

INTERNET-DRAFT  
Document: [draft-ietf-mmusic-sdp-tcpmedia-00.txt](#)  
Expires April 2001

D. Yon  
Dialout.Net  
October 2000

**TCP-Based Media Transport in SDP**  
**<[draft-ietf-mmusic-sdp-tcpmedia-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>.

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This document describes how to express TCP-based media transport using the Session Description Protocol (SDP). It defines two new protocol identifiers: TCP and RTP/AVP-TCP. It also defines the syntax and semantics for an SDP "direction" attribute that describes the TCP connection setup procedure.

## Introduction

The Session Description Protocol [[SDP](#)] provides a general-purpose format for describing multimedia sessions in announcements or invitations. SDP uses an entirely textual data format (the US-ASCII subset of [[UTF-8](#)]) to maximize portability among transports. SDP does not define a protocol, but only the syntax to describe a multimedia session with sufficient information to discover and participate in that session. Session descriptions may be sent using any number of existing application protocols for transport (e.g., SAP, SIP, RTSP, email, HTTP, etc.).

## Motivation

[SDP] describes two protocol identifiers: RTP/AVP and UDP, both of which are unreliable, connectionless protocols, an appropriate choice for multimedia streams. There are, however, applications for which the connection-oriented TCP transport is more appropriate, but [[SDP](#)] provides no way to describe a session that uses the TCP protocol.

The connection-oriented nature of TCP introduces a new factor when describing a session: not only must it be possible to express that a protocol will be based on TCP, but it must also describe the connection setup procedure.

## [1](#) Protocol Identifiers

### [1.1](#) RTP/AVP-TCP

[SDP] already specifies the RTP/AVP protocol identifier as a way to describe media that uses the Real Time Protocol and the associated Audio-Video Profiles. The assumption in [[SDP](#)] is that the underlying network transport is UDP. RTP is transport-neutral, so it is possible to transport RTP/RTCP packets using TCP rather than UDP. To describe a media session that uses RTP/AVP over TCP, the protocol identifier RTP/AVP-TCP must be specified in the m= line.

### [1.2](#) TCP

The TCP protocol identifier is similar to the UDP protocol identifier in that it only describes the transport protocol without any connotation as to the upper-layer protocol. An m= line that specifies TCP must further qualify the protocol using a fmt

identifier (see [[SDP](#)] [Appendix B](#)).

## **2 Direction Attribute**

An important attribute of a TCP connection is the setup procedure. One endpoint needs to initiate the connection and the other endpoint needs to accept the connection. The direction attribute is used to describe these roles, and the syntax is as follows:

Yon	INTERNET-DRAFT - Expires March 2001	2
INTERNET-DRAFT	TCP Media in SDP	October 2000

`a=direction:<role> <source-port>`

The <role> is one of the following:

passive:     The endpoint will accept an incoming TCP connection.

active:       The endpoint will initiate an outgoing TCP connection.

both:         The endpoint will both accept an incoming TCP connection  
              and will initiate an outgoing TCP connection.

The <source-port> is an optional value that may only be specified in the context of `direction:active` or `direction:both`.

### **2.1 Semantics of `direction:passive`**

By specifying `direction:passive`, the endpoint indicates that the port number specified in the `m=` line is available to accept a TCP connection from the other endpoint.

### **2.2 Semantics of `direction:active`**

By specifying `direction:active`, the endpoint indicates that it will initiate a TCP connection to the port number on the `m=` line of the other endpoint. The port number on its own `m=` line is irrelevant and is to be ignored by the other endpoint. Nevertheless, since the `m=` line must contain a valid port number, the endpoint specifying `direction:active` should specify a port number of 9 (the discard port) on its `m=` line. The endpoint must not specify a port number of zero, as that carries other semantics in [[SDP](#)].

The endpoint may optionally specify the port number from which it will initiate the TCP connection in the <source-port> position on the `a=` line.

### **2.3 Semantics of `direction:both`**

By specifying `direction:both`, the endpoint indicates that it will both accept a TCP connection on the port number of its own `m=` line, and that it will also initiate a TCP connection to the port number on the `m=` line of the other endpoint. As with `direction:active`, the endpoint may optionally specify the port number from which it will initiate the TCP connection in the `<source-port>` position on the `a=` line.

Since this attribute describes behavior that is similar to connectionless media descriptions in [\[SDP\]](#), it is the default value for the `direction` attribute and is therefore optional.

Endpoints may choose to specify `direction:both` for one or more of the following reasons:

Yon	INTERNET-DRAFT - Expires March 2001	3
INTERNET-DRAFT	TCP Media in SDP	October 2000

- 1) The endpoint has no preference as to whether it accepts or initiates the TCP connection, and therefore is offering the remote endpoint a choice of connection setup procedures.
- 2) The endpoints intend to use a single TCP connection to transport the media, but it is not known whether firewall issues will prevent either endpoint from initiating or accepting the TCP connection. Therefore both endpoints will attempt to initiate a TCP connection in hopes that at least one will succeed.
- 3) The endpoints intend to use two TCP connections to transport the media, and one must be initiated by the remote endpoint and the other must be initiated by the local endpoint.

If one endpoint specifies either `direction:active` or `direction:passive` and the other specifies `direction:both`, both endpoints must behave as if the latter had specified the inverse direction of the former. For example, specifying `direction:both` when the other endpoint specifies `direction:active` should cause both endpoints to behave as if the former had specified `direction:passive`. Conversely, specifying `direction:both` when the other endpoint specifies `direction:passive` should cause both endpoints to behave as if the former had specified `direction:active`.

If both endpoints specify `direction:both` then each endpoint must initiate a TCP connection to the port number specified on the `m=` line of the opposite endpoint. If only one connection succeeds, then that connection will be used to carry the media. If both

connections succeed but only one was needed (case #2 above), the following rules shall apply:

- a) Each endpoint MUST accept data from either TCP connection.
- b) Once an endpoint has transmitted data to one of the TCP connections, it MUST use that TCP connection exclusively for transmission.
- c) Once an endpoint has transmitted AND received data, if one of the TCP connections is determined to be idle, the endpoint MAY close the idle TCP connection.

### 3 Source-Port Considerations

In the cases where the endpoint is initiating the TCP connection, a source port number may optionally be specified on the a= line by that endpoint. In most environments, the source port number can be determined by binding the socket before initiating the connect, as shown in the sample C code below:

```
{
    SOCKET s_id
    SOCKADDR_IN cli_sin;
    int namelen;
```

Yon	INTERNET-DRAFT - Expires March 2001	4
INTERNET-DRAFT	TCP Media in SDP	October 2000

```
    // Create the socket
    s_id = socket(AF_INET, SOCK_STREAM, IPPROTO_TCP);

    // Bind the socket to any IP address and port
    bzero((char *)&cli_sin, sizeof(cli_sin));
    cli_sin.sin_family      = AF_INET;
    cli_sin.sin_addr.s_addr = htonl(INADDR_ANY);
    cli_sin.sin_port        = 0;
    bind(s_id, (SOCKADDR *)&cli_sin, sizeof(cli_sin));

    // Find the port number that was bound
    namelen = sizeof(cli_sin);
    getsockname(s_id, (SOCKADDR *)&cli_sin, &namelen);

    // Print the port number
    printf("Source Port = %d\n", ntohs(cli_sin.sin_port));
}
```

The motivation for specifying the source port is to allow topologies where one or more endpoints use a single, fixed TCP port for

incoming connections. Non-RTP protocols transported over TCP commonly use this technique. By specifying the source port, an endpoint avoids a potential ambiguity when more than one session is set up between two endpoints.

For example, consider two endpoints with IP addresses of 10.1.1.1 and 10.1.1.2. The endpoint at 10.1.1.1 signals the availability of a session on TCP port 2393 (passive). Before the endpoint at 10.1.1.2 has a chance to initiate the connection, events transpire that cause the endpoint at 10.1.1.1 to signal the availability of a separate session that is also found at TCP port 2393 (passive). Shortly thereafter, both entities at 10.1.1.2 initiate connections to 10.1.1.1 on port 2393.

The problem is this: how does the endpoint at 10.1.1.1 differentiate the two connections? To which entity at 10.1.1.2 does each connection correspond? By specifying the source port prior to connecting, the entities at 10.1.1.2 can avoid this ambiguity, because now the endpoint at 10.1.1.1 can simply inspect the port number from which the connection originated to determine which entity has initiated the connection.

Caution must be exercised when designing systems that rely on this feature, as not all environments are able to determine the source port prior to initiating the connection.

## **4 Examples**

What follows are a number of examples that show the most common usage of the direction attribute combined with TCP-based media descriptions. For the purpose of brevity, the main portion of the session description is omitted in the examples and is assumed to be the following:

Yon	INTERNET-DRAFT - Expires March 2001	5
INTERNET-DRAFT	TCP Media in SDP	October 2000

```
v=0
o=Me <me@ietf.org>
s=Call me using TCP
t=0 0
```

### **4.1 Example: simple passive/active**

An endpoint at 10.1.1.2 signals the availability of an audio session at port 54111:

```
c=IN IP4 10.1.1.2/127
m=audio 54111 RTP/AVP-TCP 0
```

```
a=direction:passive
```

An endpoint at 10.1.1.1 receiving this description responds with the following:

```
c=IN IP4 10.1.1.1/127
m=audio 9 RTP/AVP-TCP 0
a=direction:active
```

The endpoint at 10.1.1.1 then initiates the TCP connection to port 54111 at 10.1.1.2. Note that the TCP connection may originate from any port. The endpoint at 10.1.1.1 could have optionally committed to a source port with a simple modification:

```
c=IN IP4 10.1.1.1/127
m=audio 9 RTP/AVP-TCP 0
a=direction:active 1892
```

By adding the "1892" to the a= line, the endpoint at 10.1.1.1 must now use a source port of 1892 when initiating the TCP connection to port 54111 at 10.1.1.2.

#### **4.2 Example: agnostic both**

An endpoint at 10.1.1.2 signals the availability of an audio session at TCP port 54111, but is also willing to set up the media stream by initiating the TCP connection:

```
c=IN IP4 10.1.1.2/127
m=audio 54111 RTP/AVP-TCP 0
a=direction:both
```

The endpoint at 10.1.1.1 has three choices:

- 1) It can respond with either of the two direction:active descriptions listed in the previous example. In this case the endpoint at 10.1.1.1 must initiate a connection to port 54111 at 10.1.1.2.
- 2) It can respond with a description similar to the following:

```
c=IN IP4 10.1.1.1/127
m=audio 54321 RTP/AVP-TCP 0
a=direction:passive
```

In this case the endpoint at 10.1.1.2 must initiate a

connection to port 54321 at 10.1.1.1.

- 3) It can respond with a description that specifies direction:both, which is covered in the next example.

#### **4.3 Example: redundant both**

An endpoint at 10.1.1.2 uses the same description as the previous example:

```
c=IN IP4 10.1.1.2/127
m=audio 54111 RTP/AVP-TCP 0
a=direction:both
```

Unlike the previous example, the endpoint at 10.1.1.1 responds with the following description:

```
c=IN IP4 10.1.1.1/127
m=audio 54321 RTP/AVP-TCP 0
a=direction:both
```

This will cause the endpoint at 10.1.1.2 to initiate a connection to port 54321 at 10.1.1.1, and the endpoint at 10.1.1.1 to initiate a connection to port 54111 at 10.1.1.2. Whichever TCP connection succeeds will be used. If both succeed, one of the connections may be closed as an optimization, using the rules in [section 2.3](#).

## **5 Security Considerations**

See [\[SDP\]](#) for security and other considerations specific to the Session Description Protocol in general. There are no new security considerations introduced by these protocol identifiers and attributes.

## **6 IANA Considerations**

As recommended by [\[SDP\]](#) [Appendix B](#), the direction attribute described in this document should be registered with IANA, as should the TCP and RTP/AVP-TCP protocol identifiers.

## **Acknowledgements**

The author would like to thank Jonathan Rosenberg, Anders Kristensen, and Robert Fairlie-Cunninghame for their valuable insights.

## Appendix A: Direction Attribute Syntax

This appendix provides an Augmented BNF [[ABNF](#)] grammar for expressing the direction attribute for TCP connection setup. It is intended as an extension to the grammar for the Session Description Protocol, as defined in [[SDP](#)]. Specifically, it describes the syntax for the new "connection-setup" attribute field, which MAY be either a session-level or media-level attribute.

```
connection-setup = "direction" ":" direction-spec
```

```
direction-spec = "passive" | qualified-direction
```

```
qualified-direction = direction-ident | direction-ident port
```

```
direction-ident = "both" | "active"
```

## References

- [ABNF] D. Crocker, P. Overell, "Augmented BNF for Syntax Specifications: ABNF," [RFC 2234](#), November 1997
- [SDP] M. Handley, V. Jacobson, "SDP: Session Description Protocol," [RFC 2327](#), April 1998
- [UTF-8] F. Yergeau, "UTF-8, a transformation format of Unicode and ISO 10646," [RFC 2044](#), October 1996

## Author's Address

David Yon  
Dialout.Net, Inc.  
402 Amherst St  
Nashua, NH 03063

Phone: (603) 577-8708  
EMail: [yon@dialout.net](mailto:yon@dialout.net)

## Full Copyright Statement

Copyright (C) The Internet Society (2000). 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

Yon

INTERNET-DRAFT - Expires March 2001

8

INTERNET-DRAFT

TCP Media in SDP

October 2000

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."

