An Open, Globally-connected, Secure, and Trustworthy Internet

Version 0.9.4

1 Introduction

The Internet Society's Mission guides the work of the Internet Society. It covers several elements, and one part of it includes the following strategic goals:

> Our work aligns with our goals for the Internet to be **open**, **globally-connected**, **secure**, and **trustworthy**. We seek collaboration with all who share these goals.[3]

Since the achievement of these strategic goals is part of the Mission, it is important to understand what the terms mean. That is the purpose of this memo. It is important to emphasize that the present memo does not seek to explain the Mission overall, but only to define these terms. The Mission should be treated as an entirety when judging whether the Internet Society has the correct priorities.

1.1 A word on "The Internet"

In the Mission, the term "the Internet" is used, but it is not defined. For the purposes of this memo, that lacuna is noted but unfilled. Some discussion of the topic may be found in Annex A, but the subject is also touched upon by a separate Internet Society effort about the Internet Way of Networking. [7] The interested reader should consult both of those for additional discussion of the meaning of "the Internet", but it is not the central focus of this memo.

2 The definitions and how to use them

The purpose of these definitions is not to provide transcendental truth or to fix forever the meaning of these terms. Instead, the point of these definitions is practical: we want to know whether the Internet Society is making progress towards certain goals. To do that, we need to define the terms in ways that can be tested. Such definitions are called *operational definitions,* because they specify some kind of operation to perform to see whether a given term can be applied in a given case. In other words, a good operational definition provides

all the necessary and sufficient conditions to know whether a given term can be applied to a given state of affairs in the world.[1]

A tricky thing about operational definitions of terms like the ones under consideration is that they are not binary properties. Many operational definitions (even very complicated ones, like the definition of "voter" in a given jurisdiction) yield a binary result (e.g. one is either allowed to vote, or not allowed to vote). But the terms we are talking about here define properties that can be more or less true of the Internet as a whole. Plainly, also, the terms are at least partly compositional, in that if every part of the Internet fails to have some property, then it seems impossible that the Internet as a whole could have that property. At the same time, it is not strict composition, because (for example) even if every part of the Internet were independently secure, it would be possible for their interactions to create a globally insecure system.

The best way to think of each of these definitions, therefore, is as idealized statements of end states. In theory, if each of these properties were (somehow) achieved completely, then at least for that brief moment the Internet Society Mission as stated would be complete.[2] And yet, because the properties may be in conflict with one another, it may not even be logically possible that all four of these terms be completely true of the Internet at one time. As a result, each definition is bound to be an unrealistic totality, because it specifies a state that is likely impossible either on its own or in conjunction with other properties of the Internet. That is not a failing of the definitions, but instead a reality that we must embrace and work within, while keeping our goals in mind. It might even be that the achievement of any of the properties in absolute terms would be undesirable for the Internet as a whole, because one property might tend to prevent the realization of another.

These definitions are about network connections, because the portion of the Mission in which they appear describes the Internet, and not the uses to which the Internet is put. The reader might therefore object that there is nothing in these definitions that talks about the value of the Internet for humans. These terms are not, however, the only important parts of the Mission: it also speaks of wishing to ensure the Internet is, "…a resource to enrich people's lives, and a force for good in society." That tells us *why* the Internet Society works to achieve the goals. The enumerated goals, on the other hand, are about the Internet and are therefore defined here mostly without direct reference to people, even if people's use of the Internet is ultimately what the Mission is about.

The purpose of these definitions is to guide the Internet Society staff in making judgements about whether a given activity or objective moves the Internet

[1] Operational definitions are not magic: they can be simply wrong on the one hand, or they can provide inadequate operations. For instance, "the color blue" could be operationalized as "light with a dominant wavelength between 450-495 nanometres," or it could be operationalized as "the primary color that is not red and not yellow". If instead one offered, "Blue is the color of the sky," it would be an inadequate operation: comparing a patch of color to the sky on a rainy day would yield the wrong result. If instead one said, "blue is light with a dominant wavelength between 450-495 metres," it would simply be wrong.

[2] It is apparent that, since the Internet is not a static thing, even this imaginary state of affairs would probably not be stable, so the Mission would remain important.

Society closer to the strategic goals or does not. Therefore, these are not to be misunderstood as somehow defining the work of staff, but rather as a tool through which the Internet Society as a whole can evaluate moves intended to advance the organization's Mission. The primary audience for this document is the staff and Board of Trustees, because it provides a framework for evaluating actions at the highest level; therefore, it is not a policy of the Internet Society in the meaning of the Policy Development Process [4] and has not been thoroughly vetted via community comment. It can be considered a staff-internal document but not confidential. The reader should be aware that some members of the community are unlikely to be aware of these definitions and may not be working with them. Moreover, the reader should be prepared for a level of formality and rigor that would not be appropriate for a document with a less-formal purpose.

In each of the following sections, each term is first provided an operational definition. These definitions are formal ones, and as a result they tend to be stiff and sometimes seem complicated. That is appropriate for formal definitions of this sort, but each section contains a colloquial definition that is intended to capture the meaning of the term well enough in most circumstances, in order to make talking about the terms somewhat easier. The colloquial expressions cannot, however, be used formally.

2.1 Open

Open The Internet is completely open if and only if any person or organization is free to take any technology that could make up the Internet, and use that technology to build other things or even combine them in novel ways; and to deploy them without legal, regulatory, or other barriers; and to expect that others may choose to do the same in interoperable ways. The opposite of "open" is "closed".

Openness is multidimensional, because it is a feature of the system that tends to have echoes in other systems (even social systems) that touch the Internet but that are not strictly speaking a part of it. In part that explains why there are so many things that follow from the above definition. To begin with, barriers to access or cases where the neutrality of the network is reduced (whether that be due to business reasons or laws) are both limits on openness. Some logical barriers to openness might be physical limits: for instance, there is a limit to who can use a given part of the radio spectrum at any one time, and so if spectrum is in use by one operator in a location that necessarily limits the openness of spectrum-using technology to someone else who might want to do something with it.[3] Other barriers to openness might arise from that physical limit, however. A licensed operator might impose restrictive terms on what content a network using the licensed spectrum may carry, or might limit the services that may be offered. A network that meets all of the criteria for openness but only

[3]This is an example of the ultimate state being effectively impossible to achieve, since it is never going to be possible to allow two conflicting uses of the same spectrum at once, even though doing so would be more open.

through the payment of a lot of money might be regarded as closed, because of the barrier to entry in the form of a high price. (In this way, openness is related to the property of being globally-connected; see section 2.2. Regulatory attempts to prevent communication over the Internet, or "Internet censorship", are sometimes attacks on openness and sometimes attacks on connectedness.)

Technologies that do not interoperate with others (such as proprietary protocols and Application Programming Interfaces, or APIs) are in some sense necessarily closed, though they may behave in an open fashion if sufficiently publicly defined and with generous enough license terms. At the same time, merely being open does not guarantee interoperability: openness requires acceptance that others might not choose to interoperate, and that is no proof that the others prefer a closed network. In other words, openness does not create a promise that something will work on the network. Nevertheless, an open technology is more likely than a closed one to produce interoperable systems. This is one reason why both open standards and open source software are so often associated with the openness of the Internet: they may not be necessary for an open Internet, but they tend to increase the chances that interoperable technologies get deployed. When open source software is somehow excluded from the market that may be a bad leading indicator for the openness of the Internet too.

The meaning of of "any technology" in the above is not limited to any particular kind of technology or layer of the Internet. For instance, to the extent that people's experiences of the Internet are mediated almost entirely through closed apps that do not use standard protocols, the Internet might be judged to be more closed. To the extent that many people (even those technically unsophisticated) are able to build new, reusable functionality themselves in a social media platform, the Internet might be judged to be more open even if that platform is closed. There is also an equivocation on "technology" in this definition, because it comprises both specification (e.g. open standards) and actual implementations that interoperate.

It may be that complete openness conflicts with other socially-desirable outcomes. It might therefore be in tension both with the issues discussed in section 2.3 and with wider social security concerns that are not really Internet security issues. That is not, however, an argument that closed systems are generally, or that closing parts of the Internet would be, desirable. Closed systems are often touted as being more secure, but there is scant evidence for that claim.

Colloquially we can say that an open Internet is one where anyone may create, use, or deploy the Internet and its technologies according to their own wishes. See also the discussion of "an open architecture of interoperable and reusable building blocks" (critical property two) in [7].

2.2 Globally-connected

Globally-connected The Internet is completely globally connected if and only if all latent demand for a connection is satisfied and those connections

> are all "full" connections with no restrictions on their use, such that any arbitrary pair of nodes that are connected to the Internet may exchange packets with one another without any interference imposed by any network along the path or by any other party. The opposite of "globally-connected" is "disconnected".[4]

It is important to recognize that this definition is a definition of *potential*, not actual, communication. In other words, to have a globally-connected Internet is not to require that every part of the network be talking to every other part (even indirectly) all the time. Instead, in a globally-connected Internet nothing beyond the negotiation of connection is necessary for the connection to happen. For example, in the case of TCP, only the three-way handshake that establishes a connection is necessary – not special permission, or negotiation of new contracts, or the like. Importantly, this entails that each node on the Internet works in theory the same way, at least at the lowest level of connectivity.

It is also important to emphasize that any arbitrary pair of nodes *may* exchange packets: that does not mean they must. The decision to connect with anything else is in the hands of each node, and it is not a violation of the idea that the Internet is globally-connected that one node cannot talk to another node because one node rejects traffic (or because the nodes have no need to exchange traffic). There is admittedly a tension here, because the reasons for the rejection of the traffic matter. From the outside, there is no technical difference between, "Network N decided not to exchange traffic," and, "Network N was ordered by its government not to exchange traffic." Yet one of those implies interference with the globally-connected Internet (on the part of the government in question), whereas the other is a perfectly normal local decision made by autonomous network operators (for example, by deploying a firewall).

The use of "latent demand" here means not just that a mere connection is possible, but that the connection and resulting bandwidth be adequate to the needs of the person or thing connecting. The reason for this is plain when one considers that, for practical purposes, there is no difference among a network that is so busy (congested) that nobody can get anything through, an unreliable network link that goes up and down rapidly, and no network at all.

The global property of connectedness is an emergent property that results from the state of all network nodes, but it is desirable in its own right because it is the basis for network effects. If a network becomes more valuable in some relationship to the number of nodes it has, then it follows that an impediment to more connections is also an impediment to the value of the whole network. By the same token, active efforts to break connectivity are efforts contrary to either the globally-connected Internet, or the open Internet, or both (cf. section 2.1). That includes efforts to restrict or interrupt Internet communications for political and social reasons, whether temporarily or permanently. Such efforts might be desirable to some actors (even very important ones, such as nation

[4] *Not* globally-disconnected, note. A number of network "islands" that cannot talk to each other are connected, but not globally-connected.

states), but they are nevertheless in opposition to these goals of the Internet Society.

One challenge to the value of being globally-connected is that not every device that benefits from *some* connectivity benefits from total connectivity. For instance, a printer might benefit from being able to fetch firmware updates from the Internet, but that does not mean that it should be connected to any arbitrary node on the Internet. Many kinds of devices are really intended to be "connected to the Internet" only notionally. For example, in many Internet of Things systems, it is often only the total system of sensors, actuators, and controllers that make up the whole, and only the controller should be *really* connected to the Internet; the other devices should properly be mediated by the controller. In another way, the connectivity of some nodes can actually be a threat to other nodes, such as when a node connects for the malicious purpose of attacking others on the Internet (see section 2.3 for more on this). These challenges present practical limitations to the idea of an Internet where everything is globally-connected, and emphasize the importance of keeping in mind the reasons why a given device ought to be connected in the first place.[5]

Colloquially we can say that a globally-connected Internet is one in which everything and everyone that wants a connection can get one that is good enough to satisfy the desire, *or* that any set of nodes that wish to interconnect can do so in a manner that is good enough to satisfy the desire. See also the discussion of "an accessible infrastructure with a common protocol" (critical property one) in [7].

2.3 Secure

Secure The Internet is completely secure if and only if it does not lower the confidentiality, integrity, or availability of any information or service, whether that information or service be the Internet and its traffic, or information or services outside the Internet where the Internet is itself the source of the lowering. The opposite of "secure" is "insecure".

There are two distinguishable senses to security of the Internet. The first is the security *of the Internet itself*, which is to say the infrastructure of the Internet broadly defined as well as the traffic on the Internet. Denial of service attacks, misdirected traffic (whether through intention or accident), data that is transmitted insecurely when it ought to be transmitted securely, and accidental and deliberate severing of cables are all examples of vulnerabilities or threats the the Internet and its traffic, and so are examples of insecurity of the Internet itself.

[5]There is a related issue of Network Address Translation – NAT [13]– which generally has the property that the node behind a NAT cannot receive traffic without negotiating a connection first. NATs present challenges to global connectivity, it is certain. Given the ubiquity of NAT as a tool to extend the lifetime of the IPv4 address space, however, it would be absurd to claim that a system that relies on a NAT to connect is not really connected. This is yet another example of how these properties may be more or less true rather than being binary.

(For the purposes of clarity, let us call this sense of security *infrastructure security*, even though it is not strictly limited only to infrastructure.) Insecurity produced by the failure of *infrastructure security* can have adverse effects on the connectivity of the Internet, and mitigation attempts can cause the Internet to be more closed.

The second sense of security of the Internet arises where it is the fact of the Internet's use that creates the (in)security. It is important not to cast this net too widely, but it is a very large net anyway. Botnets that are used in phishing scams to steal banking information, for instance, are an example of this sort of insecurity: neither the means for the attack on bank accounts nor the method of that theft (through tricking people into delivering their credentials) would be available without the Internet. Impersonation and fraud are not new human behavior, but combining those into the mechanism of phishing is something novel. (This sense of security is closely related to the issue of trustworthiness, considered in section 2.4. For the purposes of clarity, let us call this second sense *use security.*)

Not every vulnerability that happens through the Internet is an "Internet security" issue in the narrow sense used here. Still, some cases clearly are correctly classified as Internet (in)security. For example, when personal information about people is leaked onto the Internet, it is not *necessarily* a problem of Internet security as such, even though it is terrible. On the one hand, imagine a case where data leaked through the exploitation of vulnerable systems that are connected to the Internet. It is reasonable to say that the system vulnerabilities fall at least in part under *use security.* (For example, a system that exists to provide validation of login credentials, to collect name and address information, and to facilitate payments for online transactions is effectively an Internet system; if people's credit card data leaks from such a system it is at least in part an Internet *use security* problem.) If, on the other hand, a system for storing classified documents is subverted so that someone carries such documents out of a secured and disconnected system and uploads them somewhere, so that the documents contained in that system appear on the Internet, that is not an Internet security problem properly speaking, even though it is a data breach. We would not think that theft of money from an unlocked and unguarded bank vault was a security problem of bank vaults themselves, even though it might be thought of as a "funds breach".[6] Similarly, if the data leaked because someone intentionally exposed data on the Internet even though it was data that should not have been so exposed, that is not different in kind to printing the data out and posting it on a public bulletin board: the Internet is actually doing what it is designed to do in that case, so it is a matter of bad judgement (such as a bad system design) or maliciousness (such as a data breach) by the person who

[6]It is a case by case evaluation whether a given case is properly considered an Internet security (*use security*) problem, or just a system security problem of something that happens to be connected to the Internet. It is unlikely that one can draw a clear distinction in general, but one indication is likely to be the extent to which the compromised system was supposed to be connected to the Internet or somehow linked to it in the first place; the evaluation must therefore fall on a spectrum.

posted the data. This distinction is important because quite often the latter cases are treated as "Internet problems", rather than something else. When bank robbers get away in a car, we do not treat the escape in a car as something that ought to be solved by the automobile industry. By way of analogy, then, if someone who has (otherwise authorized) access to data does something with that data that was not authorized, it cannot be treated as something the Internet itself should have foiled.

Similarly, social changes that make people feel less secure and that happen because the Internet has made information more widely accessible is not an issue of Internet security, even when it affects the security of societies. Technologies, both when they are invented and as they evolve, present new challenges to the societies where they are deployed; the Internet is no different. The practice of "SWATting"[7] is a terrible thing, and it is made considerably easier because of the Internet. But it is strictly speaking a vulnerability in the police protocols, which can be fooled into false responses, rather than a problem of Internet security even broadly defined.

By the same token, insecurities (in both senses) visited upon people due entirely to the Internet are clearly part of the issue of security of the Internet, even if those insecurities are not themselves intrinsically Internet issues. Identity systems and financial systems are connected to getting many things done over the Internet, and so they might each represent a danger to *use security*. In addition, when one combines more than one system, even if those systems are independently secure, one may get a system that is insecure (in either sense). Moreover, some of those insecurities are really insecurity for people.This means that some privacy issues on the Internet turn out to be *use security* issues, because it is the fact of the Internet that erodes privacy in that way. Similarly, the invention of falsehoods for political purposes may be as old as language itself, but the nature of how to publish it is affected by the Internet, so there is at least some component of *use security* in such cases.

It is important to notice that the definition talks about vulnerabilities and threats, and does not consider whether they are known to anyone. Unknown vulnerabilities are still vulnerabilities, and keeping systems closed to avoid addressing those vulnerabilities does not materially improve the security of the Internet.

It is also worth noting that the distinction between *infrastructure* and *use security* may be mostly academic in many cases, since a variety of different attack techniques might be in use at once. For practical purposes the distinction

[7]This is an attack on a person where the Special Weapons And Tactics team of local police are called to the victim's house on the pretext of a serious public security threat, when in fact there is no such threat. Note that in such an attack, the victim might have been "doxed", which is where quantities of personal information about a victim are released on the Internet. There is a difficulty in treating the whole issue as an "Internet security" problem because that conflates ways in which the Internet could be improved – guarding personal data – and areas where it is unable to influence the outcome, such as police response protocols. Popular discourse is likely to treat this as a single issue, but Internet Society analyses need to keep the different parts separate in order to distinguish ways in which the Internet could (or could not) improve.

may not matter, but for formal purposes such as those of this document the distinction is nevertheless worth making.

Colloquially we can say that the Internet is secure when it does not make any person, organization, or society overall more vulnerable than they are without it.

2.4 Trustworthy

Trustworthy The Internet is completely trustworthy if and only if it is completely resilient, reliable, accountable, and secure in a way that consistently meets users' expectations for information and services. The opposite of trustworthy is untrustworthy.

Unlike the other definitions, this definition depends not only on the state of the Internet, but also on the state of people. For it specifies that the Internet is only worthy of trust when it conforms with what people expect will happen. The definition is silent on whether those expectations are reasonable, though it implicitly assumes reasonable expectations given the audience.

It is important to note that this definition is totally independent of whether someone actually trusts the Internet. The property of trustworthiness is something that is true (or false) of the object in question: whether it is something that ought to be trusted, irrespective of the psychological state of anyone observing it. A new kind of padlock, for instance, might be utterly worthy of trust because it is constructed of some new kind of metal that is extremely strong even when slender. Someone might nevertheless be unwilling to trust it because it does not look as stout as other high-strength locks in the market. At the same time, because trustworthiness depends on meeting users' expectations, if *nobody* trusts the Internet then it cannot possibly be trustworthy.

Trustworthiness is not a matter of only one layer in the network, and it is possible that some parts of the Internet are trustworthy while other parts are not. For instance, some applications might be untrustworthy even if routing became trustworthy through the universal deployment of complete cryptographic routing validation. Since the definition is in terms of users' expectations, even behavior that is not Internet-specific (such as fraud) might undermine the trustworthiness of the Internet, if there is no way to mitigate that behavior.

To be useful, trustworthiness depends upon an informed base of users who have the tools to evaluate trustworthiness. Otherwise, their expectations may be misaligned with what the technology can provide, and the system will not be trustworthy. For this reason, trustworthiness really ought to be judged against the standard of the likely user population, and what their expectations probably will be. The expectations of a computer security researcher are unlikely to be acceptable in a technology to be delivered to a historically-unconnected audience. Counterintuitively, lowering users' expectations may be a route to trustworthiness, though that might undermine the usefulness of the Internet (and militate against other parts of the Mission). For instance, in the earliest

days of the World Wide Web, before the Secure Sockets Layer (SSL) was invented (and long before it had grown into Transport Layer Security – TLS), most users did not expect to be able to use their credit cards online, so the system was trustworthy given the expectations. But an Internet without TLS would not be trustworthy for financial data of any importance.

If the Internet is insecure, it cannot possibly be trustworthy. However, a completely secure Internet could still be untrustworthy by violating user expectations or by being unreliable or unavailable when it is needed. Trustworthiness is not only a matter of security.

Colloquially we can say that the Internet is trustworthy when it is designed, built, and deployed so that people can justifiably believe it will not be a threat and not betray their expectations.

A The Internet

As humans interact through the Internet more and more, the meaning of "the Internet" loses focus, because it means different things to different people. For some people, "the Internet" is technical infrastructure; for others, "the Internet" is the whole of the online environment. Indeed, many publications no longer capitalize the term[2], so ubiquitous is it. For the purposes of the Internet Society, however, "the Internet" is not just anything that is online.

There have been several attempts to specify what the Internet is. The Federal Networking Council (FNC) developed one such definition[1]:

> "Internet" refers to the global information system that --
>
> (i) is logically linked together by a globally unique address space based on the Internet Protocol (IP) or its subsequent extensions/follow-ons;
>
> (ii) is able to support communications using the Transmission Control Protocol/Internet Protocol (TCP/IP) suite or its subsequent extensions/follow-ons, and/or other IP-compatible protocols; and
>
> (iii) provides, uses or makes accessible, either publicly or privately, high level services layered on the communications and related infrastructure described herein.

A similarly technically-oriented but more compact formulation is attributed to Seth Breidbart: that the definition of "the Internet" is, "The largest equivalence class in the reflexive, transitive, symmetric closure of the relationship 'can be reached by an IP packet from'."[5] Definitions like this, which we might call the infrastructure view, make "the Internet" distinct from services and applications that sit atop the Internet. In this sort of formulation, the Internet is the infrastructure and applications and services are merely using the Internet as part of their style of delivery. FNC's definition requires a globally unique address space, which seems to imply that nodes that use private IP address space [12] would not qualify. Breidbart's formulation does not suffer from this issue, but it still means that many things that sure seem like important elements of the Internet do not qualify as part of it. For instance, the Domain Name System [10, 11] is clearly an important part of the Internet – it is even arguably a big part of the Internet infrastructure – but it would not qualify as "the Internet" under these definitions: only the clients and servers that make up the actually-operating DNS would qualify, and not the system itself. Such a definition of the Internet seems inadequate (at least for the purposes of talking about how the Internet affects us all). It would be like talking about transportation and mentioning only roads or airports but not vehicles or fuel supply chains.

Another way to look at the Internet is as a service [6], through which one can be more or less connected[8]. We can call this view the infrastructure-service view, because it places some emphasis on the use-value of the infrastructure itself. While this is an important element, however, it does not really suffice to say what the Internet is.

One reason the definition of the Internet is problematic is because the word does not refer to a single thing. As Ed Krol and Ellen Hoffman [9] observe, "A commonly asked question is 'What is the Internet?' The reason such a question gets asked so often is because there's no agreed upon answer that neatly sums up the Internet." In fact, this has to do with the nature of the Internet itself. The Internet is made up of many other networks, which themselves may be made of other networks (and so on). Moreover, those different networks have different purposes and so different properties: the respective networks of a home-subscriber Internet Service Provider, of a social media company, of a backhaul carrier, and of a Content Delivery Network are each different, and yet each is a part of the Internet. In addition, these boundaries are not fixed. An application that merely happens to use the Internet to carry messages, for instance, is probably not "part of the Internet" strictly speaking: it's just a user. But if that application starts to interact with, feed into, or be depended upon by other applications, then it ceases to be merely a *user* of the Internet and starts to become *part of* the functioning of the Internet. In this way, for example, social media platforms often start out *built atop* the Internet, in that they use it for carrying content but are otherwise separate from the Internet (often as "walled gardens" of networked experience). But because social media platforms are powerful and can often enable people to build new things atop the platform (hence the name), the platforms soon become *part of* the Internet because of those other things built on top.

Because this ambiguity in the meaning of "the Internet", it may not be possible to come up with an operational definition that covers all the meanings we can find even in technical contexts, much less those meanings that are used in casual conversation. It may be better to stop seeking a precise definition, and to focus on what distinguishes the Internet from other kinds of technologies. That is the strategy that the Internet Society has adopted in identifying the critical properties of the Internet, without which the Internet is not possible. These properties have been identified as part of the Internet Way of Networking. [7]

Regardless, it is essential to recognize that the Internet never refers to a single, unitary thing, but rather to a collection of other things that come together to make up the Internet. For the Internet Society, the plasticity of the meaning of "Internet" is both a danger and an opportunity. The danger is that anything that is remotely related to the Internet will be cast as an "Internet issue" when the Internet is just accidental to the thing in question. The opportunity is to get people to understand that the Internet is a profoundly human technology that, in its true form, is reconfigurable and reusable according to the needs and wants of those who use it.

Acknowledgements

This memo would not have been possible without observations and review from the Internet Society community. Jaap Akkerhuis, Peter Koch, and Paul Wilson each provided particularly helpful comments on some portions. Text in this

this memo includes significant contributions from David Belson, Katie Bengaard, Toral Cowieson, Carl Gahnberg, Katie Watson Jordan, Olaf Kolkman, Konstantinos Komaitis, Andrei Robachevsky, and Tom Stark of the Internet Society staff.

References

[1] Fnc resolution: Definition of "internet", October 1995. URL https://www.nitrd.gov/fnc/internet_res.pdf.

[2] April 2016. URL https://blog.ap.org/products-and-services/ready-to-lowercase-internet-and-web.

[3] 11 2017. URL https://www.internetsociety.org/mission/.

[4] Policy development process at the internet society, April 2020. URL https://www.internetsociety.org/about-internet-society/policy-development-process/.

[5] Seth Breidbart. The internet. Webpage, 03 2013. URL http://www.bcp38.info/index.php/The_Internet.

[6] D. Crocker. To Be "On" the Internet. RFC 1775 (Informational), March 1995. ISSN 2070-1721. URL https://www.rfc-editor.org/rfc/rfc1775.txt.

[7] Internet Society. Internet way of networking, 2020. URL https://www.internetsociety.org/issues/internet-way-of-networking/.

[8] J. Klensin. Terminology for Describing Internet Connectivity. RFC 4084 (Best Current Practice), May 2005. ISSN 2070-1721. URL https://www.rfc-editor.org/rfc/rfc4084.txt.

[9] E. Krol and E. Hoffman. FYI on "What is the Internet?". RFC 1462 (Informational), May 1993. ISSN 2070-1721. URL https://www.rfc-editor.org/rfc/rfc1462.txt.

[10] P.V. Mockapetris. Domain names - concepts and facilities. RFC 1034 (Internet Standard), November 1987. ISSN 2070-1721. URL https://www.rfc-editor.org/rfc/rfc1034.txt. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592, 5936, 8020, 8482, 8767.

[11] P.V. Mockapetris. Domain names - implementation and specification. RFC 1035 (Internet Standard), November 1987. ISSN 2070-1721. URL https://www.rfc-editor.org/rfc/rfc1035.txt. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 1995, 1996, 2065, 2136, 2181, 2137, 2308, 2535, 2673, 2845, 3425, 3658, 4033, 4034, 4035, 4343, 5936, 5966, 6604, 7766, 8482, 8490, 8767.

[12] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. Address Allocation for Private Internets. RFC 1918 (Best Current Practice), February 1996. ISSN 2070-1721. URL https://www.rfc-editor.org/rfc/rfc1918.txt. Updated by RFC 6761.

[13] P. Srisuresh and M. Holdrege. IP network address translator (NAT) terminology and considerations. Technical report, aug 1999. URL https://doi.org/10.17487%2Frfc2663.