

**The "btv" URN Scheme  
draft-finseth-btv-00.txt**

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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

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

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

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

This document expires 31 October 1999.

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

This document defines a new URI scheme, "btv". This scheme provides a mechanism by which "resources" -- which includes virtual channels,

events, applications, and data modules -- can be referenced from within a TV broadcast stream, from other broadcast streams, or from the Internet.

## Environment

The environment is television. In this environment, you have a receiver that picks up one or more signals and deal with them. (We'll use the term "receiver" to denote the box or module that does the work.) There are a couple of points to keep in mind.

First, the receiver may have only one source of signal, or it may have many. Typically, inputs can come from:

- over-the-air antenna (NTSC, ATSC, DVB, potentially other formats)
- cable feed (potentially a variety of formats)
- satellite feed (again, a variety of formats)

in any combination.

Second, the receiver may be stand alone, or it may be connected to some form of "back channel." The back channel may be proprietary or general-purpose (e.g., a network connection). The channel may be in the form of a phone line, DSL, cable modem, general LAN connection, etc.

## Description

The requirements that any TV broadcast URI scheme must meet are complex. The URI must connect a use of a resource with the definition of the resource. It must make this connection even though the resource may be transmitted in ways not foreseen by the original author.

A full summary of the requirements that such a scheme must meet is given in [\[1\]](#). This document will not repeat that discussion. Rather, this document will specify the syntax and semantics of a URI scheme that meets those requirements.

The general idea is that a URI is resolved by a TV receiver using translation information which is delivered to it in the broadcasts it receives. Thus, a URI may represent a resource which is broadcast by different terrestrial, cable, and satellite broadcasters in different locations, and which is therefore received by different TV receivers via different channels and even different broadcast protocols. These TV receivers can each correctly resolve the URI to the physical



broadcast parameters which represent the referenced resource in their circumstances by making use of translation information which is specific to their circumstances.

#### One-Paragraph Summary

Each resource -- such as a virtual channel, event, application, or data file -- is identified by a URI. The URIs are globally unique (conceptually) opaque tags and are assigned by the creator of the content. The information necessary to map the URI to the actual resource is carried in each transport stream in a transport-stream specific manner. These specifics are invisible to the application: it simply uses the URI in an accustomed manner to access the resource.

#### Syntax

The TV broadcast URIs fully comply with the syntax called out in [\[RFC2396\]](#). This document will not review the entire syntax. Rather, it will just highlight the specifics. Numbers in {}s refer to [RFC 2396](#) section numbers.

TV broadcast URIs are simultaneously URIs and URNs {1.2}. The term "URI" will continue to be used to refer to the collection. They are technically URLs as well, but only in a trivial way.

TV broadcast URIs are hierarchical and the relative forms are fully supported {1.4}.

The standard character set restrictions and operations (e.g. quoting) are used {2.}.

This document does not specify a maximum length of a TV broadcast URI. Content authors should consult the set of instantiation documents covering their target systems to determine any limitations on length.

The standard form {3.} is:

btv://<authority><path>?<query>

The <authority> part is a registration-based naming authority. The registration is the Domain Name System {3.2}. Note that the <authority> part is not expected to ever be resolved via DNS servers. Further, the <authority> part must be a name: a numeric identifier (such as IP address) is not permissible {3.2.2}.



The server-based portions (including "<userinfo>@" and ":<port>") will not be used {3.2.1}.

Path component interpretation is as specified in {3.3}.

Query components are permitted; interpretation is as specified in {3.4}.

Fragment identifiers are permitted; interpretation is as specified in {4.1}.

Relative URI references are permitted within an application (see next note). However, all URIs will be converted to fully-qualified form before being passed to the environment for lookup {5.}.

Base URIs can be obtained from the Document Context, the Retrieval URI, and the default Base URI. The Encapsulating Entity form is not meaningful and not permitted. Also, for Retrieval URIs, redirection is not meaningful and is not permitted. {5.1.}

The default Base URI is the URI explicitly associated with the invoking application. In the case where the invoking application is the system or is not clear, the default Base URI is not defined and relative forms are not permitted. {5.1.4}

## Semantics

From an application's viewpoint, a TV broadcast URI is simply an opaque string that is passed to the system for lookup. The system then uses information obtained from other sources to associate the opaque string with the actual resource. The actual resource may be specified directly (e.g., by values transmitted in mapping tables), or indirectly by, for example, a transport-specific URL. (This is exactly the same model as is used for looking up traditional "http:" URLs.)

This scheme assumes that there will be some "directory-type" information available to the system to perform the mapping. As this is presumably the same information used by the system to actually perform its other tasks, the mapping can be assumed to be complete and correct.

"Directory-type" information includes what is commonly referred to as "PSIP" information: virtual channel tables, event information tables, system description tables, etc. and possible extensions to that information. The specific names and formats of such tables will vary from one transmission method to another. In some cases, data may even



be entered manually.

TV broadcast URIs have three key properties:

1. URIs are considered to match when a case independent string comparison between the two returns True. Comparisons are always performed on the fully-qualified forms.

Note: case independent path portions were selected for ease of implementation. [RFC 2396](#) does not specify whether the path portion is case independent (the assigning authority and scheme portions must be compared case independent).

2. More than one resource may be identified by the same URI. Hence, a list of matching resources would be returned.

3. A given resource may be identified by more than one URI. Hence, the same resource may be returned by more than one request.

Note that while URIs can reference resources that are only available at particular times, the URIs themselves do not explicitly encode any time-related information: any such information is carried by other mechanisms. However, a URI may reference an "event" (e.g., a TV program) and that event does carry a time context.

#### Resources to be Named

We have identified the following types of resources that could be named with TV broadcast URIs. This list is not part of the specification, but is included to help understand how it is used.

Note that the <authority> portion identifies the entity that made the assignment.

The first two resources identify levels of transmission multiplex "bundling."

1. A transmission multiplex.

For example, the ABC network feed. These might be named:

```
      btv://abc.com/
or:   btv://abc.com/feed
or:   btv://feed.abc.com/
```

2. A particular virtual channel.





This resource identifies a particular virtual channel. Virtual channels might be named:

```
      btv://chan.abc.com/
or:    btv://abc.com/chan/
      btv://chan.kstp.com/
or:    btv://kstp.com/chan/

      btv://hbo.com/chan
```

The first two indicate the typical case of a network feed with local affiliate. (KSTP is the ABC affiliate in the Minneapolis/St. Paul area.) Both names would be attached to the same virtual channel.

Note that there is no need for the numeric form (Channel 5). If the local affiliate were to embed a URI in web content, the context would just be:

```
...<a href="btv://kstp.com/">tune to Channel 5 now</a>...
```

This same form would work on web content found in the Internet (assuming that the device had a tuner and was able to receive KSTP's signal), in content found in the TV broadcast stream, or any other content.

Since virtual channel numbers in a TV are re-used, an application can not usefully ever use the virtual channel number alone. Instead, it must refer to the context to determine how to name the desired resource. Given such a lookup, the alphabetic form of the name is no more effort to use than the numeric form.

Note that the national network can use the form:

```
...<a href="btv://abc.com/">tune to ABC now</a>...
```

and cause tuning to the local affiliate (assuming there is one) from anywhere.

While there may be some exceptions, TV broadcast URIs are intended mainly for internal use by the system and not for presentation to viewers. The following example illustrates why.

HBO (and many other networks) have multiple feeds: in the U.S., they have both an East and West coast feed, with the West coast feed ordinarily carrying the same content as the East coast feed, but with a three hour delay (live events are carried at the same time on both).



Local providers (e.g., cable systems) on the East coast ordinarily carry the East coast feed and call it "HBO". Local providers on the West coast ordinarily carry the West coast feed and also call it "HBO". Nationwide providers (e.g., satellite systems) may carry both feeds but must use different names (often "HBO" for the East coast feed and "HBOW" or "HBOP" for the West coast feed). Note that cable and satellite systems offer identical content on the West coast feed, but present that content under two user-visible names, said name depending upon the transmission path.

This URI scheme avoids problem by separating the URI naming from the user-visible naming. The provider can use the same URI whenever the content is the same. Note that the same content is just that: it's the same string of bits, even if shifted in time.

For example, a promotional spot might need a background GIF resource named:

```
btv://hbo.com/promotions/ad1453/background.gif
```

It can (and should) use the same name whenever the resource appears in the transmission. The provider can use its normal scheduling mechanisms to handle staging of the various feeds.

The next block of resources identify various combinations of event material.

### 3. A collection of event material.

For example, the set of Gilligan's Island episodes in the schedule. This resource might be defined by attaching a URI like this:

```
btv://gilligans-island.com/
```

to each episode in the schedule. A typical use of this form of URI might be on the Gilligan's Island web page. It might have a reference to:

```
...and click <a href="btv://gilligans-island.com/">here</a>  
to see a when and where it is on in <i>your</i> city...
```

### 4. A particular episode of a event series.

For example, the "Trouble with Tribbles" episode of the original Star Trek. This resource might be defined by a URI of one of these forms:

```
btv://trouble-with-tribbles.startrek.com/  
btv://startrek.com/trouble-with-tribbles
```



btv://startrek.com/2nd-season/trouble-with-tribbles

This specification does not identify which form is preferred: just as with "normal" web URLs, the preferred form is selected by the organization named in the <authority> portion.

In many cases, material with a the specific episode URI would also have a generic series URI:

btv://trouble-with-tribbles.startrek.com/  
btv://startrek.com/

#### 5. A particular showing of a event.

For example, the "Trouble with Tribbles" episode of the original Star Trek that ran on Nov 16, 1998 at 10:00pm on virtual channel 5 in Minneapolis/St. Paul Angeles. This resource might be identified by a URI of the form:

btv://kstp.com/trouble-with-tribbles/startrek.com

The URI would be placed only on the showing for the specified virtual channel and at the specified time.

The naming of this example assumes that the local station is the source of the reason for the reference. For example, there is a promotion in conjunction with a local Star Trek convention. A more typical name might be:

btv://kstp.com/promotion/star-trek/19981116

The point is, the organization creating the reference can use whatever name best suits their purpose.

#### 6. Event material that appears someplace in the time/virtual channel/provider continuum.

For example, the January 1999 Super Bowl. This event might carry the URI:

btv://super-bowl-Jan-1999.com/

or even just,

btv://super-bowl.com/

Given that only one is likely to show up in the schedule at a time. As with the nationwide network example, the reference can be embedded



as:

```
...and click <a href="btv://super-bowl.com/">here</a> to  
schedule your TV to tune to the Super Bowl!...
```

Again, this reference is only useful if the viewer can actually receive the Super Bowl broadcast and it is in the available event guide. It also assumes that references to events in the future can integrate to a scheduling system (this situation is common in satellite and cable set-top boxes).

#### 7. A particular once-only event.

All of the previous examples referred to events that were in some sense repeating or part of a series. The same forms and concepts apply to once-only events. The organization responsible assigns the URI and it is attached to the appearances.

The last two points identify "data"-oriented resources.

#### 8. A particular application.

URIs can be attached to the applications themselves. These URIs can be used in two ways.

First, they identify applications to each other for inter-process communications.

Second, the default Base URI for the application is established by the act of naming the application itself.

Example names might be:

```
btv://coke.com/commercial45/application  
btv://abc.com/promo5.html
```

#### 9. A particular data item referenced by an application.

Examples are image files (JPEGs, GIFs), sounds, data streams, etc.

For example,

```
btv://coke.com/commercial45/application/background.gif
```

### Referencing Other Data

This specification covers how a receiver learns about URIs received only over a TV broadcast stream. Receivers may use other URI schemes





as they see fit. For example:

- The "http:", "ftp:", and similar schemes can be used to access content from the Internet (given such a connection).
- The "file:" scheme may be used to access local information.
- Other TV broadcast-related schemes may coexist with this scheme (e.g., "dvb:" and "davic:"). Applications that use such schemes limit their scope of operation to those systems that implement such schemes.

### Security Considerations

This specification does not address security. It assumes that security and access control are handled by policies and procedures implemented in the systems themselves.

While this specification appears to provide a level of security by virtue of the fact that the only resources accessible by an application are those named in the broadcast stream, such is not the case. Absent other policies and procedures, an ill-behaved application can access resources by using device-specific mechanisms.

### References

[1] ten Kate, Warner; Thomas, Gomer; Finseth, Craig, "TV Broadcast URI Schemes Requirements", 11 March 1999.

<http://www.w3.org/TV/TVWeb/TVWeb-URI-Requirements-19990311>

[RFC2396] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax", [RFC 2396](#), August 1998.

### Author's Addresses

Craig A. Finseth  
U.S. Satellite Broadcasting  
3415 University Ave  
St Paul MN 55114  
Phone: +1 651-659-7162  
Email: [craig@finseth.com](mailto:craig@finseth.com)

Gomer Thomas  
LGERCA, Inc.  
40 Washington Road  
Princeton Junction, NJ 08550  
Email: [gomer@lgerca.com](mailto:gomer@lgerca.com)



## Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

This document expires 31 October 1999.

