Network Working Group
Internet-Draft
Intended status: Informational
Expires: 8 January 2026

M. Welzl
University of Oslo
E. Stephan
Orange
E. Schooler
University of Oxford
S. Rumley
HES-SO
A. Rezaki
Nokia
J. Manner
Aalto University
C. Pignataro
Blue Fern Consulting
M. Palmero
Cisco
J. Lindblad
All For Eco
S. Krishnan
Cisco
A. Ker辰nen
Ericsson
H. ElBakoury

L. M. Contreras
Telefonica
A. Clemm
Independent
J. Arkko
Ericsson
7 July 2025

Architectural Considerations for Environmentally Sustainable Internet Technology
draft-various-eimpact-arch-considerations-01

Abstract

This document discusses protocol and network architecture aspects that may have an impact on the sustainability of network technology. The focus is on providing guidelines that can be helpful for protocol designers and network architects, where such guidelines can be given.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://jariarkko.github.io/draft-eimpact-arch-considerations/draft-eimpact-arch-considerations.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-various-eimpact-arch-considerations/.

Source for this draft and an issue tracker can be found at https://github.com/jariarkko/draft-eimpact-arch-considerations.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 8 January 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust’s Legal Provisions Relating to IETF Documents (https://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction

This document discusses protocol and network architecture aspects that can have an impact on the environmental sustainability of network technology. For brevity, we will use the term sustainability in this document to refer to environmental sustainability. We do note that sustainability as a term is widely used to refer to different notions of sustainability, and the most well-known larger definition of sustainability can be seen from the United Nations Sustainable Development Goals (UN SDG) [UNSDG].

Environmental sustainability is an important consideration in society, and in networking, too. Networking technologies enable societies to operate in an environmentally sustainable manner and thereby have a positive handprint, yet networks themselves must be environmentally sustainable and attempt to minimise their negative footprint.

Fundamentally the question we try to address concerns the resource usage and the lifecycle of network equipment. The less devices are built, and energy is used, the less emissions are created. Networks are built with hardware and these in turn use electrical energy to run. Eventually, the hardware is decommissioned and some amount of the materials are recycled.

We can divide the lifecycle into three major phases (omitting some intermittent steps like shipping of products):

1. Manufacturing (including the raw material extraction and usage, the embedded chips and electronics, casing, and energy needed for these operations, etc.),

2. Use phase that is focused on the operational energy use and repairing equipment, and

3. End of life that can include both direct recycling of some of the materials or finding a new life and usage for an old product that still functions, after which it is finally recycled.

Networks also require some amount of physical construction to realize, and this construction work also creates emissions. This category of emissions is out of scope of this document because the Internet community and network engineers have limited means to impact construction work itself and the associated industry, but we can impact how networks, protocols and hardware are designed, built and operated.

All these phases create harmful emissions, into the ground and in the air, that have a negative impact on our environment and people. As the type of such emissions vary, they are often standardized as carbon dioxide equivalent (CO2e) to allow comparing sources and amounts of emissions. When discussing (carbon) emissions in this document, we generally refer to CO2e.

The manufacturing of networking hardware, both for fixed and wireless networks, is a significant source of emissions, and recycling of ICT equipment is still limited to the casing and some other minor parts. Direct energy usage of networking and the source of the energy have often been the primary concerns. There are many reports and scientific papers discussing carbon emissions of the energy used by ICT. As of today, and the foreseeable future, the difference in emissions of the electric grid between countries and regions can vary significantly. e.g. In the EU, there are 10-fold differences between countries, and similar differences exist between US states. On a global level, the differences can be over 50-fold. Yet, as the world moves towards greener energy production, the relative negative impacts related to manufacturing becomes more prominent and the importance of equipment longevity grows.

When good design and architecture can improve the sustainability of networks, they should certainly be applied to designing new protocols and building networks. Intuitively, protocol and network architecture choices can have an impact on sustainability. At the very least the right design and architecture can make it possible to have a positive impact, but of course the architecture alone is not enough. The possibilities offered by the architecture need to be realized by implementations and practical deployments.

To give an example of an architectural aspect that potentially has a sustainability impact, enabling the collection of information (e.g., energy consumption) and then using that information to make smarter decisions is one. For instance, understanding power consumption of individual nodes can be valuable input to future purchasing decisions or development efforts to reduce the power consumption. Yet, as data collection is often rather easy, it is easy to overdo it in such a way that it leads to accumulation of dark data (i.e. data that is collected and stored but never used). All data collection consumes processing power, network resources and storage space, and this can in turn increase the emissions from the network.

Other architectural examples include making it possible to scale resources or resource selection processes performed in a sustainability-aware fashion. The use of communication primitives that maximize utility in a given problem (e.g., using multicast) or the use of technologies that reduce the number or size of messages needed for a given task (e.g., binary encoding instead of textual) are some further examples.

Of course, some of these aspects may have a major impact on sustainability, where others may only have a minor effect. There are also tradeoffs, such as side-effects of architectural choices, e.g., dynamic scaling of a router network potentially impacts jitter; putting cellular base stations to sleep and activating them as capacity needs grow potentially introduces a delay in matching the needs of the data flows.

The document is intended to help engineering efforts in the IETF, provide operational guidance in the operator community, as well as to point to potential research directions in the IRTF.

The scope of the document is advice on Internet and protocol architecture, such as what architecture or capabilities new protocol designs or features should have, what kind of operational network architectures should be deployed, and how all of these can be designed to best address sustainability concerns.

The focus of this document is to provide actionable design advice to protocol designers. This document therefore addresses one aspect in the architecture question and does not claim to cover the topic exhaustively.

This document is not focused on general issues around environmental sustainability, except those that pertain to architecture or significant protocol features.

It is to be noted that networks themselves are a service, a tool, for all the applications and services on the Internet. Networks connect data, people and services. The increase in networking and size of the Internet is driven by these applications and the usage. Therefore, the emissions from networking are tied to the applications and the data they consume; with less applications or data, the Internet would have less hardware and less energy usage. The goals of this document are not to instruct application and service developers to choose what applications are worthwhile or how much content is sent. There are many forums and parties whose mission is to help these developers to implement more sustainable services, such as, the Green Software Foundation, the Green Web Foundation, Greening of Streaming, to name a few.

The next two sections present architectural and protocol design aspects that can have an impact on the sustainability of networking. Section 2 discusses those foundations that are required to prepare for sustainability improvements, and Section 3 discusses actions that can be taken to make the actual improvements. For each topic in these sections, we provide an overview, the motivation for why it would be important to consider for more sustainable networking, an analysis and recommendations for future networking professionals.

Recommendations for protocol designers are discussed throughout the document and summarized in Section 4. Finally, Section 5 discusses further work that is needed to make further concrete recommendations for the designers.

2. Understanding

It is essential to understand the current state of affairs before any improvements can be made. Thus, some level of measurement is necessary for starting to improve sustainability. In many cases measurements are also complemented by modeling. In some cases modeling needs to be used since direct measurements may not always be available. Modeling is also used to combine measurements in ways that make it more effective in aiding the understanding of the effects of the potential actions. e.g. Modeling could be used to play out multiple what-if scenarios based on the actions recommended in Section 3.

2.1. Measurement and Modeling

The key goals of measuring and modeling are to identify potential areas of improvement, and to establish a baseline to benchmark any improvements that are effected by the performed actions. This not only helps defining an objective data-driven approach to improvement, but also can illustrate what actions can cause a bigger impact. This could help prioritize what actions can be taken and when. This draft assumes that the specific semantics of sustainability-related measurements (e.g., carbon factors, device-specific formulas) are defined elsewhere and focuses instead on enabling architectures to support measurement, collection, and use.

2.1.1. Motivation

Without measurements of any kind, it is impossible to assess if the networks are functioning correctly. It is impossible to know if the system is efficient by comparing it against a baseline model. It is also impossible to check that changes aiming at optimizing something are indeed valuable.

This is particularly the case when looking at the systems as a whole in post-analysis. As discussed earlier, some level of measurements is useful input for further actions, such as deciding what parts of the network need to be targeted for further improvement.

But measurements may also be useful for some dynamic situations where power-saving decisions, for instance, depend on knowing the relative power consumption of different activities, such as when a power-off decision involves understanding the relative savings during the shutdown period vs. the power cost of shutdown and startup procedures, or the possible need to reconfigure other nodes in the network due to the shutdown.

At the same time, it is not possible (or even desirable) to measure everything. Excessive measurement collection without clear objectives can have a negative impact by itself and some considerations in this regard can be found in Section 3.5 Furthermore, any measurement must be validated. Relevance of measurements must be periodically assessed, e.g., with comparisons between measurements within a network and the aggregate numbers from the electricity provider.

Finally, measurements made in the field must be collected and structured to allow later retrieval. And measurements are counterproductive if they are endlessly accumulated without being consulted.

2.1.2. Analysis

This section discusses how measurements relate to the fabrication and usage phases and how efficiency can be measured.

While power consumption is the most commonly used sustainability metric, this document does not attempt to define energy metrics or modeling standards. Those topics are in scope for the GREEN WG (focused on operational energy) and the SUSTAIN RG (which addresses broader life-cycle impacts and carbon modeling). This section focuses on the architectural implications of enabling measurement, not metric definitions.

2.1.2.1. Measuring impacts of fabrication phase

Network infrastructure generates negative impacts principally during fabrication and usage phases. Measuring negative impacts related to fabrication falls in the activity of lifecycle analysis (LCA). LCAs are typically performed per device, either by the equipment vendor itself, or by third-party analysts. LCA involves modeling (see Section 2.1.3.6). The analysis can be done in terms of climate change (CC) but can be extended to other criteria as abiotic resource depletion (ARD), ecotoxicity (ET) or water usage (WU). LCA also involves information systems keeping an inventory of the devices uses. For many classes of devices, the embedded carbon aspects or use of raw materials are significant sustainability issues. As these measurements and inventories are external to the network architecture, they are considered out of this document scope.

2.1.2.2. Measuring impacts of usage phase

Measuring negative impacts related to the usage phase falls in the scope of monitoring. In this phase, the most obvious sustainability-related measurement is power monitoring to measure the energy consumption and estimate the related negative impacts.

Power (in Watts, that is, in Joule/s) or energy (in Joules) measurements alone are of meager use if the cause of the consumption is not measured as well. Any power/energy measurement should occur alongside other measurements that can be used to determine energy efficiency. Hence a sound measurement architecture implies the existence of an energy efficiency framework of some kind.

2.1.2.3. Measuring efficiency

In the context of carbon accounting, emission accountants are generally looking for a metric of the delivered value per unit of carbon. In networking, the most obvious delivered value is number of bits sent or received (traffic), or to the communication capacity made available during unit of time. In both cases, the unit is the bit, but the two metrics have very different meanings. In one case, one spends a Joule to send a bit. In the other case, one spends a Joule to offer a bandwidth capacity of 1 bit/s during a second. The latter is important, as often communication networks have requirements to be able to send messages when there’s a need for it, e.g., for emergency communications, even when those messages may not always be sent.

The measurement of efficiency is not restricted to energy. Traffic or offered bandwidth can be related to the carbon emitted by the device traversed by this traffic. This carbon should include the part associated with electricity, but also the part associated with fabricating the device (pro rata temporis) [LCAandUsage]. Sustainable efficiency can also be expressed in water used per traffic, for example.

2.1.3. Recommendation

The GREEN WG is chartered to define energy consumption metrics and associated frameworks. The GREEN framework provides a foundational building blocks for monitoring and optimizing energy consumption across networked devices and components.

The SUSTAIN RG addresses broader measurement questions such as embedded emissions, raw materials, and life-cycle modeling. This document assumes these efforts will define and validate the metrics themselves. Our focus is on ensuring that Internet architecture enables effective collection, transport, and use of such metrics for operational decisions and reduction of environmental impacts.

Aligning these efforts will support the development of composite metrics that connect operational energy use along with manufacturing/ end-of-life considerations in order to establish a coherent basis for sustainable digital infrastructure management.

In order to meet the needs discussed above, the following architectural design principles are proposed.

2.1.3.1. Future Proof Metrics

We recommend that any measurement framework or sustainability-related information sharing mechanism be designed to share different types of information and not limited to a single metric such as power consumption. Requirements, units, granularity and collection method specifications are sure to shift over time.

2.1.3.2. Plug-in Architecture for Collection and Control

Since the need to deliver on the use cases described is urgent, the industry has to accommodate the capabilities (and limitations) of existing equipment in the field for collecting metrics. It is recommended to apply a plug-in architecture with modules that can work with (read from and control) devices of any kind, including traditional networking hardware devices, cooling systems, software stacks, and occasionally static data sheets.

2.1.3.3. Data with Content Declaration

To make sense of the collected data, it must be possible to see exactly where all data is coming from, what it means, its precision and how it has been processed. The metadata itself must also have a formal description. YANG might be suitable for modeling the metadata schema. Keep the metadata attached to the dataflow it describes, so that the relation is clear even when components are added by other organizations at a later point in time.

2.1.3.4. Processing Flexibility and Audit Trails

The collected data passes through a pipeline from collection to decisions. By processing we mean steps to reshape the data to match further aggregation and processing steps, such as unit conversions, sample frequency alignment, filtering, etc.

Separate these pipeline stages into separate modules and use configuration to construct the pipeline. This gives flexibility, reuse and enables a full audit trail. It is essential that every data processing step can be reviewed in an audit situation without looking at code.

2.1.3.5. Aligned with Reporting Frameworks

Ensure that the system output is aligned with the measurement requirements set forth by relevant legal frameworks, e.g. ESRS (Europe), TCFD and IFRS (US, Japan), BRSR (India), etc. The responsible corporate bodies producing the corporate reports are unlikely to use any technical collection system that isn’t well aligned.

2.1.3.6. Modeling

Where power optimization choices are made, accurate information is required to decide the right choice.

The paucity of up-to-date information on equipment and system parameters, especially power consumption and maximum throughput, makes estimating the power consumption and energy efficiency of these systems extremely challenging. In addition, the rapid evolution of technology and products in ICT makes the estimation quickly outdated and possibly inaccurate. In some cases, physical measurements have to be replaced by partial measurements and mathematical modeling.

2.1.3.6.1. Power modeling

To date, two approaches to network power modeling are accepted as providing a realistic estimate of network power consumption. These approaches are referred to as "bottom-up" and "top-down". The paper [Unifying] surveys both approaches and provide a new approach which unifies both of them. The unified approach is used to estimate the power consumption of access, aggregation and core networks.

Modeling can also help address attribution aspects, such as those involved in an effort of an organization to calculate its Scope 3 emissions. Modeling can also be used to assist in establishing a baseline energy consumption, and enable later comparisons to that baseline.

Additional discussion of modeling can be found in Appendix A.

3. Actions

3.1. Dynamic Scaling

Dynamic scaling is the ability to adjust resources according to demand and possibly turn some of them off during periods of low usage. Examples include the set of servers needed for a service, how many duplicate links are needed to carry high-volume traffic, whether one needs all base stations with overlapping coverage areas to be on, etc.

Networks and communications are also critical functions of the modern digital society. The reliability of individual networking links or devices cannot always be guaranteed. As a result, various levels and forms of resiliency are often needed, for instance through redundancy. Yet, there is a question on how much redundancy is needed and how quickly a backup or resource increase can be activated due to increased demand.

Scaling can be pulled up and down by data consumption variations and more rarely by power shortage. In such situation dynamic scaling is the ability to adjust demand resources according to resources. When operating on limited backup energy sources such as batteries or generators, the architecture must support graceful adaptation before power runs out. In such situations, networks must minimize consumption to extend operational time.

3.1.1. Motivation

Outside of implementation improvements, dynamic scaling is potentially the most promising method for reducing power consumption related environmental impacts. Scaling can happen on a device-level (increasing performance as traffic levels grow) or a network segment level (increasing the number of active links or cellular base stations).

Considering current fixed networking hardware, dynamic scaling might not have an impact in situations where there’s only a single router or server serving a particular route, area, or function. Current routers and switches exhibit limited potential dynamic scaling because the focus is on high performance and a stable connectivity. There have been some recent improvements on this front as well. e.g. Energy-Efficient Ethernet (EEE) is a good example of a networking-level specification to lower energy consumption in idle mode. EEE has limited impact on a network that has continuous traffic.

Resiliency can be implemented within a single router as well, e.g. as a backup power supply, between routers and switches as multiple links between the same nodes, having different links between two end points, overlapping cellular coverage, etc. All these necessarily add more hardware to provide the same exact service. Some of that hardware can be fully operational at all times and used to serve the traffic, while other links may be in hot or cold standby depending on the use case.

Cellular networks are typically built with significant overlap in coverage areas of multiple base stations. Demand and business reasons dictate the design of the coverage, and regulations might dictate how reliable the cellular service should be. There is extensive work world-wide to optimize the operation of this overlapping coverage, e.g. by turning down some sites at night time when traffic volumes are low. A cellular base station site can consume anything from a few kWh to ten or more kWh per provider. Modern cellular base stations do implement numerous features to scale the energy consumption. In general, cellular base stations have a base energy consumption and traffic-dependent consumption, a somewhat similar behavior to what we can observe in modern CPUs.

On the network level, most large systems have significant amount of redundancy and spare capacity. Where such capacity can be turned on or off to match the actual need at a given time, significant reductions in power consumption can be achieved.

Whereas scaling down under normal conditions seeks to reduce consumption while maintaining full capabilities, power-constrained operations accept degraded performance or functionality. Operating in power backup mode introduces a shift in network behavior as it differs from network-driven auto scaling:

* Network, devices and components must reduce power usage, possibly sacrificing performance, feature sets, or redundancy.

* Each network domain (RAN, edge, and core network segments) faces its own constraints and policies in power-limited operation.

3.1.2. Analysis

Dynamic scaling could be seen as either an alternative or complementary to load stabilization, e.g., via "peak shaving". Perhaps the most realistic view is that both are likely needed.

The most rudimentary approach to dynamic scaling is just turning some resources off. However this may not be sufficient, and a more graceful/engineered approach potentially yields better results.

Network architects need to understand the impacts of scaling changes on users and traffic. These may include the fate of ongoing sessions, latency/jitter, packets in flight, or running processes, attempts to contact resources that are no longer present, and the time it takes for the network to converge to its new state.

Dynamic scaling requires an understanding of load levels for the network, so information collection is required. It also requires understanding the power, time and other costs of making changes. (See [I-D.pignataro-enviro-sustainability-architecture] for discussion of tradeoffs and multi-objective optimization.)

Understanding the resiliency requirements for a network or a piece of equipment is also important for the optimal control of resiliency, e.g., as an input to decisions on how many instances of replicated services need to be run and where.

Some of the strategies that are useful in implementing effective dynamic scaling include:

* Matching the currently used resources to the actual need, be it about traffic demand or resiliency. One way to do this is to use power-proportional underlying technologies, such as chipsets or transmission technologies. And where this is not sufficient, the ability to turn components/systems on and off is an alternative strategy.

* Using load adaptive techniques allows the capacity of the nodes to be dynamically adjusted according to the demand. Examples include Adaptive Link Rate (ALR), which dynamically adapts the link rate to suit traffic demand or power off links in Link Aggregation based on traffic demand which is empirically estimated based on traffic arrival. LACP (Link Aggregation Control Protocol) defined in IEEE 802.1AX [LinkAggregation] can be modified to power off links in an aggregation if they are not needed.

* Ability to enter "no new work" mode for equipment, to enable some resources to be eventually released/turned off.

* Ability to move ongoing tasks off to other equipment, to prevent disruption of already started tasks.

* Ability to schedule changes in advance rather than making them abruptly, with associated signaling exchanges and possible transient routing and other failures. See for instance the time-variant routing work in the IETF [RFC9657] [I-D.ietf-tvr-requirements] [I-D.ietf-tvr-schedule-yang] [I-D.ietf-tvr-alto-exposure].

* Efficient propagation of changes of new routes, new set of servers, etc. in order to reduce the amount of time where state is not synchronized across the network. The needs for the propagation solution needs to be driven by dynamic scaling and sustainability as well as other aspects, such as recovery from failures.

* Build mechanisms to deal with dynamic changes: Plan for dynamic set of resources and not expect to work with a fixed set of resources.

* Dynamic scaling requires automation in most cases, e.g., to turn on new service instances. See again [I-D.pignataro-enviro-sustainability-architecture] for a discussion of automation.

* Interaction with the energy grid can enable dynamic load shifting. For instance, a demand-response technique can be used where the system temporarily reduces its energy usage in response to pricing signals from a smart grid. The proposed demand-response technique involves deferring the load from elastic requests to later time periods in order to reduce the server demand and the current energy usage, and hence, energy costs [LoadShifting].

* Energy-aware routing. This generally aims at aggregating traffic flows over a subset of the network devices and links, allowing other links and interconnection devices to be switched off. These solutions shall preserve connectivity and QoS, for instance by limiting the maximum utilization over any link or ensuring a minimum level of path diversity. There are also algorithms for Green Traffic engineering. For instance, [Segment] employs segment routing. Experimental analysis results [Experiment] show that the resource usage for SRv6 could be more than 70% lower than that of the SPF-based forwarding, depending on the network topology.

3.1.3. Recommendation

The guidelines above need to be considered specifically for each protocol and system design. Further work in detailing this guidance would also be useful.

It is likely that there is increased attention to resiliency in the future, given for instance the increased importance of the tasks supported by networks or the potentially increasing frequency of natural disasters as a result of global warming.

Scaling steps during power shortage differ from network dynamic scaling and depend on the network domain and the events: grid outages, deployment in remote or mobile environments, extreme weather events, or any sort of enforced reductions in power usage like monthly battery testing. Nevertheless, there is a gain to have a common dynamic scaling approach that includes network-driven scaling and power-shortage scaling.

3.2. Transport

Transport protocols make it possible for communication flows to adjust themselves to the dynamic conditions that exist in the network at any given time: available bandwidth, delays, congestion, the ability of a peer to send or receive traffic, and so on. Depending on the conditions, an individual flow may carry traffic at widely different rates, may pause for some time, etc.

This behavior has an effect on sustainability, e.g., in what periods the endpoint and network systems are active or when they could be in reduced activity or sleep states. Cellular networks and mobile links can scale their energy usage based on load and enter a low-power state when a traffic flow ends. Thus, in theory, the faster the data is transferred, the faster the device transmission/reception functions can enter a low-power state.

3.2.1. Motivation

Transport behavior would have a possibility of impacting how much downtime or sleep can be had in the communication system, either on the end systems or routers or other equipment in between. The savings can be significant, at least in wireless systems.

Improvements through transport behavior are only useful if the involved systems have power proportionality.

3.2.2. Analysis

Various higher-level transport solutions may also cache or pre-fetch information. For instance, [I-D.irtf-nmrg-green-ps] lifts CDNs as one example of technology that has reduced energy consumption, by moving the needed endpoints closer to each other.

On a given set of endpoints, application behavior can impact environmental costs. For instance, [I-D.pignataro-enviro-sustainability-consid] observes the effect of protocol chattiness. 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?

Transport layer protocol behavior also has an impact. A critical issue is the tradeoff involved in sending traffic. As argued in [NotTradeOff], reducing the amount of time the endpoints and the network are active can sometimes help save energy. As a result, in general, delivering information as rapidly as possible would appear to be desirable.

On the other hand, would such as rapid transmission impact peak traffic, and as such, contribute to a need to dimension networks for higher traffic volumes? And in this case the need could be only a perceived one as a less rapid transmission would not have impacted, for instance, a user’s ability to view a video if the transmission was merely for the buffering of the rest of the video.

Furthermore, bandwidth-intensive applications can influence other applications or users by presenting a significant load on the network, and consequently reducing capacity available for others, or increasing buffering (and with it, latency) across the network path. For an application with intermittent data transfers, such as streaming video, this would seem to speak in favor of sustained but lower-rate delivery instead of transmitting short high-rate bursts [Sammy]. However, this is in contradiction with the energy-saving approach above. Thus, the tradeoff is: should data be sent in a way that is "friendly" to others (avoiding bad interference), or should it save energy by sending fast, increasing the chance for equipment to enter a "sleep" state?

At the time of writing, the common choice for video is to opt for higher rate delivery, potentially saving energy, and possibly at the expense of other traffic. For non-urgent data transfers, the IETF-recommended default approach is the opposite: the LEDBAT congestion control mechanism [RFC6817], which is designed for such use, will always "step out of the way" of other traffic, giving it a low rate when it competes with any other traffic. Alternatively, if the goal is to reduce energy, such traffic could be sent at a high rate, at a strategically good moment within a longer time interval; this would give network equipment an opportunity to enter a sleep state in the remaining time period within the interval.

A hypothesis could be made that transport protocols should take energy into account in addition to the many other inputs they decide upon. For example, it is possible that a non-urgent data transfer would send as much as possible as soon as possible when at least one of the links along the path is known to be power proportional (e.g., a cellular link), while tracking buffer growth from transmission delays to scale back if delay should occur.

Such ideas remain to be confirmed with experiments, however.

Similarly, caching and pre-fetching designs need to consider not only the likelihood of having acquired the right content in memory, but also the sustainability cost of possibly fetching too much or the timing of those fetching operations.

In general, information about the impacts of loading or not loading the network with additional traffic, and whether a certain sending pattern enables power savings through sleep modes, would be beneficial for the communicating endpoints. Mechanisms for making such information available to the endpoints would be useful.

3.2.3. Recommendation

As can be seen from the above, there are a number of complex tradeoffs merely for transport protocol behavior on a given connection.

This prompts us to give two types of advice. The first type of advice is for protocol designers: simple models are unlikely to guarantee optimal results, but as long as normal precautions such as congestion control, monitoring queue build-up, and avoiding unnecessary messages are employed, systems will operate reasonably well.

The second type of advice is for further work in the research community to better understand what strategies would actually provide the best end-user and energy performance, and whether the choice of strategy depends on other factors, such as whether sleep modes are implemented in network nodes. There is a clear need for simulations and experiments to understand this better. This may be work that fits within the IRTF SUSTAIN research group. Also, new standards may be needed if information sharing about the sustainability and sleep mode characteristics of network systems is needed for applications to make the best transport decisions.

3.3. Equipment Longevity

This section discusses the ability to extend the useful life of protocols and/or network equipment in order to amortize the embedded energy costs over a longer period, even though it may mean that the protocols/equipment may not be fully optimized for the present use. This includes devising tools to inform network administrators and their users of the potential benefits of network equipment upgrades, so that they can make better choices on what upgrades are necessary and when.

It should be noted that from an environmental sustainability perspective, it may not always be the best choice to upgrade network equipment whenever slightly less power-hungry and "greener" alternatives become available. The environmental cost of amortizing the carbon embedded inside equipment over its lifetime, including the carbon associated with the manufacturing of the equipment that is to be replaced, should be taken into consideration as well.

3.3.1. Motivation

Embedded carbon and raw materials can be a significant part of the overall environmental impact of systems. If this can be improved for devices that are manufactured in large quantities, the improvements can be significant.

The more the world moves toward low-carbon energy sources, the more the manufacturing matters in the holistic view. Today there can be an order of magnitude difference in average emissions for a kWh of electricity between two countries. Thus, any estimates that seek to compare the manufacturing and use phase emissions of a network equipment would have to be calculated per country or region, and there is no universal standard for the whole planet.

Long equipment lifetimes are only useful if the longer lifetimes can be achieved without compromising other aspects of sustainability, such as when using a high-end and power-hungry router in place of small routers. The exact moment when a hardware change is warranted for sustainability differs between countries and regions.

3.3.2. Analysis

When we engineer protocols and network equipment, we are inclined to design them in a highly optimized manner for a very specific set of requirements, use cases and context. While this is necessary in certain cases (e.g. constrained nodes with limits on processing capacity or long-lived battery powered devices), there are certainly cases where such optimized equipment is not absolutely required. Most infrastructure network nodes on the Internet utilize only a fraction of their design capacity most of the time.

Designing the equipment with an eye on longevity comes with a set of advantages:

* It allows the same equipment and protocols be reused in a different context in the future. e.g. A core router of today can become an edge router in a near future and an access router in the further future if the protocol implementations are adaptable.

* It can reduce complexity in implementations as well as in network management that are usually inherent in highly optimized systems

* It can let network equipment operate for a longer period and can reduce the frequency of hardware upgrades, in turn reducing the environmental impact associated with manufacturing, transporting, and disposing of the old/new hardware.

* One key disadvantage may be that not optimizing may result in the need for premature upgrades for capacity and this needs to be considered.

Hence, it is very likely that extending the life of protocols and equipment with higher flexibility could provide a better environmental benefit than tightly optimizing only for today窶冱 uses.

Another aspect that can play an important role in extending the longevity of equipment concerns software-defined networking, in the sense of designing networking equipment in such a way that new equipment capabilities and features can be introduced via software upgrades as opposed to requiring hardware replacement. This requires system architectures that incorporate the necessary infrastructure to support such upgrades in a secure manner that does not compromise equipment integrity.

On the other hand, it is very much possible that there could be new equipment available that is significantly more sustainable in its operation. The longevity of the existing equipment and the amortization of its embedded sustainability costs, needs to be balanced against the potential operational savings to be realized by upgrading to newer equipment over the intended lifecycle of the newer equipment.

3.3.3. Recommendation

The guidelines above should be considered for any new system design. If some aspect of protocol or network equipment design choice could be made more generic and flexible without a significant performance and sustainability impact, it needs to be studied in further detail. Specifically, the potential additional sustainability costs due to forgoing optimization need to be weighed against the potential savings in embedded carbon and raw material costs brought about by premature upgrades.

There are also cases where equipment upgrades are done to provide better peak performance characteristics (e.g. higher advertised speeds towards consumers) and these need to be viewed as well with the same tradeoffs in mind. Also, when newer more sustainable equipment is available there needs to be a cost benefit analysis made to decide whether to keep current equipment running for longer or upgrade to realize the benefits of newer equipment even though it incurs new embedded costs.

Finally, when designing networks, it is recommended to consider whether it is possible to reuse retiring equipment in a different location or for a different function (e.g. move it to lower traffic geographies, core routers become edge/access routers etc.)

3.4. Encoding

This is about considering the effects encoding methods on sustainability, such as the use of binary encodings instead of text.

3.4.1. Motivation

Better encoding can obviously reduce the length of messages sent or reduce the amount of computing required for the encoding and decoding operations. It remains a question mark how big overall impact this is, however. It should only be performed if it gives a measurable overall impact.

3.4.2. Analysis

Better encoding methods are clearly beneficial for improving the detailed-level effectiveness of communications.

The main questions are, however:

* How large are the potential remaining savings in this area, and how do they compare to other things? Particularly considering that much of the traffic on the Internet is video, which is already highly optimized and constantly updated with better encoding methods. Moran et al. argued in their 2022 paper [CBORGreener] [RFC9547] that that for a weather data example from [RFC8428] [RFC9193] there are significant savings. However, this needs more research in terms of the overall impact across different examples and the general make up of Internet traffic.

* At what layer is the compactness achieved? Are link, IP, or transport layer mechanisms that can compact some of the verbose messaging useful, or should each protocol have optimal compacting?

* Tradeoffs related to compute required to do encoding and decoding operations. These can be relatively heavy operations, particularly if compression is performed, particularly if AI-based computationally expensive methods are used.

3.4.3. Recommendation

More research is needed to quantify the likely sources of measurable impacts.

Of course, new protocols can generally be designed to work with compact encoding, unless there is a significant reason not to. But efforts to modify existing protocols for the sake of encoding efficiency should be further investigated by the above-mentioned quantification results.

One particular area of interest is the impact of AI-based compression methods and their computational and energy costs vs. achieved savings in communication efficiencies.

3.5. Sustainable by Design: Data Governance Perspective

Incorporating sustainability into the design phase of network architecture is critical for ensuring long-term environmental and operational benefits. From a Data Governance point of view, "Sustainable by Design" involves embedding sustainability principles and practices into the data management frameworks and processes from the outset.

3.5.1. Motivation

Data governance plays a pivotal role in shaping how data is collected, stored, processed, and used. By integrating sustainability into these processes, organizations can ensure that their data practices contribute to environmental goals, such as reducing carbon footprints, optimizing resource usage, and minimizing waste.

3.5.2. Analysis

Key elements of Sustainable by Design in data governance include:

* Data Minimization: Collecting only the data that is necessary and useful, reducing storage and processing requirements, which in turn lowers energy consumption.

* Efficient Data Storage Solutions: Implementing energy-efficient data storage technologies and practices that prioritize reduced power usage and cooling needs.

* Lifecycle Management: Ensuring that data is managed throughout its lifecycle in a way that minimizes environmental impact, including secure and sustainable data disposal practices.

* Transparency and Accountability: Establishing clear data governance policies that promote transparency in data usage and accountability for sustainability objectives.

3.5.3. Recommendation

Organizations should adopt data governance frameworks that incorporate sustainability as a core principle. This includes setting clear sustainability goals, measuring progress towards these goals, and continuously improving data management practices to enhance sustainability. By doing so, organizations can ensure that their data operations are not only effective but also environmentally responsible.

There is a protocol designer angle in this as well. Protocol designers should consider at least the data minimization aspects from Section 3.5.2, and may additionally consider providing mechanisms for the lifecycle management and transparency aspects.

4. Recommendations for Protocol Design

The recommendations that can be applied by protocol designers and architects have been listed in Section 2 and Section 3. Specifically:

* Measurement and modeling are a necessary foundation to understand where environmental impacts are generated, and to quantify any improvements. The recommendations related to this topic were listed in Section 2.1.3. These are primarily about ensuring that the measurement frameworks are generic enough to support data collection for an evolving set of metrics, and to prepare for the possibility that mathematical modeling may have to replace measurements in some cases.

* Dynamic scaling is the ability to respond to demand variations and resiliency requirements while optimizing energy consumption clearly has significant potential for savings. Recommendations related to this were listed in Section 3.1.3. These are about some basic techniques for being able to scale systems up and down while avoiding negative effects from these operations.

* Transport-related recommendations were listed in Section 3.2. These are about tradeoffs associated with different transport strategies.

* Longevity-related recommendations were listed in Section 3.3.3. These are primarily about how equipment can fulfill evolving roles over its lifetime, and associated tradeoffs.

* Encoding-related recommendations were listed in Section 3.4.3. These are about the effects of encoding size in protocols, and the associated compression computing impacts.

* Data governance-related recommendations were listed in Section 3.5.3. These are primarily about ensuring the right amount of data is collected, stored, and processed, in view of the effort required to do so.

5. Recommendations for Further Work and Research

There are several areas where concrete advice for protocol designers could not be given, or additional advice would be useful, but we do not understand the situation well enough to give practical advice.

These include:

* Past and ongoing work in various systems and protocols has looked at dynamic scaling extensively, but we believe work also remains. Any large-scale system likely benefits from further analysis, unless already ongoing. Guidance in Section 3.1 simple, and further work in detailing this guidance would also be useful.

* Transport-related optimizations (see Section 3.2) that enable devices to consume less power by sleeping more appear to have potential for significant savings but confirming this requires further research. Such research could be performed in the context of the recently chartered SUSTAIN research group.

* More research is needed to quantify the likely sources of measurable impacts when it comes to efficient protocol message encoding discussed in Section 3.4. Also, the tradeoffs involving the use AI-based compression methods deserve further study. Again, these are topics that the research group could take on.

6. Security Considerations

It is possible that the introduction of features and architectural properties to facilitate environmentally sustainable Internet technology introduces new attack vectors or other security ramifications.

For example, the introduction of measurements and metrics for the purpose of saving energy could be misused for the opposite effect when compromised. For example, measurements might be tampered with in order to cause an operator to waste energy. Energy measurements, when abused, might also result in compromised security, for example by allowing to infer usage profiles. They could also be abused to implement a covert communications channel in which information is leaked via tampered measurement values that are being reported.

Networking features and technology choices may have security implications regardless of why they are introduced, including for reasons of environmental sustainability. The possibility of this needs to be taken into consideration, understood, and communicated to allow for their mitigation.

7. IANA Considerations

This document has no IANA actions.

8. Informative References

[Baseline] Livieratos, S., Panetsos, S., Fotopoulos, A., and M. Karagiorgas, "A New Proposed Energy Baseline Model for a Data Center as a Tool for Energy Efficiency Evaluation", International Journal of Power and Energy Research, Vol. 3, No. 1 , April 2019.

[BenchmarkingFramework]
Mahadevan, P., Sharma, P., Banerjee, S., and P. Ranganathan, "A Power Benchmarking Framework for Network Devices", In L. Fratta et al. (Eds.): NETWORKING 2009, LNCS 5550, pp. 795窶08 , 2009.

[CBORGreener]
Moran, B., Birkholz, H., and C. Bormann, "CBOR is Greener than JSON", Position paper in the 2022 IAB Workshop Environmental Impact of Internet Applications and Systems , October 2022.

[Experiment]
Groningen, J. and C. Lung, "Green Network Traffic Engineering Using Segment Routing: An Experiment Report", 2024 20th International Conference on Network and Service Management (CNSM) , 2024.

[I-D.cparsk-eimpact-sustainability-considerations]
Pignataro, C., Rezaki, A., Krishnan, S., ElBakoury, H., and A. Clemm, "Sustainability Considerations for Internetworking", Work in Progress, Internet-Draft, draft-cparsk-eimpact-sustainability-considerations-07, 24 January 2024, <https://datatracker.ietf.org/doc/html/draft-cparsk-eimpact-sustainability-considerations-07>.

[I-D.ietf-tvr-alto-exposure]
Contreras, L. M., "Using off-path mechanisms for exposing Time-Variant Routing information", Work in Progress, Internet-Draft, draft-ietf-tvr-alto-exposure-02, 5 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-tvr-alto-exposure-02>.

[I-D.ietf-tvr-requirements]
King, D., Contreras, L. M., Sipos, B., and L. Zhang, "TVR (Time-Variant Routing) Requirements", Work in Progress,

Internet-Draft, draft-ietf-tvr-requirements-06, 7 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-tvr-requirements-06>.

[I-D.ietf-tvr-schedule-yang]
Qu, Y., Lindem, A., Kinzie, E., Fedyk, D., and M. Blanchet, "YANG Data Model for Scheduled Attributes", Work in Progress, Internet-Draft, draft-ietf-tvr-schedule-yang-05, 4 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-tvr-schedule-yang-05>.

[I-D.irtf-nmrg-green-ps]
Clemm, A., Pignataro, C., Westphal, C., Ciavaglia, L., Tantsura, J., and M. Odini, "Challenges and Opportunities in Management for Green Networking", Work in Progress, Internet-Draft, draft-irtf-nmrg-green-ps-06, 15 March 2025, <https://datatracker.ietf.org/doc/html/draft-irtf-nmrg-green-ps-06>.

[I-D.pignataro-enviro-sustainability-architecture]
Pignataro, C., Rezaki, A., Krishnan, S., Arkko, J., Clemm, A., ElBakoury, H., and S. Prabhu, "Architectural Considerations for Environmental Sustainability", Work in Progress, Internet-Draft, draft-pignataro-enviro-sustainability-architecture-02, 12 May 2025, <https://datatracker.ietf.org/doc/html/draft-pignataro-enviro-sustainability-architecture-02>.

[I-D.pignataro-enviro-sustainability-consid]
Pignataro, C., Rezaki, A., Arkko, J., Clemm, A., ElBakoury, H., and S. Prabhu, "Sustainability Considerations for Networking Protocols and Applications", Work in Progress, Internet-Draft, draft-pignataro-enviro-sustainability-consid-02, 12 May 2025, <https://datatracker.ietf.org/doc/html/draft-pignataro-enviro-sustainability-consid-02>.

[LCAandUsage]
Weppe, O., Bekri, D., Guibert, L., Aubet, L., Prvotet, J., Pelcat, M., and S. Rumley, "Carbon Topography Representation: Improving Impacts of Data Center Lifecycle", Proceedings of the 4th Workshop on Sustainable Computer Systems (HotCarbon’25) , 2025.

[LinkAggregation]
"IEEE Standard for Local and Metropolitan Area Networks--Link Aggregation", IEEE STD 802.1AX-2020 (Revision of IEEE STD 802.1AX-2014): 1333. doi:10.1109/IEEESTD.2020.9105034. ISBN 978-1-5044-6428-4 , May 2020.

[LoadShifting]
Mathew, V., Sitaraman, R. K., and P. Shenoy, "Reducing energy costs in Internet-scale distributed systems using load shifting", Sixth International Conference on Communication Systems and Networks (COMSNETS), Bangalore, India, pp. 1-8, doi: 10.1109/COMSNETS.2014.6734894 , 2014.

[Modeling] Chen, C., Barrera, D., and A. Perrig, "Modeling Data-Plane Power Consumption of Future Internet Architectures", IEEE 2nd International Conference on Collaboration and Internet Computing (CIC), Pittsburgh, PA, USA, pp. 149-158, doi: 10.1109/CIC.2016.031 , 2016.

[NotTradeOff]
Welzl, M., "Not a Trade-Off: On the Wi-Fi Energy Efficiency of Effective Internet Congestion Control", 17th Wireless On-Demand Network Systems and Services Conference (WONS), Oppdal, Norway, pp. 1-4, doi: 10.23919/WONS54113.2022.9764413 , 2022.

[RFC6817] Shalunov, S., Hazel, G., Iyengar, J., and M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)", RFC 6817, DOI 10.17487/RFC6817, December 2012, <https://www.rfc-editor.org/rfc/rfc6817>.

[RFC8428] Jennings, C., Shelby, Z., Arkko, J., Keranen, A., and C. Bormann, "Sensor Measurement Lists (SenML)", RFC 8428, DOI 10.17487/RFC8428, August 2018, <https://www.rfc-editor.org/rfc/rfc8428>.

[RFC9193] Kernen, A. and C. Bormann, "Sensor Measurement Lists (SenML) Fields for Indicating Data Value Content-Format", RFC 9193, DOI 10.17487/RFC9193, June 2022, <https://www.rfc-editor.org/rfc/rfc9193>.

[RFC9547] Arkko, J., Perkins, C. S., and S. Krishnan, "Report from the IAB Workshop on Environmental Impact of Internet Applications and Systems, 2022", RFC 9547, DOI 10.17487/RFC9547, February 2024, <https://www.rfc-editor.org/rfc/rfc9547>.

[RFC9657] Birrane, III, E., Kuhn, N., Qu, Y., Taylor, R., and L. Zhang, "Time-Variant Routing (TVR) Use Cases", RFC 9657, DOI 10.17487/RFC9657, October 2024, <https://www.rfc-editor.org/rfc/rfc9657>.

[Sammy] Bruce Spang, Shravya Kunamalla, Renata Teixeira, Te-Yuan Huang, Grenville Armitage, Ramesh Johari, and Nick McKeown, "Sammy: smoothing video traffic to be a friendly internet neighbor", In Proceedings of the ACM SIGCOMM 2023 Conference (ACM SIGCOMM ’23). Association for Computing Machinery, New York, NY, USA, 754768. https://doi.org/10.1145/3603269.3604839 , 2023.

[Segment] Lung, C. and H. ElBakoury, "Exploiting Segment Routing and SDN Features for Green Traffic Engineering", IEEE 8th International Conference on Network Softwarization (NetSoft), Milan, Italy, pp. 49-54, doi: 10.1109/NetSoft54395.2022.9844091 , 2022.

[Unifying] Ishii, K., Kurumida, J., K.-i Sato, Kudoh, T., and S. Namiki, "Unifying Top-Down and Bottom-Up Approaches to Evaluate Network Energy Consumption", In Journal of Lightwave Technology, vol. 33, no. 21, pp. 4395-4405, doi: 10.1109/JLT.2015.2469145 , November 2015.

[UNSDG] "United Nations Sustainable Development Goals", https://unstats.un.org/sdgs , 2017.

Appendix A. Modeling Approaches and Literature

The paper [Modeling] provides a model for IP Routers and the routers of other future Internet architectures (FIA) such as SCION and NEBULA. They use a generic model which captures the commonalities of IP router as well as the peculiarities of FIA routers. They conduct a large-scale simulation based on this router model to estimate the power consumption for different network architectures.

Since routers and other network devices and functions can be virtualized, this article (1) provides comprehensive "graphical, analytical survey of the literature, over the period 20102020, on the measurement of power consumption and relevant power models of virtual entities as they apply to the telco cloud." This paper A Methodology and Testbed to Develop an Energy Model for 5G Virtualized RANs IEEE Conference Publication IEEE Xplore got best paper award for GreenNet 2024, but I am not sure if we are interested to model 5G vRAN.

There is a plethora of publications on modeling communication networks and DC computing.

A.1. Customer Attribution

When organizations assess their Scope 3 emissions, they need to sum up their share of emissions from all their suppliers, one of which for example, might be a cloud hosting service. In order for the supplier to provide an emission share value back to the customer, the provider needs to develop a mechanism for attribution.

A significant challenge in accurately assessing Scope 3 emissions is avoiding Double Counting, where the same emission is reported by multiple entities. According to the GHG Protocol best practices, it is crucial to establish clear guidelines and agreements between suppliers and customers to ensure that emissions are attributed correctly and not counted multiple times. This requires transparent communication and precise emission reporting standards to ensure that all parties involved have a consistent understanding of which emissions belong to which organization.

By addressing the Double Counting issue, companies can achieve more accurate and reliable Scope 3 emissions assessments, thereby contributing to better overall sustainability reporting and improvement efforts.

A.2. Baselining and Benchmarking

Establishing a baseline is a fundamental step in the process of improving energy efficiency and sustainability of network technology. Baselining involves establishing a reference point of typical energy usage, which is crucial for identifying inefficiencies and measuring improvements over time. In this step, the controller uses only the collected data from datasheets and other reliable sources.

By establishing a baseline and using benchmarking, organizations can determine if their networking equipment is performing normally or if it is deviating from expected performance. This is the first step in identifying and guiding necessary improvements. Benchmarking involves collecting performance measurements of networking equipment under controlled conditions. This process helps establish standardized performance metrics, allowing for comparison against baselines collected during regular operational conditions.

The initial measurement of networking equipment’s energy efficiency and performance, known as Baselining, should be coordinated with vendor specifications and industry standards to understand what is considered normal or optimal performance. For example, if the baseline indicates that your switches operate at 5 Gbps per watt, while vendor specifications suggest 8 Gbps per watt and the industry standard is 10 Gbps per watt, actions should be taken to implement energy-saving measures and upgrades. Continuously tracking subsequent measurements can reveal if efficiency improves towards the benchmark of 8-10 Gbps per watt.

This practice ensures that any improvements can be quantifiably tracked over time, providing a clear measure of the effectiveness of the implemented changes and guiding further enhancements in network sustainability.

See also [Baseline] and [BenchmarkingFramework].

Acknowledgments

Everyone on the author section has contributed to the document in significant ways. The author list has been ordered in (reverse) alphabetical order.

Parts of this document extensively leverage ideas and text from [I-D.cparsk-eimpact-sustainability-considerations] and [I-D.pignataro-enviro-sustainability-architecture] and associated discussions in the IETF, IRTF, and IAB groups. We acknowledge and appreciate the many contributors whose work has enhanced its development.

Authors’ Addresses

Michael Welzl
University of Oslo
Email: michawe@ifi.uio.no

Emile Stephan
Orange
Email: emile.stephan@orange.com

Eve Schooler
University of Oxford
Email: eve.schooler@gmail.com

Sebastien Rumley
HES-SO
Email: sebastien.rumley@hes-so.ch

Ali Rezaki
Nokia
Email: ali.rezaki@nokia.com

Jukka Manner
Aalto University
Email: jukka.manner@aalto.fi

Carlos Pignataro
Blue Fern Consulting
Email: cpignata@gmail.com

Marisol Palmero
Cisco
Email: mpalmero@cisco.com

Jan Lindblad
All For Eco
Email: jan.lindblad+ietf@for.eco

Suresh Krishnan
Cisco
Email: sureshk@cisco.com

Ari Kernen
Ericsson
Email: ari.keranen@ericsson.com

Hesham ElBakoury
Email: helbakoury@gmail.com

Luis M. Contreras
Telefonica
Email: contreras.ietf@gmail.com

Alexander Clemm
Independent
Email: ludwig@clemm.org

Jari Arkko
Ericsson
Email: jari.arkko@gmail.com