
3.1 Introduction

3.1.1 Referencing Contents Unambiguously

As stated in the previous chapter, the concept of CRID is an element of prime importance in the TV-Anytime universe we are describing. In order to implement services with a high added value such as those presented in Chap. 1, it is vital to be able to reference contents unambiguously, regardless of their location, that is without knowing specific broadcast information (time, date and channel) or how to obtain them through a network, for instance, by means of a streaming service or by downloading a file from an Internet server.

Evidently, the receiver must be able to resolve these unambiguous references, that is translate them into specific data that will allow it, when the time comes, to obtain the location of that content in order to acquire it (limiting the user's intervention to a hypothetical selection among similar choices).

This would make it possible for recording processes (i.e. content-acquisition processes in general) to take place without knowing that information and even without knowing beforehand the duration or nature of the content to be recorded. Several scenarios can be pictured in which this capacity becomes vital to provide advanced features such as those that the standard intends to promote:

- The user sees the advertisement of a film that will air soon. The existence of a reference for that film would allow the user to schedule its recording immediately, even if the operator is still unaware of the broadcast information (such as time, date and channel). To achieve this, the operator would simply have to assign an unambiguous reference that would be broadcast as metadata along with the advertisement. The PDR would put this reference under watch. Once the operator has decided the specific broadcast information for the film, it would make it public and this would allow the PDR to resolve the reference into specific location information and schedule the recording.

- The user sees the advertisement of a TV series that will air soon (the date has not been determined yet) and wants the receiver to record every episode. In this case, it is even possible that the operator is unaware of the specific number of episodes to air (maybe they have not even been filmed yet), but it wishes to advertise it as soon as possible. The operator would only have to assign a generic (group) reference to the series, and the receiver would put it under watch. Once the information was available, the operator would make public the specific episode list for the series, providing unambiguous references for each of them. These references would then be under the receiver's watch. Finally, once it was known when and where the episodes air, the operator would make public the necessary information to resolve each of those references into the corresponding location information (such as time, date and channel, which may vary from episode to episode).
- The operator could be airing music videos of different genres, without the complete composition ever being known, given the fact that the selection is made on the way, day by day, while the programme lasts. In order for the user to secure the recording of videos of a given genre, the operator would only have to assign a group reference to that genre. The user would then order the receiver to record all contents that are described as belonging to that group. In this case, at no time does the operator make public a list of references or any location information; conversely, it is an ongoing dynamic process in which the receiver will be constantly monitoring the contents looking for those that belong to the group the user is interested in.

That is, the user could secure the recording of contents with little action on his part: upon seeing an advertisement for the contents, he/she would only have to indicate to the PDR that he/she wants the contents recorded. The rest of the process will be carried out with no user intervention, and he/she can just relax and let the PDR do the work.

The purpose of this chapter is, precisely, to present the way TV-Anytime defines to identify contents regardless of their location or availability, find out what the location information necessary for their due acquisition consists of and, finally, give details about the location resolution process, by means of which it is possible to find out the necessary information to acquire a given content from its reference.

3.1.2 CRID

In the TV-Anytime standard, this unambiguous reference to a given content is called CRID (*Content Reference IDentifier*), and it allows for the separation between the reference to a given content (the CRID) and the necessary information to acquire it, which is called a "locator". This separation into two information elements allows a one-to-many association relation between CRIDs and locators, that is each CRID may lead to one or more locators which will represent different copies of the same content. They may be identical copies that air on different

channels and dates or cost different prices. They may also be distinct copies with different technical parameters such as format or quality.

It may also be the case that the resolution process of a CRID provides another CRID as a result (e.g. its reference in a different network, where it has an alternative identifier assigned by a different operator) or a set of CRIDs (for instance, if the original CRID represents a TV series, in which case the resolution process would result in the list of CRIDs representing each episode).

From the above it is obvious, provided that a given content can belong to many groups (each possibly defined by distinctive qualities), that many CRIDs can lead to the same content, even within the scope of the same authority that created the CRID. In other words, several CRIDs may be resolved into the same locator.

From the perspective of a TV-Anytime system, a CRID is the result of a search and selection process, in which the user ends up choosing the content, of which only a reference (the CRID) is known. Then, the resolution process starts (see Fig. 3.1), which finishes when the receiver has gathered all the necessary information to acquire the content (i.e. a locator). Therefore, the resolution process itself involves the central part of Fig. 3.1, excluding both the content search and acquisition processes.

In this figure, it can be seen that the resolution process receives a CRID as input and may result in one or many CRIDs (e.g. if the original CRID represents a content group) or one or several locators, each of which includes the necessary information to obtain a specific copy of the content.

If several CRIDs are obtained, these would be sent back to the content selection process so that one can be chosen, which would be returned to the resolution process. If several locators are returned, one of them would be chosen to acquire the content according to the service's setup criteria.

As suggested in Fig. 3.1, this location resolution process can be autonomous or involve some interaction with the user or a receiver agent, for example, to choose among alternatives or sub-elements represented by several CRIDs or locators (choosing broadcast time, choosing among different qualities or different costs, etc.).

3.1.3 References Versus Identifiers

It is important to clarify that a CRID (a reference) is not exactly a universal, unique and exclusive identifier for a given content.

A CRID is a code designed to make it easier to locate a given content, meaning it contains or leads to the necessary information for the user to acquire the content once it has been chosen. A CRID is closely related to the authority that creates it, to the resolution service provider and to the content provider in such a way that the same content may have different CRIDs depending on the field in which they are used (e.g. a different CRID for each television operator that has the rights to air the content).

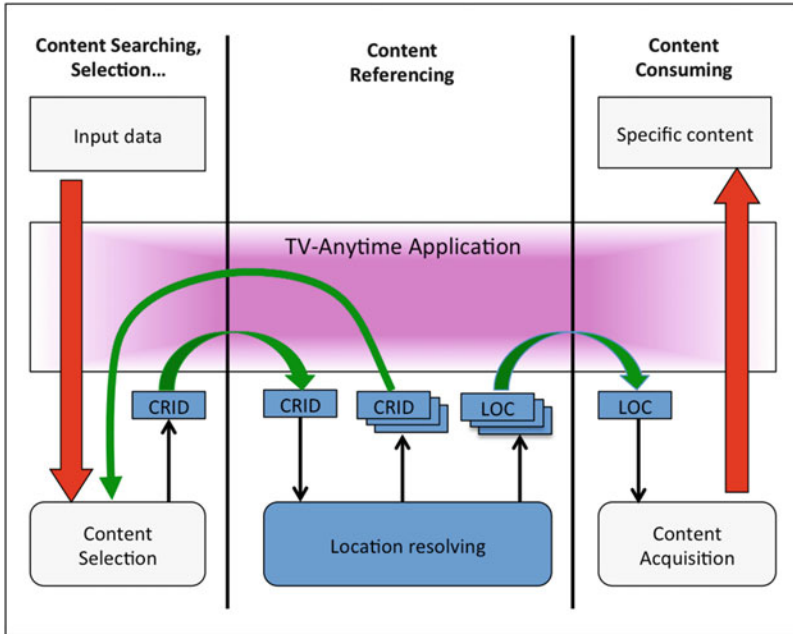


Fig. 3.1 Location resolution process

On the other hand, the identifier is a unique key which is normally assigned after the content is created so that it goes with it all the time, representing it. From this viewpoint, the identifier must be the same for a given content, regardless of the operator, aggregator or third parties involved in the commercialization. This mechanism would call for a central and global identifier resolution service aware of all the copies of the content, which would be especially difficult to agree upon and implement among all the participants in this industry.

A model based on references such as the one defined in TV-Anytime allows each entity to manage, reference and resolve the location of its contents autonomously, thus enabling the availability of the content catalogue of each operator immediately, with no business-to-business agreements.

3.1.4 Roles in the Creation and Resolution of CRIDs

As stated before, the creator of a CRID is called an *authority*, and it is the entity in charge of making sure there are no ambiguities within its resolution scope.

For all practical purposes, the authority that creates the CRID and the resolution service provider that provides the location information must be clearly distinguished. Although both tasks (CRID creation and resolution) may be performed by the same entity, they are essentially different tasks and may be carried out by different agents.

In this context, there are basically three different entities that may play a role: the content creator, the operator that aggregates the contents and distributes them and third parties that provide value-added services (e.g. a customized EPG with content recommendation). Although any of these can functionally carry out both tasks (assignment and resolution), it is true that some of the combinations are more natural and, therefore, they are expected to be more usual.

For example:

- If the content creator is the one who assigns the CRID, the resolution service is feasible and reasonable for any of the three agents. In this case, the operator and the third party would be acting as proxies for the content creator, for example, if the third party gathered information about the different operators that distribute the content.
- If it is the operator/aggregator/distributor who assigns the CRID, it can be solved naturally by the same agent and by a third party (which would be working as proxy for the operator), but it is not clear if this task may be carried out easily by the content creator.
- Conversely, if it is a third party that creates the CRID, it seems natural that this party itself resolves it. It does not seem to be appropriate for the operator to do it, and it would be very difficult for the content creator to do it.

3.2 Reference and Locator Formats

3.2.1 CRID Format

In order to secure the unambiguous identification of authorities which assign CRIDs to contents, the format of the identifier that makes reference to them is that of a string of characters according to the name system used by the Internet to identify machines instead of their IP address, the well-known *Domain Name System*, or DNS (IETF RFC 1591, 1994). According to it, valid authority names are:

www.national-geographic.com
www.telefonica.com
www.provider-z.com

From this, the complete syntax of a CRID according to what is established in the TV-Anytime standard (ETSI TS 102 822–4, 2011) is the following, where each and every part is not case sensitive:

crid://<authority>/<data>

The <authority> field represents the entity that created the CRID, and its format is that of a DNS name, as explained above. The <data> field represents a string of characters that will unambiguously identify the content within the authority scope (it is a string of characters assigned by the authority itself). For instance:

crid://www.national-geographic.com/sahara-nights
crid://www.telefonica.com/335-network-documentary
crid://www.provider-z.com/quizshow-23445-questions

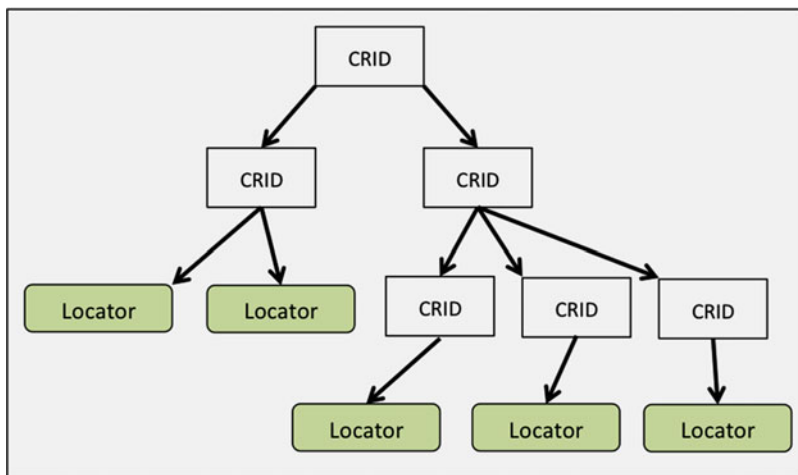


Fig. 3.2 CRID hierarchical resolution

As mentioned before, a CRID can be resolved into one or more locators or into one or more CRIDs, for example, if a CRID represents a group of contents or if an entity wants to create a CRID that makes reference to a content managed by another entity, for instance, by making reference to the CRID that the content has in that other entity.

In any event, this is a hierarchical resolution system that after many repetitions must always end in one or more content locators (Fig. 3.2).

3.2.2 Locator Format

A locator is a string of characters that contains all the necessary information for a receiver to find and acquire a given content, whether it is received through a transport stream, located in the local storage, or downloaded as a file from an Internet server or through a streaming service. For example, a DVB locator will include all the necessary parameters to identify a specific content within a transport stream: network identifier, transport stream identifier, service identifier, table identifier and/or event identifier.

Apart from the corresponding information to know how to access the content, it is also common for the locator to include information about temporal availability, which will not have time zone ambiguity, since a device can receive contents from places in different time zones.

The locator's format, as established in TV-Anytime (ETSI TS 102 822-3-4, 2011), is quite generic and simple and corresponds to:

<transport-mechanism>:<specific-data>

In the receiver's acquisition mechanism, it is usual for a specific system agent to be selected based on the locator. This agent is a manager for the corresponding "transport mechanism" that knows the details of the protocol to access information depending on the channel through which it is available and provides access to the content, hiding the communication complexities.

The first part of the locator's format (the identifier that describes the transport mechanism) must be a string of characters that is unique for each mechanism (transport stream, local file, HTTP Internet access, etc.), excluding the "CRID" string. The second part, on the other hand, must be unambiguous only within the scope of a given transport mechanism and will be standardized by the institution in charge of the regulation of the mechanism itself. In any case, the format of these *<specific-data>* will be unique within each mechanism. Many providers may share the same transport mechanism to provide users with their contents, and the locators must be absolutely clear about that.

If the locator includes specific information about the content's time availability, there are two options:

- **Planned broadcast.** Normally, the locator specifies the starting time and duration of the content.
- **On-demand access.** The content may be acquired within a specific time frame, and normally the locator indicates the beginning and end of this frame.

Some examples of valid locators are:

- In the case of an AVI format video ("*example.avi*") stored in the directory ("*videos*") of a server ("*www.server.com*") accessible through the "*ftp*" protocol: <ftp://www.server.com/videos/example.avi>
- In the case of an MHP object carousel within a transport stream:
mhp://65.1000.1;1.0.1 ~ 20111025T214000 + 0900
- And in the most frequent case that we will find in this book, a DVB locator to identify contents within the transport stream of networks that follow this standard (ETSI TS 102 851, 2010):
dvb://112.4a2.5ec;2d22 ~ 20111012T220000Z—PT01H30M
which would indicate a television show (identified by the string of characters "2d22") that airs on a channel available on a DVB distribution network identified by the address "112.4a2.5ec" (network identifier "112", transport stream identifier "4a2" and television service identifier "5ec"), on OCT/12/2011 at 10 p.m. and with a duration of 90 min.

3.2.3 Instance Metadata Identifier (IMI)

A CRID represents all the variations of a given content, regardless of certain temporary features such as location information, cost, format or encoding details. However, there are times when it may be advisable to have a mechanism to reference a given copy, precisely because we are interested in some of its particular features (format, quality, audio encoding, etc.).

For these kind of situations, there is an element of Instance Metadata Descriptions that makes it possible to assign a different identifier for each of those variations. This identifier will be unique only in the context of a given CRID.

This identifier is called *IMI* (*Instance Metadata Identifier*), which makes it possible to unambiguously identify each of the different copies of a given content which, for example, may differ in some of the abovementioned attributes. If the content's location changes, its locator will also change, but its *IMI* should stay the same.

This allows a receiver to automatically track the specific copy that the user has selected and to acquire it, even if there are any changes in the results of the location resolution process that may create ambiguity.

The format of an *IMI* is as follows:

imi : [*<name>* /] *<data>*

The *<name>* field represents the entity that creates the identifier and follows the format of a DNS name. Its inclusion is optional, and if it is not there, it is because it matches the name of the authority that created the CRID this *IMI* is associated to. The *<data>* field is a string of characters that is significant for the entity represented by *<name>*.

Some examples of correct *IMI* identifiers are:

imi:citibank.es/loans

imi:www.repsol.com/oil

imi:president

Let us assume, for example, that the user has decided to acquire a given programme *P* whose CRID is *C*. We will suppose that the receiver resolves the CRID *C* resulting in two different locators *L1* and *L2*, which differ in the quality of the video image (e.g. standard definition vs. high definition):

L1 = *transport:channel2 ~ 16:00*

L2 = *transport:channel2 ~ 18:00*

The user, probably because of his/her preferences, chooses the first one and the PDR schedules the recording. When the day of broadcast is approaching, the PDR resolves the CRID *C* again (to make sure nothing has changed) and finds changes in the airing time of both; so now locators *L3* and *L4* appear:

L3 = *transport:channel2 ~ 21:00*

L4 = *transport:channel2 ~ 23:00*

The problem this poses is that the PDR, without any help, is not capable of knowing with certainty which of the new locators (*L3* or *L4*) corresponds with the one that the user scheduled to be recorded at the time (i.e. *L1*).

In order to solve this problem, each locator may be completed with the *IMI* string of characters to distinguish among the different copies of the same content. These *IMIs* will not change, even if the original locators did.

In other words, the location resolution process would give as result in the first case:

L1 = *transport:channel2 ~ 16:00; imi:rtve.com/prog-sd*

L2 = *transport:channel2 ~ 18:00; imi:rtve.com/prog-hd*

And in the second case:

L3 = *transport:channel2 ~ 21:00; imi:rtve.com/prog-hd*

L4 = *transport:channel2 ~ 23:00; imi:rtve.com/prog-sd*

Now, the PDR knows that the copy the user had chosen (*LI*) is the one that can be acquired with the new *L4* locator, since this is the one that contains the *IMI* that was also present in the original locator.

3.3 Relevant Information in the Resolution Process

The location resolution process is a procedure by which, starting from the CRID of a given content, one or several locators of that content are obtained. As described before, resolving a CRID can be a direct process, which leads immediately to one or many locators, or it may also happen that in the first place one or many intermediate CRIDs are returned, which must undergo the same procedure to finally obtain one or several locators.

This procedure, described in detail in Sect. 3.4, involves many information elements, among which we find two data tables or structures, named *Resolving Authority Record* (RAR) and *ContentReferencingTable*, respectively. Consulting them repeatedly will take us from a CRID to one or many locators that will allow us to acquire the content.

3.3.1 Resolving Authority Record

The *Resolving Authority Record* (RAR) table is one or many data structures that provide us, for each authority that submits CRIDs, information on the corresponding resolution service provider. Among other things, it also informs about which mechanism is used to provide information to resolve the CRIDs from each authority (an example is shown in Fig. 3.3).

Each authority must have one or many RAR records that indicate the PDR where it has to go to resolve the CRIDs of that particular authority.

For example, in the first record of Fig. 3.3, there is an authority called “*tve.es*”, whose resolution service provider is the entity “*rtve.es*”, available on the URL “*dvb://1.2ef.3f5*”, which means there is resolution information¹ in the DVB transport stream identified with that URL.

Each RAR record must reach the PDR somehow, possibly depending on the transport mechanism, which is not standardized by TV-Anytime.

TV-Anytime specifications do not regulate either a representation and transport format for RAR records in unidirectional contexts. They simply describe the information fields that must be present in those records, which are the following, shown in Fig. 3.4:

- *Authority name*. The authority that created the CRID whose information appears in this record.

¹ This information is surely the second table we talked about that will be described in Sect. 3.3.2.

Field Name	RAR A	RAR B	RAR C
Authority name	tve.es	media.com	media.com
Resolution provider	rtve.es	media.com	a3tv.es
URL	dvb://1.2eef.3f5	http://media.com/lr	dvb://1.104.e5f
Version	4	2	6
Other			

Fig. 3.3 Sample RAR table

Field Name	Content
Authority name	tve.es
Resolution provider	rtve.es
Class	Secondary class
Version	4
URL	http://tva.rtve.es/locres/tve
First valid date	9:30 am October 12th 2011
Last valid date	6:00 pm October 28th 2011
Weighting	1

Fig. 3.4 Fields in a RAR record

- *Resolution provider*. The name of the entity that provides the resolution service for that authority.
- *Class*. This field indicates whether this provider can resolve all CRIDs for this authority (*class = primary*) or only some of them (*class = secondary*).
- *Version number*. A version number for the RAR record, which increases every time the record is updated.
- *URL*. The location of the site where the resolution process will take place. It can be, for instance, an Internet service (in bidirectional contexts) or a table within a transport stream (e.g. in unidirectional contexts).
- *First valid date*. The first moment when this information can be used.
- *Last valid date*. The last moment when this information can be used.
- *Weighting*. A number that will be used to assign a priority to each RAR record in case there are many entries for the same authority and the same location resolution service provider. Only in this case (same authority and same resolution provider) may this field be used to decide which one will be attempted first.

And Listing 3.1 shows a possible encoding for a RAR record in an XML structure, according to the *XML Schema* defined in the TV-Anytime standard (ETSI TS 102 822–4, 2011).

```

1 <ResolvingAuthorityRecordTable xmlns="urn:tva:ResolvingAuthority:2008">
2   <ResolvingAuthorityRecord>
3     <AuthorityName> media.es </AuthorityName>
4     <ResolutionProvider> media.es </ResolutionProvider>
5     <Class> primary </Class>
6     <VersionNumber> 1000 </VersionNumber>
7     <URL> http://www.media.es/lr </URL>
8     <FirstValidDate> 2011-08-16T00:00:00Z </FirstValidDate>
9     <LastValidDate> 2011-11-28T23:59:59Z </LastValidDate>
10    <Weighting> 1 </Weighting>
11  </ResolvingAuthorityRecord>
12  <ResolvingAuthorityRecord>
13    <AuthorityName> media.es </AuthorityName>
14    <ResolutionProvider> a3tv.es </ResolutionProvider>
15    <Class> secondary </Class>
16    <VersionNumber> 1000 </VersionNumber>
17    <URL> dvb://112.4a2.5ec </URL>
18    <FirstValidDate> 2011-07-21T00:00:00Z </FirstValidDate>
19    <LastValidDate> 2011-11-17T12:00:00Z </LastValidDate>
20    <Weighting> 3 </Weighting>
21  </ResolvingAuthorityRecord>
22 </ResolvingAuthorityRecordTable>

```

Listing 3.1 RAR table in XML format

In this listing, two entries are included to resolve the CRIDs of the “*media.es*” authority (lines 3 and 13), which possibly represents a content creator. The first one declares a resolution service of the authority itself (line 4), valid for all of its CRIDs because the *class* is characterized as *primary* (line 5) and available on the Internet URL “<http://www.media.es/lr>” (line 7). The second one indicates that an operator, entity “*a3tv.es*” (line 14), also resolves some of its CRIDs (secondary class, line 15) through the information broadcast in the transport stream “*dvb://112.4a2.5ec*” (line 17).

3.3.2 ContentReferencingTable

The second structure involved in the location resolution process is a proper resolution table which, given a content’s CRID, returns one or several locators that enable us to access an instance of that content, or to one or many CRIDs that allow us to move forward in the resolution process.

Listing 3.2 shows an example of this second structure, an XML document according to the specifications of the *XML Schema* defined in the corresponding TV-Anytime document (ETSI TS 102 822–4, 2011). In it, many different sections are included (<*Result*> elements) that structure the information that describes each resolution case:

- The first one (line 2) declares how a CRID (“*crid://C4.com/House/all*”), which corresponds to a group content that encompasses several episodes of the *House M.D.* series (three in this case, to limit the length of the example), is resolved. The result of the resolution process of this kind is a <*CRIDResult*> element

```

1 <ContentReferencingTable version="1.0">
2   <Result CRID="crid://C4.com/House/all"
3     status="resolved" complete="false" acquire="all">
4     <CRIDResult>
5       <Crid>crid://C4.com/House/S1/E1</Crid>
6       <Crid>crid://C4.com/House/S1/E2</Crid>
7       <Crid>crid://C4.com/House/S1/E3</Crid>
8     </CRIDResult>
9   </Result>
10  <Result CRID="crid://C4.com/House/S1/E0"
11    status="discard CRID" complete="true" acquire="all">
12  </Result>
13  <Result CRID="crid://C4.com/House/S1/E1"
14    status="resolved" complete="true" acquire="any">
15    <LocationsResult>
16      <Locator>
17        dvb://1.4ee2.3f4;4f5~@2011-10-12T22:00:00.00+01:00/PT01H00M
18      </Locator>
19      <Locator>
20        dvb://1.4ee2.3f4;4f6~@2011-10-19T23:00:00.00+01:00/PT01H00M
21      </Locator>
22    </LocationsResult>
23  </Result>
24  <Result CRID="crid://C4.com/House/S1/E2"
25    status="resolved" complete="false" acquire="any">
26    <LocationsResult>
27      <Locator instanceMetadataId="imi:C4.com/House/sd">
28        dvb://1.4ee2.3f4;4f5~@2011-10-19T22:00:00.00+01:00/PT01H00M
29      </Locator>
30      <Locator instanceMetadataId="imi:C4.com/House/hd">
31        dvb://1.5a2.2bc;2d3~@2011-10-20T22:00:00.00+01:00/PT01H00M
32      </Locator>
33    </LocationsResult>
34  </Result>
35  <Result CRID="crid://C4.com/House/S1/E3" status="cannot yet resolve"
36    complete="false" acquire="all"
37    reresolveDate="2011-10-20T00:00:00.00+01:00">
38  </Result>
39 </ContentReferencingTable>

```

Listing 3.2 Sample *ContentReferencingTable*

(line 4) that provides three new CRIDs, each corresponding to one of the three episodes: “*crid://C4.com/House/S1/E1*”, “*crid://C4.com/House/S1/E2*” and “*crid://C4.com/House/S1/E3*” (on lines 5, 6 and 7, respectively).

- The second *<Result>* element (line 10) resolves the CRID for Episode 0, possibly the pilot. Since this episode is no longer available, the result for the “*status*” attribute is a “*discard CRID*” value (line 11) that concludes the resolution process.
- The third *<Result>* element (line 13) resolves the CRID of the first episode. The result of the resolution process is two DVB locators (the content of the two *<Locator>* elements within the *<LocationsResult>* element in lines 17 and 20). The “*acquire*” attribute with “*any*” value in line 14 indicates that any of them are good, which means that the second one is a repetition aired a week later, after the second episode of the first season.

- The fourth *<Result>* element (line 24) resolves the CRID of the second episode of the first season, providing two DVB locators (lines 28 and 31) that indicate two channels (with different times) where that episode is aired. Each locator is accompanied by an *IMI* which describes it (distinguishing the standard definition version from the high definition one). The “*complete*” attribute with the “*false*” value in line 25 indicates that this CRID will be resolved in the future into more locators (probable future repetitions of the episode).
- The last *<Result>* element (line 35) gives information about the third episode. It indicates that it cannot be resolved yet (“*status*” attribute with the “*cannot yet resolve*” value on line 35), indicating a date on which the request for resolution information must be repeated (line 37).

3.4 Location Resolution Process

Once the user has selected a given content (identified by the corresponding CRID) to perform some action upon it, the receiver begins the location resolution process that shall lead to specific location information that allows access to a copy of the content.

This procedure depends mainly on the receiver’s connectivity. It is possible to make a basic distinction between unidirectional networks, where the receiver can only receive information through the broadcast channel, and bidirectional networks, where there is also a return channel through which the receiver can communicate with the outside world (typically an Internet access).

With more or less variety depending on the type of connectivity in hand, there will be various means to deliver the information, each with their corresponding protocols: file system access, transport stream in DVB networks, direct Internet access, etc. The receiver, in any case, must have a set of resolution handlers, software entities that know the access details through the corresponding protocol, among which it will choose the most appropriate one for carrying out the resolution process based on the information included in the CRID.

3.4.1 Resolution in Unidirectional Networks

For receivers connected only to one broadcast channel, it is clear that the resolution information must come directly from that channel or be available somehow in an existing local storage system.

After selecting a CRID, the first thing the receiver needs to do is check the information about where to find the resolution table. For this, it must find a *RAR* record associated with the authority of the selected CRID. These *RAR* records will have reached the receiver in an indefinite form, unimportant for the TV-Anytime specification, which will depend on the specific transport mechanism of the network to which the receiver is connected. Each family of standards that regulates distribution networks (DVB, ATSC, ISDB, IPTV, etc.) will have previously defined such a procedure, which will be used by devices certified according to those standards.

If no *RAR* record corresponds to the CRID's authority, the resolution process will be over, obviously failing to achieve its goal.

Once a *RAR* record corresponding to that authority is found, the receiver will know, by referring to the *URL* field, where to access (or, in this case, where to listen) the resolution information,² along with the most appropriate resolution handler for the corresponding access protocol.

The information the handler will receive through that access point will consist of a message for each of the consulted CRIDs (e.g. a *<Result>* element in the *ContentReferencingTable* of Listing 3.2). That message will begin with a field called *Status*, which will be an indicator of the type of response and which could take the following values:

- “*discard CRID*”. Unable to resolve the CRID (e.g. line 11 of Listing 3.2).
- “*resolve after date<xxx>*”. The resolution information is still not available, and a retry must take place after the date *<xxx>* (e.g. lines 35 and 37 of Listing 3.2).
- “*CRID is resolved*”. The resolution is successful and is accompanied by information on the result, which is distributed into the fields that follow the current one:
 - “*Acquisition directive*”. It could take the values “*all*” (if all the elements that follow should be captured because they are part of the content—line 3 of Listing 3.2) or “*any*” (if it is only necessary to capture one of them, because they are alternative elements—line 14 of Listing 3.2).
 - “*List CRIDs/Locators*”. A list of results of the resolution process, whether they are new CRIDs or locators (as is the case of the *<LocationsResult>* element on line 15 of Listing 3.2).
 - “*Resolution complete*”. An indicator that can take the value “*yes*” (if the previous list contains all possible results—line 14 of Listing 3.2) or “*no*” (if more results may be provided on a later date—line 25 of Listing 3.2).
 - “*Re-resolution date*”. It is a date on which the receiver should repeat the location resolution process to get more results (line 37 of Listing 3.2).

Logically, it makes sense only if the previous field's indicator was “*no*”.

According to the values taken by the described fields (“*Acquisition directive*” and “*Resolution complete*”), we may find ourselves before different situations, depending on their combination:

- “*All*”–“*No*”. All elements on the list must be acquired and more should be expected to come. In the event that the result is a list of CRIDs, this may indicate that the report includes all the identified episodes of a series, but that more are supposed to exist (e.g. series with no scheduled end). Logically, the receiver should allow the known results to be consumed. In contrast, if it is a list of locators, the receiver can wait to obtain them all in order to provide them to the user.

²This generic description does not include enhancement mechanisms that would logically exist in most systems, such as the use of cache queries on information resolved in the past.

- “All”–“Yes”. All elements on the list must be acquired, after which the process is finished.
- “Any”–“No”. One of the elements on the list must be acquired (or wait until more appear), after which the process is finished.
- “Any”–“Yes”. One of the elements on the list must be acquired, after which the process is finished.

3.4.2 Resolution in Bidirectional Networks

In the case of bidirectional networks, where the receiver has a return channel to communicate with the outside world, it is necessary to have a protocol for the PDR to initiate a connection to communicate with one or more servers and resolve a CRID.

Although the complete description of this scenario in the documents of the TV-Anytime specification is intentionally generic, here we will focus on the case of a receiver connected to the Internet, which makes its connections through a TCP/IP network.

3.4.2.1 Discovery of the Resolution Server

The starting data is, precisely, the CRID of the selected content and, specifically, the DNS name of the authority that generated it. With this name, the PDR will query a DNS name server to find the server hosting the resolution service (you can see the sequence of stages in Fig. 3.5).

This DNS server must implement the DNS extension defined in RFC 2782 (IETF RFC 2782, 2000), originally designed so that computers connected to a network could find email servers, but it can also be used to implement the resolution of other kinds of services. Specifically, the generic format used in a consultation of this type is:

```
_Service._Protocol.Name
```

In the case of a TV-Anytime location resolution service, the standardized name of the service is “_lres” (which stands for *location resolution*), and the format of a consultation that asks about the resolution server that corresponds to a CRID of a given authority would be:

```
_lres._tcp.<authority>
```

The result of this consultation would be the name of the machine (and the TCP port) where the location resolution service is provided, which would later be translated into an IP address (as shown in Fig. 3.5).

3.4.2.2 Request Format

Once the receiver has the address of the location resolution server, the format of the consultation to that server is based on the usual requests of the HTTP protocol, using a *GET* request. The generic form of that request would be:

```
http://server/path-to-server-script?[key=value]&[key=value]&...
```

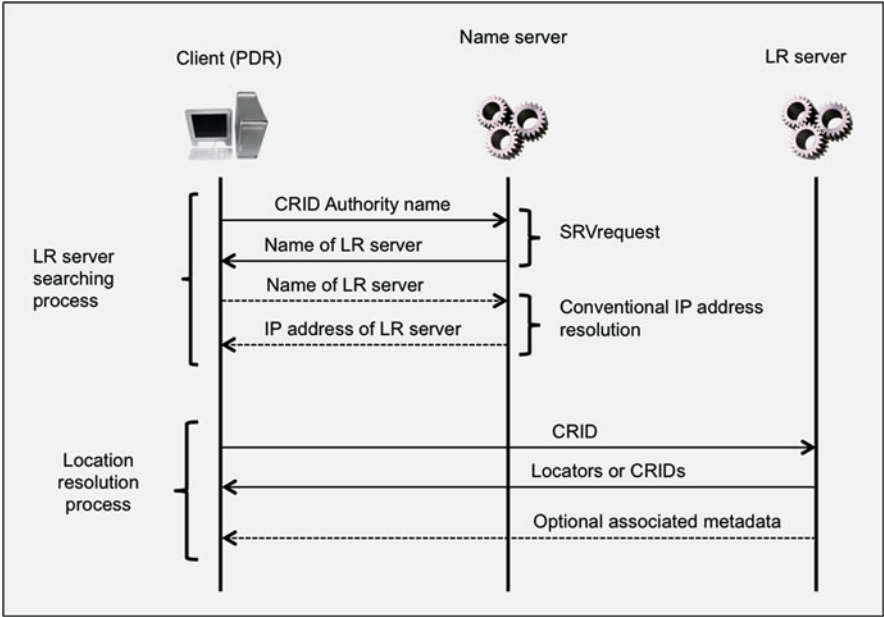


Fig. 3.5 TCP/IP-based CRID resolution

- The possible values of the parameters (the *key* fields) are:
- *Key*=“*CRID*”. The value would be the string that identifies the CRID we are asking about. It could be repeated several times if we ask about several CRIDs in one consultation.
 - *Key*=“*SubmittedCRID*”. The value could be 1 if we want the server to retrieve the metadata it has on that CRID (e.g. a *ProgramInformationTable*) or 0 if we only want the locator and not the metadata associated to the CRID.
 - *Key*=“*Result*”. The value could be 1 if we want the server to retrieve the metadata it has on that result, whether they are CRIDs or locators, or 0 if we do not want such metadata.

The only mandatory key is the first one (“CRID”) that can be repeated several times. The others can only appear once, applying to all CRIDs inquired. In case a key does not appear, it would mean that its value is 0.

3.4.2.3 Server Response

The response of the location server could be any of the following three:

- The final result of the location resolution process, consisting of the *ContentReferencingTable*. If we have requested metadata and it is available on that server, it will also provide us the *GroupInformationTable* and/or *ProgramInformationTable* structures (if the result is about CRIDs) or the *ProgramLocationTable* (if they are locators).

- A *RAR* record, where the address of another server to query is specified in the *URL* field.
- An HTTP redirection response, with the address of another server to query.

3.5 Summary

In this chapter, the content referencing mechanism defined by TV-Anytime has been presented, so that receivers can identify unambiguously any content in a generic way, regardless of the specific copy they may access later (even though it is possible to reference a concrete copy through the *IMI* tag). As stated, this mechanism offers a great flexibility to applications that provide functionalities to users, so that they can easily provide new value-added services that nowadays are not available or require an expensive and proprietary infrastructure.

Together with the information structures involved (*RAR* records and *ContentReferencingTable*), the resolution procedures that make it possible to convert references into location information, both in unidirectional contexts and in fully interconnected devices, have been presented.

In the next chapter, we will present the main element defined in the TV-Anytime documents to facilitate the implementation of the standardizing framework, that is the metadata structures available to describe contents and users, allowing for their smart classification, comparison and selection.

TV-Anytime

Paving the Way for Personalized TV

Gil Solla, A.; Sotelo Bovino, R.G.

2013, XX, 199 p., Hardcover

ISBN: 978-3-642-36765-6