

I n t e r n a t i o n a l T e l e c o m m u n i c a t i o n U n i o n

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

J.362

(11/2006)

SERIES J: CABLE NETWORKS AND TRANSMISSION
OF TELEVISION, SOUND PROGRAMME AND OTHER
MULTIMEDIA SIGNALS

IPCablecom

IPCablecom2 control point discovery

ITU-T Recommendation J.362



International
Telecommunication
Union

Documents
Document Sharing Hub

ITU-T Recommendation J.362

IPCablecom2 control point discovery

Summary

This Recommendation defines an IP-based protocol that can be used to discover a control point for a given IP address. The control point is the place where QoS operations, lawful intercept (LI) content tapping operations, or other operations may be performed.

Source

ITU-T Recommendation J.362 was approved on 29 November 2006 by ITU-T Study Group 9 (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/>.

© ITU 2007

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

CONTENTS

Page

1	Scope	1
2	References.....	1
3	Definitions	1
4	Abbreviations and acronyms	1
5	Conventions	2
6	Technical overview.....	2
6.1	Assumptions	3
7	Interface description	3
7.1	Network-based control point discovery application.....	4
7.2	NLS-TL header parameters	5
7.3	CPD request.....	5
7.4	CPD response	6
7.5	Error responses	7
8	Procedures	7
8.1	Requestor	7
8.2	Control point.....	8
9	Security considerations.....	9
	Annex A – Network Layer Signaling: Transport Layer	10
	Annex B – QoS requirements	44
	Appendix I – Open issues	45
	Bibliography.....	46

Copyright International Telecommunication Union/ITU Telecommunication Sector

ITU-T Recommendation J.362

IPCablecom2 control point discovery

1 Scope

This Recommendation defines an IP-based protocol that can be used to discover a control point for a given IP address. The control point is the place where QoS operations, lawful intercept (LI) content tapping operations, or other operations may be performed.

2 References

None.

3 Definitions

This Recommendation defines the following terms:

3.1 control point: Within the context of this Recommendation, control point refers to a point in the network that can be used to apply a function for a media flow that flows through that point. Functions described here are:

- QoS (IPCablecom multimedia [b-ITU-T J.179] or IPCablecom DQoS [b-ITU-T J.163]).
- Replication, encapsulation and transmission for the purposes of LI content tapping.

3.2 control point discovery: The act of discovering information (IP address, protocol) concerning a control point in order to allow a requestor to apply a specific controlling function.

3.3 requestor: The requestor in this context is the controller that wishes to control the control point and hence needs to discover the necessary information to do so.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

CMS	Call Management Server
CMTS	Cable Modem Termination System
COPS	Common Open Policy Service
CPD	Control Point Discovery
CR	Control Relationship
DF	Delivery Function
DNS	Domain Name Service
DQoS	Dynamic Quality of Service
ICE	Interactive Connectivity Establishment
IP	Internet Protocol
LI	Lawful Intercept
MIB	Management Information Base
NAT	Network Address Translation
NE	Network Element

NLS	Network Layer Signalling
NLS-TL	Network Layer Signalling Transport Layer
PS	Policy Server
QoS	Quality of Service
SDP	Session Description Protocol
STUN	Simple Traversal of UDP through NAT
TURN	Traversal Using Relay NAT
UDP	User Datagram Protocol

5 Conventions

Throughout this Recommendation, the words that are used to define the significance of particular requirements are capitalized. These words are:

"MUST"	This word means that the item is an absolute requirement of this Recommendation.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this Recommendation.
"SHOULD"	This word means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.
"SHOULD NOT"	This phrase means that there may exist valid reasons in particular circumstances when the listed behaviour is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behaviour described with this label.
"MAY"	This word means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

6 Technical overview

The general approach for control point discovery is illustrated by the reference model in Figure 6-1. A requestor that knows the IP address of the media endpoint sends a control point discovery (CPD) message toward that endpoint (reference point pkt-qos-2). A control point in the path between the requestor and the endpoint recognizes the CPD message and responds back with the IP address to use for the particular application. The requestor can then make the necessary request for that application.




Figure 6-1 – Control point discovery architecture

In addition to supplying the IP address to use, the CPD response indicates the protocol to use and can optionally supply which subnet the destination address of the media endpoint is contained within.

Some of the components that need to make use of CPD are as follows:

- IPCom call management server (CMS), in order to determine the IP address of the CMTS for making a DQoS request.
- IPCom multimedia policy server (PS), in order to determine the IP address of the CMTS for requesting IPCom multimedia [b-ITU-T J.179].
- IPCom delivery function (DF), in order to determine the IP address of the CMTS or aggregation device for performing a lawful intercept content tap of the media stream.

CPD is not limited to QoS and LI. It may be used for other applications as well.

The CMS, PS and DF in the above examples need to determine which CMTS or aggregation device (the IP addresses they need to control) will handle the media stream based on the information these components have (the IP address of the endpoint which was obtained via SDP).

Several approaches have been suggested, as follows:

- 1) Provisioning of subnet versus control point information within each device.
- 2) Provisioning within DNS [b-IETF RFC 4183].
- 3) Collecting subnet information over the COPS DQoS or IPCom multimedia interface.

Approaches 1 and 2 above are provisioning solutions and as such present operational difficulties. Approach 3 is an application-specific solution that uses a policy management protocol for an unintended purpose, and cannot be reused to address other requirements (e.g., LI). The approach described here is a generic approach that can be used for all three of the above applications and other applications as well.

6.1 Assumptions

The following basic assumptions are used in defining the interface specification in clause 7:

- For the majority of cases, the media endpoint is single-homed behind the control point (i.e., the media route will go through a specific CMTS, media gateway or aggregation device). For possible exceptions to this assumption, the requestor can be provisioned with alternates (i.e., when the requestor receives a CPD response with the IP address of one control point, it is able to determine the IP address of control points for alternative media paths for that media endpoint).
- For some applications (e.g., LI), it is important that the CPD message does not reach a media endpoint outside the provider's network. The assumption is that ACL mechanisms will be in place to drop CPD packets at edge devices that do not support the CPD protocol.

7 Interface description

CPD uses the network layer signalling protocol (NLS). As illustrated in Figure 7-1, NLS consists of an application layer that sits on top of an NLS transport layer (NLS-TL) protocol as defined in Annex A.




Figure 7-1 – NLS protocol

One such application is the network-based control point discovery application defined in this Recommendation.

7.1 Network-based control point discovery application

The application payload format within NLS consists of an NLS type, length, value (TLV) that consists of a 16-bit application ID followed by a payload that is opaque to the NLS transport layer. Per clause 14.2 of Annex A, the application ID for the control point discovery (CPD) application is "1".

The CPD message format is illustrated in Figure 7-2. For the network-based CPD application, a CPD messages consist of:

- A 4-bit version field. The 4-bit version field is set to "0" for the version of the CPD protocol described in this Recommendation.
- A 12-bit CPD message type. Only two CPD message types are defined:
 - CPD request: CPD message type = 1.
 - CPD response: CPD message type = 2.
- A 16-bit control relationship (CR) type. The CR TYPE identifies the type of control relationship. These values may be provisioned. The requestor **MUST** set the CR TYPE to one of the following values:
 - CR TYPE = 1: Lawful intercept content tap.
 - CR TYPE = 2: DQoS.
 - CR TYPE = 3: IPCablecom multimedia.
- A 16-bit control relationship ID (CR ID).
- A 32-bit transaction ID.
- CPD message contents for the particular CPD message type.




Figure 7-2 – CPD message format

The CR TYPE and CR ID uniquely define a specific control relationship between a group of controllers and the network elements (NEs) they control. If more than one NE along a CPD request path can respond to a CPD message with a given CR TYPE, they would have different CR IDs based on the provisioned CR ID value within the NE.

It **MUST** be possible to provision the control point with a CR ID. The requestor **MUST** set the CR TYPE to one of the following values:

- For CR TYPE = 1 (lawful intercept):
 - CR ID = 1: CMTS.
 - CR ID = 2: Aggregation router or switch in front of a CMTS.
 - CR ID = 3: Aggregation router or switch in front of media services (e.g., voice-mail).
 - CR ID = 4: Media gateway.
 - CR ID = 5: Conference server.
 - CR ID = 6: Other.
- For CR TYPE = 2 (DQoS) and CR TYPE = 3 (IPCablecom multimedia):
 - CR ID = 1: Default value.

The transaction ID is used by the requestor to relate a CPD response to a given CPD request. It is up to the requestor to pick transaction IDs that do not repeat within a time-frame that would prevent this correlation from occurring.

7.2 NLS-TL header parameters

For all network-based CPD messages, the requestor and control point **MUST** set the NLS-TL flags as follows:

- HOP-BY-HOP = 0;
- BUILD-ROUTE = 0;
- TEARDOWN = 0;
- BIDIRECTIONAL = 0.

The use of the AX_CHALLENGE and AX_RESPONSE flags are described separately in clause 9.

The requestor **MUST** set the flow-ID to a random number and the same value **MUST** be used by the control point in the response.

7.2.1 NLS-TL TLVs

The following NLS-TL TLVs are not used:

- NAT_ADDRESS;
- TIMEOUT;
- IPV4_HOP;
- IPV6_HOP.

7.3 CPD request

The CPD request message contents simply consists of a 32-bit number that contains some request control flags.




Figure 7-3 – CPD request message contents

The contents of the CPD request message are illustrated in Figure 7-3. The fields are as follows:

- Flag definitions:
 - 0x01: Subnet information request.
 - 0x02: Forward if not supported.
 - 0x04-0x80: Reserved and set to "0".
- The remaining 24 bits of the message contents are reserved and set to "0".

If the "forward if not supported" flag is set to "1", and the control relationship consisting of the CR TYPE and CR ID is not supported on this device, the control point MUST attempt to continue to forward the CPD packet towards its destination. Otherwise it MUST respond with a response code "control relationship not supported" as indicated in clause 7.5. This mechanism is helpful in several cases. A couple of examples are as follows:

- This allows for an LI to have a different control point to that used for QoS. So, for example, LI could be done on an aggregation device in front of the CMTS with the control point for QoS being on the CMTS. A QoS CPD request would be sent with the "forward if not supported" flag set to "1", while an LI CPD request would normally have this flag set to "0".
- There may be two control points along the path that are of the same CR TYPE but different CR IDs. Again, if the "forward if not supported" flag is set to "1" and the CPD request arrives at a router/switch with the correct CR TYPE but wrong CR ID, the control point MUST forward the message. If the flag is set to "0" and that control relationship is not supported, the control point MUST respond with response code "control relationship not supported".

If the "subnet information request" flag is set to "1", this is an indication to supply the value for the subnet that contains the destination IP address. Otherwise this information will not be returned in the response.

A CR ID of "0" in a CPD request message is considered to be a wild-card, i.e., it is a request to any device with any CR ID value that supports the CPD request message for the CR TYPE specified. When a wild-card value is used for CR ID in the CPD request, the control point MUST include the actual value of the CR ID assigned to the control point in the CPD response.

7.4 CPD response

This clause describes the message content format for the CPD response message.




Figure 7-4 – CPD response message contents

The contents of the CPD response message are illustrated in Figure 7-4. The fields are as follows:

- Flags (8 bits) include the following:
 - 0x01: Subnet information included.
 - 0x02: IP version of the control point: "0" for IPv4, "1" for IPv6.
 - 0x04-0x80: Reserved and set to "0".Other flag bits are not used and MUST be set to "0".
- Protocol supported flags (16 bits): interface/protocol supported by the control point for that CR type include the following:
 - 0x01: DQoS over COPS.
 - 0x02: IPCablecom multimedia over COPS.
 - 0x04: Li Tap-MIB with SNMPv3.Other flag bits are not used and MUST be set to "0".

Note that the same values can be re-used for different protocols for different CR types. Example: the 0x01 flag for DQoS could be used to describe some other protocol for some other CR type.
- Response code (8 bits) is an unsigned integer with the value "0" for a normal response. For error responses, refer to clause 7.5.
- Control point IP address is the IPv4 or IPv6 address of the control point.
- Subnet prefix is a 4-byte field representing the address prefix of the IPv4 subnet. If the "subnet information included" flag is "0", then this field as well as the prefix length field will be "0".
- Prefix length is a 1-byte unsigned integer representing the number of network bits in the subnet prefix.
- The remaining 24 bits of the message contents are reserved and set to "0".

7.5 Error responses

If an error occurs at the NLS transport level, either the `IPv4_ERROR_CODE` or `IPv6_ERROR_CODE` will be returned. If an error occurs at the CPD application level, a non-zero response code will be returned in the CPD response message. The following error codes are defined:

- Response code = 1: Control relationship not supported: this will be returned to the sender if the network element receiving the CPD request message does not support the control relationship (i.e., the CR type and CR ID specified) and the "forward if not supported" flag is set.
- Response code = 2: Poorly formed message (e.g., invalid CPD message type or invalid flags in the request).

8 Procedures

This clause provides a brief description of requestor and responder (control point) procedures.

8.1 Requestor

A requestor in this context is the controller that wants to apply some function to a control point within a network element such as a CMTS, aggregation router or media gateway. As such, it needs to obtain the IP address and protocol needed in order to apply some control over the media stream that passes through the control point.

The requestor receives information as to the IP address of the media stream. This is typically obtained via the session description information (SDP) for the media stream, i.e., the IP address is included in the "c=" line of the SDP. If NAT traversal [b-ITU-T J.160] is not used (i.e., there are no additional candidate addresses), then the requestor **MUST** send the CPD request to that media address as specified in the "c=" line of the SDP.

The requestor **MUST** send the CPD request to all candidate addresses except for a TURN server address in the case where it is able to determine that one of the candidates is a TURN server address. It may determine that a candidate is a TURN server address via one of the following mechanisms:

- by keeping a table of TURN server addresses; or
- by the existence of the "local-TURN" SDP attribute as defined in [b-ITU-T J.366.4].

When building the CPD request, the requestor **MUST** send the message to the well-known UDP port for NLS-TL (to be registered – see IANA considerations in Annex A). The requestor **MUST** send the CPD request message with CR type and CR ID set for the control point of interest. The requestor **MAY** also use the wild-card value for the CR ID.

Transaction ID selection: The requestor **MUST** pick a transaction ID that is not presently in use, which means that either:

- A response has already been received for a request with that transaction ID; or
- A response has been outstanding for some provisioned amount of time (default value, 600 seconds).

8.2 Control point

Network elements containing control points that support CPD will respond, based on the CR type that they support. A given NE may support multiple CR types. A CR ID may also be provisioned or simply left with its default value.

The NE containing the control point recognizes the NLS-TL message, based on the well-known UDP port (to be registered – see IANA considerations in Annex A). It then passes the application payload to the particular NLS application based on the application ID (in this case CPD). The CPD application looks at the CR type and CR ID to see if this control relation is supported. If not, the control point **MUST** respond as described in clause 7, based on the value of the "forward if not supported" flag.

If the control point does support the control relationship, it **MUST** respond to the source IP address of the request with a CPD response. The CPD response **MUST** contain the same CR type, CR ID, flow-id (NLS-TL) and transaction ID as was included in the CPD request message. The CPD application within the application **MUST** respond with either the IPv4 and IPv6 IP address depending on which is supported. The control point **MUST** also specify the version in the IP version flag. The protocol(s) supported **MUST** also indicate for the particular CR type.

The subnet information **MUST** be provided if requested. In the case where the subnet is provided, the control point **MUST** also indicate this in the "subnet information included" flag, corresponding to the "subnet information request" flag in the CPD request.

If the CR ID is set to a wild-card value (i.e., "0") in the request, the control point **MUST** respond with the actual CR ID value for that CR type.

9 Security considerations

The threat associated with illegitimate requests for control point information is that an attacker will have some additional information about the network element that provides that capability, as well as the networks attached to the control point. An attacker could perform some level of network discovery by requesting subnet information from control points.

The control point **MUST** be able to authenticate CPD requests before responding. The control point **SHOULD** only respond to CPD requests from authenticated sources. The NLS protocol itself provides an optional authentication mechanism that **MUST** be used. This uses a challenge mechanism as described in clause 13 of Annex A. This approach requires either pre-shared keys or some group keying mechanism for sharing a secret between the requestor and the control point. Since it does use a challenge-response mechanism, it will result in an additional round-trip delay.

Annex A

Network Layer Signaling: Transport Layer

(This annex forms an integral part of this Recommendation)

Network Working Group
Internet-Draft
Expires: May 8, 2006

M. Shore
K. Biswas
D. McGrew
Cisco Systems
November 7, 2005

Network-Layer Signaling: Transport Layer
draft-shore-nls-tl-01.txt

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of section 3 of RFC 3667. By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with RFC 3668.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on May 8, 2006.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

The RSVP model for communicating requests to network devices along a datapath has proven useful for a variety of applications beyond what the protocol designers envisioned, and while the architectural model generalizes well the protocol itself has a number of features that limit its applicability to applications other than IntServ. We are developing a modernized version that, among other things, is based on

Shore, et al.

Expires May 8, 2006

[Page 1]

a "two-layer" architecture that divides protocol function into transport and application. This document describes the transport protocol.

Table of Contents

1.	Introduction	3
1.1	Transport layer	4
2.	NLS-TL Messages	5
2.1	Message Processing Overview	5
2.2	NAT Traversal Support	6
2.3	NLS-TL Message Format	6
2.3.1	The NLS-TL Message Header	6
2.3.2	NLS-TL TLVs	7
3.	Sending NLS-TL Messages	13
4.	Messaging and state maintenance	14
4.1	BUILD-ROUTE	14
4.2	HOP-BY-HOP	14
4.3	BIDIRECTIONAL	15
4.4	Path Teardown Messages	15
4.5	Network Address Translation	15
5.	Application Interface	17
6.	NAT Interactions	18
7.	Using NLS-TL as a stand-alone NAT traversal protocol	19
8.	Discovery of non-NLS NATs, and recovery	20
9.	Endhost Processing	22
9.1	Sending	22
9.2	Receiving	22
10.	Intermediate node processing	24
11.	Using NLS-TL to support bidirectional reservations	25
12.	Proxy Considerations	26
13.	Security Considerations	27
13.1	Overview	27
13.2	Security Model	27
13.3	Cryptography	28
13.3.1	Keys	28
13.4	Datatypes	28
13.5	The Authentication Exchange (AX)	29
14.	IANA Considerations	31
14.1	NLS Application Identifiers	31
14.2	NLS TLVs	31
15.	References	31
	Authors' Addresses	32
A.	Acknowledgements	33
	Intellectual Property and Copyright Statements	34