International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Y.2611

(12/2006)

SERIES Y: GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-GENERATION NETWORKS

Next Generation Networks

High-level architecture of future packet-based networks

ITU-T Recommendation Y.2611

ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS AND NEXT-GENERATION NETWORKS

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
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
Numbering, naming and addressing	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Network control architectures and protocols	Y.2500–Y.2599
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899

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

ITU-T Recommendation Y.2611

High-level architecture of future packet-based networks

Summary

ITU-T Recommendation Y.2611 specifies a high-level architecture for future packet-based networks (FPBNs). This Recommendation also specifies the relationship between an FPBN and the NGN strata and the interfaces in an FPBN.

In order to be able to provide a full suite of services (examples of which include data, video and voice telephony services) to their customers, operators may need to utilize both connectionless packet switched (cl-ps) and connection-oriented packet-switched (co-ps) transport modes. This is because each mode is well suited to the transport of some services and not so well suited to the transport of others.

FPBNs provide the topmost layer(s) of the transport stratum as defined in ITU-T Recommendation Y.2011. The services mentioned above form part of the service stratum as defined in ITU-T Recommendation Y.2011.

Source

ITU-T Recommendation Y.2611 was approved on 14 December 2006 by ITU-T Study Group 13 (2005-2008) under the ITU-T Recommendation A.8 procedure.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. 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, 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 TSB patent database at http://www.itu.int/ITU-T/ipr/.

CONTENTS

ITU-T Recommendation Y.2611

High-level architecture of future packet-based networks

1 Scope

This architecture for an FPBN addresses both connectionless packet switched (cl-ps) and connection-oriented packet switched (co-ps) layer networks. Connection-oriented circuit switched (co-cs) layer networks used to provide the lower layer(s) of the transport stratum are outside of the scope of this Recommendation. The definition and specification of specific services is left to other NGN Recommendations and is outside of the scope of an FPBN and this Recommendation.

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.

[ITU-T G.805] ITU-T Recommendation G.805 (2000), *Generic functional architecture of transport networks*.

[ITU-T G.809] ITU-T Recommendation G.809 (2003), *Functional architecture of connectionless layer networks*.

[ITU-T X.200] ITU-T Recommendation X.200 (1994), *Information technology – Open Systems Interconnection – Basic Reference Model: The basic model*.

[ITU-T X.800] ITU-T Recommendation X.800 (1991), *Security architecture for Open Systems Interconnection for CCITT applications*.

[ITU-T Y.1710] ITU-T Recommendation Y.1710 (2002), *Requirements for Operation & Maintenance functionality for MPLS networks*.

[ITU-T Y.1711] ITU-T Recommendation Y.1711 (2004), *Operation & Maintenance mechanism for MPLS networks*.

[ITU-T Y.2001] ITU-T Recommendation Y.2001 (2004), *General overview of NGN*.

[ITU-T Y.2011] ITU-T Recommendation Y.2011 (2004), *General principles and general reference model for Next Generation Networks*.

[ITU-T Y.2111] ITU-T Recommendation Y.2111 (2006), *Resource and admission control functions in Next Generation Networks*.

[ITU-T Y.2601] ITU-T Recommendation Y.2601 (2006), *Fundamental characteristics and requirements of future packet-based networks*.

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 address: See [ITU-T Y.2601].

3.1.2 authentication: See [ITU-T X.800].

3.1.3 client/server relationship: See [ITU-T G.805].

3.1.4 connection: See [ITU-T G.805].

3.1.5 flow: See [ITU-T G.809].

3.1.6 identifier: See [ITU-T Y.2601].

3.1.7 trail: See [ITU-T G.805].

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

3.2.1 availability: A measure of the capability of a given entity (for example, a layer network, connection, flow, etc.) to maintain connectivity with the associated performance criteria that have been guaranteed by the entity.

3.2.2 name: A name is the identifier of an entity (e.g., subscriber, network element) that may be resolved/translated into an address.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

ATM	Asynchronous Transport Mode
cl-ps	connectionless packet switched
CPE	Customer Premises Equipment
co-cs	connection-oriented circuit switched
co-ps	connection-oriented packet switched
CV	Connectivity Verification
E-NNI	External Network-to-Network Interface
FPBN	Future Packet-Based Network
FT_Sk	Flow Termination Sink
FT_So	Flow Termination Source
HRX	Hypothetical Reference Connection
I-NNI	Internal Network-to-Network Interface
IP	Internet Protocol
L2TP	Layer 2 Tunnelling Protocol
MPLS	Multi-Protocol Label Switching
MTNM	Multi-Technology Network Management
MTOSI	Multi-Technology Operations Systems Interface
NGN	Next Generation Network
NMS	Network Management System
NNI	Network-to-Network Interface
OAM	Operations, Administration and Maintenance
OSI BRM	Open Systems Interconnection Basic Reference Model

OSS	Operations Support System
PM	Performance Management
PPP	Point-to-Point Protocol
PSTN	Public Switched Telephone Network
p-t-mp	point-to-multipoint
p-t-p	point-to-point
QoS	Quality of Service
RACF	Resource and Admission Control Functions
RPT	Reference Point Type
SDH	Synchronous Digital Hierarchy
SLA	Service Level Agreement
TCP	Termination Connection Point
TFP	Termination Flow Point
TMF	TeleManagement Forum
TN	Transport Network
TT_Sk	Trail Termination Sink
TT_So	Trail Termination Source
UNI	User-to-Network Interface
VC-4	Virtual Container Level 4

5 Conventions

None

6 High-level architecture of future packet-based networks

6.1 FPBN architecture

A future packet-based network (FPBN) is composed of packet-based path layer networks (as defined in [ITU-T G.805] and [ITU-T G.809]) in the transport stratum (the functionality is similar to layers 2 and 3 in [ITU-T X.200]). An overview of [ITU-T G.805] and [ITU-T G.809], and the relationships to the open systems interconnection basic reference model (OSI BRM), is provided in Appendix I. The transport stratum is depicted in Figure 1 of [ITU-T Y.2011]. Each layer network 'system' in an FPBN consists of a user plane, a control plane and a management plane and each of the planes within a layer network will have its own traffic forwarding component which may belong to the same layer network (if the planes are not isolated from each other) or different layer networks (if the planes are isolated from each other).

It is a requirement identified in clauses 6 and 7.8 of [ITU-T Y.2601] that an FPBN is expected to:

a) *completely secure the internal control and management plane traffic from external attack and ensure that it remains secure and stable under situations of extreme stress (clause 6)*;

b) *provide mechanisms to protect the control plane communications from security threats (clause 7.8).*

An identical requirement also exists for protecting the FPBN management plane from security threats. The user, control, and management planes (of each layer network) should be segregated

from each other in order to keep the performance, security and reliability of each plane (and that of the other planes) from being violated. Techniques for doing so include (but are not limited to) isolation between planes or special treatment of traffic belonging to the different planes. How a particular NGN network maintains the integrity of its planes is up to it, so long as the requirements detailed in [ITU-T Y.2601] are met.

It is a requirement identified in clause 6 of [ITU-T Y.2601] that an FPBN: *should support off-path control and management planes* and therefore isolation is the preferred 'default' mechanism that can meet the requirements for protecting the user, control and management planes (of each layer network) from each other. The user, control and management planes can be isolated from each other by the allocation of independent connection-oriented packet switched (co-ps) or connection-oriented circuit switched (co-cs) server layer network trails. The type of isolating technology is determined by several factors, such as location (e.g., access or core), network status, etc. It is up to the operator to decide to what degree they wish to operate their control and management planes off-path. Another motivation for isolating the control and management planes from the user plane is to ensure that the FPBN control and management planes continue to operate even if the FPBN user plane is overloaded or faulty.

An FPBN should seek to harmonize functional components (e.g., control and management plane design and operation) across the networking modes as far as practically possible.

Figures 6-1 and 6-2 show functional diagrams that depict the user plane of the FPBN architecture. The connectionless packet switched (cl-ps) network is drawn using G.809 conventions and the co-ps network is drawn using G.805 conventions.

The transport stratum may be implemented by multiple discrete layer networks that form client/server relationships. A different networking mode (cl-ps, co-ps and co-cs) may be used for each of the layer networks (this is not shown in Figures 6-1 and 6-2). The number of layer networks and networking modes used is a choice for the particular operator deploying the transport stratum and is beyond the scope of this Recommendation.

In Figures 6-1 and 6-2, the cl-ps and co-ps layer networks are shown separately. This separation may be physical or logical. The cl-ps layer network may use co-cs server layer network trails that are separate from the server co-cs layer network trails used by the co-ps layer network. Alternatively, the separation may be logical; i.e., the cl-ps and co-ps layer networks share the same server layer network trails. There may be some strict logical partitioning between them so that bandwidth sharing is impossible.

Similarly, the cl-ps layer network may use physically separate networking equipment (e.g., routers) to the co-ps layer network or both layer networks may use the same physical networking equipment but that equipment will be logically partitioned between the cl-ps and co-ps layer networks.

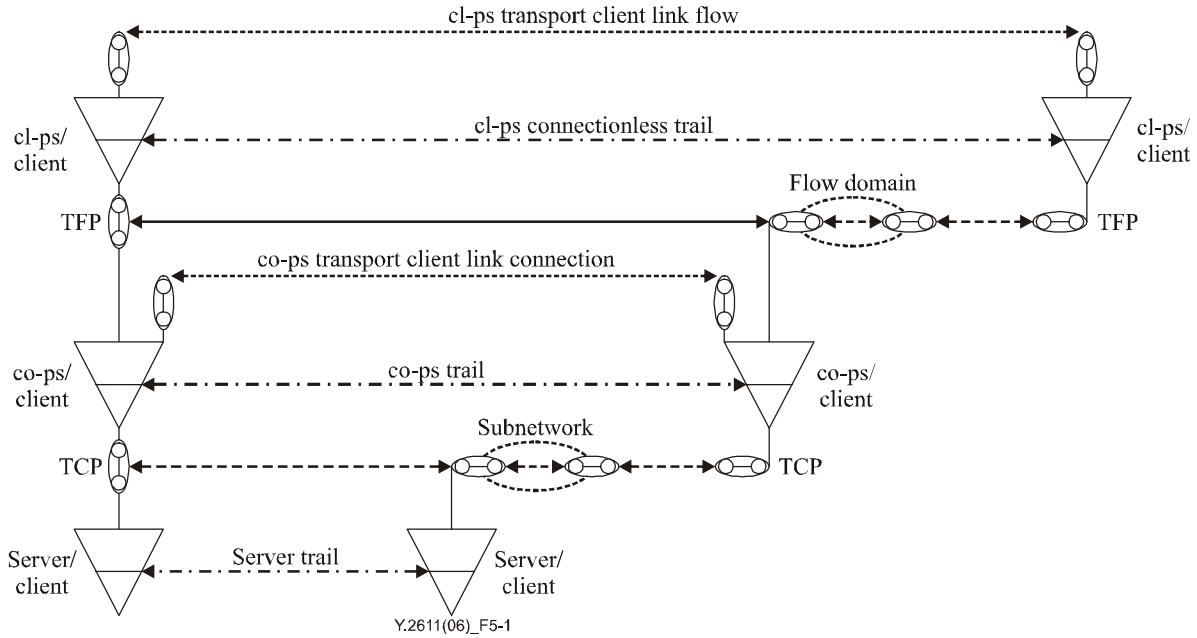

Figure 6-1 – Functional diagram depicting the user plane of the FPBN architecture (cl-ps transport over co-ps layer network trails)

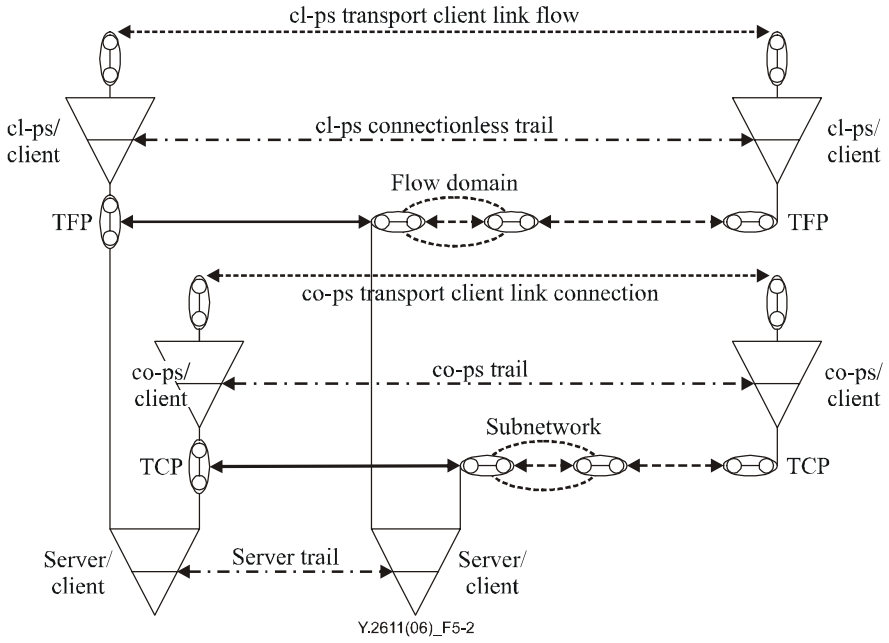

Figure 6-2 – Functional diagram depicting the user plane of the FPBN architecture (cl-ps transport over server layer network trails)

In Figures 6-1 and 6-2, the server layer network trail may be provided by any technology, switched or un-switched. Further client/server relationships may exist below the server layer network trail; however, it should be noted that client layers inherit the impairments of their server layer networks and that this inheritance is recursive down to the duct.

In co-cs layer networks, each client is explicitly allocated a dedicated amount of bandwidth from the server layer network trail. The clients are fully isolated and therefore one client's loading cannot impact the performance of another client. This makes it simple to guarantee dedicated bandwidth for a client.

In co-ps layer networks, each client is allocated bandwidth from a server layer network trail. However, as the clients are only logically isolated, one client's loading may directly impact the capacity available to another client. The appropriate allocation of bandwidth and the use of ingress admission control and policing make it possible to guarantee dedicated bandwidth for a client.

In cl-ps layer networks, flows are not normally explicitly allocated to server layer network trails. Therefore, the capacity available to one client flow may be impacted by the loading of other client flows. This may be mitigated by engineering the appropriate capacity in the server layer network (i.e., over-provisioning) or by establishing resource reservation state per-hop and pinning routes. This makes it possible to guarantee dedicated bandwidth for a client. This procedure is implicit in a co-ps layer network. However, these techniques are not generally used for the majority of traffic in cl-ps layer networks.

NOTE – Due to the different characteristics of each networking mode it is generally advisable to stack modes that less efficiently provide dedicated bandwidth on top of modes that more efficiently provide dedicated bandwidth.

Looking at the top of Figure 6-1 and working down the model shows cl-ps transport over cl-ps layer network connectionless trails, which are transported over co-ps layer network trails which are in turn transported over server layer network trails. co-ps transport is provided over co-ps layer network trails, which are in turn transported over server layer network trails.

Looking at the bottom of Figure 6-1 and working up the model shows a server layer network trail providing transport for a co-ps layer network. The co-ps layer network in turn provides transport for co-ps services as well as providing transport for the cl-ps layer network. Then the model shows that the cl-ps layer network provides transport for cl-ps services.

Looking at the top of Figure 6-2 and working down the model shows cl-ps transport over cl-ps layer network connectionless trails, which are in turn transported over server layer network trails. co-ps transport is provided over co-ps layer network trails, which are in turn transported over server layer network trails.

Looking at the bottom of Figure 6-2 and working up the model shows a server layer network trail providing transport for a co-ps layer network and a cl-ps layer network. The cl-ps layer network provides cl-ps transport and the co-ps layer network provides co-ps transport.

An operator may choose to use either of the options depicted in Figures 6-1 and 6-2 above (i.e., co-ps transport or other server layer network trails) in order to support a cl-ps layer network. Alternatively, an operator may choose to mix the above options (i.e., co-ps transport and other server layer network trails), so, for example, an operator may choose to use co-ps transport for some cl-ps connectionless trails and other server layer network trails for other cl-ps connectionless trails. One reason for mixing these options is that some cl-ps links within the operator's network may require the larger/coarse bandwidth granularities provided by server layer network trails, whereas other cl-ps links may require a finer bandwidth granularity. In order to maximize the utilization of the larger/coarse bandwidth granularities provided by server layer network trails, the operator may wish to utilize the co-ps layer network as a method of mediation between the server layer network trails and cl-ps layer networks.

For the cl-ps layer network depicted in Figures 6-1 and 6-2, the choice of using co-ps transport or other server layer network trails or both will be a decision taken by the operator and will be dependent on a number of factors both economic and technical including (but not limited to):

- the operator's local policy;
- the traffic level guarantees the operator has made to their customers;
- the level of bandwidth granularity a given service requires;

- the volume of cl-ps traffic which is being aggregated, i.e., small volumes are likely to be better served off the co-ps mode, whilst larger volumes are likely to be better served off the co-cs mode.

The specific encapsulation format used for an FPBN user plane is independent of network mode. It is the control plane or management plane that commonly determines the network mode. Therefore, operators may use the same encapsulation format for both cl-ps and co-ps network modes even though the forwarding behaviour of each mode is different.

6.2 User plane

User plane resources may be allocated to different service classes, so as to adapt to the open market, competitive circumstance, services implementation and evolution.

Resources of service classes will be allocated on demand. Resources allocated to service classes are independent of each other. Different service classes have different attributes. For example, some service classes may guarantee the packet loss ratio and delay of packet transport, some may guarantee packet "importance", some may guarantee much higher security for packets, some may provide guaranteed throughput for packet streams, and some others may provide combinations of these above attributes or even combinations with some other attributes.

It is not necessary to provision all services in a service class in the same way in an FPBN. The control plane may set up some of them, while the management plane may set up others.

As the service stratum may require a large number of service classes with different attributes, an FPBN should provide service classes in an extensible way. There are many advantages to doing so; for example, voice service can be put into an independent service class so that traditional PSTN carriers can provide consistent voice service characteristics. As another example, an FPBN could provide "carrier of carriers" services so that the transport carrier and the service carriers can be different operators, etc.

It is a requirement identified in clause 7.11 of [ITU-T Y.2601] that an FPBN is expected to support:

a) *point-to-point transport stratum services without adaptation;*

b) *point-to-point transport stratum services including adaptation functions;*

c) *point-to-multipoint transport stratum services including adaptation functions.*

Such transport stratum services may support link connections (or link flows) within the service stratum or within other layer networks within the transport stratum. Such link connections (or link flows) may be operated by entities other than the entity that operates the FPBN layer network that is providing the transport stratum service upon which those link connections (or link flows) are built. It is clear that a client/server relationship exists between a link connection (or link flow) and the transport stratum service that supports that link connection (or link flow). It is also clear that in order for an FPBN to be able to support different entities operating different layer networks within the transport or service strata, the client and server layer networks within such a client/server relationship must be separated such that the server layer network can provide transparent (and client agnostic) transport to the client layer network.

NOTE – When a client link connection (or link flow) extends beyond an FPBN transport stratum service without adaptation, the transport stratum service only provides transport for part of that link connection (or link flow), and adaptation is provided outside of the FPBN.

6.3 Control plane

The control plane configures the user plane to forward traffic from its source to its destination. The control plane will set up and maintain user plane service classes by allocating and scheduling FPBN resources according to the requirements of the services that an FPBN supports.

To support NGN services that require quality of service (QoS), the FPBN control plane should support a resource and admission control function (RACF) [ITU-T Y.2111].

The identifier space of the control plane may be independent of any other identifier spaces in an FPBN; see clause 6.10 for more details.

The control plane of a layer network should be physically or logically segregated from the other planes of that layer network. Control plane communications may use user plane trails or may use logically or physically segregated trails.

The user plane may rely on control plane mechanisms in order to provide survivability and robustness against failures. Therefore, the survivability design of the control plane is likely to be different to the survivability design of the user plane. In the case where the user plane relies on control plane mechanisms in order to provide survivability and robustness, then the diversity of the topology of the control plane communications should be at least as great as the diversity provided to the user plane.

An FPBN may provide both cl-ps and co-ps user planes in order to provide both cl-ps and co-ps transport stratum services. The cl-ps user plane will be independent of the co-ps user plane and each user plane will have its own control plane.

Although the control plane of the cl-ps user plane will be isolated from the control plane of the co-ps user plane, it is likely that there will be some overlap between the functions and features provided by both control planes. For example, both control planes may use a routing protocol to distribute the topology of the user plane that they are controlling. An FPBN should reuse as many functions and features as possible where such functions and features are required in both control planes. For example, if both control planes require a routing protocol, then they should both use the same routing protocol; however, the exact syntax and semantics of the routing protocol messages may differ between the two networking modes as the topology information that needs to be distributed by each mode and the requirements placed on each mode are not identical.

6.4 Management plane

The management plane provides configuration, fault reporting, billing, security, and performance management for an FPBN.

The identifier space of the management plane may be independent of any other identifier spaces in an FPBN; see clause 6.10 for more details.

The management plane of a layer network should be physically or logically segregated from the other planes of that layer network. Management plane communications may use user plane trails or may use logically or physically segregated trails.

An FPBN may provide both cl-ps and co-ps user planes in order to provide both cl-ps and co-ps transport stratum services. The cl-ps user plane will be independent of the co-ps user plane and each user plane will have its own management plane.

6.5 OAM, performance management and availability

It is a requirement as identified in clauses 6 and 7.4 of [ITU-T Y.2601] that an FPBN is expected to:

a) offer the appropriate operations, administration and maintenance (OAM) functions for each plane (clause 6);

b) support network performance monitoring (PM) including availability, packet loss, delay and jitter between any two points in the network (clause 7.4).

OAM, performance monitoring and availability are related and this clause discusses aspects of each function individually and then goes on to discuss the relationships between them.

A layer network has two basic states: fully working, or broken to some degree. However, a specific client (service or layer network) of that layer network will only see either a working service (perhaps with some level of impairment) or a broken service.

If the layer network is a co-ps layer network, then its trails have two basic states: available (and working within its performance objectives) or unavailable. Both these states are deterministic and can be fully specified. However, it is not possible to describe a cl-ps layer network so easily because cl-ps layer networks do not have a trail construct and therefore cl-ps layer networks can have a far wider range of what may be considered as impaired versus broken behaviour.

Within a well designed and well engineered network, defects and performance degradation should be rare. However, there will be failures and/or performance problems from time to time and therefore OAM is required in order to detect and manage such problems. There are two broad categories of OAM: proactive fault detection ('always on') OAM and reactive fault location/diagnostic ('on demand') OAM.

Proactive OAM is generally responsible for the rapid detection of defects (for example, by using connectivity verification (CV) flows) and initiating the necessary consequent actions. Proactive OAM should be as simple as possible so that the cost of continuously processing OAM flows is minimized. This cost of processing includes operational as well as capital costs (historically, the operational cost of enabling continual OAM monitoring has been very high for some networking technologies which has resulted in operators disabling the proactive OAM in some of their layer networks). Proactive OAM should not be burdened with the complexity required for fault diagnosis or fault location identification. The role of proactive OAM is simply to detect defects in a layer network and perform the necessary consequent actions (which may include triggering reactive OAM).

Reactive OAM is responsible for providing and performing the more complex OAM functions that the proactive OAM does not perform, for example performance management measurements, defect diagnosis, defect location identification and tracing functions. These more complex OAM functions are not normally performed by proactive OAM for a number of reasons, including (but not limited to): complex OAM functions need not be performed continuously and the additional cost that they would add, to the proactive OAM component, is considered to be too large.

Performance monitoring (or performance management) is the measurement of transfer performance for a given trail when that trail is in the up state. As noted previously in clause 6.1, client layer networks inherit the impairments of their server layer networks and this inheritance is recursive down to the duct. Therefore, the performance of a given trail is defined by the performance impairments inherited from its server layer networks plus the additional impairments introduced by the trail itself (from the layer network it is part of). This inheritance between client and server layer networks leads to a requirement that a server layer network's performance criteria is expected to be at least as stringent as its most stringent client layer network in order for the server layer network to be able to meet the performance criteria demanded by its client layer networks.

The availability of a given layer network is essentially a measure of the capability of that layer network to maintain connectivity in spite of one (or more) defects or failures. As noted previously, because a link connection (or a link flow) in a client layer network is supported by a trail (or a connectionless trail) in that client's server layer network, a client layer network inherits certain characteristics (such as link diversity) from its server layer network and this inheritance recurses down to the duct. This means that regardless of where in the network stack a given layer network is situated, its ability to effect disjoint routing is closely coupled to the available physical duct topology. Therefore, it is impossible to achieve routings in a client layer network that are more diverse than the physical duct topology.

In order to efficiently manage a layer network, that layer network's OAM, performance management and availability must be designed and processed in a logical order so that that layer

network's OAM, performance management and availability mechanisms are extensible without adverse impact to that layer network or to the operator that 'owns' that layer network.

The recommended logical ordering is as follows. The differences between the needs and requirements of proactive OAM and reactive OAM should be well understood and then the network mode (i.e., cl-ps or co-ps for FPBNs) that the OAM will be operating in must be identified. This is because each of the two packet switched networking modes has different characteristics, defects and consequently different OAM requirements.

For each mode it is necessary to define suitable and appropriate OAM for defect detection and handling (i.e., proactive OAM), including defining which defects can occur in that networking mode. For example, there is a common requirement for both packet switched networking modes to provide a connectivity verification (CV) mechanism and therefore both packet switched networking modes must provide a mechanism that allows a trail's termination sink to identify that trail's termination source. In the cl-ps mode, the CV function effectively 'comes for free' because each and every packet contains a source address. However, verifying the connectivity of a cl-ps layer network that only supports transit flows requires some additional proactive OAM functionality. In addition to its other functions, a periodic CV flow between a pair of flow or trail termination points can be used to distinguish whether that flow or trail is quiescent or broken.

For each defect identified in a given networking mode, it is necessary to define a set of entry and exit criteria (for the available and unavailable states) based on defect persistency as well as a set of consequent actions for that defect. The exact entry and exit criteria and consequent actions will depend upon the nature of the defect and the networking mode it applies to.

Once the available defects, their entry and exit criteria and any consequent actions have been defined (for the networking mode being considered), only then is it possible to start to address mechanisms for taking performance management measurements and assessing a given trail, connection, flow or layer network against any performance management service level agreements (SLAs) that have been agreed. This is because performance measurements, at least for SLA purposes, are only meaningful when the network entity considered is in the available state.

It should be noted that it is not just performance management that is dependent on the correct order of processing as outlined above. Some other examples include:

- network element specification (in terms of registers and threshold crossing exception reports);
- network management systems/operations support systems (NMSs/OSSs) that have to process network element collected data about defects, availability and performance management;
- the definition of hypothetical reference connections (HRXs) and suitable end-to-end and apportioned availability and performance management objectives;
- the definitions of consistent network services with measurable SLAs.

6.5.1 OAM, performance management and availability of co-ps layer networks

A layer network requires some mechanism (or mechanisms) to enable it to differentiate the up state from the down state for a given p-t-p connection, which in turn allows that layer network to measure against any performance SLAs that have been agreed for a given connection in that layer network. There is a requirement for the up and down states to have been clearly identified before we can consider performance management because performance SLAs are only meaningful when the connection they refer to is in the up state.

NOTE 1 – If a co-ps service is being guaranteed by the transport stratum, then by implication the transport stratum has a "call admission policy" to prevent over-subscription and the consequent performance degradation.

NOTE 2 – A p-t-mp trail can be considered as a set of p-t-p connection instances between a source and a specific sink. From the perspective of a given client instance, the only thing of concern is whether that client's source/sink p-t-p connection is working or not. Therefore, p-t-mp connectivity can be discussed in terms of p-t-p connectivity behaviour.

NOTE 3 – In general, the server layer network's protection or restoration is designed such that it can recover the connection in the event of a failure before the connection is declared to be unavailable.

NOTE 4 – In general, a transport network is monitoring a transit connection i.e., the service trail terminations are not within the scope of the transport network.

The minimum set of possible defects within a co-ps layer network that proactive co-ps OAM should be capable of detecting is as follows.

Loss of connectivity – This defect occurs when traffic originating from the co-ps trail termination source does not arrive at the corresponding co-ps trail termination sink. For example, for a co-ps trail between trail termination source A (TT_So A) and trail termination sink A (TT_Sk A), traffic originating from TT_So A does not arrive at TT_Sk A.

NOTE 5 – Due to congestion or packet loss, a certain degree of lost connectivity may be deemed acceptable within a co-ps layer trail. Consequently, a loss of connectivity defect should only be raised once connectivity has been lost for a sustained period of time as defined in the entry and exit criteria for the loss of connectivity defect.

Incorrectly connected connection – This defect occurs when, for whatever reason (for example, failures or incorrect operator configuration), a given trail termination source is connected to the incorrect trail termination sink. For example, a trail that should connect TT_So A to TT_Sk A is instead connected to TT_Sk B.

Incorrectly merged connections – This defect occurs when, for whatever reason (for example, failures or incorrect operator configuration), traffic in one trail is 'leaking' into another trail. For example, for a co-ps trail between TT_So A and TT_Sk A, traffic arriving at TT_Sk A is originating from both TT_So A and TT_So B.

Entry and exit criteria for the above defects are not defined in the FPBN architecture. However, definitions for defect entry and exit criteria along with definitions for when a connection is considered available or unavailable must be defined in order to allow HRXs with performance apportionments to be specified.

[ITU-T Y.1710] specifies requirements for OAM functionality in MPLS networks and [ITU-T Y.1711] specifies an operation and maintenance mechanism for MPLS networks. Although specific to MPLS networks, the principles contained in [ITU-T Y.1710] and [ITU-T Y.1711] can be generalized and applied to any co-ps layer network and co-ps OAM mechanisms in an FPBN should reuse the general principles of [ITU-T Y.1710] and [ITU-T Y.1711] as appropriate to the specific co-ps layer network technology used.

For services that provide bidirectional connectivity between two communicating entities, if one direction enters the down state then the service (i.e., both directions) should enter the down state (i.e., the service should be considered unavailable). Therefore, the collection of performance management measurements for a bidirectional connection must be suspended in both directions even if only one direction of the connection is defective (i.e., in the down state).

The availability of a given connection is essentially a measure of the capability of that connection (or more precisely the layer network that that connection belongs to) to maintain connectivity (with the associated performance criteria that that connection has guaranteed) in spite of one (or more) defects or failures.

6.5.2 OAM, performance management and availability of cl-ps layer networks

In general, it is not feasible to individually monitor the state of all flows within an FPBN cl-ps layer network. In addition, it is also not feasible for an FPBN to individually monitor the state of all

service stratum sessions. This is in part due to the large number of flows that may exist at any one time and the short-lived nature of many of those flows.

It is however feasible to monitor the connectivity (i.e., the ability to transfer packets between two points) in a cl-ps layer network. A cl-ps layer network therefore requires some mechanism (or mechanisms) to enable it to differentiate whether or not connectivity exists between two points within that cl-ps layer network. This in turn allows a cl-ps layer network to measure any guarantees that have been agreed for connectivity between two points in that cl-ps layer network. Additionally, an FPBN should be able to detect when packets are delivered to an unintended destination(s)/egress(es).

The minimum set of possible defects within a cl-ps layer network that proactive cl-ps OAM should be capable of detecting is as follows.

Loss of connectivity – This defect occurs when traffic originating from a cl-ps flow termination source does not arrive at the corresponding cl-ps flow termination sink. For example, for a cl-ps flow between flow termination source A (FT_So A) and flow termination sink A (FT_Sk A), traffic originating from FT_So A does not arrive at FT_Sk A.

NOTE – Due to congestion or packet loss, a certain degree of lost connectivity may be deemed acceptable within a cl-ps layer flow. Consequently, a loss of connectivity defect should only be raised once connectivity has been lost for a sustained period of time as defined in the entry and exit criteria for the loss of connectivity defect.

Packets within a cl-ps layer network always contain a unique (within the context of that layer network) source address and therefore cl-ps layer network packets are always self-identifying with respect to their source. This means that cl-ps layer networks only multiplex, and never merge, flows and therefore a cl-ps layer network cannot experience misconnected flow defects or mismerged flow defects.

Entry and exit criteria for the above loss of connectivity defect are not defined in the FPBN architecture. However, definitions for defect entry and exit criteria along with definitions for when a flow is considered available or unavailable must be defined in order to allow HRXs with performance apportionments to be specified.

Flows are always unidirectional; however, many services require bidirectional connectivity and therefore it is often necessary to monitor the connectivity of both directions between two points in a cl-ps layer network. For services that provide bidirectional connectivity between two communicating entities, if one direction loses connectivity, then the service (i.e., both directions) should enter the down state (i.e., the service should be considered unavailable). Therefore, the collection of performance management measurements between two points in a cl-ps layer network must be suspended in both directions even if the loss of connectivity is only in one of the directions.

The availability between two points in a cl-ps layer network is essentially a measure of the capability of that layer network to maintain connectivity with the associated performance criteria that have been guaranteed.

6.6 Relationship between layer networks and the OSI BRM

The X.200 model and the G.805/G.809 model are useful in describing different aspects of the transport stratum. In general the X.200 model is most useful when describing the horizontal relationships (between peered layers) and functions between layers within a single stack. The G.805/G.809 model is most useful when describing the recursive interlayer relationships in multilayer transport networks. The term *layer* is used when applying the X.200 model and the term *layer network* is used when applying the G.805/G.809 model. The definition of 'layer network' used in G.805/G.809 is not the same as the definition of 'layer' used in X.200. Both X.200 and G.805/G.809 are widely used within the industry to describe networks. A brief overview of the X.200 model and the G.805/G.809 model is provided in Appendix I.

6.7 Relationship with other strata

See Figure 6-3.

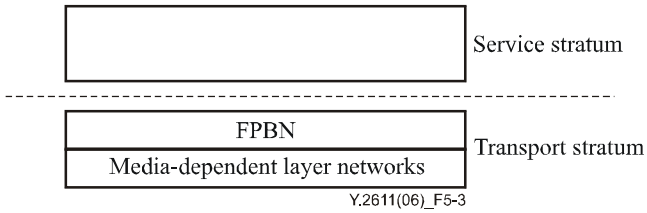

Figure 6-3 – Relationship between an FPBN and the transport and service strata

An FPBN is located between the service stratum and the lower part of the transport stratum from a interlayer point of view (as defined in [ITU-T G.805] and [ITU-T G.809]). An FPBN may provide co-ps and/or cl-ps transport stratum services. The FPBN may be implemented with multiple layer networks as described in clause 6.1.

For transparency, an FPBN is independent of any lower (server) layer networking (media dependent) technologies. The lower (server) layer network provides the necessary adaptation functions and transport services required in order to interconnect FPBN nodes.

FPBN packets may be adapted onto (i.e., encapsulated in) both present and future server layer networking technologies.

6.7.1 Relationship between an FPBN and its client (service or layer network)

As required by [ITU-T Y.2001] an FPBN should act as a server layer and therefore must be independent of its client layers. The client layer packets, whether they are user packets, management packets or control packets, are all treated as the payload of an FPBN user plane.

Client layers can be carried over cl-ps, co-ps or both transport modes as long as the service requirements of the client layers are satisfied.

One service can be mapped into one or more than one service class.

6.7.2 Relationship between an FPBN and its server layer network

An FPBN should act as a client of its underlying server layer and therefore the server layer must be independent of the FPBN.

6.8 Relationship between an FPBN and existing networks

It is a requirement identified in clause 6 of [ITU-T Y.2601] that an FPBN is expected to *interwork and co-exist with current cl-ps and co-ps packet networks*. In order for an FPBN to interwork with current cl-ps and co-ps networks, it may be necessary to perform address translation and other functions at the boundary of an FPBN.

6.9 Interfaces in an FPBN

An FPBN can serve as a core network and/or an access network, which may belong to different operators. An FPBN can interconnect remotely with another FPBN and/or connect with other heterogeneous transport networks.

Consider the network interconnection scenarios depicted in Figure 6-4 in clause 6.10 below, in which the reference points of an FPBN are defined. In this figure, the core transport network may be connected to one or more access transport networks, and each access transport network may be connected to one or more user networks.

FPBN A is interconnected with the adjacent FPBN B; at the same time, it is also interconnected with FPBN C. FPBN D belongs to a different operator to the operator that owns FPBNs A, B and C. FPBN D is connected with FPBN A but is not trusted by it. Another transport network (marked 'Other transport network' in Figure 6-4) is heterogeneous but connected with FPBN A and it is also not trusted by FPBN A.

6.10 Reference points in an FPBN

The reference points for a layer network within an FPBN are classified as types a, b, c, d, e, or f. The network interfaces include user-to-network interfaces (UNIs), internal network-to-network interfaces (I-NNIs), and external network-to-network interfaces (E-NNIs).

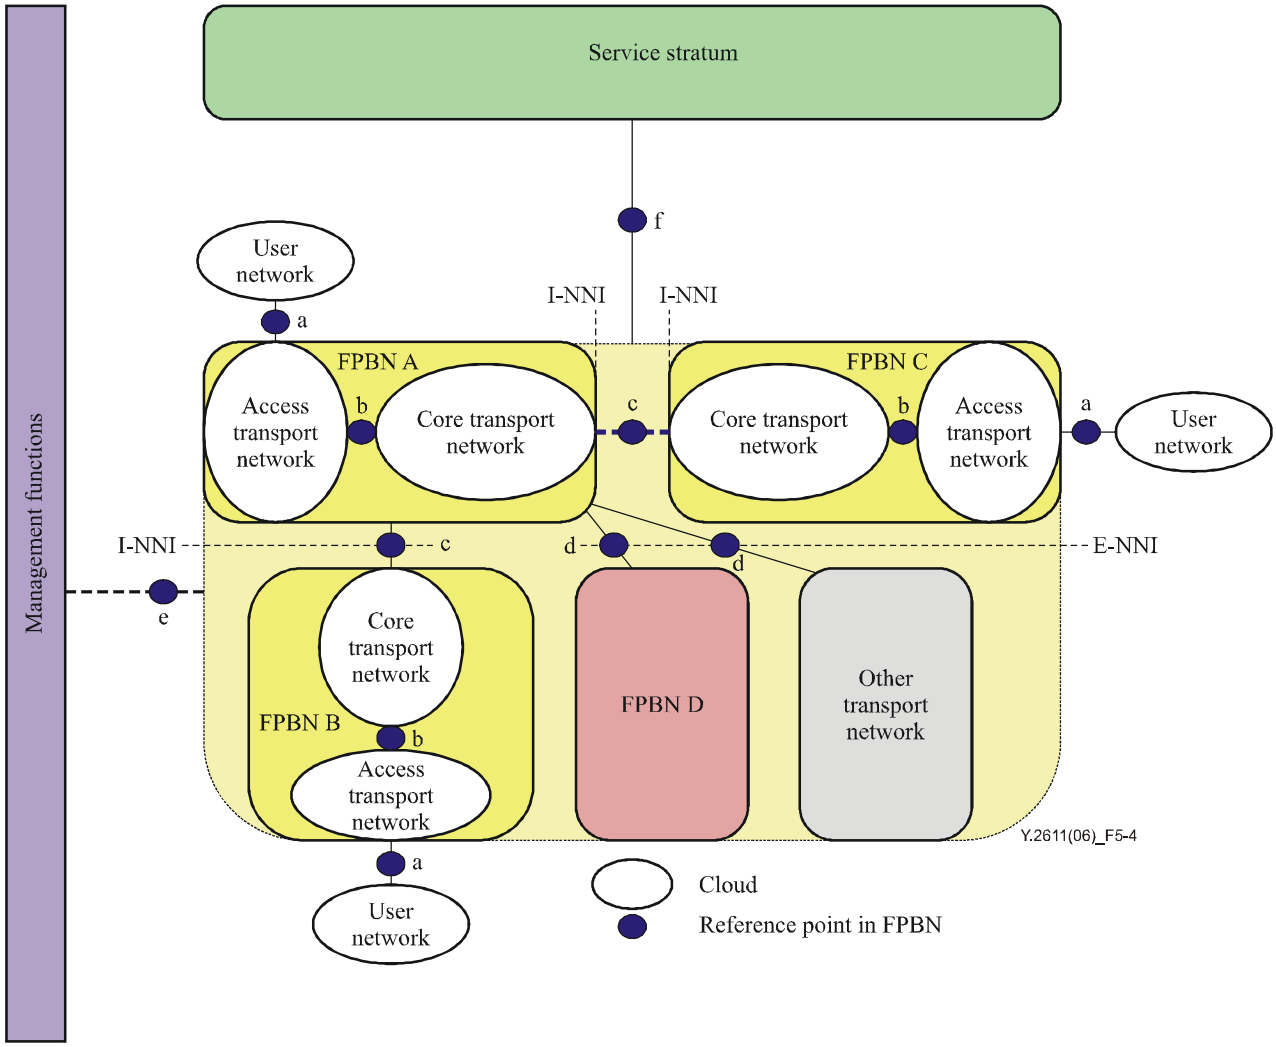

Figure 6-4 – Reference points in an FPBN

In Figure 6-4 each FPBN shown consists of an access transport network and a core transport network. However, the access transport network or the core transport network may be null. In other words, an FPBN may only support an access transport network or a core transport network but not both.

A user network could be a home network, an enterprise network, or some other network.

Reference point type (RPT) exists between a user network and an FPBN access transport network. It allows the user to transfer and receive user data, OAM and signalling information.

RPT a may support more than one service instance within an NGN.

RPT b is located between an FPBN access transport network and an FPBN core transport network. It acts as an aggregation point for the FPBN core transport network.

RPT c represents an FPBN I-NNI and is located between two adjacent FPBN core transport networks. A single FPBN I-NNI may support more than one service instance destined to different destinations.

RPT d represents an FPBN E-NNI and is located between two FPBNs that belong to different operators, or an FPBN and a heterogeneous transport network. A single FPBN E-NNI may support more than one FPBN service instance destined to different destinations in either operator's network.

RPT e represents the management interface between the management plane of a layer network that belongs to the transport stratum and any network management functions which are outside of that layer network's management plane.

RPT f represents the interconnection point between the transport and service strata within an NGN.

6.11 Naming and addressing in an FPBN

An FPBN needs an addressing mechanism to identify a node, a link, an interface or other entities.

Identification is required in each layer network of the NGN transport stratum. A given entity will be assigned one or more identifiers depending on the function of that entity. FPBN layer network identifiers are independent of any client (and any server) layer network identifiers even if they share the same syntax or structure. At the boundary of a layer network, mapping and/or translation mechanisms are required in order to set up relationships between the identifier used by the client layer network and the identifier used by the server layer network.

NOTE – Identifiers could be determined from multiple discontinuous fields. The global uniqueness of an identifier may be provided by the context as well as the identifier itself.

Whether a given identifier is considered to be a name or an address is dependent on several factors including the perspective (and location) of the entity that is using (or mapping to) that identifier. The same identifier can be considered an address to one entity and a name to a different entity because their perspective is different.

An FPBN may require multiple identifier spaces, for example user, management and control plane identifier spaces. Each identifier space may be independent of the other identifier spaces (even if they use the same syntax or structure). Additional identifier spaces may also be used, for example to allow independent identification of the components that implement control plane functions.

Each resource at a network boundary of the user plane of each layer network will have a name (from the user plane name space of that layer network) which is visible to the exterior of the network. These names may need to be translated into topologically significant addresses (from the user plane address space of that layer network) on the interior of the layer network boundary. In other words, resources on the interior of a given layer network use addresses. When these resources are made visible to entities on the exterior of that layer network, a name may be provided instead of the interior address.

Identifiers within a layer network are administered by the owner of that layer network and must be unique within that context. Any identifiers that are made externally visible are administered within the boundaries of the enclosing network to ensure that they are unique within that context.

6.12 Security considerations

An FPBN requires mechanisms to make it safe or "trusted" by its client layer.

An entity can be said to 'trust' a second entity if the first entity assumes that the second entity will behave as expected by the first entity. Such an assumption of trust relies on the identity of the second entity being authenticated.

Appendix I

Relationship between layer networks and the OSI BRM

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

This appendix highlights and clarifies the key differences between the G.805/G.809 model and the X.200 model in order to assist a practitioner of one model to achieve an appreciation of the other model. This appendix is not intended to be a complete description of either the G.805/G.809 or X.200 models.

I.1 The OSI BRM (X.200) model

The Open Systems Interconnection Basic Reference Model (OSI BRM) [ITU-T X.200] is normally applied to describe a single 'network stack' from the application layer to the physical transport layer. [ITU-T X.200] describes a single network in terms of the logical functions that form the network and the hierarchy that exists between those logical functions at different levels within the network.

When describing a network, [ITU-T X.200] assumes that there is only a single 'network stack' (a single open system), and that this contains a hierarchy of (up to seven) different layers that are named and organized according to their functions: Applications, Presentation, Session, Transport, Network, Data Link and Physical. Generally, the transport stratum in the NGN architecture could be represented by the lower three layers in the OSI BRM, i.e., the Network, Data Link and Physical layers.

The Network layer plays an important role in providing the interface between the service stratum and the transport stratum. The core function in the Network layer is routing and relay. It provides the service stratum with connection-oriented mode (co-ps) or connectionless mode (cl-ps) layer network services. The layering of the transport stratum based on the X.200 model is shown in Figure I.1.

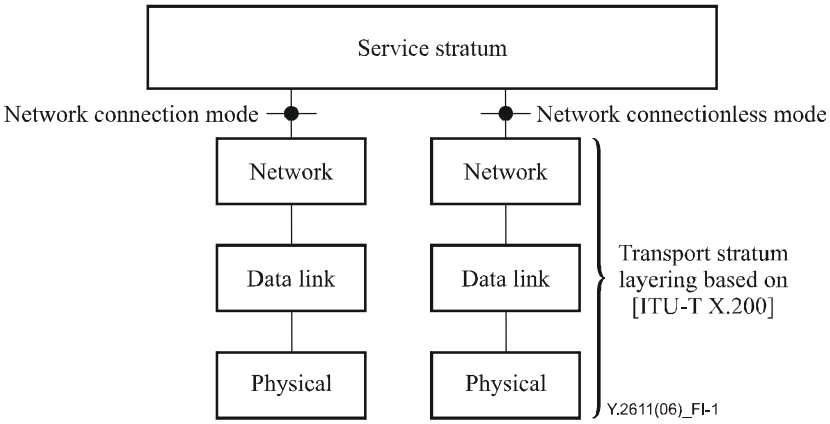

Figure I.1 – Transport stratum layering based on the X.200 model

I.2 The G.805/G.809 model

The G.805/G.809 model is used to describe "layer networks" within the transport stratum. Thus, the G.805/G.809 model includes the concept of recursion, i.e., one layer network can be the client of another layer network. This is known as a client/server interworking relationship. G.805/G.809 provide a set of tools and rules that allows us to visualize complex transport networks that are multi-operator and multi-technology.

I.3 Comparing the two models

In contrast to X.200, G.805/G.809 assumes that a single layer network may contain all of the functions described in [ITU-T X.200]. In G.805/G.809, a layer network may be one of many that co-exist in parallel (either completely independently of each other or nested in client/server relationships), each of which have their own set of functions that map to the functions that are described by the OSI BRM (termed "layers" in [ITU-T X.200]). [ITU-T G.805]/[ITU-T G.809] does not restrict the functions that can exist within a layer network, which allows the G.805/G.809 model to describe a layer network (or stack of layer networks) to whatever level of abstraction is most appropriate. Similarly, [ITU-T G.805]/[ITU-T G.809] does not restrict the number of layer networks that can exist within a 'network stack', which allows G.805/G.809 models to describe a possibly infinite number of client/server relationships between layer networks in the 'network stack'.

A single layer network as described by [ITU-T G.805]/[ITU-T G.809] does not map directly to a single layer as described in [ITU-T X.200]. In fact, client/server relationships between G.805/G.809 layer networks allow for them to function independently, and each layer network has its own instantiation of the OSI BRM which is distinct from any instantiation of the OSI BRM in any parallel layer network. This includes both horizontally and vertically parallel layer networks. However, layer networks (as described by [ITU-T G.805]/[ITU-T G.809]) need not instantiate all seven layers of the OSI BRM.

This is not to say that functionality resembling that described in the OSI BRM is not present in layer networks (as defined by [ITU-T G.805]/[ITU-T G.809]), but rather that the functionality may be distributed quite differently, say across a fewer or greater number of functions, or just simply distributed differently, and not 'layered' in the same rigid hierarchical fashion as that specified in the OSI BRM.

NGN architectures require a greater amount of flexibility than was envisaged when [ITU-T X.200] was developed. Further details can be found in clause 6 of [ITU-T Y.2011] where the relationship between NGNs and X.200/OSI BRM is discussed in more detail. Annex A of [ITU-T Y.2011] identifies some areas of X.200 which are either too restrictive and/or insufficient to accommodate recent, emerging or expected future technologies. Additionally, Annex B of [ITU-T Y.2011] contains a detailed list of items retained from [ITU-T X.200] (since they are applicable to NGN) and a list of items not retained (since they are not applicable to NGN) from [ITU-T X.200].

Figure I.2 shows how each layer network (as described by [ITU-T G.805]/[ITU-T G.809]) has its own instantiation of the OSI BRM which is distinct from any other instantiation of the OSI BRM that exists in any parallel layer networks. Figure I.2 shows a scenario in which an Ethernet layer network is supported by an MPLS layer network that is, in turn, supported by a SDH layer network. Each layer network is depicted using the diagrammatic conventions described in [ITU-T G.805]/[ITU-T G.809]. Alongside each layer network, an instantiation of the X.200/OSI BRM is shown to highlight that each of the three layer networks (Ethernet, MPLS and SDH) co-exist (nested in client/server relationships) and each of them has their own set of functions that map to the functions that are described by X.200/OSI BRM.

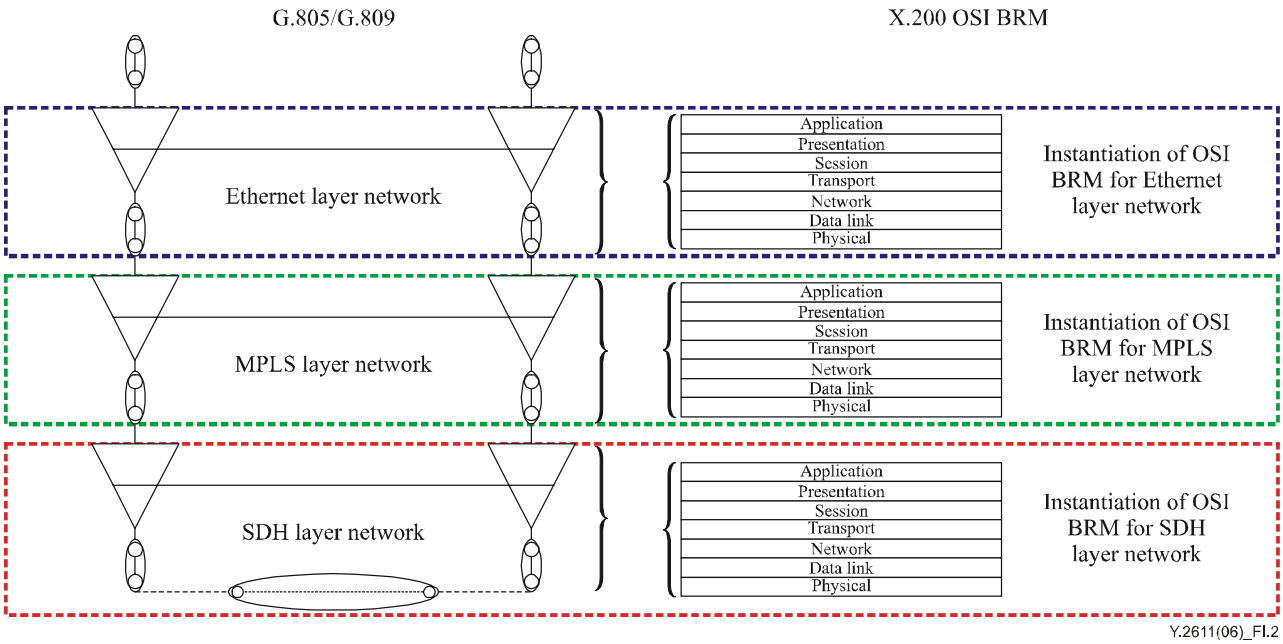

Figure I.2 – How each G.805/G.809 layer network has its own instantiation of the OSI BRM

Note that a layer network does not necessarily instantiate all seven layers of the OSI BRM (for example MPLS would not instantiate the OSI BRM physical layer). It is also worth noting that Figure I.2 shows a hierarchy of layer networks at a given level of abstraction. The G.805/G.809 model allows a layer network to be described at any level of abstraction, so for example the diagram could be expanded in order to decompose the SDH network into its constituent layer networks (VC-4, multiplex section, regenerator section, etc.).

In addition to providing a model for describing layer networks (and their layering and interactions), the G.805/G.809 model can also be mapped into detailed equipment specifications (for example [b-ITU-T I.732] provides an ATM equipment specification and [b-ITU-T G.783] provides an SDH equipment specification) as well as management information models (for example TeleManagement Forum (TMF) Multi-Technology Network Management (MTNM) specifications TMF 513, TMF 608, TMF 814 and TMF Multi-Technology Operations Systems Interface (MTOSI) TMF 517 and TMF 608).

Detailed equipment specifications are considered important by equipment manufacturers as they provide a detailed formal specification of what components a piece of transport equipment should contain, how those components should interact and how the piece of equipment itself should behave. Management information models are considered important by network operators (and management standardization organizations such as the TeleManagement Forum (TMF)) because they formally define and describe the reference points that the operator's OSS system must interact with in order to manage a piece of transport equipment (and ultimately the transport network itself).

Figure I.3 shows a single SDH path layer network (e.g., VC-4) at the most abstract level (the highest level of partitioning), i.e., it is depicted as a single subnetwork bounded by its access points. This SDH path layer network is used to support various "network stacks". Note that the SDH network is itself decomposed into multiple layer networks (e.g., VC-4, Multiplex section, Regenerator section, Wavelength, etc. down to the duct level). Figure I.3 illustrates that [ITU-T G.805]/[ITU-T G.809] allows us to describe a single server layer network that may support multiple (different) client layer networks (it is not possible to do the same with the OSI BRM because the OSI BRM assumes a single 'network stack'). Figure I.3 also shows how [ITU-T G.805]/[ITU-T G.809] supports recursion (through client/server relationships) and demonstrates that layer networks are not always stacked according to the rigid model provided by

X.200/OSI BRM. A wide variety of network stacks may be modelled using [ITU-T G.805]/[ITU-T G.809]. This is illustrated in Figure I.3.

Working from left to right, the network stacks shown in Figure I.3 are:

- PPP over L2TP over IP over MPLS over Ethernet over SDH;
- IP over Ethernet over MPLS over Ethernet over SDH;
- IP over SDH;
- ATM over MPLS over SDH.

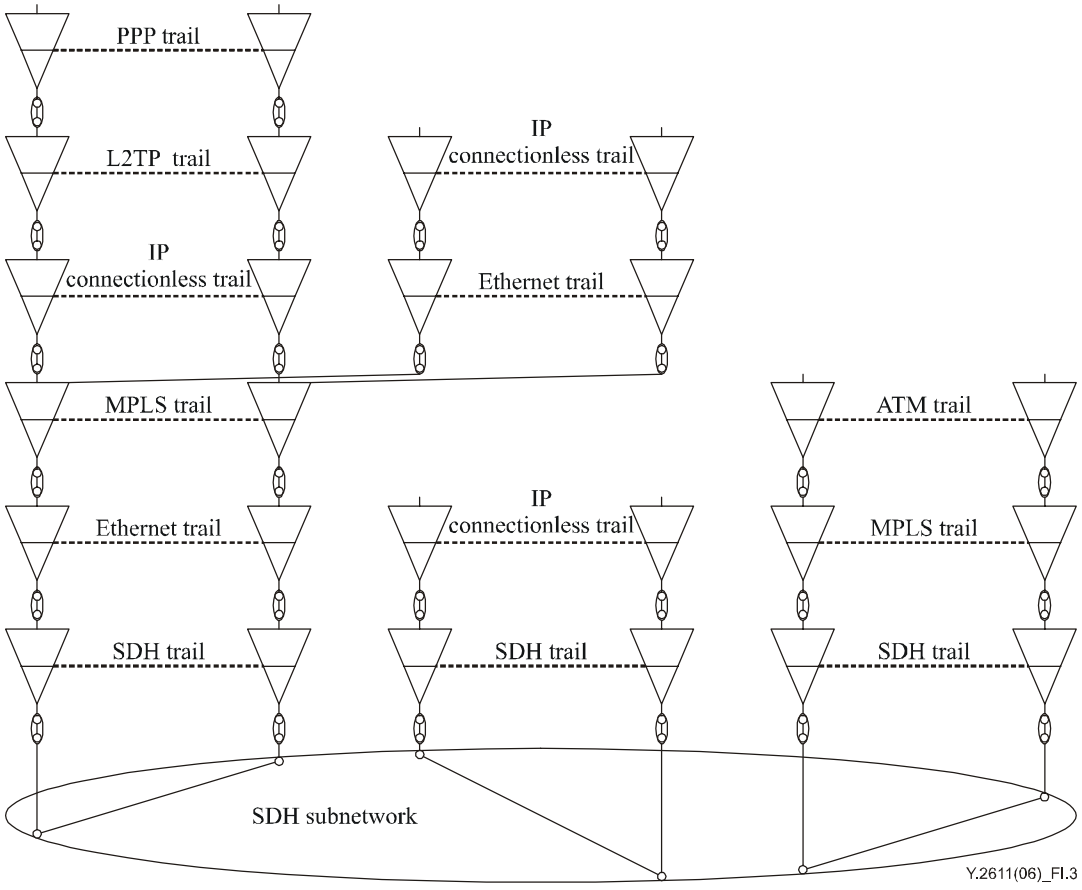

Figure I.3 – Illustration that G.805/G.809 allows to describe a server layer network that may support multiple (different) client layer networks including client/server recursion

Bibliography

[b-ITU-T G.783] ITU-T Recommendation G.783 (2006), *Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks*.

[b-ITU-T I.732] ITU-T Recommendation I.732 (2000), *Functional characteristics of ATM equipment*.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
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	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
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 and next-generation networks**
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2007