example, the European Union Agency for Network and Information Security (ENISA) recently articulated their concerns about Europe’s “digital sovereignty” [ENISA19]. They point out that the top 15 Internet companies in the world (e.g., Google, Facebook, and Alibaba) are either from the US or from China and not one of them from Europe. In addition, they highlight that European tech companies often get acquired by non-European companies (e.g., 53 were bought by US “tech titans” in 2011-2016). The risks they associate with these developments include European service providers and citizens losing control over their data and cybersecurity facilities, Europe no longer being able to meet their citizens’ norms and expectations, reduced competitive power, and drain of technical expertise [ENISA19].

While European policy makers are trying to curb this problem through new policy proposals (e.g., for Artificial Intelligence, 5G cellular networks, and the Internet of Things) [EIT20] and initiatives such as a European federated cloud service [GAIAX20a], we observe that the Internet infrastructure has not received much attention yet in this context, except in an ad-hoc way, such as following reports on alleged security weaknesses in 5G equipment [NYT20]. We believe this is a serious omission, because ultimately trust and sovereignty also require service providers and product manufacturers to be in control of their dependencies on the Internet infrastructure, specifically when it comes to security and resilience. This is particularly relevant for critical service providers (e.g., power grids, transportation systems, mobile networks, and manufacturing facilities), which have become increasingly dependent on computer networks [WODC19]. For example, such providers want to know if they are routing their traffic through networks with equipment that might have backdoors [NYT20].

To fill this void, we propose the novel notion of a **responsible Internet**, which aims to provide a higher degree of trust and sovereignty for critical service providers and all kinds of other users by making the Internet more transparent, accountable, and controllable at the network-level. This means users have (1) insight into the security attributes of chains of network operators (e.g., ISPs, data centers, and DNS operators) that carry or could potentially carry their data flows (transparency and accountability) and (2) are able to use these details to send their data flows through certain classes of network operators or request network operators to change the way they handle these flows (controllability), perhaps by changing the infrastructure itself with the help of policy makers. We believe that improving the Internet’s transparency, accountability, and controllability is key for users to trust the network and to be in control of their dependencies on the Internet infrastructure (and thus to be sovereign).

Without a responsible Internet, users will continue to be subjected to the current “black box Internet”, which has weak transparency, accountability, and controllability properties. For critical service providers, for example, this means they will mostly not know which network operators their services depend on and they will have little control over which classes of operators they would prefer to carry their traffic (e.g., based on these operators’ security posture). As a result, they may prefer to continue to use their own dedicated networks rather than a shared global Internet, which ultimately limits their flexibility and increases costs.

From a technical perspective, a responsible Internet introduces two new distributed and decentralized systems. The first is the **Network Inspection Plane** (NIP), which improves the Internet’s transparency and accountability by allowing a wide range of users to request measurement-based descriptions of the Internet infrastructure in terms of the chains of network operators that handle their data flows, their interrelations, and their properties (e.g., their jurisdiction and if they use open source router software). The second is the **Network Configuration Plane** (NCP), which increases the Internet’s controllability by allowing users to specify how they expect the infrastructure to handle their data based on the NIP’s network descriptions.

We complement the NIP and NCP with a set of **policies** (POL) that help shape the network in the longer term, for instance to incentivize operators to use open source control and data plane software for security reasons or to join the NIP.

We make **two contributions**. The first is that to the best of our knowledge we are the first to define the concept of a responsible Internet, its properties (transparency, accountability, and controllability) and its key architectural components (NIP, NCP, and policies). Our second contribution is that we discuss research directions and starting points to realize a responsible Internet by combining three currently largely disjoint research areas: large-scale measurements for the NIP, open source-based programmable networks for the NCP, and policy making (POL) using insights gained from the NIP.

Our properties of a responsible Internet are inspired by those of responsible Artificial Intelligence (AI) [Dignum17]. The parallel with a responsible Internet is that policy makers worry about society’s level of control over AI systems as well [EC19] because they are black boxes similar to the Internet and because they also shape societies [Dignum17] (e.g., through AI-based parole and air pollution decisions [Rudin19]). The European Commission recently embraced a concept similar to responsible AI in their four principles for trustworthy AI [EC19], which they call the “explicability principle”.

Similar to responsible AI, a responsible Internet introduces a wide range of cross-disciplinary challenges, for instance in the fields of ethics, education, legislation, business models, and technology. While our paper focuses on technical challenges, our **goal** is to trigger and facilitate a wider, cross-discipline dialogue on a responsible Internet among stakeholders such as researchers, citizens, operators, and policy makers as well as to provide guidance for new research directions.

We think of a responsible Internet as the next stage in the evolution of the Internet, building on earlier and ongoing developments to increase the Internet infrastructure’s security (e.g., through technologies such as DNS-over-HTTPS, DNS security extensions, and a public key infrastructure for the routing system), resilience (e.g., through anycast), and privacy-friendliness (e.g., through DNS query name minimization). The concept can also be applied to “clean slate” infrastructures, such as based on SCION [SCION17] or RINA [Maffione16].

We emphasize that our vision of a responsible Internet continues to follow the Internet’s open, bottom-up, and multi-stakeholder nature. Our notion of sovereignty is about service providers and individuals being in control of their dependencies on the Internet infrastructure and is explicitly not about creating government-controlled or even isolated national networks (cf. the “Beijing Internet” or the “Moscow Internet” [O’Hara20]), nor is it about excluding technologies from specific regions [EIT20].

In the rest of this paper, we first briefly outline scenarios to illustrate the added value of a responsible Internet for various types of users (Section 2). Next, we discuss the design of a responsible Internet (Section 3) and the research directions we identified for the NIP (Section 4), the NCP (Section 5), and the new policy mechanisms (POL) that a responsible Internet enables (Section 6). We continue with a discussion on the “Internet trust transition” that we think a responsible Internet facilitates (Section 7) and wrap up with an overview of related work (Section 8) and our conclusions (Section 9).

2 Illustrative Examples

The purpose of a responsible Internet is to provide a higher degree of trust and sovereignty for a broad range of users. In this section, we illustrate what this entails using four simple scenarios: critical infrastructure providers (Section 2.1), policy makers (Section 2.2), network operators themselves (Section 2.3), and individuals (Section 2.4).

We envision that critical infrastructure providers, policy makers, and network operators will initially benefit the most from a responsible Internet. Individuals might benefit as well but will need novel user interfaces and additional guidance to enable them to navigate the network descriptions that a responsible Internet provides.

2.1 Critical Infrastructure Operators

One of the key beneficiaries of a responsible Internet are critical infrastructure providers such as power grid operators and providers of intelligent urban transport systems. They benefit because a responsible Internet gives them more control over their dependencies on the network, which is essential to protect the security of their services and prevent large-scale incidents such as data breaches and safety risks for large groups of citizens.

As an example, consider a provider of a smart grid that sends flows of instructions to remote field stations to control power line switches [Chromik19]. In a responsible Internet, the grid provider can request the network to provide a description of how these flows travel through different networks, what type of equipment is used along the path, who operates the networks, and if any operations are outsourced to other networks (transparency and accountability).

In addition, the grid provider can request the responsible Internet to select an alternative network path [SCION17] (controllability) based on the network descriptions it obtained earlier, for instance because they reveal that some network operators use equipment that might have software vulnerabilities (e.g., alleged back doors [NYT20]). It can also use these descriptions to work with policy makers to request enduring changes through regulatory action (see Section 2.2).

2.2 Enabling New Internet Policies

The network descriptions that a responsible Internet provides enable new types of policy making (developing the principles for a policy), policy mediation (translating the principles laid out in a policy to concrete, actionable steps), and policy enforcement

(ascertaining that the steps are indeed implemented), which are three typical steps of policy development.

Policy making. While classic policy making relies on consolidating input from stakeholders and taking interests and capabilities into account, a responsible Internet enables policy makers to take a more data-driven and proactive approach based on network descriptions. For example, a responsible Internet enables national policy makers to: (1) analyze risk areas in their local Internet infrastructure (e.g., concentrations of power or single points of failure [Arkko20]) based on historical data analysis; (2) infer models that help them play out realistic what-if scenarios; and (3) develop new regulatory strategies (e.g., to protect Europe’s digital sovereignty [ENISA19] [EIT20]).

Policy mediation. We expect that a responsible Internet will enable regulators to act much faster upon emerging problems and risks, saving costs in litigation. For example, they could feed network descriptions into a platform that facilitates evidence-based feedback between parties. Also, critical infrastructure providers such as power grids and transportation systems can voice their concerns based on network descriptions obtained from a responsible Internet. For example, they can indicate that more network operators need to adhere to the “Mutually Agreed Norms for Routing Security” (MANRS) [MANRS] to properly protect their services. In turn, regulators can judge by the outcomes (e.g., tracked configuration changes) to determine whether further investigation or intervention is required.

Policy enforcement. Regulators benefit from network descriptions because they support data authenticity through cryptographic proof (see Section 4), which will help solve enforceable liability with respect to operators and third-party vendors (e.g., for operators of IoT services [Leverett17]). Regulators are able to understand in which society the operator is embedded (e.g., in terms of safety, privacy, freedom of speech, and laws for corporate and state surveillance [Brown10]). The network descriptions present a useful interactive map for stakeholders who can efficiently identify the issues and the associated responsible parties, especially when Internet infrastructure is attacked. Law enforcement authorities can use the graph to map out key operators and identify areas for further investigation.

2.3 Enabling Cross-Network Operator Incident Analysis

Another class of users of a responsible Internet are network operators themselves. For example, an operator that measures the properties of the DDoS attacks that it handles (e.g., Mirai-based DDoS attacks [Mirai17] or incidents similar to the 2015 DDoS attack on the DNS root [Moura16]) can include the metadata of these datasets in network descriptions along with a usage license. The advantage is that it becomes much easier for other operators to find such datasets and the licenses to use the data [Koning19]. This enables them to collaboratively combine and learn from each other’s measurements, which improves their collective incident response capabilities.

Ultimately, we envision that a responsible Internet enables the real-time sharing of measurements across network operators, allowing them to collaboratively fend off security incidents as they occur (e.g., by dynamically moving scrubbing functionality to a specific part of their network using Network Function Virtualization (NFV) [Koning19]) or even proactively before they can cause real harm. Network operators

could share the actual measurements in various ways, such as directly from their own servers or through a shared platform in which multiple operators upload their measurements (e.g., DDoS fingerprints [Ceron20] [Coreflow18]).

2.4 Giving Individuals more Insight in and Control over their Data

In the long term, we expect individuals to benefit from a responsible Internet as well. For example, people using video conferencing services (e.g., Zoom) could request a network description from a responsible Internet, which enables them to verify where their video flows end up and potentially change their endpoint to a data center in another region.

The Covid-19 pandemic of 2020 illustrated the relevance of this type of scenario. With lockdowns enacted in many countries, Zoom [Zoom20] emerged as one of video communication tools of choice. Confronted with a list of security issues, governments soon warned against using the software [Zoom20a]. Among the cited concerns, the storage of cryptographic material in data centers outside “friendly” jurisdictions was considered problematic. Zoom reacted to this by allowing their customers to choose which data centers they wanted to connect to.

In a responsible Internet, these kinds of facilities would be built into the network infrastructure and thus be available for all applications, including Zoom.

3 Designing a Responsible Internet

In this section, we present the outlines for the design of a responsible Internet, which builds on two new distributed and federated systems.

We discuss our notion of the qualifier “responsible” (Section 3.1), our proposed design goals (Section 3.2), our high-level architecture to realize these goals (Section 3.3), and the technical blueprint of network operators in a responsible Internet (Section 3.4).

This section is the starting point for the research directions that we foresee and elaborate on in Sections 4 through 6.

3.1 Origin and Meaning of “Responsible”

Our notion of a responsible Internet is inspired by work of the responsible Artificial Intelligence (AI) community, which focuses on giving people more insight in how AI systems reach decisions and why [Rudin19]. This is important because AI systems can have a profound impact on people’s lives. For example, there have been known cases in which AI systems incorrectly denied people parole or miscalculated air pollution levels [Rudin19]. Responsible AI extends the design and operation of AI systems with three design goals (transparency, accountability, and responsibility [Dignum17]) that help researchers, developers, and AI operators to consider the impact of their work “by design” (e.g., in terms of ethics) and not only focus on the predictive performance of their algorithms (e.g., in terms of accuracy) .

The parallel with the Internet infrastructure is that the latter is a complex black box as well, that much of the focus of its development has been on its performance characteristics (e.g., response times, security, and resilience), and that it also may affect

people's lives in unpredictable ways, albeit more indirectly because it is a communications substrate that applications build upon [Arkko19]. For instance, a power grid operator may be reluctant to remotely control power lines at field stations over the Internet because it does not know the properties of the chain of network operators responsible for enabling the communication and cannot control them. Another example is that the opaqueness of the Internet infrastructure may lead to concentrations of power going unnoticed, resulting in individuals and businesses becoming overly dependent on large commercial players they have little influence over [Arkko20]. Finally, individuals typically do not know if their data passes through network operators they would not trust or that their employer would disallow for certain classified types of work.

The difference to AI systems is that the Internet has only one high-level task, which is to securely and reliably provide end-to-end communications. In addition, a large part of the Internet's complexity stems from its decentralized architecture with distributed ownership and control [Lehr19], whereas in AI the complexity is in the decision making algorithms. Finally, the need for a responsible approach emerged relatively quickly in the field of AI, likely because the effects of its classification algorithms are more visible to users.

Similar to responsible AI, a responsible Internet **extends** the design of the Internet infrastructure with four design goals, which we discuss next.

3.2 Design Goals

Inspired by responsible AI, we propose to update the design of the Internet so that its infrastructure becomes more transparent, accountable, and controllable at the network-level, which is how we define a responsible Internet. In addition to these three design goals, we also formulate a fourth one, which is that the functions that reinforce the Internet's transparency, accountability, and controllability properties need to be usable by a wide range of end-users. Together, we think of our four design goals as extending the Internet's original design goals, such as federating autonomously administered networks and survivability of failures [Clark95].

Transparency. Transparency is the ability of a responsible Internet to describe its internal structure in terms of network operators, their properties (e.g., their jurisdiction and technical infrastructure), and their relations with other network operators. By network operator we refer to an administrative entity that operates a network, such as an access network, a transit network, a data center network, or a Content Distribution Network (CDN).

We distinguish two types of transparency:

- 1) Data transparency describes which network operators transport a particular data flow (e.g., instructions to configure a power grid's field station) and how they process these flows. Data transparency typically requires network operators to track how they process data flows, for instance where a flow entered their network, which types of routers handled the flow, and where the flow left the network. Data transparency for instance enables power grid operators to track how flows of instructions reach field stations (see Section 2.1). A flexible implementation requires advanced network functions such as inband telemetry in open programmable networks (see Section 5).

2) Infrastructure transparency describes the infrastructure properties and relationships of network operators (e.g., their servers, routers, their geolocation, and the open source software they use), independent of specific data flows. Infrastructure transparency is based on self-declarations by network operators about their properties and relationships (e.g., the third parties they use) and on independent observers that map networks using large-scale measurement systems (see Section 4). Infrastructure transparency for instance enables policy makers to study the concentrations of power in an ecosystem such as the DNS (cf. Section 2.2).

Our notion of transparency is similar to that of responsible AI, which is about “the need to describe, inspect and reproduce the mechanisms through which AI systems make decisions and learn to adapt to their environment, and to the governance of the data used or created” [Dignum17]. Responsible AI does however not distinguish the concept of a flow, which is specific to computer networking.

Accountability. Similar to transparency, we consider two types of accountability:

1) Data accountability is about network operators explaining that they process specific data flows in a certain way, such as that they made certain routing decisions or that an intermediate network operator (e.g., a CDN) terminates TLS connections rather than the intended endpoint.

2) Infrastructure accountability is about network operators explaining that they designed their infrastructures in a certain way. These details can pertain to why they outsource parts of their operation (e.g., to flexibly provide DNS services in different parts of the world [SIDN18]) or why they use particular open source software.

Accountability requires actors to explicitly describe the norms (or ground rules) they use for decision-making. For example, network operators could indicate that they prefer to route their traffic through certain groups of operators, such as those that implement the MANRS rules for secure routing (e.g., to actively prevent the propagation of incorrect routing information) [MANRS]. Similarly, a global cloud provider could inform its users that its default policy is to connect users to a local data center for performance reasons. As a result, Europeans traveling to the US would know that they will be using a US-based data center, which they may then ask the cloud provider to change (see controllability).

The norms are our equivalent of the “representation of the moral values and societal norms holding in the context of operation, which the agent uses for deliberation” in responsible AI [Dignum17].

A responsible Internet captures both transparency and accountability details in so-called network descriptions (see Section 4).

Controllability. Controllability is about the ability of users (e.g., critical infrastructure operators) to specify how they expect chains of network operators to handle their data based on descriptions of the Internet’s internal operation (see transparency and accountability). For example, a smart grid provider could use the Intent Definition Language [Riftadi19a] to indicate that it only wants to send instructions to remote field stations via certain classes of network operators, such as those that have certain security properties, are in particular jurisdictions, or that use verified open source stacks or certain types of routers.

Controllability requires new network functions such as path control based on multiple parameters, which we propose to implement using open source-based programmable networks (see Section 5).

A different form of controllability is through policy making, which operates at longer timescales and requires policy instruments (e.g., about allowed levels of outsourcing) rather than new network functions.

Usability is the ability of a responsible Internet to realize the other three design goals in an easy to use way for a wide range of users. This is important because users such as smart grid providers and policy makers will typically not be network experts and because the Internet infrastructure and the technologies it uses are complex, even for network experts. A responsible Internet therefore needs to provide transparency and accountability details about its internal workings at a high level of abstraction and in a machine-readable way so they can be interpreted by automated tools (e.g., to analyze network descriptions). Responsible AI offers some contrast in that usability is not an explicit design goal.

3.3 High-level Architecture

Figure 1 shows our high-level architecture of a responsible Internet, using a power grid provider as an example user. Our architecture realizes the four design goals of Section 3.2 through two new distributed and decentralized systems (NIP and NCP), and a set of policies.

Network Inspection Plane (NIP). The NIP improves the transparency and accountability properties of the Internet in a usable way. It allows users such as smart grid operators to query a responsible Internet for details about its internal operations in terms of network operators (interaction ①), which may include ISPs, DNS operators, and cloud providers. These **network descriptions** cover network operator properties such as their jurisdictions, technical infrastructure (e.g., routers, switches, servers, and their security posture), and relations with other network operators (e.g., outsourcing relations). Network descriptions can be based on a language such as the Network Description Language (NDL) [NDL07] [NDL13] or the recently proposed GAIA-X self-descriptions [GAIA-X20b].

The added value of a network description is that it abstracts away from the underlying technical mechanisms to obtain the details about network operators (e.g., through large-scale measurement systems), which makes it useful for a wide range of non-expert users. For example, the power grid provider in Figure 1 can use the NIP’s network descriptions to assess how instructions for remote power switches flow to field stations, while policy makers can use it to spot concentrations of power.

The network descriptions that the NIP returns can pertain to a specific data flow such as flow F in Figure 1 (data transparency and accountability) or to the infrastructure irrespective of a particular flow (interaction ③). The latter type of information is relevant, for example, for policy makers (see Section 2.2).

The NIP populates network descriptions using various sources, including heterogeneous large-scale measurements from independent observers (e.g., using a system such as OpenINTEL for the DNS [Rijswijk16]) and open programmable telemetry functions in the infrastructure of network operators [Hill18] (see Section 5). It also uses self-declarations from network operators (see Section 4), similar to the “self-descriptions” of GAIA-X operators [GAIA-X20b] or the “cybersecurity labels” that large companies such as Deutsche Telekom, Ericsson, and Thales recently advocated for service providers and manufacturers [CS16].

The NIP provides mechanisms that enable users to verify the data source (e.g., similar to DNSSEC) that provides details about a particular network operator. This is important because it helps users trust the network descriptions that the NIP provides.

An equivalent of the NIP does not exist in the current Internet, because measurement systems are mostly not standardized and typically require scarce technical expertise of people such as network operator staff and security researchers.

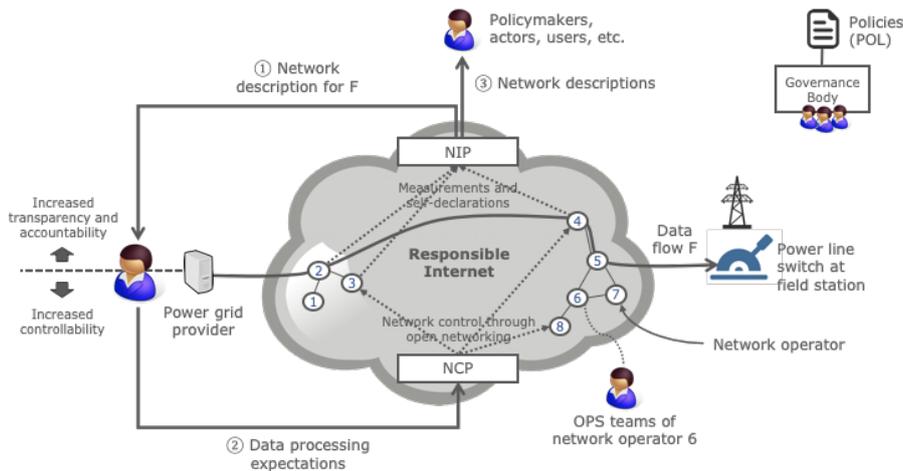


Fig. 1. A responsible Internet's key components (NIP, NCP, and policies) using a power grid provider as an example.

Network Control Plane (NCP). The NCP increases the controllability property of the Internet in a usable way. It is the counterpart of the NIP and enables users to specify how they expect chains of network operators to handle their data based on network descriptions (interaction ②). For example, the operator of the smart grid in Figure 1 can use the NCP to indicate that instructions for power switches at remote field stations [Chromik19] should only pass through network operators in a certain jurisdiction or through network operators that have open sourced their data and control plane software. Similarly, customers of video services such as Zoom could use the NCP to select a video server on a data center in their own jurisdiction rather than a differently situated, default video server (see Section 2.4).

The NCP consists of a set of control and data plane services for open programmable network equipment that map users' expectations to programmable network functions. It also contributes to the transparency property of a responsible Internet through open programmable telemetry functions (see Section 5).

The level of control that we envision for the NCP is much richer than in the current Internet, where control across networks is relatively one-dimensional.

Policy framework (POL). A responsible Internet also requires a set of policies that define the norms that network operators need to adhere to in terms of transparency, accountability, controllability, and usability. This includes auditing to check if requirements continue to be met. For example, the policy framework could define a basic level of responsibility that only requires network operators to publish rudimentary

details such as their legal jurisdictions. Higher levels of responsibility could amount to network operators also sharing details on their relations with other operators, data plane telemetry, the geolocation of their servers, the source code of their data and control planes, and audits of data plane software.

The policy framework needs to be managed by a governance body, for which we envision a lightweight, multi-stakeholder model such as MANRS [MANRS] for routing security. More “heavy weight” models are possible as well, such as a governing body that is part of ICANN, RIPE, or a national government.

For simplicity, we omitted the fourth interaction in Figure 1, which is between policy makers (top) and the network operators. These interactions for instance involve the former incentivizing the latter to change their infrastructure to share details about their operation through the NIP (see Section 2.2), perhaps based on citizen-supplied network descriptions.

3.4 Network Operator Architecture

Figure 2 provides an overview of the architecture of a network operator in a responsible Internet, using Figure 1 as an example. The numbers in Figure 2 (①, ②, and ③) correspond to the interfaces in Figure 1.

The main components of the operator architecture are:

- NIP server: locally stores the description of an operator and shares it with the global NIP. A NIP server also collects measurements from within the operator and acts as a NIP client to obtain descriptions of other operators from the global NIP. This includes enhancements of the operator’s own description with measurements from independent observers.
- Open telemetry functions: control and data plane functions for open programmable networks that collect fine-grained telemetry, such as network paths and routing table versions [Hill18]. The telemetry functions provide input to the operator’s NIP server.
- NCP server: invokes networking functions that enable users to influence how a network operator processes their data flows (cf. Section 2.1) by calling programmable network functions.
- Open networking functions: predefined control and data plane modules that enable network operators to program their white box network equipment (routers, switches, etc.).

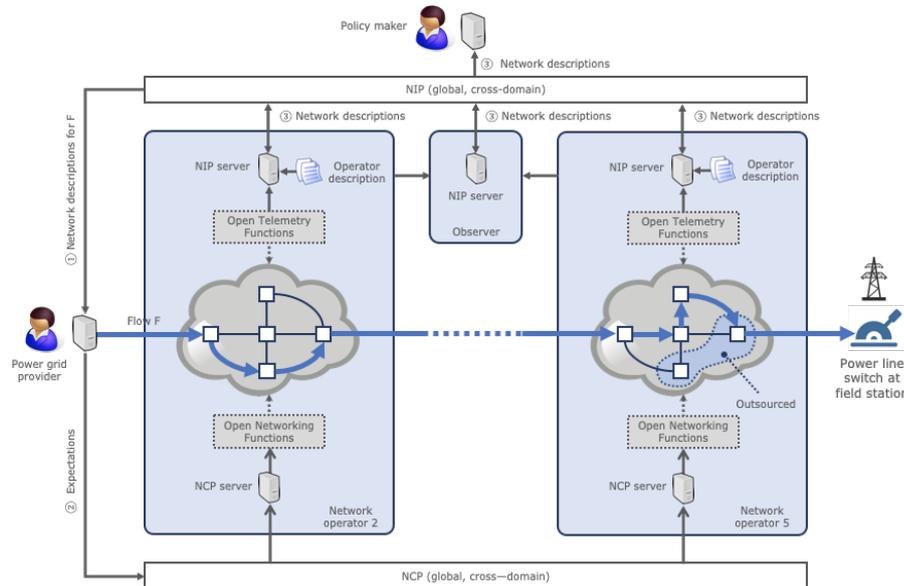


Fig. 2. Architecture of a network operator in a responsible Internet.

Observers implement a NIP server as well and use their “outside-in” measurements to enhance network operator descriptions.

Our key challenge is how to design, implement, operate, and evaluate the components of the architecture, which we will discuss in the next three sections.

4 More Transparency and Accountability through the Network Inspection Plane

The Network Inspection Plane (see Section 3.3) increases the transparency and accountability of the Internet through high-level, measurement-based network descriptions of network operators (e.g., ISPs, DNS operators, and cloud operators), their relations, and their attributes. The NIP creates and populates these descriptions, which brings about many challenges.

We first discuss the concept of a network description in more detail (Section 4.1). Next, we present a first set of research challenges we identified to develop and evaluate the NIP (Section 4.2) and several measurement systems that can act as starting points to address these challenges (Section 4.3).

4.1 Network Descriptions in more Detail

A network description is a machine-readable specification of the properties and relations of a group of interrelated network operators. A network description consists of network operator descriptions, which cover operator attributes such as:

- Services the operator provides (e.g., transit, DNS services, or CDN services)
- Types of relations with other network operators (e.g., delegation)

- Infrastructure (e.g., autonomous systems, router types, geolocations)
- Data and control plane details (e.g., software/hardware attributes)
- Data control capabilities (e.g., path control or geo-based end-point selection)
- Internet security incidents handled (e.g., domain or routing hijacks)
- Available measurements (e.g., DDoS traces, data plane telemetry)
- Norms used (e.g., MANRS or security audits)
- Applicable jurisdictions
- Support for security functions such as RPKI

Figure 3 shows a simple example based on Figure 1, where the network description consists of network operator descriptions NOD1 through NOD8. The dashed lines between operator descriptions represent inter-operator relationships (e.g., between NOD1 and NOD2), while the dashed lines between a description and an operator indicates that the description pertains to that operator (e.g., NOD2 is the description of network operator 2).

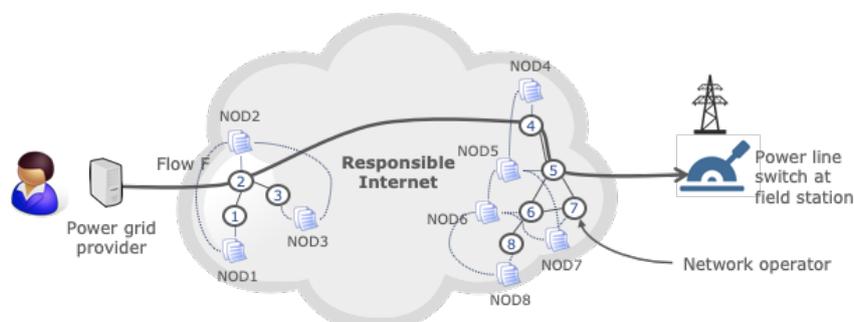


Fig. 3. Example of a network description (based on Figure 1).

The NIP populates a network operator description using two types of sources. The first are independent observers that carry out infrastructure measurements. For example, an observer like the OpenINTEL system [Rijswijk16] regularly obtains the DNS records of a wide range of TLDs, which enables it to map authoritative name servers to the networks where they reside and include these details in the descriptions of the TLD operators. Similarly, an observer such as the RIPE ATLAS measurement network can derive that a DNS operator outsources parts of its operations because its measurements reveal that clients end up at different anycasted DNS servers in different networks depending on the client's geographic region.

The other source of information are network operators themselves. For example, they can add details on what other operators they peer with (cf. PeeringDB [Böttger17]), to which operators they outsource part of their operations, details about the types of equipment they use, their certification levels, and what measurements they carry out on their infrastructure that are available for collaborative incident analysis (see Section 2.3). The latter can for instance be based on the output of telemetry functions in open programmable networks [Hill18] (see Section 5).

A network operator description captures these attributes at an abstract level (e.g., using the Network Description Language [NDL07] or the recently proposed GAIA-X

self-descriptions [GIAIX20b]), thus hiding the details of the various data sources from applications and making them easy to use.

4.2 Challenges for the Research Community

We identified a first set of 7 high-level research challenges for the NIP, labeled NIP-RQ1 through NIP-RQ7.

NIP-RQ1: how do we model network operators? The NIP needs an abstract data model to describe the properties of constellations of network operators like the one in Figure 3. Sub-questions we identify are:

- How do we model a network operator? For example, what attributes do we need to capture (see Section 4.1), what are their semantics, and what is the appropriate breadth and depth of the model? This is important because network operators can have diverse properties. Also, should the scope of an operator description be all-encompassing or do we opt for a modular approach? For example, should we include details about physical Internet infrastructure as part of one monolithic model or as a separate module?
- How do we model relations between network operators? For example, they may delegate activities to other network operators, they may collaborate, but also commit to policies (e.g., MANRS for routing decisions). Accountability requires careful organization of information on dependencies, including unique identifiers for network operators but also standardized yet extensible relationship descriptions. All participating systems must be able to recognize these appropriately and, if necessary, update definitions. Accountability is also closely linked to verifiability, which we discuss in NIP-RQ3.
- How do we keep the data model forward-compatible and generic to also support clean slate Internet architectures such as SCION [SCION17]?

NIP-RQ2: how do we populate network descriptions? Another challenge is how network operators and observers together populate network descriptions. This is a challenge because of the wide range of measurement systems that they use (e.g., passive and active systems), which are currently not standardized in terms of methodologies and output semantics. It is also a challenge because the measurement data will come from multiple vantage points (e.g., home networks or backend systems) and may conflict with each other [Clark03]. Solving this challenge is important because the value of network descriptions is that they abstract away from measurement details so they become useful to a wide range of users.

A related question is how network operators update their operator description. For example, network operators should be able to automatically update their description through their NIP server (Figure 2) when they change their infrastructure (e.g., when they outsource part of their operations to a third party). Also, operations teams should be able to verify their operator description before publishing it in the NIP.

Similar facilities to provide such details exist today, but in a scattered and unstructured way. For example, Internet operators provide assertions of network peerings through PeeringDB [Böttger17]. Similarly, some DNS operators blog about the software types [SIDN16] and third-party DNS operators [SIDN18] they use.

NIP-RQ3: how do we validate the authenticity of network descriptions? The values in a network description (e.g., relations, policies, and available measurements)

will typically originate from network operators and various independent observers (see NIP-RQ2), so we will need some way of validating their authenticity. In some cases, there will be a trust anchor, such as sources adding an RPKI-based signature to the value they provide. In other situations, we will need to revert to measurements from a variety of vantage points because there is no evident trust anchor.

One possible research direction is to augment the NIP with an append-only log that stores measurements similar to Certificate Transparency [Trillian20], which makes it possible to establish a causal chain of measurements documenting an event or configuration. The advantage is that these logs do not require universal verifiability and that attempts to tamper with results of previous measurements can be detected.

NIP-RQ4: how do we design the NIP? A key factor in the design of the NIP is the expected usage patterns of its two main services: (1) enabling users to look up the descriptions of groups of interrelated network operators and (2) allowing network operators and independent observers to update operator descriptions. For example, we could design the NIP as a hierarchical system similar to the DNS if its access pattern consists of relatively few updates of network descriptions and many lookups. If these patterns are more symmetrical, then a peer-to-peer design might be more appropriate.

Getting an indication of these usage patterns before building the system will require longitudinal measurements of how constellations of network operators and their attributes currently change over time, for instance in the DNS and in the routing system. A potential approach for the DNS is to study the changes in the DNS ecosystem, such as in a ccTLD¹.

Another factor is the expressiveness of NIP queries, which should allow users such as the grid operator to search for network operators with certain properties (e.g., those with certain packet forwarding policies or security posture), amongst others.

Addressing this challenge requires a flexible system design, which is particularly important because the NIP is a cross-layer system: it provides a network-level service, but it populates network descriptions using measurements and declarations from different levels (e.g., network-level peerings as well as properties of equipment). This includes developing open standards that the NIP requires, for instance to facilitate interactions between NIP servers or express network descriptions (e.g., using the Network Description Language [NDL07]).

NIP-RQ5: how can we validate the usefulness of the NIP? Validating the added value of the NIP will require the development and evaluation of tools and algorithms that analyze network descriptions for various real-world use cases, such as those of Section 2. For example, these tools could query the NIP to calculate sector-specific “responsibility scores” of network operators based on raw network descriptions (e.g., for power grid providers or for citizens). Similarly, network operators could develop tools to automatically and collaboratively triangulate measurements of the same security incident (e.g., a DDoS attack) from multiple vantage points.

Validation also requires the development of target group-specific user interfaces (e.g., visualizations) that enables users such as power grid providers to easily and intuitively browse the NIP’s network descriptions and correctly interpret them for their specific use case. This will likely require new user interaction mechanisms, for instance

¹ Three of the authors of this paper are affiliated with the ccTLD operator of the Netherlands, .nl.

to represent infrastructure concepts for users unfamiliar with networking and enable them to explore network descriptions at different levels of granularity.

The evaluation of the NIP will be a multi-disciplinary effort, requiring extensive consultation between domain experts, developers, and designers. This effort needs to be at the core of making a responsible Internet a reality.

NIP-RQ6: How do we incentivize network operators to join the NIP? Network operators will need an incentive to join a responsible Internet, because this will likely require significant investments from them, for instance in terms of technical facilities to join the NIP and support open networking (see Section 5) and training their staff.

For the NIP, one possibility is to create a demand for the network descriptions that the NIP provides. A potential strategy to explore if such a demand exists is that a network operator interest group such as RIPE collaborates with “industry verticals” (e.g., critical service providers or consumers interest groups) to understand what kind of descriptions they would like to obtain from network operators. A small group of network operators could then set up a basic version of the NIP to pilot how this would work in practice, both for the network operators as well as for the users of the network descriptions. They could include their lessons learned in a set of implementation guidelines for other network operators to use, similar to the MANRS [MANRS].

As part of such a pilot, other types of users might develop new services based on network descriptions, such as a reputation system that calculates the “responsibility score” of a network operator. This would enable critical service providers and other types of users to easily discover network operators with “good” responsibility scores and prefer them to carry their data flows using services that the NCP provides (see Section 5). In this way, an initial small set of network operators may stimulate adoption towards a larger group because the transparency that the smaller group offers makes their services more appealing to users such as critical services providers (competitive advantage). A reputation system like this is similar to internet.nl, a site that enables users to check the security features of their ISP’s connection, amongst others.

NIP-RQ6 is related to the adoption of open networking and policy making, which we will discuss in sections 5 and 6, respectively.

NIP-RQ7: how to balance transparency and security? This last research question is perhaps the most important one for the NIP. Our notion of transparency and accountability implies that network operators share a certain amount of detail about their operations through network descriptions, but this may offer attackers quicker and more effective reconnaissance methods for possible targets. For example, sharing details on software versions might make a network operator more vulnerable to exploits. While a fast-paced patching cycle would solve this problem, this is not necessarily an option for everyone. For example, many organizations need to first test patches thoroughly for correct functionality before they deploy them in their production environment (e.g., in the financial industry).

We will thus need to look into the right level of detail to be included in network descriptions, which also ties into NIP-RQ1. A further direction to explore are ways to encode information in such a way that a malicious actor hardly profits, but a querying NIP user still receives useful results. There is precedence in the DNS: the DNSSEC NSEC3 record confirms the non-existence of a domain name while making it very hard for an attacker to enumerate those domains that do exist.

Solutions that address this research question will also require some form of standardization to achieve buy-in from network operators, for instance in the form of a cross-operator framework.

4.3 NIP Starting Points

We discuss a few recent research results in the field of Internet measurements that can act as starting points for the challenges that we identified in Section 4.2. The measurement community developed these systems over the years because data availability and diversity is crucial to further our understanding of the Internet ecosystem (e.g., for DDoS characterization [Jonker19]). In this section, we discuss a few of these systems and how they could contribute to the NIP. We are involved in some of them.

OpenINTEL. OpenINTEL [Rijswijk16] has the long-term goal of collecting a daily comprehensive dataset on a significant part of the global domain namespace. It currently covers around 65% of the global namespace and collects over 3.7 billion data points every day.

The data collected by OpenINTEL can form the basis for an independent observer that covers large parts of the DNS and adds its measurements to network operator descriptions. It can also be used to perform a retroactive study of dependency developments in the DNS (cf. [Abhishta19]), which provides details on inter-operator relations. Its reverse DNS dataset can augment transparency data on the IP layer.

Certificate Transparency (CT). CT introduces logs to the certification process. These are neutral parties that can be run by anyone, although a spread across many different jurisdictions is desirable. Logs issue cryptographically verifiable receipts for every certificate they receive. Browsers can verify that the certificate they receive in a TLS connection has been correctly logged. CT’s notion of “transparency” has since been generalized in Google’s ongoing project Trillian [Trillian20].

The concept of transparency logs can be used for network descriptions to log measurements about operators or their relations from different sources. However, they may need to be scaled up because CT is designed for an ecosystem of just a few hundred actors (Certification Authorities). A small number of well-known logs is enough to enable this. A responsible Internet has operators orders of magnitudes larger (there are currently around 70.000 autonomous systems (AS-es) [CDIR20]), which usually are not aware of each other. Logs can in principle be run by any such operator, but an additional mechanism will be needed for the existence of logs to be communicated. Measurements are needed to validate information in the logs.

BGP Hijacking Event Analysis (HEAP). HEAP [Schlamp16] attempts to detect the cause of anomalies in the Internet’s routing system, such as legitimate inter-AS traffic engineering or attacks on an Autonomous System (AS). HEAP accomplishes this by combining a feed of “hijacking reports” with publicly available routing information and measurement data from Internet-wide scans. Routing hijacks are incidents in which an AS announces a route to an IP range that is a sub-prefix of a BGP announcement by a different AS. Such prefixes are generally globally accepted by all ASes and result in all traffic taking the new route instead of the old one. Ultimately, this is possible because BGP does not offer any security itself (RPKI has some deployment but is not widely used to filter routes).

HEAP uses descriptions taken from RIRs to reason about legitimate business relationships between ASes. For example, RIRs such as RIPE store relationships of the form “maintained by” between ASes. This indicates that an incident is most likely benign as one AS has informally described an outsourcing of responsibility to another AS. The transparency in the responsible Internet we envisage would be a superset of such descriptions.

MADDVIPR. MADDVIPR² (Mapping DNS DDoS Vulnerabilities to Improve Protection and Prevention) aims at comprehensively analyzing the DDoS landscape targeting the DNS (e.g., in terms of characteristics of DDoS traffic). The project stems from the observation that DDoS attacks on the DNS can have devastating effects [Abhishta19] [Mirai17], as it effectively leads directly to loss of connectivity for end users and services.

MADDVIPR can contribute to the creation of network descriptions because it is able to shed light on the DDoS weak points of the DNS landscape. For example, it is able to map single points of failure in the global DNS and vulnerabilities [Sommese20] in DNS deployment that DDoS attacks can exploit. MADDVIPR also aims at mapping DNS DDoS “hotspots”, in terms of attackers, attacks and targets, which is relevant for network descriptions as well.

5 More Internet Controllability through the Network Configuration Plane

The NCP (see Figure 1) consists of a set of control and data plane services for open programmable network equipment, which network operators use for two purposes: to enable users such as grid operators to express a limited number of high-level data processing preferences (controllability) and to provide new data plane telemetry functions (transparency and accountability).

We envision that open networking (networking based on open source software and open programmable networks) will play an important role to realize a responsible Internet, but we identify a number of open problems.

We discuss our notion of open networking (Section 5.1), our research challenges for the NCP (Section 5.2), and starting points to address them (Section 5.3).

5.1 Open Networking and a Responsible Internet

We define open networking as network equipment that uses open source software (e.g., based on OpenSourceNetworking³) and open hardware modules (e.g., based on the Open Compute Project⁴). Open networking is important for a responsible Internet because it enables users such as the grid operator of Figure 1 to verify the security of these modules, which enables higher levels of trust and sovereignty. Network operators in a responsible Internet share details about the software and hardware they use through network descriptions (see Section 4) and their local NIP server (Figure 2).

² <https://maddvipr.org/>

³ <https://www.linuxfoundation.org/projects/networking/>

⁴ <https://www.opencompute.org/>

Open networking requires network equipment that can be programmed. Until a few years ago, networking hardware (routers/switches) were the proverbial black boxes that came with vendor-specific software that could be configured to some degree but could not be changed (re-programmed). Moreover, the hardware integrated both the equipment’s control plane (protocols and algorithms needed to make routing decisions) and its data plane (forwarding of packets). This hampered innovation, as adopting any new protocol basically required purchasing a new device.

While programmable networks have been studied since the 1990s [Astuto14], Software-defined Networking (SDN) [Kreutz14] introduced a new type of networking hardware that separated the control and data plane functions, allowing the control logic to be programmed (by the user) in software and the corresponding rules to be installed in the data plane. In addition, new types of programmable network hardware allow engineers to flexibly develop custom hardware-based packet processors for the data plane, for instance to extend IPv6 packets with custom headers [2STiC20] [Hill18] or implement a “clean slate” protocol such as SCION [SCION17]. A popular language that supports this kind of programmability is P4 [Boss14].

From a functional perspective, control and data plane programmability is important for a responsible Internet because it enables network operators to develop and standardize new network functions that allow users such as grid operators to indicate their data processing preferences for chains of operators (see Section 2.1), thus increasing the controllability of the Internet infrastructure. Operators enable users to express these preferences in a language such as the Intent Definition Language and implement them on programmable hardware (e.g., using P4) [Riftadi19a].

Programmable networks are also important for the NIP (Section 4) because they enable operators to add fine-grained telemetry from the data plane to network operator descriptions. Finally, programmable networks enable operators to manage their networks in new ways, such as through custom traffic management functions to handle DDoS attacks.

5.2 Challenges for the Research Community

We identify 8 open networking-related research challenges to realize the NCP, which complement the NIP challenges that we discussed in Section 4.2.

Our research questions cover aspects related to exposure of telemetry data (NCP-RQ1 and NCP-RQ2), the security implications of such transparency (NCP-RQ3 and NCP-RQ4), the effects on users of open networking techniques (NCP-RQ5 and NCP-RQ6), and future extensions (NCP-RQ7 and NCP-RQ8).

NCP-RQ1: What open telemetry measurements are useful for network descriptions? This is important because open networking and in particular programmable data planes allow for an unprecedented level of telemetry [Knossen19] [Hill18]. Examples of measurements include the path that a flow takes through an operator’s infrastructure, the version of the routing table that each router uses to make routing decisions, the source and type of open source software used, and the operations that a router applies to the packets in a flow (e.g., forward, decrypt, sinkhole). The NIP can for instance use these details to enable data transparency and accountability (see Section 3.2).

A related research question is the appropriate granularity and frequency of the measurements, which is an important consideration for routers and a network operator's NIP server (see Figure 2). For example, the sampling frequency of the telemetry system needs to be such that it still allows for line speed forwarding of large numbers of flows.

Another key question is how to enable users to verify the authenticity of data plane measurements, which is related to NIP-RQ3 (Section 4.2).

NCP-RQ2: How do we get data plane measurements into network descriptions?

This is important for data transparency, so users get insight into which operators processes their data flows. One possible solution is that routers include measurements in the packets themselves (e.g., using IPv6 extension headers) and forward it to the next hop ("packet forwarding state" [SCION17]). For example, the border router of operator 2 (Figure 2) could add its measurements for flow F to outgoing packets in extension headers so that operator 5 can upload them into the NIP through its NIP server for the whole operator chain. While this could be a feasible approach (SCION's path transparency functions work similarly [SCION17]), its downside is that it increases message size which is a disadvantage on wireless networks, which are typically bandwidth-constrained.

An alternative is that each network operator retrieves the data plane telemetry from their routers and adds it to network descriptions through their local NIP server. The potential downside of this approach is that it requires extra state in the network operator's control plane, which increases its complexity.

NCP-RQ3: How to protect the integrity of open source data and control plane software? Similar to the NIP (see NIP-RQ7), a major challenge for the NCP is how to balance the openness of data plane and control plane software and their security in terms of vulnerabilities that can be exploited.

One potential direction is to enhance network control programs with run-time attestation of these programs' binaries [Abera16], which enables network operators to verify the integrity of execution paths in the code and that they have not been changed by attackers. Similarly, network operators can also use static attestation to check the integrity of binaries by computing a hash over it at boot time and making them available for lookup [Abera16], for instance as part of a network description.

Another way to protect the integrity of open data and control plane programs is through auditing. For example, the set of policies in the overall architecture of Figure 1 could not only define responsibility levels but could also set requirements for open data and control plane software that network operators need to adhere to. Ultimately, such auditing mechanisms could become part of operational security best practices such as ISO270001.

NCP-RQ4: How do we promote adoption of open networking systems? Similar to network operators joining the NIP (see NIP-RQ6), another challenge is to develop mechanisms and incentives to stimulate the adoption of open networking so the concept of a responsible Internet can be rolled out gradually. Without open networking, a responsible Internet would require a complete overhaul of all Internet equipment and software, which would be virtually impossible.

A major challenge is to enable network operators to understand what joining a responsible Internet means for their business model. For example, they might need to redimension their infrastructure because their "responsibility score" results in users sending additional traffic through their network. This may be an advantage if the users

are paying customers, but it might be a disadvantage if they are some other operator's customer. In this case, a responsible Internet will likely also have an impact on the business relations between network operators.

As part of developing a business case for open networking, network operators will also need to understand what investments they will need to make to change their operations (e.g., in terms of new equipment, educating staff members, and operational costs). Early adopters of the concept could include such lessons learned in a set of implementation guidelines (cf. NIP-RQ4), which the governing body (see Section 3.2) could further promote.

NCP-RQ4 is related to policy making, which we will discuss in Section 6.

NCP-RQ5: How can open networking take advantage of the additional insights that network descriptions offer? Open networking allows for a large degree of flexibility, which can be driven by the details that the NIP provides. The challenge for network operators is how to map network descriptions to the network control programs of the NCP.

For example, open programmable devices allow for adaptive rerouting of data flows among various public and private entities for crowd management applications. These changes can be driven by devices dynamically analyzing network descriptions to find weaknesses in the network that are a problem for this specific type of application and that they can therefore solve more effectively.

Similarly, network operators can proactively change their network (e.g., using VNF for fine-grained adaptations) because network descriptions of operators [CSCloud19] [Coreflow18] provide them with a more comprehensive view on what is going on in the network (cf. Section 2.3). The decision where to place a network function [Ochoa18] will play a big role in how network operators perform in a responsible Internet. The ability for operators to take autonomous decisions in response to security incidents [Koning19] will require further investigation.

NCP-RQ6: What will be the effect of NCP on users? NCP users will likely also require novel interaction concepts (cf. NIP-RQ5), specifically to express their preferences on how chains of network operators should handle their data flows. For example, users such as grid operators might need an extension of their control room clients to include such network controls.

NCP-RQ7: How can we leverage open networking to evaluate and extend the concept of a responsible Internet for other internet architectures? Programmable networking systems enable researchers to more quickly experiment with non-IP architectures, such as SCION [SCION17], RINA [Maffione16], and NDN [Adrichem15]. An open challenge is how to specialize the concept of a responsible Internet for these architectures. With the uptake of languages such as P4, we expect an increase in new architectures that will coexist with the current IP-based Internet. Recent efforts such as [Cauley19] identify the emergence of network virtualization and network programmability as the components that will allow the development of future Internet infrastructures. We build our work on the same insight.

NCP-RQ8: What are the scalability limits of the information exchange required by the NCP? A key challenge is how to scale a responsible Internet to large numbers of users. This will for instance require a careful design of the mechanisms that maintain the additional state that a responsible Internet requires. For example, open telemetry will enable operators to summarize how they processed a user's data flow but sharing

these details as packet forwarding state across operators (see NCP-RQ2) might result in too much overhead at the IP-level.

Similarly, users might need to express their data processing preferences at different levels of granularity in order for the system to scale. For example, users wanting to reroute their flows through different network operators might need to choose from several predefined paths like in SCION path control [SCION17] because full per-flow routing will not scale.

5.3 NCP Starting Points

We discuss a few recent research results in the field of open networking that can act as starting points for the NCP, which we are involved in.

The Netherlands’ national P4 testbed. A consortium of 3 Dutch universities and 5 companies (e.g., the Dutch national research and education network and two Internet exchange points) called “2STiC”⁵ recently set up the first national multi-domain P4-programmable network in the Netherlands (see Figure 4). The testbed uses switches and NICs that can be programmed through P4. It consists of six different sites interconnected by a star-shaped optical network, which can be configured to use different topologies.

The programmability of the 2STiC network makes it very suitable to experiment with novel network functions, such as the ones that the NCP needs.



Fig. 4. Netherlands’ national P4 testbed (March 2020).

SCION-in-P4. “SCION-in-P4” is a P4 implementation of the SCION data plane protocol. SCION (Scalability, Control, and Isolation on Next-Generation Networks) [SCION17] is a clean slate internet architecture that, for instance, aims to enable users to control the inter-domain routes their traffic takes (i.e., which autonomous systems their traffic passes through). We are currently testing SCION-in-P4 on (parts of) the 2STiC testbed of Figure 4.

⁵ www.2stic.nl

This work is relevant to extend the concept of a responsible Internet to other types of network infrastructures.

Data plane telemetry. We developed and experimented with several P4 telemetry mechanisms. For example, Sequential Zeroing [Turkovic20] is a heavy-hitter (i.e., big flow) detection solution for P4-programmable hardware. It operates at line rate, which leads to new types of optimization problems because P4 programs need to adhere to the stringent memory access rates of programmable hardware.

Also, in [Turkovic18] we enable programmable switches to (1) track processing and queuing delays of latency-critical flows and (2) react immediately in the data plane to congestion by rerouting the affected flows. Another example is our in-band telemetry implementation for the 2STiC testbed (Figure 4), which appends node identifiers to IPv6 extension headers [Knossen19] so that destination nodes can extract the full path a packet took from these headers.

These data plane telemetry mechanisms are a starting point for the NCP’s open telemetry functions, for instance to add flow-specific details to network descriptions.

Network slicing. We also used P4-programmable switches to dynamically create, discard and switch network slices (i.e., reserved resources, with known latency characteristics, dedicated to a specific type of application) [Polachan20]. This approach demonstrates how Quality-of-Service (QoS) can be attained for dynamic applications that require stringent latency constraints, such as remote surgery, which is relevant for the NCP.

While this network slicing approach takes its decisions based on real-time measurements from the data plane, it could also be extended to incorporate more information from the NIP.

P4 code generation. We also experimented with the automatic generation of P4 code (based on intents) [Riftadi19a] [Riftadi19b], thus providing first steps towards networks that can adapt themselves with only a few high-level commands from the users or operators (self-programming networks).

This work is relevant for a responsible Internet to manage the quality of the P4 code that the NCP uses and to develop a (standardized) repository of P4 network control software that network operators can choose from.

6 Policy Mechanisms

A responsible Internet not only introduces technical challenges (see sections 4 and 5), but policy challenges as well, such as how a responsible Internet enables better policies and how to incentivize network operators to join the NIP and adopt the NCP.

We first provide a short background on how policies are typically being developed (Section 6.1) and then discuss the research challenges we identify (Section 6.2). We conclude with an overview of policy forums that are potential candidates to address these challenges (Section 6.3).

6.1 Policy Development Background

Policies are made for users with very diverse technical knowledge and skill sets: from savvy, advanced users to late adopters. The policy community needs to understand the

information available about a responsible Internet at each stage: policy making, policy mediation, and policy enforcement (see Section 2.2).

A common deficiency in governing the Internet is that policy makers, especially regulators, have difficulties following the pace at which technical developments occur. For example, the recent (and possibly short-lived) explosion of cryptocurrencies as “regular” forms of payment, or smart contracts as semi-autonomous, self-executing code, led to enormous uncertainties in terms of how to regulate cryptocurrency exchanges or code in smart contracts [Xue19].

Other research on multi-stakeholder governance approaches also highlighted the increasing importance of the private sector in Internet policy, visible for example in standardization and data protection regulation (e.g., private companies such as Google and Facebook affecting legislation drafts at an early stage) and proposals for public-private actions to fight botnets [Silva19].

6.2 Challenges for the Policy Community

We identified four policy-related research questions, labeled POL-RQ1 through POL-RQ4:

POL-RQ1: How to incentivize (large) network operators to join the NIP and adopt the NCP? A responsible Internet will need to grow incrementally, like the Internet itself did. However, it might not result from market pressure alone, which is unlike communications-driven properties such as lower latencies and increased bandwidths that improve applications such as video conferencing [IAB20].

As a result, an important challenge to deploy a responsible Internet is to develop incentives that stimulate a few “first movers” to join the NIP and adopt the NCP. This is a challenge because it requires network operators to invest in changing their infrastructure, for instance to switch to open programmable networks and train their staff (see NIP-RQ6 and NCP-RQ4). Policies that provide these incentives might range from voluntary similar to MANRS [MANRS] to mandated by national regulators.

Another challenge is how to give network operators equally fair possibilities to participate in a responsible Internet. This is important to stimulate competition, which is a driving force in innovation and the inclusion of more diverse network operators should help a future responsible Internet thrive as well. A related challenge is how to get support from existing standardization bodies to encourage more operators to develop and adopt the NIP and the NCP. A potential hurdle to take is that large corporations often play important roles in numerous standardization bodies [Froomkin03] (e.g., W3C, IEEE, and the IETF).

Another dimension is how to incentivize large organizations (e.g., large ISPs or content providers) to join a responsible Internet. In such cases, there is a risk of “regulatory capture” [Carpenter14], which means that a few large dominant Internet actors use their economic power to shape potential legislation aimed at stimulating a responsible Internet in favor of their own interests [Wu18]. This type of risk is real and has been described for diverse scenarios of today, in particular cloud services, modern AI, and data-driven businesses.

POL-RQ2: How do we ensure that a responsible Internet represents the interests of the public, particularly in critical infrastructures? International governmental organizations often advocate values that they would like to see reflected

in the development of a future Internet. For example, the United Nations highlights nine values: inclusiveness, respect, human-centeredness, human flourishing, transparency, collaboration, accessibility, sustainability, and harmony [UN19]. The EU envisions the next-generation Internet as more human-centric, supporting openness, decentralization, inclusiveness, and the protection of privacy, while giving control back to the end-users, in particular with respect to their data [EC20]. The EC also articulated these kinds of values for AI [EC19].

It is significant to continue to research how these values are reflected in critical infrastructures (e.g., energy grids or transportation systems) that use a responsible Internet in countries that favor diverging values or that prioritize them differently. For example, how should critical infrastructure operators across countries incorporate network descriptions (see Section 4) in their services to reflect the above-mentioned common values?

POL-RQ3: How can the international policy community collaborate towards a global responsible Internet at a time of fragmentation? Internet fragmentation, sometimes referred to as “Balkanization”, refers to nation states applying territorial control to their networks. This development has been long debated [Gold06] [Mueller17] and it is well-known that several countries contribute to this process by deploying topic and domain-based filtering at large scale.

Internet fragmentation along territorial borders forms a major challenge for a global responsible Internet, for instance to incentivize national policy makers with varying expectations regarding fragmentation to collaborate. At the EU level, it will be important to examine how member states will utilize transparency features such as network descriptions in an effort to harmonize regulation and strengthen the Digital Single Market (DSM).

POL-RQ4: How do we adequately translate policy to different target groups (e.g., those in Section 2) to ensure the values and function of the future internet? To implement new policies on a largely privately owned and operated Internet, policy makers also need to help service providers translate the values of POL-RQ2 to responsibility profiles for network operators. For example, service providers and other users need to be able to make sense of the new kinds of details that a responsible Internet provides and make the right decisions in their own context. This particularly applies to individuals that are marginalized in the social-economic spectrum.

We expect this kind of research to grow into several smaller research areas that rely on empirical analysis and investigate actual implementation and impact. This breakdown is important because of the complexity of the work. For example, the EU now has complex policies related to their digital agenda [EP20] and it is unclear whether these policies are sufficient to address the issues raised in the context of a responsible Internet (e.g., in sectors such as energy, finance, and medical care).

It will also be important to evaluate how responsible Internet technologies are actually used when different social groups participate (e.g., policy makers and regulators, service providers of critical infrastructures, and individuals). A strategic approach is also required to initiate public and private partnerships and cross-disciplinary research to understand how a responsible Internet is used in different social-legal-cultural contexts.

6.3 Policy Starting Points

Our research questions illustrate that a concerted effort will be necessary to make a responsible Internet a reality. No single country or organization is able to determine future standards alone, at least if they are to be used by a majority of service providers, device manufacturers, and operators. Fortunately, there are a number of forums and consortia that may serve as good starting places and that already have similar items on their agendas.

Internet Governance Forum (IGF). The IGF is a venue to increase awareness of Internet governance, foster conversation, and educate the market.

The themes of a responsible Internet are represented across three of the IGF's current four core policy agendas: an #OnlinePeaceFramework, a #DigitalMarshallPlan, and #OnlineRights4all. The fourth policy agenda, #ResponsibleAIStewardship, is the pendant of the responsible Internet in AI [IGF19].

High-Level Panel on Digital Cooperation. The United Nations High-Level Panel on Digital Cooperation also highlighted better ways of digital governance in its report "The Age of digital interdependence" [UN19]. The report proposed three mechanisms to support an inclusive approach for global collaboration on Internet governance.

- Internet Governance Forum Plus (IGF+) aims to bring more representatives together and promote more actionable outcomes from discussions.
- The Distributed Co-Governance Architecture (COGOV) is dedicated to building a network to design and promote digital norms which policy makers can use as a blueprint to develop their policy, regulation, and laws.
- The Digital Commons Architecture (DCA) works on key issues around social harms to promote established digital commons.

These policy mechanisms can be utilized to foster conversations with a wide range of users and to initiate educational programs to increase awareness of a responsible Internet. These dialogues lead to strong international discourse and help increase users' awareness of their rights in the digital space.

Council of Europe's strategy of Internet Governance. At a regional level, the Council of Europe's (CoE) strategy of Internet Governance concentrated activities in three areas: building democracy online, ensuring online safety and security for all, and respecting and protecting the human rights of everyone in the digital world [CoE19]. Along this strategy, the CoE initiated partnership agreements with eight leading Internet companies including Apple, Deutsche Telekom, Facebook, Google, Microsoft, Kaspersky Lab, Orange, Telefonica, and Cloudflare, as well as six international associations. The goal was to tackle issues including bioethics, data protection, disinformation, and cybercrime.

The starting approach could be, for example, to join the partnership and collaborate with the CoE to address high-priority issues concerning safeguards for internet critical infrastructure.

7 Internet Trust Transition

We believe a responsible Internet enables a global Internet that is trusted by billions of non-expert users with widely varying norms and expectations [Brown10]. This is a move away from the Internet's original 1960s-1970s trust model, which revolved

around a then-local Internet and a relatively small and homogenous group of expert users trusting each other (personal trust [Gomm15]) [Clark95].

We call this change the Internet trust transition. Figure 5 shows that we think of it as a movement across the layers of a Maslow-like “pyramid of Internet needs”. A responsible Internet is at the top of the pyramid and builds on the advances achieved in roughly the previous two decades (the second stage of Figure 5), in which the research and operational communities significantly increased (and are still increasing) the security, stability, and privacy-friendliness of the Internet. Examples of technologies that they developed, standardized, and deployed for this purpose include certificate automation (e.g., through Let’s Encrypt), certificate transparency [Trillian20], DNS security and privacy (e.g., through DNS security extensions and query name minimization), and routing security (e.g., using the Resource Public Key Infrastructure). The lower layer of the pyramid represents the period from the inception of the Internet in the late 1960s/early 1970s to the end of the 20th century, which focused on sufficiently maturing the Internet as a technology and getting it ubiquitously deployed [Ammar18].

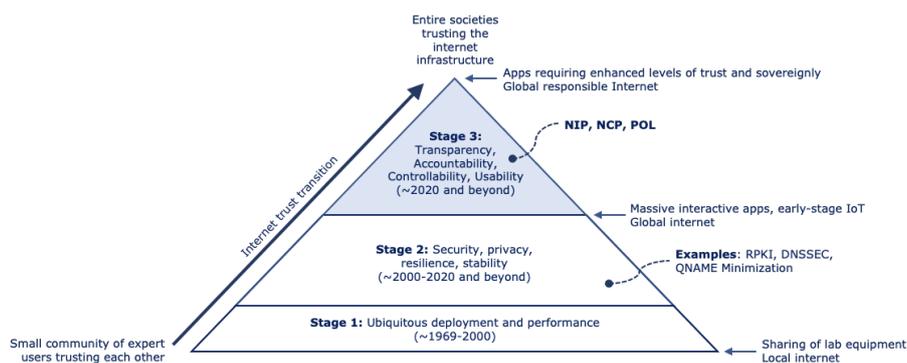


Fig. 5. Internet trust transition (left arrow going up).

A secondary transition that we believe a responsible internet needs to facilitate is what we call the **value transition**: from a relatively homogeneous set of norms and expectations of the community that governed the design and operation of the Internet in the early days (stage 1 in Figure 5), to a broader and much more heterogeneous set enhanced with the norms, laws, and expectations of the different societies in which the Internet is embedded today (stages 2 and 3) [Brown10]. The need for such a transition for the Internet is exemplified by recent work in Europe, where they are already actively seeking to align technological developments with European norms and expectations, for instance in cybersecurity [ENISA19] and AI [EC19].

We believe a responsible Internet enables this transition because it adds transparency and accountability as first class citizens to the Internet, which are important values in many societies as well as in Internet governance bodies (e.g., at ICANN and the IETF). It does mean, however, that the Internet increasingly embodies human values and that it becomes even less “value free”, which is a well-known tussle [Brown10].

8 Related Work

To the best of our knowledge, we are the first to propose and define the concept of a responsible Internet and to provide research directions for it. Our work is also unique because we join three existing but largely disjoint research areas: large-scale measurements (through the NIP, see Figure 1), open networking (through the NCP), and policy making (using the NIP). The related work we did find addresses isolated aspects of our proposal.

Internet evolution. The need to evolve the Internet architecture has been forefront in the networking research community for a long time: from seminal work such as that of Chowdhury et al. [Chowdhury10] where network virtualization was proposed as solution to the Internet ossification, to a very recent proposal such as Trotsky by McCauley [McCauley19] that puts the focus on the use of programmability to allow multiple Internet architectures to coexist. Our work, however, goes one step further because we also include the policy perspective, which is unlike these efforts that primarily have a technological focus.

GAIA-X. GAIA-X [GAIAX20b] is an ongoing project to create a cloud infrastructure and data ecosystem to improve Europe’s data sovereignty.

Similar to our responsible Internet, GAIA-X also advocates “responsible” design goals such as transparency and accountability. Another similarity is GAIA-X’s concept of self-descriptions, which is similar to our network operator descriptions. Nodes (an abstract term that can refer to elements such as data centers, network services, and virtual machines) can self-describe their characteristics, which are meant as inputs for users (consumers and providers) to select the level of data security they need. Self-descriptions can be certified by trusted parties and may refer to self-descriptions of other GAIA-X actors, effectively creating a self-description graph.

The key differences with our work are that we focus on the network-level. For example, GAIA-X currently allows for the self-description of network operators that cloud operators directly connect to (e.g., PoPs and transit providers), but unlike our work they do not consider transparency of the end-to-end communications path, nor do they outline the corresponding measurement systems. Another difference is that they do not consider an equivalent of our NCP. We thus consider our work complementary to GAIA-X.

Open Internet Order. Lehr et al. [Lehr15] discuss the FCC’s Open Internet Order (OIO in short, superseded by the FCC’s Restoring Internet Freedom Order⁶ on June 11, 2018), which aimed to promote open broadband Internet access. Lehr et al. argue that Internet information disclosure and transparency (D&T) are important for different actors, such as ISPs, regulators, and customers. Their D&T policies involve information disclosure along 5 dimensions, such as why disclosure is needed and what data needs to be disclosed. From a network perspective, D&T may pertain to disclosing operational practices such as congestion management and application-specific traffic engineering.

Lehr et al.’s suggestions include creating a D&T Coordinator (a kind of meta-tool) and advocating the use of an independent third-party measurement infrastructure. They

⁶ <https://www.fcc.gov/restoring-internet-freedom>

also refer to examples of voluntary transparency reports by Google⁷ and Automattic⁸. They also attempted to set up an independent measurement platform that can provide new disclosure capabilities, but unfortunately the project seems to have been discontinued and no record of its results were found by us. Also the D&T Coordinator was only presented as a conceptual model and, as far as we could tell, was not implemented.

While our vision bears similarities with the objectives of the D&T Coordinator, the key differences is that we follow a more distributed approach towards D&T that is more fleshed out as well (e.g., because we propose key components and provide a set of starting points).

Transparency vs. anonymity. The need to balance the respect for the privacy of Internet users and the desire to have increased transparency into the operations of the network is a tussle that has been studied the past years. The work of [Pappas15] was one of the first to describe the importance of Internet transparency and possible approaches towards realizing it without sacrificing (too much) anonymity. They were particularly focused on addressing the relation between transparency and net-neutrality, and their conclusion was that the focus should move from the latter (neutrality) to the former (transparency) as this would ultimately benefit users. This has been followed by a number of proposals all centered around privacy-preserving data collection in networks. For example, [Fragkouli19] have recently proposed an algorithm to provide aggregated insight into network flows, even in settings with limited number of flows.

We address this dichotomy in our work too, but we do go beyond the traditional flow-based approach. For example, for NIP-RQ7 (“how to balance transparency and security?”) we will need to look at the effect of transparency on all Internet actors, including users (e.g., grid operators and citizens) and network operators. We also move beyond the current solutions running on traditional hardware because we exploit network programmability, for instance to address NCP-RQ1 (“What open telemetry measurements are useful for network descriptions?”). The additional telemetry we have access to provides us with metrics that enable the right ratio of openness and protection.

Defining the Internet. Lehr et al. [Lehr19] posit that how to define “the Internet” is not easily answered and rather that it should be viewed from the following different perspectives: (1) the architectural building blocks, (2) the enterprises that use that architecture to offer services, and (3) the customer experience. To illustrate the point that a single definition of the Internet should not be pursued, they provide several thought-provoking examples related to (policies for) an open Internet. For example, security problems are not solely caused by weaknesses in the Internet architecture, but often arise from the applications used. As such, securing the Internet is a shared responsibility.

The three perspectives and examples put forth in this paper illustrate that answering research question POL-RQ2 (“How do we ensure that a responsible Internet represents the interests of the public and the digital society?”) is not trivial and requires balancing the interests of many actors.

Internet Knowledge Plane. Clark et al. [Clark03] describe their vision of a Knowledge Plane (KP), a globally distributed system that extends the Internet with

⁷ <https://transparencyreport.google.com/government-removals/overview>

⁸ <https://transparency.automattic.com>

advanced network management capabilities. The goal of the KP is “a network that can configure itself, that can explain itself, that can repair itself, and does not confound the user with mysteries”. The KP accomplishes this through AI techniques that automatically decide how to configure different parts of the network based on measurements from multiple vantage points, which may be conflicting or incomplete.

The similarity with our work is that a responsible Internet is a global extension of the Internet as well, with the NIP also measuring the network from multiple vantage points like the KP. Another similarity is that the KP supports accountability through an abstract “why” command (returning why something broke in the network) and controllability through a “fix” command (repairs faults in the network).

The main differences with our work are that we focus on providing higher levels of trust and sovereignty for users rather than on automating network management. Also, in our vision (and that of responsible AI), the KP and the network would have to be designed in a transparent and accountable way, which the KP does not consider. Our proposal is furthermore based on open networking, which the KP does not consider.

SpoVNet. Measurements have been used to verify, validate, and improve the functionality of the Internet since its earliest days. Some projects, however, have taken the idea considerably further. The SpoVNet project [Wald10] developed the notion of application-specific overlay networks, where applications would communicate over network paths that were specifically designed and created according to their performance and security needs. An important component was a measurement framework that would run on every participating node and could be invoked by any SpoVNet application to optimize the overlay [Haage09].

Our NCP shares the aspect of controllability as it provides for adjusting parameters and settings for optimized traffic flows. The responsible Internet is not restricted to using overlay networks to achieve this, however. It is mostly agnostic to the Internet’s current or future architecture and enables controllability for any of them. We naturally share the idea of using ongoing measurements. However, in a responsible Internet, they are not only used for optimization but also as a tool to validate information stored in the NIP.

SCION. SCION provides the user with control over the paths that their network traffic takes, on an AS-level [SCION17]. In order to achieve this the user is provided with different paths to the desired destination, if available. This gives transparency of the possible paths and the topology of the ASes, not only to the user but also to the network operators.

SCION’s path transparency and control are excellent building blocks within our proposal for a responsible Internet, as it can provide both input for the network descriptions, through the discovered topology, and enables control over how traffic flows through the Internet based on the analysis of the network descriptions.

The key difference with our work is that our approach is more generic: (1) our network descriptions capture a wide range of security-related attributes at the level of an entire network operator and not just of a specific flows and (2) we enable any user to verify these descriptions, not just the entities on a communications path.

OKN-KISMET. [OKN19] The OKN-KISMET project aims to prototype a knowledge network to improve the security and functioning of three key Internet core systems (naming, addressing, and routing). In particular, the project aims for a reduction and mitigation of abuses of IP address space, routing, and DNS operations.

Their plan is to gather data that can help inform decisions to the end of improved Internet security.

The similarity with our work is that our motivation for the network descriptions has similarities with theirs: they also observe that many potentially useful data sources on the structure of the Internet exist, at various levels of abstraction, but that it remains difficult to extract meaningful information from them to gain knowledge on the structure and evolution of the Internet and to inform policy.

OKN/KISMET focuses on security. It aligns with our work on the NIP in terms of conceptualizing and representing measurements.

9 Conclusions

A responsible Internet takes the Internet into the 2020s because it fulfills the widely supported demand for higher levels of trust and sovereignty for critical infrastructure operators and other types of Internet users. We expect this demand to only increase in the near future as economies and societies are moving online at a further accelerated pace, for instance as a result of the Covid-19 pandemic.

Realizing a responsible Internet is an ambitious undertaking with a wide range of challenges lying ahead, as we have illustrated in this paper. However, we believe it is an attainable goal because several building blocks already exist (e.g., various measurement systems) that can be used as a stepping-stone for the development of a responsible Internet's main components (NIP, NCP, and policies). Also, we expect further thrusts from the close collaboration of practitioners and researchers from different disciplines (e.g., measurements, open networking, and policy development) and from the lessons learned in other complementary areas of technology where the relevance of topics such as trust, sovereignty, and transparency is increasing as well (e.g., cloud services and AI).

While a responsible Internet will put Internet users such as critical infrastructure operators, policy makers, and individuals in the driver's seat, it will also require them to think differently about the Internet: no longer as a black box, but as a crucial piece of machinery that everyone's daily life depends on and that we therefore need to have some level of insight in and control over.

We are looking forward to a wider dialogue with the community to make a responsible Internet the new global communications vehicle of the future.

Acknowledgements

We thank Jeroen van de Lagemaat (NDIX), Valerie Frissen (Leiden University), Victor Reijs (SIDN Labs), Thymen Wabeke (SIDN Labs), and Caspar Schutijser (SIDN Labs) for their valuable feedback on the draft version of this paper.

This work is part of the 2STiC research program (Security, Stability, and Transparency for inter-network Communications), a collaboration of AMS-IX, Delft University of Technology, NDIX, NLnet Labs, SIDN Labs, University of Amsterdam, and University of Twente. Website: <https://www.2stic.nl/>

SIDN and the University of Twente were partly funded by the European Union's Horizon 2020 Research and Innovation program under Grant Agreement No 830927. Project website: <https://www.concordia-h2020.eu/>

University of Amsterdam was funded by the Dutch Science Foundation in the Commit2Data program (grant no: 628.001.001). Project website: <https://dl4ld.nl/>

References

- [Kreutz14] D. Kreutz, F. M. Ramos, P. E. Verissimo, C.E. Rothenberg., S. Azodolmolky and S. Uhlig (2014). Software-defined networking: A comprehensive survey. *Proceedings of the IEEE*, 103(1), 14-76.
- [Knossen19] S. Knossen, J. Hill and P. Grosso, "Hop Recording and Forwarding State Logging: Two Implementations for Path Tracking in P4," 2019 IEEE/ACM Innovating the Network for Data-Intensive Science (INDIS), Denver, CO, USA, 2019, pp. 36-47, doi: 10.1109/INDIS49552.2019.00010.
- [Hill18] J. Hill, M. Aloserij and P. Grosso, "Tracking Network Flows with P4," 2018 IEEE/ACM Innovating the Network for Data-Intensive Science (INDIS), Dallas, TX, USA, 2018, pp. 23-32, doi: 10.1109/INDIS.2018.00006.
- [Ochoa18] Ochoa-Aday, L., Cervelló-Pastor, C., Fernández-Fernández, A., & Grosso, P. (2018). An online algorithm for dynamic NFV placement in cloud-based autonomous response networks. *Symmetry*, 10(5), 163.
- [Koning19] Koning, R., de Graaff, B., Polevoy, G., Meijer, R., de Laat, C., & Grosso, P. (2019). Measuring the efficiency of SDN mitigations against attacks on computer infrastructures. *Future Generation Computer Systems*, 91, 144-156.
- [Coreflow18] Koning, R., Buraglio, N., de Laat, C., & Grosso, P. (2018). CoreFlow: Enriching Bro security events using network traffic monitoring data. *Future Generation Computer Systems*, 79, 235-242.
- [CSCLoud19] Koning, R., Polevoy, G., Meijer, L., de Laat, C., & Grosso, P. (2019, June). Approaches for Collaborative Security Defences in Multi Network Environments. In 2019 6th IEEE International Conference on Cyber Security and Cloud Computing (CSCloud)/2019 5th IEEE International Conference on Edge Computing and Scalable Cloud (EdgeCom) (pp. 113-123). IEEE.
- [Dignum17] V. Dignum, "Responsible Artificial Intelligence: Designing AI for Human Values", *ITU Journal: ICT Discoveries*, Special Issue No. 1, 25 Sept. 2017
- [SCION17] A. Perrig, P. Szalachowski, R. Reischuk. L. Chuat, "SCION: A Secure Internet Architecture", Springer International Publishing, October 2017
- [Chowdhury10] Chowdhury, NM Mosharaf Kabir, and Raouf Boutaba. "A survey of network virtualization." *Computer Networks* 54.5 (2010): 862-876.
- [Cauley19] McCauley, J., Harchol, Y., Panda, A., Raghavan, B., & Shenker, S. (2019). Enabling a permanent revolution in internet architecture. In

- Proceedings of the ACM Special Interest Group on Data Communication (pp. 1-14).
- [Ammar18] M. Ammar, “Ex uno pluria: The Service-Infrastructure Cycle, Ossification, and the Fragmentation of the Internet”, ACM SIGCOMM Computer Communication Review, Vol. 48, Issue 1, January 2018, <https://ccronline.sigcomm.org/2018/ccr-january-2018/ex-uno-pluria-the-service-infrastructure-cycle-ossification-and-the-fragmentation-of-the-internet/>
- [Clark95] D. Clark, “The Design Philosophy of the DARPA Internet Protocols”, ACM SIGCOMM Computer Communication Review, January 1995
- [Gomm15] L. Gommans, J. Vollbrecht, B. Gommans - de Bruijn, C. de Laat, “The Service Provider Group Framework; A framework for arranging trust and power to facilitate authorization of network services”, Future Generation Computer Systems, Vol. 45, blz. 176-192, maart 2015, <http://www.delaat.net/pubs/2015-j-2.pdf>
- [IAB20] IAB Liaison Statement on “New IP, Shaping Future Network”, March 2020, <https://datatracker.ietf.org/liaison/1677/>
- [Böttger17] T. Böttger, F. Cuadrado, S. Uhlig, “Looking for Hypergiants in PeeringDB”, ACM SIGCOMM Computer Communication Review, Volume 47, Issue 1, January 2017
- [Moura16] G. Moura, R. de O. Schmidt, J. Heidemann, W. de Vries, M. Müller, L. Wei, and C. Hesselman, “Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event”, ACM Internet Measurement Conference (IMC 2016), Nov 2016
- [SIDN16] “Czech Knot DNS added to .nl name servers”, October 2016, <https://www.sidn.nl/en/news-and-blogs/czech-knot-dns-added-to-nl-name-servers> [Accessed April 14, 2020]
- [SIDN18] “Further extension of anycast network for .nl”, February 2019, <https://www.sidnlabs.nl/en/news-and-blogs/further-extension-of-anycast-network-for-nl> [Accessed April 14, 2020]
- [Boss14] P. Bossharty, D. Daly, G. Gibby, M. Izzardy, N. McKeownz, J. Rexford, C. Schlesinger, D. Talaycoy, A. Vahdat, G. Varghesex, and D. Walker, “P4: Programming Protocol-Independent Packet Processors”, ACM SIGCOMM Computer Communication Review, Volume 44, Issue 3, July 2014, Pages 87-95
- [Filsfils15] Filsfils, C., Nainar, N. K., Pignataro, C., Cardona, J. C., & Francois, P. (2015, December). The segment routing architecture. In 2015 IEEE Global Communications Conference (GLOBECOM) (pp. 1-6). IEEE.
- [Arkko19] J. Arkko, “Centralised Architectures in Internet Infrastructure”, draft-arkko-arch-infrastructure-centralisation-00, November 2019
- [GAIAX20a] “GAIA-X: The European project kicks off the next phase”, June 2020, https://www.data-infrastructure.eu/GAIAX/Redaktion/EN/Publications/gaia-x-the-european-project-kicks-of-the-next-phase.pdf?__blob=publicationFile&v=6

- [GAIAX20b] “GIAI-X: Technical Architecture”, June 2020, <https://www.data-infrastructure.eu/GAIAX/Redaktion/EN/Publications/gaia-x-technical-architecture.html>
- [Rudin19] Cynthia Rudin, “Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and use Interpretable Models Instead”, *Nature Machine Intelligence*, 2019
- [Zoom20] “Coming April 18: Control Your Zoom Data Routing”, <https://blog.zoom.us/wordpress/2020/04/13/coming-april-18-control-your-zoom-data-routing/> [Accessed May 9, 2020]
- [Zoom20a] “Zoom security issues: Zoom buys security company, aims for end-to-end encryption” (timeline of security issues in Zoom) <https://www.cnet.com/news/zoom-security-issues-zoom-buys-security-company-aims-for-end-to-end-encryption/> [Accessed May 14, 2020]
- [Arkko20] J. Arkko, “The influence of internet architecture on centralised versus distributed internet services”, *Journal of Cyber Policy*, 5:1, 30-45, DOI: 10.1080/23738871.2020.1740753
- [TT20] <https://trans-trust.verisignlabs.com/?z=es>. (queries for Spain)
- [Adrichem15] N. van Adrichem and F.A. Kuipers, NDNFlow: Software-Defined Named Data Networking, Proc. of the 1st IEEE Conference on Network Softwarization (IEEE NetSoft 2015), London, UK, April 13-17, 2015.
- [Carpenter14] Daniel Carpenter and David A. Moss (ed.), Preventing regulatory capture: special interest influence and how to limit it. Cambridge University Press, 2014.
- [IGF19] Internet Governance Forum Berlin, Towards a Global Framework for Cyber Peace and Digital Cooperation: An Agenda for the 2020s, [https://www.hans-bredow-institut.de/uploads/media/default/cms/media/05sk62i_Kleinwa%CC%88chter.%20Kettemann,%20Senges%20\(eds.\),%20Global%20Framework%20for%20Cyber%20Peace%20\(2019\).pdf](https://www.hans-bredow-institut.de/uploads/media/default/cms/media/05sk62i_Kleinwa%CC%88chter.%20Kettemann,%20Senges%20(eds.),%20Global%20Framework%20for%20Cyber%20Peace%20(2019).pdf), 25-29 November 2019.
- [MANRS] Mutually Agreed Norms for Routing Security, <https://www.manrs.org/> [Accessed: May 20, 2020]
- [Jonker19] Jonker, M. (2019). DDoS Mitigation: A Measurement-Based Approach. Enschede: University of Twente. <https://doi.org/10.3990/1.9789036548687>
- [Rijswijk16] van Rijswijk-Deij, R., Jonker, M., Sperotto, A., & Pras, A. (2016). A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements. *IEEE Journal on Selected Areas in Communications*, 34(7), 1877–1888. <https://doi.org/10.1109/JSAC.2016.2558918>
- [Abhishta19] Abhishta Abhishta, Roland van Rijswijk-Deij, and Lambert J. M. Nieuwenhuis. 2019. Measuring the impact of a successful DDoS attack on the customer behaviour of managed DNS service providers. *SIGCOMM Comput. Commun. Rev.* 48, 5 (January 2019), 70–76. DOI:<https://doi.org/10.1145/3310165.3310175>
- [Sommese20] R. Sommese, G. Moura, M. Jonker, R. van Rijswijk-Deij, A. Dainotti, k. claffy, and A. Sperotto, "When parents and children disagree: Diving into DNS delegation inconsistency", in *Passive and Active Measurement Conference (PAM)*, Mar 2020.

- [OKN19] K. Claffy and G. Voelker, OKN-KISMET Project Proposal, 2019, https://www.caida.org/funding/okn-kismet/okn-kismet_proposal.xml
- [Schlamp16] Johann Schlamp, Ralph Holz, Quentin Jacquemart, Georg Carle, Ernst W Biersack, “HEAP: reliable assessment of BGP hijacking attacks”, IEEE Journal on Selected Areas in Communications (34:6), 2016.
- [Turkovic20] B. Turkovic, J. Oostenbrink, F.A. Kuipers, I. Keslassy, and A. Orda, Sequential Zeroing: Online Heavy-Hitter Detection on Programmable Hardware, Proc. of IFIP Networking 2020, Paris, France, June 22-25, 2020
- [Turkovic18] B. Turkovic, F.A. Kuipers, N. van Adrichem, and K. Langendoen, Fast network congestion detection and avoidance using P4, Proc. of the ACM SIGCOMM 2018 Workshop on Networking for Emerging Applications and Technologies (NEAT 2018), Budapest, Hungary, August 20, 2018
- [Polachan20] K. Polachan, B. Turkovic, T.V. Prabhakar, C. Singh, and F.A. Kuipers, Dynamic Network Slicing for the Tactile Internet, Proc. of the 11th ACM/IEEE International Conference on Cyber-Physical Systems (ACM/IEEE ICCPS 2020), Sydney, Australia, April 22-24, 2020
- [Riftadi19a] M. Riftadi and F.A. Kuipers, P4I/O: Intent-Based Networking with P4, Proc. of the 2nd IEEE International Workshop on Emerging Trends in Softwarized Networks (ETSN 2019), June 28, 2019
- [Riftadi19b] M. Riftadi, J. Oostenbrink, and F.A. Kuipers, GP4P4: Enabling Self-Programming Networks, arXiv:1910.00967
- [CoE19] Council of Europe, Strategy on Internet Governance (2016-2019) Final report, 2019, <https://rm.coe.int/168098fc7d>
- [EC20] European Commission, Horizon 2020 Work Programme 2018-2020 Information and Communication Technologies, 2020, https://ec.europa.eu/research/participants/data/ref/h2020/wp/2018-2020/main/h2020-wp1820-leit-ict_en.pdf
- [Froomkin03] Froomkin, A. Michael , Habermas@discourse.net: Towards A Critical Theory of Cyberspace, Harvard Law Review Vol. 116 January 2003 Number 3.
- [Gold06] Goldsmith, Jack and Wu, Tim, “Who Controls the Internet? Illusions of a Borderless World.” Oxford University Press, 2006.
- [Mueller17] Mueller, Milton, “Will the Internet Fragment? : Sovereignty, Globalization and Cyberspace.” Polity Press, 2017.
- [UN19] UN, The age of digital interdependence Report of the UN Secretary-General’s High-level Panel on Digital Cooperation, 2018, <https://www.un.org/en/pdfs/DigitalCooperation-report-for%20web.pdf>
- [Wu18] Wu, Tim, The Curse of Bigness: Antitrust in the New Gilded Age, Columbia Global Reports, 2018.
- [Xue19] Xue, Janet, Hui and Holz, Ralph, Applying smart contracts in online dispute resolutions on a large scale and its regulatory implications, In: Massimo Ragnedda, Giuseppe Destefanis (eds) Blockchain and Web 3.0 Social, Economic, and Technological Challenges, Routledge, London, 2019, pp.177-191.
- [Trillian20] Google. Trillian. A transparent, highly scalable and cryptographically verifiable data store. <https://github.com/google/trillian>

- [Wald10] Oliver P Waldhorst, Christian Blankenhorn, Dirk Haage, Ralph Holz, Gerald G Koch, Boris Koldehofe, Fleming Lampi, Christoph P Mayer, Sebastian Mies, Spontaneous Virtual Networks: On the Road Towards the Internet's Next Generation. *it-Information Technology* (50:6), 2010.
- [Haage09] Dirk Haage, Ralph Holz, Heiko Niedermayer, Pavel Laskov, “CLIO—a cross-layer information service for overlay network optimization”. *Proc. Kommunikation in Verteilten Systemen (KiVS)*, 2009.
- [Maffione16] V. Maffione, F. Salvestrini, E. Grasa, L. Bergesio, and M. Tarzan, “A software development kit to exploit RINA programmability”, 2016 IEEE International Conference on Communications (ICC), Kuala Lumpur, Malaysia, May 2016
- [Abera16] T. Abera, N. Asokan, L. Davi, F. Koushanfar, A. Paverd, A. Sadeghi and G. Tsudik, “Things, Trouble, Trust: On Building Trust in IoT Systems”, Design Automation Conference (DAC), 2016
- [Lehr15] William Lehr, Erin Kenneally, and Steven Bauer, “The Road to an Open Internet is Paved with Pragmatic Disclosure & Transparency Policies,” TPRC 43: The 43rd Research Conference on Communication, Information and Internet Policy, 2015
- [Lehr19] William Lehr, David Clark, Steve Bauer, Arthur Berger and Philipp Richter, “Whither the Public Internet?,” *Journal of Information Policy*, Vol. 9 (2019), pp. 1-42
- [Pappas15] C. Pappas, K. Argyraki, S. Bechtold, A. Perrig, “Transparency Instead of Neutrality,” *HotNets '15*, November 16–17 2015, Philadelphia, PA, USA
- [Fragkouli19] G. Fragkouli, K. Argyraki, and B. Ford, “MorphIT: Morphing Packet Reports for Internet Transparency,” *Proceedings on Privacy Enhancing Technologies*; 2019 (2):88–104
- [WODC19] J.M. Ceron, J.J. Chromik, J.J. Cardoso de Santanna, A. Pras, “Online discoverability and vulnerabilities of ICS/SCADA devices in the Netherlands”, Tech Report, University of Twente, June 2019
- [EIT20] “European Digital Infrastructure and Data Sovereignty -- A Policy Perspective”, EIT Digital Report, June 2020, <https://www.eitdigital.eu/newsroom/news/article/new-report-on-european-digital-infrastructure-and-data-sovereignty/>
- [ENISA19] “Consultation paper: EU ICT Industrial Policy: Breaking the Cycle of Failure”, ENISA, July 2019, <https://www.enisa.europa.eu/publications/enisa-position-papers-and-opinions/eu-ict-industry-consultation-paper>
- [O’Hara20] K. O’Hara and W. Hall, “Four Internets: Considering the merits of several models and approaches to Internet governance”, *Communications of the ACM*, Vol. 63, No. 3, March 2020
- [CS16] “European Cybersecurity Industry Leaders Recommendations on Cybersecurity for Europe”, January 2016, http://ec.europa.eu/newsroom/dac/document.cfm?action=display&doc_id=13326
- [NYT20] “E.U. Recommends Limiting, but Not Banning, Huawei in 5G Rollout”, *New York Times*, January 2020,

- <https://www.nytimes.com/2020/01/29/world/europe/eu-huawei-5g.html>
[Accessed: June 18, 2020]
- [NDL07] Jeroen van der Ham, Paola Grosso, Ronald van der Pol, Andree Toonk, Cees de Laat, “Using the Network Description Language in Optical Networks”, Tenth IFIP/IEEE International Symposium on Integrated Management (IM 2007), 21-25 May 2007, Munich, Germany. May 21 2007, ISBN: 1-4244-0799-0, page 199-205
- [NDL13] M. Ghijsen, J. van der Ham, P. Grosso, C. Dumitru, H. Zhu, Z. Zhao, C. de Laat, “A Semantic-Web Approach for Modeling Computing Infrastructures”, Journal of Computers and Electrical Engineering, Elsevier, Volume 39, Issue 8, November 2013, Pages 2553-2565
- [CIDR20] Plot AS Count, CIDR website, <https://www.cidr-report.org/as2.0/>
[Accessed May 7, 2020]
- [2STiC20] P. Grosso, C. Hesselman, L. Hendriks, J. Hill, R. van der Pol, V. Reijs, R. van Rijswijk-Deij, J. de Ruiter, C. Schutijser, and H. Steenman, “A national programmable infrastructure to experiment with next-generation networks”, blog, March 2020, <https://www.2stic.nl/national-programmable-infrastructure.html>
- [Astuto14] B. Nunes Astuto, M. Mendonca, X. Nam Nguyen, K. Obraczka, T. Turetli, “A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks”, IEEE Communications Surveys & Tutorials, Vol. 16, Issue 3, Third Quarter 2014
- [Ceron20] J. Ceron” DDoS Clearing House”, CENTR Jamboree 2020, May 2020, <https://www.sidnlabs.nl/downloads/5NMQ6vFghWvzFFnRikGaAS/ead763585ff6b4de90a0839bac89dc92/SIDN-labs-CENTR-ddos-talk-ceron.pdf>
- [Brown10] Ian Brown, D. D. Clark, and D. Trossen, “Should Specific Values Be Embedded In The Internet Architecture?”, ACM ReArch 2010, November 30, 2010, Philadelphia, USA
- [Chromik19] J. J. Chromik, “Process-aware SCADA traffic monitoring: A local approach”, Ph.D. thesis, University of Twente, 2019
- [Leverett17] E. Leverett, R. Clayton, and R. Anderson, “Standardisation and Certification of the ‘Internet of Things’”, 16th Annual Workshop on the Economics of Information Security (WEIS2017), USA, June 2017
- [Mirai17] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z., Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, “Understanding the Mirai Botnet”, 26th USENIX Security Symposium, 2017
- [Clark03] D. Clark, C. Partridge, J.C. Ramming, and J.T. Wroclawski, “A Knowledge Plane for the Internet”, SIGCOMM’03, August 25–29, 2003, Karlsruhe, Germany
- [EC19] European Commission, Ethics Guidelines for Trustworthy AI, 2019, <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>

- [EP20] European Parliament, Digital Agenda for Europe, 2020, <https://www.europarl.europa.eu/factsheets/en/sheet/64/digital-agenda-for-europe>
- [Silva19] K. e Silva, “Mitigating botnets: Regulatory solutions for industry intervention in large-scale cybercrime”, Ph.D. thesis, Tilburg University, the Netherlands, Dec 2019