Workgroup: INTAREA Working Group
Internet-Draft:
draft-bryant-arch-fwd-layer-uc-04
Published: 6 August 2022
Intended Status: Informational
Expires: 7 February 2023
Authors: S. Bryant
University of Surrey 5/6GIC
U. Chunduri
Intel
T. Eckert
Futurewei Technologies Inc.
A. Clemm
Futurewei Technologies Inc.

Forwarding Layer Use Cases

Abstract

This document considers the new and emerging use cases for IP. These use cases are difficult to address with IP in its current format and demonstrate the need to evolve the protocol.

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 7 February 2023.

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

There is an emerging set of new requirements that exceed the network and transport services of the current Internet, which currently only delivers "best effort" service. While many controlled or private networks include further services, such as other DiffServ QoS in addition to best effort and traffic engineering with bandwidth guarantees, the solutions used today only support walled gardens and are thus they are not available to application service providers and consumers across the Internet.

The purpose of this document is to look at current, evolving and future use cases that need to addressed by the Internet forwarding layer. In parallel with this use case study, a study of the gaps between the capability of the existing IP forwarding layer and the requirements described in this use case study is provided in [I-D.bryant-arch-fwd-layer-ps]. It is thus the purpose of this text to provide the wider context for the forwarding layer problem statement.

The purpose of this text is thus to stimulate discussion on the emerging contexts in which the forwarding layer will need to operate in the future.

1.1. Forwarding Layer

The term "forwarding layer" is used in this document to indicate that that development work will likely need to reach down to layer 2.5 in order to ensure that packets are handled correctly down to the physical layer, and that it is equally it is possible that development work will need to reach into the transport layer. This is described in more detail in [I-D.bryant-arch-fwd-layer-ps].

2. New Use Cases for packet networks

This section summarizes the use case areas that have been observed by the authors, and are considered relevant to any analysis of the gaps in forwarding layer capabilities.

This section is structured into sub-sections discussing either group of use cases directly or the work of specific groups that are identifying use cases and that may also work on identifying issues and or proposed architectures or solutions for them.

Subsections are ordered from what might be considered to be the most near-term use cases to the potentially most far reaching ones.

2.1. Role of Fixed Networks in 5G and Beyond 5G

The 5G and beyond 5G (B5G) services are not meant to be limited to the 5G-NR (new-radio). In fact for those services relating to uRLLC, and mMTC packet networks have evolve along with the radio technologies. While 5G-NR protocol stack has evolved to provide per-frame reliability and latency guarantees, the IP/MPLS transport network by-and-large remains best-effort. It is no longer possible to solve network problems simply by increasing the capacity [SysArch5G]. The expectations 5G devices have of 5G networks, can not be met without improving IP/MPLS based back-haul networks. For example, the 5G based systems involve machine to machine communications, generally using command-based smaller payloads. In this case the overheads of packet headers and overlays become apparent when computing latency budget of such packets.

The IETF has produced a large body of work on the deterministic needs of network applications [RFC8578]. These range from refinements and expansions of above summarized Audio/Video and AR/VR use cases over gaming into many more "industrial" use cases. Industrial use cases generally involve industrial controllers for high-precision machinery and equipment, such as robotic arms, centrifuges, or manufacturing equipment for the assembly of electronic components.
These use cases have in common that they require delivery of packets with very precise and "deterministic" performance characteristics, as the controlled equipment and the control loops involved have very exact timing requirements and are not tolerant of any latency variations, as otherwise control loop issues and other undesired effects may occur.
Specifically, the use cases involve curtailing maximum latency that could be incurred. However, deterministic networking, by itself, does not appear to be sufficient to meet all of the emerging needs.

2.2. Convergence of Industrial Control Networks

Industrial control networks exist to serve specialist applications and are deployed in well controlled networks subject to tight timing and reliability constrains and tight security constraints. They mostly use bespoke, application specific proprietary protocols. There is a desire to achieve economy of scale by using a single

protocol, and to integrate the production network with the back-office network. The obvious protocol to use would be IP, but to be deployed in this mixed application environment IP needs to satisfy the non-negotiable needs of the industrial control network such as timing, reliability and security.

2.3. Cloud Based Industrial Automation

Future industrial networks are significantly different from best effort networks in terms of performance and reliability requirements. This is discussed in [NET2030SubG1]. These networks need more than basic connectivity between the back office and the factory floors, instead they require integration from devices all the way through to the business systems. This permits many new types of UI and full automatic operation and control of industrial processes without significant human intervention. These networks need to deliver better than best effort performance, and require real-time, secure, and reliable factory-wide connectivity, as well as inter-factory connectivity at large scale.

Such systems typically require low end-to-end latency to meet closed loop control requirements. Such system also need low jitter connectivity. IIoT systems, as an example contain many control sub-systems that run at cycle times ranging from sub-ms to 10 ms. In such systems, end-to-end control requires in-time signaling delay at the same cycle time level, without malfunctions. These low latency requirements of IIoT applications are increasingly not only relevant to internal system communications, but also becoming essential for the interconnection of remote systems.

As another example, it is a fundamental requirement for multiple-axis applications to have time synchronization in order to permit cooperation between various devices, sometimes remotely. In order to recover the clock signal and reach precise time synchronization, the machine control, especially the motion control sub-system, requires very small jitter at sub- microsecond level, and such small jitter is expected to have bounded limits under some critical situations.

In some IIOT systems a service availability of 99.999999% is needed, as any break in communications may be reflected as a million-dollar loss. At the same time, as part of the Industry 4.0 evolution, operational technologies (OT) and information technology (IT) are converging. In this model control functions traditionally carried out by customized hardware platforms, such as Programmable Logic Controllers (PLC), have been slowly virtualized and moved onto the edge or into the cloud in order to reduce the CAPEX and OPEX, and to provide increased system flexibility and capability and to allow 'big data' approaches. This move of industrial system to the cloud places higher requirements on the underlying networks, as the

latency, jitter, security and reliability requirements previously needed locally have to be implemented at larger scales.

2.4. Volumetric Data Transmission

Volumetric Data refers to cases where very large data sets need to be transferred continuously in real time. One example is Immersive AR/VR media transmission Section 2.6. Another example is V2X with many sensors continuously generating data which needs to made available for, amongst other reasons, technical analysis by the manufacturer as part of product development, and insurance purposes.

2.5. ITU-T Focus Group Network-2030

The ITU-T has been running a Focus Group (FG) Network-2030 [FGNETWORK2030] to analyze the needs of networks in the period post 2030. This work started in July 2018 and submitted it report to ITU-T Study Group 13 in June 2020. It has been an open process with contribution by a cross-section of the networking industry. Because this is non-IETF work, this section summarizes the currently finalized key findings of the ITU-T Focus Group Network-2030 to make it easier for the reader to better understand the work. Note that this work is still ongoing and additional findings may be published.

The Focus Group Network 2030 considered a number of use cases that it was postulated would need to be addressed in the 2030 time-frame and the technology gaps that need to be bridged in order to address these needs. It then considered a number of new network services that would be needed to support these services.

An ongoing piece of work on the architecture of the network post 2030 has not yet been completed at the time of writing and is only partially discussed in this document.

The reader is referred to [WP], [NET2030SubG2], [UC] for information beyond that provided in this summary.

ITU-T FG NET2030 Sub-group Sub-G1 (Sub-G1) considered a number of use cases that it considered to be representative of the network needs post 2030. These needs are legitimate needs in their own right, but as is always the case act as poster-children for new applications that will inevitable conceived in the light of the new network capabilities that we postulate to be necessary.

*Holographic-type communications (HCT)

*Tactile Internet for Remote Operations (TIRO)

*Network and Computing Convergence (NCC)

*Digital Twin (DT)

*Space-Terrestrial Integrated Networks (STIN)

*ManyNets

*Industrial IoT (IIoT) with cloudification.

Further information on these use cases is provided in Section 7, and in the ITU documents [UC] and [WP].

Note to the reader: Unlike ITU-T Study Groups which are restricted to members, ITU-T Focus Groups are open to anyone without payment. At the time of writing, ITU-T Focus Group Network-2030 material that is not available for anonymous download, is accessible for free by joining the Study Group.

2.6. Emerging and New Media Applications

Audio/Video streaming for production, entertainment, remote observation, and interactive audio/video are the most ubiquitous applications on the Internet and private IP networks after web-services. They have grown primarily through an evolution of the applications to work with the constraints of todays Internet and adopting pre-existing infrastructure such as content caches: best-effort streaming with adaptive video, no service guarantees for most services, and co-location of caches with large user communities. In environments where more than best-effort services for these applications are required and deployment of current technologies to support them is feasible, it is done. Examples include DiffServ or even on or off-path bandwidth reservations in controlled networks.

Networked AR/VR is a very near term set of use cases, where solution models are very much attempting to use and expand existing solution approaches for video network streaming but where the limits of above current best practices are also amplified by the larger bandwidth requirements and stricter latency and jitter requirements of AR/VR.

To ensure a good user experience, for live Virtual Reality (VR), a much higher resolution than 8K video is required. In addition to the high bandwidth requirements of VR, there needs to be a supporting transmission network to provide a communications path with bounded low latency as well. This stringent VR latency requirement is a challenge to existing networks.

In cellular networks, even though the the air interface link latency needed is significantly reduced e.g. with New Radio (5GNR), the end-to-end (E2E) requirements for live VR is harder to meet. This is because of the fixed L2/IP/MPLS networks in front/mid/backhaul

components, and because of the best effort nature of the packet delivery systems in these networks.

3. Deployment Models

In this section we look at a number of network deployment models. We group these deployment models into three types:

- *The traditional deployment models
- *Emerging deployment model models
- *Envisioned new deployment models

The service requirements demanded from the networks and security implications vastly differ in these different deployment models.

A few general observations are useful in providing context to this section:

- *End to End traffic over the Internet backbone is becoming minority traffic.
- *Commercial deployments do not operate the way they used to when many of the original Internet protocols and invariants were established.
- *The application trajectory is for the applications to be hosted on (protected) servers a few hops from the user.
- *Applications are becoming self-contained and use their own stack which is tunneled over UDP/IP to the server.

3.1. Traditional Deployment Models

In this section we look at the traditional deployment models that have been in existence for many year and formed the foundation of Internet.

3.1.1. Best-effort Internet

In this model connectivity is edge-to-edge, and in the general case the edge connectivity is provided by a service provider who peers with a transit provider that provides connectivity to other service providers possibly via other transit providers. This is shown in Figure 1.

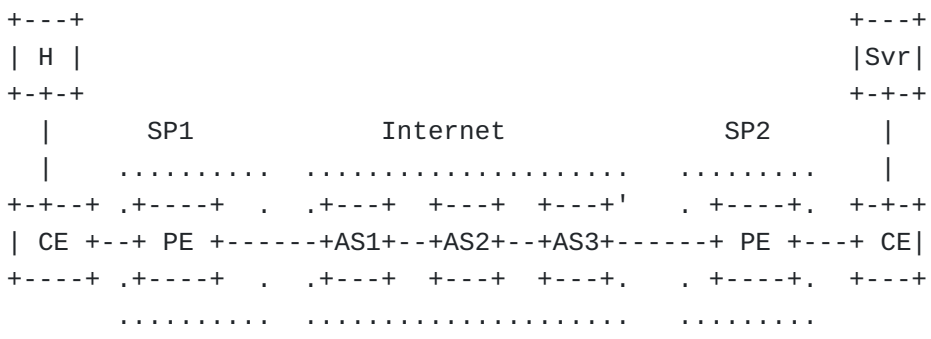

Figure 1: An Edge-2-Edge Classical Internet

This service is generally known as "best-effort" in that each element of the service path undertakes to do no more that try its best to provide equitable service to all traffic. These are traditional E2E deployments where communication endpoints of the data traffic on different provider networks with regional, transit network providers through Internet Exchange Providers (IXPs) providing the global inter connection. The term lower-common-denominator might be a better term in that the service quality is the service of the worst element of the path on a packet by packet basis.

This model is in the process of being replaced by a model in which the most popular and important service are provided at the edge with Internet transit traffic being used where there is no alternative.

In this case the provider controls only the path to the CE and can certify the correct operation of the service according to contract from that point but the user is responsible for providing the required service characteristics into their own network.

In this network environment it is difficult to support any form of enhanced service since it is unlikely that the whole path is know to support extended capabilities in the forwarding plane. It is not infeasible, and it would be possible to set up such paths in principle given suitable enhancements to the routing system. However such a scenario must be considered infeasible for the foreseeable future.

3.1.2. Enhanced Service

This is the traditional service provider deployment where various network services (VPN, security, Bandwidth..) are offered to the endpoints of the communication and other providers. Such capabilities are purchased through contract with the service provider and are typically expensive.

These networks predominantly use MPLS technology though native IP (IPv4/IPv6) with GRE and IPv6 with routing extension headers with SRv6 are being deployed recently.

Figure 2: An Edge-2-Edge Network

In this case there is a single provider network in which E2E offerings and host session are initiated and terminated with in the single provider network.

3.1.3. Over-the-top (OTT) Providers

In this model the endpoints of the communication (virtual or physical hosts) consuming services through with in the OTT provider network servers (Cloud and Data Center (DC) networks); where the other endpoint can be in the same server form or on the DC Gateway or on the other end of the DC Server Farm connected through Data Center Interconnect (DCI).

The local provider is thus just a connectivity provider to opaque traffic with no ability to enhance the service. However the corollary to this is that whilst the the OTT provider has full control of what happens whilst the user data is within their network they have no control over how the user traffic transits to them across the "public" network.

3.1.4. Cooperating Providers

Where two providers interconnect with no Internet Transit Network: Another variant of the E2E connectivity can be seen as evolving comprises only endpoints provider (access) network and receiver access provider network with global transit provided by one ISP. This case is more tractable provided there is co-operation between the providers.

3.2. Emerging Deployment Models

The emerging model is to provide the service close to the user by embedding that service with the service provider network. This has three advantages, firstly that the service latency is lower, secondly that that there is less transit traffic that the network provider needs to manage or pay for, and thirdly that the service

availability and reliability is in the hands of the network provider that the customer is contracted to.

3.2.1. Embedded Service

The industry move is towards content and application service providers embedding themselves within the edge network. This is currently done to save bandwidth and improve response time. As the need for high precision low latency networking develops the need for edge computing rises since the closer the client and the server the less the scope for network induced performance degradation.

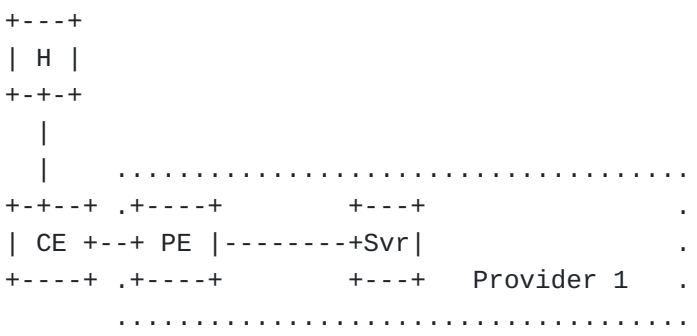

Figure 3: An Edge-2-Provider

In this network the server S (owned by the content and applications provider) has a contractual relationship with provider 1 and is thus able to negotiate the network characteristics needed to meet its service requirement. This model in which the server brokers the user to network interface (UNI) requirements removes many of the objections to the classical UNI model in which the client requests the service requirements. In this model the host authenticates itself with the server, having formed a previous business relationship (for example by purchasing a holographic conferencing service). The server has a relationship with Provider1, and thus is a trusted party able to request that the service be set up between itself and and its client, paying as necessary. As this is a requested paid service traversing a limited distance over a defined network, a bespoke packet protocol can, if necessary, be used with in a contained and constrained way.

How the server communicates with any other part of the application domain is out of scope for this document and possibly out of scope for Provider 1.

This takes us to consider the embedded global service described in {#EGS}.

3.2.2. Embedded Global Service

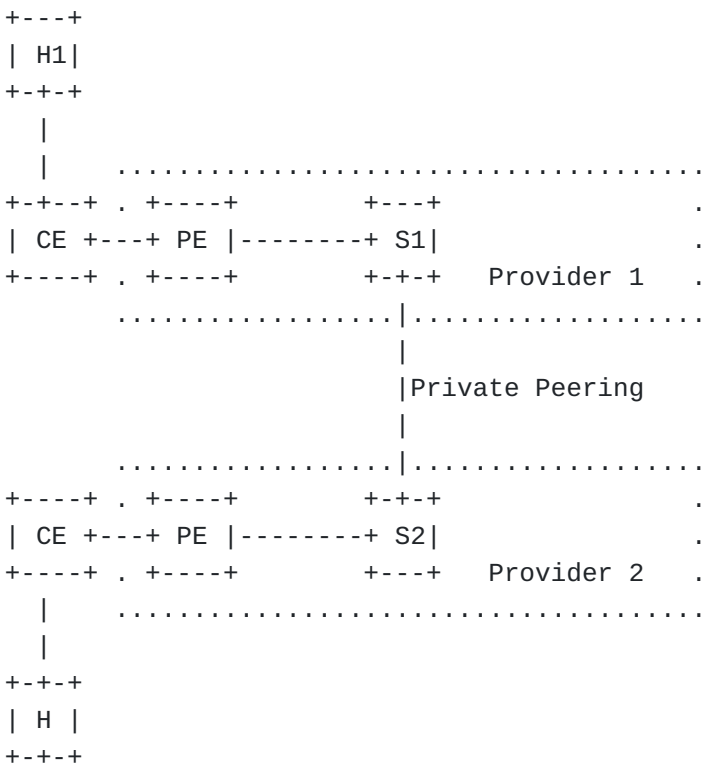

Figure 4: Edge-2-Edge via Provider

In this network model, the server S1 (owned by the content and applications provider) has a contractual relationship with provider 1 and is thus able to negotiate the network characteristics needed to meet its service requirement. It is servicing the needs of host H1.

Similarly that same provider has a contractual relationship with provider 2 where it is servicing the needs of host H2.

By a method outside the scope of this document and outside the scope of the global Internet the contents and applications provider has a private path between S1 and S2.

This scenario shown in Figure 4 is important because it removes the overwhelming issues associated with providing enhanced service across the global Internet. Furthermore it describes a model where there is commercial incentive, at scale, for the edge providers (Provider 1 and 2 above) to invest in providing and enhanced access service.

3.2.3. Changing Fixed Access Models (1 or 2 Providers)

The preceding sections are the basis for a change in the network fixed access model.

The access network either connects to a data center gateway or one is embedded in the access network. This gateway either passes the traffic to a locally connected data center that provides the required service or passes it over a private global data center interconnect to a partner data center for service provision. Such a connection provides service model in which the required service level cane be more readily addressed.

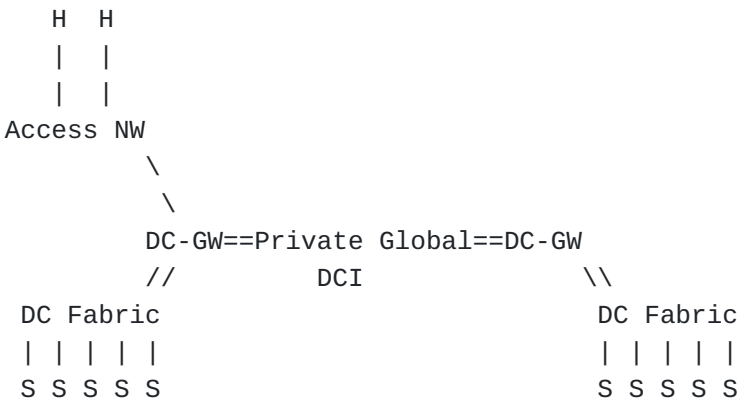

Figure 5: Changing Fixed Access Model

3.2.4. Single "Underlay" provider E2E for 5G/B5G network (Cellular/ Access Networks)

The preceding sections are the basis for the emerging change in the structure of the 5B and Beyond 5G (B5G) network design.

Endpoints (UE's) connecting to the provider wireless or wired networks, where service is terminated inside the provider network end points. Based on the service offerings connection termination can happen close to the Radio/access nodes with multi-access edge computing (MEC) clouds or in the provider core network (core-cloud) before going to the Internet eventually. Example of these deployments include BNG, 4G and 5G wireless access/RAN/backhaul networks.

Thus in Figure 6 user equipment connects to the customer site provider edge via the radio network. This in turn is connected to the aggregation PE which in turn determines if the traffic should be routed to a local data center for processing, or passed to a core data center. At the core DC the traffic may be processed locally, passed out to the Internet, passed to a peer DC via a private DCI, or processed locally with the help of resources access via that external interconnects.

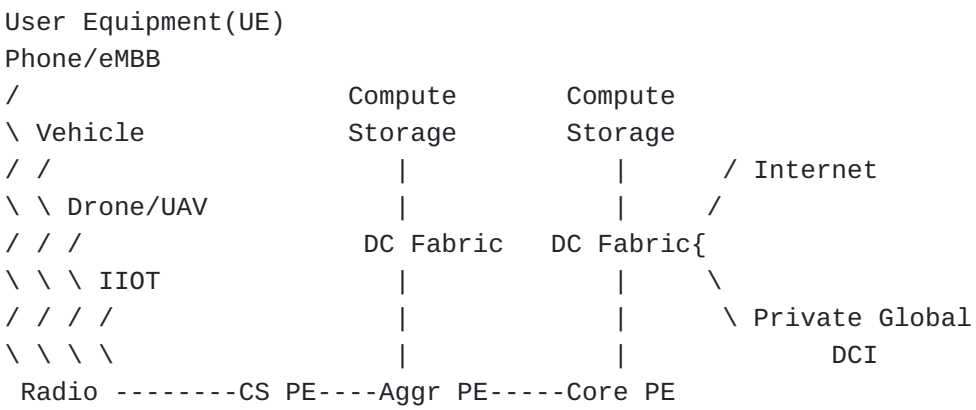

Figure 6: 4G and 5G underlay provider network

3.3. Envisioned New Deployment Models

The emerging network deployment models are a potential vector for fundamental change in the way the network operates.

3.3.1. Network Slicing

Network slicing is a method of creating a private subset of a public network. Unlike VPNs it is not a simple over the top approach, instead it is more integrated with the base network in terms of the way the base network provides services and allocates resources. A network slice provides significant isolation between one slice and another and between the slice and best effort users of the network. In an ideal slice, the users of one slice have no way of knowing anything about the traffic in any other slice. Such a service could be offered through statistical multiplexing techniques with real bandwidth permanently allocated to each slice, but this would not easily offer the statistical multiplexing that make packet networking so economic and so flexible. In particular it would not be easy to transparently "borrow" unused committed bandwidth in a way that was undetectable. It seems likely that to create a high fidelity slice will require new properties in the packet layer, either through extension of the existing packet protocols, or through the introduction of an alternative design. A useful discussion of network slicing relevant to this context can be found in [I-D.ietf-teas-enhanced-vpn].

Largely popularized as part of 5G the concept of network slicing has wider applicability.

3.3.2. Private 5G Networks

A use case is emerging for 5G technology in private networks. The interest is in the protection and security that comes with the use of licensed spectrum. Unlicensed spectrum offers no protection against other users of that spectrum and thus another aspect of best

effort comes into play, not only is the network best effort with respect to traffic within the network (an addressable problem) but the radio is best effort with respect to radio traffic from adjacent networks. Without extensive radio shielding of the facility a user cannot know if the spectrum is available for their use at any time, and they have to suffer interference from adjacent users, who may be benignly using the spectrum for legitimate purposes, as is their equal right, or may be using it to cause service disruption to a commercial enterprise.

5G runs on licensed and hence protected spectrum. In return for the paying the license fee the spectrum owner has a statutory protection against interference.

Thus it is interesting to note that a major UK car plant just announced the use of 5G to provide connectivity for equipment at their manufacturing facility.

Such applications of 5G are not as architecturally constrained as public 5G deployments and thus have the ability to make different fundamental choices regarding their packet protocols.

3.4. Limited Domains

[RFC8799] provides a useful insight into the emergence of limited domains in which fewer (or different) constraints on protocol design and operation apply. Limited domains offer an opportunity to deploy specialist forwarding layer protocols, designed to meet specific objectives, which are not readily addressed by general purpose protocols such as IPv4|6 without the need to worry about inter-working and inter-operation across the big I Internet.

Such domains can be considered sandboxes in which new proposals can be deployed without the wider concerns of full-scale Internet deployment.

4. New Network Services and Capabilities

In order to support the use cases presented in Section 2, a number of new network services will be needed. Likewise, a number of additional more general network capabilities will becoming increasingly important. Neither services nor capabilities are sufficiently supported to the degree that will be required by Internet technology in use today.

This section describes these services and capabilities at a high level. It builds on a corresponding analysis that was conducted at ITU-T FG-NET2030; readers are referred Section 8 for further detail and, of course, to output produced by that group [NET2030SubG2] for a more complete explanation of their considerations.

4.1. New Services

[NET2030SubG2] identifies a number of network services that will be needed to support many of the new use cases. These network services are divided into two categories:

- Foundational Services (FS) require which dedicated support on some or all network system nodes which are delivering the service between two or more application system nodes.

- Compound Services (CS) are composed of one or more foundational services, and are used to make network services easier to consume by certain applications or categories of use cases. An example of a CS would be a Tactile Internet Service which consisted of tactile control channel and a haptic feedback channel.

The following are a set of Foundational Services :

- High-Precision Communications Services: services with precisely defined service level objectives related to end-to-end latency. Three high-precision communications services that have so far been proposed:

  - In-time Services: services that require end-to-end latency within a quantifiable limit. This service is similar to the service provided by DetNet [RFC8655] but with more demanding applications which need to be satisfied over IP.

  - On-time Services: services require end-to-end-latency to be of an exact duration.

  - Coordinated Services: Coordinated services require multiple interdependent flows to be delivered with the same end-to-end latency, regardless of any (potential additional) service level objective.

- Qualitative Communication Services: services that are able to suppress retransmission of less relevant portions of the payload in order to meet requirements on latency by applications that are tolerant to this.

These are described in more detain in Section 8.1.

4.2. New Capabilities

[NET2030SubG2] identifies also a number of network capabilities that will become increasingly important going forward, in addition to the support for any particular services.
A number of those need to be taken into consideration from the very

beginning when thinking about how future data-planes need to evolve. These capabilities are described in more detail in Section 8.2.

*Manageability: Many of the services that need to be supported in the future will require advances in measurements and telemetry will be required in order to monitor and validate that promised service levels are indeed being delivered. These will requires advanced instrumentation that is ideally built.

*High Programmability and Agile Life-cycle: Methods to provide operators need to be able to rapidly nd easily introduce new network services and adapt to new contexts and application needs.

*Security and Trustworthiness: New mechanisms are needed to authorize packets to enter the network from a host or from another network, and for them to then receive the required premium service that can operate. This must operate without impacting the latency and MTU requirements. This security mechanism has to protect both the network, the user data and the user privacy, but still expose sufficient information to the network that the correct premium service can be delivered.

*Resilience: Ultra-low-latency requirements and the huge increase of bandwidth demands of new services such as holographic type communication services make retransmission as a mechanism to recover data that was lost in transit increasingly less feasible. Therefore, network resilience and avoidance of loss becomes more importance that it is for best effort networks.

*Privacy-Sensitive: There is a growing awareness of the lack of privacy in the Internet and its implications.
New network services have to be sensitive to and comply with heightened user privacy expectations.
At the same time, the need for privacy needs to be balanced with legitimate needs of network providers to operate and maintain their networks, which requires some visibility into what is happening on the network and how it is being used. There are a variety of privacy-related requirements that ensue, such as:

  -Anonymization

  -Opaque User data

  -Secured Storage

  -Flow anonymization

*Accountability and Verifiability: Provision of the methods to account for an verify delivery of premium services.

5. IANA Considerations

This document does not request any allocations from IANA.

6. Security Considerations

Security is likely to be more significant with the applications being considered in this work. With interest in tightly controlled access and latency, and contractual terms of business it is going to be necessary to have provable right of access to network resources. However heavyweight security is a contra-requirement to the light-weight process needed for power efficiency, fast forwarding and low latency. Addressing this will require new insights into network security.

Further information on the issue of providing security in latency sensitive environments can be found in [RFC9055] which are a sub-set of the considerations applicable to the new use cases considered in this text.

7. Appendix 1: Expanded Summary of Sub-G1 Use Cases

7.1. Holographic-type communications

This work projects that we will move towards a holographic society where users remotely interact with the physical world over the network. In industry the digital twin model will enable the control of real objects through digital replicas. Tele-presence will move to a new level with multi-site collaborations becoming much closer to physical meetings that can take place without the time and environmental cost of physical travel. 3D medical scans will become full 3D views rather than the body/ organ slices that too many of us are regrettably familiar with. It is easy to imagine that this technology will take message delivery to a completely new level.

Analysis of these concepts results in the conclusion that the following key network requirements are necessary:

*Ultra-high bandwidth (BPS class)

*Ultra-low latency (sub-ms)

*Multi-stream synchronization

*Enhanced network security

*Enhanced network reliability

*Edge computation

7.2. Tactile Internet for Remote Operations

Two cases were proposed as examples of this class of application. The first is remote industrial management which involves the real-time monitoring and control of industrial infrastructure operations. The second involves remote robotic surgery. Remote robotic surgery within an operating suite complex is a standard practice today, however there are cases where it would be desirable to extend the range of this facility.

Analysis of these concepts results in the conclusion that the following key network requirements are necessary:

*Ultra-high bandwidth (Tbps class)

*Ultra-low latency (sub-ms)

*Sensory input synchronization

*Enhanced network security

*Enhanced network reliability

*Differentiated prioritization levels

7.3. Space-Terrestrial Integrated Networks

The game-changer in the area of space-terrestrial networking is the active deployment of huge clusters of cheap Low Earth Orbit (LEO) satellite constellations. These LEOs have a number of properties that make them attractive, but arguably the most important is that they combine global coverage with low latency. Studies [Handley] show that for distances over 3000Km latency via a LEO cluster is lower than the latency of terrestrial networks. The up-link to a LEO cluster has to constantly change the point of attachment to the cluster as the satellites that form the cluster rapidly move across the sky relative to both the ground and relative to the satellites in other orbits. In this scenario a number of access and connection models need to be considered.

Analysis of these concepts results in the conclusion that the following key network requirements are necessary:

*A suitable addressing and routing mechanism to deal with a network that is constantly in flux.

*Sufficient bandwidth capacity on the satellite side to support the new application needs

*A suitable satellite admission system

*Edge computation and storage

7.4. ManyNets

There is evidence that there is a change in direction from the Internet as a single hetrogenious network back to a true Internet, that is an interconnection of a number of networks each optimized for its local use but capable of inter-working.

For example, satellite and the terrestrial networks adopt different protocol architecture, which causes the difficulty to internetwork between them, yet the common goal is to provide access to the Internet. Secondly, there will be a massive number of IoT-type devices connecting to the networks but the current interconnection schemes are too complex for these services. There are further trends in 5G/B5G back-haul infrastructure, requiring diverse set of resource guarantees in networks to support different industry verticals. The collection of such special purpose networks, existing together and requiring interconnection among themselves are called ManyNets.

Much closer the traditional Internet model is the move to edge computing services in which the client traffic is terminated at a compute node very close to access edge. [DOT] Any resultant application traffic is a private matter between the application on the edge server and the servers it communicates with in the fulfillment of those needs. Furthermore the application on the client may be using a tunnel to the edge compute server. In such a network the protocol used inside the tunnel and the protocol used between the servers executing the service is a private matter.

The ManyNets concept aims to support flexible methods to support the communication among such heterogeneous devices and their networks.

8. Appendix 2: Expanded Summary of Sub-G2 New Network Capabilities and Services

This appendix expands on the ITU-T Sub-G2 new network capabilities and services introduced in Section 4 It builds upon the analysis that was conducted at ITU-T FG-NET2030; readers are also referred to output produced by that group [NET2030SubG2] for more detail.

8.1. New Services

[NET2030SubG2] identifies a number of network services that will be needed to support many of the new use cases. These network services are divided into two categories:

*Foundational Services (FS) require which dedicated support on some or all network system nodes which are delivering the service

between two or more application system nodes. FS cannot be decomposed into other services. For example, IP packet routing and forwarding are is a (pre-existing) foundational network services.

*Compound Services (CS) are composed of one or more foundational services. CS are "convenience services" that make network services easier to consume by certain applications or categories of use cases, but do not by themselves introduce new network services or requirements into network system nodes. One example would be a Tactile Internet Service which consists of two communications channels, one for tactile control and the other for haptic feedback.

The following sections focus on Foundational Services only, as these are the ones that provide the basic building blocks with which the needs of all other services can be addressed, and which are the ones that potentially introduce new foundational requirements on network system nodes.

8.1.1. High-Precision Communications Services

High-Precision Communications Services refers to services that have precisely defined service level objectives related to end-to-end latency, in many cases coupled with stringent requirements regarding to packet loss and to bandwidth needs. These requirements are in stark contrast to the best effort nature with related to existing network services.
Of course, existing services often go to great lengths in order to optimize service levels and minimize latency, and QoS techniques aim to mitigate adverse effects of e.g. congestion by applying various forms of prioritization and admission control. However, fundamentally all of these techniques still constitute patches that, while alleviating the symptoms of the underlying best-effort nature, do not address the underlying cause and fall short of providing service level guarantees that will not be just of a statistical nature but that will be met by design.

The high-precision communications services that have been identified are described in the following three sub-sections.

8.1.2. In-time Services

In-time services require end-to-end latency within a quantifiable limit. They specific a service level objective that is not to be exceeded, such as a maximum acceptable latency (putting a hard boundary on the worst case). In-time services are required by applications and use cases that have clear bounds on acceptable latency, beyond which the Quality of Experience would deteriorate

rapidly, rendering the application unusable. An example concerns use cases that involve providing tactile feedback to users. Creating an illusion of touch requires a control loop with a hard-bounded round-trip time that is determined by human / biological factors, beyond which the sense of touch is lost and with it the ability to e.g. operate a piece of machinery from remote. Because many such use cases are mission-critical (such as tele-driving or remote surgery), in addition any loss or need for retransmission is unacceptable.

This service is similar to the service provided by DetNet [RFC8655] but with more demanding applications which need to be satisfied over IP.

8.1.3. On-time Services

On-time services require end-to-end-latency to be of an exact duration, with the possibility of a small quantifiable variance as can be specified either by an acceptable window around the targeted latency or by a lower bound in addition to an upper bound. Examples of use cases include applications that require synchronization between multiple flows that have the same in-time latency target, or applications requiring fairness between multiple participants regardless of path lengths, such as gaming or market exchanges when required by regulatory authorities. The concept of a lowest acceptable latency imposes new requirements on networks to potentially slow down packets by buffering or other means, which introduces challenges due to high data rates and the cost e.g. of associated memory.

8.1.4. Coordinated Services

Coordinated services require multiple interdependent flows to be delivered with the same end-to-end latency, regardless of any (potential additional) service level objective. Use cases and applications include applications that require synchronization between multiple flows, such as use cases involving data streams from multiple cameras and telemetry sources. In the special case where an on-time service is required, no additional service is needed (as synchronization occurs by virtue of the fact that each flow adheres to the same SLO), but coordination may also be required in cases where no specific end-to-end latency is required, as long as all flows are serviced with service levels that are identical.

8.1.5. Qualitative Communication Services

Qualitative communication services (QCS) are able to suppress retransmission of portions of the payload that are deemed less relevant when necessary in order to meet requirements on latency by

applications that are tolerant of certain quality degradation. They may involve the application of network coding schemes.

QCS is a new service type that is needed to support AR/VR, holographic-type communications Industrial Internet and services such as autonomous driving. This needs the support of a new network capability that is as yet to be developed.

8.2. New Capabilities

[NET2030SubG2] identifies also a number of network capabilities that will become increasingly important going forward, in addition to the support for any particular services. These were introduced in Section 4.2. A number of these capabilities need to be taken into consideration from the very beginning when thinking about how future data-planes need to evolve.
While many of those capabilities are well known, the past has shown that retrofitting data-planes with such capabilities after the fact and in a way that is adequate to the problem at hand is very hard.

8.2.1. Manage ability

Many of the services that need to be supported in the future have in common that they place very high demands on latency and precision that need to be supported at very high scales, coupled with expectations of zero packet loss and much higher availability than today.

In order to assure in-time and on-time services with high levels of accuracy, advances in measurements and telemetry will be required in order to monitor and validate that promised service levels are indeed being delivered. This requires advanced instrumentation that is ideally built-in all the way to the protocol level.

For example, the ability to identify and automatically eliminate potential sources of service-level degradations and fluctuations will become of increasing importance. This requires the ability to generate corresponding telemetry data and the ability to observe the performance of packets as they traverse the network. Some of the challenges that need to be addressed include the very high volume of data that gets generated and needs to be assessed, and the effects of the collection itself on performance. In general, greater emphasis will need to be placed on the ability to monitor, observe, and validate packet performance and behavior than is the case today. For seamless support, these capabilities will be inherently integrated with the forwarding function itself, for example delivered together with the packets. Today's solution approach, IOAM, is a promising technology currently that points in the right direction, and that also highlights some of the challenges - from

MTU considerations due to extending packet sizes to the ability to customize and obtain the "right" data. It will therefore be not sufficient by itself. Data to be generated from the network will need to be "smarter", i.e. more insightful and action-able. This will require additional abilities to process data "on-device". In additional, the need for new management functions may arise, such as functions that allow to validate adherence with agreed-upon service levels for a flow as a whole, and to prevent data or privacy leakage as well as provide evidence for the possibility or absence of such leakage.

8.2.2. High Programmability and Agile Life-cycle

Operators need to be able to rapidly introduce new network services and adapt to new contexts and application needs. This will require advances in network programmability. Today's model of vendor-defined (supporting service features via new firmware or hardware-based networking features) or operator-defined (supporting service features via programmable software-defined networking (SDN) controllers, virtualized network functions (VNF) and Network Function Virtualization (NFV), and service function chaining (SFC) will no longer be sufficient.

Software Defined Networking and Network Function Virtualization (NFV) have opened up the possibility to accelerate development life-cycles and enable network providers to develop new networking features on their own if needed. Segment Routing is being evolved for that purpose as well. Furthermore, network slicing promises more agility in the introduction of new network services. However, the complexity of the associated controller software results in its own challenges with software development cycles that, while more agile than life-cycles before, are still prohibitive and that can only be undertaken by network providers, not by their customers. Rapid customization of networking services for specific needs or adaptation to unique deployments are out of reach for network provider customers. What is lacking is the ability for applications to rapidly introduce and customize novel behavior at the network flow level, without need to introduce application-level over-the-top (OTT) overlays. Such a capability would be analogous to server-less computing that is revolutionizing cloud services today. In addition, it should be noted that softwarized networks are built on relatively stable (and slowly evolving) underlying physical commodity hardware network infrastructure. This is insufficient to deliver on new high-precision network services, which require hardware advances at many levels to provide programmable flow and QoS behavior at line rate, affecting everything from queuing and scheduling to packet processing pipelines.

The evolution of forwarding planes should allow development life-cycles that are much more agile than today and move from "Dev Ops" to "Flow Ops" (i.e. dynamic programmability of networks at the flow level).
This requires support of novel network and data-plane programming models which can possibly be delivered and effected via the forwarding plane itself.

8.2.3. Security

The possibility of security threats increases with complexity of networks, the potential ramifications of attacks are growing more serious with increasing mission-criticality of networking services and applications.
The forwarding plane plays a large role in the ability to thwart attacks.
For example, the fact that source addresses are not authenticated in existing IP is at the root of a wide range security problems from phishing and fraudulent impersonation designed to compromise and steal user assets to amplification attacks designed to bring down services.
Going forward, it is absolutely critical, then, to minimize the attack surface of the forwarding plane as it evolves.

A key security aspects needed from the network point of view includes to verify if the packet is authorized to enter into the network and if it is sufficiently integrity protected. However, when packets are emitted from the host for these new communication services, the network portion of the packet (e.g., an extension header or an overlay header) should not be encrypted because network nodes may need to interpret the header and provide the desired service.
Lack of encryption and integrity validation, of course, would at the same time increase the threat surface and open up the possibility for attacks.
Mechanisms for authorization and integrity protection must be developed to meet the line rate performance as services delivered can be time sensitive. At the same time, the size of packets should not be significantly increased to avoid negative impact on utilization and overhead tax.
This limits the options for additional security collateral that can be included with packets.

Homomorphic forms of encryption may need to be devised in which network operations can be performed in privacy-preserving manner on encrypted packet headers and tunneled packets without exposing any of their contents.

Another dimension to security arises when the end to end service that needs to be delivered crosses the administrative boundary of the originating host. For those cases, additional mechanisms need to be specified to sufficiently ensure the privacy and confidentiality of the network layer information. While there are lot of avenues to tackle these issues and some aspects are being looked into by various Standards Development Organizations, e.g. IRTF PANRG on Path-Aware Networking, comprehensive solutions are yet to be worked out.

Any mechanisms specified for authorization, integrity protection, and network header confidentiality should be orthogonal to the transport layer and above transport layer security mechanisms set in place by the end host/user. Regardless of whether or not the latest security advances in transport and layers above (e.g. TLS1.3, QUIC or HTTPSx) are applied on the payload, network nodes should not have to act on this information to deliver new services to avoid layer violations.

8.2.4. Trustworthiness

As future network services are deployed, deployment scenarios will include cases in which packets need to traverse trust boundaries which are under different administrative domains. As the forwarding plane evolves, it should do so in such a way that trustworthiness of packets is maintained - i.e. integrity of data is protected, tampering with packet meta-data (such as source authentication or service level telemetry) would be evident, and privacy of users is guarded.

8.2.5. Resilience

Ultra-low-latency requirements and the huge increase of bandwidth demands of new services such as holographic type communication services make retransmission as a mechanism to recover data that was lost in transit increasingly less feasible. Therefore, network resilience and avoidance of loss becomes of paramount importance.

There are many methods for providing network resilience. This includes providing redundancy and diversity of both physical (e.g. ports, routers, line cards) and logical (e.g. shapers, policers, classifiers) entities. It also includes the use of protocols that provide quick re-convergence and maintain high availability of existing connections after a failure event occurs in the network. Other techniques include packet replication or network coding and error correction techniques to overcome packet loss.
As the forwarding plane evolves, mechanisms to provide network resilience should be inherently supported.

8.2.6. Privacy-Sensitive

Today, there is a growing awareness of the lack of privacy in the Internet and its implications.
New network services have to be sensitive to and comply with heightened user privacy expectations.
At the same time, the need for privacy needs to be balanced with legitimate needs of network providers to operate and maintain their networks, which requires some visibility into what is happening on the network and how it is being used.
Likewise, mechanisms to provide privacy must be provided in such a way to not compromise security, such as allowing anonymous attackers to prey on other users.

An evolved forwarding plane must provide mechanisms that ensure users privacy by design and prevent illegitimate exposing of personally-identifiable information (PII), while preventing abuse of those mechanisms by attack exploits and while affording network providers with legitimate visibility into use of their network and services.

There are a variety of privacy-related requirements that ensue, such as:

- Anonymization: To prevent tracking by eavesdropper by packet capture, visible information in packets such as source and destination addresses should be difficult (ideally: impossible) to directly correlate to PII.

- Opaque User data: Networks must not rely on the user data to provide or improve the service. However, network providers may use specific service-visible data in packets.

- Secured Storage: Some services may require the network to slow down the delivery of the packets, implying the possibility that packets are temporarily buffered on the router. The storage of those packets must be secured and prevented from extraction for deep inspection or analysis.

- Flow anonymization: Flows of information should be randomized in a dynamic manner so that it is difficult through traffic analysis to deduce patterns and identify the type of traffic.

Potential mechanisms to consider include (but are not limited to) avoiding the need for long-lived addresses (to prevent trackablity) and the use of homomorphic encryption for packet headers and tunneled packets (in addition to traditional payload encryption) that allow to perform network operations in privacy-preserving manner without exposing meta-data carried in headers.

8.2.7. Accountability and Verifiability

Many new services demand guarantees instead of being accepting of "best effort".
As a result, today's "best effort" accounting may no longer be sufficient.

Today's accounting technology largely relies on interface statistics and flow records.
Those statistics and records may not be entirely accurate.
For example, in many cases their generation involves sampling and is thus subject to sampling inaccuracies.
In addition, this data largely accounts for volume but not so much for actual service levels (e.g. latencies, let alone coordination across flows) that are delivered.
Service level measurements can be used to complement other statistics but come with significant overhead and also have various limitations, from sampling to the consumption of network and edge node processing bandwidth.
Techniques that rely on passive measurements are infeasible in many network deployments and hampered by encryption as well as issues relating to privacy.

Guarantees demand their price. This makes it increasingly important both for providers and users of services to be able to validate that promised service levels were delivered on.
For example, proof of service delivery (including proof of service level delivery) may need to be provided to account and charge for network services.
This will require advances in accounting technology that should be considered as forwarding technology evolves, possibly providing accounting as a function that is intrinsically coupled with forwarding itself.

9. Informative References

[DOT] Huston, G., "The Death of Transit and Beyond", n.d., <https://hknog.net/wp-content/uploads/2018/03/01_GeoffHuston_TheDeath_of_Transit_and_Beyond.pdf>.

[FGNETWORK2030] "Focus Group on Technologies for Network 2030", n.d., <https://www.itu.int/en/ITU-T/focusgroups/net2030/Pages/default.aspx>.

[Handley] Handley, M., "Delay is Not an Option: Low Latency Routing in Space", n.d., <http://nrg.cs.ucl.ac.uk/mjh/starlink-draft.pdf>.

[I-D.bryant-arch-fwd-layer-ps] Bryant, S., Chunduri, U., Eckert, T., and A. Clemm, "Forwarding Layer Problem Statement", Work

in Progress, Internet-Draft, draft-bryant-arch-fwd-layer-ps-05, 5 August 2022, <https://www.ietf.org/archive/id/draft-bryant-arch-fwd-layer-ps-05.txt>.

[I-D.ietf-teas-enhanced-vpn] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Network (VPN+) Services", Work in Progress, Internet-Draft, draft-ietf-teas-enhanced-vpn-10, 6 March 2022, <https://www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-10.txt>.

[NET2030SubG1] ITU-T FGNet2030, "FG NET-2030 Sub-G1 Representative use cases and key network requirements for Network 2030", January 2021, <http://handle.itu.int/11.1002/pub/815125f5-en>.

[NET2030SubG2] ITU-T FGNET2030, "New Services and Capabilities for Network 2030: Description, Technical Gap and Performance Target Analysis", October 2019, <https://www.itu.int/en/ITU-T/focusgroups/net2030/Documents/Deliverable_NET2030.pdf>.

[RFC8578] Grossman, E., Ed., "Deterministic Networking Use Cases", RFC 8578, DOI 10.17487/RFC8578, May 2019, <https://www.rfc-editor.org/info/rfc8578>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/info/rfc8655>.

[RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <https://www.rfc-editor.org/info/rfc8799>.

[RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <https://www.rfc-editor.org/info/rfc9055>.

[SysArch5G] "System architecture for the 5G System (5GS)", n.d., <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>.

[UC] ITU-T FGNET2030, "Use Cases and Requirements for Network 2030 Summary report "Representative use cases and key network requirements for Network 2030"", January 2020, <https://www.itu.int/en/ITU-T/focusgroups/net2030/Documents/Technical_Report.pdf>.

[WP]
"Network 2030 - A Blueprint of Technology, Applications, and Market Drivers towards the Year 2030 and Beyond, a White Paper on Network 2030, ITU-T", May 2019, <https://www.itu.int/en/ITU-T/focusgroups/net2030/Documents/White_Paper.pdf>.

Authors' Addresses

Stewart Bryant
University of Surrey 5/6GIC

Email: sb@stewartbryant.com

Uma Chunduri
Intel

Email: umac.ietf@gmail.com

Toerless Eckert
Futurewei Technologies Inc.

Email: tte@cs.fau.de

Alexander Clemm
Futurewei Technologies Inc.

Email: ludwig@clemm.org