ETSI TR 101 803-10 V1.1.1 (2004-07)

Technical Report

Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN); Dynamic synchronous Transfer Mode (DTM); Part 10: Routeing and switching of IP flows over DTM

Reference

DTR/TISPAN-03003-DTM

Keywords

addressing, DTM, IP, switching

ETSI

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

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

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

Important notice

Individual copies of the present document can be downloaded from:
http://www.etsi.org

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
http://portal.etsi.org/tb/status/status.asp

If you find errors in the present document, send your comment to:
editor@etsi.org

Copyright Notification

Contents

Intellectual Property Rights

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

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

Foreword

This Technical Report (TR) has been produced by ETSI Technical Committee Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN).

The present document is part 10 of a multi-part deliverable. Full details of the entire series can be found in part 1 [5].

Introduction

Dynamic synchronous Transfer Mode (DTM) is a time division multiplex and a circuit-switched network technique that combines switching and transport. The recommendation specifying the DTM system and protocols are divided into 13 parts.

This part (Part 10) describes the method by which IP (as defined within IETF) traffic is carried over DTM. The topics of the other parts are as follows:

Part 1 introduces DTM and describes the service over a unidirectional data channel.

Part 2 includes system aspects that are mandatory or optional for nodes from different vendors to interoperate. These system aspects are addressing, routing, synchronization and channel management. The interworking granularity should be at node level, such that nodes from different vendors can interoperate with regard to well-defined functions.

Part 3 specifies the physical layer protocol for 8b/10b encoding based physical links.

Part 4 specifies the physical layer protocol for SDH VC4 container based physical links.

The transport of various tributary signals is specified for PDH **(Part 5)**, SDH **(Part 6)**, Ethernet **(Part 7)**, Frame Relay **(Part 8)**, ATM **(Part 9)**, IP **(Part 10),** Mapping of MPLS over DTM (**Part 11**), and video streaming **(Part 12)**. Note that DTM can either run over SDH or carry it as a tributary.

Finally, management aspects are standardized in **Part 13.**

1 Scope

The present document is a technical report keeping the state of the early development done on the transport of IP over DTM. It gives insight into how things can be done, but is not complete enough to become a standard, so the status is reported so that further work can make use of this knowledge.

The present document describes how IP traffic can be carried over a DTM network. Specifically, the mapping between IP and DTM is described and the system environment in which it operates. It further describes two services to carry IP traffic over a DTM, topology based forwarding and flow based forwarding.

The first service is based on hop-by-hop IP forwarding where the DTM transport is used as a flexible transport technology. The connectivity between the IP/DTM interworking functions is established (where IP routing functions are connected to DTM networks) by network management. This achieves an IP network overlaid on the DTM infrastructure. The IP overlay network is semi permanent as is usually the case for transport networks. For this service, the present document covers the adaptation of the IP traffic on to the connections and mechanisms for establishing the connections. User IP traffic will then be routed over the overlay IP network by the IP routers.

In the second service, the IP/DTM interworking function sets up a specific connection to other IP/DTM interworking functions "on-demand" across a DTM network to provide a dedicated connection for streaming IP traffic across the DTM network. These are referred to as dynamic channels. This service provides the capability to ensure, for example, traffic isolation, low latency and low packet loss for a specific IP flow.

The standard describes the adaptation of IP traffic on to the channels, channel setup and connection modification. The present document currently only defines the mechanisms for setting up channels with one sender and one receiver.

2 References

For the purposes of this Technical Report (TR) the following references apply:

[1] IETF RFC 791: "Internet Protocol".

[2] IETF RFC 2328: "OSPF Version 2", J. Moy, April 1998.

[3] IETF RFC 2460: "Internet Protocol, Version 6 (IPv6) Specification".

[4] IETF RFC 2543: "SIP Session Initiation Protocol".

[5] ETSI ES 201 803-1: "Dynamic synchronous Transfer Mode (DTM); Part 1: System Description".

[6] ETSI ES 201 803-7: "Dynamic synchronous Transfer Mode (DTM); Part 7: Ethernet over DTM Mapping".

[7] IETF RFC 2205: "Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification".

[8] IETF RFC 261: "PPP over SONET/SDH", A. Malis, W. Simpson. June 1999.

3 Definitions and abbreviations

3.1 Definitions

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

access node: a node that supports an external network interface, contains an interworking function for an external network and uses the DTM service

channel: set of slots allocated from one source Access node to one or more destination access nodes in a network

NOTE: The source and destination nodes can be the same, where the channel is internal to the node.

control channel: channel used for control signalling

data channel: channel used for transport of user data

domain: DTM network or part of a network that is managed by a particular commercial or administrative entity (carrier/operator)

DTM network client: client of the DTM network service, i.e. mapping functions

DTM network: set of connected DTM nodes

NOTE: A DTM network may be single-domain, or multi-domain.

dynamic channel: channel set up between two access nodes for specific purposes

NOTE: When used to transport IP traffic it is not visible to the IP routing protocols as being additional connectivity.

node: network element that contains DTM functions

node address: DTM network layer address of a node

3.2 Abbreviations

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

ARP	Address Resolution Protocol
CMI	Channel Multiplexing Indentifier
CMMI	Channel Management Module Interface
DCAP-1	DTM Channel Adaption Protocol 1
DCP	DTM Channel Protocol
DST	DTM Service Type
DSTI	DTM Service Type Instance
DTM	Dynamic synchronous Transfer Mode
IP	Internet Protocol
IPOD	IP Over DTM
LSA	Link Signalling Activity
PDU	Protocol Data Unit
RFC	Request For Comment (IETF document)
RSVP	Resource ReSerVation Protocol
RIP	Routing Information Protocol
POS	Packet Over Sonet (/SDH)
VPN	Virtual Private Network
NHRP	Next Hop Reservation Protocol

4 Service overview

Two different methods for transporting IP traffic across a DTM network are specified in the present document.

The DTM network may be used to provide a transmission layer between adjacent IP routers which provide a service based on traditional hop-by-hop forwarding. This method is normally used for best effort IP traffic. In this case, the channels are set up by management actions and provide transmission links available to the IP routers and connectivity that is visible to the IP routing protocols. The IP router functionality may be included in the same equipment as the DTM Switch, but such IP functionality is not the subject of the present document.

Using the second method, the DTM network may provide unidirectional channels across the network between ingress and egress IP routers/hosts on demand. This mechanism provides dedicated channels (dynamic channels) across the DTM network and can be used for IP flows between IP/DTM interworking functions (e.g. DTM aware IP routers). The IP flows could contain, for example, real-time traffic, Video, VPN or Database transfers.

The DTM address and DSTI allocated to the IPOD interface and used by the DTM layer is provided to the IP layer across the IPOD interface such that the IP layer may resolve the addresses using one of its protocols such as OSPF opaque LSA, NHRP, RSVP, etc.

The most significant difference between IP and DTM is that IP is connectionless, while DTM is a connection-oriented technology. IP maintains no state information; every packet has complete address information identifying the destination and each packet transported is handled individually in the process of forwarding packets from sender to receiver. DTM, on the other hand, is connection-oriented; meaning that data is transported using an established connection. During the establishment of the connection sufficient state information is stored in the switches along the path from source to destination to forward the data without the need for each data item to carry information specifying the destination.

As a result of the connection-oriented properties of DTM, channels need to be set up between IP/DTM Interworking functions before transporting the IP traffic over the DTM network.

5 System overview

IP over DTM (IPOD) is designed for building large IP- networks overlaid on DTM networks. IPOD is designed to efficiently provide a best-effort service for bursty traffic, as well as transport of QoS demanding traffic. IPOD supports both types of traffic in the same infrastructure by treating them differently. Normally, best-effort traffic is handled with topology-based forwarding and traffic that requires a high QoS is handled with flow based forwarding.

IPOD includes the following:

- A base set of channels to provide the connectivity for a hop-by-hop IP overlay network with the possibility to aggregate traffic as in normal router based IP networks. These channels are established at startup and are semi permanent.
- An efficient IP topology. IPOD utilizes the multi channel characteristics of DTM to have many adjacent routers meaning that there can be channels to many adjacent routers using the same physical DTM link. This means that many IP routers can be reached directly thus appearing to the IP network as adjacent which reduces the number of router hops in the IP network.
- Efficient resource usage in the network. By using the capabilities to adjust the capacity of already established channels, IPOD can adjust the size of channels to fit them to the amount of data currently transported on them. This adjustment can be achieved manually by network management system or automatically by using a channel manager.
- The capability to add channels dynamically for transport of flows from ingress to egress points of the DTM network without passing through any intermediate routers. This can be used either for guaranteed transport (i.e. to offer a high quality service), or because large amounts of data are to be transported between designated points. The dynamic channels are established and removed on demand using the control aspects of the IPOD interface.

5.1 Topology based forwarding

The hop-by-hop forwarding technique traditionally used in IP-based networks has proved very effective at transporting best-effort traffic where traffic from many sources can be aggregated to efficiently use network resources (i.e. links and switches). With DTM, a flexible transport network is provided to build a hop-by-hop IP network that efficiently utilizes the underlying infrastructure (e.g. fiber). The multi-channel DTM service can be used to establish direct paths between several routers with only a single physical interface. The possibility to change the bitrate of established channels provides the means for efficient use of resources to adjust to changing traffic patterns.

NOTE: The DTM layer could make the "cost" of the router to router links available for use by OSPF (or whatever IP routing is used). However, this is considered not to be necessary and could lead to other problems, and is therefore not included.

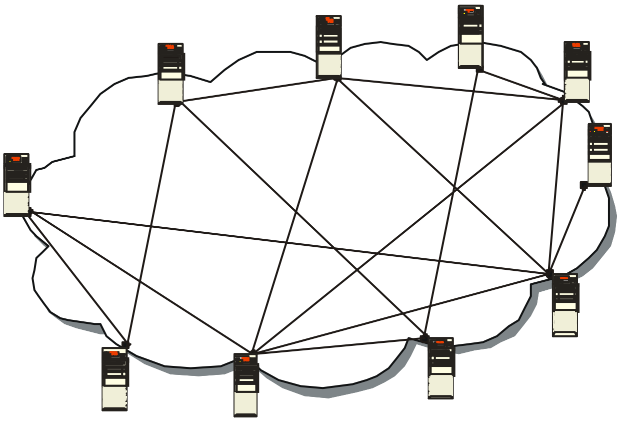

Figure 1: IPOD network

Figure 1 shows an example of an IPOD network where a number of routers and IP hosts are connected in an arbitrary structure. Base channels are always allocated in pairs to establish bi-directional connectivity between routers that are directly connected via base channels. All the channels shown are semi permanently set up .The capacity of each channel can be changed during its existence either via a management interfaces to IPOD or as a result of an automatic mechanism requesting a change using the control aspects of the IPOD interface. This information is then conveyed from IPOD to the channel management. The connectivity of the IPOD network is discovered by the IP routing protocols (such as RIP or OSPF) that traditional IP networks use to exchange information with peer routers.

For small networks with few routers, a complete mesh of channels can be established. When the network grows in number of routers, partial meshes can be configured. A mesh or partial mesh is considered an IP subnet to the IP layer. For even larger networks, it is inevitable to have more router hops through the network for two reasons. Firstly, the number of channels in a large network will be large in a big network. Secondly, for best effort type of traffic in a large network, it is desirable to aggregate the traffic to more efficiently utilize the capacity of the links. This is accomplished by establishing a number of partial meshes and interconnecting the partial meshes with a number of routers. It is possible to have several interconnection points between meshes to avoid a single point where a lot of traffic is transported.

The use of indirect DTM channels between IP routers gives freedom from the normal need for the large number of physical links between the routers in large networks.

5.2 Flow based forwarding

DTM has inherent support for end-to-end transport with complete traffic isolation resulting in no data loss and predictable delay. This support is used by IPOD to establish direct channels across the DTM network to carry IP traffic, thus short cutting intermediate routers making the IP transport over the network one hop regardless of the size of the network. This reduces the load on routers in the network and provides transport where traffic is not aggregated on the IP layer.

The flow-based forwarding (i.e. the decisions on which packets are forwarded on the dynamic channel and how the IP layer decides to initiate the establishment of dynamic channels) is outside of this IPOD specification. The initiation of the dynamic channels can for example be from IP signalling protocols such as RSVP or by configuration using the network management system. The flow-based forwarding can be based on flow classification that determines which traffic is transported on the dynamic channel.

When the IP layer decides to establish a dynamic channel, the IP signalling protocol is used to signal this using the control aspects of the IPOD interface. The IP signalling protocol messages are transported through the IPOD network and only the ingress and egress nodes need to process the RSVP messages and keep information on the RSVP state regarding the dynamic channel (both on the IP and DTM layer). Intermediate nodes only forward the IP signalling protocol messages.

Dynamic channels are always established as unidirectional channels. If a bi-directional dynamic channel is needed, for example for a videoconference that requires guaranteed transport in both directions, a separate dynamic channel should be used from each source.

5.3 IPOD interfaces

An IPOD interface represents an IP interface to the IPover DTM service. The IPOD interface provides access to a virtual multicast and broadcast capable IP subnet being specified as an IPOD segment. Configured base connections provide the basic connectivity within the IPOD segment while dynamic channels can be setup for identified flows. In contrast to the base channels, the dynamic channels are not bound to the IPOD segment (i.e. the sub network). Rather, the dynamic channels can terminate either within or outside the IPOD segment. The forwarding of IP packets on the dynamic channels can be done based on matches on destination IP address, source IP address, port numbers, TOS fields, DS code points in DiffServ, etc. The forwarding of packets to dynamic channels violates the best prefix match, normally used in IP networks, to allow for special treatment of traffic with requirements on, for example, QoS. This is achieved by using a forwarding table function, which is not only constructed by routing protocol information but also information from the network management system and/or from signalling protocols (such as RSVP) at the IP layer.

Each node participating in an IPOD segment must have exactly one IPOD interface connected to that IPOD segment. An IPOD interface is a logical entity within a node. An IPOD interface is not bound to a single DTM physical interface. There can be several IPOD interfaces associated with a single DTM physical interface and node.

An IPOD interface is identified in the DTM network by the DTM address of the node in which it resides and a DSTI number that is unique for this IPOD interface in this node. To establish a channel to an IPOD interface, the sender must know the DTM address of the node, the DST for IPOD (which is a well-known number) and the DSTI for the IPOD interface.

An IPOD interface must also have an IP address that can be used to identify the IPOD interface in the IP network.

The dynamic channels of IPOD use the IPOD interfaces for the base connections to originate and terminate the dynamic channels. However, the dynamic channels are not visible to the overlay network formed by base channel.

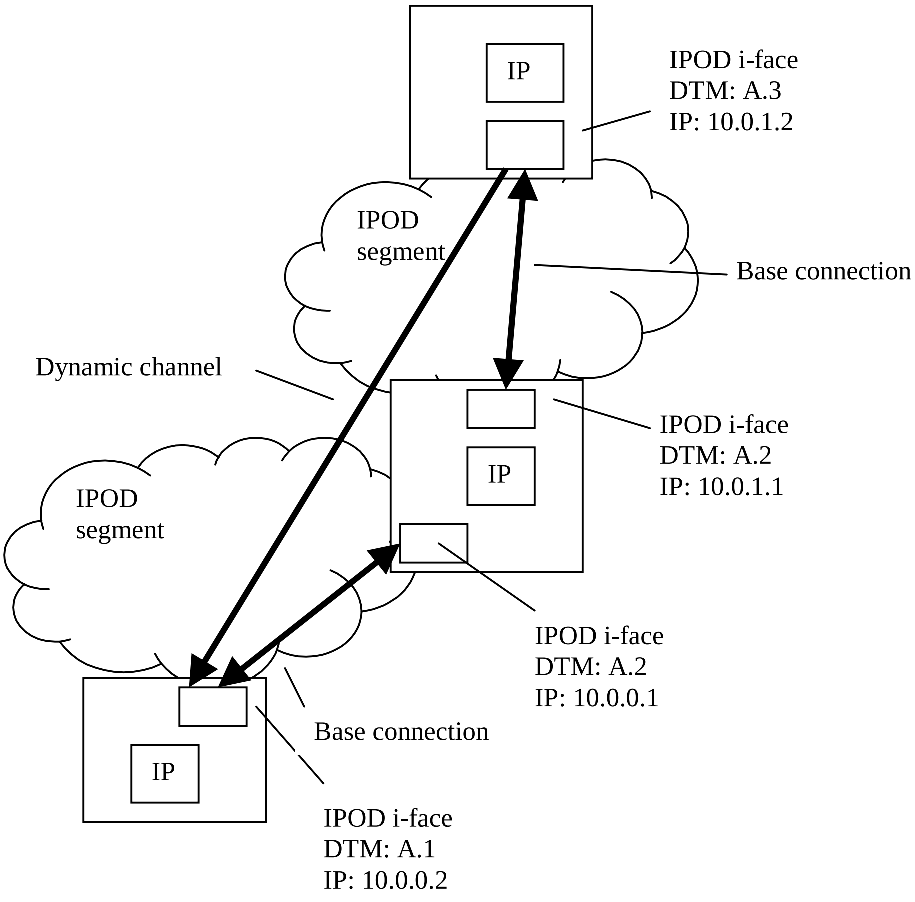

Figure 1a: IPOD interfaces and sub networks

In figure 1a, there are two IPOD segments implemented in the same DTM network. In the example, a dynamic channel is established from DTM node A.3 to DTM node A.1. This dynamic channel originates from the IPOD interface with IP address 10.0.1.2 and terminates on the IPOD interface with IP address 10.0.0.2 thus using the same IPOD interfaces that are used for the IPOD segments.

NOTE 1: This clause should include identification of both traffic service interfaces and control service interfaces with a diagram showing the relationships.

NOTE 2: The two virtual interfaces, namely the one used for best effort and the one used for streamed traffic should be identified here.

NOTE 3: More work is needed from here on.

5.4 IPOD segments

An IPOD segment consists of a number of IPOD interfaces and thus forms an IP subnet emulating a broadcast medium. Each IPOD interface is located in a physical node. There can only be one IPOD interface for a specific IPOD segment in a single node, but one physical node can have several IPOD interfaces connected to different IPOD segments. IP routing protocols that support point-to multipoint data links should be configured to use that facility over an IPOD segment.

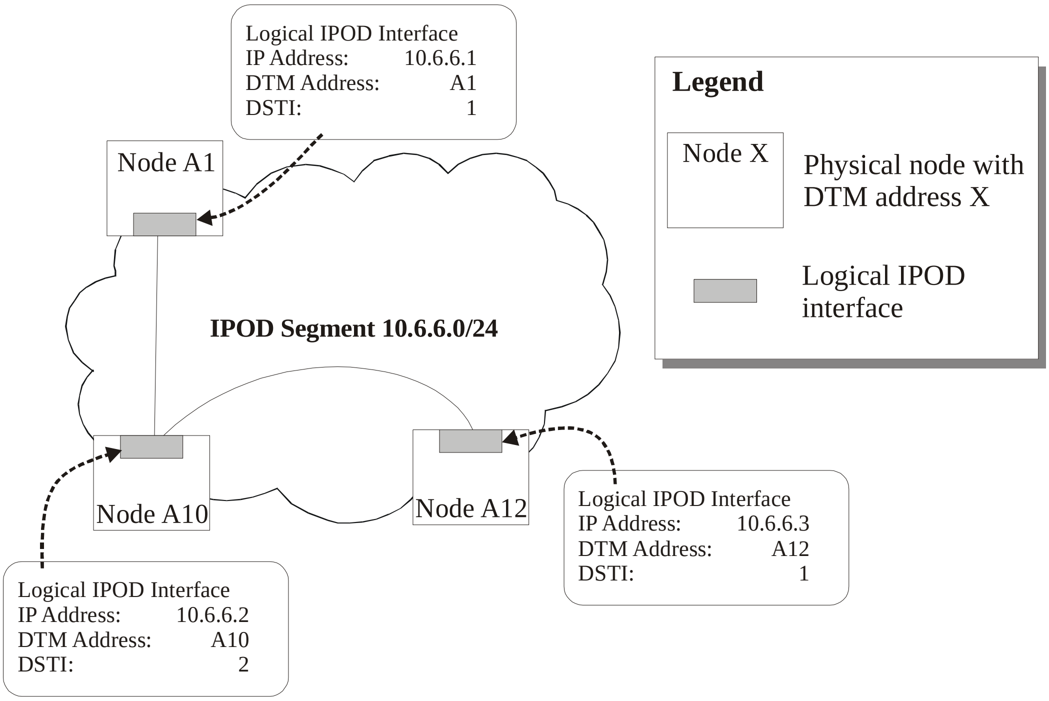

Figure 2: Example of an IPOD segment with three participating nodes

It is possible to establish several IPOD segments in a DTM network. This can be done if:

- there are several different IP access nodes connected to the same DTM network and IP communication between the customers are only allowed through firewalls if it is allowed at all;
- the IPOD network contains too many routers to be configured as a single routing domain. Then the IPOD network can be broken down into several different IPOD segments configured as separate routing domains.

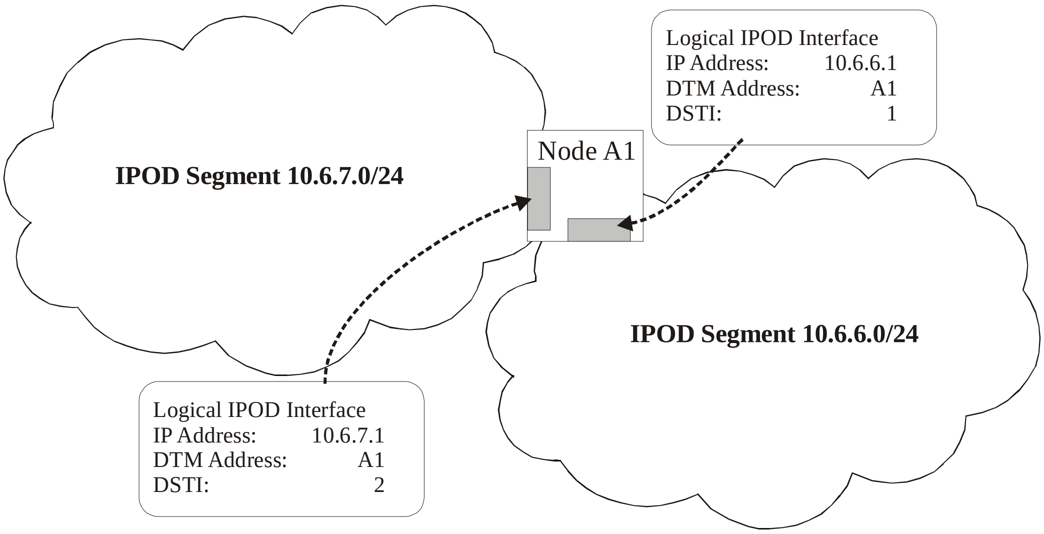

Figure 3: One physical node with one DTM interface and two logical IPOD interfaces attached to two different IPOD segments

Figure 3 shows one DTM node, A1, which has a single DTM interface. Via this interface, it is connected to two different IPOD segments, with one logical IPOD interface per segment. The node A1 can for example act as a router between these two IPOD segments or as a firewall between these two IPOD segments. Note that both IPOD segments are overlaid on the same DTM network.

5.5 Channels in IPOD

All packets in IPOD are transmitted on channels using DCAP-1 encapsulation with its CRC-32 checking of the data content. The IPOD channels are used for sending both IP datagrams and IPOD control messages. The CMI is used to demultiplex the IP datagrams and IPOD control messages at the destination. The CMI value of the data item being transported is specified in the service interface of DCAP-1. The CMI information is then transported by DCAP-1 in its header.

It is allowed in IPOD to establish several channels between the same two IPOD interfaces. These channels can be used for example for traffic with different priority levels.

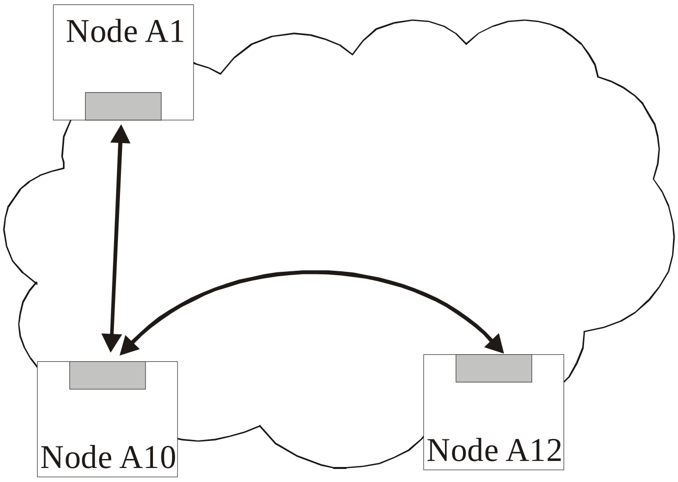

Figure 4: Three IPOD interfaces connected with two base channels

As described earlier there are two types of channels in IPOD, base channels and dynamic channels. The base channels are always established in pairs to allow for bi-directional communication as in figure 4. They are used for building the hop-by-hop routing network between all the routers connected to the IPOD segment.

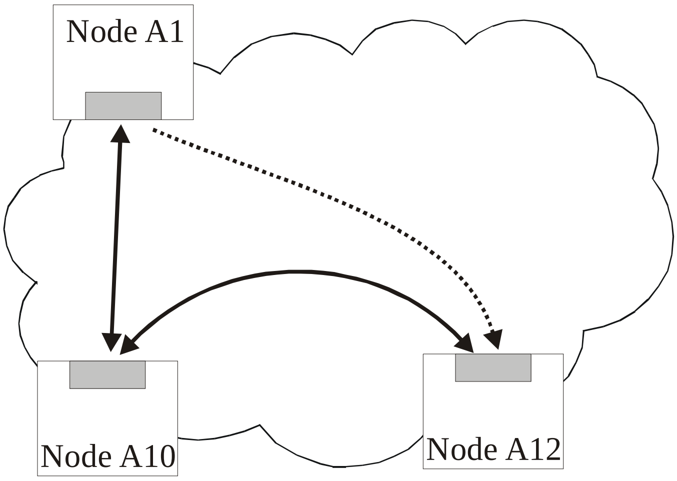

Figure 5: Addition of a dynamic channel between two of the IPOD interfaces

The dynamic channels are established to forward a specific IP flow (or some other well-defined set of packets) directly from source to destination or from ingress to egress router. In figure 5, a dynamic channel has been established from the IPOD interface in A1 to the IPOD interface in A12. The dynamic channel is dedicated to a single flow, described by source and destination IP addresses and possibly UDP/TCP port numbers.

The establishment of a dynamic channel is initiated at the transmitting IPOD interface; the receiving IPOD interface simply accepts the channel provided it has enough resources to do that.

Dynamic channels are always unidirectional, i.e. they are not established in bidirectional pairs. If bi-directional communication is needed, two separate dynamic channel channels should be established.

Only packets that meet the restrictions for an established dynamic channel can be forwarded on the dynamic channel.

NOTE: The two virtual interfaces, namely the one used for best effort and the one used for streamed traffic should be referenced here. As described earlier, it is actually the same interface.

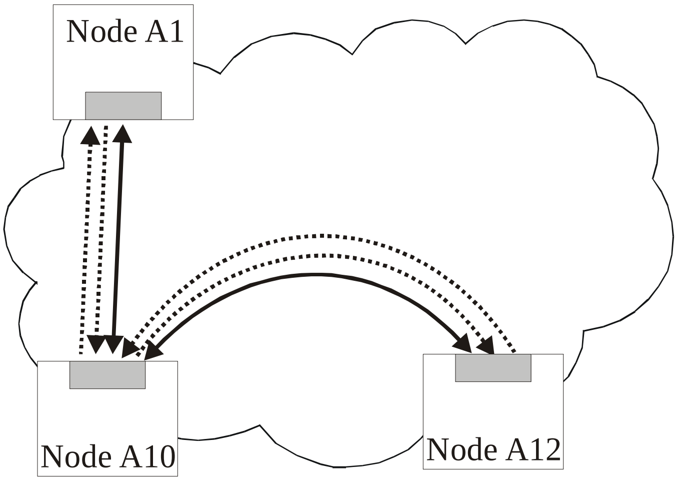

Figure 6: Dynamic channels in parallel with the base channels

The dynamic channels in figure 6 can for example have a restriction that only packets with certain TCP/UDP port numbers can be sent on them, all other packets are sent on the base connections. This means that we can prioritize certain types of traffic.

Dynamic channels can also be established in parallel with base channels. When a flow is requested with endpoints in the DTM network that are already connected with a base channel. Although we already have a direct channel between the two nodes, a new channel is established as a result of the request. This new channel separates the new flow from the rest of the traffic and provides the QoS guarantees needed by the new flow.

5.6 Base channel capacity

NOTE: This clause should be moved to another document since it is common to IPOD/DLE/DLT and it should also be re-written to allow for several different bitrate specification schemes. Agreed, will move when we have identified another location. We must have something concerning BW mgmt here since IPOD instructs the channel management block to set the size of the channel.

Base channels should always be established, but the IPOD interface is allowed to change the bitrate of an outgoing base channel to match the actual flow of traffic over the channel. There are a number of parameters that govern the bitrate allocated to an outgoing base channel:

- STARTBR The initial bitrate for base channel x.
- MINBR The minimum bitrate for base channel x.
- MAXBR The maximum bitrate for base channel x.
- MINACCEPTABLEBR The minimum amount of bitrate that can be accepted for the base channel.
- All bitrate parameters are expressed in number of slots.

MINACCEPTABLEBR must always be less than or equal to MINBR. STARTBR must be greater than or equal to MINBR and less than or equal to MAXBR.

When an IPOD-client tries to establish a basechannel, it should try to get the bitrate specified by STARTBR. If this is not possible, it should try to get as much as possible, but at least MINACCEPTABLEBR. If it cannot reserve MINACCEPTABLEBR, it should not establish any channel and signal that it was unable to establish the channel for the base connection.

The IPOD client should continuously monitor the amount of traffic sent on each basechannel and adjust the bitrate of the basechannel accordingly. This adjustment should be done within the limits specified by MINBR and MAXBR. It might happen that the IPOD interface is unable to allocate MINBR for the channel because of lack of capacity. This is acceptable as long as the bitrate of the basechannel is more than or equal to MINACCEPTABLEBR. If the bitrate is less than MINACCEPTABLEBR, the basechannel should be closed and an alarm raised.

5.7 Interaction with routing protocols

The IPOD base channel topology can be implemented using partly connected meshes that add requirements on the IP routing protocol. To obtain simple configuration and IP routing stability, the IP routing protocol should support multipoint interfaces as in for example OSPF.

5.7.1 OSPF

The base connection is used by OSPF RFC 2328 [2] to build a topology map of the network. The IPOD segment corresponds to an OSPF routing domain.

As defined in RFC 2328 [2] clause 12.4.1.4, an IP interface should be identified in an OSPF router by its IP address. Each base connection from the OSPF router should be identified by the IP address of the destination IPOD interface. It should not be necessary to add IP subnet information per base channel.

Dynamic channels should not be used by OSPF. The dynamic channels should only be used to forward the packets that they were intended for. They should therefore not be taken into account when calculating the shortest path between nodes, only the base channels should be used for this.

6 Service Interface

This is structured as in the Ethernet over DTM standard in terms of Traffic interface and Control interface.

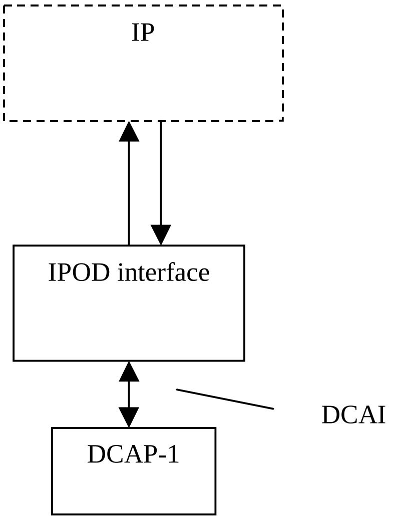

Figure 7: Transport view of IPOD

To transport the IP packets over DTM, the Ethernet/DTM Interworking function uses the DCAI service interface of DCAP-1.

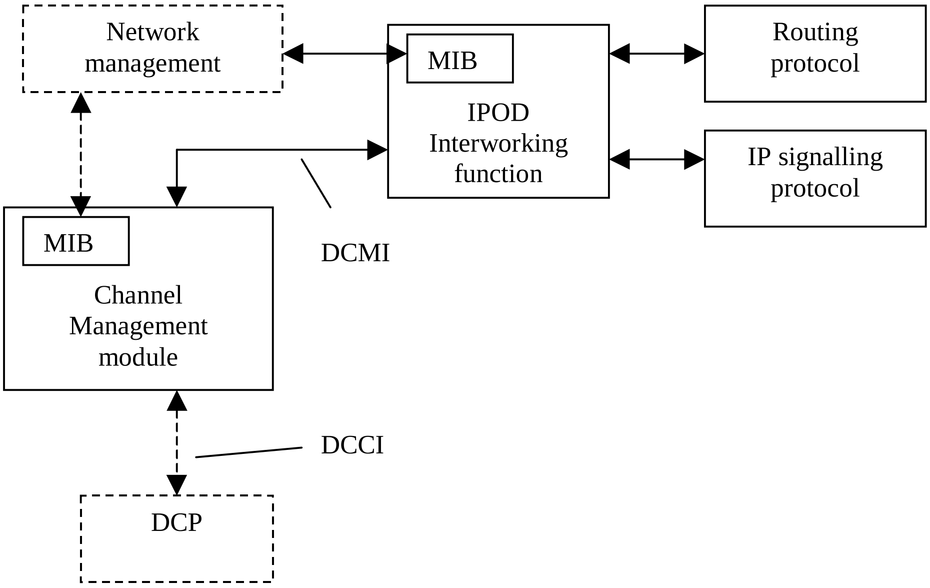

Figure 8: Control view of Ethernet Interworking function

For the control of the IPOD Interworking function there are interaction with the Network management system and the Channel Management module. The Channel Management module handles the set up, removal and capacity modifications of channel. The interface where the IPOD interworking function instructs channel operations is called Channel Management Module Interface (CMMI) [7]. The network management interfaces towards the Interworking function is used to provisioning the services apart from the normal management functions such as status gathering and configuration. There is also interaction with the IP routing protocol and the IP signalling protocols to construct forwarding tables and to establish dynamic channels.

6.1 Service provided

The IPOD provide a management interface to the upper layers. The details of this interface are not part of the present document.

6.2 Service required

The Ethernet channels use the DTM Channel Management Interface (DCMI) to establish channel, modify capacity of channel and remove channel. The size of the channels is also defined through the DCMI.

7 Detailed protocol description

This clause starts with a description of the entities involved in IPOD: interfaces and incoming and outgoing channels. It continues to describe the actions performed by an IPOD interface.

Each node participating in an IPOD segment has one IPOD interface for that IPOD segment. Each IPOD interface is associated with a number of incoming and outgoing channels.

7.1 IPOD interface parameters

The following parameters are associated with an IPOD interface.

Table 1: IPOD interface parameters

Parameter	Description
IP address	The IP address of the IPOD interface. This should be unique, both within the node and within the IPOD segment that this interface is connected to.
Segment IP address	The IPOD segment has an IP address prefix associated with it. All IPOD interfaces belonging to the segment must be configured with an IP address with this prefix. The prefix is described by the Segment IP address and netmask.
Segment IP netmask	See clause 8.1.2
DTM address	This is the DTM address of the node that this IPOD interface is located in.
DSTI	The DSTI where the IPOD interface should listen for connection attempts. The DSTI in combination with the DST value for IPOD should be unique within the node. To establish a channel to this IPOD interface, the sender needs to know the DTM address and DSTI of the IPOD interface as well as the DST value for IPOD (which is a well-known number). These three parameters uniquely identify the IPOD interface in the DTM network.

7.2 Incoming channel properties

Each incoming channel to an IPOD interface is associated with the following parameters.

- Source DTM address.
- The DTM address of the sending IPOD interface.
- Source DSTI.
- The DSTI of the sending IPOD interface.

7.2.1 Channel type

This parameter identifies the type of the outgoing channel. It can have the following values:

- CONFIGURED_BASE: This channel is part of a base connection that was established because of a local configuration.
- AUTOMATIC_BASE: This channel is part of a base connection that was established because of a remote configuration of a base channel, which was received from a peer. It has not been configured in this node.
- DYNAMIC CHANNE: This channel is a dynamic channel.

7.2.2 Source IP address

The IP address of the sending IPOD interface.

7.2.3 Filter specifications

In RSVP (RFC 2205)[7] this is called "filter specification" or "filter spec" for short. I propose we use that terminology over "flow restriction" which in my mind has a bad ring to it.

A filter specification can be associated with a dynamic channel. The filter specification has the following properties:

- Source IP address;
- Destination IP address;
- IP protocol type (UDP or TCP);
- Source port number;
- Destination port number.

The properties of the filter specification describe the packets that can arrive on the channel. If there are any filter specifications associated with a channel, only packets matching one of these restrictions can be sent on the channel. See the description of the Flow Restriction Extension for an explanation of the content of the properties. This information can be used to speed up the forwarding decision at the receiver in some implementations of IPOD and also to allow the operator to know what each incoming dynamic channel is used for.

NOTE: In RSVP (RFC 2205) [7] the terms "classifier" and "filter spec" is being used. We should use them. The "filter spec" tells the "classifier" how to identify received packets as belonging to a "flow". The "packet scheduler" will then schedule the identified "flows" according to their corresponding "flow spec". In the IPOD case the scheduler will schedule the packets onto base and dynamic channels. The base channels will use the best effort scheduling while dynamic channels will be scheduled as defined in their "flow spec". It should be noted that we can (and should) allow for multiple flows to be scheduled onto the same dynamic channel. Allowing multiple flows per dynamic channel allows for some multiplexing where the individual streams as so small that they would waste capacity if setup as separate channels. The "flow spec" might be upgraded over time, and if so necessary a new dynamic channel may be requested or two dynamic channels may be merged. Allowing for this without enforcing it can provide a useful tool for implementers to provide optimization. The downside of allowing multiple flows on a channel is that this enforces a flow classifier at the receiver end. However, the flow classifier is only required when the flows take on a different path or to see different scheduling rules than previously.

7.3 Outgoing channel parameters

Each outgoing channel has the following parameters associated with it.

Table 2: Outgoing channel parameters

Parameter	Description
Destination DTM address	The DTM address of the destination IPOD interface.
Destination DSTI	The DSTI of the destination IPOD interface.
Channel type	This property identifies the type of the outgoing channel. It can have the following values: CONFIGURED_BASE: This channel is part of a base connection that was established because of a local configuration. AUTOMATIC_BASE: This channel is part of a base connection that was established because of a remote configuration of a base channel, which was received from a peer. It has not been configured in this node. DYNAMIC CHANNEL: This channel is a dynamic channel.
Maximum bitrate	The maximum bitrate requested for this channel.
Minimum bitrate	The minimum bitrate requested for this channel.
Current bitrate	The bitrate currently allocated for this channel.

7.4 IPOD interface operation

7.4.1 IPOD interface startup

When an IPOD interface is started, it should perform the following steps:

1) Start to accept channels on the configured DSTI.

2) Start to establish the channels for the base connection. The base connection are configured at one end of the bi-directional base connection, if configured in the node, the node starts to establish an originating channel and waits for the remote node to set up the corresponding channel to achieve a bidirectional pair of base channels. If the remote node initiates the set up of base channels, the node simply accepts the incoming base channel and establish a corresponding channel back to the remote node as locally configured and as required by the acceptance of base channels due to remote configuration.

3) As soon as at least one base connection has been established for bi-directional communication, the IPOD interface should register with the routing entity and declare itself available.

4) Continue to establish base channels and dynamic channels until all configured channels have been established.

7.4.2 IPOD interface shutdown

When an IPOD interface is shut down, it should perform the following steps:

1) Deregister with the routing entity to avoid having more packets sent to it.

2) Stop accepting new channels.

3) Close all incoming and outgoing channels.

7.4.3 Establishing an IPOD channel

To establish a channel from IPOD interface A to IPOD interface B, A must know the DTM address and destination DSTI of B. This information can be found either in the configuration of A or it can be obtained from the channel establishment message received from B if B has already established a channel to A. The establishment of channels is handled by DTM Channel Protocol through the service interface of the channel manager.

The following steps are performed in the channel establishment:

1) A initiates the creation of the channel by using the channels manager. The IP/DTM interworking function of A issues a DCMI_CONNECTION_ESTABLISH containing the DTM address and DSTI of B, the DST equal to the DST_IPOD and the DTM address and DSTI of A.

2) B receives a DCMI_ICHANNEL_RECEIVED for the channel from A, including the DTM address and DSTI of A. B accepts the channel by issuing a DCMI_ICHANNEL_ACCEPT back to A.

3) When A receives the DCMI_OCHANNEL_ESTABLISHED as a result of that B has accepted the channel, A sends an IPOD_REGISTER message on the channel. The IPOD_REGISTER message contains the IP address of A and possibly a Flow destination and/or a Flow priority parameter for the channel. The IPOD_REGISTER can also contain an Authenticate extension if necessary. More details on this message are given in clause 9.2.

4) B receives the IPOD_REGISTER message. It checks the authentification authentication contained in the Authenticate extension (if any) and that the supplied IP address falls within the netmask for the IPOD interface and decides if it can receive packets from A. If B decides to accept packets from A, it should create a new Incoming Channel Entry with the Source DTM address, DSTI and IP address as well as the Flow destination and Flow priority properties if these are supplied. Otherwise it should close the channel and raise an alarm for erroneous or illegal connection attempt.

If B receives IP packets on the channel before receiving the IPOD_REGISTER message (see note) it must discard the received IP packets and close the channel. It can then try to re-establish the channel again.

NOTE: This can happen for two reasons; The IPOD_Register message was lost because of a transmission failure (possible, but very rare) or the sender is a malicious user trying to circumvent the security policies in IPOD.

Error handling is covered in the following two clauses.

7.4.4 Establishing base connection

A base connection consists of exactly two channels between two IPOD interfaces, one in each direction. Both channels are established as described in clause 7.4.3. Base channels should always be created with the base connection bit set in the IPOD_REGISTER message to signal to the receiver that a base connection is requested to be established. This means that the receiver must establish a channel back to the sender if its configuration allows it to.

NOTE: For the pair of channels, we better use "base connection" than "base channel". I prefer that we always use "channel" to mean a unidirectional channel when talking about DTM functionality (obviously can other technologies being discussed have a different meaning of channel within their scope). One needs to defined a "base connection" to be valid when we have a bidirectional connectivity by a pair of matching base channels. The benefit of this is that we can more easily separate the discussion of unidirectional channel properties (such as when setting up or tearing down) with the perceived bidirectional properties.CB, I agree not to use channel. I think that it might be useful to use a complete different word for base channel: Maybe adjacent connection or something else?

Base connections establishment is performed based on configuration data.

When the channels are established, IPOD will register the channels as IPOD channel resulting in that if a channel fails to be established, the channel manager module will retry to establish the channel with a certain interval. The reestablishment interval between retries will increase exponentially up to a certain maximum level. This is done to prevent excessive signalling in the network. When the algoritn has reached the max interval, there will be an alarm generated for each unsuccessful retry.

7.4.5 Establishing dynamic channels

Dynamic channels are created either based on configuration of the IPOD interface or in response to a request from a function outside the IP/DTM interworking. This function can for example be an IP signalling protocol. Dynamic channels are always uni-directional and the sender creates them. The receiver must accept incoming dynamic channels if it has enough resources to do so and the IPOD_REGISTER message contains the correct authentication information.

If the establishment of a dynamic channel fails and the dynamic channel establishment was initiated by a request from a function outside the IP/DTM interworking function, the IPOD interface issues a service primitive to the function that requested the dynamic channel indicating the failure. The requesting function is then allowed to try to establish the dynamic channel again. Applications are not allowed to establish channels more frequently than what is done when establishing IPOD base connections.

7.4.6 Receiving a channel

When a new channel is received, the IPOD interface is issued a DMCI_ICHANNEL_RECEIVED primitive from the channel manager. The IPOD interface must accept the channel if it has enough resources to do so. It accepts the channel by issuing a DMCI_ICHANNEL_ACCEPT service primitive to the channel manager. When the channel has been established, the IPOD interface should wait for an IPOD Register message. If an IP packet arrives before the IPOD_REGISTER message the packet should be discarded and the receiving IPOD interface should close the channel.

When the IPOD Register message is received, the authentication extension should be checked first. If the authentication is correct, the Base Channel parameter should be checked to see if the channel is part of a base channel or if it is a dynamic channel.

7.4.6.1 Receiving a channel for a base connection

If the new channel is a base channel, the receiving IPOD interface should check if it already has an entry for this base channel as identified by the DTM address and DSTI of the sending IPOD interface. If so, the channel should be added as the incoming channel for that base connection. The base connection is now bi-directional and the routing entity should be notified of its existence.

If the IPOD interface has no entry for this base connection, it should check to see if IPOD_BASECHANNEL_ACCEPT_UNCONFIGURED is true. If it is true, and the IPOD_REGISTER message has the CONFIGURED bit set, it should add a new base connection and attach this channel to it. The creation of a new base connection also triggers the IPOD interface to try to establish a channel back to the sender of the current connection.

If the IPOD interface has no entry for this base connection and either if the Configured bit is not set or IPOD_BASECHANNEL_ACCEPT_UNCONFIGURED is false, the channel should be closed immediately and an alarm raised.

7.4.6.2 Receiving a dynamic channel

If the received channel is a dynamic channel, it should be accepted.

7.4.7 Remote teardown of an incoming channel for a base connection

If an incoming channel for a base connection is torn down by the sender or by a network failure, the receiver should signal to the routing entity that the base connection is no longer available.

If the base connection is a configured base connection, the receiver should raise an alarm and wait for the other node to re-establish the channel. When the channel has been re-established, the base connection should be registered with the routing entity again and the alarm cleared.

If the base connection does not have the configured bit set, the sender should tear down the outgoing channel belonging to the same base connection and remove the entry for the base connection.

7.4.8 Remote teardown of an outgoing channel for a base connection

If an outgoing channel of a base connection is torn down by the receiver or by a network failure, the sender should raise an alarm and signal to the routing entity that this base connection no longer exists.

If the base connection has the configured bit set (which means that it was configured locally), the sender should try to establish the channel again. When the channel has been re-established, the base connection should be registered with the routing entity again.

If the base connection is not configured in the node (meaning that the base channel exists because it was remotely configured), the sender should teardown the incoming channel belonging to the same base connection and remove the base connection altogether. It is then up to the remote node to try to establish the base connection again.

7.4.9 Remote teardown of an incoming dynamic channel

If an incoming dynamic channel is torn down, the receiving IPOD interface should do nothing except removing the entry for the channel.

7.4.10 Remote teardown of an outgoing dynamic channel

If an outgoing dynamic channel is torn down by a network failure or by the receiver, the transmitting IPOD interface should notify the entity that requested the dynamic channel. It is then up to that entity to request that the dynamic channel should be re-established.

7.4.11 Removal of a base channel because of re-configuration

If an IPOD interface is reconfigured so that a base connection is removed from the configuration, the corresponding channels should be closed and the base connection shall be removed.

Note that the base connection can be re-established by the IPOD interface at the other end if that IPOD interface still has a configuration entry for the base connection.

8 Messages

This clause describes the different control messages that can be sent on an IPOD channel.

All IPOD channels should be established to the DST for IPOD. IPOD control messages should be sent with CMI=1 and IP packets should be sent with CMI=2.

8.1 General message format

All IPOD messages have the following format:

```
   6         5         4
3210987654321098765432109...
```

VER	MT	Data…	Extension…

Figure 9: General IPOD message format

The VER field is 4 bits and identifies the version of IPOD used. The current version number is 0.

The IPOD_MSG_TYPE (MT in figure 9) is 4 bits. The following values for IPOD_MSG_TYPE have been defined.

Table 3: Defined values for IPOD_MSG_TYPE

IPOD_MSG_TYPE	Message Type
0	IPOD_REGISTER

Data is the data for the message. The length of the Data-section is determined by the IPOD_MSG_TYPE. See each individual message below for details.

Each message type can have zero or more extensions. Extensions have the following format.

EXT_ TYPE	IPOD_EXT_LEN	Data...

Figure 10: General format of IPOD message extensions

The IPOD_EXT_TYPE is 8 bits.

The IPOD_EXT_LEN is 16 bits. It contains the length of the extension Data-field in bytes.

If a receiving IPOD interface does not know how to interpret an incoming extension, the extension should be ignored and the message should be handled as if the extension did not exist.

Below is a description of each IPOD message. Grey fields are reserved for future use and should be set to zero on send and ignored on receive. All values should be stored in little-endian order.

8.2 IPOD_Register

This message is used when establishing a channel from one IPOD interface to another IPOD interface.

8.2.1 Normal message format

```
   6         5         4         3         2         1         0
3210987654321098765432109876543210987654321098765432109876543210
```

0	0	IPOD_REG_ FLAGS		SOURCE_IP_ADDRESS

Figure 11: IPOD_Register message format

7	6	5	4	3	2	1	0
B	C	U	U	U	U	U	U

Figure 12: IPOD_REGISTER_FLAGS

8.2.1.1 Source IP Address

The IP address of the IPOD interface that sent the message.

8.2.1.2 Base Connection

Bit 7 (marked with a B) of IPOD_REGISTER_FLAGS.

This bit indicates the type of channel being established.

0 dynamic channel.

1 base connection

8.2.1.3 Configured

Bit 6 (marked with a C) of IPOD_REGISTER_FLAGS.

If this is a channel of a base connection (as indicated by the Base Connection bit above), this parameter tells whether the sender has this dynamic channel configured or not.

0 Not configured (i.e.the channel was established because a channel of a base connection was received).

1 The sender is configured with this base connection.

If the Channel Type is dynamic channel, this parameter should be set to 0 on send and ignored on receive.

8.2.2 Flow Restriction Extension

The Flow Restriction Extension tells the receiver that only packets matching the specified flow specification will be sent on the channel.

```
     6         5         4         3         2         1         0
3210987654321098765432109876543210987654321098765432109876543210
```

0	14		PROTOCOL	PROTOCOL_SPECIFIC
DESTINATION_IP_ADDRESS				SOURCE_IP_ADDRESS

Figure 13: Flow Restriction Extension general format

```
     6         5         4         3         2         1         0
3210987654321098765432109876543210987654321098765432109876543210
```

0	14		PROTOCOL	SOURCE_PORT	DEST_PORT
DESTINATION_IP_ADDRESS				SOURCE_IP_ADDRESS	

Figure 14: Flow Restriction Extension format when PROTOCOL is TCP or UDP

8.2.2.1 Extension type

The Flow Restriction Extension has extension type 0.

8.2.2.2 Extension length

The Flow Restriction Extension is 14 bytes long, excluding the Extension Type and Extension Length fields.

8.2.2.3 Source IP address

If this is non-zero, it indicates that all packets arriving on this channel will have this source IP address. Other packets should be discarded? A zero value indicates that the sender gives no guarantees on which source IP addresses can arrive on this channel.

8.2.2.4 Destination IP address

If this is non-zero, it indicates that all packets arriving on this channel will have this destination IP address. A zero value indicates that the sender gives no guarantees on which destination IP addresses can arrive on this channel.

8.2.2.5 Protocol

If this is non-zero, it indicates that only packets with this protocol number in the IP header will arrive on this channel. A zero value indicates that the sender gives no guarantees on which protocols can arrive on this channel.

The value of the Protocol-field affects how the PROTOCOL_SPECIFIC field should be interpreted. If protocol is UDP or TCP, then the PROTOCOL_SPECIFIC field contains source and destination port-numbers as shown in figure 14. The contents of the PROTOCOL_SPECIFIC for other values of PROTOCOL have not yet been specified. An IPOD interface that receives a flow-restriction extension with a value of PROTOCOL for that it does not know how to interpret the PROTOCOL_SPECIFIC field may safely ignore the PROTOCOL_SPECIFIC field, since it only contains information to simplify packet decisions in the IPOD interface and provide feedback to the management application.

8.2.2.6 Source port

If this field is non-zero, it indicates that this channel will only be used to send packets with the indicated source port number. If the Source Port is zero, it indicates that packets with any source port number can arrive on this channel.

8.2.2.7 Destination port number

If this field is non-zero it indicates that this channel will only be used to send packets with the indicated destination port number. If the Destination Port is zero, it indicates that packets with any destination port number can arrive on this channel.

8.2.3 Priority restriction extension

The priority restriction extension tells the receiver that only packets with a priority equal to or higher than specified in the extension will be sent on the channel.

8.2.4 Authentication extensions

One or several different Authentication extensions will be added later. They are used to verify that the sender is a valid member of the IPOD segment.

9 IP encapsulation

All IP packets are transported by DCAP-1 using its CRC 32 bit check on its data content. The IP datagram is the service data unit which forms the data content of the DCAP-1 PDU.

NOTE: For multi-byte fields, the most significant byte is always transferred first, i.e.the most significant byte of the packet length can be found in bit-position 40 to 47.

Annex A: MIB parameters

A.1 Configuration parameters

This clause contains all parameters that should be possible to configure for an IPOD interface. Note however that not all these parameters need to be user configurable.

A.1.1 IPOD_IP_ADDRESS

The IP address of this IPOD interface.

A.1.2 IPOD_DSTI

The DSTI that this IPOD interface should listen to.

A.1.3 IPOD_BASECHANNEL_STARTBR_DEFAULT

The default initial bitrate of a basechannel. Measured in slots.

A.1.4 IPOD_BASECHANNEL_MINBR_DEFAULT

The default minimum bitrate of a basechannel. Measured in slots.

A.1.5 IPOD_BASECHANNEL_MAXBR_DEFAULT

The default maximum bitrate of a basechannel. Measured in slots.

A.1.6 IPOD_BASECHANNEL[x].IP_ADDRESS

The IP address of the receiving IPOD interface for basechannel x.

A.1.7 IPOD_BASECHANNEL[x].DTM_ADDRESS

The DTM address of the receiving IPOD interface for basechannel x.

A.1.8 IPOD_BASECHANNEL[x].DSTI

The DSTI of the receiving IPOD interface for basechannel x.

A.1.9 IPOD_BASECHANNEL[x].STARTBR

The initial bitrate for base channel x. Measured in slots.

A.1.10 IPOD_BASECHANNEL[x].MINBR

The minimum bitrate for base channel x. Measured in slots.

A.1.11 IPOD_BASECHANNEL[x].MAXBR

The maximum bitrate for base channel x. Measured in slots.

A.1.12 IPOD_INTERVAL_MIN

The algorithm to calculate the intervals between retries to establish base connections uses this parameter.

A.1.13 IPOD_INTERVAL_MAX

The algorithm to calculate the intervals between retries to establish base connections uses this parameter.

A.1.14 IPOD_BASECHANNEL_NORETURN_TIMEOUT

The number of milliseconds an IPOD interface must wait for a return channel from another IPOD interface.

A.1.15 IPOD_BASECHANNEL_ACCEPT_UNCONFIGURED

Whether this IPOD interface should accept basechannels that it has not been configured with. True means that it should accept them and establish a channel back, false means that it should reject them.

History

Document history		
V1.1.1	July 2004	Publication