Workgroup: Network Working Group
Internet-Draft:
draft-cparsk-eimpact-sustainability-considerations-07
Published: 24 January 2024
Intended Status: Informational
Expires: 27 July 2024

Authors: C. Pignataro, Ed. A. Rezaki S. Krishnan
NC State University Nokia Cisco
H. ElBakoury A. Clemm
Independent Consultant Futurewei

Sustainability Considerations for Internetworking

Abstract

This document defines a set of sustainability-related terms to be used while describing and evaluating environmental impacts of Internet technologies. It also describes several of the design tradeoffs for trying to optimize for sustainability along with the other common networking metrics such as performance and availability.

Embedding sustainability considerations at the design of new protocols and extensions is more effective than attempting to do so after-the-fact. Consequently, this document also gives network, protocol, and application designers and implementors sustainability-related advice and guideance. This document recommends to authors and reviewers the inclusion of a Sustainability Considerations section in IETF Internet-Drafts and RFCs.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Over the past decade, there has been increased awareness of the environmental impact produced by the widespread adoption of the Internet and internetworking technologies. The impact of Internet technologies has been overwhelmingly positive over the past years (e.g., providing alternatives to travel, enabling remote and hybrid work, enabling technology-based endangered species conservation, etc.), and there is still room for improvement.

This document describes some of the tradeoffs that could be involved while optimizing for sustainability in addition to or in lieu of traditional metrics such as performance or availability. It also proposes some common terminology for discussing environmental impacts of Internet technologies, and gives network and protocol designers and implementors sustainability-related advice and guideance. Further, it discusses how Internet technologies can be used to help other fields become more sustainable.

Specifically, this document is organized with the following outline:

*Section 2 includes a "Definition of Terms"

*Sections 3 through 7 detail sustainability and environmental impact considerations, and their implications to Internet protocols, architectures, and technologies.

*Section 8 lists "Sustainability Guidelines for Protocol and Network Designers and Implementers"

The ultimate objective of this document is to detail guidance regarding aspects of sustainability and environmental impact that authors and reviewers of Internet protocol and architecture documents should consider in a "Sustainability Considerations" section.

2. Definition of Terms

Given that the term 'considerations' is well known within the IETF community, it is fair to start by defining 'sustainability'. The 1983 UN Commission on Environment and Development had important influence on the current use of the term. The commission's 1987 report [UNGA42] defines it as development that "meets the needs of the present without compromising the ability of future generations to meet their own needs". This in turn involves balancing economic, social, and environmental factors.

This section defines sustainability-specific terms as they are used in the document, and as they pertain to environmental impacts. The goal is to provide a common sustainability considerations lexicon for network equipment vendors, operators, designers, and architects.

Notwithstanding the most comprehensive set of definitions of relevant terms readers can find at [IPCC], this section contributes the application and exemplification of the terminology to the internetworking domain and field. The terms are alphabetically organized.

Appropriate technology:

formerly referred to as 'intermediate technology', it refers to technology that is adapted to the local needs of its users, that is affordable, sustainable, and usually small scale and decentralized. Globally impactful technology is to be adaptable to local contexts it is used in. Regarding internetworking, there could be linkages to centralization / decentralization challenges, as well as maintainability & deployability aspects. Considering the diversity of local contexts, from developed countries with remote/rural coverage/ access issues, to developing countries with unstable electricity grids as well as literacy and technology usability/accessibility issues, internetworking technology needs to be designed, developed and operated according to these local requirements, also supporting small scale business models to make impact.

Biodiversity loss:

Biological diversity is a measure of the abundance and variety of life on earth. Biodiversity loss is the depletion of this diversity due to human activity, notably through the destruction of natural ecosystems and through the cascading effects of climate change, materials extraction, waste disposal and pollution, among other impacts, on the living world and species.

CO2e / CO2eq / CO2-eq:

Carbon dioxide equivalent, is the unit for measuring the climate change impact of non-CO2 gases as compared to CO2, which is selected as a benchmark.

Carbon awareness:

is being mindful of the carbon intensity of the electricity being used and prioritizing the use of low carbon intensity electricity in network set-up and operations. As carbon intensity is location and time dependent, carbon awareness requires dynamic monitoring and response, such as carbon aware routing and networking. This is a form of "demand shaping" which aims to match the use of energy with the supply of clean energy.

Carbon intensity (CI):

also referred to as emission intensity and emission factor, is a measure of the carbon-equivalent emission of consumed electricity, i.e., grams of carbon-equivalent per kilowatt hour (gCO2e/KWh). When the supplied energy mix is purely from renewable sources such as sun and wind, carbon intensity is practically 0, when coal and gas-powered electricity generation gets in the mix, carbon intensity increases. Carbon intensity could change instantaneously or predictably based on the time and location of electricity use. Prioritizing electricity use when carbon intensity is low is a target.

Carbon offset and credit: is a reduction of GHGs from the atmosphere as compensation for GHGs produced elsewhere and the credit generated and used respectively. This reduction in GHG emissions can be an increase in carbon storage through land restauration, or an actual removal of GHG. For example, certified forestation projects that absorb carbon dioxide produce carbon credits that an airline can use to offset its GHG emissions by using (purchasing) these credits. There are accredited carbon trading mechanisms to facilitate this exchange. This is generally regarded as a non-scalable solution, and activities such as the reduction of GHG emissions and the shifting of electrical energy production to renewables are a primary focus.

Circularity (circular economy): is a model or system where material resources and products are kept in use for as long as possible through long life cycles, reuse, repair, refurbishing and recycling, thereby reducing materials use, waste, and pollution as well as biodiversity and geodiversity loss. Keeping internetworking equipment in longer use through modularity, serviceability, upgradeability, maintainability are strategies to improve circularity.

Climate change (climate emergency, global warming): can be summarized as the increase in the global average temperatures and its destructive impact on life on Earth. The climate emergency refers to the ongoing and projected impacts of rising global temperatures and the narrow time window we have to limit temperature increases to a threshold determined by the Paris Climate Agreement (2015) to avoid the permanent destabilization of Earth life-support systems.

Climate change adaptation: are the measures we can take to adjust ourselves to the already happening and projected future adverse effects of climate change. This notably includes raising the resilience of internetworking solutions to higher operating temperatures and other impacts of climate change, as well as the use of internetworking technology to increase the resilience of societies and nature itself.

Climate change mitigation: encompasses all measures to reduce the impact of climate change. More specifically, any measures that reduce the amount of GHGs in the atmosphere can be considered as climate change mitigation through reduced inflow of GHGs into the atmosphere (such as burning of fossil fuels) or increasing the impact of carbon sinks such as forests and oceans. Reducing the carbon footprint of internetworking and increasing its carbon handprint by helping other sectors to decarbonize are mitigation efforts.

CUE: Carbon usage effectiveness [CUE] is a metric that helps determine the amount of greenhouse gas (GHG) emissions produced per unit of IT energy consumed within a data center. It provides an effective way to measure operational carbon footprint and thus the environmental impact of data center operations. The CUE is the ratio of the total CO2 emissions caused by total data center energy consumption, divided by the energy consumption of IT equipment. To calculate CUE when using electricity from the grid, carbon emissions can be based on published data. See also "PUE".

Doughnut economics: is a visual framework for sustainable development. It attempts to find a safe operational space within planetary boundaries and complementary (yet seemingly opposing) social boundaries, thereby meeting the needs of human societies without pushing earth environmental boundaries to their tipping points [Doughnut]. The significance of this model for interworking is that it demonstrates how to conceptualize and position boundaries in our designs that are seemingly opposing, to create a balanced approach, for example between energy efficiency and performance or resiliency and materials efficiency. It is not one or the other, but to find a space where both can be achieved without crossing boundaries in respective domains.

Embodied emissions: also referred to as embodied carbon and embedded carbon, refers to the amount of GHG emissions associated with upstream phases - raw material extraction, production, transportation (of materials and of product), and manufacturing- stages of a product's lifecycle. Some initiatives also consider disposal.

Energy, power, and their measurement: In physics, energy is defined as the capacity or ability to do work. For a system to provide an output, the quantitative property of energy is transferred to it. The energy measurement unit in the International System of Units (SI) is the joule (J). Power is energy used per second, measured in the International System of Units in watts (W), equivalent to the rate of one joule per second (J/s). In other words, energy is the integration of power over time. As such, Kilowatt-hour (kWh) is also a measure of energy, equivalent to 1 kW of power maintained for 1 hour, which is equal to 3.6 MJ (million joules).

Energy efficiency (EE):

increased energy efficiency can be summarized as doing the same task with less energy use, that is, providing a useful output/impact with as little energy as possible, eliminating energy waste. Switching to more efficient power supplies and silicon or developing more efficient transmission or signal processing algorithms improves EE. Developing energy efficiency metrics for internetworking and associated measurement methodologies and conditions as well as consistently collecting this data over time are essential to demonstrating EE improvements. An example of a common outcome-oriented metric is energy consumption per data volume or traffic unit, in Wh/B [Telefonica]; this particular metric has however also been criticized for being easy to misinterpret, falsely indicating that systems are energy proportional even when they are not (see "Energy proportionality".)

Energy equity:

Energy equity aims to minimize the negative impacts of energy systems and maximize the benefits for all energy users. Historically, these impacts and benefits haven't been equitably distributed. Energy equity recognizes that disadvantaged communities have been historically marginalized and overburdened by pollution, underinvestment in clean energy infrastructure, and lack of access to energy efficient housing and transportation.

Energy proportionality:

is the correlation between energy used and the associated useful output. For internetworking this is generally interpreted as the proportionality of traffic or traffic throughput and energy used. This concept is broadly applicable to networking infrastructure, data center, and other communication architectures. It is not a given that there is a one-to-one correlation between traffic and energy use, notably due to the materially significant idle power use by devices, as well as the overall network capacity being allocated to serve at times of highest traffic utilization.

Energy savings / conservation (ES):

is the avoidance of energy use, by eliminating a task altogether, when possible. Shutting down unused ports on a networking equipment is energy savings/ conservation.

Footprint (environmental/ecological):

in general terms is the impact we have on the planet. It can be divided into subcategories as carbon footprint, water footprint, land footprint, biodiversity footprint, etc. Related to the climate emergency, we are mostly focused on our carbon footprint, however, it has been shown that sub-categories of footprint are not entirely independent of each other. For example, our carbon footprint has a proven impact on the climate emergency through rising global temperatures, cascading significant impact on forest cover in warming areas since tree species adapted to certain climates vanish, thereby reducing biodiversity in that region, in-return impacting the carbon sink properties of the environment and exacerbating climate change. A holistic approach to our environmental footprint would therefore provide the best opportunity to create impact.

GHGs:

Greenhouse gases are types of gases that trap heat from the sun in earth's atmosphere, thereby increasing average global temperatures and creating the climate emergency. Carbon dioxide (CO2) is one of the most common (and referenced) greenhouse gases. There are others such as methane (CH4 - a much more potent GHG than CO2) and sulfur hexafluoride (SF6 - an artificial electrical insulator with tens of thousands of times more warming effect than CO2).

GHG Emissions Scopes:

According to the Greenhouse Gas (GHG) Protocol [GHG-Proto], Chapter 4, the emissions scopes are defined as below:

*Direct GHG emissions are emissions from sources that are owned or controlled by the company.

*Indirect GHG emissions are emissions that are a consequence of the activities of the company but occur at sources owned or controlled by another company.

The GHG protocol [GHG-Proto], Chapter 4, also includes the following descriptions of emissions scopes for accounting and reporting purposes:

*Scope 1 Emissions: Direct GHG emissions - Direct GHG emissions occur from sources that are owned or controlled by the company, for example, emissions from combustion in owned or controlled boilers, furnaces, vehicles, etc.; emissions from chemical production in owned or controlled process equipment.

*Scope 2 Emissions: Electricity indirect GHG emissions - Scope 2 accounts for GHG emissions from the generation of purchased electricity consumed by the company. Purchased electricity is defined as electricity that is purchased or otherwise brought into the organizational boundary of the company. Scope 2 emissions physically occur at the facility where electricity is generated.

*Companies shall separately account for and report on scopes 1 and 2 at a minimum.

*Scope 3 Emissions: Other indirect GHG emissions - Scope 3 is an optional reporting category that allows for the treatment of all other indirect emissions. Scope 3 emissions are a consequence of the activities of the company, but occur from sources not owned or controlled by the company. Some examples of scope 3 activities are extraction and production of purchased materials; transportation of purchased fuels; and use of sold products and services.

In telecommunications networks, Scope 3 emissions include the use phase of the sold products in operations, and is currently the largest part by far, of the whole GHG emissions (Scopes 1, 2 and 3), depending on the carbon intensity of the energy supply in use.

GWP:

Global warming potential, is the potential impact of GHGs on climate change, measured in CO2e.

Geodiversity:

is the variety of the nonliving parts of nature, that is, the materials constituting Earth, including soils, water (rivers, lakes, oceans), minerals, landforms and the associated processes that form and change them. The materials used in the production of internetworking equipment as well as their manufacturing and operational processes themselves, have impact (footprint) on geodiversity. Materials efficiency as well as circularity improvements help mitigate this impact.

Handprint (environmental/ecological):

is a concept developed in contrast to footprint, to quantify and demonstrate the positive environmental/ecological impact of technologies, products or organizations. Through a LCA (life cycle assessment) approach, the use of a technology or the products and services of an organization would have both a footprint and handprint usually denoted by the terms "X for sustainability" (handprint) and "Sustainable X" (footprint). What is important is that handprint impact does not compensate for footprint impact. They are to be calculated and reported independently; footprint to be minimized as much as possible, and handprint maximized as much as possible, which are by definition different activities anyway. Otherwise, this might be construed as "greenwashing". A popular seesaw figure in common sustainability literature depicting handprint and footprint sitting on opposite ends of a seesaw, one going up while the other is going down is a misguided representation.

LCA (Life Cycle Assessment):

is a comprehensive methodology to measure the environmental impact of a product, service, or process over its complete lifecycle, from the extraction and procurement of materials, through design, manufacturing, distribution, deployment, operations (use), maintenance/repair, decommissioning, refurbishment/reuse, recycling and disposal (waste), considering the full upstream and downstream supply chains as well. It is an extremely complicated process and there are multiple methods used worldwide, which might not produce same/similar results. LCA covers full footprint aspects, not only covering carbon, but also materials and biodiversity. Please refer to Section 6.6 for additional details on "Attributional and Consequential Models".

Materials efficiency and reuse:

is the concept of using less primary and (more) recycled materials to provide the same output. A networking equipment that provides the same function with less aluminium used is more materials efficient. Reuse of materials in manufacturing, thereby reducing primary materials extraction is a cornerstone of circularity, reducing environmental footprint and promoting geodiversity.

Net-zero:

in general, is to bring down GHGs as close to zero as possible. It is generally recognized that it may not be possible to get GHGs to 0 in many contexts and the balance is said to be covered by carbon offset. For example, many organizations and countries have net-zero targets by certain dates and typically what they mean is that they will reduce their GHGs by more than 90% and the remaining up to 10% will be offset.

PUE:

Power usage effectiveness, is a data centre energy efficiency metric. The PUE is defined by dividing the total amount of power entering a data center by the power used solely to run the IT equipment within it. PUE is expressed as a ratio, with the overall power usage effectiveneess improving as the quotient decreases towards one. See also "CUE".

Planetary boundaries:

is a concept that defines 9 environmental boundaries that, if not crossed, provides a safe space for humanity to live. This was developed and tracked by the Stockholm Resilience Centre [Planet-B]. Their latest report indicates that 6 out of the 9 boundaries have already been crossed. This translates to the increased risk of irreversible environmental change, the so-called tipping points. Climate change is one of these boundaries, represented as carbon dioxide concentration in the atmosphere (ppm by volume) and others are biodiversity loss, land use, fresh water, ocean acidification, chemical pollution, ozone depletion (one boundary that has been successfully mitigated), atmospheric aerosols and biogeochemical (nitrogen in the atmosphere and phosphorus in oceans).

Rebound effect:

is the reduction in the potential benefits of more efficient technologies and solutions to reduce resource use, due to the increased demand they might trigger as costs might decrease, in return even increasing the overall resource use. This is known as Jevons paradox: efficiency leading to increased demand. In internetworking, this can manifest itself when more energy and resource efficient systems reduce the cost for infrastructure build and operations and when this is reflected to customers as reduced cost, customers respond by increased use of telecommunications services which pushes infrastructure build and operations upwards, thereby negating the projected gains from efficiency measures. Another descriptive source for this phenomenon can be found at [Frontiers].

Tipping points:

are critical environmental thresholds, which when crossed likely lead to irreversible state changes in climate systems that might push the overall earth system out of its stable state that supports life on Earth. For example, there are tipping points defined for the Antarctic and Greenland ice sheets disappearing, the Arctic sea-ice loss, Siberian permafrost loss or the dieback of the Amazon and Boreal forests. As planetary boundaries are crossed, the likelihood of the tipping points being reached also increases. When the tipping points are hit, notably simultaneously, the overall impact to the global Earth system might be catastrophic, as another stable state which no longer supports life could be reached.

UN SDGs:

United Nations Sustainable Development Goals are 17 global objectives that collectively define a framework for a sustainable global system where people and the planet collectively thrive and live in peace, prosperity and equity. They were adopted in 2015 and most of them have a target achievement date of 2030 [UN-SDG]. They are part of the so-called UN 2030 Agenda. The International Telecommunications Union (ITU) has published on how our technology could help meet the UN SDGs [ITU-ICT-SDG]. Notably, most UN SDGs provide guidance for the handprint impact of internetworking technologies, while some are also related to potential action for footprint reduction. The 17 SDGs are:

Goal 1	No poverty
Goal 2	Zero hunger
Goal 3	Good health and well-being
Goal 4	Quality education
Goal 5	Gender equality
Goal 6	Clean water and sanitation
Goal 7	Affordable and clean energy
Goal 8	Decent work and economic growth
Goal 9	Industry, innovation and infrastructure
Goal 10	Reduced inequalities
Goal 11	Sustainable cities and communities
Goal 12	Responsible consumption and production
Goal 13	Climate action
Goal 14	Life below water
Goal 15	Life on land
Goal 16	Peace, justice and strong institutions
Goal 17	Partnerships for the Goals

The SDG Academy [SDG-Acad] also provides useful information on the topic, as well as progress to date.

3. 'Sustainable X' versus 'X for Sustainability'

Every technology solution, system or process has sustainability impacts, as it uses energy and resources and operates in a given context to provide a [perceived] useful output. These impacts could be both negative and positive w.r.t sustainability outcomes. With a simplistic view, the negative impact is termed as footprint and the positive impact is handprint, as defined in the "Definition of Terms" section. Again, generally speaking, footprint considerations of a technology are grouped under "Sustainable X" and the handprint considerations are covered under "X for Sustainability".

Additionally, when sustainability impacts are considered, not only environmental but also societal and economic perspectives need to be taken into account, both for footprint and handprint domains. A systems perspective ensures that the interactions and feedback loops are not forgotten among different sub-areas of sustainability.

Another fundamental sustainability impact assessment requirement is to cover the complete impact of a product, service or process over its full lifetime. Life Cycle Assessment (LCA) starts from the raw materials extraction & acquisition phases, and continues with design, manufacturing, distribution, deployment, use, maintenance, decommissioning, refurbishment/reuse, and ends with end-of-life treatment (recycling & waste). It is imperative that we consider not only the design and build stages of our technologies but also its use and end-of-life phases. An equally essential way of ensuring a holistic perspective is the supply-chain dimension. When we consider the footprint impact of a technology we are building, we need to consider the full supply chain that the technology is part of, both upstream, what it inherits from the material acquisition, components and services used, to downstream for wherever the technology is used and then decommissioned. Further, this includes transportation of materials or products, and the carbon-friendliness of the means and routes chosen. What this implies is that we are responsible for the direct and indirect impacts of our activity, both on demand and supply directions.

Below, we cover the "Sustainable Internetworking" and "Internetworking for Sustainability" perspectives in more detail.

3.1. Sustainable Internetworking

Sustainable internetworking is about ensuring that the negative impacts of internetworking are minimized as much as possible.

In the environmental / ecological sustainability domain, the sub-areas to be considered are:

- *Climate change,
- *materials efficiency, circularity, preservation of geodiversity, and
- *biodiversity preservation.

Climate change considerations in internetworking by and large translate to energy sourcing, consumption, savings and efficiency as this impacts the GHGs of the internetworking systems directly, when mostly non-renewable energy sources are used for the operations of the networks. When the carbon intensity of the energy supply used in operations decreases (more renewable energy in the supply mix), then the use phase GHGs also proportionally decrease. This might put the GHG emissions of the manufacturing and materials extraction and acquisition phases ahead of the use phase. These are called the embodied emissions.

However, energy is not the only aspect to consider: materials efficiency and circularity are key considerations to limit the resource use of our technologies, thereby reducing the scarcity of materials but also the destruction of many ecosystems during their extraction and manufacturing, polluting water and land with waste, which might also impact directly or indirectly the abundance and health of the species on the planet, namely biodiversity. While it is significantly more difficult to quantify and measure the impact of our technologies in these domains, the planetary boundaries framework provides helpful guidance.

For the societal and economic footprint of our technologies, we need to be mindful about the potential negative effects of our technologies w.r.t. the social boundaries, as depicted in the so-called doughnut economics model, that includes education, health, incomes, housing, gender equality, social equity, inclusiveness, justice and more. What we need to realize is that our technology has direct and indirect impacts in these aspects and the challenge is not only to meet environmental sustainability targets but social and economic ones as well. There are very practical considerations, for example: are there partial or total barriers to accessing the Internet or its services? what is the impact of biases in artificial intelligence (AI), as it pertains gender biases, when those AI models are used in job selection? More technology doesn't always mean better outcomes for all and can we mitigate this impact? Admittedly, a quantitative approach to the societal and economical aspects is more challenging; thinking in terms of profit, people, and planet, as well as the Key Values (KV) / Key Value Indicators (KVIs) approach described in Section 4 bring some relief.

3.2. Internetworking for Sustainability

When it comes to the positive impact of internetworking in tackling the sustainability challenges faced, we are in the "internetworking for sustainability" realm. This is a very diverse topic covering innumerable industrial and societal verticals and use cases. Essentially, we are asking how our technology can help other sectors and users to decarbonize, and to reduce their own footprints and to increase their handprints in environmental, societal and economic dimensions. These are induced or enablement effects. Examples are how internetworking is being used in smart energy grids or smart cities, transport, health care, education, agriculture, manufacturing and other verticals. While efficiency gains are usually a basis, there are also other impacts through ubiquitous network coverage, sensing, affordability, ease of maintenance and operation, equity in access, to name a few.

Climate change mitigation and climate change adaptation, as defined in the "Definition of Terms" section, are particular focus areas where internetworking could help create more resilience in our societies and economies along with sustainability.

Essentially, handprint considerations are asking us to think about how our technology could be used to tackle sustainability challenges at first, and second, to generate feedback on how to create enablers and improvements in our technology for it to be more impactful. The usual Key Performance Indicators (KPIs) related to technical system parameters would be largely insufficient for this purpose. Supporting this effort, the Key Values (KV) and Key Value Indicators (KVIs) concepts have been developed, to be used in conjunction with use cases to develop impactful solutions. KV and KVIs are the subject of Section 4.

The following are some examples of internetworking for sustainability. This is not a comprehensive list; many more such examples can be found. Leveraging internetworking for sustainability usually involves special requirements, which are listed along with the examples.

Smart Grid:

The Smart Grid [RFC6272] generally refers to enhancements to traditional electrical grids that offer additional features such as two-way flows of electricity (e.g., accommodating solar panels, electrical batteries) and granular control of the grid (e.g., allowing to selectively turn off certain consumers such as Heating, Ventilation, and Air Conditioning (HVAC) units during certain times.) The Smart Grid aims to improve sustainability by facilitating concepts such as peak shaving (i.e., lowering peak usage to reduce the amount of excess generation of electricity that is not needed during non-peak periods), and encouraging residential homes and business to invest in renewable energy sources such as solar, for example offering credit for feeding surplus energy being generated back into the grid. For this to work, the Smart Grid requires support by networking technology that enables the required control loops as well as visibility into grid telemetry. This, in turn, requires the support of new requirements, including aspects of security (since a critical infrastructure is at stake), adherence to high precision service levels and ultra-low latency communication (e.g., to mitigate sudden spikes in voltage), and special provisions to ensure data privacy (given that data from private households, electrical vehicles, and personal devices is involved.)

Smart Cities:

Many applications for smart cities involve optimizations to make cities more sustainable. Examples include smart garbage disposals that reduce the number of truck rolls (and associated emissions) to collect garbage only when needed, and guidance systems for smart parking that reduce the amount of vehicle traffic used to find parking spots. These applications are enabled by networking. Again, special requirements need to be supported for networks to support those applications, such as the ability to deploy equipment in harsh urban environments, or monitoring for vandalism.

Smart Agriculture:

Smart agriculture involves minimizing usage of resources such as fertilizer and water in the production of agricultural output. This also helps minimize the area set aside for farming and reclaim land for other purposes including biodiversity. Similarly, networking is an enabler for environmental sustainability. Special requirements for applications in this space include aspects such as the ability to support networking equipment without the need to run power lines (e.g., using battery or solar), and support for intermittent communications.

4. Key Values and Key Value Indicators

In the context of sustainability, key values are what matters to societies and to people when it comes to direct and indirect outcomes of the use of our technology. While KPIs help us to build, monitor and improve the design and implementation of our technologies, key values and their qualitative and quantitative indicators tell us about their usefulness and value to society and people. As we want our technology to help tackle the grand challenges of our planet, their likelihood of usefulness and impact is a paramount consideration. KVs and KVIs help set our bearings right and also demonstrate the impact we could create. The main idea is shifting from measuring performance to measuring value.

While key values could be universal, like for example the United Nations Sustainable Development Goals (UN SDGs) [UN-SDG], how they are measured, or perceived (KVIs) could be context dependent and use case specific. To give a simplified example, UN SDG 3, "good health and well-being" is a key value for any society and individual. Then, when we consider the use case of providing health care and wellness services in a remote, rural community which doesn't have any hospitals or specialist doctors, a key value indicator could be how fast a patient could access health care services without having to travel out of town, or the successful medical interventions that could be carried out remotely. Then the next step is to identify which parts of our technology could help enable this and design our technology to create impact for the KVs as per KVIs. In this case, universal network coverage, capacity and features to integrate a multitude of sensors, low-latency and jitter communication services could all be enablers with their own design targets and KPIs defined. Subsequently, we would track the KVIs and the KPIs together for successful outcomes.

Admittedly, this might not be a straightforward task to carry out for each protocol design. Yet, such analyses could be included in design processes along with use case development, covering a group of technology design activities (protocols) together. There are ongoing efforts in mobile networking research to use KVs/KVIs efficiently [M6G-SOCIETAL-KV-KVI] [M6G-VALUE-PERF] [Hexa-X_D1.2].

While we find ourselves trying to optimize seemingly contradicting parameters or aspects such as reducing latency and jitter and increasing bandwidth and reach targets with sustainability parameters or aspects such as reduced energy consumption and increased energy efficiency, key values and key value indicators would help keep our eyes on the targets that matter for the end users and communities and societies. Considerations for such potential design trade-offs, which are at the heart of our engineering innovations, are the topic of the next section.

4.1. Key Value Enablers

Between the design and creation of a technology, and realization of the value generated by its deployment and use, there are a number of enablers and blockers of its usage. We generally refer to them as KV Enablers. These are the key factors that would scale and spread use cases or block their deployment.

Technical enablers are the features needed for the technical capabilities and feasibility of the use cases, like the network features being deployed to support the use case. Beyond the technical aspects, there are also criteria at the system level which determine the context in which the technology will be used as well as the actions of the use case stakeholders. These might affect the level of adaptation to a particular society or ecosystem, such as cost of connectivity and Internet service access, availability of services, security, and privacy. While technical enablers are in more direct control of protocol and network designers, system-level enablers might in second-order, indirect, or beyond control, depending on the actions of other stakeholders and the existing environment.

An important corollary is that KV enablers can be used to derive technological requirements, KPIs and advancements to maximize key value.

5. Implications to the IETF

This section describes the implications of sustainability to the IETF. Specifically, the high-level relevant areas on which the IETF can act upon, and a rough prioritization. These potentially include use cases, protocols, metrics, etc.

A key area to understand the relevance and implication is regarding IETF Protocols.

6. Sustainability Considerations - How Will the Natural Environment be Impacted?

6.1. Design Tradeoffs

Traditionally, digital communication networks are optimized for a specific set of criteria that proxies for business metrics. A network operator providing services to their customers intends to maximize profits, by increasing top-line revenue and decreasing bottom-line associated costs. This directly translates to goals of optimizing performance and availability, while reducing various costs.

Most recently, various forces elevate the need for sustainability in networking technologies and architectures, to quantify and minimize negative environmental impact.

Optimizing only network availability (e.g., by having excess capacity and backup paths) or optimizing only performance (e.g., by increasing speeds selecting paths based on delays only) can seemengly be in opposition to optimizing sustainability objectives. For a given application, use-case, or vertical realization of technology, a Pareto-efficient choice can potentially improve sustainability without sacrificing availability or performance beyond the application tolerance. That is, a win-win.

Consequently, network architects and designers are presented with a set of new design tradeoffs: a multi-objective optimization that satisfies border requirements and global optima for availability, performance, and sustainability simultaneously. This is not unlike the doughnut economics model concept introduced in the "Definition of Terms" section.

6.2. Multi-Objective Optimization

To understand this new model, we can analyze a simplified example. Assume the following topology, passing traffic from A to B:

```
                         A
                         |
                    +----------+
                    | Router 1 |------------+
                    +----------+            |
                     | | | | |        +----------+
                     | | | | |        | Router 3 |
                     | | | | |        +----------+
                    +----------+            |
                    | Router 2 |------------+
                    +----------+
                         |
                         B
```

Figure 1: Simplified Network for Multi-Objective Optimization

Router 1 is directly connected to Router 2 through five parallel links, of 10 Gbps each. Router 1 can also reach Router 2 through Router 3 with 40 Gbps links between Router 1 and Router 3, and between Router 3 and Router 2. Let's assume that the capacity-planned traffic between A and B equals 15 Gbps.

In this scenario, a topology optimized for performance and availability/resiliency would have all links and routers on, and would likely forward traffic using two of the parallel links. Utilizing the path through Router 3 might lower performance, but it serves as a backup path.

On the other hand, when we add sustainability as a consideration, different options are presented. One of them is to remove from the topology Router 3 and associated links, and shutdown links and optics in two or three of the parallel links. Another option is to completely shutdown all the parallel links and route traffic through Router 3 (i.e., not maximizing performance alone, but maximizing at the time performance, availability and resiliency, and sustainability.) The choice between these two options will depend on the aggregate sustainability metrics of network elements in each of the two topologies.

Another option is to use flexible Ethernet, where the five links combined are aggregated into one active virtual link which has 15 Gbps, and another inactive link of 35 Gbps of capacity -- although a physical link cannot be half-active and half-inactive as far as PHY and optics are concerned.

6.3. How Much Resiliency is Really Needed?

When we add sustainability considerations, resiliency is not the single objective to optimize.

There are many methods to improve network resiliency, including a design eliminating single-points-of-failure, performing software safe-release selections and upgrades, deploying network real-time testing systems (including operations, administration, and maintenance (OAM) tools, monitoring systems (e.g., [RFC8403]), chaos-based testing, and site reliability engineering (SRE) principles), and utilizing redundancy across network elements as well as across a topology. Each one of these methods incurs also a sustainability cost. Yet, the functions for resiliency improvement and sustainability cost for each of these methods are not the same. Considering sustainability means quantifying its impact in the decision of how to improve resiliency, and how much is needed.

6.3.1. Redundancy and Sustainability

Let's first explore redundancy. For example, consider the ratio of overall network capacity (in bandwidth, compute power, etc.) over the used network capacity, and let's call it "Redundancy Index". If this number is one, there's no redundancy; and as the ratio grows, so does the potentially unused capacity that could be utilized in a failure event. Similarly, consider the values of sustainability metrics for when the Redundancy Index is one and for when it is two. These border points might give an indication of the slope for each objective function.

Adequate Redundancy:
In order to determine how much redundancy needs to be built into the overall network capacity, which can be referred to as "adequate redundancy to avoid network outings", it will be important to (1) measure the bandwidth of attacks against the overall network capacity; and (2) understand how quickly "high bandwidth" attacks can be detected and diverted. Measuring these results over time may lead the "adequate redundancy" to become higher over time.

Justified Redundancy:
Traditionally, network operators would be much less worried about energy use than about the possibility that the network would have brownout or backout outages - thus the measuring will help better balance the "adequate redundancy" against the related energy use, resulting in turn in "justified redundancy": a balance between costs and benefits, with energy use as well as material use as a clear cost factor.

Please note that "justified redundancy" may be higher than "adequate redundancy" when we manage to organize the redundancy in a multi-layer fashion: (1) capacity that is "always on" and always uses energy; (2) capacity that can turn on quickly when needed; (and possibly (3) capacity that is "on the shelf" (even in the box) but ready to be deployed quickly when needed.)

6.4. How Much are Performance and Quality of Experience Compromised?

The fields of performance and quality of experience have the benefit of significant study and standardization of metrics. As with resiliency, a degradation of performance and Quality of Service parameters, such as bandwidth, latency, jitter, etc., can be observed and measured, as a variation of sustainability metrics. The relative slopes of improvement of each goal would hint as to where the balance lies.

6.5. End-to-End Sustainability

The networking industry is in the starting phases of addressing this objective. We are seeing a sprinkling of sustainability features across the networking stack and components of devices, whether it is on forwarding chips, power supplies, optics, and compute. Many of those optimizations and features are typically local in nature, and widely scattered across different elements of a network architecture. An opportunity for maximizing the positive environmental impact of these technologies calls for a more cohesive and complementary view that spans the complete product lifecycle for hardware and software, as well as how some of these features work in unison.

For example, features that provide energy saving modes for devices can be dynamically managed when the network utilization is such that performance would not significantly suffer. A core router, instead of becoming obsolete due to the need for higher throughput in the core, could become a future edge/access router. That is an example of reuse and repurpose, before recycling. There are benefits of macro-optimizations by clustering in specific features, versus micro-optimizing locally without awareness of the network context.

6.6. Attributional and Consequential Models

Many of the subtleties and nuances of the measurement of GHG and environmental impacts stem from the very important distinction between attributional and consequential models. Detailed definitions can be found at [UNEP-LCA].

Attributional:
Also referred to as Allocational models, start by allocating or attributing quantities (e.g., GHG emissions) to entities (e.g., a router, a building, a town), and performing

comparisons between the measurements (or estimates) of the quantity by the entities.

Consequential:
Perform the measurement of the quantity by establishing a baseline scenario (e.g., before feature introduction) and a modified scenario (e.g., after the feature introduction).

While both models are quite different, they do use the same terms and frames of references, measures, and language. Without explicit clarifications, they are prone to confusion.

For example, measuring the carbon footprint attributed to a batch process or a workload based on its energy efficiency would not consider that the hardware is still there running. When is it most effective to charge battery-powered devices, during the night when there's less load, or during the day when there's solar energy? In other words, if a person who commutes by train to their office five days a week starts working from home two days a week, there could be an attributional reduction of GHG emissions, yet the train still continues running equally. However, if that person commutes by combustion-engine car alone, the consequences are different.

Considering the attributional versus consequential distinction, there are some implications and a potential corollaries:

- For an environmental-impact analysis, it is critical to explicitly cite the model used, as well as clearly define the boundary.

- The activities that we embark upon as internetworking and protocol designers - including the ones targeting reduction of negative environmental impacts - have an energy footprint of themselves.

- "Do no harm" in the context of improving sustainability of networks is to look beyond bounded attributions and consider (both intended and unintended) consequences.

6.7. The Role of Network Management and Orchestration

Deployment and operational aspects play a critical role in making networks more sustainable. A detailed explanation of that role, the associated challenges, as well as an outline of solution approaches is provided in [I-D.irtf-nmrg-green-ps]. Here are some areas in which network management can help make networks more sustainable; for a more extensive treatment, please refer to that document.

Dimensioning:
Networks should be deployed and configured with sufficient capacity to serve their intended purpose. At the same

time, overprovisioning and providing too many resources should be avoided, as this results in waste and unnecessary environmental impact. Network management can facilitate proper dimensioning of networks by providing the methods and tools that allow to assess network usage, determine required capacities, identify trends to allow to proactively accommodate traffic growth and new services.

Network Optimization:
Network management applications can help solve difficult network optimization problems involving multiple parameters, multiple and sometimes conflicting objectives, and mitigation of tradeoffs. Network optimization examples include maximization of utilization or of aggregate QoE scores, minimization of the possibility of SLA violations with a given amount of network resources, or optimization of the cost of configured paths. Network metrics related to sustainability are another set of parameters that can be optimized.

Rapid Discovery and Provisioning Schemes:
One of the biggest potential opportunities in reducing environmental impact of networks concerns the ability to power resources such as equipment or line cards down when they are momentarily not needed due to swings in traffic demands. In general, this involves fully automated management control loops with very short time scales. Network management can enable such schemes, involving algorithms that determine and control the rapid de- and re-commissioning of networking resources, as well as the necessary control protocols that facilitate aspects such as rapid resource discovery, reprovisioning, or reconvergence of management state.

6.7.1. Metrics for Sustainability

A sustainability quantification framework is paramount for understanding the sustainability posture of a system, as well as its potential for aid in sustainability outcomes.

7. Sustainability Requirements and Phases

The considerations and advice for sustainability described in the "Sustainability Considerations - How Will the Natural Environment be Impacted?" and "Sustainability Guidelines for Protocol and Network Designers and Implementers" sections and their associated goals cannot always be achieved at the same time and we expect the following high level phases:

1. Visibility: In this phase we focus on the measurement and collection of metrics.

2. Insights and Recommendations: In this phase we focus on deriving insights and providing recommendations that can be acted upon manually over large time scales.

3. Self-Optimization via Automation: In this phase we build awareness into the systems to automatically recognize opportunities for improvement and implement them.

7.1. Phase 1: Visibility

Visibility represents collecting and organizing data in a standard vendor agnostic manner. The first step in improving our environmental impact is to actually measure it in a clear and consistent manner. The IETF, IRTF and the IAB have a long history of work in this field, and this has greatly helped with the instrumentation of network equipment in collecting metrics for network management, performance, and troubleshooting. On the environmental-impact side though, there has been a proliferation of a wide variety of vendor extensions based on these standards. Without a common definition of metrics across the industry and widespread adoption we will be left with ill-defined, potentially redundant, proprietary, or even contradicting metrics. Similarly, we also need to work on standard telemetry for collecting these metrics so that interoperability can be achieved in multi-vendor networks.

7.2. Phase 2: Insights and Recommendations

Once the metrics have been collected, categorized, and aggregated in a common format, it would be straightforward to visualize these metrics and allow consumers to draw insights into their GHG and energy impact. The visualizations could take the form of high-level dashboards that provide aggregate metrics and potentially some form of maturity continuum. We think this can be accomplished using reference implementations of the standards developed in "Phase 1: Visibility". We do expect vendors and other open projects to customize this and incorporate specific features. This will allow identifying sources of environmental impact and address any potential issues through operational changes, creation of best-practices, and changes towards a greener, more environmentally friendly equipment, software, platforms, applications, and protocols.

7.3. Phase 3: Self-optimization and Automation

Manually making changes as mentioned in "Phase 2: Insights and Recommendations" works for changes needed on large timescales but does not scale to improvements on smaller scales (i.e., it is impractical in many levels for an operator to be looking at a dashboard monitoring usage and making changes in real-time 24x7). There is a need to provision some amount of self-awareness into the network itself, at various layers, so that it can recognize opportunities for improvement and make those changes and measure the effects by closing the loop. The goals of the consumers can be stated in a declarative fashion, and the networks can continually use mechanisms such as machine learning (ML), deep learning (DL), and artificial intelligence (AI) with an additional goal to optimize for improvements in the environmental impact. These include, for example:

- Discovery and advertisement of networking characteristics that have either direct or indirect environmental impact,
- greener networking protocols that can move traffic onto more energy efficient paths, directing topological graphs to optimize environmental impacts, and
- protocols that can instruct equipment to move under-utilized links and devices into low-energy modes.

7.3.1. Cycle of Phases

The three phases run in an iterative fashion, such that after phases 1, 2, and 3 are completed for an interation, there will be an added awareness of what (else) to collect back to phase 1.

Further, sustainability-aware self-optimization is something to explore in Autonomic Networking.

8. Sustainability Guidelines for Protocol and Network Designers and Implementers

This section renders the Sustainability Considerations into specific guidelines and advice for the design and development of networking technologies.

These specific items are labeled so as to follow and reference as a check-list.

a. General:
The section title "Sustainability Considerations" should be used to detail the environmental-impact implications of protocols, architectures, and Internet technologies.

a.1.

For each of the items covered, explicitly state the "boundary of analysis" considered. For example, this can include a scope, time boundary, or lifecycle phases.

a.2. Consider attributional versus consequential analysis methods, explaining environmental impact benefits.

a.3. Clearly state the units used for each magnitude in every analysis (e.g., gCO2e/KWh.)

b. Network Management:

Several areas of network management have direct relationship with sustainability.

b.1. Metrics:

Instrument equipment, network elements, and networks with a set of relevant and meaningful metrics that provide visibility into sustainability and environmental-impact attributes (e.g., power and energy consumption.) This is the foundation for any mechanisms to improve and optimize sustainability.

b.2. Managed Elements:

Facilitate, extend, and enrich the manageability of network elements and sub-elements which have environmental impact, such as Power Supplies. For example, provide visibility into sourced power, e.g. energy mix, and allow to account for the "dirtiness" of power being consumed to obtain a truer picture of sustainability than can be gained by visibility into power consumption alone.

c. Energy Management:

Minimizing energy consumption is a critical consideration in making networks more sustainable. Minimizing energy consumption typically comes also with important economic side benefits associated with reducing operational expenses and making network providers more competitive.

To facilitate energy efficiency schemes, designers of networking devices and protocols should examine and consider the following considerations:

c.1. Energy linearity. In many cases, the amount of power drawn by a device is not in linear proportion to the volume of traffic that is passed. Instead, energy consumption when idle generally accounts for a very significant percentage of the energy consumption when under full load. The implication of this is that the volume of traffic by itself is of relative consequence to energy consumption, as long as the volume does not get to the point where additional equipment needs to be added to the network to handle peak loads.

c.2. Power saving modes. Similarly, many devices and resources support power saving modes that can be entered when idled. Similarly, during periods of exceedingly low traffic, some links may support downspeeding associated with energy savings. As a result, a big opportunity for energy savings involves schemes in which resources are temporarily put into power saving modes, including almost shut-down, at times when they are not needed.

c.3. Chattiness of protocols. For a given protocol, what are the message exchange patterns? does the protocol rely on periodic updates or heartbeat messages? Could such message patterns result in preventing links or nodes from going to sleep (absent other communications), and in such a case, would an alternative pattern be feasible?

c.4. Exploiting burstiness versus smoothening of traffic. Is it feasible to design the protocol in such a way that traffic is sent with a smoother traffic pattern with lower traffic volumes that are sent continuously, as opposed to a way that traffic is bulked up and then sent in one fell swoop?

c.5. Rapid discovery and convergence. Does the protocol involve the exchange of state and information about other systems? In that case, how can the protocol be designed in such that any such information can be discovered quickly and protocol synchronization reconverged efficiently? Does the protocol design support stateful schemes that might accelerate this? In cases where there is a possibility of nodes going to sleep, the associated overhead of going offline and coming back online should be minimized. By shortening the time interval needed to go offline and come back online, it might be possible to have enter sleep mode in situations where it would otherwise not be feasible.

c.6. Encoding schemes. How much computational effort goes into encoding and decoding? Assess the tradeoff between encoding efficiency and computational effort, which directs into carbon for cycles to perform coding operations.

d. Carbon Awareness:
See the definition in Section 2.

  - **d.1.** Consider Carbon Intensity (CI) / Emission Factor (EF) as an attribute. For example, CI is used to optimize for lower-carbon sources of electrical energy (e.g., using renewables.) Prioritizing electricity use when carbon intensity is low is a target, and, for that, this attribute needs to be accessed or advertised.

  - **d.2.** Consider embodied emissions (i.e., embedded carbon) with any new product. For example, a new generation of networking device might significantly improve energy efficiency, and a replacement migration would include the embedded emissions (of producing and transporting the new product as well as disposing of the old one), and hence there's a break-even point (BEP).

e. Beyond Carbon:
Characterize and note full-spectrum environmental impacts, beyond GHG emissions, and into water usage, raw materials usage, circularity in supply chain, repurpose, reuse, and recycle, etc.

  - **e.1.** WIP

  - **e.2.** WIP

9. Conclusion

The pre-eminent message in this document is to elevate the need and sense of urgency of including sustainability considerations in our protocol and system design, and to provide editors with a sustainability lexicon, definitions, and priorities to carry out that task. As an added benefit, by including sustainability considerations, it will be possible to optimize for not only performance parameters but also sustainability ones, through respective trade-offs in our protocols and systems.

We also envision that on top of minimizing the environmental impact of our technologies and helping consumers identify and reduce the environmental impact of their use, we can also make a positive impact on other systems. E.g., use our technologies to choose greener and more efficient sources of power, control HVAC systems efficiently, etc.

9.1. Call to Action

The intention of this document is multifaceted: establish definitions and a lexicon for sustainability, characterize environmental implications of internetworking technologies, and provide specific guidelines for designers and implementors.

Making these objectives actionable involves:

1. Familiarize yourself with the terms defined in Section 2,

2. understand the sustainability considerations (Section 3 through Section 7) and their implications to protocol and architecture, and

3. consider, qualify, quantify, and explain the specific guidelines in Section 8 as you develop protocols, extensions, and architectures.

10. Security Considerations

Sustainable practices offer many environmental, economic, and social benefits, and security is a route to sustainability rather than a hurdle to clear.

> The creation of sustainability features for an element or a system should not weaken or compromise their security posture, nor should it increase the surface of attack or create attack vectors.
>
> - Sustainability metrics and data models ought to describe how to secure the sustainability data exposed and surfaced through telemetry.
>
> - Sustainability control capabilities, as for example for power management, should consider potential attacks leveraging those controls. Setting a device on low-power or power-save modes during peak traffic can be a denial-of-service attack vector, negatively impacting end-to-end services.
>
> The development of security features should, in turn, balance the environmental impact and sustainability considerations detailed in this document.
>
> - Computational increase on cryptographic operations can result in higher power use. Since generally the increase of energy required is not linear with the increase of computational complexity, there's a desire to satisfy security requirements while minimizing environmental impact.

- Proof-of-Work schemes' and AI models' energy consumption also grows non-linearly as a function of the precision achieved. In these, perfect is the enemy of good, and bounding precision through specifications supports sustainable compute considerations.

11. Acknowledgements and Contributions

The subject of sustainability considerations for internetworking sits in the intersection of several disciplines, benefiting from the collaboration of diverse educational, experiential, and exposure backgrounds. The authors are not only grateful but also inspired by the open collaboration and expertise of many individuals, including:

* Maarten Botterman provided text on network redundancy, and definitions for justified redundancy balancing adequate redundancy.

* Dom Robinson shared ideas and text on attributional and consequential methods, in turn inspired by a post from Chris Adams.

* Michael Welzl provided a very comprehensive and critical review of the complete document, and highlighted several fixes and improvements.

12. Informative References

[CUE] Belady, C., Azevedo, D., Patterson, M., Pouchet, J., and R. Tipley, "Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric", 2 December 2010, <https://www.thegreengrid.org/en/resources/library-and-tools/241-Carbon-Usage-Effectiveness-%28CUE%29%3A-A-Green-Grid-Data-Center-Sustainability-Metric>.

[Doughnut] Wikipedia, "Doughnut (economic model)", 13 October 2023, <https://en.wikipedia.org/wiki/Doughnut_(economic_model)>.

[Frontiers] Frontiers, "The Rebound Effect and the Jevons' Paradox: Beyond the Conventional Wisdom", 2023, <https://www.frontiersin.org/research-topics/6598/the-rebound-effect-and-the-jevons-paradox-beyond-the-conventional-wisdom#overview>.

[GHG-Proto] "The Greenhouse Gas Protocol - A Corporate Accounting and Reporting Standard, revised version", Geneva: World Business Council for Sustainable Development, Washington, DC: World Resources Institute, 2004, <https://ghgprotocol.org/sites/default/files/standards/ghg-protocol-revised.pdf>.

[Hexa-X_D1.2]
"Expanded 6G vision, use cases and societal values - including aspects of sustainability, security and spectrum", Deliverable D1.2, 30 April 2021, <https://hexa-x.eu/wp-content/uploads/2021/05/Hexa-X_D1.2.pdf>.

[I-D.irtf-nmrg-green-ps]
Clemm, A., Westphal, C., Tantsura, J., Ciavaglia, L., Pignataro, C., and M. Odini, "Challenges and Opportunities in Management for Green Networking", Work in Progress, Internet-Draft, draft-irtf-nmrg-green-ps-02, 30 December 2023, <https://datatracker.ietf.org/doc/html/draft-irtf-nmrg-green-ps-02>.

[IPCC] "Intergovernmental Panel on Climate Change (IPCC). (2022). Annex I: Glossary. In Global Warming of 1.5 degrees Celsius: IPCC Special Report on Impacts of Global Warming of 1.5 degrees Celsius above Pre-industrial Levels in Context of Strengthening Response to Climate Change, Sustainable Development, and Efforts to Eradicate Poverty (pp. 541-562)", Cambridge University Press, doi 10.1017/9781009157940.008, 2022, <https://www.ipcc.ch/sr15/chapter/glossary/>.

[ITU-ICT-SDG] ITU, "Digital technologies to achieve the UN SDGs", 2023, <https://www.itu.int/en/mediacentre/backgrounders/Pages/icts-to-achieve-the-united-nations-sustainable-development-goals.aspx>.

[M6G-SOCIETAL-KV-KVI]
Wikström, G., Schuler Scott, A., Mesogiti, I., Stoica, R., Georgiev, G., Barmpounakis, S., Gavras, A., Demestichas, P., Hamon, M., Hallingby, H., and D. Lund, "What societal values will 6G address?", 17 May 2022, <https://doi.org/10.5281/zenodo.6557534>.

[M6G-VALUE-PERF] Ziegler, V. and S. Yrjola, "6G Indicators of Value and Performance", 4 May 2020, <https://doi.org/10.1109/6GSUMMIT49458.2020.9083885>.

[Planet-B] Azote for Stockholm Resilience Centre, based on analysis in Richardson et al 2023, "Planetary boundaries", 2023,

<https://www.stockholmresilience.org/research/planetary-boundaries.html>.

[RFC6272] Baker, F. and D. Meyer, "Internet Protocols for the Smart Grid", RFC 6272, DOI 10.17487/RFC6272, June 2011, <https://www.rfc-editor.org/info/rfc6272>.

[RFC8403] Geib, R., Ed., Filsfils, C., Pignataro, C., Ed., and N. Kumar, "A Scalable and Topology-Aware MPLS Data-Plane Monitoring System", RFC 8403, DOI 10.17487/RFC8403, July 2018, <https://www.rfc-editor.org/info/rfc8403>.

[SDG-Acad] "SDG Academy", 2023, <https://sdgacademy.org>.

[Telefonica] Telefonica, "Consolidated Annual Report 2021", 2021, <https://statics.telefonica.com/en/financial-reports/2021/Consolidated-Annual-Report/#i7111116644124bd299d124820baf28f7_394>.

[UN-SDG] "Sustainable Development, the 17 Goals", 2023, <https://sdgs.un.org/goals>.

[UNEP-LCA] "Global guidance principles for life cycle assessment databases : a basis for greener processes and products", 2011, <https://www.lifecycleinitiative.org/library/global-guidance-principles-for-lca-databases-a-basis-for-greener-processes-and-products/>.

[UNGA42] UN, "Report of the World Commission on Environment and Development : "Our common future"", 4 August 1987, <http://digitallibrary.un.org/record/139811>.

Authors' Addresses

Carlos Pignataro (editor)
North Carolina State University
United States of America

Email: cpignata@gmail.com, cmpignat@ncsu.edu

Ali Rezaki
Nokia
Germany

Email: ali.rezaki@nokia.com

Suresh Krishnan
Cisco Systems, Inc.
United States of America

Email: sureshk@cisco.com

Hesham ElBakoury
Independent Consultant
United States of America

Email: helbakoury@gmail.com

Alexander Clemm
Futurewei
2220 Central Expressway
Santa Clara, CA 95050
United States of America

Email: ludwig@clemm.org