Draft for discussion

WD 21558-x

Future Network — Architecture — Part x: Switching and routing

Source: UK expert

Contents

Page

Introduction

ISO/IEC TR 29181-1 describes the definition, general concept, problems and requirements for the Future Network (FN).

ISO/IEC TR 29181-3 examines the requirements for carrying data over digital networks, and identifies those that are not satisfied by the current Internet. It also notes some expected characteristics of new systems that are better able to satisfy the requirements, and specifies a model which supports both the existing system and the new systems. This will enable a migration to the new systems; it is also intended to make networks of all sizes easier to manage.

Relationship to other documents listed in clause 2 of protocols draft to be explained here, including a diagram.

Also relationship to the 7-layer model in ISO 7498.

This International Standard specifies the FN architecture, which is designed to meet the requirements identified in TR 29181-3. Protocols to support this architecture are specified in ISO/IEC 21559-x.

FN is a packet network which, as well as carrying data between computers, also meets the rather different requirements of live digital audio and video, which form an increasing proportion of the traffic on today's networks.

Whereas in IP all addressing information needs to be present in the packet headers, in FN the information needed to route packets is carried separately from the packets themselves. This reduces the size of the packet header by an order of magnitude, and simplifies the process of forwarding the packets in switches.

Most importantly, it allows different addressing mechanisms to be used without changing the packet format, and supports mobility without needing artificial devices such as IP-in-IP tunnels.

FN offers two main services: an ultra-low-latency “AV” service tailored to the needs of constant-bit-rate traffic such as audio and video, and a best-effort “IT” service suitable for the kind of unpredictable demand for which IP was intended. The AV service can also be used for file transfer, where it eliminates the need for the kind of empirical throughput testing that is a feature of TCP.

Some details of the services (particularly the slot size for the AV service, which was originally envisaged as being much more flexible) are a result of experimentation with a prototype implementation.

1 Scope

This International Standard specifies the switching and routing architecture of the Future Network.

2 Normative references

To be added when the main body of text is completed; must include the "architecture" document (ISO/IEC 21558-x), ISO/IEC 62365, ISO/IEC 62379-5-2, and AES51. Or should we make an updated form of AES51 into a new sub-part of IEC 62379-5? Or simply specify the relevant information explicitly in this document (see EdNote at end of ***Error! Reference source not found.****)?*

ISO/IEC TR 29181-1, *Information technology -- Future Network -- Problem statement and requirements -- Part 1: Overall aspects*

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO/IEC TR 29181-1 and the following apply.

To be added when the main body of text is completed.

4 Abbreviations

To be added when the main body of text is completed.

5 Topology

Links between network elements are classified as "physical" and "virtual". Physical links implement "tight synchronisation" as specified in clause 9, using the protocols specified in ISO/IEC 21559-x. Virtual links are those that do not meet that requirement; they are typically implemented over non-FN packet technologies.

An FN "island" is a group of network elements connected to each other by physical links, or a single network element which is capable of supporting physical links.

An FN "cloud" is a group of islands connected to each other by virtual links, or a single island.

An FN "client" is a network element which is not capable of supporting physical links.

An "internal" virtual link connects two islands. An "external" virtual link connects a client to an island.

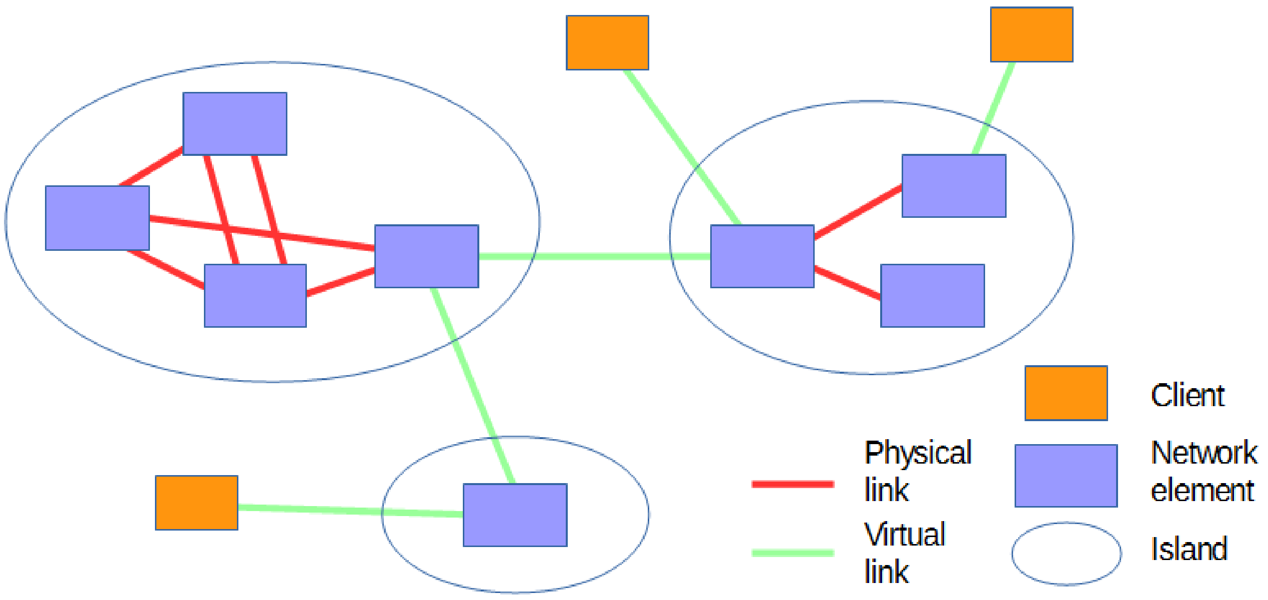

Figure 1 – FN cloud consisting of three islands, with three clients

6 Service interfaces

6.1 Service provided to the layers above

The service provided by FN switching and routing shall be delivery of packets.

Each packet shall consist of a header and a payload, and shall be delivered to one or more destinations, defined by the flow of which it is a part.

The payload shall be an octet string which shall not be processed by intermediate network elements and shall be delivered verbatim except for the effect of transmission errors.

The header shall code information necessary for the forwarding of the packet, as specified in clauses 8 and 9.

6.2 Service required from the layers below

The service required from the layers below shall be as follows.

In the case of a point-to-point physical link, the service shall be conveyance of a stream of data octets, along with formatting which allows the position of each octet in the stream to be identified.

NOTE: One possible kind of formatting is the division of the octet string into frames; another is the periodic insertion of non-data symbols.

Links over broadcast or multicast media, such as wireless, are for further study.

In the case of virtual links, the service is specified in ISO/IEC 21559-x.

7 Flows

7.1 General

Each packet shall be part of a “flow”, which follows a well-defined route through the network. The action to be taken when a packet arrives at a network element shall be defined by the flow to which it belongs.

Globally-significant identifiers for flows, and procedures for setting up flows, shall be as specified in ISO/IEC 62379-5-2.

NOTE: Flow identifiers are used in signalling messages, but not in packet headers.

Three types of flow are supported by FN, as specified in the following subclauses.

7.2 AV flows

An AV flow shall carry AV packets (see clause 9). It shall have a single source, and one or more destinations. Each packet transmitted by the source shall be delivered to all the destinations. There shall be no upper limit to the number of destinations that an AV flow can support.

NOTE 1: This is similar to the service provided by cross-point audio and video routers.

NOTE 2: Because there is only a single source, the sender of each packet arriving on an AV flow is well-defined.

7.3 Connection-oriented IT flows

A connection-oriented IT flow shall carry IT packets (see clause 8). It shall have a single source and a single destination.

NOTE: Because there is only a single source, the sender of each packet on a connection-oriented IT flow is well-defined.

Control plane procedures should support bi-directional flows that are implemented as two connection-oriented IT flows, one in each direction.

7.4 Connectionless IT flows

A connectionless IT flow shall carry IT packets (see clause 8). It shall have one or more sources, and a single destination.

NOTE 1: Connectionless flows are intended for carrying packets from other network technologies such as IP. This International Standard does not specify how the sender of each packet on a connectionless IT flow is identified.

NOTE 2: The form of connectionless service envisaged in ISO/IEC TR 29181-3, with “uplink” and “downlink” paths, is no longer thought to be useful.

8 IT services

8.1 Packet format

The header of an IT packet shall code the payload length and a label; the format shall depend on the formatting of the link across which the flow passes.

The label shall identify the flow; this identification shall be local to each link over which the flow passes, and shall be chosen by the receiving network element.

NOTE 1: A possible value to use for the label is the address of the flow's entry in the receiving network element's routing table.

NOTE 2: The label value, and possibly also the header format, will need to be changed at each switch through which the packet passes.

The header format should include appropriate error detection for the length and the label. The packet format should not include error detection for the payload.

8.2 QoS considerations

The transit time for packets on an IT flow may be undefined. Packets may be lost for lack of space in forwarding buffers.

Error detection in the header should be used to reduce the likelihood of packets being misrouted.

The probability of corruption of payloads should be low, and an estimate of it should be included in the signalling messages that set flows up. On links that use an unreliable physical medium, the data link layer shall include measures to ensure an acceptable level of service.

NOTE: End-to-end detection of errors in payloads, and correction methods appropriate to the application, can be implemented in the transport layer.

8.3 Carriage over other technologies

IT flows may be tunnelled over non-FN technologies by encapsulating the IT packets according to the underlying technology. Encapsulation according to AES51 should be used.

9 AV services

9.1 Synchronisation domains

Two kinds of AV service are specified in this International Standard: “tightly synchronised” and “loosely synchronised”.

Each network element shall ensure that all the physical links to which it is connected are phase-aligned, using the protocols specified in ISO/IEC 21559-x.

The tightly synchronised AV service, which provides the lowest possible latency, can only be used within an island.

The loosely synchronised AV service is used for flows which pass over virtual links. The latency is greater, and de-jitter buffers are required at the receiving end. Islands within a cloud shall be frequency-locked to a common reference, but do not need to be phase-aligned.

9.2 Packet format

The header of an AV packet shall code the payload length; the format should be as specified in ISO/IEC 21559-x.

9.3 QoS considerations

The transit time for packets on a tightly-synchronised AV flow shall be well-defined, and shall be reported in the signalling messages that set the flow up.

Signalling messages that set up a loosely-synchronised flow shall report that the flow is not tightly-synchronised, and should include an estimate of the maximum and minimum transit time and whether packets on virtual links are likely to be dropped.

Packets on either kind of AV flow shall not be lost for lack of space in forwarding buffers in FN network elements.

The probability of corruption of packets should be low, and an estimate of it should be included in the signalling messages that set flows up. On links that use an unreliable physical medium, the data link layer shall include measures to ensure an acceptable level of service.

NOTE 1: The packet header is not used for routing, so corruption of the header cannot cause a packet to be misrouted, though it can cause data to be lost from the end of the packet or spurious data to be added.

NOTE 2: For many applications it is better to deliver a corrupted AV stream than to delete a whole packet for what may be a single bit error.

If a flow is re-routed around a failure, the revised transit time should be reported as soon as possible. If a change in the point of attachment of a mobile device changes the transit time, the new transit time should be reported before the change occurs.

Measures such as forward error correction or retransmission may be implemented in the higher layers, although they increase the transit time. The probability of corruption should be low enough that they are unnecessary for most applications.

9.4 Slots and allocations

Each physical link shall be formatted into “allocation periods” which recur at an interval which shall be (0,99968 ± 0,0001) ms or an integer multiple or submultiple thereof.

NOTE 1: The interval of 0.99968 ms can be regarded as a nominal 1 ms with a margin to accommodate signals whose reference clocks are up to 320 ppm faster than the timing reference used for the allocation periods.

Each allocation period shall contain a fixed sequence of “slots”, each 64 octets in size and thus able to carry one AV packet.

NOTE 2: ISO/IEC TR 29181-3 in its 6.2.2.2 and 6.2.2.3 envisaged the provision of a wide range of slot sizes and packet sizes; however, subsequent experimentation and prototyping has shown that a fixed slot size is preferable.

Each AV flow routed across the link shall be allocated one or more slots per allocation period. The flow to which an AV packet belongs shall be identified by the slot in which the packet arrives.

NOTE 3: An AV packet can therefore be routed without examining its header.

Any slot that does not carry an AV packet, either because it is not allocated to a flow or because the flow does not have a packet available, is an “empty slot”. If the link format does not provide a specific indication that a slot is empty, each empty slot shall be coded with a “null packet” which has the hexadecimal value 40 in its first octet.

NOTE 4: A null packet is thus a zero-length packet with f=1. If f is used as specified in 9.2, inserting or dropping null packets in a flow will have no effect.

Link formats should allow all octets in a slot that are not occupied by an AV packet or “empty slot” header to be used for other services.

9.5 Virtual links

A virtual link shall consist of a tunnel across a non-FN technology which carries IT packets, including FramePhase packets as specified in ISO/IEC 21559-x.

AV packets shall be encapsulated in IT packets; an IT packet may carry more than one AV packet. The AV flow to which each AV packet belongs shall be defined by the IT flow and its position within the IT packet.

NOTE 1: It will usually be convenient to have a 1:1 relationship between IT flows and AV flows, in which case the position of the AV packet in the IT packet is irrelevant. However, this definition also allows an IT flow to be defined such that each IT packet carries, say, two packets from AV flow A followed by one from AV flow B, with any further AV packets being from AV flow C.

NOTE 2: Details of the encapsulation will be established by the protocols that set up the virtual link or that set up the IT flow across it. Such details will include whether each AV packet always occupies a 64-octet slot, or only the length signalled in its header (possibly rounded up to a convenient word size).

10 Signalling service

The signalling service shall be used to carry control plane messages between adjacent network elements.

NOTE 1: Control plane messages include the signalling messages that set up and clear down flows, as well as the protocols used to set up physical and virtual links and maintain synchronisation across them.

Signalling packets shall have the same format as IT packets. On each link they shall have a specific flow label value, which shall be established as part of the protocol which sets up the link.

NOTE 2: On physical links a fixed value could be used, in the same way as VCI 5 in ATM, but where there are several virtual links connected through the same physical port it is convenient to be able to assign a different label value to each.

Outgoing signalling packets shall take priority over IT packets.

Processing of incoming signalling packets at a physical port of a network element shall not be prevented by an overload of incoming IT packets on the same port, nor of any traffic on other ports.

NOTE 3: This is intended to provide a defence against denial-of-service attacks, and can be achieved by having a buffer for incoming signalling packets from each port, separate from the buffers used for forwarding IT packets, and servicing the buffers in rotation.

NOTE 4: It is assumed that a physical port can only carry one physical link, but when connected to a non-FN network can carry multiple virtual links. Whether signalling traffic on those links can be protected from an overload of other traffic depends on the non-FN network.

Carriage of control plane messages between non-adjacent network elements (e.g. to and from a network element with similar functionality to an SDN controller) is for further study.