# Draft ETSI GS NGP 013 V0.0.6 (2018-07)


## Next Generation Protocol, (NGP);
## Flexilink: efficient deterministic packet forwarding in user plane for NGP;
## Packet formats and forwarding mechanisms

**Disclaimer: This DRAFT is a working document of ETSI ISG NGP. It is provided for information only and is still under development within ETSI ISG NGP. ETSI and its Members accept no liability for any further use/implementation of this Specification.**

**Non-published NGP drafts stored in the "Open Area" are working documents, these may be updated, replaced, or removed at any time**

**Do not use as reference material.**

***Disclaimer***

**Do not cite this document other than as "work in progress".**

The present document has been produced and approved by the Next Generation Protocols (NGP) ETSI Industry Specification Group (ISG) and represents the views of those members who participated in this ISG.
It does not necessarily represent the views of the entire ETSI membership.

**Approved and published Specifications and reports for implementation of the MEC system shall be obtained via the ETSI Standards search page at:**
**http://www.etsi.org/standards-search**

Reference
DGS/NGP-0013

Keywords
3GPP access; core network; fixed network, flexilink; IP; MPLS; next generation protocol; non 3GPP access; NR; QoS; switching

***ETSI***

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

# Contents

# Intellectual Property Rights

### Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

### Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

# Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification Group (ISG) Next Generation Protocols (NGP).

# Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

# Introduction

ISG NGP is tasked with finding a system of packet routing that does not suffer from the problems operators have experienced with LTE and also optimizes the performance, efficiency, and scalability of new services proposed for 5G.

Internet Protocol (IP) is not fully able to meet these requirements for a number of reasons which have been widely documented. Many of the constraints that applied when IP was developed, around 1980, (such as memory size) are no longer an issue, while other considerations that are important now (such as mobility and latency) were not an issue then. Furthermore, back then all processing had to be done by code running on a CPU, whereas now many tasks can be done more efficiently by dedicated logic in a System-on-a-Chip (SoC).

Most packets are part of a "flow" such as a TCP session or a video stream. Increasingly, there is a separation between the processes of deciding the route packets will follow and of forwarding the packets, for example in Software Defined Networking (SDN) and Control and User Plane Separation (CUPS). This is taken to its logical conclusion by specifying a system in which routing decisions are taken per-flow by control plane code, and per-packet processing is made simple enough that it can be implemented entirely in logic.

The system provides two separate services, a "basic" service suitable for traditional statistically multiplexed packet data, and a "guaranteed" service providing the lowest possible latency for continuous media, not only audio and video but also signals used in industrial automation and newer services proposed for 5G such as tactile feedback. Implementing these two services and multiplexing them together on communication links is less complex than attempting to serve both kinds of traffic with a single service.

The control plane procedures for managing flows are not specified in the present document. The system has been prototyped using the signalling messages specified in ISO/IEC 62379-5-2 [i.1]; messages containing similar information in other formats can also be used.

# 1 Scope

The present document specifies user plane packet formats and routing mechanisms for 5G core and access networks, based on the requirements documented in ETSI GS NGP 012 [i.3] and taking technologies from ETSI GR NGP 003 [i.2] as appropriate.

It does not specify the control plane procedures for managing routes.

# 2 References

## 2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at https://docbox.etsi.org/Reference.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced document is necessary for the application of the present document.

[1] IEEE 802.3™-2015: "Local and metropolitan area networks - Specific requirements Part 3: Carrier sense multiple access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications".

## 2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] ISO/IEC 62379-5-2:2014: "Common control interface for networked digital audio and video products - Transmission over networks - Signalling".

[i.2] ETSI GR NGP 003: "NGP Next Generation Protocol; Packet Routing Technologies".

[i.3] ETSI GS NGP 012: "KPIs for Next Generation Protocols Basis for measuring benefits of NGP".

# 3 Definitions and abbreviations

## 3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

**basic flow:** flow with no guarantees as to latency or whether packets will be dropped

**basic service:** packet transmission service carrying basic flows

**flow:** specification for how packets are forwarded from source to destination(s), and interpreted at the destination

**guaranteed flow:** flow offering a specified latency and specified probability of packet loss for packets which are transmitted at a specified rate

**guaranteed service:** packet transmission service carrying guaranteed flows

## 3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

| | |
|---|---|
| API | Application Program Interface |
| FCS | Frame Check Sequence |
| ls | least significant |
| MAC | Media Access Control |
| ms | most significant |
| PHY | PHYsical layer, or hardware that implements an interface to it |
| SFD | Start-of-Frame Delimiter |
| UDP | User Datagram Protocol |

# 4 Packets and flows

A packet shall consist of a payload and encapsulation. The payload shall be an octet string which is conveyed unaltered (apart from the effect of transmission errors and equipment faults) from source to destination. The encapsulation shall depend on the medium over which the packet is conveyed.

NOTE 1: Thus in many cases the encapsulation will change as the packet progresses through the system.

Each packet shall be associated with a flow.

NOTE 2: Management of flows is handled by the control plane, and is thus out of scope for the present document.

The action to be taken when a packet is received at a network element shall be defined by the flow with which it is associated; it shall not (except at a destination of the flow and as provided in clause 5.3.2.3) depend on the content of the payload.

NOTE 3: The route taken by packets associated with the flow is established, and any required resources reserved, before the first packet is sent. This allows security checks to be made before any data is transmitted.

# 5 Basic service

## 5.1 Service characteristics

The service specified in clause 5 is intended for the kind of communication between software processes that commonly occurs in computer systems. Transmission will typically be intermittent, so that reservation of capacity would be inappropriate and thus latency is undefined and packet loss due to queue overflow can occur.

Each flow shall have exactly one destination. It may have more than one source.

NOTE 1: A flow with more than one source would be similar to an MPLS Forwarding Equivalence Class.

NOTE 2: Multicasting not supported by basic flows, to simplify implementation of the forwarding process. Typically, basic flows will be used in pairs, one in each direction, so that messages can be acknowledged and dropped packets repeated. Multicasting is supported by the guaranteed service (see clause 6).

Switches may support facilities similar to those specified for IP networks to prioritize particular classes of traffic. Such facilities would be invoked by control plane messages (with the traffic class being a property of the flow and not signalled in packet headers) and are thus out of scope for the present document. Consideration should be given to whether they are necessary, or whether all traffic for which they might be useful can use the guaranteed service.

NOTE 3: These facilities would add complexity, but could be useful for sharing out capacity in the context of network slicing, and possibly also for compatibility with legacy systems.

## 5.2 Encapsulation

The details of packet formats, including the encapsulation, shall be defined separately for each type of link (see clause 8).

The payload should be preceded by a label which serves as an identification of the flow which is local to the link.

The encapsulation shall indicate the number of payload octets, either directly or by marking the start and end of the packet or of the payload.

## 5.3 Payload

### 5.3.1 Size

The payload of a packet transmitted by an end system shall be any number of octets in the range 1 to 2 000 inclusive.

NOTE 1: The lower limit is much less than for Ethernet. Zero-length payloads are not supported because support for them has been found to increase complexity in the forwarding logic.

NOTE 2: The upper limit supports tunnelling of maximum-size Ethernet envelope frames; per-packet overheads are minimal, so there is no incentive to support larger packets. A fixed upper limit eliminates the need for path MTU discovery.

Switching equipment shall support payload sizes up to 2 016 octets.

NOTE 3: That allows for up to 16 octets of inner labels.

### 5.3.2 Format

#### 5.3.2.1 General

The payload of a basic service packet shall consist of a user message and an integrity check.

The form taken by the integrity check shall be signalled in the control plane messages that set the flow up. The user message shall consist of all payload octets that are not part of the integrity check.

#### 5.3.2.2 User message

The syntax and semantics of the user message shall be negotiated by the source and destination entities as part of the process of setting the flow up.

The network shall treat the user message as an unstructured octet string, and shall not take any action that is based on the value of any octet in the string.

#### 5.3.2.3 Integrity check

Integrity check formats supported by the control plane messages shall include:

- No integrity check; all payload octets are part of the user message.
- Last two payload octets are coded with a nonzero value such that the sum of the set of 16-bit numbers derived as follows is a multiple of 65535. If the payload is an odd number of octets, a zero octet is inserted before the penultimate octet (i.e. at the end of the user message). Then, the payload is divided into pairs of octets, each holding a 16-bit number, with the more significant half of each number being in the first octet of the pair.
- Last four payload octets hold a CRC calculated as specified in clause 3.2.9 of IEEE 802.3-2008 [1].

- Last four payload octets hold longitudinal parity, calculated as specified in clause 8.2.2.

NOTE 1: The integrity check covers the user message but not the header, which has its own integrity check, and is changed at each switching point.

The integrity shall be checked by the destination entity. It may be checked at other points in the network, and the packet discarded if the check fails.

NOTE 2: In many cases corruption of a packet will be sufficiently unlikely that checking the integrity at intermediate nodes is not worth the effort.

#### 5.3.2.4 Aggregated flows

A flow may be created that gathers together a bundle of other flows, by prepending a label that identifies the flow within the bundle (the "inner" label") to the payload of each packet.

NOTE 1: This is similar to an MPLS "push" operation, but unlike with MPLS the existence of the inner label is a property of the outer (aggregated) flow and not indicated in the packet format; the format of the inner label only needs to be understood by the network elements at the ends of the aggregated flow, not by any intermediate switches through which it passes.

Integrity check fields within the payload should not be adjusted to include the inner label. The format of the inner label shall include its own integrity check.

NOTE 2: It will usually be appropriate for an aggregated flow to be signalled as "no integrity check", so that switches through which it passes do not perform any checks on the payloads.

# 6 Guaranteed service

## 6.1 Service characteristics

The service specified in clause 6 is intended for continuous media such as digital audio and video, tactile feedback, and position of vehicles and industrial robots. However, it can also be used for other traffic.

NOTE 1: If the guaranteed service is used for file transfers, the data rate can be negotiated via control plane messages and packets will not be dropped due to buffer overflow, so the transport protocol can be greatly simplified.

Each flow shall have exactly one source. It may have more than one destination.

Each flow shall support a defined number of packets per second, negotiated by the control plane protocols when it is set up, and should have resource allocated such that latency is bounded and packets will not be lost due to congestion.

The service on each link shall be either synchronous as specified in clause 6.3 or asynchronous as specified in clause 6.4.

NOTE 2: A synchronous service can achieve lower latency than an asynchronous service.

## 6.2 Packet format

Each guaranteed service packet shall consist of a header and 0 to 63 payload octets.

The header shall be a single octet formatted as:

ms bit: odd parity

1 bit: flag *f* (see below)

ls 6 bits: *n*, the number of payload octets

The flag *f* shall not be used for forwarding, but shall be available for use by end systems; if it is used to guide reassembly of longer messages, it should be set to 0 in the last fragment of a message and 1 in others.

NOTE 1: Unlike basic service packets, the header does not change when the packet is forwarded and does not need to be read in order to identify the flow; this is potentially useful for forwarding in the optical domain.

NOTE 2: See clause 5.2.2 of ETSI GR NGP 003 [i.2] for the rationale behind the payload size.

## 6.3 Synchronous service

### 6.3.1 Slots

Each link between network elements that is capable of carrying a synchronous service shall be formatted into "slots", and the slots grouped into "allocation periods". Each slot shall be 64 octets in length. The link may also carry symbols (e.g. framing) that are not part of any slot.

Each guaranteed flow shall be allocated one or more slots per allocation period. The framing on the link shall show where each allocation period starts, and a flow's allocation shall be of the same set of slots in each allocation period.

An allocation period shall be nominally 0,49984 ms multiplied by *m*, which shall be a power of two in the range 1 to 64 inclusive.

NOTE 1: If $m = 2$, as in the system described in clause 5 of ETSI GR NGP 003 [i.2], the minimum allocation for a flow (one slot per period) is 1 000 packets per second plus a tolerance for the source of the data having a clock that is up to 320 ppm faster than the reference used by the source of the frame timing.

NOTE 2: As remarked in ETSI GR NGP 003 [i.2], the size of the routing table is proportional to the length of the allocation period (i.e. to *m*) and also to the bit rate on the link. In the case of the prototype implementation $m = 2$, the bit rate is 1 Gb/s, and the table has just under 2 K entries. On a 10 Gb/s link it would have 20 K entries with $m = 2$, 10 K with $m = 1$.

A slot that does not contain a packet shall be coded with $f = 1$ and $n = 0$.

NOTE 3: If *f* is coded as specified in clause 6.2, adding empty slots (which can occur when changing the transmission rate of an existing flow, or if a flow is routed between links with different values of *m*) will have no effect.

The content of the last $(63 - n)$ octets of the slot is undefined as far as the synchronous service is concerned.

NOTE 4: However, it can be defined for other purposes; see clause 8.2.1.

The flow with which a guaranteed service packet is associated shall be identified by the location in the allocation period of the slot which it occupies.

### 6.3.2 Alignment of allocation periods

Network elements to which links carrying a synchronous service are connected shall maintain a fixed phase relationship between incoming and outgoing allocations, to within a defined tolerance.

When a link which can carry a synchronous service first comes up, negotiation via control plane messages shall establish whether the two sides have a common reference for frame alignment; if not, a further exchange of control plane messages shall cause the subnetwork on one side of the link to take its frame timing from the other. The control plane protocols used, and the process whereby the source of the frame timing is chosen, are out of scope of the present document.

Guaranteed flows shall not be connected across a link using a synchronous service until the phase relationship is stable.

NOTE: See also clause 5.3.4.2 of ETSI GR NGP 003 [i.2].

## 6.4 Asynchronous service

### 6.4.1 General

Where it is not possible to route a guaranteed flow over links that support the synchronous service, other services may be used. In that case, guaranteed service packets shall be encapsulated in basic service packets as specified in clause 7.1, and if necessary the basic service packets shall be further encapsulated as specified in clause 8.3.

QoS guarantees should, if possible, be negotiated with each network over which the flow is routed.

### 6.4.2 Alignment of allocation periods

Network elements between which links carrying an asynchronous service are connected shall maintain a fixed phase relationship between incoming and outgoing allocations, to within a defined tolerance. The control plane protocols used, and the process whereby the source of the frame timing is chosen, are out of scope of the present document.

NOTE 1: In most cases, the tolerance will be significantly greater than for the synchronous service (see clause 6.3.2). However, it can be reduced by relating allocation periods to time as distributed by protocols such as PTP.

NOTE 2: See also clause 5.7.2 of ETSI GR NGP 003 [i.2].

# 7 Carriage of one service in another

## 7.1 Guaranteed service packets as basic service payloads

Guaranteed service packets on links or subnetworks that do not support a synchronous service shall be carried in basic service packets on a basic flow that is signalled (in the control plane) as carrying guaranteed service packets.

NOTE 1: If the synchronous service is not supported, the flow to which a guaranteed service packet belongs cannot be established by its position in an allocation period.

The payload of each basic service packet on the flow shall consist of one or more guaranteed service packets.

The information relating to the flow shall specify the guaranteed flow for each guaranteed service packet, the guaranteed service packets being identified by their position in the basic service packet. It shall also specify how to establish how many octets each guaranteed service packet occupies, based on its payload length. The coding of this information, and the control plane protocols used to convey it, are out of scope of the present document.

NOTE 2: Options can include always carrying all 64 octets of the slot, or only carrying those octets specified in clause 8.2.1 as being foreground octets.

## 7.2 Basic service packets as guaranteed service payloads

This clause specifies a method for carrying basic service packets in a guaranteed flow, for example to provide defined QoS to the basic flow.

The sender shall form a data unit by prepending the length to the basic service packet's payload. The length should be coded in two octets as specified in clause 8.2.2. The data unit shall be divided into segments, each of not more than 63 octets, and successive segments shall be transmitted in guaranteed service packets. The last segment shall be coded with $f = 0$, all others with $f = 1$. Packets coded with $f = 1$ and $n = 0$ may be sent at any time.

The recipient shall accumulate incoming data until a guaranteed service packet with $f = 0$ is received. It shall then check whether the length signalled at the beginning of the data unit corresponds to the number of data octets received, and discard the data unit if not.

NOTE 1: The flow to which the data unit belongs is identified in the same way as for other guaranteed flows; thus, no label is included. If the flow is an aggregate of basic flows, the inner labels will be present but will be seen as part of the payload.

NOTE 2: The checking of the length is similar to ATM AAL5. If the basic service packet's payload includes an integrity check, that can also be checked.

## 7.3 Legacy packets as basic service payloads

The format carried by a basic flow is signalled in the messages that set it up; these messages are out of scope for the present document, but should include code points for complete Ethernet packets, complete IP datagrams, and packets or datagrams with certain fields stripped at ingress and restored at egress.

NOTE: This provides a mechanism to tunnel legacy flows across a Flexilink network.

In the case of communication across a Flexilink network between an application running in a mobile device and a gateway to an IP network, as much as possible of the API should run remotely in the gateway, to minimize the size of headers sent over the air interface.

EXAMPLE: When using UDP via the Berkeley Sockets interface, a `connect()` call can be implemented as setting up a flow connected to the packet gateway, and `send()` can then be implemented by transmitting the data on that flow, with the packet gateway adding the IP etc headers for onward transmission on the IP network. Similarly, for `recv()` the gateway can process the IP etc headers and forward the data to the application. There can be a separate flow for each socket, or a single flow similar to an aggregated flow (see clause 5.3.2.4) with the socket handle acting as an inner label. These actions are not expected to be any more onerous than the Network Address Translation used in current systems. The flow is cleared down when `close()` is called.

## 7.4 Legacy packets as guaranteed service payloads

This can be implemented as a combination of clauses 7.3 and 7.2.

# 8 Encapsulation formats

## 8.1 General

Any link that provides a continuous (or nearly continuous) full-duplex byte stream, such as a point-to-point copper or fibre link, should provide a synchronous service, transmitting a continuous sequence of frames as specified in clause 8.2 in each direction.

Other links should provide a service appropriate to the underlying technology, as specified in clause 8.3.

Control plane protocols should report the level of service that is provided to each flow.

## 8.2 Formats for continuous transmission

### 8.2.1 Common elements

A continuous sequence of frames shall be transmitted. Each frame shall include timing information as specified below and slots as specified in clause 6.3.1.

Timing information shall consist of the time by the sender's clock at which a particular event connected with transmission of the frame occurred, coded as an integer number of seconds and an integer number of nanoseconds, the latter being in the range 0 to 999 999 999 inclusive. If timing information is not available, the field should be coded with binary 1 in every bit. Other code points where the number of nanoseconds is more than 999 999 999 are reserved. The reference to which the timing values relate should be signalled as part of the process of setting the link up.

Within the frame, each octet shall be classified as framing or foreground or background. The background octets in successive frames shall form an intermittent octet stream which shall be used to carry basic service packets.

Within each slot, the first *k* octets shall be foreground and the remainder background, where *k* is *n* + 1 rounded up to a multiple of *w*, *n* is as signalled in the header (see clause 6.2), and *w* is a characteristic of the link which is signalled as part of the process of setting the link up.

NOTE 1: Typically, *w* will be the width of the data path in the MAC logic, in octets.

NOTE 2: Specification of the process of setting the link up is out of scope for the present document.

## 8.2.2 1 Gb/s Ethernet PHY

### 8.2.2.1 Frame format

Each frame shall consist of:

- 2 octets: preamble, each coded as hexadecimal 55
- 1 octet: start of frame delimiter (SFD), coded as hexadecimal D5
- 1 octet: frame type
- 4 octets: timing information
- 7 744 octets: 121 slots (see clause 8.2.1; *w* = 1)
- 40 octets: "trailing" background octets
- 4 octets: longitudinal parity
- 12 to 16 octets: inter-frame gap

All octets other than the slots and the trailing octets are framing octets.

Every 512$^{th}$ frame shall have its frame type coded as binary 0101 0000. All other frames shall have the frame type coded with 0100 in the ms four bits and the ls four bits 1 more (modulo 16) than in the previous frame.

The timing information shall be coded with the time at which the frame type octet was passed to the PHY, with the number of seconds (modulo 4) in the ms 2 bits and the number of nanoseconds in the remaining 30 bits.

Longitudinal parity shall be used only to monitor the quality of the link. Checking of the parity shall not delay or otherwise affect processing of foreground or background octets. Each octet of longitudinal parity shall contain the bit-reversed ones complement of the result of an XOR operation applied to every fourth preceding octet, not including the preamble or SFD.

NOTE: Longitudinal parity can be calculated with the same circuit that is used for the Ethernet FCS, by suppressing the feedback terms.

Where frame timing is taken from an external reference, there shall be a signal which shows when the next frame is to start, and gap octets shall be transmitted until that signal appears, subject to the restriction that there shall be at least 12 gap octets and not more than 16. The average number of gap octets shall be monitored, and the same average maintained (for example, using a binary rate multiplier) if the reference is lost.

### 8.2.2.2 Background octet stream

The background octet stream shall carry basic service packets and idle octets.

An idle octet shall be a single octet coded as hexadecimal FF, and shall be transmitted whenever there is no basic service packet available for transmission.

A basic service packet shall consist of:

- 2 octets: length, coded with the value *l* - 1

- 2 octets: label
- *l* octets: payload ($0 < l \leq 2016$)

The length and label fields shall each contain a 13-bit value followed by 3 bits which shall contain the ones complement of the remainder of the division (modulo 2) by the generator polynomial $x^3 + x + 1$ of the product $x^3$ multiplied by the 13-bit value.

NOTE 1: If the bits are numbered from d15 (ms) to d0 (ls) then the ls 3 bits can be calculated as: d2 = not d13^ d12^d11^d9^d6^d5^d4, d1 = not d15^d12^d11^d10^d8^d5^d4^d3, and d0 = not d14^d13^d12^d10^d6^ d5^d3.

NOTE 2: The length fits in eleven bits, so the ms two bits of the first octet will always be zero; for an idle octet these two bits will always be ones.

The label shall be coded with a value identifying the flow; this value shall be chosen by the recipient of the stream and signalled to the sender as part of the process of setting up the flow.

## 8.3 Formats for intermittent transmission

### 8.3.1 General

Before two network elements begin to exchange Flexilink packets over a medium which does not support the continuous transmission specified in clause 8.2, they shall set up an association between them which is referred to as a "virtual link". The process of setting up a virtual link is out of scope for the present document.

In cases where clause 8.3 does not specify a format for guaranteed service packets, guaranteed service packets may be encapsulated in basic service packets as specified in clause 7.1 provided allocation periods are aligned as specified in clause 6.4.2.

### 8.3.2 Basic service packets in legacy networks

A basic service packet may be carried directly over Ethernet MAC or over UDP.

NOTE 1: The Ethertype value in the former case, and the destination port number in the latter, are for further study.

The Ethernet MAC client data or UDP data shall consist of:

- 1 octet: hexadecimal 02
- 1 octet: hexadecimal 26
- 4 octets: timing information
- remainder: basic service packet

NOTE 2: Other code points in the first octet are reserved. Other code points in the second octet are used for management of the "virtual link" formed by the path across the Ethernet or IP network.

NOTE 3: The code point in the first octet is inherited from an earlier standard; it is retained because removing it would mean the basic service packet is aligned on an odd byte boundary.

In the Ethernet case, the packet shall have padding added if necessary, as specified in clause 3.2.8 of IEEE 802.3-2008 [1].

The format of the basic service packet shall be negotiated as part of the process of setting up the virtual link. It should be similar to the format specified in clause 8.2.2.2.

### 8.3.3 Wireless links

Transmission over wireless links is for further study.

# Annex A (informative): Assessment against KPIs

## A.1 KPIs for Naming and Addressing

Names and addresses do not appear in packets, only in control plane messages which are out of scope for the present document. A wide variety of addressing schemes, and other ways of identifying the remote entity, can be supported.

## A.2 KPIs for Performance

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| Per1 | Does the scheme require the address to be encoded in every packet of a flow? | Yes/no | No. |
| Per2 | Latency: the delay between the encapsulation of application data into a network protocol datagram by the sending endpoint; the forwarding of those datagrams to the destination endpoint; and the subsequent decapsulation of the datagram to extract the application data. | Time (ms) | Forwarding latency for the guaranteed service is fixed for each flow and is typically less than 10 µs per hop. Guaranteed service packets are maximum 63 octets, which places a limit on the time to fill a packet from data such as digital audio. Encapsulation and decapsulation of digital media can be done entirely by logic, eliminating software delays. |
| Per3 | Predictability/reliability: the ability of the protocols to deliver datagrams without loss or corruption; and to deliver datagrams in order as required. | Lost/corrupted packets as a % of the flow total | No loss or corruption on wired links; frame format includes check on link integrity, allowing unreliable links to be taken out of service. Performance on wireless links is for further study. |
| Per4 | Jitter: any variation in latency over time. Lower jitter would indicate a more predictable network protocol. | Standard deviation from expected latency | Jitter on a guaranteed flow is less than a slot time (512 ns for a 1 Gb/s link) if a single slot is allocated. If more than one slot, jitter is fixed; it depends on how evenly-spaced the slots are, and can be reported to the application when the flow is set up. |
| Per5 | Prioritization: the ability of the network protocol to support both prioritization and non-prioritization when processing flows from different sources. | Yes/no | Yes: guaranteed flows are delivered as scheduled, basic flows are statistically multiplexed. |

## A.3 KPIs for Mobility

Details of handover between cells are for further study.

A flow U - R1 - S - R2 - E is rerouted as U - R3 - S - R2 - E (where S is a switch and Rn are routes which may pass through further switches) by: set R3 up; rewrite routing table entry in U and/or S; clear R1 down.

## A.4 KPIs for buffering

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| Buf2 | Drop/queue support<br>The ability of the protocol to request that the network either drop or queue packets under resource contention. | Yes/No | The guaranteed service is a synchronous service so there is no contention at packet level and no queuing or packet drop.<br>The basic service can support this feature in the messages that set flows up (out of scope for the present document) if required. |
| Buf3 | Queue occupancy support when choosing optimal route. | Yes/No | Procedures for setting basic flows up (out of scope for the present document) are expected to support this. |
| Buf4 | Support for configurable scheduling - queuing for a configurable time. | Yes/no | As Buf2. |

## A.5 KPIs for Multihoming

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| MH1 | Do the protocols name the node, and not the network interface? This allows native multihoming and reduces complexity, improves scalability, load balancing and session continuity. | Yes/No | Procedures for setting flows up (out of scope for the present document) are expected to support this. |
| MH2 | Do the protocols support aggregation of content from different sources, to provide resilience? | Yes/No | Yes. |

## A.6 KPIs for Protocol and Energy Efficiency

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| PE1 | Protocol efficiency: The ratio of useful data in the payload to overhead has a direct financial impact on communication links. More performant protocols will deliver a higher value per second. NGP protocols should minimize header complexity and overhead. | Application bits as a ratio of total bits. | Up to 98,4 % for guaranteed service packets (minimum is 50 % for a single data octet; value is 80 % for 4 data octets)<br>Up to 99,8 % for basic service packets (minimum is 20 % for a single data octet; value is 50 % for 4 data octets)<br>Details of transmission over radio links are for further study. |
| PE2 | Processing overheads: instructions<br>The number of instructions required to process the protocol headers.<br>If software, how many machine instructions.<br>If logic, how many gates. | Number of processing steps (Integer) | No overhead for guaranteed service packets.<br>Single-cycle lookup in routing table for basic service packets.<br>Where clause 8.2 is used there is also parsing of the byte stream to demultiplex the two services; in the prototype implementation in Xilinx Spartan 6, MAC logic for one port, (receive and transmit, including above-mentioned multiplexing and demultiplexing, supporting Ethernet MAC as well as Flexilink) takes 613 registers and 729 LUTs, also two RAM blocks for the packet FIFOs for the basic service. |
| PE3 | Processing overhead: primary storage<br>The size of the information to be stored and processed. A higher information size will use up more memory bandwidth and buffer space. | Bytes | In the prototype implementation, the forwarding plane's tables (in fast sRAM, accessed per packet) are 18 bits per flow for the basic service and 9 bits per slot for the guaranteed service; the control plane tables (in dRAM, accessed only when required for management) occupy about 400 bytes per flow |

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| PE4 | Increase in space in routing tables<br>An efficient protocol will minimize increase in routing table size under multihoming, aggregation and traffic engineering. | Routing table entry insertions following a multihoming event or a mobility event (Integer). | Control plane routing information is out of scope for the present document.<br>Forwarding plane entry for the old route is released immediately after a mobility event. |
| PE5 | Connection establishment overhead<br>For connection-oriented protocols: How many round trips are required to establish a connection. Note, the latency of round trips should be considered the same when comparing two protocols/<br>For connectionless protocols: the instructions required to bind the flow to a sender/receiver. | Integer | One, unless certain optional features (such as authentication, or requiring the user to agree to a charge) are invoked, in which case one and a half or two. |
| PE6 | Retransmission of already-queued data<br>Endpoints should not retransmit information which is already queued upstream in the network path. | Yes/no | Not applicable; this is a feature of the next layer up. |
| PE7 | Flow Control loops<br>Reaction to loss or resource contention is most efficiently done at the point it occurs. | Number of network hops to report and react to congestion; number of decapsulations required to detect congestion signals (integer) | In the case of the guaranteed service, congestion does not occur.<br>If implemented for the basic service, this is a feature of the next layer up.<br>When a link is lost, flows can be rerouted locally. |
| PE8 | Overhead of security: the transmission and processing burden of encrypting, including the process of securing a flow, decrypting and integrity checking the application bits | Processing overhead, Bytes overhead per PE2 and PE3. | Not applicable; this is a feature of the next layer up. However, in some cases the ability to perform checks before setting up forwarding plane routing may mean encryption is not required. |
| PE9 | Is header re-encapsulation and modification required, such as checksum recalculation? | Yes/no | Guaranteed service packet headers are not modified. Labels on basic service packets are replaced at each hop, but they include their own data protection so no recalculation is required. |

## A.7 KPIs for Security and privacy

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| SEC1 | Security by default<br>Security achieved without overlays. | Yes/No | Yes, though details of the mechanisms are out of scope for the present document. |
| SEC2 | Crypto-agility for algorithms and key management independent of function invocation.<br>Whilst 'security by default' should identify a requirement for crypto-agility, this should be implemented in such a way that a change of the crypto solution should not impeded the functional capability of the NGP. | Yes/No | Yes, though details of the mechanisms are out of scope for the present document. |
| SEC3 | Reporting of security events to a recognized standard. | Yes/no | Attacks are easier to identify and resist than with connectionless technologies; details of reporting mechanisms are out of scope for the present document. Conformance to TC CYBER is for further study. |

# A.8 KPIs for traffic Management

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| NET1 | Latency of traffic identification. | Time (ms) | Typically < 10 ms.<br>Traffic identification is signalled during flow set-up, before any packets are transmitted.<br>Time required for flow set-up depends on the efficiency of the control plane software, the power of the processor it runs on, and the number of separate control plane entities involved. |
| NET2 | Volume of data to be inspected for traffic identification<br>Lowest volume of data in order to identify traffic in the early stage; this includes control plane bits if used. | Bits | One bit, in the control plane message setting a flow up, shows whether the flow is guaranteed or basic.<br>Further fields in the message carry additional identification; details (including coding efficiency) are out of scope for the present document.<br>All identification occurs before the first packet is transmitted. |
| NET3 | Real-time traffic identification of traffic<br>What is the latency incurred in identifying traffic classes? | Time (ms) | None.<br>All identification occurs before the first packet is transmitted. |
| NET4 | "Accuracy" in identifying the proper class of traffic<br>Based on tests that compare the 'perceived' traffic class from the actual traffic class. | Percentage | 100 % - all identification is explicit. |
| NET5 | QoS support and levels. | Integer | No limit other than link capacity and granularity of allocations: each flow has its own allocation. |
| NET6 | Scalability of management policies.<br>The intention is to reduce the complexity to manage policies. | Integer maximum number of network locations to apply traffic management | No management required for guaranteed flows other than allocation of slots when setting up a new flow.<br>Basic flows are not expected to require traffic engineering as any flows with QoS requirements are expected to use the guaranteed service. Any such facilities that are provided are out of scope for the present document. |
| NET7 | Capabilities of traffic management policies (i.e. expressivity). | Integer, Number of operations and number of parameters per operations. | See NET6. |

# A.9 KPIs for Interoperability

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| INT1 | Ability to support TCP/IP applications via interoperability. | Yes/no | Yes; see clause 7.3. |
| INT2 | Interworking with 3GPP R15/16 with minimal complexity. | Yes/no | Yes. |

# A.10 Deployment effort

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| COE1 | Effort required for integration with existing infrastructure. | 1 = None<br>2 = Minor<br>3 = Major<br>4 = Not possible | 2) Minor |
| COE2 | Re-use of existing infrastructure. | Percentage | Existing transmission links (fibres, etc.) can be used. For maximum benefit, new switches are required; however, transmission formats that can be supported by existing hardware are also possible though with reduced performance. Interworking with existing infrastructure is supported (see clauses 7.3 and 8.3.2). |

# A.11 Revenue opportunities

| ID | Definition and rationale | Metric | Value |
|---|---|---|---|
| BBE1 | Business market needs:<br>Type of benefits the NGP proponents expect to deliver to their possible Business customers compared to existing solutions. | Textual (List) | Support for lowest possible latency as standard.<br>More efficient use of spectrum.<br>More secure: easier to resist attack.<br>Easier to manage.<br>Lower power.<br>Simpler, cheaper equipment.<br>New business models, e.g. payment for connection to media streams. |
| BBE2 | Business impact. | Small/Medium/Large<br>How NGP is impacting Business customers compared to existing solutions. | Potentially large. |

# Annex B (informative): Authors & contributors

The following people have contributed to the present document:

**Rapporteur:**
John Grant, Nine Tiles

# History

| Document history | | |
|---|---|---|
| V0.0.3 | June 2018 | Clean-up done by ***editHelp!***<br>E-mail: mailto:edithelp@etsi.org |
| V0.0.6 | July 2018 | Clean-up done by ***editHelp!***<br>E-mail: mailto:edithelp@etsi.org |
| | | |
| | | |
| | | |