International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Q.3061

(02/2022)

SERIES Q: SWITCHING AND SIGNALLING, AND ASSOCIATED MEASUREMENTS AND TESTS

Signalling requirements and protocols for the NGN – Network signalling and control functional architecture

---

Signalling requirements for service function paths load balancing traceroute in service function chaining

Recommendation ITU-T Q.3061

ITU-T

ITU-T Q-SERIES RECOMMENDATIONS

SWITCHING AND SIGNALLING, AND ASSOCIATED MEASUREMENTS AND TESTS

SIGNALLING IN THE INTERNATIONAL MANUAL SERVICE	Q.1–Q.3
INTERNATIONAL AUTOMATIC AND SEMI-AUTOMATIC WORKING	Q.4–Q.59
FUNCTIONS AND INFORMATION FLOWS FOR SERVICES IN THE ISDN	Q.60–Q.99
CLAUSES APPLICABLE TO ITU-T STANDARD SYSTEMS	Q.100–Q.119
SPECIFICATIONS OF SIGNALLING SYSTEMS No. 4, 5, 6, R1 AND R2	Q.120–Q.499
DIGITAL EXCHANGES	Q.500–Q.599
INTERWORKING OF SIGNALLING SYSTEMS	Q.600–Q.699
SPECIFICATIONS OF SIGNALLING SYSTEM No. 7	Q.700–Q.799
Q3 INTERFACE	Q.800–Q.849
DIGITAL SUBSCRIBER SIGNALLING SYSTEM No. 1	Q.850–Q.999
PUBLIC LAND MOBILE NETWORK	Q.1000–Q.1099
INTERWORKING WITH SATELLITE MOBILE SYSTEMS	Q.1100–Q.1199
INTELLIGENT NETWORK	Q.1200–Q.1699
SIGNALLING REQUIREMENTS AND PROTOCOLS FOR IMT-2000	Q.1700–Q.1799
SPECIFICATIONS OF SIGNALLING RELATED TO BEARER INDEPENDENT CALL CONTROL (BICC)	Q.1900–Q.1999
BROADBAND ISDN	Q.2000–Q.2999
SIGNALLING REQUIREMENTS AND PROTOCOLS FOR THE NGN	Q.3000–Q.3709
General	Q.3000–Q.3029
Network signalling and control functional architecture	**Q.3030–Q.3099**
Network data organization within the NGN	Q.3100–Q.3129
Bearer control signalling	Q.3130–Q.3179
Signalling and control requirements and protocols to support attachment in NGN environments	Q.3200–Q.3249
Resource control protocols	Q.3300–Q.3369
Service and session control protocols	Q.3400–Q.3499
Service and session control protocols – supplementary services	Q.3600–Q.3616
Service and session control protocols – supplementary services based on SIP-IMS	Q.3617–Q.3639
VoLTE/ViLTE network signalling	Q.3640–Q.3655
NGN applications	Q.3700–Q.3709
SIGNALLING REQUIREMENTS AND PROTOCOLS FOR SDN	Q.3710–Q.3899
TESTING SPECIFICATIONS	Q.3900–Q.4099
PROTOCOLS AND SIGNALLING FOR PEER-TO-PEER COMMUNICATIONS	Q.4100–Q.4139
SIGNALLING REQUIREMENTS AND PROTOCOLS FOR IMT-2020	Q.5000–Q.5049
COMBATING COUNTERFEITING AND STOLEN ICT DEVICES	Q.5050–Q.5069

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

Recommendation ITU-T Q.3061

Signalling requirements for service function paths load balancing traceroute in service function chaining

Summary

Recommendation ITU-T Q.3061 specifies the signalling requirements for service function paths (SFPs) load balancing traceroute in service function chaining (SFC). The signalling is used for tracing and figuring out the set of load balanced SFPs more efficiently.

History

Edition	Recommendation	Approval	Study Group	Unique ID*
1.0	ITU-T Q.3061	2022-02-13	11	11.1002/1000/14850

Keywords

Load balancing, service function chaining, service function path, signalling requirements, traceroute.

* 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

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure, e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation 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 Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents/software copyrights, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the appropriate ITU-T databases available via the ITU-T website at http://www.itu.int/ITU-T/ipr/.

Table of Contents

Recommendation ITU-T Q.3061

Signalling requirements for service function paths load balancing traceroute in service function chaining

1 Scope

The scope of this Recommendation consists of:

– Overview of service function paths (SFPs) load balancing traceroute in service function chaining (SFC);

– SFC traceroute interface (SFCtri) reference model;

– Signalling procedure for SFCtri;

– Signalling requirements for SFCtri.

The appendix to this Recommendation also provides:

– An example of signalling procedure of SFPs load balancing traceroute.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[IETF RFC 7665] IETF RFC 7665 (2015), *Service Function Chaining (SFC) Architecture.*

[IETF RFC 8300] IETF RFC 8300 (2018), *Network Service Header (NSH).*

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 network service header (NSH) [IETF RFC 8300]: A type of encapsulation for service function controllers that are designed to encapsulate an original packet or frame and, in turn, be encapsulated by an outer transport encapsulation (which is used to deliver the NSH to NSH-aware network elements).

3.1.2 service function chain [b-ITU-T Y-Sup.41]: A chain that defines an ordered set of abstract service functions and ordering constraints that must be applied to packets and/or frames and/or flows selected as a result of classification and/or policy.

3.1.3 service function forwarder (SFF) [IETF RFC 7665]: A service function forwarder is responsible for forwarding traffic to one or more connected service functions according to information carried in the SFC encapsulation, as well as handling traffic coming back from the SF. Additionally, an SFF is responsible for delivering traffic to a classifier when needed and supported, transporting traffic to another SFF (in the same or different type of overlay), and terminating the service function path (SFP).

3.1.4 service function path (SFP) [IETF RFC 7665]: The service function path is a constrained specification of where packets assigned to a certain service function path must go. While it may be so constrained as to identify the exact locations, it can also be less specific. The SFP provides a level of indirection between the fully abstract notion of the service chain as a sequence of abstract service functions to be delivered, and the fully specified notion of exactly which SFF/SFs the packet will visit when it actually traverses the network. By allowing the control components to specify this level of indirection, the operator may control the degree of SFF/SF selection authority that is delegated to the network.

3.2 Terms defined in this Recommendation

None.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

NSH Network Service Header

OAM Operation Administration and Management

SF Service Function

SFC Service Function Chaining

SFCtri Service Function Chaining Traceroute Interface

SFF Service Function Forwarder

SFP Service Function Path

5 Conventions

In the body of this Recommendation, the words should and may sometimes appear, in which case they are to be interpreted, respectively as, is recommended and can optionally. The appearance of such phrases or keywords in an appendix or in material explicitly marked as informative is to be interpreted as having no normative intent.

{A}: indicates that the parameter A is mandatory.

6 Overview

Load balancing or multipathing is an attempt to balance traffic across a network by allowing the traffic to use multiple paths. It has several benefits including easing capacity planning, help absorb traffic surges by spreading them across multiple paths and allowing better resilience by offering alternate paths in the event of a link or node failures.

An SFC defines an ordered set of abstract service functions and ordered constraints that must be applied to packets/frames. There are more than one service function path's (SFPs) for one SFC. It creates the chance to balance the SFC traffic for bandwidth usage optimization. However, due to SFC maintenance, there is no operation administration and management (OAM) tool to traceroute the SFPs of one dedicated SFC.

The OAM tool "traceroute" is originally used in the IP layer. It can gather the IP address and latency of each node on the way and consequently trace the traffic path by sending and receiving echo request/response packets. In the IP layer, the traffic path is chosen by the transmission information of the packets. All the packets in one flow have the same transmission information, e.g., five tuples (source MAC address, destination MAC address, source IP address, destination IP address and port number). Hence one flow, has only one path to take despite its number of packets.

For SFC packet flow, since there is an appended NSH header upon the IP header, it creates the opportunity to balance the flow which cannot be realized in the IP layer by carrying a new type of transmission information in the NSH header other than the five tuples. Correspondingly, the SFC OAM tool could be developed to have the capability to trace the different SFPs of a balanced SFC.

In conclusion, it is essential to specify the signalling requirements of an SFP load balancing traceroute mechanism.

7 Interface SFCtri reference model

The reference model of interface SFCtri is shown below:

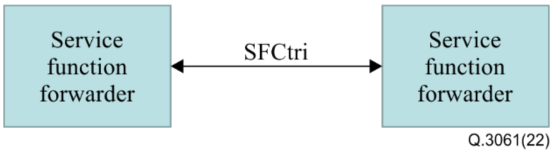

Figure 7-1 – The reference model of interface SFCtri

The interface SFCtri is between the SFFs. This interface is responsible for sending and receiving traceroute messages between SFFs.

The SFFs located in different locations on the path perform different actions. For example, the first SFF usually inserts the necessary information to the NSH header for SFC load balancing. The intermediate SFFs are relay forwarders and the last SFF responds its information to the first SFF.

8 Signalling procedure of interface SFCtri

When the SFC traffic flows through the network, to release the traffic pressure of a specific link, the traffic loads are distributed among different interfaces of a specific device and thus form different SFPs. The NSH header should carry an "entropy field" as an extension. Then the first SFF and the SFC classifier, performs a hash computation of the traffic transmission information to obtain the entropy value and inserts the entropy value to the NSH entropy field of the traffic packet.

When the following SFF receives the data packets:

(1) It reads out the value of the "entropy field" of the NSH header. The entropy value is determined by the transmission information of the packet flow.

(2) It takes the "modulus operation" action on the entropy value and figures out the remainder based on the modulus number equal to the number of equivalent paths and determines the packet forwarding path according to the entropy value.

(3) It uses the remainder to find out via which path (outbound interface) the packets should be transferred and forwards the packets to the determined path.

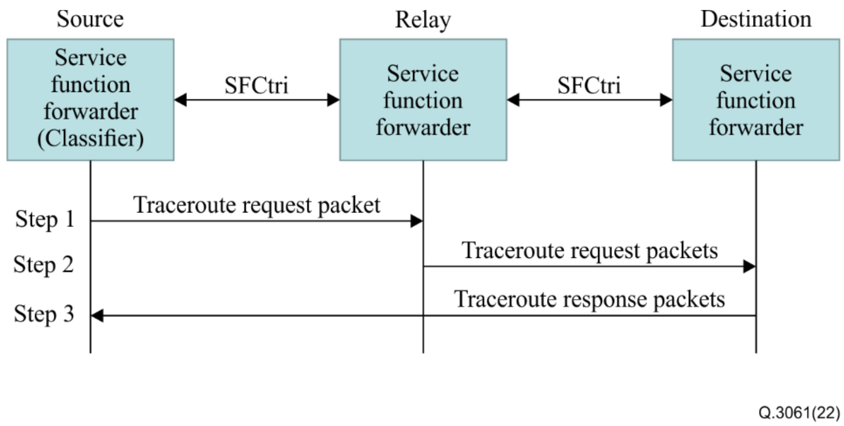

Figure 8-1 – The signalling procedure of interface SFCtri

The signalling procedure is illustrated in Figure 8-1 according to the simplest SFC which is composed of three SFFs, including the source SFF, the relay SFF and the destination SFF.

Step-1 The source SFF sends bunches of the traceroute request packets with different entropy values and source IP address to its following SFF (the relay SFF) through the interface SFCtri.

Step-2 When the following SFF receives the traceroute packet through interface SFCtri, it has the following actions:

1. Minus TTL value by 1.
2. If the TTL > 0, this SFF is not the destination SFF and it is responsible for relaying the packets to the destination SFF, it has the following actions:
   2.1 It reads out the value of the "entropy field" of the NSH header.
   2.2 It takes the "modulus operation" action on the entropy value and figures out the remainder.
   2.3 It uses the remainder results to find out via which downstream SFF the traceroute packets should be transferred.

Step-3 When the following SFF receives the traceroute packet from upstream SFF through interface SFCtri, it has the following actions:

1. Minus TTL value by 1.
2. Find out TTL = 0 and it sends the response messages to its source SFF by carrying the entropy value in the entropy field and its IP address.

9 Signalling requirements of interface SFCtri

9.1 Overview

The load balanced SFP tracing related signalling messages are exchanged over the interface SFCtri. The signalling messages may be extensible markup language (XML)-based messages over (or carried by) transmission control protocol (TCP), user datagram protocol (UDP), stream control transmission protocol (SCTP), transport layer security (TLS), etc. All of the messages consist of the message header and the message body.

The message format is described in Figure 9-1.

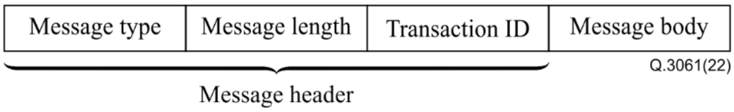

Figure 9-1 – Message composition

The message header field contains the following information:

– Message type: uniquely specifies the type of message.

– Message length: specifies the length of the message body.

– Message transaction ID: generated by the sender of the message. If there is a response message for the request message, the transaction IDs of the request and response messages are the same.

The message body field contains the message contents.

9.2 The SFCtri-RQ/RS message

The SFP traceroute request message is defined as the SFCtri-RQ message.

The SFCtri-RQ message, indicated by the message type in the message header field, is sent by the upstream SFF to the downstream SFF in order to collect the destination SFF's information.

Message format:

```
<SFCtri-RQ-Message> ::= < Message Header >
                 {Entropy-value}
                 {TTL-Value}
                 {Source-IP-address}
```

Meanings and explanations:

The detailed information indicates but is not limited to:

1. `Entropy-value` uniquely specifies the entropy value carried by the NSH packet.
2. `TTL-Value` uniquely specifies the time of live parameter of the NSH packet.
3. `Source-IP-address` uniquely specifies the IP address of source SFF.

NOTE – The `Source-IP-address` carried in the SFCtri-RQ-Message is different from the source IP address carried in the IP header which is changed from time to time.

The traceroute response message is defined as an SFCtri-RS message.

The SFCtri-RS message, indicated by the message type in the message header field, is sent from the downstream SFF to the upstream SFF.

Message format:

```
<SFCtri-RS-Message> ::= < Message Header >
                 {Entropy-value}
                 {Destination-IP-Address}
```

Meanings and explanations:

The detailed information indicates but is not limited to:

1. `Entropy-value` uniquely specifies the entropy value carried by the NSH packet.
2. `Destination-IP-Address` uniquely specifies the IP address of the destination SFF.

Appendix I

An example of signalling procedure of SFPs load balancing traceroute

(This appendix does not form an integral part of this Recommendation.)

This appendix provides an example based on the rules of signalling procedure described in clause 8.

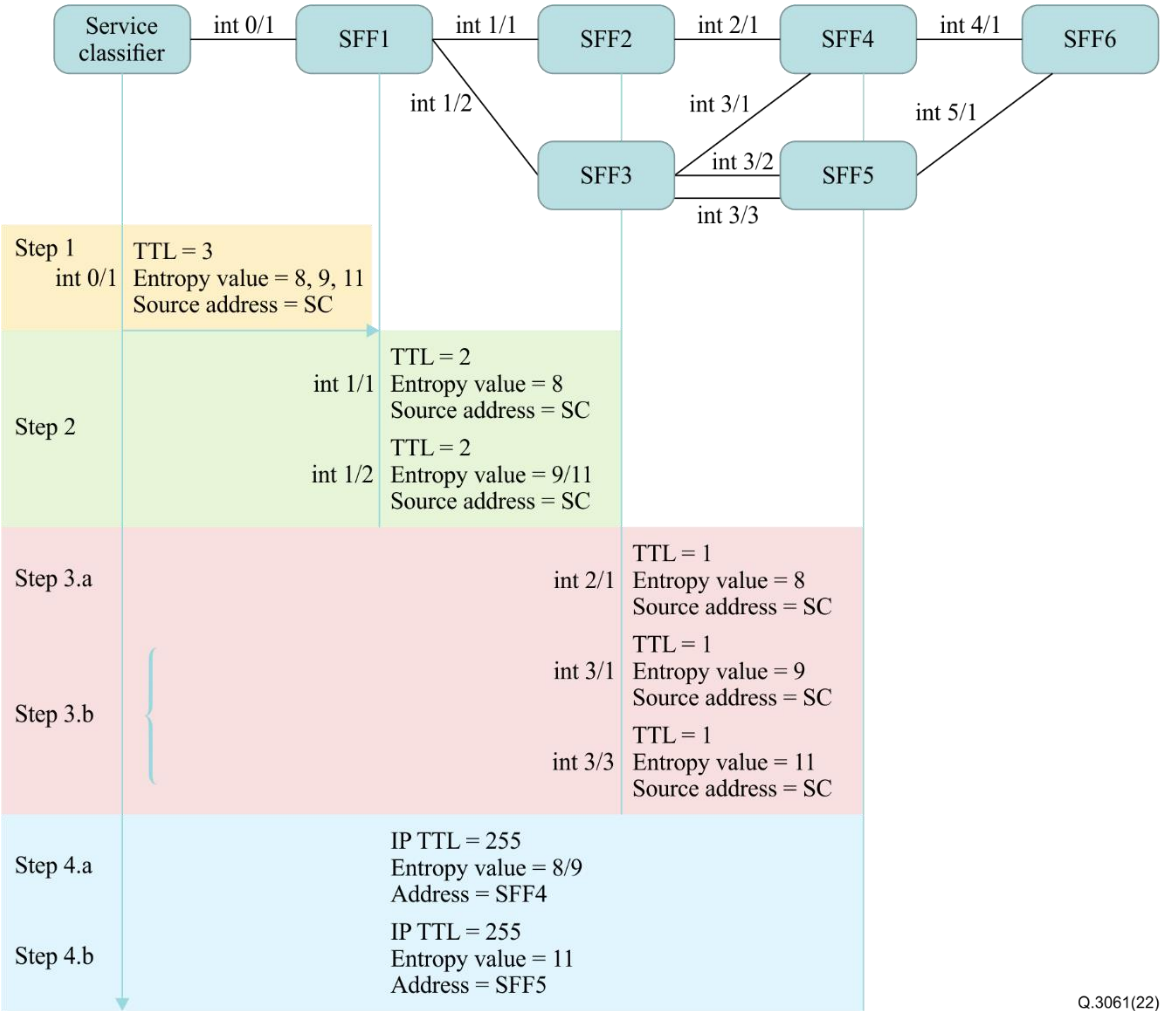

Figure I.1 – Example of signalling procedure of SFC load balancing traceroute

The load balanced SFP traceroute should be implemented as follows:

The source SFF sends bunches of traceroute packets with different TTL values in sequence until the SFC destination address responds back to the source SFF. For example, if each equivalent SFP of a dedicated SFC has 5 SFFs from the source to destination, the TTL value which is equal to 1, 2, 3, 4 will be sent in sequence until the address of the SFF6 is replied to the source.

For example, take TTL = 3 to test the load balancing path from SFF1 to SFF6:

Step-1 The service classifier sends 3 traceroute packets with 3 different entropy values (for example, entropy value = 8, 9, 11) to SFF1.

Step-2 When SFF1 receives the traceroute packet, it has the following actions in sequence:

(1) Minus TTL by 1, and TTL is equal to 2.

(2) TTL > 0, SFF1 reads out the value from the "entropy field" of the NSH header:

A. If entropy value = 8. Because there are 2 outbound interfaces int 1/1 and int 1/2, SFF1 takes the "modulus operation" action on the entropy value and figures out remainder = 0 (8 mod 2 = 0).

B. If entropy value = 9 and 11. SFF1 takes "modulus operation" action on the entropy value and figures out remainder = 1 (9 mod 2 = 1, 11 mod 2 = 1).

(3) SFF1 transfers the traceroute packet with entropy value = 8 via int 1/1 to SFF2 and the traceroute packets with entropy value = 9 and 11 via int 1/2 to SFF3.

Step-3.a When SFF2 receives the traceroute packet, it has the following actions in sequence:

(1) Minus TTL by 1, and TTL is equal to 1.

(2) TTL > 0, SFF2 has only 1 outbound interface int 2/1 for downstream SFF (SFF4), so it transfers the traceroute packet to SFF4.

Step-3.b When SFF3 receives the traceroute packets, it has the following actions in sequence:

(1) Minus TTL by 1, and TTL is equal to 1.

(2) TTL > 0, SFF3 reads out the values 9 and 11 from the "entropy field" of NSH headers.

(3) SFF3 has 3 outbound interfaces int 3/1, int 3/2 and int 3/3. So SFF3 takes "modulus operation" of mod 3 on entropy values of 9 and 11. The remainder of 9 mod 3 is 0 and the remainder of 11 mod 3 is 2.

(4) SFF3 transfers the packets with entropy value 9 (remainder = 0) via int 3/1 to SFF4 and the packets with entropy value 11 (remainder = 2) via int3/3 to SFF5.

Step-4.a When SFF4 receives the traceroute packet, it has the following actions in sequence:

(1) Minus TTL by 1 and find out TTL = 0.

(2) TTL = 0, SFF4 reads outs the "source IP address" from the SFCtr-RQ message header and encapsulates this "source IP address" as the destination IP address of the IP layer and set IP TTL = 255 to form the response SFCtr-RS message. These response packets which carry different entropy values 8, 9 and the IP address of SFF4 are sent to the source SFF.

Step-4.b When SFF5 receives the traceroute packet, it has the following actions in sequence:

(1) Minus TTL by 1 and find out TTL = 0.

(2) TTL > 0, SFF5 reads outs the "source IP address" from the SFCtr-RQ message header and encapsulates this "source IP address" as the destination IP address of the IP layer and set IP TTL = 255 to form the response message. These response packets which carry different entropy values 11 and the IP address of SFF5 are sent to the source SFF.

When the source SFF collects all the IP addresses of the SFFs along the path, it will construct the load balancing network topology using this information.

Bibliography

[b-ITU-T Y-Sup.41] ITU-T Y-series Recommendations – Supplement 41 (2016), *Deployment models of service function chaining*.

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

Printed in Switzerland
Geneva, 2022