

International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Series Y
Supplement 66
(07/2020)

SERIES Y: GLOBAL INFORMATION
INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS,
NEXT-GENERATION NETWORKS, INTERNET OF
THINGS AND SMART CITIES

**ITU-T Y.3000-series – Network 2030 services:
Capabilities, performance and design of new
communication services for the Network 2030
applications**

ITU-T Y-series Recommendations – Supplement 66

ITU-T



ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS, NEXT-GENERATION NETWORKS, INTERNET OF THINGS AND SMART CITIES

GLOBAL INFORMATION INFRASTRUCTURE

General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899

INTERNET PROTOCOL ASPECTS

General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	Y.1300–Y.1399
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	Y.1700–Y.1799
Charging	Y.1800–Y.1899
IPTV over NGN	Y.1900–Y.1999

NEXT GENERATION NETWORKS

Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Enhancements to NGN	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Packet-based Networks	Y.2600–Y.2699
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899
Carrier grade open environment	Y.2900–Y.2999

FUTURE NETWORKS

CLOUD COMPUTING	Y.3000–Y.3499
-----------------	---------------

BIG DATA	Y.3500–Y.3799
----------	---------------

QUANTUM KEY DISTRIBUTION NETWORKS	Y.3800–Y.3999
-----------------------------------	---------------

INTERNET OF THINGS AND SMART CITIES AND COMMUNITIES

General	Y.4000–Y.4049
Definitions and terminologies	Y.4050–Y.4099
Requirements and use cases	Y.4100–Y.4249
Infrastructure, connectivity and networks	Y.4250–Y.4399
Frameworks, architectures and protocols	Y.4400–Y.4549
Services, applications, computation and data processing	Y.4550–Y.4699
Management, control and performance	Y.4700–Y.4799
Identification and security	Y.4800–Y.4899
Evaluation and assessment	Y.4900–Y.4999

For further details, please refer to the list of ITU-T Recommendations.

Supplement 66 to ITU-T Y-series Recommendations

ITU-T Y.3000-series – Network 2030 services: Capabilities, performance and design of new communication services for the Network 2030 applications

Summary

Supplement 66 to ITU-T Y.3000 series Recommendations aims to assess the new services and capabilities required by Network 2030. This Supplement includes common terminology and definitions necessary for describing new services. It also analyses gaps in existing communication technology to provide the reasoning behind the new communication services that are proposed in this Supplement.

History

Edition	Recommendation	Approval	Study Group	Unique ID*
1.0	ITU-T Y Suppl. 66	2020-07-31	13	11.1002/1000/14385

Keywords

Coordination guarantee, haptic communication, high precision communication, holographic type communication, in-time guarantee, many-to-many, Network 2030, one-to-many, on-time guarantee, qualitative communication, qualitative communication.

* To access the Recommendation, type the URL <http://handle.itu.int/> in the address field of your web browser, followed by the Recommendation's unique ID. For example, <http://handle.itu.int/11.1002/1000/11830-en>.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

This is an informative ITU-T publication. Mandatory provisions, such as those found in ITU-T Recommendations, are outside the scope of this publication. This publication should only be referenced bibliographically in ITU-T Recommendations.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this publication may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the publication development process.

As of the date of approval of this publication, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this publication. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2020

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

Table of Contents

	Page
1 Scope	1
2 References.....	1
3 Definitions	1
3.1 Terms defined elsewhere	1
3.2 Terms defined in this Supplement	1
4 Abbreviations and acronyms	2
5 Conventions	3
6 Overview of Network 2030 services	3
7 Motivation for Network 2030 services	4
8 Network 2030 services: Foundational services	6
8.1 In-time and on-time services	6
8.2 Coordinated services	13
8.3 Qualitative communication service	19
9 Network 2030 services: Compound services.....	25
9.1 Haptic communication.....	25
9.2 Holographic-type communications (HTC) services.....	29
10 Other aspects and capabilities of future networking services.....	32
10.1 Network service interfaces	32
10.2 High programmability and agile lifecycle.....	33
10.3 Manageability	34
10.4 Security.....	35
10.5 Resilience	36
10.6 Loss-lessness	37
10.7 Privacy.....	38
10.8 Trustworthiness	40
10.9 Accounting, accountability, validation of delivered services.....	40
Bibliography.....	42

Introduction

This Supplement describes new services and capabilities required by network 2030 in order to meet the requirements and challenges imposed by future network applications. These services are referred to as "Network 2030 services" and they are intended to accommodate the use cases and meet the requirements that have been analysed by the ITU-T Focus Group on Network 2030 (FG-NET 2030) in ITU-T TD375.

IMT 2020 [b-IMT2020] describes the emergence of immersive media enabled by high bandwidth and of new critical applications enabled by massive machine-type communications (mMTC) and ultra-reliable low-latency communications (uRLLC). Over the next ten years, further advances in multimedia can be expected that will make it even richer and more immersive and interactive, involving, for example, holographic communications and tele-haptics. To enable this, not only will abundant bandwidth and ubiquitous connectivity be necessary, but networks will also need to provide new capabilities that are not supported today. This includes, for example, the ability to deliver on stringent latency guarantees and to provide precise coordination across many concurrent data streams and communication channels.

Current internetworking infrastructure provides network services that are fundamentally built on the basis of "best effort". While differentiated services allow for the prioritization of traffic and the reservation of resources, and while transport-layer protocol can add reliability via retransmission schemes, all of these mechanisms are associated with significant trade-offs and limitations. In order to support new applications, Network 2030 services need to move beyond best effort and support a new concept of "high precision": high precision in terms of quantifiable latency guarantees, in terms of synchronization of packet flows across multiple communication channels and communicating parties, in terms of behaviour in light of congestion and resource contention.

Network 2030 services can be categorized into foundational and compound services. Foundational services are those that cannot be decomposed further. For example, a service with a certain bandwidth capacity or a certain guaranteed resiliency is a foundational service. Compound services are those assembled from other foundational network services. An example of a compound service is a holographic type communication service that would provide assurance (each an instance of a foundational service) of both latency and throughput. Some compound services are closely related to application-layer services but focus on those aspects that require support by networking infrastructure that cannot simply be addressed by endpoints themselves or overlays.

This Supplement is structured as follows:

- Clause 2 provides a list of references.
- Clause 3 contains a set of definitions and acronyms that are used in this Supplement.
- Clause 4 re-iterates what is covered as part of the scope of this Supplement.
- Clause 5 describes conventions.
- Clause 6 give an overview of Network 2030 services.
- Clause 7 describes the motivation and summarizes the drivers for Network 2030 services. It summarizes requirements and gaps in existing network technology that need to be addressed.
- Clauses 8 describes new foundational networking services and their properties. The focus is on foundational (basic) services that cannot be decomposed further, which may serve as building blocks also for compound services and applications.
- Clause 9 describes compound networking services that build on foundational services as introduced in clause 8. Using several examples, it shows how foundational future networking services as described in clause 8 can be used to construct future service offerings at a higher layer, such as the tactile Internet or holographic-type communications (HTC), and how networking applications shape the requirements for such services.

- Clause 10 discusses additional aspects of future networking services, such as architectural implications of service interfaces or aspects related to "soft" properties of services, such as manageability, programmability, and security.

Supplement 66 to ITU-T Y-series Recommendations

ITU-T Y.3000-series – Network 2030 services: Capabilities, performance and design of new communication services for the Network 2030 applications

1 Scope

This Supplement to ITU-T Y.3000 series Recommendations focuses on new network-layer services, "network services" for short, and tries to assess the new network services and capabilities required based on Network 2030 vision. This Supplement describes new communication services for Network 2030, provides gap analysis, and specifies performance targets for different types of new services and capabilities.

2 References

[ITU-T Y.3112] Recommendation ITU-T Y.3112 (2018), *Framework for the support of network slicing in the IMT-2020 network*.

3 Definitions

3.1 Terms defined elsewhere

None.

3.2 Terms defined in this Supplement

This Supplement defines the following terms:

3.2.1 compound (or composite) network service: A network service that is provided through a combination of other network services.

3.2.2 coordinated service: A network service that involves multiple concurrent flows that exhibit some inter-flow dependency, such as a timing or ordering dependency that needs to be observed and coordinated respectively by the network.

3.2.3 foundational network service: A network service that cannot be decomposed or provided in terms of other services.

3.2.4 haptic communication service: A compound network service that serves the needs of haptic applications that convey a sense of "touch" to end users, involving both tactile data (e.g., surface texture, pressure points) and kinaesthetic data (positioning awareness).

3.2.5 high-precision network services: Network services that support stringent service level objectives at very high precision that is explicitly specified, such as in-time and on-time latency guarantees.

3.2.6 holographic-type communication service: A compound network service that serves the needs of applications which need to convey holographic data.

3.2.7 in-time service: A service that ensures delivery of packets with a required latency that is not to be exceeded.

3.2.8 on-time service: A network service that ensures delivery of packets with a required latency within a specified time window.

3.2.9 qualitative communication service: A network service that can differentiate between payload chunks of different relative priority (assigned by users or applications, not involving payload

inspection by the network) and apply differentiated treatment to those chunks, for example with regard to selective dropping and retransmission of chunks.

4 Abbreviations and acronyms

This Supplement uses the following abbreviations and acronyms:

AC	Admission Control
AI	Artificial Intelligence
API	Application Programming Interface
AR	Augmented Reality
CBR	Constant Bit Rate
CC	Congestion Control
CPS	Cyber-Physical-Systems
CRC	Cyclic Redundancy Check
DB	Delay Bound
DetNet	Deterministic Networking
Diffserv	Differentiated Services
DSCP	Differentiated Services Code Point
EF	Expedited Forwarding
FRER	Frame Replication and Elimination for Reliability
FRR	Fast Re-Route
HTC	Holographic-Type Communications
Intserv	Integrated Services
IP	Internet Protocol
mMTC	Massive Machine Type Communications
MPLS	Multi-Protocol Label Switching
NFV	Network Function Virtualization
PHB	Per Hop Behaviour
PII	Personally Identifiable Information
PIM	Protocol Independent Multicast
PREOF	Packet Replication, Elimination, and Ordering Functions
QoE	Quality of Experience
QoS	Quality of Service
QUIC	Quick User datagram protocol Internet Connections
RSVP	Resource Reservation Protocol
SDN	Software-Defined Networking
SLA	Service Level Agreement
TCP	Transmission Control Protocol

TSN	Time Sensitive Network
uRLLC	Ultra-Reliable Low-Latency Communications
VBR	Variable Bit Rate
VR	Virtual Reality

5 Conventions

None.

6 Overview of Network 2030 services

Networks are usually layered into physical layer (layer 1), link layer (layer 2), network layer (layer 3), transport layer (layer 4) and application layer (layer 7). New functionalities are sometimes added between two layers. For example, "MPLS" can be viewed as a "thin" layer between the link layer and the network layer. Each layer provides a service for its immediate upper layer to use/consume. The transport layer and application layer are traditionally implemented inside host operating systems. More recently, solutions have started to appear where aspects of the transport layer are also supported by hardware in routers inside the network. However, the network layer is mostly implemented by routers and sometimes by switches. This Supplement is within the scope of the network layer.

The Supplement also focuses on the definition of new network-layer services, "network services" for short, in order to support emerging applications and vertical industries in the year 2030 and beyond. Usually, network-layer services are provided by routers, and sometimes by switches on the data plane or user plane.

It should be emphasized that this Supplement focuses only on *new* services, not on services which are already supported today and that are expected to continue to (co-)exist. The new services that are defined here will not necessarily replace today's network services, nor will the network needed to support these new services necessarily replace today's network. Instead, it should be anticipated that new services will be added and provided in addition to existing services, which in many cases will continue to be offered.

This Supplement uses the term "foundational" and "compound" (or "composite") services according to the illustrative diagram in Figure 6-1.

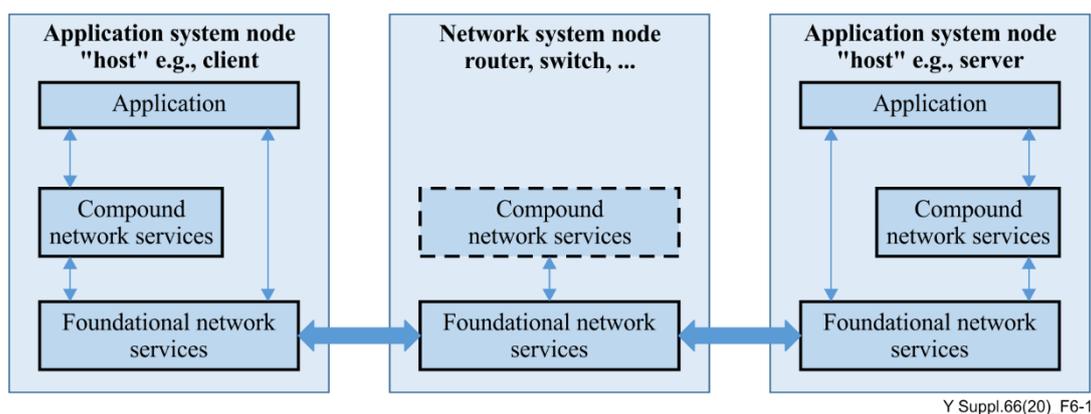


Figure 6-1 – Foundational and compound network services

A foundational network service is one that requires dedicated support on some or all network system nodes that deliver the service between two or more application system nodes. For example, IP packet routing and forwarding is a (pre-existing) foundational network service.

A compound (or composite) network service is one that can be composed of one or more foundational services. Compound network services of interest in this Supplement are those that require at least one new foundational service (and any number of pre-existing foundational network services), but a compound network service itself does not necessarily introduce any new network service or requirements into the network system nodes.

To enable these new network services on the data plane, new transport-layer services and new capabilities on the control plane may be required but are not discussed in this Supplement.

The Supplement also discusses several new compound services that will leverage new foundational network services. For example, tactile Internet services [b-Tactile] and holographic-type communications (HTC) services that are two compound services which will enable many Network 2030 use cases. For such services, encodings and compression techniques for tactile or holographic information [b-Schwarz] are important building blocks but are outside the scope of this Supplement, as they are expected to be performed by applications or endpoints without the involvement of network nodes (which should be agnostic to the contents/payload). However, support for those techniques requires the use of Network 2030 foundational services, for example to meet latency guarantees and to coordinate the delivery of packets across multiple flows.

In summary, this Supplement defines and describes the following new network-layer services on the data plane:

- High-precision communications services with service level objectives that relate to packet travel time, or latency, in networks:
 - In-time services, which require latency to be within a quantifiable limit.
 - On-time services, which require latency to be of an exact duration, with the possibility of a small quantifiable variance.
 - Coordinated communication services (network services that require coordination between multiple flows, with interdependencies between the service levels that need to be delivered across these flows).
- Qualitative communication services, new services that suppress retransmission of portions of the payload that are deemed less relevant in order to meet requirements on latency by applications that are tolerant of certain quality degradation.
- Coordinated and generalized multicast services which involve high-precision services between a sender and a group of receivers, between a group of senders and a single receiver, or between groups of senders and groups of receivers.
- Network services to support haptic applications.
- Network services to support holographic-type communications (HTCs).

7 Motivation for Network 2030 services

Future advances in networking technology will be driven by future networking applications. The following applications areas are seen among the primary drivers for new Network 2030 services:

- a) **Industrial and robotic automation:** Machine-to-machine communication for industrial and robotic automation is at the heart of the next industrial revolution that is commonly referred to as "Industry 4.0" [b-Brettel] [b-blueoceans] [b-Forbes]. This type of machine-to-machine communication requires very fine-grained timing accuracy for the dispersion of control commands and for the collection of telemetry data. Without this capability, the envisioned high-precision control loops quickly break down. Network 2030 services therefore need to support critical grade reliability and extremely low as well as highly precise latency for the delivery of packets.

- b) **Emergence of holographic media and other advances in multimedia technology:** Holograms, haptics, and other sensory data will provide immersive and "real" user experience, enhancing the experience when media is consumed and facilitating interactions of users in a world where the line between what is real and what is virtual becomes increasingly blurred. For this to happen, very high data throughput involving tight coordination across bundles of streams among multiple stream sources and sinks will be necessary, as well as the ability to rapidly prioritize data items within and between streams per guidance from applications. This will be coupled with advances in the way in which the environment is captured and rendered by endpoints. The emergence of holographic media requires a new type of communication over networks: HTCs, which is characterized by very high throughput, timely delivery, (sometimes) tolerance of quality degradation, and coordination when multiple parties join the same holographic streaming application.
- c) **Autonomic and critical infrastructures:** Network services that enable mission-critical applications such as self-driving vehicles, drones, automated traffic control systems, all communicating with one another and their environment, need to be failsafe so that infrastructure can rapidly adapt and react to unexpected events. Likewise, extremely high demands will be placed on such services to avoid the possibility of tampering and ensuring that trust and accountability are maintained. Without it, such applications might quickly devolve from a perceived blessing into a safety hazard. For autonomic and mission-critical applications, time-guaranteed packet delivery is often required and/or favoured.
- d) **Diversity of applications and their needs:** An explosion of new applications that consume networking services can be expected. Many of those applications may be driven by artificial intelligence (AI) and depend on myriads of data feeds; they may also involve novel mixes of humans, machines, and IT systems communicating with one another. Some of the resulting networking needs may differ in unexpected ways, for example, in the communication patterns that need to be supported, in the interdependencies between data streams, in their needs to account for and validate that communication services have been delivered, in the service levels that they require, and in the trade-offs (for example, between throughput, reliability, and latency) that they are willing to make. This will require the ability to not just deliver but also to dynamically adapt associated network services.
- e) **Accountability for services delivered:** Often stringent service requirements directly impact operational and maintenance costs as seen through managed service models. Accountability in the form of evidence that actions were taken in delivering a service in conformance with the agreement benefits both the functional and business aspects of a service. It incentivizes service providers to offer new type of service delivery models and allow innovations in applications to incorporate such capabilities. Assured packet service delivery with auditable evidence is expected to unlock new opportunities for both service providers and their customers.
- f) **Expectations for varying degrees of distortion tolerance:** Applications that can absorb intermittent or partial loss of data and still function normally are said to be distortion tolerant. While many Network 2030 applications are characterized by their need for high precision networking services, other classes of applications may in fact be distortion tolerant to a degree. In some such cases, applications may demand novel abilities to differentiate contents that is tolerant to loss, and articulate more sophisticated ways to deal with such loss than is the case today – for example, the ability to apply network coding schemes or the ability to specify content-dependent prioritization and protection schemes that are supported by networking infrastructure, instead of just endpoints connecting to the network.

Many of the drivers listed above point to precise timing and latency of packet delivery, coupled with the ability to provide precise control of that latency, as a critical enabler of Network 2030 applications. This leads to the need for in-time and on-time services (discussed in clause 7), which allow to quantify precise latency objectives given various constraints (such as required throughput and acceptable loss).

In addition, many drivers involve applications that will require multiple concurrent flows with various interdependencies, which in turn require close packet delivery coordination. Communication patterns include not only one-to-one, but also one-to-many (or even many-to-many), all of which need to be supported by a consolidated set of network services. Those requirements are addressed by coordinated services, discussed in clause 8.2.

Driver (f) points to a need for services that allow applications to differentiate between payload that must not be "distorted" and needs to be protected by the network at all cost, and payload for which loss is an acceptable trade-off to other service level parameters such as latency and reservation cost. These needs are addressed by qualitative communication services, as described in clause 8.3.

Drivers (a) and (b) also point to the special relevance of tactile network services and holographic-type communications in Network 2030 applications. Those are examples of compound services (building on the more foundational services introduced in clause 8) and discussed further in clause 9.

Finally, many drivers (e.g., e, f) point to other trends and aspects that will affect how Network 2030 will be delivered. Those aspects include the ability to account for proper usage of services, to facilitate rapid customization of network services by users and applications, and advances needed to properly manage at very high precision the delivery of Network 2030 services. Those aspects will play an important role in the commercial adoption and success of many applications and are discussed in clause 10.

8 Network 2030 services: Foundational services

8.1 In-time and on-time services

8.1.1 Introduction and motivation

One important category of Network 2030 services concerns communication services that adhere to stringent quantifiable latency objectives. For example:

- Haptic applications (as outlined in clause 7 item b) require end-to-end networking latencies with an upper bound on the order of 5 ms or less. This low latency is needed in order to allow for round-trip control loops that allow haptic applications to communicate feedback in well under 10 ms, even as low as 1 ms in some cases [b-Fettweis]. If such guarantees cannot be met, not only does the quality of experience for users deteriorate, but the applications themselves may become unusable. This is the case because to the end user, the illusion of remotely "touching" something and the ability to, for example, remotely operate machinery based on haptic feedback is lost. Latencies that are merely "as short as possible" are insufficient; instead, quantified exact latency requirements must be met.
- Autonomic mission-critical infrastructure (as outlined in clause 7 item c) relies on similar latency objectives. For example, latency must be extremely short to avoid collisions between vehicles that are operated and controlled remotely, etc.; at the same time there is no tolerance for packet loss. Again, merely making the latency "as short as possible" as is done in the current Internet is not sufficient. Instead, quantified objectives must be met; otherwise autonomic mission-critical applications cannot be supported.
- Industrial and robotic automation (as outlined in clause 7 item a) requires not only "not-to-exceed" latency, but latency that is in effect "deterministic", with packets not only not exceeding a certain latency, but also not being delivered any sooner. This is because some industrial controllers require very precise synchronization and spacing of telemetry streams and control data, facilitating (for example) precise operation of robotic effectors along multiple degrees-of-freedom.

Network 2030 services therefore need to support "high-precision" communication services, where "high-precision" refers to a precise latency that packets may incur, which is explicitly specified. Those

services are also referred to as "in-time" and "on-time" services, with respect to the latency objectives that are imposed on the packets that deliver those services.

Contrary to existing technology, in which networks can be engineered and optimized for "low" latency, but the actual latency that is obtained still needs to be measured, latency objectives in Network 2030 should be provided as a specific parameter for the service.

It should be noted that in cases where ultra-low latencies are required, physical limitations related to the propagation speed of light come into play which may prohibit communication over long geographical distances while meeting latency objectives at the same time. In those cases, network services need to be complemented with application architectures that, for instance, place content and computation close to the edge.

8.1.2 Description

There are many existing definitions of the term "latency", including definitions that are used in ITU-T. For the purposes of this Supplement and in the following discussion, *latency* refers to the time that elapses from when a packet is sent by a sender (i.e., from when the transmission of the first bit of the packet is started) until the packet (i.e., the last bit of the packet) is received by a receiver across the network.

In this discussion, the focus is on latency that is incurred from a sending point to a receiving point. A sending/receiving point can be a host or a border router of a domain (usually an autonomous system). When no confusion arises, such latency is regarded as end-to-end latency. Latency that is incurred between individual hops on individual path segments along the path from sender to receiver is not taken into consideration. End-to-end latency is the aggregate of multiple component latencies, including latency that is incurred by physical propagation and processing of packets along individual hops for queuing, packet serialization, and packet processing.

"*Required latency*" refers to an objective for the latency of a packet. "*Actual latency*" refers to the latency that a packet physically experiences.

In-time services are services that ensure delivery of packets with a required latency that is not to be exceeded. Packets may be delivered at any time before or until the latency deadline. Multimedia applications supporting buffering capabilities are typical applications that use in-time services.

A client application requesting an in-time service will specify:

- The required maximum latency that is not to be exceeded.
- Optionally, constraints under which the required latency is to be delivered:
 - The expected bandwidth (e.g., packet rate, possibly differentiated by sustained and burst rates).
 - The acceptable miss rate (i.e., the ratio of packets that are dropped or do not meet the required latency versus the total number of packets). Miss rate refers to the ratio of the number of packets in a flow that miss a service level objective (in the context of this Supplement: required latency), including packets that are lost, to the total number of packets in a flow over a given time interval.

On-time services are services that ensure the arrival of data within a specific time window. Like in-time services, they impose a maximum latency that is not to be exceeded. In addition, they indicate a minimum latency. A packet must be delivered no later than the upper bound of the time window, but also no earlier than the lower bound of the time window. The window can be specified either in terms of specifying lower and upper bounds, or in terms of a latency target representing the midpoint of the window and the size of the window. A special case of an on-time service is the case when the time window is nominally 0 (with the lower bound equalling the upper bound), resulting in latency which is deterministic within the bounds of the clock uncertainty.

A client application requesting an on-time-service will specify:

- The required latency (specified, for example, using a target latency midpoint and time window, or lower and upper latency bounds, or even a target delivery time which is converted into latency by the network).
- Optionally, constraints under which the required latency is to be delivered (as with in-time services).

Figure 8-1 summarizes the difference between in-time and on-time data delivery and shows the latency with which a packet is expected to be received. Packets whose actual latency falls outside the range that is depicted in green (i.e., packets that are late or early) are considered out of compliance and contribute to the miss rate. The miss rate specifies the ratio of packets that fail to be delivered per the required latency, i.e., whose actual latency falls outside the required latency range, or that are lost entirely. The miss rate must approach 0 as close as possible.

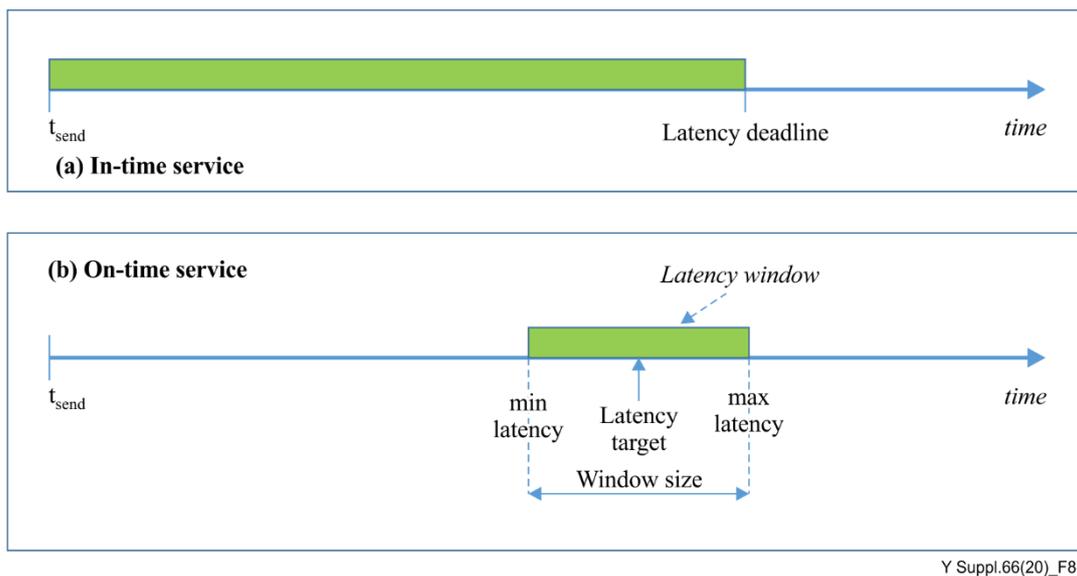


Figure 8-1 – In-time vs on-time services

An "in-time service" can be considered as a special case of an on-time service in which the acceptable latency window extends all the way from zero. However, solutions that support an on-time service involve additional challenges beyond solutions that support merely an in-time service. As packets must not be delivered earlier than a minimum latency, the network needs to be able to buffer packet or defer their delivery when needed.

In mathematical terms, for a minimum latency pl_{min} , a maximum latency pl_{max} , and a latency pl_i of a given packet p_i the following must hold for all packets p_i in the flow for an on-time-service:

$$\forall p_i: pl_{min} \leq pl_i \leq pl_{max}$$

And for an in-time-service:

$$\forall p_i: 0 \leq pl_i \leq pl_{max}$$

Additional considerations

Granularity

Latency targets and window bounds are specified with a certain granularity. It should be noted that a latency target of "10 ms" is not the same as a latency target of "10 000 us". In the former case, any packet of an on-time service that arrives with a 9 500 to 10 499 us latency would be deemed as meeting its latency target (as the latency in each case translates to 10 ms). In the latter case, only packets arriving with 9 999.5 to 10 000.5 us latency would actually meet the target.

Accuracy

Latency can sometimes be difficult to assess accurately. Limitations in accuracy should be accounted for in the definition of latency bounds. An on-time service should therefore define a corresponding latency window taking into account the limitations of the nodes' timing inaccuracies. For example, if it is known that the time accuracy between nodes in the network is ± 1 ms, an on-time service would need to define a corresponding latency window that accounts for the possibility of inaccuracies.

Constraints

In order to deliver to a required latency, networks need to make certain assumptions and impose certain constraints. Specifically, it makes a difference if a sender expects to send a single packet or a flow of packets at a high rate, as this imposes different demands on the resources that need to be available. Likewise, it makes a difference whether an application is sensitive to packet loss since this may impose additional robustness considerations.

For a network to be able to provide any type of commitment or guarantee to meet a required latency, constraints or assumptions for those commitments must be specified, specifically bandwidth (or packet rate and packet size), and acceptable packet loss rate (possibly zero). These become additional input parameters for in-time and on-time service requests, in addition to the required latency.

On the difference between "latency" and jitter

It should be noted that latency and jitter (inter-packet delay variation) are not the same. Jitter refers to differences in latency that are experienced for different packets of a flow. More formally, with plv denoting the packet latency variation and pl_{min} and pl_{max} denoting the maximum respectively minimum latency of packets in a flow:

$$plv = (pl_{max} - pl_{min})$$

The window size in an on-time service provides an upper bound for possible jitter, but the actual jitter that is experienced could be smaller. Likewise, high latency does not imply high jitter. It is conceivable to specify separate jitter objectives, which designate the permissible variation in latency independent of the actual latency.

Concluding notes

Networks that support in-time and/or on-time services need to be "latency aware", in order to be able to determine whether packets comply with latency targets and to react accordingly. In addition, they need to be able to assess whether they can meet a latency target before giving any type of guarantee, before admitting traffic for which such a guarantee is granted. Furthermore, in order to meet the lower latency bounds before which packets must not be delivered, networks that provide an on-time service need to be capable of buffering or "slowing down" packets before they reach their destination.

For traditional applications, tolerable latency from the time data is sent to when it is played out to a user by a receiving endpoint can be in the order of a hundred milliseconds. For networks in 2030, acceptable latency may be much lower depending on the application. For example, tactile Internet applications may call for deadlines not to exceed 10 milliseconds. Packets that are late might as well not be delivered at all; in fact, late delivery might even be harmful. This means that contrary to the behaviour of current networks, support for latency deadlines that are quantified needs to be provided.

On-time services will typically be required to operate at a fine-grained timescale granularity (e.g., microseconds) and should be able to offer deterministic latencies. For instance, in industrial applications, a controller might need to send commands to a sequence of devices. Each device should receive the command and operate at a precise time in a streamlined (or quasi-synchronized) manner.

8.1.3 Gap analysis

This clause describes the requirements that in-time and on-time services must meet and gives an overview of architectures and services that were proposed by the research community and the standardization groups in order to offer in-time/on-time guaranteed services. A description of the gaps that remain is given in the final paragraph of clause 8.2.2.

8.1.3.1 Requirements

As mentioned in clause 8.1.2, in-time services must provide the ability to support a quantifiable end-to-end latency for packet delivery across a network that must not be exceeded, given a set of constraints (which include a rate at which packets can be sent, and a loss rate that would be acceptable).

- Services may be required to ensure that clients adhere to agreed constraints (e.g., a not-to-exceed packet rate). For this purpose, they may perform admission control (AC) or rate limiting as needed. Alternatively, services may simply monitor agreed-to-constraints to warn users in case violations occur. Any high-precision commitments given by the provider of a high-precision network service will no longer apply in case of violation of constraints – while the network may still deliver the demanded latencies if it is possible to do so, it is not committed to do so.
- "Miss rate" is defined as the ratio between the number of packets that do not meet the latency objective (including packets that are lost), and the total number of packets. In other words, a low-precision packet which misses its latency objective is considered the same as a "lost" packet. In an extreme case, it is possible that a miss rate could be specified as zero, in which case no misses would be acceptable. However, it should be noted that in reality, a miss rate (and loss rate) of zero will be impossible to achieve at all times and under all circumstances (for example in case of occurrence of a cosmic event), although it can be asymptotically approached and guarantees can be given that in the presence of, e.g., single device and single link failures, no misses will occur.

On time-service must additionally support a quantifiable end-to-end latency that must be met within a given window. The window boundaries define a latency that is not to be exceeded (as in the case of an in-time service), as well as a minimum acceptable latency.

Accordingly, an in-time service request is characterized by the following parameters:

- Required latency
- Constraints:
 - Packet rate (possibly refined further, e.g., sustained vs burst)
 - Miss rate

An on-time service request is characterized by the following parameters:

- Required latency
- Latency window size
- Constraints:
 - Packet rate
 - Miss rate

8.1.3.2 Existing mechanisms requirements and their gaps

Traditional networks support multiple mechanisms to reduce and optimize latency. What those mechanisms have in common is that by and large they do not support latency objectives that can be quantified in advance. While today's networks can be engineered with certain latency outcomes in mind (applying various schemes to dimension, allocate, reserve resources and prioritize traffic

[b-Bennett] [b-Boudec]), latency is fundamentally still measured, not delivered on by design, and has to be accepted as such. The gap that needs to be addressed for Network 2030 services concerns the ability to deliver on latency objectives that are precisely quantified as part of the service request.

8.1.3.2.1 Internet QoS architecture

The IETF defined two complementary high-level quality of service (QoS) architectures for IP networks: Integrated Services (IntServ) and Differentiated Services (Diffserv). These architectures can also be combined.

IntServ [b-IETF RFC 1633] includes two services: The Guaranteed Service ([b-IETF RFC 2212]) and Controlled Load Service ([b-IETF RFC 2211]). Of particular relevance here is the Guaranteed Service, for which the term "IntServ" will be used interchangeably in this Supplement. IntServ provides per-flow fixed bandwidth guarantees and is based on the concept of reserving resources in advance for a given flow, which are for exclusive use by packets of that flow and not shared with other flows. To maintain bandwidth guarantees, IntServ traffic is shaped at the ingress network edge as necessary so the flow does not consume more resources than have been reserved. To support latency guarantees, flows need to be re-shaped on every hop. Without shaping, collisions and resource contention between packets could occur, which would lead to the possibility of loss and unpredictable variations in latency. IntServ-type solutions are also referred to as "admission controlled" (AC).

IntServ is a precursor for the IEEE L2 time sensitive networking (TSN) solution [b-TSN] and recent IETF deterministic network (DetNet) solutions [b-DetNet], which are described separately below. Both TSN and DetNet are based on different variations of the reservation principle and support fundamentally the same type of services.

DiffServ is a multiplexing technique that is used to manage bandwidth between different classes of traffic, including IntServ-style admission-controlled traffic as well as other traffic, e.g., traffic that is subject to congestion control.

For congestion-controlled traffic, no resource reservations are made in advance, which leads to the possibility of network congestion. This congestion can be mitigated in several ways; the corresponding techniques are referred to as congestion control (CC). Congestion control leads to the dynamic adjustment of flow bitrates based on the available bandwidth resources in the network. In the worst case, congestion can lead to loss which CC cannot avoid. In that case, retransmission is normally used for recovery. This becomes an issue for low-latency services, which often cannot afford retransmissions because this would result in the target latency being exceeded. While some applications that require low latency may be able to deal with low probability random packet loss resulting from transmission media bit error rates, very few can cope with typical 50 msec interruptions resulting from equipment or link failures when reactive protection such as fast re-route (FRR) mechanisms are used.

The Internet QoS architecture is insufficient to meet the needs of Network 2030 for a number of reasons, including the following:

- The need for per-flow admission control makes IntServ expensive to support and scale, even if performed out-of-band via software-defined networking (SDN).
- The inability to dynamically adjust the bitrate under varying network utilization makes this model too inflexible even for current, let alone future networks. Originally built for non-IP voice/video applications that required fixed bandwidth and support for constant bit rates (CBR), Network 2030 applications require support for variable bit rates and elastic bandwidth. This is not adequately supported by the existing Internet QoS architecture.
- No mechanisms exist to support application-defined upper and lower bounds for the desired latency independent of the path round trip time.
- There are no mechanisms to slow down packets based on the desired earliest delivery time.
- Queuing cannot prioritize packets based on their desired end-to-end latency.

The same reasons fundamentally apply also to TSN and DetNet, described in clauses 8.1.3.2.2 and 8.1.3.2.3, respectively.

8.1.3.2.2 Time-sensitive networking

The time-sensitive networking (TSN) [b-DetNet] [b-SIEMENS-In] is a set of updates to the IEEE Ethernet standard that aims to empower standard Ethernet with time synchronization and deterministic network communication capabilities. For the purpose of discussing gaps for latency control, TSN can be best understood as an Ethernet layer 2 variation of the IntServ service, but with two important enhancements:

- With 802.1QCH (cyclic queuing), TSN supports a model for deterministic shaping that does not require per-flow state on transit nodes. It requires strict time synchronization and its throughput deteriorates with increasing network size. With 802.1QCR, TSN also introduces "Asynchronous Traffic Shaping" (ATS) in the style of IETF IntServ to avoid the need for time synchronization.
- With 802.1CB frame replication and elimination for reliability (FRER), TSN introduces 1:n (n typically 1) path protection where packets are replicated n+1 times on ingress, sent across failure disjoint paths and then the replicas are eliminated on egress. This so-called pro-active path protection supports close-to-zero loss in the face of link or equipment (node or linecard) failure. In contrast, reactive mechanisms such as L2 or L3 fast reconvergence or fast-reroute typically require up to 50 msec to patch the failure caused interruption – which is too long for low-latency traffic.

As a collection of layer 2 Ethernet services, TSN aims to provide deterministic service inside a LAN over a short distance, and is thus not routing-capable. TSN does not aim to provide on-time guaranteed service over large-scale networks and over longer distances. Like IntServ, TSN is geared towards CBR traffic, not variable bitrates (VBR) traffic, and does not support the slowing down of packets based on the required earliest delivery time.

8.1.3.2.3 Deterministic networking architecture (DetNet)

The deterministic networking architecture (DetNet) [b-Finn] is an architecture that has been proposed by the IETF DetNet Working Group in order to ensure a bounded latency and low data loss rates within a single network domain.

The DetNet architecture intends to provide per-flow service guarantees in terms of (1) the maximum end-to-end latency (called bounded delay in DetNet) and bounded jitter, (2) packet loss ratio, and (3) an upper bound on out-of-order packet delivery. Some options considered in DetNet may in the future also be able to provide bounded delay-variation between packets of a flows. These service guarantees are ensured thanks to three techniques used by DetNet. The first of these techniques involves resource reservations (to avoid the possibility for resource contention) as well as per-hop re-shaping of traffic to avoid accumulation of bursts further downstream. The second technique involves protection against loss caused by random media errors and equipment failures. It is based on the Packet replication, elimination, and ordering functions (PREOF) mechanism. PREOF is similar to the TSN FRER mechanism and is based on duplicating single flows into multiple flows that traverse disjoint paths, then re-combining them and dropping any duplicates near the egress point. The third technique concerns the use of explicit routing to take advantage of engineered paths with specific bandwidth/buffering properties that are disjoint to other paths required for PREOF.

Although DetNet provides efficient techniques to ensure deterministic latency, scalability remains a challenge. In particular, implementing the DetNet techniques requires the data plane to keep track of per-flow state and to implement advanced traffic shaping and packet scheduling schemes at every hop. This is not scalable because core routers can receive millions of flows simultaneously. In the control plane, if resource reservation protocol (RSVP) is used, every hop needs to maintain per-flow resource reservation state, which is also not scalable.

The most fundamental limitation of DetNet, similar to IntServ, is in its targeted scope of CBR reservations. Future applications may have highly variable bitrates (VBR). Lower latency bounds, as required for on-time services, are also not directly supported in DetNet. While arguable bounded jitter in effect also imposes a lower bound in the case of CBR traffic, the same is not true for VBR. There is still much more work to be done in order to design a solution that is both effective in ensuring deterministic latency for all types of traffic (not just at constant bitrates) and scalable to support a large number of simultaneous flows.

Additional limitations apply regarding the combination of PREOF and admission control, as applicable by IntServ as well as DetNet: fully distributed solutions such as distributed Maximum Redundant Trees cannot calculate the optimum paths and do not work well together with admission control. At the same time, centralized solutions have no scalable method to instantiate their desired paths in the network forwarding plane, and existing forwarding plane mechanism based on loose path steering can cause unexpected path traffic under failure.

8.1.4 Performance and design target

Network 2030 applications may impose required latencies as low as 5 ms, for example, for tactile Internet applications.

Granularity that is specified and measured in microseconds (for end-to-end latency) may need to be supported (for example, for certain Industrial Internet applications).

Likewise, accuracy on the order of 1 microsecond may need to be supported.

8.2 Coordinated services

8.2.1 Introduction and motivation

A new category of Network 2030 services concerns communication services that adhere to stringent quantifiable coordination objectives. For example:

- Multi-sensory communications: In today's networks, audio/visual information in augmented reality/virtual reality (AR/VR) media is sent as a single flow. As a variety of sensory experiences may get integrated with holographic communications, it may become necessary to transmit different views or sensor feeds over different flows and possibly over different paths that are subject to different latency and bandwidth constraints. However, to provide a fully immersive experience, delivery of multi-sensory information needs to be synchronized across all sources when delivered to the user. Furthermore, the characteristics of the streamed data vary in terms of encodings, packet sizes, and user's perception times. This can have an impact on buffering, scheduling, and traffic shaping mechanisms in the network, making the coordinated delivery of dependent information challenging.
- Virtual orchestra and/or concerts: Imagine an instrument ensemble in which the holographic life-size 3D projections of musicians, each in a different place in the world come together and perform live on the stage in front of the audience. Further, assume a conductor on the stage directing the sound of the ensemble with his gestures. These gestures must be received at the same time by the remote musicians at different locations to play their instruments at a specific time with a specified tempo. Similarly, the music transmitted from those locations to the stage must be played together with the same beats and tempo. Any delay or early arrival of the sound from any one instrument can cause the ensemble to go out of tune and destroy the entire performance. Furthermore, performing an ensemble with multiple participants separated by large as well as varying distances (from less than a mile, to 1000 miles) is quite difficult for applications due to varying path and latency characteristics. Therefore, the network needs to support the coordination of directions from the conductor to all of the musicians and the audio/visuals from musicians to the stage. In particular, in a large-scale ensemble when many instruments are involved, in order to preserve the integrity of performance, it may be necessary to allow for the dropping of sound

and hologram streams of a musician that cannot arrive at the same time as the others and to provide mechanisms for subsequent fast synchronization.

- Multi-party holographic communication: In a multi-party remote holographic media conferencing application, or in an interactive multi-player immersive game (such as virtual tennis), a near real-time placement of a virtual object for different receivers is required. Any changes in the position of objects should be rendered in other locations simultaneously; otherwise, the receivers will operate on different views of the digital scenario as they are completely unaware of each other's behaviour. Such different views of the position of objects will certainly happen because the end-to-end path latencies will vary for each sender-receiver pair. To enable fully synchronized operation and to cope with the heterogeneity of delivery paths in the network, some mechanism to provide all parties with the required information at the same time is required.

The Internet is a spatial-temporal heterogeneous environment, yielding different content delivery behaviours in time and space. No two paths (or even different flows on the same path) can be assumed to have identical properties in terms of latency, jitter, and bandwidth. However, as discussed above, many emerging applications require timely delivery of dependent information carried over multiple flows and/or multiple paths.

Network 2030 services therefore need to support "coordinated" communication services, where "coordinated" refers to the need for cooperation among multiple flows with respect to inter-dependent constraints such as time, ordering or any application defined property.

Currently, any effort to support this in the networks is not feasible. Moreover, when left up to applications to manage flow dependencies, they cannot always guarantee absolute time constraints due to unpredictable changes in network conditions. Coordinated services can be used to minimize the complexity at the endpoints and to facilitate dependent delivery from the network.

Contrary to in-time/on-time service technology in which networks need to follow strict guarantees of time, coordinated services need not necessarily be engineered in a similar manner. It is only necessary for this service to guarantee that multiple flows meet the dependency or constraint criteria, these dependencies should be provided as service objectives in Network 2030 as a specific parameter for the service.

8.2.2 Description

A coordinated service provides a guarantee of delivery of multiple flows in a dependent manner. We refer to these co-dependent flows as co-flows for short. Each flow in the co-dependent flow set is referred to as a member-flow. The co-flows may express different kinds of dependencies or relationships. A coordinated service should be able to coordinate delivery of co-flows over different categories of group communications. The mechanisms to support coordinated services in network requires new capabilities referred to as network coordination functions.

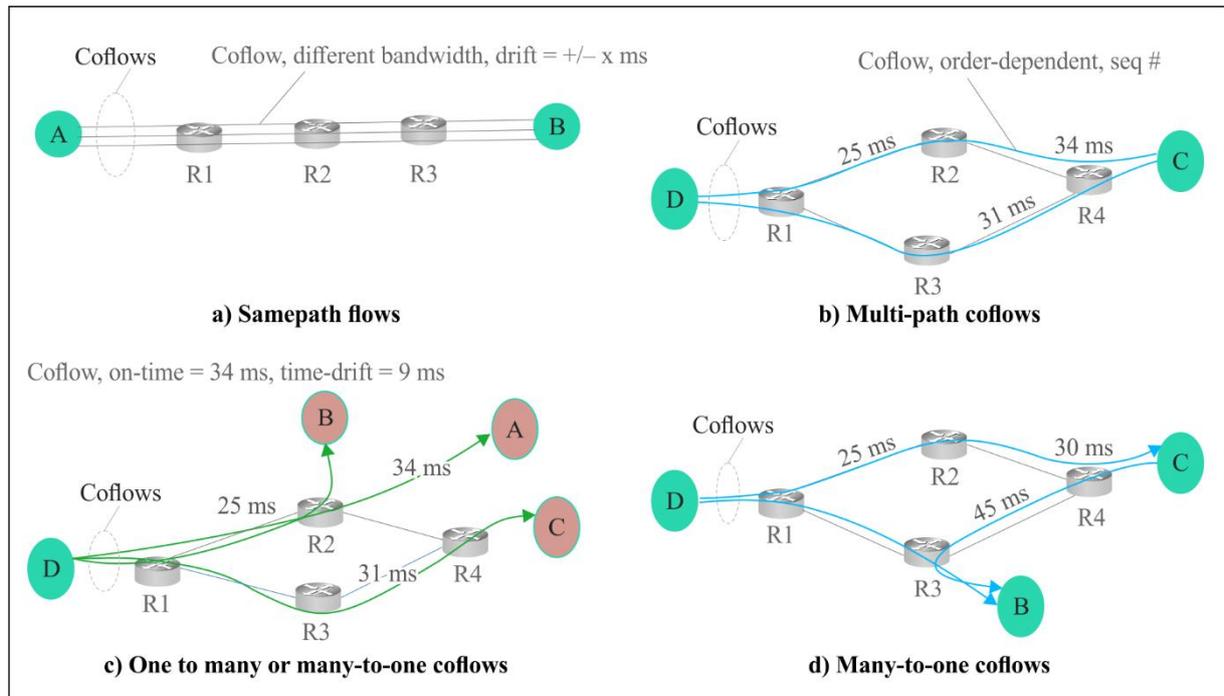
The coordination constraints need to be met end-to-end across the network by all members in co-flows. The following types of dependency objectives need to be met:

- "*Time-based dependency*" is a dependency that requires that the co-flows will meet different time related guarantees. This includes the simultaneous arrival of co-flows at the same or different destinations depending on the category of group communication. This is the guarantee of coordinated delivery and not always concern with meeting hard delivery times as in case of on-time guaranteed services.
- "*Ordering dependency*" is required when member-flows or any dependent part of member flows need to be delivered in a specific sequence.
- "*QoS fate sharing*" is a relationship constraint that specifies that member-flows of co-flows that all need to be consistent with the type of QoS requested by the client. For example, if

one member-flow experiences quality-degradation, then (and only then) it might be acceptable for other members of co-flows to be subjected to the same reduced service level.

Coordinated services also need to consider the various types of communication patterns of member-flows that are required. These patterns are necessary to describe how constraints are met across the network by all members in co-flows. The following types of patterns are considered:

- "*One-to-one communication*", when member flows in co-flows are sent from a single source to a single destination. These member flows may go through the same path or multiple paths as shown in Figure 8-2 (a) and (b), respectively.



Y Suppl.66(20)_F8-2

Figure 8-2 – Different communication patterns for co-flows: (a) samepath, (b) multipath, (c) incast and multicast, and (d) multiparty co-dependent flows

- "*Many-to-one and one-to-many*" Communication, when member flows in co-flows are sent from multiple sources to a single destination or from a single source to multiple destinations, respectively. These two scenarios, referred to as *many-to-one* (from A, B, C to D in Figure 1(c)) and *one-to-many communications* (from D to A, B, C in Figure 8-2(c)), are typical incasting and multicasting, respectively.
- "*Many-to-many communication*", when multiple member flows in co-flows are sent from multiple sources to multiple destinations. This is a fully cooperative multi-party environment, where each party may be both sender and receiver at the same time and may need to comply with dependencies in bidirectional manner. This may be also referred to as groupcasting as shown in Figure 8-2(d).

A client requesting a coordinated service will specify:

- The dependency constraint that needs to be coordinated, such as time-based (same time, relatively sooner or later) or a fate-shared QoS parameter.
- Co-flow identifier under which the required constraint is to be delivered:
 - The expected communication pattern.
 - The behaviour to be executed on failure to meet the coordination objective.

In-network coordination: Coordinated services are a new kind of network service and their requirements emerge from transmission of immersive, virtual, and holographic type applications in a

consistent manner. The network needs to address heterogeneity of available resources and delivery in a consistent manner such that there is one and only one instance of virtual scenario.

In order to realize the Network 2030 coordinated service in networks, in-transit and edge nodes will need:

- "*Knowledge of co-flows*" is needed on border nodes that serve as exit or entry points to a coordinated network. They perform the function of distribution of data transmitted between co-flows and coordinated aggregation and dependency constraint validation on incoming flows. Transit nodes in a coordinated network would perform constraint-based forwarding and efficient replication. Other required network functions may also include signalling in order to exchange rules for the coordination and the management of dynamic membership in co-flows.
- "*Ingress buffering*" may be needed for a certain coordination duration to send co-flows together to the receiving endpoint. Additionally, on egress, *shaping to pace* flows in compliance with the dependency may be required for certain member-flows due to difference in path latencies. This ensures that applications receive co-flows always after the dependency constraints have been met in the network.
- *Awareness of path heterogeneity* (i.e., awareness of differences between paths, such as latency and bandwidth characteristics and constraints) allows the coordinated network to decide whether to slow down or speed up transmission of member flows of co-flows along the path. Alternatively, path heterogeneity must be considered for a member flow also. For example, when the path of a member flow changes, the impact on co-flows needs to be mitigated and coordination among flows maintained.

In addition, the specific type of coordination that is needed may change over the duration of a co-flow. In order to accommodate this, coordinated services could provide specific markers to indicate what coordination needs to occur for which parts of flows.

Concluding notes

Networks that support coordinated services need to be aware of various constraints that span across multiple flows. These services differ from in-time and/or on-time services in that coordinated services may only have relative delivery time constraints (for example, specifying that the latency of member flows must be the same, without specifying a specific latency value). Coordinated networks help in computing the parameters relating to inter-dependencies. For example, the path determined to be of the highest latency may serve as the basis for a coordinated time dependency that specifies member flows of co-flows needed to be delivered at the same time.

8.2.3 Gap analysis

In this clause, the specific requirements for coordinated services are described. Subsequently, an overview of complexity with available techniques in order to deliver coordinated guaranteed services is provided.

8.2.3.1 Requirements

- Coordinated services must provide the ability to support a quantifiable end-to-end coordination of multiple member flows in co-flows delivery across a network.
- Services may be required to ensure that coordinated network elements can perform the tasks necessary to compute, buffer and synchronize flows. Furthermore, services may be required by a client to verify that coordination occurred as expected.
- Services shall support efficient multicast replication in the network, i.e., instead of sending multiple copies from the sender, replication is done on the network nodes where path towards multiple receivers diverge. The coordinated service requires that dependent constraints must also be met with replication.

A dynamic capability to join/withdraw membership from co-flows is required to be supported.

Accordingly, a coordinated service request is characterized by the following parameters:

- Dependency constrain:
 - o time-based,
 - o ordering,
 - o sequencing,
 - o QoS fate-shared.
- Membership:
 - o co-flow identification.

8.2.3.2 Limitations and gaps of existing mechanisms

8.2.3.2.1 Application-level complexity

The need for coordinated services emerges from a number of scenarios where full responsibility left up to the applications to orchestrate and manage coordinated flows can lead to massive complexity in the endpoints while still not being able to guarantee the coordinated delivery of co-flows.

To support coordinated service, and co-flow membership, current applications need to form and manage groups of endpoints and determine and monitor the characteristics of the paths between them for themselves. These applications also need to accurately measure the time of delivery from sender to receiver at every endpoint for all sources and destinations in co-flows.

If coordinated services were to be implemented in hosts, then each host would need to keep track of runtime network state with respect to the dependency constraints. This would further require that the senders would need to manage complex scheduling when transmitting information to different receivers in order to manage transmission times to each receiver corresponding to end-to-end latency over each path, i.e., sending on slower links sooner than the faster ones. The receiver side also would need to provide complex buffer management to buffer received data at the receiver until it is ready to consume it. This would lead to the suboptimal use of memory at endpoints while waiting for other member flow data to arrive.

The fact that many dynamic changes occur inside the network compounds those challenges and creates numerous difficulties for applications to manage coordinated services at the endpoints. An ability to support coordinated network services within the network itself, managing the delivery of member flows according to their inter-dependencies and coordination requirements as a function of the network, therefore becomes critically important to support applications that depend on co-flows.

8.2.3.2.2 Server-based coordination

Nowadays, applications such as teleconferencing require coordination across multiple users. Thus, each member-flow suffers from triangular routing, having to go through the server and then to the receiver, leading to path inefficiency as a consequence. In addition, most server-based applications use unicast forwarding. This inefficient use of bandwidth then becomes a prominent drawback in holographic type future applications.

It should also be considered that a server is a device in the middle that terminates connection before relaying it. Therefore, in such approaches the content will be less secure than when end-to-end direct coordinated communications are used.

In terms of scalability server-based approach requires fewer $(n+n)$ connections than n^2 for peer to peer connections. In-network coordination can still be utilized with server, by providing 2-segment coordination, between endpoints to server and server to endpoints. An alternate server-less approach is discussed in clause 8.2.4.

8.2.3.2.3 Multicast and incast guarantees

Group communication in the network is supported through multicast routers such as protocol independent multicast (PIM) routers in the best-effort manner. The PIM protocol builds distribution trees for multicast forwarding, validates reverse forwarding path, exchanges group membership, and performs replication functions on multicast routers. The protocol neither carries any dependency information nor actively performs coordination. This includes a lack of feedback among members of co-flows as they get delivered at different times.

8.2.3.3 Difference from on-time services

Most examples of coordinated services involve time-based dependency constraints. This raises the question of what coordinated services involving time-based dependency constraints provide as well as what could not be accomplished with in-time and on-time services. After all, an application requiring coordinated timing across multiple flows could choose to utilize on-time services for each member-flows in order to support coordinated service.

However, this approach will suffer from drawbacks mentioned earlier in clause 8.2.2 in that it will now fall upon the application to take into account member flow interdependencies. For example, in order to ensure simultaneous delivery, an application could assess minimum end-to-end latency across multiple paths and subsequently use the longest-latency path as an on-time parameter for each member flow. Subsequently, the application needs to verify that each member flow in the co-flow meets its guarantee so that overall coordination guarantees are still met. At the same time, commitment to a particular latency might not actually be required from an application perspective if all that is required is the same or coordinated latency from each flow, regardless of the particular latency across the co-flow.

In contrast, in-network coordination can be realized by new type of network elements inside the coordination network. By signalling co-flow specifications between network elements, the need to use on-time meta-data (e.g., timestamps or sequence numbers carried in the packets) can be eliminated. This helps with simplifying the maintenance of co-flow state in applications. However, when realizing coordinated services, some of the mechanisms of on-time services (e.g., slowing down/speeding up transmissions based on the current network state) can still be used to provide time-aware scheduling in the networks.

8.2.4 Performance and design target

The following performance considerations apply:

- In order to provide coordinated services across co-flows, each member flow must be capable of meeting the minimal performance criteria of the application.
- In some cases, in-time services may be needed when certain upper bounds are requested by the application. Coordinated services will still be useful over on-time services because they take care of co-flow dependencies.

Design target

Figure 8-3 illustrates coordinated services supported in the network through two scenarios of coordination for multi-sender and multi-receiver coordination, respectively.

- Co-flows originating from multiple senders U4 and U2 are to be received by destination U1 with the dependency constraint, for example "together" without any hard limits of time.
- A member-flow of co-flows from U1 may also need to be received by destinations U2 and U3 with user-defined constraints, e.g., at the same time.

To enable coordinated services, some network nodes may take a new role as coordinated points. These nodes may be the first gateway node for an end point or at a coordinated service enabled node in the network.

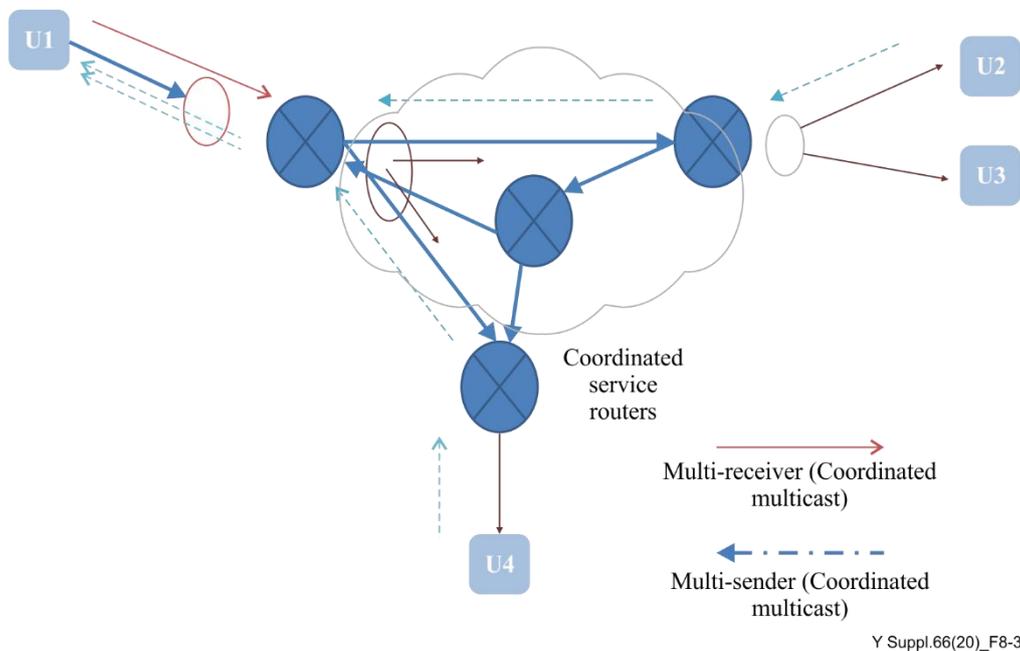


Figure 8-3 – A coordinated service network

8.3 Qualitative communication service

8.3.1 Introduction and motivation

A packet is a minimal, self-contained unit of delivery that gets transmitted, classified, or discarded in its entirety by the network nodes. Whether a packet is a single packet or one of multiple packets in a flow, it is treated as an atomic unit over which network actions are performed.

Where a process of reliable packet delivery is in use, a packet that fails to reach its destination is retransmitted in response to an explicit or implicit indication that the packet was not received. This process takes time as the receiver has to deduce that a packet that was sent was not received, and then must send a request for a replacement, which must be sent over the network. A packet lost at the first hop will thus typically incur latency of more than one round trip time, plus the time for detection and processing, before a replacement packet is available at the receiver. This process is continued until a replacement packet is received at the receiver and can result in significant delay until a critical data fragment is received. This process is similar regardless of whether the packet is being carried over a reliable transport protocol such as transmission control protocol (TCP), or reliability is being introduced at the application layer over a datagram transport protocol such as user datagram protocol.

Packets are lost in a network for three reasons:

- Congestion discard
- Equipment (including link) error or failure
- Bit errors on the links

Of these three causes, congestion discard is by far the most common, but as applications demand extreme reliability mitigation of the other two causes becomes important.

With the migration towards very high bandwidths, there is a tendency to move to larger packets to provide line-rate data transmission. Whether a packet is small or large, its header needs to be processed, hence larger packets offer certain efficiency gains. However, as a result, each discarded packet results in a larger quantum of data that needs to be retransmitted than in the past. Alternatively, in some circumstances, packets can be fragmented into parts to avoid maximum transmission unit issues. This process considers all fragments of equal value and all of the fragments are forwarded to the destination. If one or more of the fragments fails to reach the fragment reassembly point, then the

whole of the packet is discarded. In either case, efforts made by the hosts and routers is wasted as data need to be retransmitted that otherwise would not have to be the case. This violates the guiding principle of "work conservation", which states that systems should perform any work only once if at all possible, ensuring that any outcomes are preserved once they have been achieved to avoid having to redo the same processing steps multiple times.

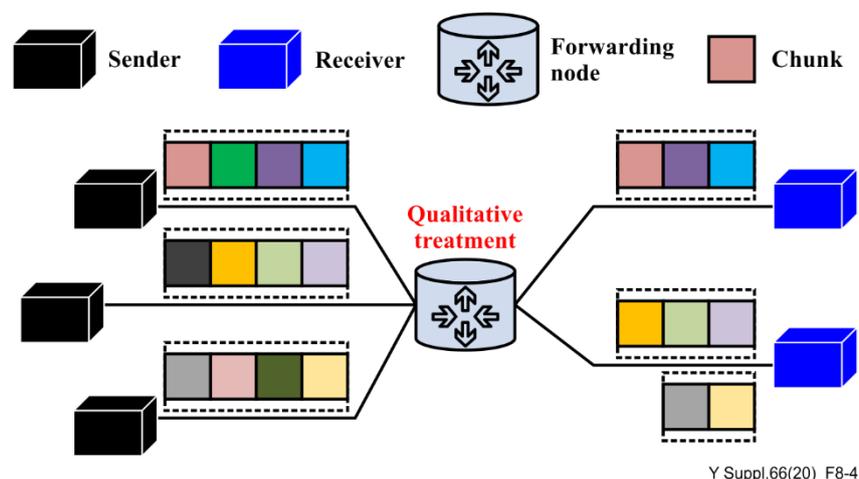
In all the current approaches to congestion avoidance through packet discard, an assumption is made that all portions of the packet are of equal relevance. However, in practice, some payload portions may be more important to applications than others. The qualitative networking approach exploits this fact by allowing senders to group payload within a packet by relative priority, then allowing the network to selectively discard portions of lesser priority when needed.

Specifically, qualitative communication services allow applications to differentiate between different portions of packet payload, referred to as "chunks", and describe their relative priority to the network. Packets carry the necessary metadata needed to describe those chunks. If needed, a lower priority chunk can be dropped from the packet payload while the higher priority chunk can be preserved to continue to their destination. This way, congestion can be reduced and continuity of delivery of critical data to the application, while minimizing the need for retransmission, can be ensured. Qualitative communication thereby addresses both the latency and work conservation issues associated with the approach taken by the established reliable transport protocols.

Qualitative communication does not entirely eliminate the need for re-transmission since it cannot mitigate against irrecoverable loss of critical elements of the packet. However, the amount of information needing retransmission and the frequency of retransmission will be dramatically reduced compared to transport protocols such as TCP, since only critical information would need to be retransmitted which would now also be much less prone to discards than before.

Some applications can tolerate degradation in quality in exchange for timely data delivery. In case of a holographic type live video conferencing scenario under a congested network, it is not tolerable for the end-users to look at a frozen video display. Instead, in order to keep the holographic image smooth and timely, qualitative services could be used to remove the least-significant part(s) of the payload, as indicated by the metadata included in the packet by the source application, e.g., significance level of different parts of the payload, relationship among them, etc.

For example, the source application may choose to deprioritize data that represents the ambient environment, the enhancement video layers, or simply some parts of each image. With proper packetization methods, the network nodes may be able to understand the significance or relationship of bytes in the packet. Then, based on the current state the network nodes can decide which byte(s) in the packet can be dropped with the qualitative treatment, as depicted in Figure 8-4.c



Y Suppl.66(20)_F8-4

Figure 8-4 – Qualitative treatment of packets

8.3.2 Description

As stated in clause 8.3, qualitative communication service allows the users of the service to distinguish among portions of the payload. When network circumstances arise that would currently require the whole packet to be dropped, this mechanism allows the network to drop those less significant or lower priority portions when otherwise it would have been necessary to discard the whole packet. The packet payload can thus be protected by using just enough re-transmission to avoid critical loss.

The qualitative communication service is not an approach that can completely replace the normal existing congestion detection, avoidance, reduction, and notification mechanisms. There will always be circumstances when a router's buffers become full and the dropping of packets is inevitable. However, the qualitative communication service is a substantial improvement over the entire packet dropping, and allows the partial, yet timely delivery of a packet.

The Network 2030 qualitative communication service understands a packet or a flow as a collection of information, where different pieces of information may have different significance or functional relationship with each other. A packet under the qualitative communication service is called a qualitative packet. The qualitative communication service is enabled by two categories of grouping: (1) by breaking down the packet payload into smaller units (called chunks); (2) by grouping parts of a flow into segments such that each segment can have different traffic treatment criteria. Based on the local congestion state, a network node can take forwarding decisions on a chunk or a segment basis. These decisions could be to drop, buffer, or forward one or more of the chunks within the packet. In a more advanced qualitative communication system, if a packet is (partially) corrupted, a node supporting qualitative communication might be able to recover a packet based on the associated recovery function, on the relationship with remaining chunks in a packets, or on data carried in other packets of the flow.

The qualitative communication service therefore enables a much finer granularity of bandwidth regulation than is possible with the current atomic packet approaches to congestion and packet corruption. The meta-information required by forwarding nodes to decide how to treat a packet may be carried with the packet, may be programmed for a flow out of band, or may be intelligently and independently decided by the individual forwarding node.

As shown in Figure 8-5, in a traditional network approach the chunks need to either be grouped together in a single packet or carried individually. The issues with these approaches have been described in clause 8. With the deployment of a qualitative communication service, the chunks are carried in a single packet to reduce packet header, and packet forwarding overhead, but is done in such a way that the important chunks can be salvaged in conditions where a packet discard would otherwise be required. This mitigates the impact on both the application and the network in circumstances where a full packet discard would otherwise be needed.

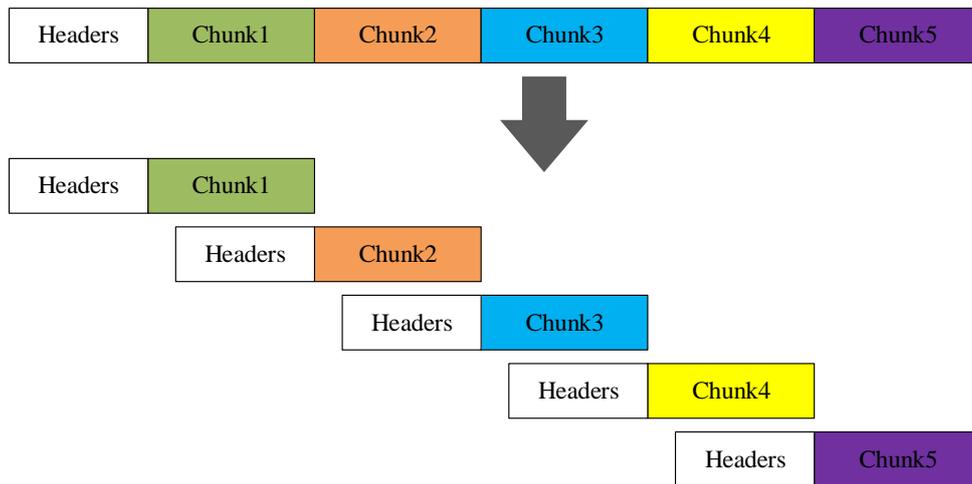


Figure 8-5 – Packetization of groups of chunks vs individual chunk transmission

From the discussion above it can be seen that Network 2030 qualitative communication service provides the following benefits:

- Packet re-transmission may not be required if the receiver has the capability to continue to operate by processing the remnant of the original packet after the removal of less critical chunks from the packet payload by the intermediate network nodes. In this case, the receiver can acknowledge the acceptance of the packet. It may also usefully inform the sender that one or more chunks were dropped by the network thereby allowing the sender to tailor future packets to the congestion state of the network. This reduces wasted bandwidth in the pre-congestion stages of the path and hence frees up resources that would otherwise be wasted. These reclaimed resources are then available to other network users. Network resource usage can thus be reduced and better prioritized for the delivery of other packets.
- The throughput for an individual data flow is the amount of data moved successfully from one place to another in each time period. The qualitative communication service allows the network to deliver more important information in the packets to the destinations, by preventing whole packets from being completely discarded in light of congestion. Effective throughput as perceived by the user is thus less harmed when the network is congested, resulting in a higher effective throughput rate.

To enable and implement qualitative communication service, support from both the application and the network is required.

New packetization: The Network 2030 qualitative communication service makes a paradigm shift from packet-level to chunk-level services. To do this, the first requirement is a new packetization method in which the payload is constructed as a series of chunks and the information needed to extract, prioritise and process the chunks is carried in the packet header. A qualitative packet may carry metadata such as a function or significance parameters that allow the network nodes to know which chunks to drop, and the threshold beyond which a packet must be discarded rather than be further degraded.

In an advanced qualitative communication service, the packet may carry enough error detection and correction information such that the useful chunks may be extracted from a packet that is partially corrupted and would otherwise be discarded due to a cyclic redundancy check (CRC) error. The error rates in optical networks are such that this would rarely be required, but with the higher error rates other types of transmission media might cause such a capability to be of use.

Source application function: The characterization of what information is qualitatively more significant (namely qualitative context) is decided and assigned by the source application or its proxy/delegator. It is necessary that the source application understands the encoding of the user data

in the payload, so that the qualitative context of the chunks can be indicated in the packet. A qualitative context includes a function selected by the source application which can be used to identify the relationship, degree of significance of chunks and/or to help the recovery of lost chunks. This context allows the network to operate on a qualitative packet without needing to look inside the chunk payload.

Optionally, the source application may further re-arrange the positions of the chunks in the payload according to the qualitative context, e.g., significance of the chunks, which helps the network nodes to run the qualitative communication service by, for example, always dropping the chunks from the tail of the packet payload when it is necessary to reduce the size of the packet.

Forwarding-node function: The network forwarding nodes need to perform a packet editing operation by which the chunks with lower significance are dropped from a packet upon congestion while retaining as much information as possible.

Destination application function: Upon receiving a qualitatively treated packet, the destination application needs to decide whether to accept/acknowledge the packet qualitatively; an indication to the sender node of the qualitative outcome, and if/how to recover the original packet through the available metadata and payload chunks. The receiver may send the feedback about its satisfaction level concerning the received packet, whether more information is needed, and which piece of information needs to be retrieved from the source application or from the caching locations.

8.3.3 Gap analysis

The existing transport solutions only operate on full packets and use retransmission mechanisms to maintain the completeness of a data stream discarded due to data loss because of congestion discard, or link errors. In the case of a media streaming application, this packet loss may feed into the codec to cause it to reduce the load offered to the network. However, throughput is a factor of round-trip time and packet loss ratios in the network. Minimizing both leads to a higher effective throughput.

8.3.3.1 Requirements for existing mechanisms and their gaps

Some mitigation against link or equipment failure is possible with techniques such as fast re-route (FRR), but this is limited in capability by the failure detection time and the time to reconfigure the paths. The needs of the demanding applications considered in this Supplement are more stringent than can be met by FRR.

Bit error loss, equipment failure loss and to some extent congestion loss can be mitigated by one plus one system in which a packet is duplicated and sent over two or more paths. This approach is expanded in the IEEE TSN and IETF DetNet approach which perform in network duplication and duplicate reduction. However, all these schemes expand the bandwidth consumed on the network by at least a factor of two which can be problematic with the extremely large bandwidth demands required by some of the applications that Network 2030 seeks to address, for example, holographic networking. Additionally, in many cases the path used to transmit the duplicate packet will be longer than the shortest path available to the primary packet.

Congestion discard can be avoided by using a mixture of ingress traffic shaping and traffic engineering to ensure that there is no over subscription at any point in the network, but this means designing the network for the worst possible case which is not a cost-effective approach, particularly for applications deployed at scale.

The congestion discard problem is further exacerbated by transport protocols such as TCP and (quick user datagram protocol Internet connections) QUIC which in the process of avoiding network congestion drive the network to congestion to detect the maximum data-rate available to their application. They work on the assumption that the occasional congestion discard is worth the price in exchange for operating the network as a whole at its maximum capacity. Some of the demanding applications that Network 2030 considers find even this loss unacceptable. As noted earlier in this

clause, this approach not only introduces significant additional delays for the discarded packet, but it is not work-conserving (i.e. effort made to transmit the packet so far).

It may be possible to avoid retransmissions by delivering partial, yet useful packet fragments to the end user, which is the fundamental characteristic of the Network 2030 qualitative communication service.

8.3.3.2 Additional considerations

1. **Payload opacity:** A cardinal rule of networking is to never look at the contents in the packet and only perform forwarding functions based on packet headers. This rule arises mainly from packet forwarding performance considerations but has become a necessity with the ubiquitous use of payload encryption. Qualitative communication must operate with this level of necessary payload opacity.
2. **Overhead:** Qualitative communication service comes at the cost of adding metadata information in the packet to identify the chunks, and to assist the network nodes in deciding how to edit the packet to reduce its size when this is needed. Thus, on the one hand, the qualitative communication approach reduces the payload size to deal with congestion. However, on the other hand, it is necessary to add additional information in the header to enable this network service. This results in a trade-off between bandwidth utilization in the normal case vs the error/congestion case. Although a complete analysis of this is an area for research, this overhead is likely to be minimal in future networks where the packet payload size is expected to significantly increase.
3. **Compatibility:** The digital representation of multimedia content data, such as image, music, and video communicated in the network is encoded and formatted according to the appropriate standard, for example: MPEG-4, ITU-T H.264, ITU-T H.265. The operation of the qualitative communication service needs to take into consideration that the relationship between the chunks in the packet must be provided to the service by the application. A particular challenge to the introduction of a qualitative service is the need for the SDOs responsible for the media format to develop coding standards that encode the relationship between the chunks in the packet. This needs to be done for both the existing multimedia content and for the emerging holographic applications.
4. **Encryption considerations:** The approach to encryption in a qualitative service will need to be different from the current approach of encrypting the complete packet as a single unit. This is because the network will have to be able to remove lower priority chunks from the packet without destroying the privacy or the integrity of the payload. Clearly at a cost of increased packet overhead and processing cost, each chunk could be individually encrypted without reducing its security. What is more difficult is to cryptographically assure the integrity of the complete packet, i.e., to make sure that only chunks legitimately removed are missing from the packet. This is an area that needs further study and possibly further technical research.

8.3.4 Performance design target

A guideline for the design performance target is that under similar network conditions qualitative networking exceeds the effective throughput delivered to an application by TCP whilst at the same time not driving the network into greater congestion than would have occurred if protocols of reliable packet delivery, such as TCP or QUIC, had been used.

The instrumentation, design and specification of Network 2030 qualitative service performance metrics is for further study

9 Network 2030 services: Compound services

9.1 Haptic communication

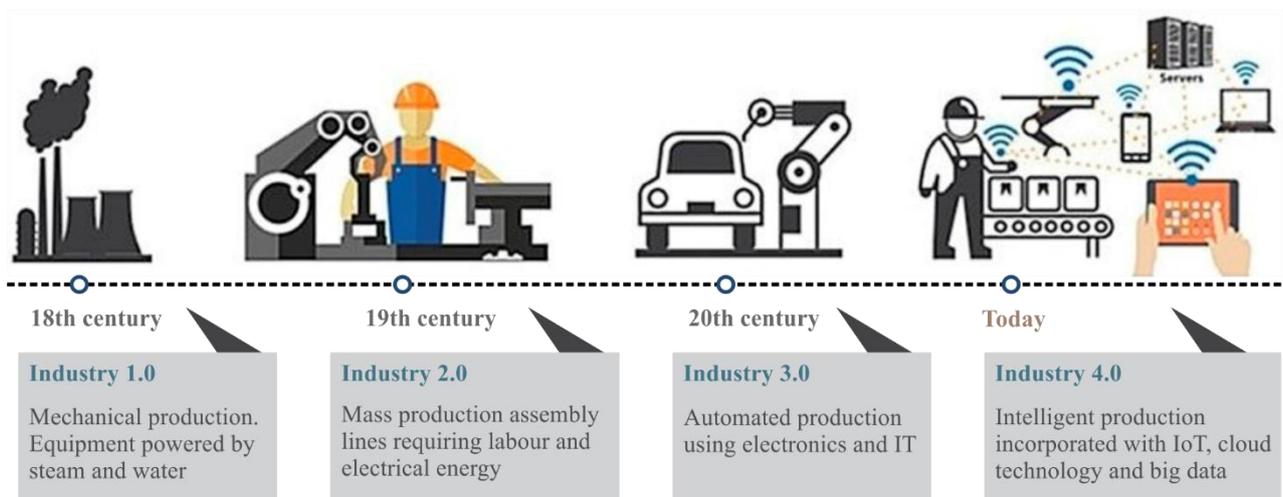
9.1.1 Introduction and motivation

ITU defines the tactile Internet as the network that combines ultra-low latency with extremely high availability, reliability and security [b-Tactile]. The tactile Internet envisions real-time monitoring, management, and control of remotely located infrastructure and devices involving haptics.

In some sense, the term "Tactile Internet" may be a slight misnomer, as tactile is only one of two types of haptic feedback, referring to things that one can feel when touching a surface, such as pressure, texture, vibration, temperature. The other type of haptic feedback is kinesthetic, referring to forces (e.g., gravity, pull) that act on muscles, joints, and tendons in an "actuator" such as an arm, contributing to (among other things) a sense of position awareness. Both types of haptic feedback are important for tactile networking applications. Communication involving one or both types of haptic feedback accordingly is referred to as "Haptic Communication".

Haptic communication is expected to form the backbone of the Industry 4.0 [b-Forbes] along with other application domains such as tele-health, online immersive gaming, remote collaboration, etc. The tactile Internet envisions the creation of a paradigm shift from content delivery to skill set/labour-delivery networks. While traditional networks support audio-visual communication, the tactile Internet will enable haptic communication, i.e., providing a medium to transport the sense of touch (tactile) and actuation (kinesthetic) in real time. Haptic communication accentuate true immersive steering and control in remote environments along with novel immersive audio/video feeds.

As depicted in Figure 9-1, the three previous revolutions in manufacturing were all triggered by technical innovations – mechanization powered by water and steam in the first revolution to mass production and assembling using electricity in the second to the adoption of programmable logic controllers for automation in the third. The next revolution will be triggered by networks that facilitate communication between humans and machines in cyber-physical-systems (CPS) over substantially large networks. Industry 4.0 envisions communication between connected systems, thereby making decisions without human intervention. In order to bring that vision of a "smart" factory into reality, collaboration among the CPS, the Internet of things and the Internet of systems is necessary. The tactile Internet forms the core of such collaboration.

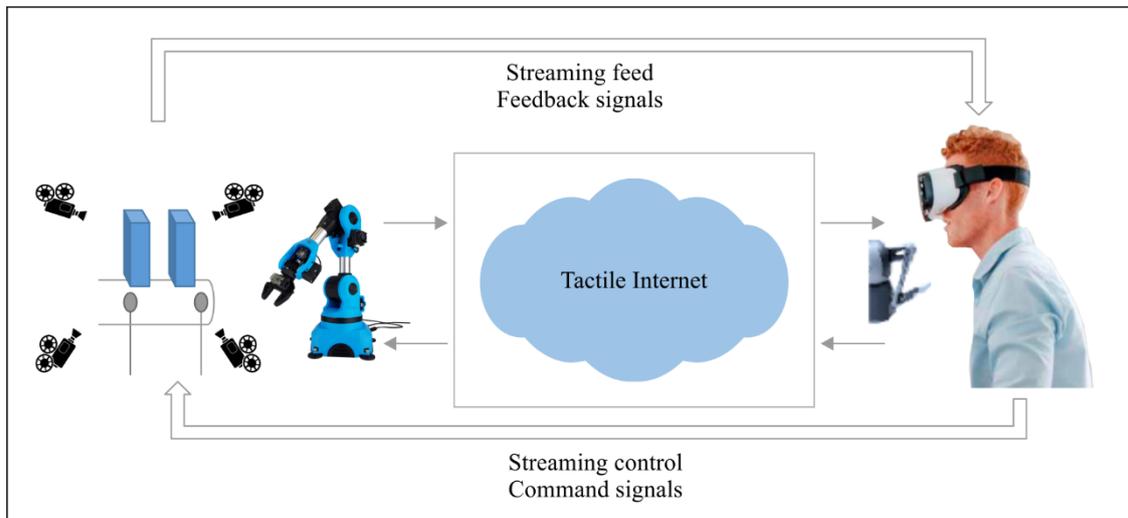


Y Suppl.66(20)_F9-1

Figure 9-1 – Evolution of industry

The stringent ultra-low latency required by haptic communication coupled with novel immersive audio-visual feeds, opens avenues for a plethora of application domains. One use case example

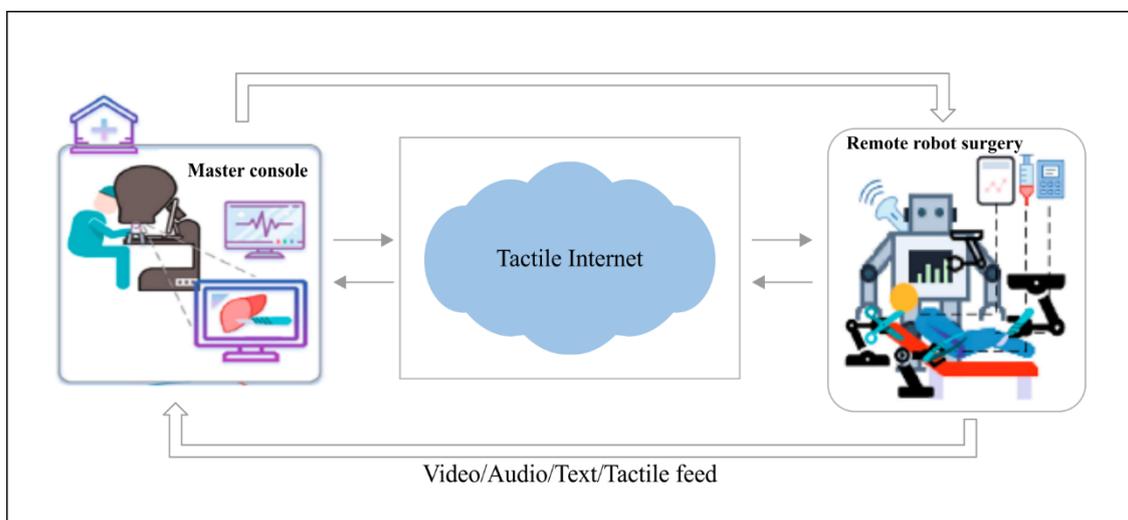
involves remote industrial management. Remote industrial management involves real-time monitoring and control of the industrial infrastructure operation. This will allow a human operator to monitor a remote machine aided by immersive audio-visual feeds, such as virtual reality (VR) video streaming or HTC, and to control the machinery by means of their kinesthetic feedback involving haptic devices, as depicted in Figure 9-2.



Y Suppl.66(20)_F9-2

Figure 9-2 – Remote monitoring and management of an assembling task

Another use case example for haptic communication is remote robotic surgery (Figure 9-3). A surgeon operating from a remote site gets a real-time audio-visual and telemetry feed of the patient and operating room. The surgeon operates remote actuators to perform surgery on the patient, using the audio-visual feed as well as haptic information fed back from the actuator.



Y Suppl.66(20)_F9-3

Figure 9-3 – Remote surgery

Common to each of these use cases is the need for communication channels that are characterized by extremely low latency [b-Aijaz]. There is a strict time-budget for the round-trip time from when an actuator is operated by a human until the tactile feedback is provided. This is on the order of 5 ms or even less. Anything longer and the ability to confidently operate the machinery remotely breaks down rapidly. (Of course, due to the physical limitation of the speed of light, such applications require Network 2030 services to be complemented with application architectures that provide application

service functionality, compute, and content close to the network edge. Efficient support for this may impose additional requirements on networking infrastructure).

While time budgets are slightly longer for audio-visual feedback, the same considerations apply there. Furthermore, because applications are mission critical and retransmission of packets is not an option due to latency concerns, packet loss is not tolerable. As multiple data feeds are involved for data that needs to be rendered and acted on in unison, there is also a need for precise synchronization [b-Torres].

9.1.2 Description

A haptic networking application in general involves two channels that provide a tactile control loop (Figure 9-4):

- A haptic feedback channel, used to communicate haptic data from one or more remote haptic sensors (for example, sensors in a robotic arm) to a haptic effector (for example, a "data glove" rendering tactile sensations to a user). Haptic data includes tactile data, such as surface texture and pressure points, and kinesthetic data, such as force feedback and location/positioning awareness.
- A control channel, used to operate a remote actuator (for example, a robotic arm).

In addition, such applications can involve additional "channels":

- Live visual feed(s) from the remote location (e.g., high-resolution video, immersive video/VR, holograms).
- Live audio feed(s) from the remote location.
- Live telemetry feed(s) from the remote location.

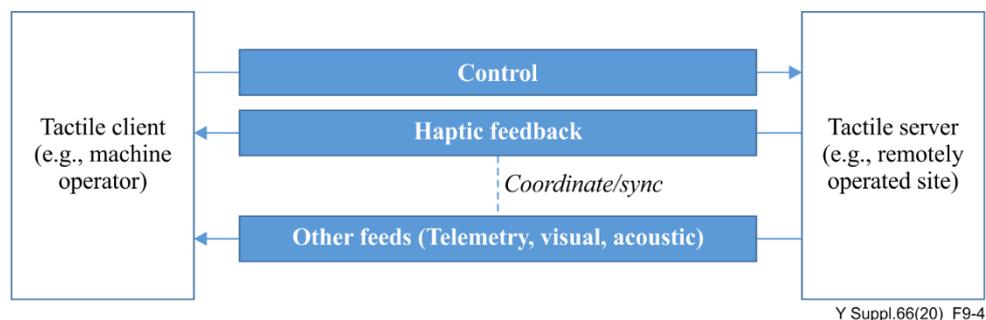


Figure 9-4 – A tactile Internet application

Each channel by itself can be mapped to individual communication flows, e.g., instances of an in-time/on-time service. To improve synchronization among those channels, instances of a coordinated service can be used. However, instead of requiring applications to manage and orchestrate those channels themselves, a network service composition function can be offered, providing a haptic networking service.

A haptic networking service provides a set of coordinated services (clause 8.2), consisting of:

- A haptic feedback channel,
- A control channel,
- Optionally, channels for additional feeds that are bundled with haptic feedback and need to be synchronized with it.

The haptic networking service will ensure that the requirements of a corresponding haptic networking application are met. It will allow clients to specify parameters such as the following:

- Round-trip haptic control latency.

- Haptic codecs being used and/or bandwidth and miss rate requirements associated with the haptic codec.
- Additional channels/feeds to bundle with the haptic control loop, which need to be synchronized with haptic feedback.

This composite service is subsequently mapped onto foundational network services as follows:

- Automatic selection of proper latency parameters for in-time service instances for haptic control and haptic feedback channels. Round-trip latency requirements for the tactile application as a whole are broken down into individual one-way latency requirements for control and feedback channels. Typically, one-way latency requirements for both channels will be the same and together add up to the round-trip latency requirement, but other mappings are conceivable.
- Automatic selection of parameters such as packet rate and acceptable miss ratio for those channels, tuned to the needs of the specific haptic codecs and encodings.
- Additional instances of in-time/on-time services for additional channel feeds. Note that some of those instances can themselves be instances of composite services, as in the case of a holographic feed.

9.1.3 Gap analysis

The main gap that currently exists concerns the ability of networks to support foundational services with sufficiently low latency to realize a tactile feedback channel, coupled with extremely low miss rates to account for the high reliability of tactile applications. In other words, the gap consists of the absence of foundational services that meet the performance design targets as specified below.

A haptic networking service is an example of a composite service that would be reasonably straightforward to provide once foundational services for Network 2030 become a reality, but which cannot be provided by networks at present due to the lack of such services.

The challenge, and gap, lies in enabling haptic communications across larger networks and across wider geographic areas. Performance design targets for ultra-low latency of very few milliseconds, as specified below, quickly run into physical limitations due to signal propagation that cannot exceed the speed of light (300 km one-way in one millisecond, or 200 km in an optical fibre). For those reasons, the distances across which haptic communication services can be offered will still be bounded, and tactile networking applications may need to leverage additional techniques (such as bringing intelligence and compute close to the network edge) to mitigate those limitations.

9.1.4 Performance design targets

1. **Ultra-low latency:** Latency is most crucial for the future high precision networks. The maximum latency that goes unnoticed by the human eyes is 5 milliseconds [b-Fettweis]. For the operation to be smooth and immersive, the new paradigm even proposes sub-millisecond end-to-end latencies for tactile feedback. Although it varies depending upon the use case, in general the tactile Internet envisions an end-to-end latency in the range of (1-5) milliseconds [b-Miao].
2. **Ultra-low packet loss:** In such critical applications loss of information means loss of reliability on the system. In addition, retransmission is generally not an option due to latency concerns. Hence loss should be as close to zero as is practical.
3. **Ultra-high bandwidth:** The bandwidth requirement is especially important in case of remote monitoring as increasing the complexity of the visual feed (from 360 degrees video to holograms) makes the required bandwidth grow drastically as well. A bandwidth up to 5 Gbit/s is required for VR feeds [b-Brettel] and it increases up to 1 Tbit/s for holograms [b-Torres-3D]. The complexity increases with the numbers of streams.

4. **Strict synchronization:** The human brain has different reaction times to different sensory inputs (tactile (1 ms), visual (10 ms) or audio (100 ms)). By themselves, some streams (e.g., audio) might thus allow for slightly higher latency than others (e.g., tactile). Nonetheless, synchronization is important, even in the presence of ultra-low latency, as synchronization needs to be on time scales still significantly shorter than latency. This means that tolerable latency for e.g., video might be lower in scenarios when the visual information needs to be synchronized with tactile feedback than in other scenarios where no tactile feedback is involved.

In addition, the network should be capable of prioritizing streams based on their immediate relevance. Since the visual feed involves multiple views and angles for immersive media, the relevance of such different streams should be considered and the ones with higher importance to the operator's view and current task should be given higher priority.

9.2 Holographic-type communications (HTC) services

9.2.1 Introduction and motivation

The use of holograms as a means for users to interact with computing systems has long captured people's imagination, as evidenced in movies such as "Star Wars" or "Minority Report". As holographic display technology has made significant advances, holographic applications are well on their way to becoming a reality. Many such applications will involve network aspects, specifically the ability to transmit and stream holographic data from remote locations across the network to render it on a local holographic display.

Examples of such applications abound. For example, holographic telepresence will project remote participants as a hologram to local meeting participants in a room. Remote troubleshooting and repair applications will allow technicians to interact with holographic renderings of artefacts located in a remote location. Training and education can provide users the ability to dynamically interact from remote with ultra-realistic holographic objects for teaching purposes. Audio-visual feeds for robotic tele-surgery, as mentioned in clause 9.1.1, can involve holograms as well. Then there is immersive entertainment, gaming, sports, and much more.

It is easy to foresee that the majority of those applications will involve HTC, i.e., the ability to transmit and stream holographic data across networks. Rather than representing simply yet another media type, there are several unique aspects about holographic data that pose significant challenges to networks.

The following background is intended to help explain some of these challenges. In a hologram, the same image is captured from different viewpoints, tilts, and angles. Depending on the position of the viewer relative to the image, a different "field" in an array of images is seen, with each image depicting the same "object" or "scene" from a slightly different viewpoint. For example, in the case of a smooth holographic representation that features differentiated images for every 0.5 degree variation in viewer angle, a two-dimensional array of 800 separate images is needed to accommodate 20 degrees differences in viewing angle over 10 degrees of tilt. The raw amount of bandwidth required is enormous; however, clever compression/decompression schemes across the image array allow the encoding and rendering systems to exploit the fact that individual images in the array include only minute differences. Another option of representing holographic data is through use of point clouds consisting of volumetric data. In this case, objects are represented as "point clouds", i.e., sets of three-dimensional "volume pixels", or voxels, in a conceptual three-dimensional box. Instead of streaming arrays of images, volumetric media data is streamed. The actual image can then be dynamically rendered from any viewing angle at the local endpoint, placing the point cloud object into a scene or even rendering multiple point cloud objects simultaneously.

In order to reduce the amount of holographic data to be streamed, applications are expected to take advantage of techniques such as user interactivity prediction schemes. The goal is to minimize the volume of data that needs to be transmitted while maintaining acceptable quality. This occurs by

focusing on the data that will likely have the highest effect on quality first, for example transmitting image data of fields that are in focus at the highest quality, while transmitting other images at lesser quality (e.g., reducing resolution, frame rate) or not at all (e.g., dropping certain tilts and angles). Since the user may change viewpoint or position, supporting such schemes requires highly adaptive and ultra-low latency control schemes to be able to adapt streamed holographic contents as needed.

9.2.2 Description

HTC services will provide a set of network services that can be used to transmit streams with holographic data, i.e., data that can be used to render holographic images.

There are different flavours of holographic data streams that may need to be supported:

- Point cloud based, i.e., the sender sends volumetric data objects from which holograms are rendered at the receiver side. In many cases, a volumetric data object can be decomposed into multiple, smaller volumetric data objects, e.g., "3D tiles". Depending on the viewpoint and position of the end user on the receiver side, some data objects may be obstructed at any one point in time. To preserve bandwidth while maintaining high image quality, HTC services need to support rapid "switching" between different data objects as they come in and out of view. This allows the pool of available bandwidth to be preferably applied to those data objects that will be in view and in focus.
- Image array based, i.e., the sender sends an array of images instead of a point cloud. Analogous to point clouds, depending on the viewpoint and position of the end user on the receiver side, different fields in the image array may need to be prioritized. Again, HTC services need to support rapid "switching" between different fields in the array, prioritizing the image quality of some feeds over that of others as they come in and out of user view.
- Multiple camera feeds, i.e., a set of senders send a series of two-dimensional images, possibly coupled with depth information. In that case, HTC data is sent "raw", not pre-processed, and feeds get combined at the receiver side to result in one holographic image/point cloud.

In each case, a resulting communication service can be characterized as follows:

- It can involve multiple channels of holographic data (e.g., one per component point cloud or 3D tile in case of volumetric data, one per field in an image array, one per camera feed).
 - Each of these channels may map to a separate flow with stringent in-time requirements to ensure an internally consistent/synchronized holographic rendering.
 - Some channels may have differing resilience requirements – a drop of some data, while not desirable, may result in a slight degradation in quality of experience (QoE) for users but still yield an acceptable result. In some cases, a drop of data in one channel may lead to the data in that channel to be deprioritized completely – it may be preferable to deprioritize one channel versus other channels (or drop it completely) instead of having uniform slight degradation across channels. (However, multi-dimensional compression across different fields in the image array can occur. In such a case, resilience requirements may be dramatically increased, and different prioritization schemes may apply.)
 - Aggregate resources for the totality of holographic data may be shared (resulting, e.g., in a requirement for "aggregate bandwidth") and may need to be continuously reallocated among the channels (as optimization schemes continuously adapt which contents to stream based on user interactivity and parts of holographic images coming into and going out of user focus).
- It may involve an additional channel of "manifest data" that indicates how to compose the holographic image from the multiple feeds. This data needs to be especially protected, as any corruption of data may render other holographic data useless.

- It will involve a "back channel" to control transmission and prioritization between 3D tiles or image array fields, as end user viewpoints shift, and different parts of the holographic data come into and out of view.

An HTC network service will allow clients to specify parameters such as the following:

- The number of channels for the holographic data.
- The aggregate bandwidth that can be allocated among the channels.
- The acceptable end-to-end latency specified as an in-time requirement that must be met by all holographic channels as well as any manifest channel.
- The latency that is needed for the back channel (which determines how much in advance user interactivity and changes in user viewpoint need to be predicted and adjustments of individual channel feeds needs to occur).

The HTC network service can be composed from a set of coordinated services, consisting of:

- A set of channels to carry holographic data from holographic source to destination/rendering endpoint. Each of those channels will share the same in-time requirement. In addition, the aggregate bandwidth of each channel must not exceed the overall bandwidth allocated for the coordinated service.
- A channel to carry manifest data from source to rendering endpoint. The latency of this channel must not exceed the latency of any of the holographic data channels.
- A control channel in the opposite direction (from rendering endpoint to holographic source), to adjust manifest and streamed data as needed. (This channel can be provided through a separate instance of an in-time service and does not need to be included as part of the coordinated service.

9.2.3 Gap analysis

The main gap that exists concerns the ability of networks to support foundational services with sufficiently low latency (in-time services with quantifiable latency) and sufficiently high bandwidth. In addition, existing technology does not facilitate the notion of aggregate bandwidth shared across and dynamically reallocated among a set of flows.

Like tactile networking services, an HTC networking service is an example of a composite service that would be reasonably straightforward to provide once foundational services for Network 2030 become a reality, but that cannot be currently provided by networks today due to lack of such services.

Another challenge lies in the ability to deliver holographic-type data with very low latency. A lack of low latency can partially be traded off against an increase of bandwidth: Higher latency implies that the time horizon of user interactivity prediction needs to be longer, respectively that there is enough additional data provided to "tide the user" over while adaptations among the data channels occur (so that different fields in the array or different 3D tiles in the point clouds can be transmitted in higher quality).

9.2.4 Performance design target

- Low latency: The latency requirement is on the order of 10 ms to allow instant viewer position adaptation at 60 frames/second. However, the latency requirement can be relaxed e.g., for lower frame rates and at the expense of higher bandwidth. It can become as low as conventional interactive video (on the order of 100 ms, gated by latency requirement for interacting with the remote party, not by latency requirement regarding viewpoint prediction).
- Ultra-high bandwidth: Required bandwidth may start from roughly 1 Gbit/s and increase up to 1 Tbit/s [b-Finn] (Figure 9-5) but depends heavily on encoding and trade-offs regarding bandwidth and compute for optimization schemes. A feed from a current commodity RGB-D

sensor like Intel Real Sense or Microsoft Kinect generates roughly 2 Gbit/s of raw data (for 512*424 pixels with 2 Bytes of depth data) but can be compressed further.

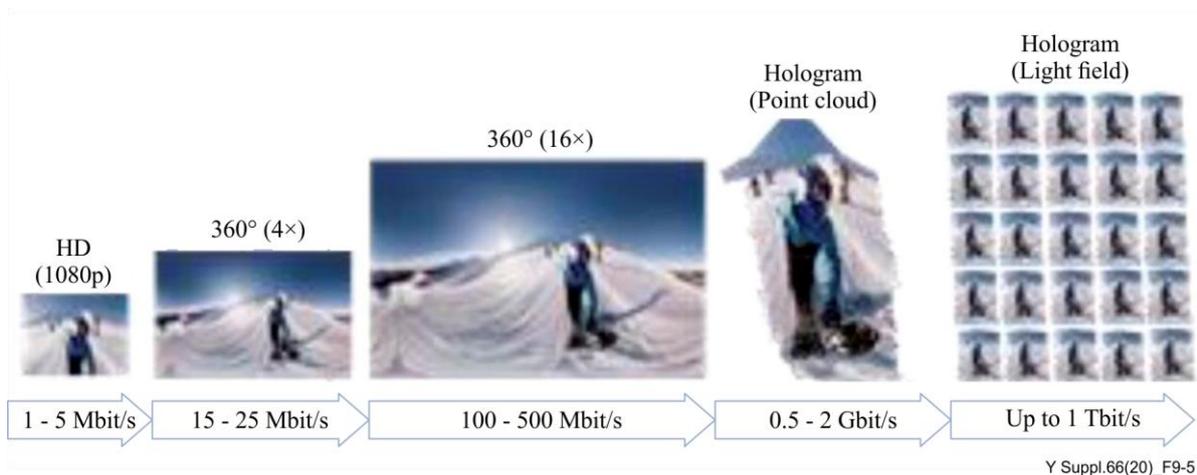


Figure 9-5 – Bandwidth requirements for Holographic-type communications

- Strict synchronization. At 60 frames/second, latency variation across channels should not exceed 7 ms (duration for half a frame).
- Support for concurrent flows. Depending on point cloud and image array dimensions, on the order of 1 000 concurrent flows may need to be supported.
- Ultra-low miss rates, specifically in the presence of strong compression techniques.

In addition, the network should be capable of prioritizing streams based on dynamic and varying criteria (related to viewing position and user focus).

10 Other aspects and capabilities of future networking services

In addition to supporting new services, other capabilities will need to be provided and requirements addressed by Network 2030. While the focus of this Supplement lies on the services, these other aspects need to also be taken into consideration. Clauses 10.1 to 10.9 provide a brief discussion of many of those aspects.

10.1 Network service interfaces

[b-IETF RFC 1633] states that while both the network and applications will evolve, the need for compatibility requires that service interfaces remain relatively stable. This principle remains true today: while underlying networking technology will continue to make strides and novel Network 2030 applications are emerging, network service interfaces should remain stable and evolve in ways that preserve compatibility.

This implies that Network 2030 service interfaces will take an evolutionary approach and support well-known interface patterns including sockets. Of course, network service interfaces will need to allow applications to access and take advantage of Network 2030 services, requiring these interfaces to evolve to provide support, e.g., for high-precision services (Figure 10-1).

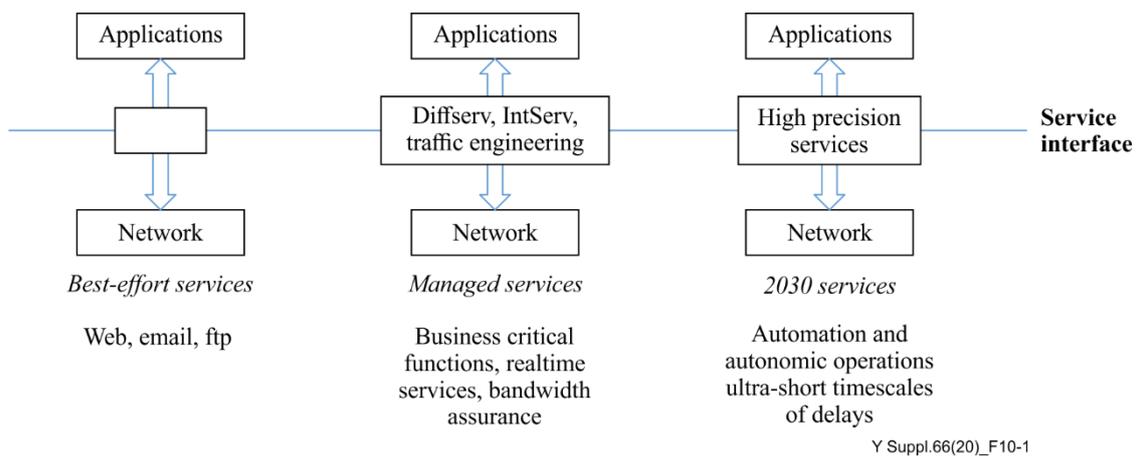


Figure 10-1 – Evolution of network service interfaces

Consequently, certain advances will be required. Specifically, network service interfaces will need to account for new network service parameters. For example:

- Latency targets (for in-time and on-time services) need to be dynamically negotiated between application and network.
- Co-flows (for coordinated services) need to be composed and component flows referred to.
- Payload relevance (what parts can be dropped, and what cannot be – for qualitative services) needs to be designated.

In addition, network service interfaces may need to account for additional properties, for example, to be able to properly account for and validate services delivered across the interface, to accommodate additional trust mechanisms, and to support the additional instrumentation required for manageability purposes. These aspects will be discussed further in clauses 10.3 to 10.8.

10.2 High programmability and agile lifecycle

Network providers in 2030 will need to be able to rapidly introduce new network services, or network services with properties that need to be rapidly adapted to new contexts, deployments, and application needs. Likewise, the business landscape may require users of network services to be able to rapidly adapt services to their 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 SDN controllers, virtualized network functions and network function virtualization (NFV), and service function chaining will no longer be sufficient.

Supporting new networking services through custom network appliances, as is still customary today, suffers from lengthy product development cycles for hardware and firmware. These lifecycles typically take several years and can be undertaken only by equipment vendors, not by network providers and the users themselves. This stifles progress and slows the introduction of new network services. Assuming that somehow network providers can get their requirements addressed, they still have no choice but to wait through lengthy vendor development lifecycles. More recently, SDN and NFV have opened the possibility of accelerating development lifecycles and enabling network providers to develop new networking features on their own. Segment routing is being evolved for that purpose as well. Furthermore, network slicing described in [ITU-T Y.3112] promises more agility in the introduction of new network services. Programmable packet processing technology such as P4, despite many limitations, facilitates the rapid introduction of new protocols in support of new services.

However, the complexity of the associated controller software results in its own challenges with software development cycles that, while more agile than lifecycles 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 behaviour at the network flow level, without the need to introduce application-level over-the-top overlays. Such a capability would be analogous to lambda programming 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 the network services described in this Supplement, which require hardware advances at many levels to provide programmable flow and QoS behaviour at line rate, affecting everything from queuing and scheduling to packet processing pipelines.

Network 2030 will require advances in the programmability of networks that allow development lifecycles that are much more agile than those 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 programming models. It also requires the development of new programmable networking hardware, which offers the ability to map novel network programming primitives into packet processing pipelines and the execution of custom logic at line rate.

10.3 Manageability

Support for Network 2030 services will require advances in manageability to be able to successfully provide and operate such services. Such advances will need to proceed in lock step with advances in the services themselves.

Manageability for assurance of high-precision Network 2030 services

Network 2030 services have in common that they will 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. This is needed in order to monitor and validate that promised service levels, such as latency and miss rate, are being delivered. Among other things, this will require the need to support instrumentation with much higher precision than today:

- Accuracy: Where in the past measurement accuracy in the millisecond range has been sufficient, accuracies on the order of 10s of microseconds will likely be required.
- Coverage: Coverage of what is being measured needs to be dramatically improved. Past sampling techniques will run into limitations when it comes to, for example, detecting violations of service levels that may occur only in one in 10^{12} packets. Instead, measurements may need to be incorporated as a basic feature of network services themselves.

Active measurements that make use of synthetic traffic that is generated specifically for measurement purposes are a common practice today. However, given accuracy and coverage constraints, active measurements alone will be even less sufficient in Network 2030 than is the case today. Passive measurements based on the observation of production traffic can equally be applied, but can face various limitation, from the encryption of traffic to legal considerations associated with the snooping of production traffic that belongs to users, not the network provider. Currently, hybrid active/passive measurement techniques are being proposed [b-IETF RFC 8321] and analytics techniques is making rapid advances, all of which promise to mitigate those problems, but further advances will be needed, including advances in the instrumentation of networking devices for these purposes.

By the same token, the ability to identify and eliminate potential sources of service level degradations and fluctuations will become of increasing importance. This requires further advances in instrumentation, telemetry generation, and tracing capabilities. IOAM [b-Brockners] is a promising

technology currently under development that will allow to collect packet telemetry for that purpose and that points in the right direction, but will not be sufficient. Some of the challenges that will need to be addressed include the very high volume of data that gets generated (one data item per parameter, per traversed node) and the effects of the collection on performance (data records get piggy-backed onto packets during network traversal, potentially leading to higher network resource consumption, latency, and jitter).

In general, greater emphasis will need to be placed on the ability to monitor, observe, and validate the compliance of actual with expected network behaviour than is the case today. Data to be generated from the network will need to be "smarter", i.e., more insightful and actionable. This will require additional abilities to process data "on-device". In addition, the need for new management functions may arise that will require novel capabilities, such as:

- Functions that allow to validate adherence with agreed-upon service levels.
- Methods that prevent data or privacy leakage.
- Methods that provide evidence for the possibility or absence of such leakage.

Manageability for fulfilment and "operation-at-scale" of Network 2030 services

Another challenge will involve enabling operators and users to manage network 2030 at scale. This will require further automation and the closing of management control loops. In the past, where possible and where routine tasks are involved, human operators have been increasingly taken out of the loop and replaced with management systems and controllers that were in most cases hosted in a central location or in the cloud. The ever-increasing need for shorter control loops means that management services will increasingly need to migrate closer to the edge of the network and indeed into devices themselves.

However, despite all those advances, networks will not become clairvoyant and will still need to be given guidance for certain tasks and require some degree of human interaction. For this reason, advances in abstractions will be required to facilitate the ways in which operators can interact with networks. These abstractions are needed for productivity reasons (operate at greater scale) and to constrain complexity (greater heterogeneity, growing number of interdependencies which are becoming less understood, etc.).

Technologies such as Intent-based networking [b-Clemm], which will allow networks to be managed by defining outcomes rather than prescribing rules or procedures, are expected to provide significant contributions here. While vendors frequently tout their controller interfaces and policy frameworks as "intent interfaces", true intent technology is still in its infancy. For example, intent technology will require novel human/machine interfaces that allow to iteratively infer and refine intent. It will also require advances in the application of AI and Machine Learning technology that are able to automatically define and continuously refine plans of actions that generate desired outcomes.

Furthermore, in order to meet scalability challenges, novel management architectures may need to be supported that support greater management functionality in a distributed or decentralized manner across the network, as opposed to relying solely on centralized management systems and controllers as predominantly the case today.

10.4 Security

Network 2030 services will need to be secured. Network services such as coordinated communication services or qualitative communication services defined in this Supplement can be originated at a network ingress point and consumed from an end host or network egress point. Additional security mechanisms are needed beyond those that are provided by traditional transport mode IP Security.

A key security aspect needed from the network point of view concerns the need 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 unless network nodes can still 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 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 tunnelled packets without exposing any of their contents.

Some Network 2030 services provide additional security challenges, for example:

- Coordinated services (clause 8.2): by attacking a single member flow, the co-flow as a whole as well as other member flows could be compromised. For example, an attack that introduces additional latency on a member flows might also slow other flows depending on inter-dependencies.
- Qualitative communication service (clause 8.3): with this service, it is no longer enough to secure packets and ensure their integrity as a whole. Since chunks, i.e., certain portions of the payload, might be legitimately dropped, packets and payload need to also be secured at the individual chunk level.

Another dimension to security arises when the end to end service that needs to be delivered crosses the administrative boundary of the originating host. Here, apart from the above considerations regarding authorization and integrity protection for a single domain service, additional mechanisms need to be specified to sufficiently ensure the privacy and confidentiality of the network layer information. While there are lots of avenues to tackle these issues and some aspects are being investigated 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 security mechanisms set in place by the end host/user at the transport layer and above. 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 information that is applied by those layers in order to deliver new services. This way, layer violations are avoided.

10.5 Resilience

Resilience is the ability of a network or system to provide and maintain an acceptable level of service in the face of various failures and challenges (e.g., attacks) to normal operations. At the network (topology) level, resilience amounts to preserving loss, jitter, and latency as successfully as possible for a given service – all these QoS metrics can be compromised if failures/attacks occur and if there is a lack of resilience mechanisms to remediate/mitigate them [Sterbenz, 2010]. At the service level, the relevant QoE metrics that need to be maintained are availability and reliability.

Resilience takes on additional importance for Network 2030 services, because in many cases these services are used for mission-critical applications and require high-precision, moving beyond "best effort" that was acceptable for many applications in the past.

- Network 2030 services are characterized by the need for high precision timing (e.g., in-time and on-time services) and synchronization between large numbers of flows (coordinated services). Any network degradation puts these services in jeopardy and makes the applications that rely on them infeasible.
- Where degradations are acceptable, the mechanisms and extent of degradation need to be controlled more precisely than in today's networks (e.g., qualitative services). Hence there

will be a much higher demand for resilience (and how resilience is integrated into the network service).

The ultra-low-latency requirements, and the huge increase of bandwidth demands of Network 2030 services such as holographic type communication services, make an unrecovered failure a significant loss for network operators. Therefore, network resilience becomes of paramount importance to maintain the network QoS, high availability, and reliability of these new and extremely demanding services.

There are many methods for providing network resilience. The first is to provide redundancy and diversity of logical and physical entities. Logical entities include network paths and functional entities such as data plane functions, including shaping, policing, classification, and scheduling. Physical entities include ports, routers, and router line cards. The second method is to use protocols to provide quick re-convergence and to maintain high availability of existing connections after a failure event occurs in the network. Among the other techniques is the use of packet replication, network coding, and error correction to overcome packet loss.

To support the ultra-low-latency and lossless networking requirements of Network 2030, the switchover from the primary entity to the backup entity must be very fast in the order of micro-seconds or even less.

When traffic is rerouted from the primary path to a new path, the new path should provide the same network high precision communications services that are available in the primary path.

Although redundancy and diversity enable high availability and reliability, they impose higher costs for realization of the network service. In order to keep such costs at an acceptable level, the addition of redundant instances needs to be driven by the target resilience level that needs to be achieved. Cost effectiveness with regard to network service implementation must be kept in mind. Furthermore, applying redundancy and diversity might impose the additional complexity of managing the redundant instances and updating their states in order to maintain their ability to promptly take over the functionality of faulty instances.

Service level agreements (SLAs) for Network 2030 are expected to cover appropriate resilience objectives that indicate the importance of the service in terms of expected availability and reliability. This statement of intent will be mapped into the additional resilience measures to be taken to avoid violating the SLAs. In stating SLA requirements, a business will have to specify matching resilience requirements not only for the network, but also for their applications. The application might need to be available despite a catastrophic failure in a specific region. The application might have to respond to user requests in a specific amount of time. Since network or application failure can also be caused by security attacks, the network operator needs to apply relevant security policies and provide necessary tools to detect and mitigate these attacks or prevent them.

The assurance of resilience in future networked systems (viz. in Network 2030 applications) will be addressed in the work of network management/orchestration (including application management).

10.6 Loss-lessness

An aspect of resilience that is of special importance concerns the avoidance of loss. Avoiding loss and achieving networking services that are lossless is an important objective for mission-critical applications that require high-precision and low latency and that, as a result, cannot afford to rely on retransmission and reliability schemes provided at the level of the application. While no service will be able to guarantee zero loss due to the possibility of some catastrophic cosmic event, loss as a result of single equipment or link failures should be ruled out.

A major source of packet loss is tied to the classical problem of congestion and limitations in network resources (bandwidth, buffer spaces) as a result of competition between too many concurrent packets and flows. One way to avoid loss is to avoid oversubscription of resources to ensure that congestion cannot occur. However, in general this leads to poor network utilization, as most network traffic does

not occur at constant rates that are known in advance and the economic benefit of statistical multiplexing can no longer be taken advantage of. Accordingly, a trade-off needs to be made between loss-lessness and cost: achieving loss-lessness can result in high cost, which may in some cases be prohibitive. This problem is compounded if, in order to protect against the possibility of link or equipment failures, network traffic is sent redundantly over multiple paths.

The challenge thus concerns how to achieve loss-lessness while keeping cost acceptable.

Possible mitigation techniques include machine learning and AI techniques, to recognize the possibility of resource contention early and to dynamically adjust packet forwarding in such a way that congestion and thus loss are avoided even at high utilization levels.

10.7 Privacy

In recent years, there is a growing awareness by the general public of the lack of privacy on the Internet. Any new network service introduced must comply with heightened user privacy expectations. Network 2030 services will need to take those concerns into account and address them, balancing the rights of network users for privacy with the legitimate needs of network providers to operate and maintain their networks.

The definition of privacy of a user is currently still a grey area and few users are aware how their personally identifiable information (PII) is being tracked, shared and with whom. The recent regulations that went into effect had the benefit of widely publicizing that PII is being shared and to require that users consent to the sharing. However, many regulations have proven to be complex and inconvenient for end users, led to inconsistencies across countries, and are difficult to enforce.

An added difficulty is that breach of privacy for a user may take several forms depending on whether it is an observer (authorized or legitimate entity) or an eavesdropper (unauthorized) or both.

- Access to the user data in the packet (both): The user data may be encrypted and be opaque to an eavesdropper to prevent this type of breach. However, data flow analysers can recognize patterns of the type of information exchanged by analysing unencrypted packet headers, as well as observable packet properties such as packet size and traffic characteristics such as packet arrival patterns.
- Trackability of a device location of user by observer (observer): For example, the Internet service provider access point or wireless authentication to the network have some trackability needs for billing purposes that are unavoidable to provide a service.
- Trackability of a user location and patterns by correlation (eavesdropper): The observable packet headers, readily available information (whois), cookies combined with cross-sharing of PII is insidious.

In general, the normal routing services interpret non-user data information in the packets to provide the service. However, the combination of data analytics, with the PII such as addressing and third party sharing of information create an opportunity to track the user and observe patterns. The challenge of maintaining privacy is that PII may be part of the packet that needs to be interpreted. The solution to greater privacy and protection of the PII will need both a technical solution at the network level as well as regulatory solution.

From this, we derive the following requirements with regards to privacy for Network 2030 services:

1. Anonymization: To prevent tracking by eavesdropper by packet capture, the visible information in packets such as source and destination addresses should be more difficult to directly correlate to PII.
2. Opaque user data; Network 2030 services must not rely on the user data to provide the service but rather on specific service-visible data in the packet. For example, this information may be the service level parameters for the data in the packet. These parameters are distinct from the user data which need to be opaque.

3. Secured storage: For some Network 2030 services which may require the network to slow down the delivery of the packets, this implies that the packets are temporarily buffered on the router. The storage of those packets should be secured in such a way that it is not easily duplicated or stored for later deep inspection or analysis.
4. Flow anonymization; Data should be obfuscated but the flow 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. Services such as qualitative analysis may provide more fluidity in the traffic patterns for hard correlation.

To meet these requirements, the following gaps with regards to today's technology need to be addressed:

1. The current architecture and best practices implement long lived address allocation that makes it easy to track a user (specifically if the user is using his home internet access). The conjunction of this PII and the third-party cookies allows using multiple websites to exchange information and glean much more information without the control of the user.
2. The user data is expected to be encrypted in future networks as they are today and any service level parameter should not be considered as user data but rather a qualifier for the user data which may be interpreted by the network devices. However, today most of information in the headers are in the clear and therefore observable by eavesdroppers. The information in the clear should be kept to a minimum or encrypted as well when possible.
3. Routers handling packets should be able to secure the packets stored in their buffers and prevent misappropriation of this information. The storage buffers should be encrypted or at least be protected by mechanisms that prevent access from outside the router.
4. Traffic data analysis requires consistency to determine patterns. Today it is possible with DPI telemetry to determine the type of traffic. These analysers rely on distinct flow characteristics and packet sizes. For data to be truly anonymized, randomness in the flow and packet sizes is required.
5. The definition of what constitute a PII should be consistent across the network and the applications so as to better protect the user.
6. Regulations should evolve more rapidly in response to the technological changes to close loopholes more efficiently.

Furthermore, users should be able to reconstruct what happened to their packets and data:

- Which systems had access to it?
- Which geographies were traversed?
- Were packets buffered or stored along the path?
- Could packets have been subjected to being copied or diverted, by whom?
- Can a network provider provide provable evidence of the presence or absence of such privacy-related occurrences?

Unfortunately, this information also reveals information about networks which network providers may not be willing to share, and which introduces the possibility for attacks on the network or the user traffic traversing it. Clearly, the ability to provide such information, and to do it in a way that does not compromise security concerns of network providers, exceeds today's technical possibilities and points to research challenges.

The following hints at some of the possibilities of solutions that could be pursued:

- Information in data packets should aim to have less significance to an eavesdropper by stronger encryption and/or change in packet format. The service level parameters for new services may be described by meta-data that is opaque to the eavesdropper.
- Avoiding the need for long lived addresses in deployments in order to prevent trackability.

- Regulation regarding PII definition, ability, and permission of owner of data for further dissemination and closing of existing regulatory loopholes.
- Uses of homomorphic forms of encryption for packet headers and tunnelled packets, in addition to traditional payload encryption, that allow to perform network operations in privacy-preserving manner without exposing metadata carried in headers.

10.8 Trustworthiness

As future end-to-end communication are deployed with new network services, packets will traverse several trust boundaries which are under different administrative domains. New Network 2030 services rely on the trustworthiness of the different nodes in order to protect the integrity of the data, handling of the packets and concurrently guarding the privacy of the users.

Additional mechanism will be required to verify that the nodes traversed are indeed trustworthy in handling the packets. There is also the need to ensure that the packets are themselves trustworthy. One possible approach involves the introduction of trustworthiness scores and rules on how to act when the different actors are outside the boundaries of what is acceptable.

The trustworthiness scores of nodes and packets should rely on security mechanism to verify, authorize, and ensure packet integrity. The study of this and related approaches is a subject for future documents.

10.9 Accounting, accountability, validation of delivered services

Many Network 2030 services place very high demands on the network in terms of required service levels, demanding guarantees instead of being accepting of "best effort".

Guarantees demand their price, making it increasingly important to be able to validate that promised service levels were delivered. This will require advances in accounting technology. For example:

- Measurements of service levels will need to be accurate enough to account for service performance targets.
- Proof of service delivery (including proof of service level delivery) may need to be provided to account and charge for network services. This is particularly the case as network services move from a best effort basis to a guaranteed basis and are used for mission-critical applications. Guarantees should be expected to have their price, and best effort accounting may no longer be sufficient for 2030 Networks.
- Advances in accounting protocols, in order to enable new incentive-based schemes to deliver services. For example, using prepay models, applications would no longer be able to just demand a network service with a certain network service level and rely on the goodwill of the network provider to provide it, but give network providers an incentive to deliver on them. Conversely, network providers will be able to allocate their resources more effectively than today in ways that best support economic goals. This can enable new business models and communication service supply chains that in turn foster further innovation for Network 2030.

In contrast, 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 rely largely on active measurement techniques that also have limitations related to sampling. Furthermore, it comes with significant overhead, including the consumption of network bandwidth as well as additional processing on edge nodes. Techniques that rely on passive measurements are unfeasible in many

network deployments and hampered by encryption, as well as issues relating to privacy, the concerns for which are expected to increase further.

New challenges arise from novel network services such as qualitative communication, as accounting may not be sufficient at the level of packets and flows, but a new level of chunks is introduced. Likewise, coordinated communication will require the development of techniques that account not just for individual packets and flows being delivered, but also for their coordination. Combined, this results in interesting challenges for accounting to be addressed in Network 2030.

Bibliography

- [b-IETF RFC 1633] IETF RFC 1633 (1994), *Integrated Services in the Internet Architecture: An Overview*.
- [b-IETF RFC 2211] IETF RFC 2211 (1997), *Specification of the Controlled-Load Network Element Service*.
- [b-IETF RFC 2212] IETF RFC 2212 (1997), *Specification of Guaranteed Quality of Service*.
- [b-IETF RFC 8321] IETF RFC 8321 (2018), *Alternate-Marking Method for Passive and Hybrid Performance Monitoring*.
- [b-Aijaz] Aijaz A, M. Dohler, AH Aghvami, V. Friderikos, M. Frodigh: *Realizing the tactile internet: Haptic communications over next generation 5G cellular networks*, IEEE Wireless Communications. 2017 Apr; 24(2): 82-9.
- [b-Bennett] Bennett J. C. R., K. Benson, A. Charny, W.F. Courtney, J. Y. L. Boudec, *Delay Jitter Bounds and Packet Scale Rate Guarantee for Expedited Forwarding*, IEEE/ACM Transactions on Networking, vol. 10, no. 4, pp. 529–540, Aug 2002.
- [b-Boudec] Boudec J. Y. L. and A. Charny, *Packet Scale Rate Guarantee for Non-FIFO Nodes*, IEEE/ACM Transactions on Networking, vol. 11, no. 5, pp. 810–820, Oct. 2003.
- [b-Brettel] Brettel M, N. Friederichsen, M. Keller, M. Rosenberg: *How virtualization, decentralization and network building change the manufacturing landscape: An Industry 4.0 Perspective*, International Journal of Mechanical, Industrial Science and Engineering. 2014 Jan;8(1): 37–44.
- [b-Brockners] Brockners F, S. Bhandari, C. Pignataro, H. Gredler, J. Leddy, S. Youell, T. Mizrahi, D. Mozes, P. Lapukhov, R. Chang, D. Bernier, J. Lemon: *Data Fields for In-situ OAM*, IETF Internet Draft draft-ietf-ippm-ioam-data, July 2019.
- [b-blueoceans] *Industry 4.0*, <https://www.blueoceans.com/home/events/news/industry-4-0/>
- [b-Clemm] Clemm A, L. Ciavaglia, L. Granville, J. Tantsura: *Intent-Based Networking – Concepts and Overview*, Internet Draft, IETF, July 2019.
- [b-DetNet] IETF: *Deterministic Networking*, <https://datatracker.ietf.org/wg/detnet/about/>
- [b-Fettweis] Fettweis GP: *The Tactile Internet: Applications and challenges*, IEEE Vehicular Technology Magazine. 2014 Mar;9(1):64–70.
- [b-Finn] Finn N, P. Thubert, B. Varga and J. Farkas: *Deterministic Networking Architecture (DetNet)*, Internet Draft draft-ietf-detnet-architecture-13, IETF, May 2019.
- [b-Forbes] *What is Industry 4.0? Here's A Super Easy Explanation For Anyone*, <https://www.forbes.com/sites/bernardmarr/2018/09/02/what-is-industry-4-0-heres-a-super-easy-explanation-for-anyone/#495df4d29788>
- [b-IMT2020] ITU: *IMT for 2020 and beyond*, <https://www.itu.int/en/ITU-R/study-groups/rsg5/rwp5d/imt-2020/Pages/default.aspx>.
- [b-ITU-T TD375] *Representative use cases and key network requirements for Network 2030*.

- [b-Miao] Miao Y, Y. Jiang, L. Peng, MS Hossain, G. Muhammad: *Telesurgery Robot Based on 5G Tactile Internet*, Mobile Networks and Applications. 2018:1-0.
- [b-Schwarz] Z. Li, J. Llach, K. Mammou, R. Mekuria, O. Nakagami, E. Siahhan, A. Tabatabai, A. Tourapis, V. Zakharchenko: *Emerging MPEG Standards for Point Cloud Compression*, IEEE Journal on Emerging and Selectec Topics in Circuits and Systems Vol 9 No 1, March 2019.
- [b-SIEMENS-In] Frequently asked questions on the topic "*Time-Sensitive Networking*", <https://support.industry.siemens.com/cs/ww/en/view/109757263>
- [b-Tactile] ITU-T Technology Watch Report (2014), *Tactile Internet*, https://www.itu.int/dms_pub/itu-t/opb/gen/T-GEN-TWATCH-2014-1-PDF-E.pdf
- [b-Torres] Torres Vega M, T. Mehmlı, J. van der Hooft, T. Wauters, F. De Turck, *Enabling Virtual Reality for the Tactile Internet: Hurdles and Opportunities*, 1st International Workshop on High-Precision Networks Operations and Control (HiPNet), Rome, Italy, IEEE, Nov 2018.
- [b-Torres-3D] Xu X, Y. Pan, PP Lwin, X. Liang: *3D holographic display and its data transmission requirement*, 2011 IEEE International Conference on Information Photonics and Optical Communications (IPOC), Oct 2011.
- [b-TSN] IEEE 802.1, *Time-Sensitive Networking (TSN) Task Group*, <https://1.ieee802.org/tsn/>

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	Tariff and accounting principles and international telecommunication/ICT economic and policy issues
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Environment and ICTs, climate change, e-waste, energy efficiency; construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling, and associated measurements and tests
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects, next-generation networks, Internet of Things and smart cities
Series Z	Languages and general software aspects for telecommunication systems