

ETSI TS 102 822-4 V1.1.2 (2004-10)

Technical Specification

Broadcast and On-line Services: Search, select, and rightful use of content on personal storage systems ("TV-Anytime Phase 1"); Part 4: Content referencing

European Broadcasting Union



Union Européenne de Radio-Télévision

EBU·UER



Reference

RTS/JTC-TVA-PH1-09

Keywords

broadcasting, content, TV, video

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, please send your comment to one of the following services:

http://portal.etsi.org/chaicor/ETSI_support.asp

Copyright Notification

No part may be reproduced except as authorized by written permission.
The copyright and the foregoing restriction extend to reproduction in all media.

© European Telecommunications Standards Institute 2004.

© European Broadcasting Union 2004.

All rights reserved.

DECT™, **PLUGTESTS™** and **UMTS™** are Trade Marks of ETSI registered for the benefit of its Members.
TIPHON™ and the **TIPHON logo** are Trade Marks currently being registered by ETSI for the benefit of its Members.
3GPP™ is a Trade Mark of ETSI registered for the benefit of its Members and of the 3GPP Organizational Partners.

Contents

Intellectual Property Rights	5
Foreword.....	5
Introduction	5
1 Scope	6
2 References	6
3 Definitions, abbreviations and conventions	7
3.1 Definitions	7
3.2 Abbreviations	8
3.3 Conventions of specification syntax.....	8
4 Introduction	9
5 Key Concepts and Features	9
6 Location Resolution	11
7 The Authority	11
8 CRID	11
9 Locators.....	12
10 Instance Metadata Identifiers	14
11 Resolving Authority Record (RAR).....	16
11.1 Contents of Resolving Authority Record	16
12 Protocols for Location Resolution.....	17
12.1 Common Features and Interfaces	17
12.1.1 Resolution Handler	17
12.2 Location Resolution Over Unidirectional Networks	18
12.2.1 Guidance on use of Resolution Status Flags.....	20
12.3 Location Resolution Over Bi-directional Networks	20
12.3.1 Generic Bi-directional Resolution Server Discovery.....	21
12.3.2 Generic Bi-directional Request to Resolution Server	21
12.3.3 Generic Bi-directional Response From Resolution Server	22
12.3.4 Dynamic Behaviour of PDR and Location Resolution Server.....	23
12.3.4.1 Requirements for PDR dynamic behaviour.....	23
12.3.5 TCP/IP Based Resolution Server Discovery.....	24
12.3.5.1 Internet Service Discovery Query (RFC 2782).....	24
12.3.5.2 <i>TV-Anytime</i> Location Resolution Service Query	24
12.3.6 TCP/IP Based Request to Resolution Server	25
12.3.6.1 Further Requirements on a PDR HTTP Client.....	26
12.3.7 TCP/IP Protocol for Response from Resolution Server.....	26
12.3.7.1 Case 1: Returning the Result of Resolving CRIDs	26
12.3.7.2 Case 2: Returning an HTTP Re-direct	27
12.3.7.3 Case 3: Returning a Resolving Authority Record.	27
12.3.7.4 Encoding Server Response.....	27
Annex A (normative): XML Schema for Content Referencing	28
A.1 Schema Definition.....	28
A.1.1 Location Resolution Schema.....	28
A.1.2 Resolving Authority Record Schema	30
A.2 Example Instance Documents	31

Annex B (informative):	Example dynamic communication behaviour between a PDR and a Remote Location Resolution Server	33
Annex C (informative):	Bibliography	34
History		37

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 ETSI 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 Specification (TS) has been produced by Joint Technical Committee (JTC) Broadcast of the European Broadcasting Union (EBU), Comité Européen de Normalisation ELECTrotechnique (CENELEC) and the European Telecommunications Standards Institute (ETSI).

The present document is part 4 of a multi-part deliverable covering Broadcast and On-line Services: Search, select and rightful use of content on personal storage systems ("*TV-Anytime* Phase 1"), as identified below:

- Part 1: "Phase 1 Benchmark Features";
- Part 2: "System description";
- Part 3: "Metadata";
- Part 4: "Content referencing";**
- Part 5: "Rights management";
- Part 6: "Delivery of metadata over a bi-directional network";
- Part 7: "Bi-directional metadata delivery protection".

Introduction

The present document is based on a submission by the *TV-Anytime* forum (<http://www.tv-anytime.org>).

"*TV-Anytime* Phase 1" (TVA-1) is the first full and synchronized set of specifications established by the *TV-Anytime* Forum. TVA-1 features enable the search, selection, acquisition and rightful use of content on local and/or remote personal storage systems from both broadcast and online services.

The features are supported and enabled by the specifications for Metadata, Content Referencing, and Bi-directional Metadata Delivery Protection and Rights Management, TS 102 822-3 sub-parts 1 [10] and 2 [11], TS 102 822-4 (the present document), TS 102 822-5 [12] TS 102 822-6 [13] and TS 102 822-7 [14]. All Phase 1 Features listed in TS 102 822-1 are enabled by the normative *TV-Anytime* tools specifications. This list of Phase 1 Features is to be used as guidance to manufacturers, service providers and content providers regarding the implementation of the Phase 1 *TV-Anytime* specifications.

There will be further *TV-Anytime* phases published and Business Models for Post-Phase 1 are currently being defined to include Private and public domains, portable recordable media, super distribution (legal sharing of content between consumers), peripheral device support and mobile devices, amongst others.

1 Scope

The present document is the fourth in a multi-part series of "TS 102 822 series" documents produced by the *TV-Anytime* Forum. These documents establish the fundamental specifications for the services, systems and devices that will conform to the *TV-Anytime* standard, to a level of detail which is implementable for compliant products and services.

As is common practice in such standardization efforts, these specification documents were preceded by requirements documents which define the requirements for the *TV-Anytime* services, systems, and devices.

These specifications are parsed into three major areas, each described in a separate document of the series: TS 102 822-3 sub-parts 1 [10] and 2 [11] and TS 102 822-4 (the present document).

The other two documents are intended to define the environment and system architecture in which the standards in TS 102 822-3-1 [10], TS 102 822-3-2 [11] and TS 102 822-4 (the present document) are to be implemented. The first document in the series (TS 102 822-1 [8]) provides benchmark business models against which the *TV-Anytime* system architecture is evaluated to ensure that the *TV-Anytime* standard enables key business applications. The next document in the series (TS 102 822-2 [9]) presents the *TV-Anytime* System Architecture. These two documents are placed ahead of the other three for their obvious introductory value; TS 102 822-1 [8] and TS 102 822-2 [9] are both informative.

Although each of the TS 102 822 series documents is intended to stand alone, a complete and coherent sense of the *TV-Anytime* system standard can be gathered by reading all of the specification documents in numerical order.

The scope of the present document of the specification comprises the location independent identification of content, the location information, and the process of finding the location when an identifier is given. The content referencing process begins after a content item has been selected by a consumer through to, but not including, the actual acquisition of the desired content item.

The content referencing specification provides the mechanisms for location resolution in unidirectional network environments as defined in Business Model 1 of TS 102 822-1 [8], as well as support for location resolution in a bi-directional environment to fulfil Business Models 2 and 3.

2 References

The following documents contain provisions which, through reference in this text, constitute provisions of the present document.

- References are either specific (identified by date of publication and/or edition number or version number) or non-specific.
- For a specific reference, subsequent revisions do not apply.
- For a non-specific reference, the latest version applies.

Referenced documents which are not found to be publicly available in the expected location might be found at <http://docbox.etsi.org/Reference>.

- [1] IETF RFC 2396: "Uniform Resource Identifiers (URI): Generic Syntax".
- [2] IETF RFC 1591: "Domain Name System Structure and Delegation".
- [3] IETF RFC 2068: "Hypertext Transfer Protocol - HTTP/1.1".
- [4] IETF RFC 1945: "Hypertext Transfer Protocol - HTTP/1.0".
- [5] IETF RFC 2660: "The Secure HyperText Transfer Protocol".
- [6] IETF RFC 2782: "A DNS RR for specifying the location of services (DNS SRV)".
- [7] HTML 4.01 Specification, D. Ragget, A. LeHors, L. Jacobs.

NOTE: Available at: <http://www.w3.org/TR/html4/>.

- [8] ETSI TS 102 822-1: "Broadcast and On-line Services: Search, select, and rightful use of content on personal storage systems ("TV-Anytime Phase 1"); Part 1: Phase 1 Benchmark Features".
- [9] ETSI TS 102 822-2: "Broadcast and On-line Services: Search, select, and rightful use of content on personal storage systems ("TV-Anytime Phase 1"); Part 2: System description".
- [10] ETSI TS 102 822-3-1: "Broadcast and On-line Services: Search, select, and rightful use of content on personal storage systems ("TV-Anytime Phase 1"); Part 3: Metadata; Sub-part 1: Metadata schemas".
- [11] ETSI TS 102 822-3-2: "Broadcast and On-line Services: Search, select, and rightful use of content on personal storage systems ("TV-Anytime Phase 1"); Part 3: Metadata; Sub-part 2: System aspects in a uni-directional environment".
- [12] Void.
- [13] ETSI TS 102 822-6 (all sub-parts): "Broadcast and On-line Services: Search, select, and rightful use of content on personal storage systems ("TV-Anytime Phase 1"); Part 6: Delivery of metadata over a bi-directional network".
- [14] ETSI TS 102 822-7: "Broadcast and On-line Services: Search, select, and rightful use of content on personal storage systems ("TV-Anytime Phase 1"); Part 7: Bi-directional metadata delivery protection".

3 Definitions, abbreviations and conventions

3.1 Definitions

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

acquisition: retrieval of content

authority: organization that creates CRIDs

capture: storing the acquired content (e.g. to local storage)

content: anything the viewer would like to access (movies, games, TV programmes, radio programmes, etc.).

content creator: producers of content

content provider: entity that acts as the agent for and is the prime exploiter of the content

content reference: pointer to a specific content item

Content Reference Identifier (CRID): identifier for content that is independent of its location

Domain Naming System (DNS): system used on the Internet to register names that can then be mapped into IP addresses using a DNS server

Digital Video Broadcasting (DVB): set of standards used for European digital TV broadcasting

Electronic Programme Guide (EPG): means of presenting available content to the consumer, allowing selection of desired content

instance metadata identifier: identifier associated with a locator to bind to an instance metadata description

Internet Protocol (IP): generic name for the network protocols used on the Internet

location resolution: process of establishing the address (location and time) of a specific content instance from its CRID

locator: time and place where a content item can be acquired

metadata: generally, data about content, such as the title, genre and summary of a television programme as well as metadata also includes consumer profile and history data.

Programme Delivery Control (PDC): European system used by some broadcasters to accurately indicate when programmes start and stop

resolution handler: functional unit that provides location resolution on a specific transport mechanism

resolving authority: body which provides location resolution

Resolving Authority Record (RAR): information needed for retrieving the location resolution data for the given authority

service provider: aggregator and supplier of content which may include gateway and management roles

System Information (SI): collection of information tables used in DVB

3.2 Abbreviations

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

CRID	Content Reference Identifier
DVB	Digital Video Broadcasting
DNS	Domain Naming System
EPG	Electronic Programme Guide
IP	Internet Protocol
IPR	Intellectual Property Rights
MX	Mail eXchange
PDC	Programme Delivery Control
PDR	Personal Digital Recorder
RAR	Resolving Authority Record
SI	System Information
SRV	Search for seRVice
URI	Uniform Resource Identifier
URL	Uniform Resource Locator

3.3 Conventions of specification syntax

Within the present document, syntax is defined for various textual items.

Items in **bold** are textual characters that must be used exactly as specified.

Items in *<angle brackets and italics>* are replaced by a suitable value, which is specified beside the syntax definition.

An item in square brackets ("[" and "]") denote optional items. The entire syntax definition within the square brackets can be omitted, subject to the rules specified.

For example a syntax specification of:

Hello *<name>* how are you?

<name> is your name

The correct reading would be to take the text "Hello" and "how are you?" and insert your name between them.

Another example of a syntax specification is:

Hello *<name>* [*how are you?*]

<name> is your name

The correct reading would be to take the text "Hello" and append your name after it, and then optionally append the text "how are you?"

4 Introduction

The purpose of content referencing is to allow acquisition of a specific instance of a specific item of content. For example, if a consumer sees an announcement on TV saying "There'll be a new series of "The Agatha Christie Murder Mysteries" next year", he/she may want to instruct their Personal Digital Recorder (PDR) to record the whole series, but cannot since he/she does not know when the episodes are going to be broadcast. In fact, the broadcaster may not know yet either. Still the viewer will want to make sure at this point that he/she does not miss the opportunity to acquire the content.

To provide the capability desired by the consumer, the ability is needed to refer to content (in this example a series of programmes) independent of its location, whether that location is on a particular broadcast channel on some date and time, or on a file server connected to Internet, or wherever.

In this example, the PDR system would be provided with a reference for the series. In due time, the body who assigned the reference would provide the information required to link this reference to the individual episodes, and subsequently to a specific date and time for each episode so that the PDR would be able to acquire all of them.

This example demonstrates the purpose of content referencing - to provide the ability to refer to content independent of its location, and the ability to subsequently resolve such a reference into one or more locations where the content can be obtained.

Of course, *content* can refer to many types of information. In addition to the television programmes in the example above, it may include radio programmes, audio tracks, MPEG-4 objects, scenes, images, music, etc.

5 Key Concepts and Features

The content referencing process begins after a content item has been selected by a consumer through to, but not including, the actual acquisition of the desired content item. In figure 5.1 the scope of the content referencing process is shown. Some parts outside the scope of content referencing are included in figure 5.1 to give it a global context.

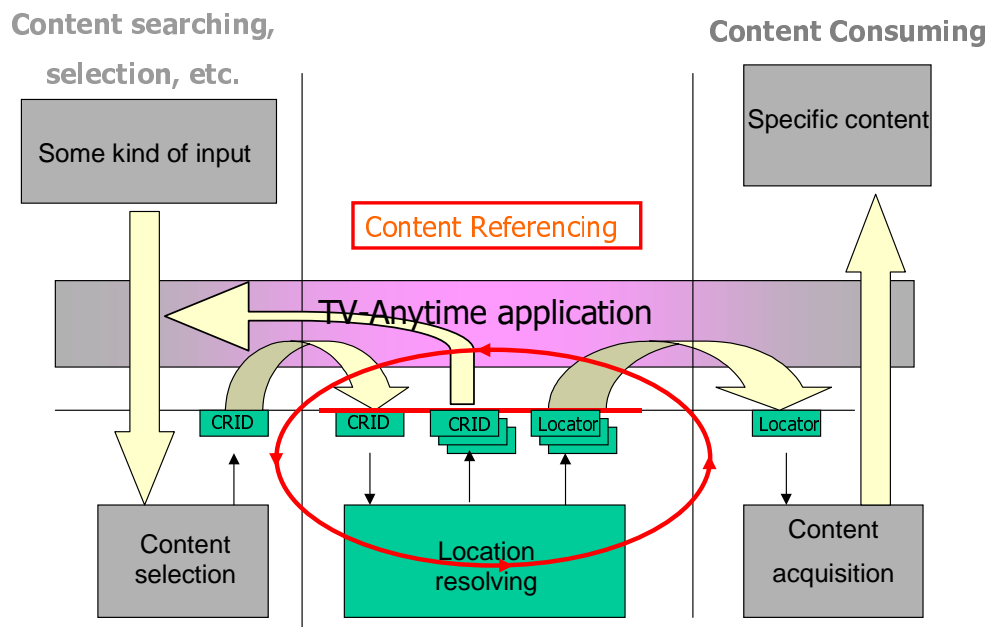


Figure 5.1: Content Referencing Environment

In this model we can identify three areas of interest:

- 1) The content selection process resulting in a Content Reference Identifier (CRID).
- 2) The location resolving service - the core of content referencing - resulting in one or more pointers to instances of that content (locators) or one or more content reference identifiers.
- 3) The retrieval process using one of these locators to subsequently acquire the content.

Before the process of content referencing can be employed, a selection process must supply a CRID. This CRID identifies a specific content item, but does not specify where that content item can be found. More detailed characteristics of the content, such as the price, the quality, the presence of commercials, etc., may not be known before the CRID is produced.

Given a CRID, there can still be numerous instances of the same desired content item (for example a broadcast may be repeated, a web site may be mirrored, the content item may be available from different content providers, via different networks, etc.). Some iteration may be required to narrow down the selection.

A content item that a CRID refers to might be a grouping of other content items, for example a CRID for referring to an entire series of programmes.

The key concept of content referencing is the separation of the reference to a content item - the CRID - and information that is related to its retrieval.

The separation provided by the CRID enables a one-to-many mapping between content references and the locations of the deliverables.

Content is a general term. It is the context of a *TV-Anytime* service that determines what content will be delivered upon following a locator. In that sense a content instance is any deliverable that can be acquired (via a respective locator) by a consumer. New types of content delivery merely require new types of locators.

In the same way, it is the *TV-Anytime* location resolution service provider who declares what can be considered content items (e.g. programmes, serials, etc.).

Content referencing, in conjunction with applications and additional metadata, may be able to deliver "first time" success, but may also involve some iteration enabling the overall system to deliver features including:

- Selection between alternatives.
- Selection of sub-elements.
- Selection between near matches.
- Selection of time of delivery.
- Selection ahead of time of release.
- Selection based on coding quality.
- Selection based on cost of delivery chain.
- Selection based on acceded rights.
- A point of reference - not necessarily unique - for the content item and any associated metadata.

Content referencing is:

- a form of content identification data and a specification for a set of rules for resolving this data into content locations from where this content can be retrieved.

Content referencing is not:

- the selection or retrieval process (it is in between);
- the definition of metadata used in any inquiry.

6 Location Resolution

Location resolution is the process of translating a CRID into other CRIDs or locators. Location resolution involves mapping a location-independent content reference (the CRID) to its location in time (e.g. scheduled transmission time in a broadcast system) and space (e.g. TV channel, IP address). In the context other present document, locations in time and space are referred to as "locators".

The process of location resolution may happen inside the PDR device (for example in a broadcast only system) or using a physically remote server (e.g. a server on the Internet).

7 The Authority

An authority, as defined in the present document, is the body that creates a CRID, which they guarantee to be unambiguous. An authority will also provide the ability for the CRID to be resolved into locators or other CRIDs.

The CRID shall ultimately be resolved either directly by the CRID authority or by another party. If another party is providing resolution, the ability to resolve the CRID requires the flow of some information from the authority to the resolution provider, in order to tie the CRID to its resolution.

In most PDR implementations, there will be multiple authorities that the PDR must be able to distinguish between. To distinguish between authorities, each authority has a unique name. This *TV-Anytime* specification uses the Domain Naming System (DNS) to provide unique names for each authority.

The syntax of an authority name is:

<DNS name>

<DNS name> is a registered Internet domain name. (See RFC 1591 [2] for DNS name registration). The **<DNS name>** is case insensitive and must be a fully qualified name according to the rules given by RFC 1591 [2].

Some example authority names are:

- www.broadcaster.com
- ISP.net
- www.commerce.com

8 CRID

A CRID is the output of the search and selection process and is an unambiguous identifier that refers to a piece of content, however multiple CRIDs may refer to that same piece of content.

It is permissible for a CRID to resolve into one or more other CRIDs. This CRID to CRID functionality can be used for two purposes:

- A CRID can resolve into multiple CRIDs for grouping content items such as an entire series of programmes.
- A CRID may also resolve into one or more CRIDs to allow one authority to refer to CRIDs of another authority.




Figure 8.1: Example of a Tree-Structured CRID

The syntax of the CRID is:

CRID://<authority>/<data>

<authority> uses the *TV-Anytime* authority naming rules given in clause 7 to assure uniqueness.

<data> is a free format string that is Uniform Resource Identifier (URI) compliant, and is meaningful to the authority given by the <authority> field. The <data> portion of the CRID is case insensitive.

In its entirety, the CRID is URI compliant. (See RFC 2396 [1] for URI compliance specification). As per RFC 2396 [1], the CRID:// part of the syntax is case insensitive.

To ensure interoperability it is recommended that all characters not within the range of characters allowed in a URI must be encoded into UTF8 and included in the URI as a sequence of escaped octets. An escaped octet is encoded as a character triplet, consisting of the percent character "%" followed by the two hexadecimal digits representing the octet code. For example, the character A would be represented as "%A", the character LATIN CAPITAL LETTER A WITH GRAVE would be represented as "%C3%80", and the character KATAKANA LETTER A would be represented as "%E3%82%A2".

Examples of syntactically valid CRIDs are:

Table 8.1: Examples of CRID

CRID	Description
CRID://company.com/foobar	CRID created by "company.com" authority, with a data part of "foobar".
crid://broadcaster.co.jp/wibble	CRID created by "broadcaster.co.jp" authority, with a data part of "wibble".
crid://broadcaster.co.jp/%E3%82%A8%E3%82%A4%E3%82%AC	CRID created by "broadcaster.co.jp" authority, with a data part of "E" "I" "GA" represented as KATAKANA LETTERS (Japanese characters) meaning "movie"

9 Locators

An instance of content may be located on various media such as local storage, live broadcast stream, data broadcast stream, data file on the Internet and data stream via the Internet.

A locator specifies a location, and possibly time of availability, where a content item can be acquired. There will be many formats of locators as there are many different means by which a PDR can acquire content. It is a requirement of a locator to ensure that it is possible for the PDR to parse enough of the locator to be able to decide if it has the ability to use the relevant transport mechanism.

The locator will be parsed and used by media-dependent methods to identify the content location and to acquire the content using the media or transport specific protocol. For example, a DVB locator will contain location parameters for a DVB stream, such as transport stream ID, service ID, table ID and event ID.

The syntax for a locator is:

<transport mechanism>:<transport system specific>

<transport mechanism> shall be unique for each mechanism. The string "CRID" shall not be used as the name for a <transport mechanism>.

<transport system specific> will be defined by the creator of the <transport mechanism>.

In its entirety, the locator is URI compliant. (See RFC 2396 [1] for URI compliance specification.)

For each <transport mechanism>, there will be only one format for the syntax of the <transport system specific> section.

The <transport system specific> section shall provide the following information:

- Location - This provides the location where the content can be acquired. It is possible that a PDR can receive content from many different providers which all share the same <transport mechanism>. For this reason, it is a requirement of a *TV-Anytime* locator that it is unambiguous between multiple providers using the same <transport mechanism>.
- Type of availability - It is possible that some schemes will be used for both schedule-based and on-demand acquired content. Content that is available at a specific time at a specific location (e.g. a broadcast TV programme, a web-cast) is schedule based. Schedule-based content must be retrieved at the time given by the locator. Content that can be received at any time between two limits (e.g. content that is on a server for one month) is on-demand based. On-demand based content can be acquired at any time that it is available.

For schedule-based content:

- Start time - This provides information about when the content is scheduled to start. It is required that start time be unambiguous with respect to local time zone as a PDR may be able to receive content from many different time zones.
- Duration of content - The length in time of the content.

For on-demand content:

- Start of availability - This optional field, when present gives the first moment in time when the content is available. This time shall be unambiguous in terms of time zone as a PDR may be able to receive content from multiple time zones.
- End of availability - This optional field gives the first moment in time that the content will become unavailable. This time shall be unambiguous in terms of time zone as a PDR may be able to receive content from multiple time zones.

In defining a syntax for the <transport system specific> section of a locator associated with a <transport mechanism>, there is an assumption about the environment the PDR exists within. For each <transport mechanism> a PDR will need a certain amount of information in order to receive content from this system. This information may be provided within the transport mechanism, or by any other means appropriate to the target PDR.

As an example of the information within the environment, a PAL Western European locator might use the network and channel identifier within its syntax. The mapping of network and channel identifier to a physical channel uses information carried in the vertical-blanking interval.

The transport mechanism may provide a more accurate timing system than the start time, which the PDR may decide to use for accurate content capture (e.g. PDC information, DVB event IDs).

10 Instance Metadata Identifiers

TS 102 822-3 sub-parts 1 [10] and 2 [11] provides the ability to describe specific instances of content. Whilst *TV-Anytime* created the CRID for location independent selection of content, it recognizes that there will be occasions where the consumer may wish to acquire a location dependent version of a piece of content.

To enable this scenario, an optional identifier may be assigned to each locator, and may be signalled in the instance description metadata. When the location of the content changes the locator might change, but the instance metadata identifier will not change.

An instance metadata identifier is only guaranteed to be unique within the scope of the CRID to which it has been assigned. It is permissible to assign the same identifier to different CRIDs.

A PDR may use the instance metadata identifier to track changes in the location of an instance of a piece of content. A PDR will need to use both the CRID and the identifier to track changes (as opposed to just the identifier) because the identifier is only guaranteed to be unique within a given CRID.

The identifier can also be used to link the instance description metadata with the information retrieved by the location resolution process.

Each instance metadata identifier will need to be unique within the scope of a given CRID. In order to enable multiple parties to provide instance identifiers, it is a requirement for every instance metadata identifier to begin with a name that is guaranteed to be unique. To ensure that this name is unique, it must be a registered Internet domain name. (See RFC 1591 [2] for DNS name registration). The name portion of an instance metadata identifier specifies the body that created the instance metadata identifier.

For a given instance identifier creator, there shall be only one instance metadata identifier per CRID - locator pair.

The syntax of an instance metadata identifier is:

imi:[<name>/]<data>

<name> is a registered Internet domain name. (See RFC 1591 [2] for DNS name registration). The **<name>** is case insensitive and must be a fully qualified name according to the rules given by RFC 1591 [2]. If the **<name>** portion of an instance metadata identifier is the same as the authority name of the CRID, the name and the forward slash may be omitted.

<data> is a free format string (with the exception that the forward slash character is prohibited) that is Uniform Resource Identifier (URI) compliant, and is meaningful to the body specified by the **<name>** field. The **<data>** portion of the instance metadata identifier is case insensitive.

Examples of syntactically valid instance metadata identifiers are:

Table 10.1: Examples of Instance Metadata Identifiers

Instance Metadata Identifier	Description
imi:company.com/foobar	Instance metadata identifier created by "company.com", with a data part of "foobar".
imi:broadcaster.co.jp/broodjeham	Instance metadata identifier created by "broadcaster.co.jp", with a data part of "broodjeham".
imi:meaning	Instance metadata identifier created by the CRID authority (the <name> portion of the instance metadata identifier was omitted, therefore the CRID authority is used), with a data portion of "meaning".

It is optional for location resolution information to contain instance metadata identifiers. Therefore a PDR shall not assume that these identifiers are always available in its acquisition procedures.

The following example demonstrates the use of an instance metadata identifier to track the change in location of a piece of content.

Table 10.2: Example of a resolved CRID

CRID	Locator	Identifier
crid://example.net/abc123	dvb://123.5ac.3be;3e45@2001-12-07T12:00:00.00+01P02:10	imi:def.com/1
	ftp://example.net/mirror/def123.mov	imi:def.com/2

Table 10.3: Example of the same resolved CRID after a change in location identified by "imi:def.com/1"

CRID	Locator	Identifier
crid://example.net/abc123	dvb://123.5ac.100;1e4a@2001-12-07T15:00:00.00+01P02:10	imi:def.com/1
	ftp://example.net/mirror/def123.mov	imi:def.com/2

In the following example the name portion has been omitted from the instance metadata identifier, because it is the same as the name of the CRID authority. In this example, the identifier "imi:1" is equivalent to writing "imi:example.net/1".

Table 10.4: Example of an Instance Metadata Identifier with no name portion

CRID	Locator	Identifier
crid://example.net/abc123	dvb://123.5ac.100;1e4a@2001-12-07T15:00:00.00+01P02:10	imi:1
	ftp://example.net/mirror/def123.mov	imi:2

Table 10.5 shows a more complex set of permissible combinations of CRID and instance metadata identifier.

Table 10.5: Example of a more complex resolution table

CRID	Locator	Identifier	Note
crid://example.net/abc123	dvb://123.5ac.100;1e4a@2001-12-07T15:00:00.00+01P02:10	imi:1	1
	ftp://example.net/mirror/def123.mov	imi:2	
crid://example.net/je98	dvb://123.6ef.200;5e23@2002-01-31T14:20:00.00+01P00:30	imi:mdprov.com/01	2
	dvb://123.6ef.200;1c24@2002-02-14T14:20:00.00+01P00:30	imi:mdprov.com/02	
crid://example.net/ja90	dvb://2a3.faa.100;8ee9@2002-01-29T01:20:00.00+01P01:30	imi:mdprov.com/01	3
	dvb://c01.ad3.400;003c@2002-02-14T18:00:00.00+01P01:00	imi:mdprov.com/02	
crid://broadcaster.co.uk/0203	dvb://c01.ad3.400;003c@2002-02-14T18:00:00.00+01P01:00	imi:1	4
		imi:mdprov.com/01	
NOTE 1: Resolution for CRID "crid://example.net/abc123" with two instance metadata identifiers. No name is provided for the instance metadata identifier, therefore the CRID authority of "example.net" is used.			
NOTE 2: Resolution for CRID "crid://example.net/je98" with two instance metadata identifiers. In this example an instance metadata identifier provider of "mdprov.com" has been specified.			
NOTE 3: The same as the previous two examples, showing the resolution of a CRID to two locators with associated instance metadata identifiers. Note that this example uses the same instance identifiers as the previous example. The instance identifiers are disambiguated by the CRID.			
NOTE 4: This example shows two different instance metadata identifiers being assigned to the same CRID and locator. This is only allowed because the identifiers are from different identifier providers.			

11 Resolving Authority Record (RAR)

The Resolving Authority Record (RAR) is an essential element in the location resolution process for both the unidirectional and bi-directional networks. It provides the information needed for retrieving the location resolution data for a given authority.

Each resolution authority will require one or more resolving authority records to exist in the PDR for location resolution to take place. Each resolving authority record will need to be placed inside some sort of transport specific container, which allows the PDR to know that this is a resolving authority record.

In the case of multiple records for the same authority, for each location resolution to be done, the PDR can choose to use any of them.

11.1 Contents of Resolving Authority Record

The present document does not define a transmission format for the RAR in uni-directional networks. The present document does however define the information that must be carried by the encoded RAR. A transmission format is specified for the bi-directional content referencing process over TCP/IP, which is given in clause 12.3.

A *TV-Anytime* compliant RAR shall contain at least the following items of information:

Resolution Provider: The name of the body that is providing location resolution. It is possible that different bodies are providing location resolution for a single authority, for example a broadcaster might be providing a secondary resolution service for a content creator. These different location resolution providers need to be able to identify themselves for such purposes as updating their resolving authority records. The name of the resolution provider follows the naming rules given in clause 7.

Authority name: The name of the *TV-Anytime* CRID authority as described in clause 7.

Class: The class field defines whether this authority record defines a resolution authority that can resolve all CRIDs for this authority name (class = primary) or only resolves some CRIDs for this authority name (class = secondary). The reason for providing secondary resolving authority records is to allow one resolution provider to provide a partial resolving service for a content referencing authority. An example of this is a broadcaster providing partial resolution of another authority (such as the content creator) for the content they broadcast.

Version number: A number that is incremented each time the resolution provider wishes to update its resolving authority records for a given authority name. The set of authority records the PDR should update is based on the combination of the authority name and the resolution provider. When a new version number is received for an authority by a resolution provider, all old resolving authority records for this authority name and resolution provider combination will be discarded by the PDR. When the version number reaches $2^{32} - 1$, the next version number shall be zero. Tables are considered equivalent if they have the same values of Resolution Provider, Authority name, version number and URL.

URL: The location where resolution information can be found. The URL could point to a broadcast stream, or to a server on the Internet or any other place where location resolution information can be found. The syntax of the URL is that of the locator as given in clause 9.

First valid date: The first date when this authority can be used, in a form that is unambiguous with respect to time zone.

Last valid date: The first date when this authority cannot be used, in a form that is unambiguous with respect to time zone.

The reason for providing start and end dates for resolution is so that resolution providers can move their resolution URLs and be sure all PDRs have switched to the new URL once the last valid date of the old resolution record has passed.

Weighting: The weighting field is used to give a hint to the PDR as to the order to try multiple records for an authority from the same resolution provider by providing the largest weighting number to the URL that should be tried first. The weighting field is only used to provide ordering between resolution provider records for the same combination of resolution provider and authority name and not for ordering one provider over another.

Table 11.1: An example Resolving Authority Record

Field Name	Contents
Class	Secondary class
Weighting	1
First valid date	9:30 am 26 September 2000
Last valid date	6:00 pm 28 November 2000
Resolution Provider	tva.resprov.com
Authority name	autnam.com
URL	http://www.resprov.com/lr/autnam

12 Protocols for Location Resolution

12.1 Common Features and Interfaces

This clause defines the features common to location resolution using unidirectional or bi-directional networks. It is included here to provide a context for the details that follow in clauses 12.2 and 12.3, and provides the descriptions of the terms that will be used in these clauses.

12.1.1 Resolution Handler

In the *TV-Anytime* environment, content referencing services can be delivered through various delivery systems such as IP networks or broadcast TV. Figure 12.1 shows the concept of a modular resolution system, with multiple resolution handlers as required for specific location resolution transport mechanisms.

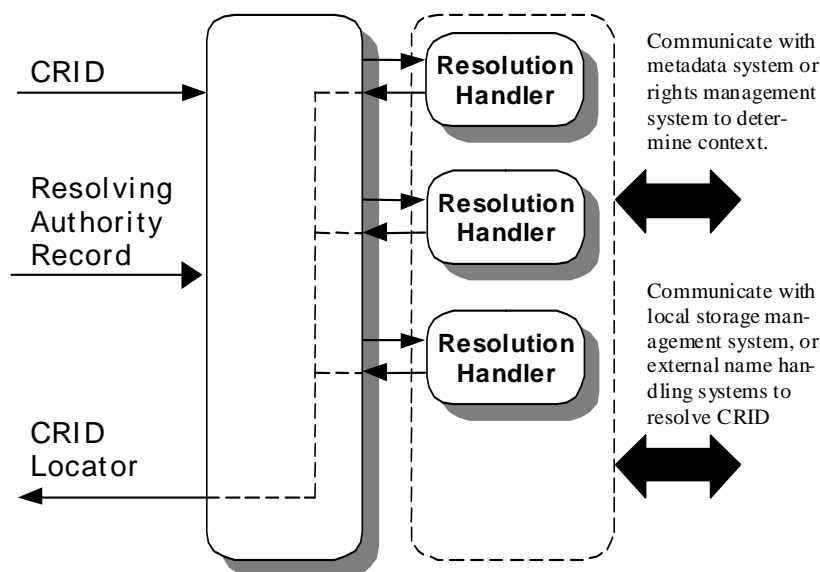


Figure 12.1: CRID Resolution Architecture

Figure 12.1 provides a network-transparent method for resolution, with multiple resolution handlers providing the network and protocol-dependent CRID resolution. For example, one resolution method is resolving CRIDs locally, by co-operating with the local storage management system to resolve the location of the locally stored content. Another resolution system might resolve a CRID using external name handling systems via a back channel or an Internet connection. Another system may refer to System Information (SI) tables, which contain mapping tables between CRIDs and locators and are transported in a digital broadcast stream.

It is envisaged that the resolution handler is an extensible resolution mechanism so that in the future when a new location resolution transportation method is developed, it is possible to extend the resolution process by implementing and plugging in a new resolution handler.

Steps in the content referencing process are now summarized:

- 1) The CRID is used to decide which resolution handlers to invoke to resolve this CRID.
- 2) The resolution request is forwarded to the appropriate resolution handlers.
- 3) Each chosen resolution handler tries to resolve the CRID into locators or another set of CRIDs. The resolution process depends on the implementation of resolution handler. As part of this resolution process, the resolution handler may need to communicate with an external system. Some example processes between a resolution handler and an external system are:
 - Resolve the CRID using a mapping table located in the PDR. This method is appropriate for locally recorded contents or cached information from broadcast or IP networks.
 - Resolve the CRID using the broadcast stream.
 - Resolve the CRID via Internet or back channel.

12.2 Location Resolution Over Unidirectional Networks

This clause of the specification defines the common features of resolution handlers that work with unidirectional networks. Each resolution handler using a unidirectional network will have a dynamic behaviour similar to the one given in the flow chart figure 12.3.

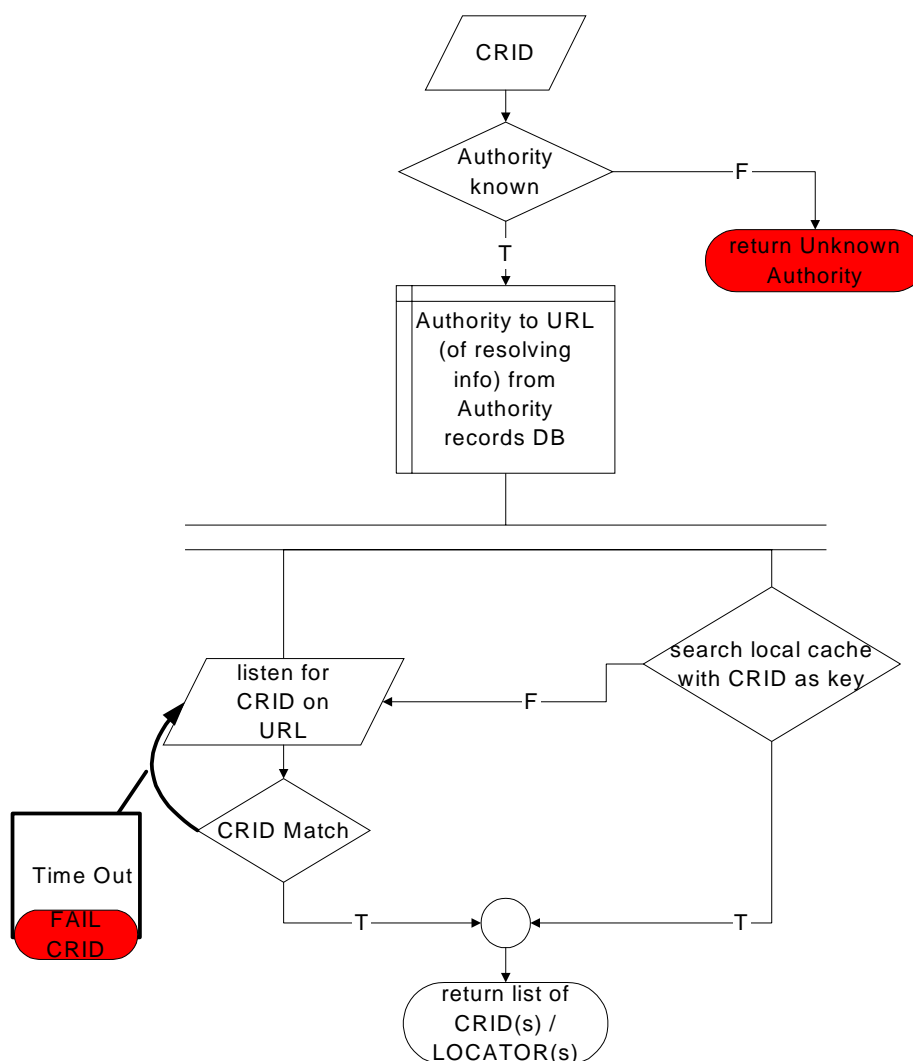


Figure 12.2: Dynamic Behaviour of Resolution Handler using Unidirectional Network

The first step of location resolution in a unidirectional system is for the PDR to learn where it can receive location resolution information. This location is provided by the resolving authority record, which will need to be broadcast to the PDR at some known location. Failure to find any authority records for the given CRID's authority will result in a failure to resolve CRIDs created by that authority.

Once this resolving authority record has been located, the PDR will know where to listen for location information for the given CRIDs' authority (by using the URL field of the appropriate authority record).

The PDR will have to choose one or more resolution handlers to use for resolving the desired CRID, as it will need to pick the resolution handlers that can understand the protocols used to carry the location resolution information. An example of this choice is the PDR choosing the DVB resolution handler if the resolution record says that resolution information is being sent on a DVB transport stream.

It is likely that a PDR implementation will also use the local resolution handler in case the content the CRID refers to is already available locally.

The information that will be sent in the unidirectional location resolution stream shall be in the form of a table that consists of CRID to message mappings. Each input CRID will output to a message that shall contain a status field. When the status field contains a value to indicate that the input CRID is valid, the message shall contain one or more CRIDs, or one or more locators.

A *TV-Anytime* compliant unidirectional location resolution stream shall contain a stream of matched pairs of:

CRID	Message
------	---------

Each message shall at least contain:

Table 12.1: Location Resolution Message Format for Unidirectional System

Field	Description
Status	"CRID is resolved" (resolution list follows). "discard CRID" (E.g. CRID is no longer valid). "resolve after date <xxx>" (keep CRID, try later).

If Status = CRID is resolved:

Field	Description
Acquisition directive	"all" (all items of the following list must be acquired) "any" (any item from the following list may be acquired as they are alternative locations for the same content).
A list of CRIDs or a list of Locator(s)	CRIDs will conform to the syntax given in clause 8. Locators will conform to the syntax given in clause 9. Optionally, each locator can have an associated instance metadata identifier.
Resolution complete	Is the list complete? yes (CRID is completely resolved) (e.g. this is the last episode of the series) no (CRID might resolve into more items at a later date)
Re-resolution date	Date after which the PDR should attempt to re-resolve the CRID. This field is only meaningful when the Resolution Complete flag is set to "no". This date shall be unambiguous with respect to time zone

If Status = "Resolve after date"

Field	Description
Date	The date and time on or after which the PDR should try to re-resolve the CRID. This field shall be unambiguous with respect to time zone.

Table 12.2 describes the acquisition behaviour of a PDR in response to the status flags

Table 12.2: Interpretation of Status Flags

Acquisition Directive	Resolution Complete	Description
All	No	Acquire all items in list and await further items.
All	Yes	Acquire all content items, after which acquisition of this CRID is complete.
Any	No	Select any of current item list (after which the acquisition is complete) or await additions to list.
Any	Yes	Select any one of the list items after which the acquisition is complete.

12.2.1 Guidance on use of Resolution Status Flags

CRIDs that when resolved translate into one or more CRIDs, with the acquisition directive of "all" and have the resolution complete flag set to "no" can be used for grouping of content that changes over time, for example a TV series. Such a group CRID may continue with its resolution complete flag set to "no" for a long period of time if the series has no planned end. It is permissible for a PDR to allow the user to view the content the PDR has acquired for this incomplete group CRID.

A CRID that resolves into one or more locators should not be used for an ongoing group (such as a series) as the PDR will assume that when the acquisition directive is set to "all" it must acquire all the parts of the content specified by the list of locators before that content is fully acquired, and hence viewable by the user.

Table 12.3: PDR behaviour in response to acquisition directive

Resolution type	Acquisition directive	Description of PDR behaviour
CRID to CRID(s)	All	All result CRIDs should be acquired. Each result CRID can be considered either as its own content item or as part of the group.
CRID to CRID(s)	Any	Any of the result CRIDs can be acquired as they are considered equivalent by the authority that created this CRID.
CRID to Locators(s)	All	All items must be acquired before the content is complete. It is an implementation choice as to whether a PDR will allow viewing of incomplete content.
CRID to Locators(s)	Any	Select any one of the locators as they are considered equivalent by the authority that created this CRID.

It is an implementer's option as to whether the PDR will always go to the unidirectional stream for location resolution, or provide some local caching mechanism. This caching mechanism might be to cache resolved CRIDs, or to cache the unidirectional stream in case it is needed later. This cached information might be used by the resolution handler that handles the unidirectional stream or by creating another resolution handler that uses locally cached data.

12.3 Location Resolution Over Bi-directional Networks

In order for a PDR to use the location resolution services over bi-directional networks it is necessary to define a protocol to allow the PDR to initiate a connection and then transfer requested data between itself and a resolution service located at a remote server.

This clause specifies how a PDR can discover the location of such a server on a bi-directional network and the TVA protocol to achieve the appropriate *TV-Anytime* data transfers over such a network.

This clause does not specify how content is retrieved over a bi-directional network.

12.3.1 Generic Bi-directional Resolution Server Discovery

This discovery stage may not be required by all network implementations.

Given a CRID to resolve, the first step is to find a server that might be able to resolve this CRID. The process of server discovery is based on using the authority name from the CRID.

The present document makes the assumption that the CRID the PDR has been asked to resolve is from an authority completely new to the PDR, so that it has no prior knowledge of where to get this CRID resolved. In an actual implementation there may be some sort of caching of previous server discoveries, but the exact nature of this caching is implementation specific.

In the following example, an intermediate server is contacted to discover the address of the location resolution server.

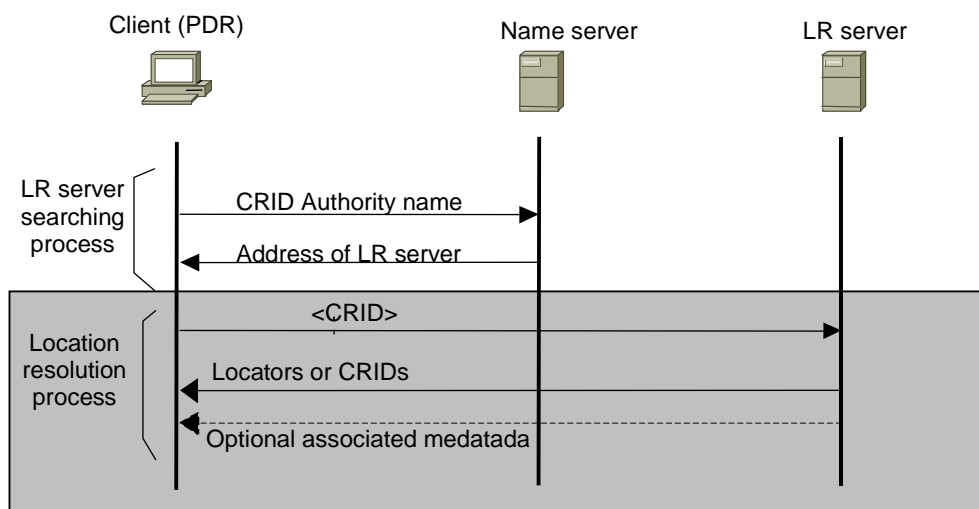


Figure 12.3: Example of Location Resolution Server Discovery

12.3.2 Generic Bi-directional Request to Resolution Server

Once the resolution server has been discovered, the next step is to communicate with this resolution server. The inputs to the server will be a list of one or more CRIDs to resolve, and optional flags to specify how the response should be created.

The *TV-Anytime* defined optional flags are:

- 1) SubmittedCRID.
- 2) Result.

Table 12.4: Submitted CRID Flag Description

Value of SubmittedCRID flag	Description
0	No descriptive metadata about the CRIDs being submitted shall be returned with the resolution information.
1	Instances of the ProgramInformationTable or GroupInformationTable schemas shall be returned which describe the submitted CRIDs, if the location resolution server has this information.
All other values	Reserved.

Table 12.5: Result Flag Description

Value of Result flag	Description
0	No descriptive metadata about the results from the location resolution shall be returned with the resolution information.
1	If the submitted CRID resolves into further CRIDs, instances of the ProgramInformationTable or GroupInformationTable schemas shall be returned for each resolved CRID, if the location resolution server has this information. If the submitted CRID resolves into locators, instances of the ProgramLocationTable schema shall be returned for each locator, if the location resolution server has this information.
All other values	Reserved.

If any of the optional flags are not present they are assumed to have a value of zero.

12.3.3 Generic Bi-directional Response From Resolution Server

The resolution server will respond with one of three possible types of information:

- 1) The result from resolving the CRID. The response will contain instance(s) of XML schemas defined in TS 102 822-3-1 [10] and in annex A of the present document. The permitted response schemas are listed later in this clause.
- 2) A Resolving Authority Record (RAR). The PDR should store this RAR using the rules given in the unidirectional model, and then contact the server given by the URL field of the RAR.
- 3) A re-direct. The resolution server will return a message that gives the address of another location resolution server to contact.

In case (1) and case (3) where a PDR does not receive a RAR, the PDR shall assume the location resolution server is a primary class server, and follow the appropriate rules given in the present document for primary class resolution servers.

The permitted instantiated elements of the XML schemas defined in TS 102 822-3-1 [10] and in annex A of the present document are:

- GroupInformationTable.
- ProgramInformationTable.
- ProgramLocationTable.
- ContentReferencingTable.

When an XML instance document is returned by the location resolution server, an instance of a ContentReferencingTable that contains CRIDResult or LocationsResult for each CRID submitted shall be returned. When indicated by the SubmittedCRID and Result flags, instances of GroupInformationTable, ProgramInformationTable and ProgramLocationTable may also be returned.

The results may be returned in any order the location resolution server desires and not necessarily the order in which the PDR specified in the request to the server.

12.3.4 Dynamic Behaviour of PDR and Location Resolution Server

For the case where the PDR is connected to a bi-directional network there is a need to specify some aspects of the dynamic behaviour of location resolution requests in a *TV-Anytime* compliant PDR. This is required in order to avoid large numbers of PDRs all trying to contact the same location resolution server at the same moment in time.

As well as a PDR having a prescribed manner to avoid excess loading of a location resolution server, it is advisable that a location resolution provider takes certain steps in their implementation to aid the reduction of large instantaneous loads on their servers. When a server returns a response that indicates a CRID should be resolved again at a future date, it is advisable for the server not to return the same time and date information to all clients requesting this CRID, as they will then all try to contact the server at the same moment in time.

It should be noted that a PDR, in addition to receiving location information from a bi-directional network, might also be receiving information from one or more broadcast networks. When receiving information from a broadcast chain, there is the potential for many PDRs to all receive the same time and date for re-resolution.

In order to reduce the effect of instantaneous server loading, the PDR will need to implement a dynamic behaviour that does not cause overloading of a location resolution server by frequent repeated access to the same server.

12.3.4.1 Requirements for PDR dynamic behaviour

- 1) When the date and time for re-resolution arrives, the PDR shall wait a further random amount of time before contacting the location resolution server. This is required in order to reduce the chances of the location resolution server being swamped by many PDRs contacting it at the same moment in time.
- 2) If the location resolution server returns a re-resolution time and date that is in the past, the PDR shall wait a random amount of time before contacting the location resolution server again.
- 3) If the location resolution server is unavailable, the PDR shall assume the response for the CRIDs being resolved is "resolve again at a later date". The date and time at which to re-resolve is based on the current date and time plus a random amount of time.
- 4) When a location resolution server returns information to specify that the CRID is unknown and this location resolution server is a secondary class server, and the PDR wishes to try this server again, the PDR shall wait a random amount of time before trying again.
- 5) When a location resolution server returns information to specify that the CRID is unknown and this location resolution server is a primary class server, the PDR should cease trying to resolve this CRID as it will never be resolvable.
- 6) The random time interval generator algorithm inside a given manufacturer's PDR should not be configured in such a manner that all their PDRs have identical random time interval sequences. Testing compliance to this requirement is not specified in the present document.
- 7) The standard deviation of the random time interval generator shall be at least 10 minutes. Testing compliance to this requirement is not specified in the present document.

Annex B describes a candidate for the dynamic behaviour of the PDR when communicating with a remote location resolution server.

12.3.5 TCP/IP Based Resolution Server Discovery

As the <DNS name> part of the authority name is a name that is a registered Internet domain name, the mechanisms defined for DNS name lookup can be used as part of the server discovery phase.

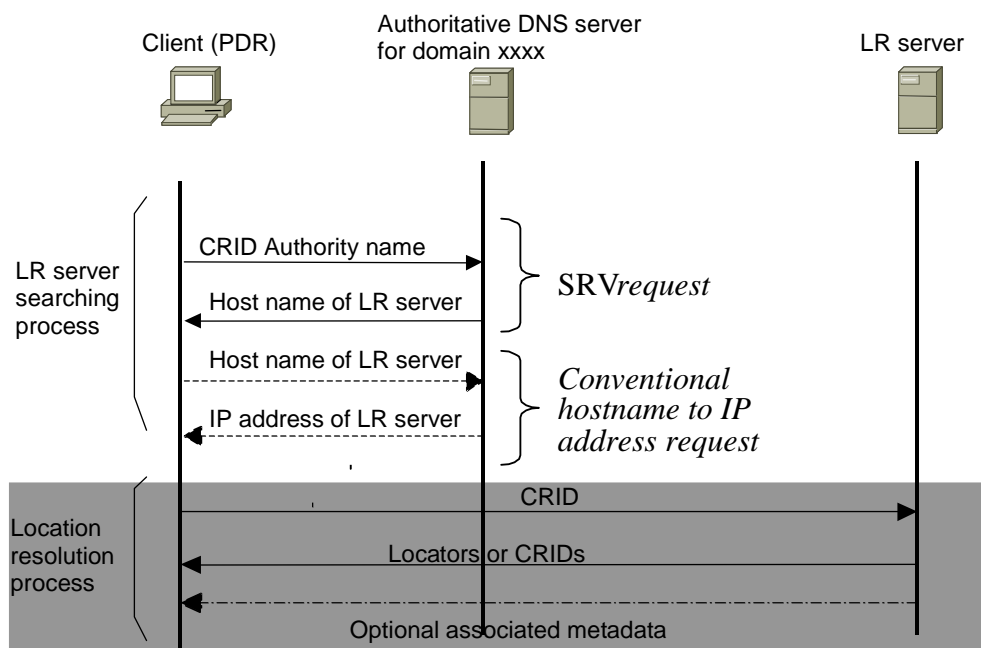


Figure 12.4: Stages of TCP/IP Based Resolution Server Discovery

12.3.5.1 Internet Service Discovery Query (RFC 2782)

RFC 2782 [6] provides an expansion of the DNS (see RFC 1591 [2]) system that is used to allow Internet connected machines to find mail servers. In addition to being able to search for Mail eXchange (MX) records, it is also possible to Search for seRVice (SRV) records.

An RFC 2782 [6] compliant query is made up of several parts, namely:

_Service._Protocol.Name

For example a query for an HTTP server for example.com would be "_http._tcp.example.com". The DNS server will respond with the host name and port number corresponding to the network location at which the requested service can be found. In the previous example, the return might be "webserver2.example.com on port 80".

12.3.5.2 TV-Anytime Location Resolution Service Query

The name for the *TV-Anytime* location resolution service is "lres" which is a shortened version of "location resolution". The use of a shortened name is adopted, as there is a limit of 512 characters for a DNS response in some DNS client implementations.

The complete query name will look like:

_lres._tcp.<CRID authority>

For example, given the CRID "cid://europe.example.com/9afc2", the query string would be "_lres._tcp.europe.example.com" which would be sent to a DNS server that provides lookups for "europe.example.com".

Another example, would be, given the CRID "cid://example.co.uk/9afc2", the query string would be "_lres._tcp.example.co.uk" which would be sent to a DNS server that provides lookups for "example.co.uk".

12.3.6 TCP/IP Based Request to Resolution Server

The protocol for sending a request to the location resolution server is based on the HTTP protocol (see RFC 2068 [3]). The format of the query string shall follow that generated by the submission of an HTML form using a GET request:

http://<Path_to_server_script>?[key=value]&[key=value]

where the key/value pair is repeated as required. The key is case sensitive and shall be represented using the exact case given for each key as specified in this clause.

The precise specification for encoding the key-value pairs into the HTTP URL is given in clause 17.13 of the HTML 4 specification [7]. It is the option of the location resolution service provider to implement this service using any server side technology they wish (CGI scripts, Java servlets etc.).

Each key shall be one of:

Table 12.6: Key definitions for HTTP URL encoding

Key	Description	Allowable value
CRID	The CRID to resolve	A CRID inside quotes
SubmittedCRID	Used to specify whether metadata on the CRID being resolved is desired. See clause 12.3.2 for semantics of this flag.	A number. See clause 12.3.2 for allowed values.
Result	Used to specify whether metadata is required for each of the results from resolving this CRID. See clause 12.3.2 for semantics of this flag.	A number. See clause 12.3.2 for allowed values.

It is permissible to resolve multiple CRIDs in one HTTP request by using multiple "CRID" keys in the URL, but the "SubmittedCRID" and "Result" keys may only be specified once in a request as they affect all CRIDs being resolved.

For the first connection to a location resolution server after the DNS based server location phase, the <path to server> is the hostname, a colon character, the textual representation of the port number followed by a slash. If the port number is 80, the colon and port number may be omitted.

For example, if the DNS server returned host name "computer2.example.com" on port 1234, the <path to server> would be:

computer2.example.com:1234/

When a location resolution server provides a re-direct using an HTTP redirect, the <path to server> is the URL returned by the "Location" header of the HTTP redirect response.

For example, if the HTTP response was:

Location http://redirect.example.com/tva/lr

The <path to server> would be:

redirect.example.com/tva/lr

When a location resolution server provides a re-direct using the RAR, the <path to server> is the URL field from the RAR.

For example, if the URL field of the RAR contained the value:

http://kaas.example.nl/scripts/resolution.cgi

the <path to server> would be:

kaas.example.com/scripts/resolution.cgi

Examples of valid complete URLs are:

`http://computer2.example.com:1234/?CRID="crid://example.com/abc123"`

`http://broadcaster.com/?CRID="crid://broadcaster.com/abc123"&CRID="crid://broadcaster.com/def456"`

`http://kaas.example.com/scripts/resolution?CRID="crid://example.com/abc123"&Result=1`

`http://redirected.example.com/tva/lr?CRID="crid://example.com/abc123"&SubmittedCRID=1`

12.3.6.1 Further Requirements on a PDR HTTP Client

The PDR shall at least implement the HTTP v1.0 specification (see RFC 1945 [4]) for issuing a GET request to the server. In addition to the requirements of HTTP v1.0, the PDR shall also send the HTTP v1.1 header of "host" and the HTTP v1.0 header "user-agent".

The HTTP client in the PDR shall support at least the following MIME type:

- text/xml

EXAMPLE 1: The HTTP client will need to send an accept command with at least the following components:

Accept: text/xml

In order to permit the secure transfer of resolution requests from the PDR to the location resolution server and secure results from the location resolution server, the PDR and location resolution server may negotiate the secure HTTP protocol as given in RFC 2660 [5].

If the PDR supports decoding of an encoded response from the resolution server (e.g. decompressing a compressed response), the PDR shall indicate this by sending an HTTP "Accept-Encoding" header.

EXAMPLE 2: A HTTP client capable of decompression would send a header containing:

Accept-Encoding: x-compress; x-zip

12.3.7 TCP/IP Protocol for Response from Resolution Server

The response from the location resolution server will be based on the HTTP v1.0 specification (see RFC 1945 [4]).

The response from the server can be one of three possible types:

- 1) The result of resolving the CRIDs sent to the server.
- 2) An HTTP re-direct to allow the distribution of services amongst a number of machines.
- 3) A standard Resolving Authority Record (RAR) to facilitate PDR caching, server load balancing, dynamic server administration and cross platform capability.

The use of a MIME type given by the "Content-Type" HTTP header shall be used to indicate which of the two possible server responses (type 1 or type 3) is being returned.

12.3.7.1 Case 1: Returning the Result of Resolving CRIDs

If the response from the location resolution server is the result of resolving the CRIDs requested by the PDR, one or more instances of the content referencing XML Schema as defined in the present document or in TS 102 822-3-1 [10] shall be returned.

The MIME type returned by the location resolution server shall be text/xml.

EXAMPLE: One of the response lines from the HTTP server will be:

Content-Type: text/xml

12.3.7.2 Case 2: Returning an HTTP Re-direct

Use of the HTTP redirect commands (HTTP error codes 300 to 399) may be used by the location resolution server to indicate that the PDR should disconnect and connect to a different location resolution server. The reason for providing this functionality is so that a location resolution provider can redirect their resolution requests based on the CRID being resolved rather than just the authority name (which can be re-directed during the DNS lookup phase of CRID resolution).

The "location" response header shall be used to indicate where the PDR should contact.

EXAMPLE: Location: <http://www.example.com/tvaresolve>

When the PDR has been re-directed from its initial location resolution server, it should provide a HTTP v1.0 "Referer" header containing the location of the server it has been redirected from, to the new server.

12.3.7.3 Case 3: Returning a Resolving Authority Record.

In the case returning a resolving authority record, the MIME type shall be text/xml and the response shall consist of an instance of a ResolvingAuthorityRecordTable conforming to the syntax specified in clause A.1.2.

12.3.7.4 Encoding Server Response

It is permissible for the response from the server to be encoded, for example by compressing or encrypting the XML instance document. The "Content-Type" response header does not change, but rather the "Content-Encoding" gives the encoding that has been performed on the data. The exact form of encoding used is not specified in the present document and it is the responsibility of the HTTP client and server to negotiate commonly understood encoding systems.

EXAMPLE: Content-Type: text/xml
Content-Encoding: x-zip

Annex A (normative): XML Schema for Content Referencing

A.1 Schema Definition

In this clause, we define the normative content referencing schema. Instances of this schema are used during the location resolution process via bi-directional networks.

A.1.1 Location Resolution Schema

```
<?xml version="1.0" encoding="UTF-8"?>
<schema targetNamespace="urn:tva:ContentReferencing:2004"
xmlns:metadata="urn:tva:metadata:2004" xmlns="http://www.w3.org/2001/XMLSchema"
xmlns:CR="urn:tva:ContentReferencing:2004">
<element name="ContentReferencingTable" type="CR:ContentReferencingTableType">
  <annotation>
    <documentation>A document conforming to the TV Anytime content referencing
specification</documentation>
  </annotation>
</element>
<complexType name="ContentReferencingTableType">
  <sequence>
    <element name="Result" type="CR:ResultType" minOccurs="0" maxOccurs="unbounded"/>
  </sequence>
  <attribute name="version" type="float" use="required"/>
</complexType>
<simpleType name="AcquisitionDirectiveType">
  <restriction base="string">
    <enumeration value="all"/>
    <enumeration value="any"/>
  </restriction>
</simpleType>
<simpleType name="ResolutionStatusType">
  <restriction base="string">
    <enumeration value="resolved"/>
    <enumeration value="discard CRID"/>
    <enumeration value="cannot yet resolve"/>
    <enumeration value="unable to resolve"/>
  </restriction>
</simpleType>
<complexType name="ResultType">
  <choice>
    <sequence>
      <element name="CRIDResult" type="CR:CRIDResultType" minOccurs="0"
maxOccurs="unbounded"/>
    </sequence>
    <sequence>
      <element name="LocationsResult" type="CR:LocationsResultType" minOccurs="0"
maxOccurs="unbounded"/>
    </sequence>
  </choice>
  <attribute name="CRID" type="metadata:CRIDType" use="required"/>
  <attribute name="complete" type="boolean" use="required"/>
  <attribute name="acquire" type="CR:AcquisitionDirectiveType" use="required"/>
  <attribute name="status" type="CR:ResolutionStatusType" use="required"/>
  <attribute name="reresolveDate" type="dateTime" use="optional"/>
</complexType>
<complexType name="CRIDResultType">
  <sequence>
    <element name="Crid" type="metadata:CRIDType" maxOccurs="unbounded"/>
  </sequence>
</complexType>
```

```

<complexType name="LocatorType">
  <simpleContent>
    <extension base="anyURI">
      <attribute name="instanceMetadataId" type="metadata:InstanceMetadataIdType"
use="optional"/>
    </extension>
  </simpleContent>
</complexType>
<complexType name="TimeAndURLType">
  <simpleContent>
    <extension base="anyURI">
      <attribute name="start" type="dateTime" use="required"/>
      <attribute name="duration" type="duration" use="optional"/>
      <attribute name="end" type="dateTime" use="optional"/>
      <attribute name="instanceMetadataId" type="metadata:InstanceMetadataIdType"
use="optional"/>
    </extension>
  </simpleContent>
</complexType>
<complexType name="LocationsResultType">
  <sequence maxOccurs="unbounded">
    <choice>
      <element name="Locator" type="CR:LocatorType"/>
      <element name="DecomposedLocator" type="CR:TimeAndURLType"/>
    </choice>
  </sequence>
</complexType>
</schema>

```

Name	Definition
ContentReferencingTable	The top level element within which all content referencing results are instantiated.
ContentReferencingTableType	The syntax definition for the ContentReferencingTable element.
Result	The content referencing information for each CRID being resolved are contained within this element.
Version	The syntax version for this XML Schema. For instances conforming to the Schema defined in the present document, this field shall contain the value 1.0.
ResultType	This type provides the container for all possible resolutions of a CRID.
CRID	The CRID that is being resolved.
Complete	True if the resolution of this CRID is complete. If false, the CRID may resolve into further CRIDs or locators in the future.
Acquire	The grouping type for the list or CRIDs or Locators. This field is only meaningful when status equals resolved.
status	The status of resolving this CRID.
rerresolveDate	If status equals "cannot yet resolve", Or, Complete equals false, this field contains the date and time when re-resolution should next be attempted. This date and time shall be unambiguous with respect to time zone.
CRIDResult	An element for instantiating a result that represents a CRID resolving into one or more CRIDs.
LocationsResult	An element for instantiating a result that represents a CRID resolving into one or more locations.
CRIDResultType	When a CRID resolves into one or more CRIDs, an instance of CRIDResultType shall be used.
Crid	One of the "output" CRIDs.
LocatorType	This extension to the uriReference type holds the URL pointing to the content and has an optional attribute that contains the instance metadata identifier.

Name	Definition
instanceMetadataID	An attribute of Locator to provide a binding reference to the Instance Description Metadata, clause 10 Instance Metadata Identifiers".
TimeAndURLType	This extension to the uriReference type holds the URL pointing to the content and has attributes that contain the timing information required for acquisition.
start	If this is scheduled content, the date and time when the content will start. For on demand content, this is the time and date at which the content is first available.
duration	For scheduled content, this is the duration of the content. This element shall not be used for on demand content.
End	For scheduled content, this attribute shall not be used. For on demand content, this attribute contains the first time and date when the content is no longer available.
LocationsResultType	When a CRID resolves into one or more locators, an instance of LocationsResultType shall be used.
Locator	One of the "output" locators. This element is used when the locator contains the necessary timing information (such as start time, duration) as part of its syntax.
InstanceMetadataID	An attribute of Locator to provide a binding reference to the Instance Description Metadata (see Instance Metadata Identifiers).
DecomposedLocator	One of the "output" locators. This element is used when the locator does not contain the necessary timing information (such as start time, duration) as part of its syntax.
AcquisitionDirectiveType	When a CRID resolves into a list of CRIDs or locators, the AcquisitionDirectiveType type describes what sort of group this list represents.
All	All items in the list must be acquired. All items in the list combine to create the entire content assigned to the CRID being resolved.
Any	One of the items from the list should be acquired. All items in the list are equivalent.
ResolutionStatusType	Type indicating the result of a resolution request.
resolved	The CRID has been successfully resolved.
discard CRID	The CRID should be discarded.
cannot yet resolve	The CRID cannot be resolved yet, and should be resolved again at a later date.
unable to resolve	The CRID cannot be resolved.

A.1.2 Resolving Authority Record Schema

```
<?xml version="1.0" encoding="UTF-8"?>
<schema targetNamespace="urn:tva:ResolvingAuthority:2004"
xmlns="http://www.w3.org/2001/XMLSchema" xmlns:RAR="urn:tva:ResolvingAuthority:2004"
elementFormDefault="qualified">
  <element name="ResolvingAuthorityRecordTable"
type="RAR:ResolvingAuthorityRecordTableType">
    <annotation>
      <documentation>A document conforming to the TV Anytime content referencing
specification</documentation>
    </annotation>
  </element>
  <complexType name="ResolvingAuthorityRecordTableType">
    <sequence>
      <element name="ResolvingAuthorityRecord" type="RAR:ResolvingAuthorityRecordType"
minOccurs="1" maxOccurs="unbounded"/>
    </sequence>
  </complexType>
  <simpleType name="ProviderClassType">
    <restriction base="string">
      <enumeration value="primary"/>
      <enumeration value="secondary"/>
    </restriction>
  </simpleType>
</schema>
```

```

</simpleType>
<complexType name="ResolvingAuthorityRecordType">
  <sequence>
    <element name="ResolutionProvider" type="string"/>
    <element name="AuthorityName" type="string"/>
    <element name="Class" type="RAR:ProviderClassType"/>
    <element name="VersionNumber" type="unsignedLong"/>
    <element name="URL" type="anyURI"/>
    <element name="FirstValidDate" type="dateTime"/>
    <element name="LastValidDate" type="dateTime"/>
    <element name="Weighting" type="integer"/>
  </sequence>
</complexType>
</schema>

```

Name	Definition
ResolvingAuthorityRecordTable	The top level element within which the RARs are instantiated.
ResolvingAuthorityRecord	The element within which the RAR information is instantiated.
ResolutionProvider	The string containing the name of the provider of this resolution service.
AuthorityName	The string containing the name of the CRID authority that is being resolved by this resolution service.
Class	The class of this resolution service (primary or secondary).
URL	The location from which resolution information is available.
FirstValidDate	The first time and date at which this RAR becomes valid.
LastValidDate	The first time and date at which this RAR becomes invalid.
Weighting	The weighting of this RAR record with respect to the other RAR records provided by this authority for the given CRID authority.

A.2 Example Instance Documents

The following is an example instance document conforming to the location resolution XML Schema specified in the present document.

```

<?xml version="1.0" encoding="UTF-8"?>
<ContentReferencingTable version="1.0" xmlns="urn:tva:ContentReferencing:2004">
  <!-- Example of a CRID resolving to other CRIDs -->
  <Result CRID="crid://broadcaster.co.uk/akdsjdkjdf" status="resolved" complete="true"
  acquire="all">
    <CRIDResult>
      <Crid>CRID://example.com/greatstuff</Crid>
      <Crid>CRID://nextcrid.com/lkj kj</Crid>
    </CRIDResult>
  </Result>
  <!-- Example of a CRID that is no longer valid -->
  <Result CRID="crid://isp.net/868457549845f" status="discard CRID" complete="true"
  acquire="all"/>
  <!-- Example of a CRID resolving to other CRIDs and is incomplete -->
  <Result CRID="crid://example.co.uk/wibble" status="resolved" complete="false"
  acquire="all">
    <CRIDResult>
      <Crid>CRID://example.com/stuff</Crid>
      <Crid>CRID://nextcrid.com/broodje</Crid>
    </CRIDResult>
  </Result>
  <!-- Example of a CRID resolving to locators -->

```

```

<Result CRID="crid://broadcaster.com/ajcnd" status="resolved" complete="true"
acquire="all">
  <LocationsResult>
    <Locator instanceMetadataId="imi:1">dvb://1.4ee2.3f4;4f5@2001-03-
27T18:00:00.00+01:00</Locator>
    <DecomposedLocator instanceMetadataId="imi:metadataProv.com/2" start="2001-03-
29T18:00:00.00">ftp://myserver.example.com/directory12/hello.mp3</DecomposedLocator>
    <DecomposedLocator start="2001-03-29T18:00:00.00" end="2001-04-
03T18:00:00.00">ftp://myserver.example.com/directory12-
backup/hello.mp3</DecomposedLocator>
  </LocationsResult>
</Result>
</ContentReferencingTable>

```

The following is an example instance document conforming to the RAR XML Schema specified in the present document.

```

<?xml version="1.0" encoding="UTF-8" ?>
<ResolvingAuthorityRecordTable xmlns="urn:tva:ResolvingAuthority:2004">
  <ResolvingAuthorityRecord>
    <ResolutionProvider>autnam.com</ResolutionProvider>
    <AuthorityName>autnam.com</AuthorityName>
    <Class>primary</Class>
    <VersionNumber>1000</VersionNumber>
    <URL>http://www.autnam.com/lr/</URL>
    <FirstValidDate>2000-09-06T09:30:00Z</FirstValidDate>
    <LastValidDate>2000-09-28T18:00:00Z</LastValidDate>
    <Weighting>1</Weighting>
  </ResolvingAuthorityRecord>
  <ResolvingAuthorityRecord>
    <ResolutionProvider>tva.resprov.com</ResolutionProvider>
    <AuthorityName>autnam.com</AuthorityName>
    <Class>secondary</Class>
    <VersionNumber>1000</VersionNumber>
    <URL>http://www.resprov.com/lr/autnam</URL>
    <FirstValidDate>2002-09-26T09:30:00Z</FirstValidDate>
    <LastValidDate>2002-10-28T18:00:00Z</LastValidDate>
    <Weighting>3</Weighting>
  </ResolvingAuthorityRecord>
</ResolvingAuthorityRecordTable>

```

Annex B (informative): Example dynamic communication behaviour between a PDR and a Remote Location Resolution Server

The following is an informative description of how the dynamic communication behaviour between a PDR and a remote Location Resolution Server may be implemented on the PDR.

If the response to location request is "resolve again after a given date and time":

- 1) If the date and/or time is in the future:
Wait until this date and time arrives. Wait for a further random amount of time before contacting the server.
- 2) Else If the date and time has already past:
Wait for a random amount of time before contacting the server

If the result from contacting the server again is that the re-resolve date and time is still in the past, double the range of the random delay from its current range and try again.

When delay \geq 1 day, the PDR can either:

- a) Try each day.
- b) Continue doubling delay until delay \geq one week, at which point the delay stays fixed at one week.

Else If the LR server is unavailable:

Wait a random amount of time and then try again. If server is still unavailable, double the range of random delay from its current range and try again.

When delay \geq 1 day, the PDR can either:

- a) Try each day.
- b) Continue doubling delay until delay \geq one week, at which point the delay stays fixed at one week.

Else If response is "CRID is unknown" and server type = secondary:

Wait for a random amount of time before contacting the server. If the result from contacting the server again is that the CRID is still unknown, double the range of the delay from its current range and try again.

When delay \geq 1 day, the PDR can either:

- a) Try each day.
- b) Continue doubling delay until delay \geq one week, at which point the delay stays fixed at one week.

Else If response is "CRID is unknown" and server type = primary:

- Either: a) Assume CRID is invalid (as if LR server returns CRID = invalid).
- Or: b) Follow delay system above. After one day assume CRID is invalid.

Annex C (informative): Bibliography

- ETSI TS 102 822-5: "Broadcast and On-line Services: Search, select, acquisition, and rightful use of content - on personal storage systems ("*TV-Anytime* Phase 1"); Part 5: Rights management".

List of figures

Figure 5.1: Content Referencing Environment.....	9
Figure 8.1: Example of a Tree-Structured CRID	12
Figure 12.1: CRID Resolution Architecture.....	17
Figure 12.2: Dynamic Behaviour of Resolution Handler using Unidirectional Network	18
Figure 12.3: Example of Location Resolution Server Discovery	21
Figure 12.4: Stages of TCP/IP Based Resolution Server Discovery	24

List of tables

Table 8.1: Examples of CRID	12
Table 10.1: Examples of Instance Metadata Identifiers	14
Table 10.2: Example of a resolved CRID.....	15
Table 10.3: Example of the same resolved CRID after a change in location identified by "imi:def.com/1"	15
Table 10.4: Example of an Instance Metadata Identifier with no name portion.....	15
Table 10.5: Example of a more complex resolution table	15
Table 11.1: An example Resolving Authority Record	17
Table 12.1: Location Resolution Message Format for Unidirectional System.....	19
Table 12.2: Interpretation of Status Flags	20
Table 12.3: PDR behaviour in response to acquisition directive.....	20
Table 12.4: Submitted CRID Flag Description	21
Table 12.5: Result Flag Description.....	22
Table 12.6: Key definitions for HTTP URL encoding	25

History

Document history		
V1.1.1	October 2003	Publication
V1.1.2	October 2004	Publication