

Low Energy Audio – Basic One-Way Unicast Audio

Bluetooth® Informational Publication

- **Version:** v1
- **Version Date:** 2024-06-11
- **Prepared By:** Generic Audio Working Group

If you have any questions or comments on this informational publication or Bluetooth LE Audio specifications, please submit a ticket: <https://www.bluetooth.com/leaudio-spec-feedback/>

Abstract:

This document gives an introduction to Bluetooth LE Audio by explaining the basics of a use case where stereo one-way audio is transmitted from one device to another.



Version History

Version Number	Date (yyyy-mm-dd)	Comments
v1	2024-06-11	Approved by the Bluetooth SIG Board of Directors.

Acknowledgments

Name	Company
Asbjørn Sæbø	Nordic Semiconductor ASA
Georg Dickmann	Sonova AG
Nick Hunn	GN Hearing A/S
Andrew Estrada	Sony Group Corporation
Oren Haggai	Intel Corporation
Chris White	Dolby Laboratories, Inc.
Bjarne Klemmensen	Demant A/S



This document, regardless of its title or content, is not a Bluetooth Specification as defined in the Bluetooth Patent/Copyright License Agreement (“PCLA”) and Bluetooth Trademark License Agreement. Use of this document by members of Bluetooth SIG is governed by the membership and other related agreements between Bluetooth SIG Inc. (“Bluetooth SIG”) and its members, including the PCLA and other agreements posted on Bluetooth SIG’s website located at www.bluetooth.com.

THIS DOCUMENT IS PROVIDED “AS IS” AND BLUETOOTH SIG, ITS MEMBERS, AND THEIR AFFILIATES MAKE NO REPRESENTATIONS OR WARRANTIES AND DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING ANY WARRANTY OF MERCHANTABILITY, TITLE, NON-INFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, THAT THE CONTENT OF THIS DOCUMENT IS FREE OF ERRORS.

TO THE EXTENT NOT PROHIBITED BY LAW, BLUETOOTH SIG, ITS MEMBERS, AND THEIR AFFILIATES DISCLAIM ALL LIABILITY ARISING OUT OF OR RELATING TO USE OF THIS DOCUMENT AND ANY INFORMATION CONTAINED IN THIS DOCUMENT, INCLUDING LOST REVENUE, PROFITS, DATA OR PROGRAMS, OR BUSINESS INTERRUPTION, OR FOR SPECIAL, INDIRECT, CONSEQUENTIAL, INCIDENTAL OR PUNITIVE DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, AND EVEN IF BLUETOOTH SIG, ITS MEMBERS, OR THEIR AFFILIATES HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

This document is proprietary to Bluetooth SIG. This document may contain or cover subject matter that is intellectual property of Bluetooth SIG and its members. The furnishing of this document does not grant any license to any intellectual property of Bluetooth SIG or its members.

This document is subject to change without notice.

Copyright © 2024 by Bluetooth SIG, Inc. The Bluetooth word mark and logos are owned by Bluetooth SIG, Inc. Other third-party brands and names are the property of their respective owners.



Contents

1	Introduction	5
1.1	Scope.....	5
1.2	MSC conventions.....	6
2	One-way audio	7
2.1	Initial conditions.....	8
2.2	Establishing an ACL connection.....	10
2.3	Codec and CIG configuration.....	12
2.4	QoS configuration.....	14
2.5	Enabling ASEs and setting up CISEs.....	16
2.6	Starting streaming.....	18
2.7	Ending streaming and cleaning up.....	20
2.8	Considerations on setup latency.....	22
3	References	23



1 Introduction

This document is intended as a first introduction to unicast¹ Low Energy (LE) Audio. It gives an overview of a simple scenario where stereo one-way audio is transferred from one device to another.

A Message Sequence Chart (MSC), subdivided into multiple parts, is used to explain the steps for the Bluetooth LE Audio use case of a user listening to music from one device (e.g., a PC) on another device (e.g., a headset) connected to the first one. While each of the LE Audio specifications addresses only individual aspects of a given use case, this document shows how the concepts from multiple LE Audio specifications are combined.

The content of this document is based upon the following specifications:

- Basic Audio Profile (BAP) [1]
- Published Audio Capabilities Service (PACS) [2]
- Audio Stream Control Service (ASCS) [3]
- Common Audio Profile (CAP) [4]

These specifications are considered required reading for anyone implementing LE Audio. In addition, Bluetooth SIG publishes other informational documents to help understand the specifications.

This is not a Bluetooth specification, therefore, the established Bluetooth SIG specification language conventions for use of the words **shall**, **shall not**, **must**, **should**, **should not**, **may**, and **can** do not apply to this document.

1.1 Scope

The purpose of this document is to introduce the basics of LE Audio by showing one possible implementation using the features within Bluetooth LE Audio specifications. The MSC used in this document shows many of the concepts of LE Audio specifications including the:

- Use of Context Types [5]
- Setup of data paths
- Configuration of Audio Stream Endpoints (ASEs)
- Use of BAP [1] Announcements

The solution shown is not the only possible solution; developers can use alternative approaches based on other considerations.

¹ «Unicast» means that the audio data are transferred over a connection between the devices, as opposed to “broadcast” transfer, where the audio data are not transferred over a connection.



1.2 MSC conventions

The MSC in this document uses the conventions illustrated in Figure 1.1.

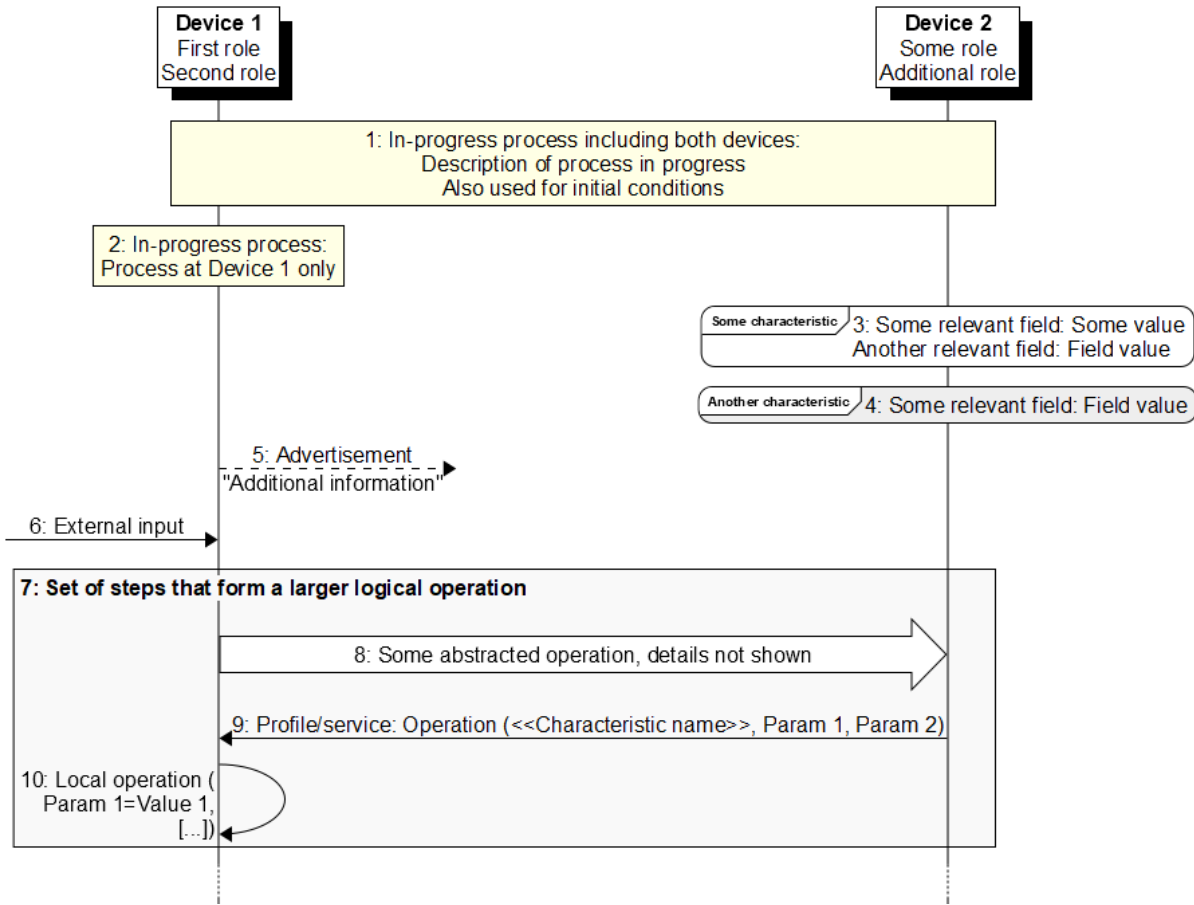


Figure 1.1: MSC conventions

- Descriptions that refer to the in-progress process are shown in yellow boxes that span across the relevant devices. The same type of box is also used for the initial conditions.
- PACS [2] characteristics, with relevant values, are shown in rounded white boxes.
- ASE characteristics, with the state and relevant other values, are shown in rounded grey boxes.
- Advertisements are shown as dashed arrows starting from the transmitters. The direction of the arrows does not imply that the advertisement is received by any specific device.
- BAP procedures (or other sets of steps that together form a larger logical operation or process) are surrounded by large light-gray boxes. A tag in the upper-left corner of each box describes the procedure.
- Some parts of the MSC, where the details are not important to the topic of this document, are shown abstracted as white block arrows.
- To refer to GATT characteristic X, <<X>> is used. For example, as in <<Sink PAC>>, which refers to the PACS Sink PAC characteristic.
- The steps of the MSC are simply numbered according to the order in which they occur in the MSC; the numbering does not imply that this is the only possible ordering of the steps.

2 One-way audio

The MSC in [Figure 2.1](#) through [Figure 2.7](#) shows an overview of one-way unicast transmission of audio from a computer to a headset. (Instead of a computer, the transmitting device could be, for example, a tablet or a smartphone, and instead of a headset, the receiving device could be e.g., a loudspeaker. The steps in the MSC would still be the same.)

The MSC is simplified to show only the fundamental operations and parameters required to set up the audio transmission. Additional operations and parameters, such as using media control operations for controlling a media player and details on the use of Context Type values are out of scope for this document.

The MSC begins with the situation where the two devices have previously paired and bonded but are not connected. It is also possible to start from an already-connected state; in that case, Steps 7 through 22 shown in the MSC would not apply.

As shown in the MSC, the computer in this setup is a CAP Initiator. The computer is also:

- A BAP Unicast Client
- A Generic Attribute Profile (GATT) Client
- In the Generic Access Profile (GAP) Central role
- In the BAP Audio Source audio role

The MSC also shows that the headset in this setup is a CAP Acceptor. The headset is also:

- A BAP Unicast Server
- A GATT Server
- In the GAP Peripheral role
- In the BAP Audio Sink audio role



2.1 Initial conditions

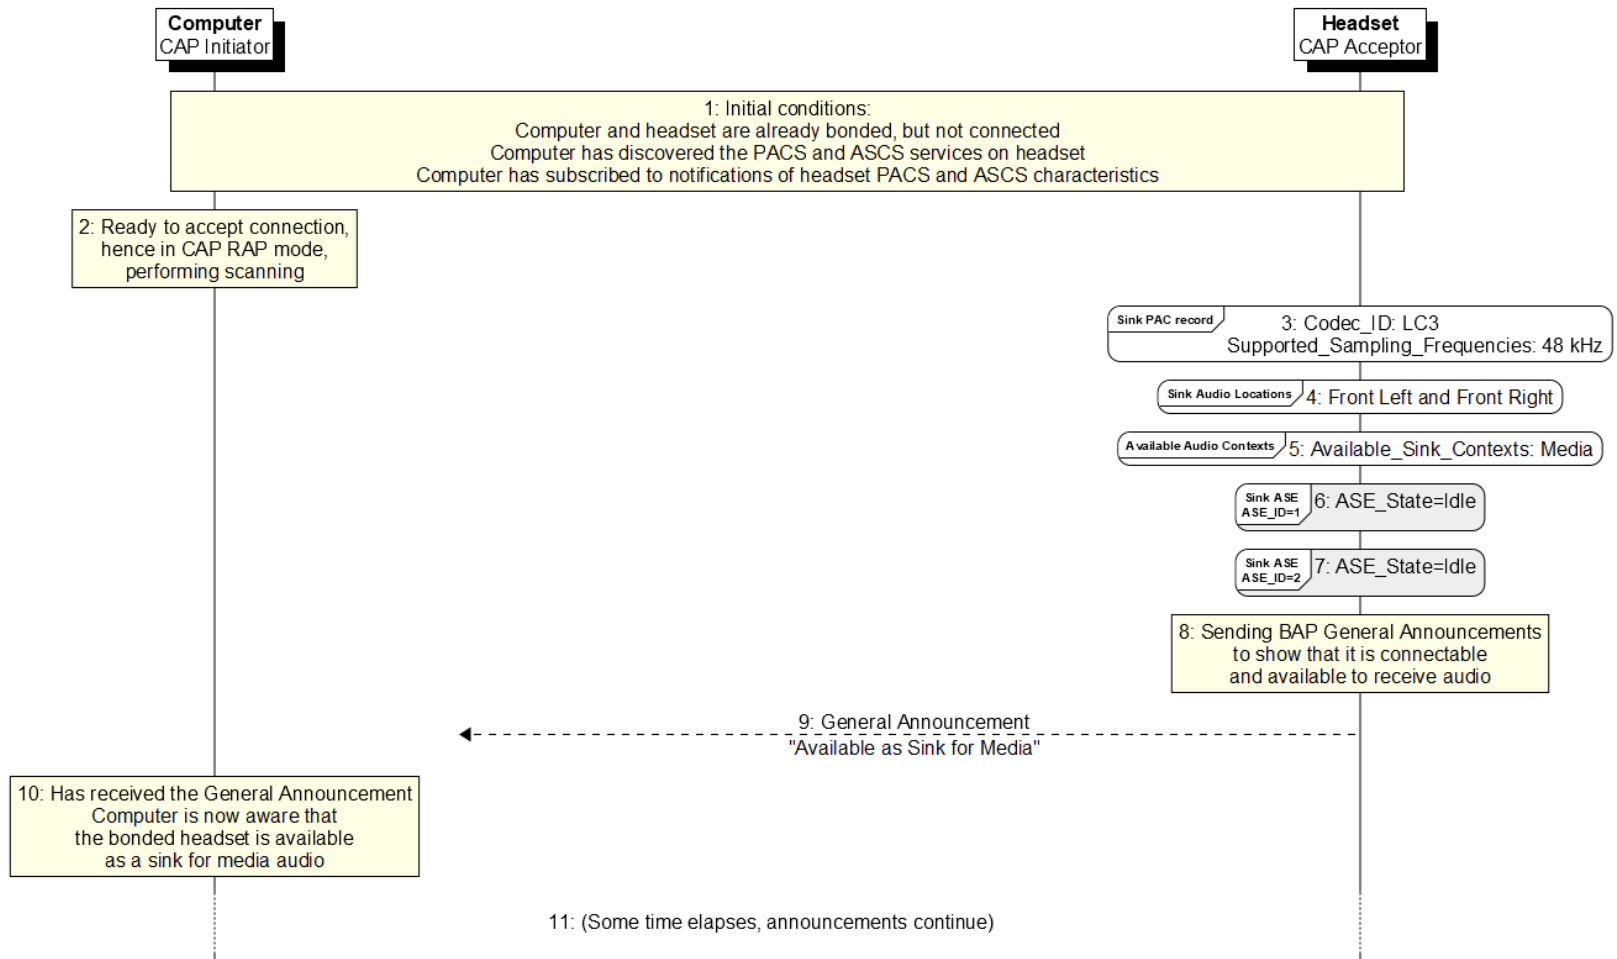


Figure 2.1: One-way unicast streaming from a computer to a headset (Part 1 of 7)



Figure 2.1 shows the initial conditions:

- The computer and the headset are bonded but not connected. The computer has previously discovered the services on the headset and subscribed to notifications for the service characteristics.
- The computer is ready for a connection for LE Audio – it is in the CAP Ready for Audio related Peripheral (RAP) mode (see Section 8.2 in [4]). It is therefore scanning for announcements.
- The headset exposes two sink ASEs. This corresponds to BAP Audio Configuration 6(i). In BAP, this configuration is mandatory to be supported by Unicast Servers in the Audio Sink role that expose more than one audio location. The configuration is also mandatory to be supported by Unicast Clients in the Audio Source role. This configuration can, for example, be used both for banded headphones and for LE Audio devices representing a pair of physical devices (e.g., a pair of earbuds represented by one of the earbuds).
- The headset is announcing that it is available.

2.2 Establishing an ACL connection

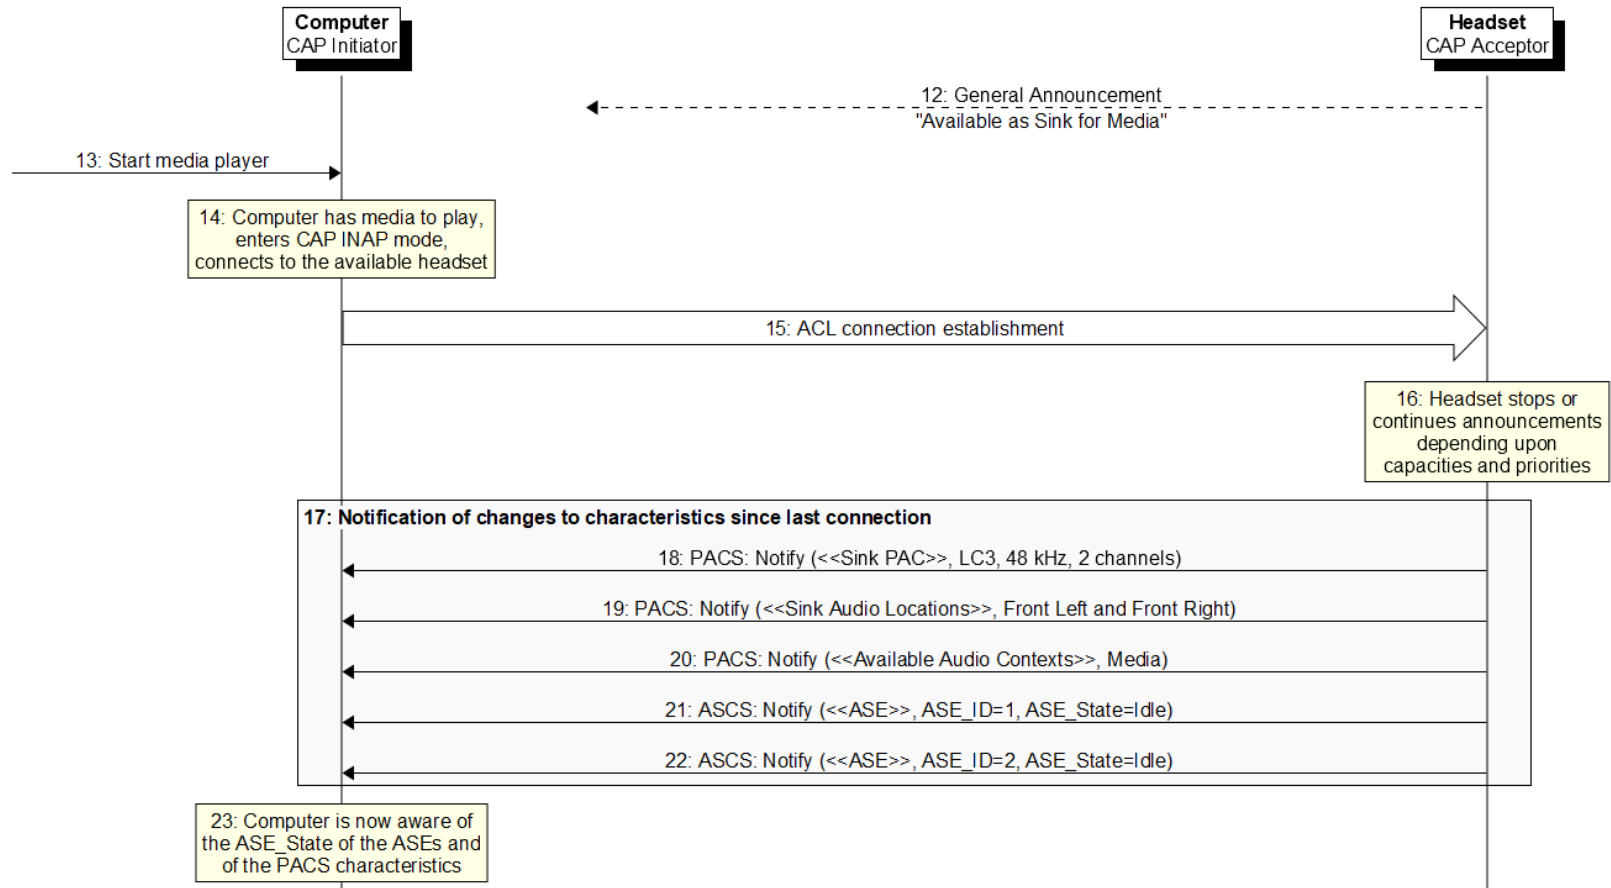


Figure 2.2: One-way unicast streaming from a computer to a headset (Part 2 of 7)



Figure 2.2 shows the following steps:

- Some external input (e.g., a button press from a user) starts a media player.
- The computer now has a need for a sink for media. It is in the CAP Immediate Need for Audio related Peripheral (INAP) mode (see Section 8.2.2 in [4]). It knows from the announcements that the headset is available for this and connects to the headset. (This illustrates the “unconnected model” possible under LE Audio, where a device does not connect until there is a need to transfer audio.)
- The headset notifies the computer of any changes to the sink PAC record and the state of the ASEs since the computer was last connected to the headset. (See [2] and [3] for details on when it is mandatory to send such notifications.) Alternatively, the computer may read these characteristics from the headset.
- As a result, the computer is aware of the current state of the ASEs and of the PACS characteristics of the headset.

2.3 Codec and CIG configuration

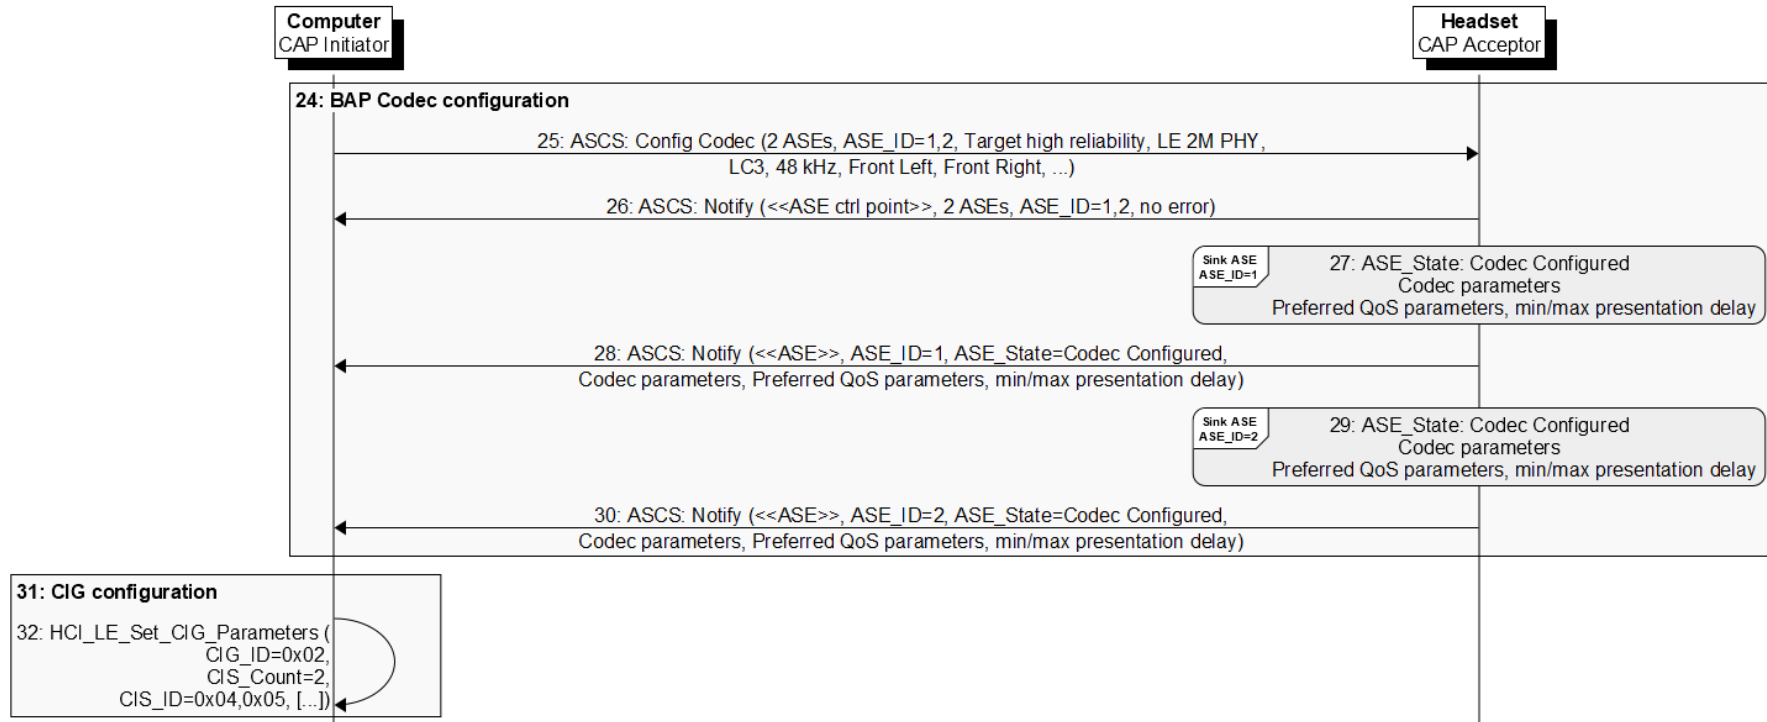


Figure 2.3: One-way unicast streaming from a computer to a headset (Part 3 of 7)



Figure 2.3 shows the following steps:

- The computer performs Configure Codec operations for ASE 1 and ASE 2. Some of the codec configuration values are sent as LTV structures (see Section 6.12.5 in [6]).
- The headset transitions the ASEs to Codec Configured state and notifies the new state and the associated parameters.
- The computer then proceeds to configure the Connected Isochronous Group (CIG). This may also be done earlier (implementation specific) but must be done before proceeding to QoS configuration.

2.4 QoS configuration

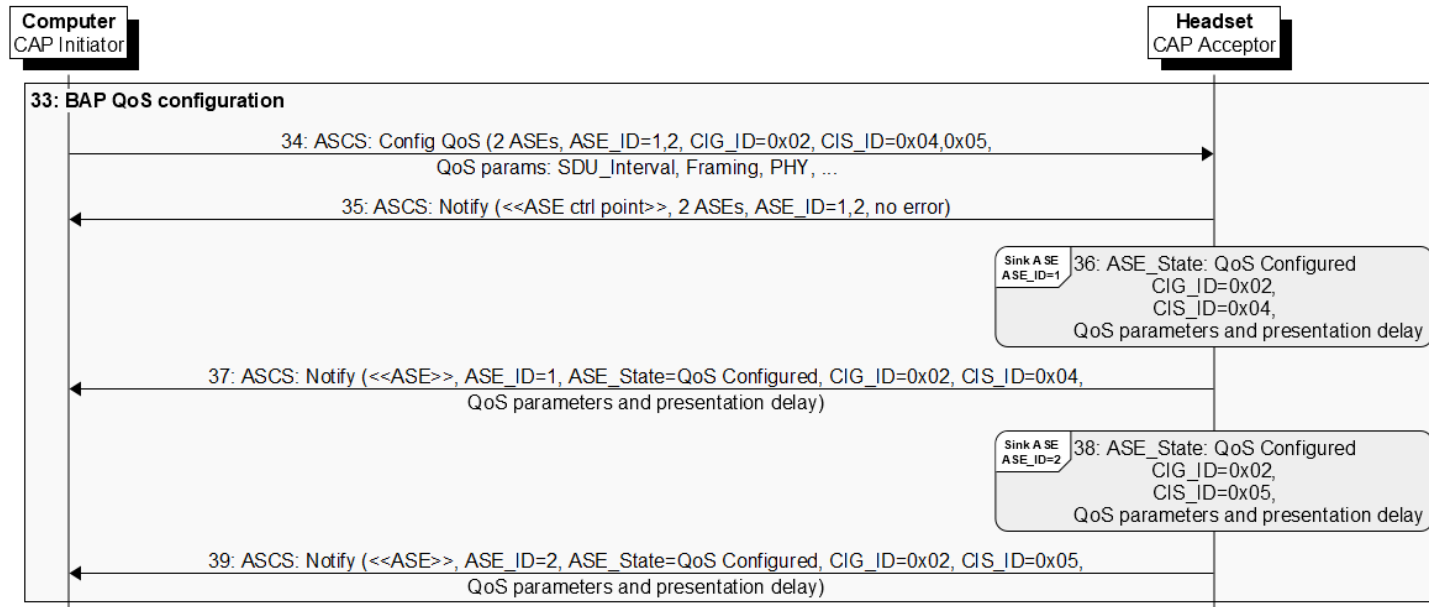


Figure 2.4: One-way unicast streaming from a computer to a headset (Part 4 of 7)



Figure 2.4 shows the following steps:

- The computer configures the QoS settings.
- The headset notifies the new state (QoS Configured) and the associated parameters.
- After the successful Config QoS operation, the Connected Isochronous Streams (CISes) identified in the Config QoS operation are now considered bound to the ASEs.
- Note that the QoS parameters in Steps 36 to 39 are the ones written by the computer host in Step 34. The actual QoS parameters chosen by the controller may still differ from these (see Volume 4, Part E, Section 7.8.97 in [7]).

2.5 Enabling ASEs and setting up CISes

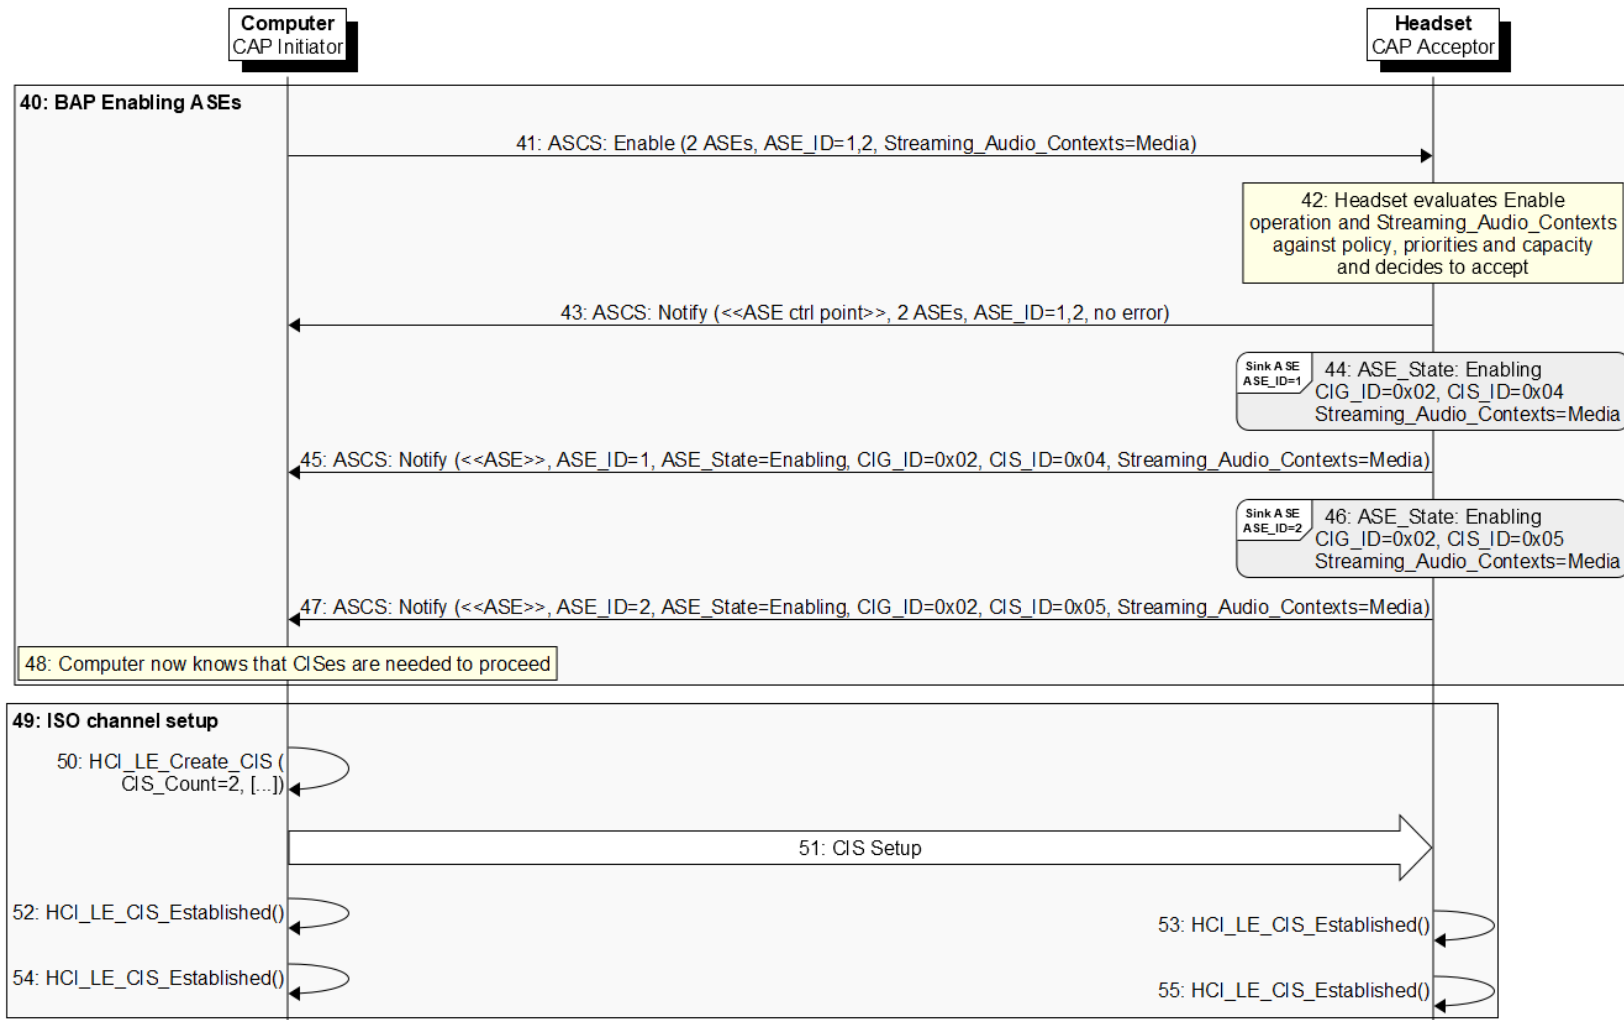


Figure 2.5: One-way unicast streaming from a computer to a headset (Part 5 of 7)



Figure 2.5 shows the following steps:

- The computer enables the ASEs. At this point, the headset knows both the configuration of, and the use case (context) for, the CISes connected to its ASEs. At this point, the headset may also drop other CISes to free up capacity.
- The headset notifies the new state (Enabling) and the associated parameters.
- The computer proceeds to set up the required CISes. This triggers CIS setup on both sides, ending when the host on each side has received the two HCI_LE_CIS_Established events, one for each CIS. The CIS setup may also be done earlier, but if so, the headset may reject, not knowing the use case for the CISes. By doing the CIS setup here, after the Enable operation has been accepted, the computer can be confident that the headset will not reject the CISes. (See Section 5 in [1] for details on CIS setup.)
- After the CISes are established, they are considered coupled to the ASEs.

2.6 Starting streaming



Figure 2.6: One-way unicast streaming from a computer to a headset (Part 6 of 7)



Figure 2.6 shows the following steps:

- The headset sets up the ISO data path for each of the CISEs to route data from the controller to where it will be rendered.
- The headset initiates the Receiver Start Ready operation, transitioning the ASEs to the Streaming state.
- The headset notifies the new state (Streaming) and the associated parameters.
- The computer sets up the ISO data path for each of the CISEs to route data from where it is coming to the controller.
- Audio is now flowing from the computer to the headset.



2.7 Ending streaming and cleaning up

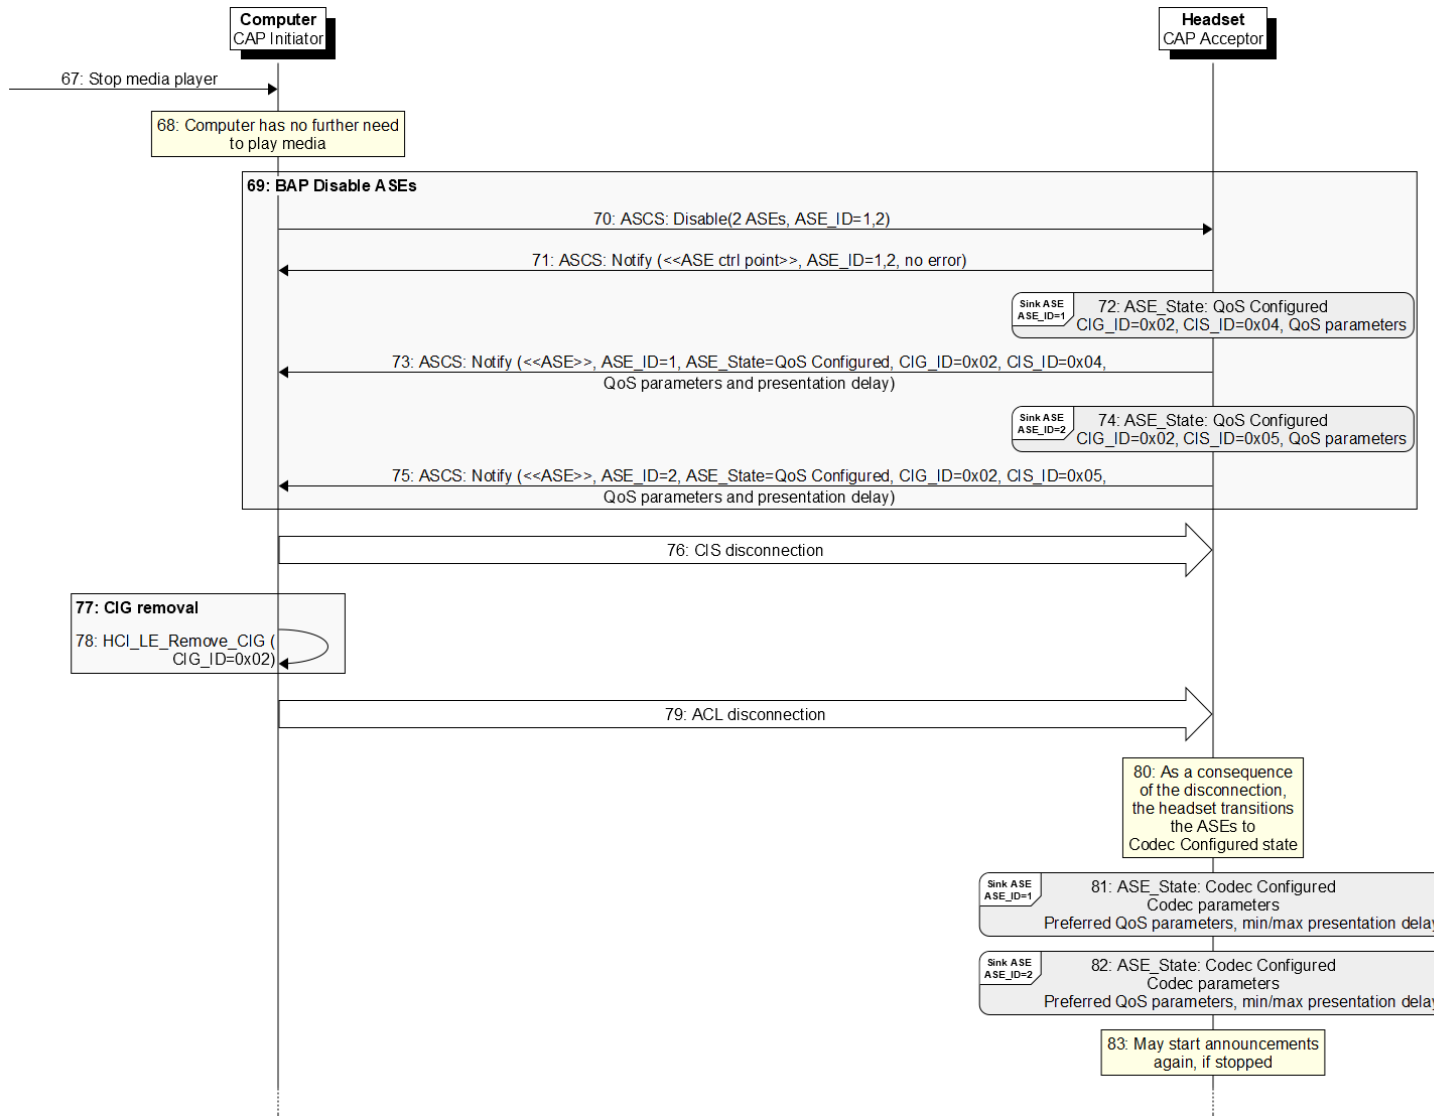


Figure 2.7: One-way unicast streaming from a computer to a headset (Part 7 of 7)



Figure 2.7 shows the following steps:

- Some external input (e.g., a button press from a user) stops the media player.
- The computer no longer has a need to stream audio and disables the ASEs.
- The headset decouples the ASEs from the CISes and notifies the new state.
- The computer and the headset disconnect the CISes (indicated by the block arrow) and remove the CIG.
- The ACL is disconnected. The computer may also choose to not disconnect, depending upon the use case or other circumstances.
- The headset transitions the ASEs to Codec Configured state. The headset may also transition the ASEs further, to Idle state (via the Releasing state), by doing an autonomous Release operation.
- The headset may start the announcements again (if it stopped them in the first place).



2.8 Considerations on setup latency

The number of GATT transactions characterize the setup latency. This section summarizes the number of GATT transactions.

For the example in this document, from the point in time where an ACL connection is established until the point in time at which audio transmission starts:

- There are 3 GATT transactions from the computer to the headset (Codec configuration, Quality of Service (QoS) configuration, and enabling of ASEs, with corresponding result notifications).
- There are from 8 to 13 additional notifications from the headset to the computer, depending upon whether the state of the sink PAC record and the ASEs has changed since the last connection.

3 References

- [1] Basic Audio Profile, Version 1.0.1 or later
- [2] Published Audio Capabilities Service, Version 1.0.1 or later
- [3] Audio Stream Control Service, Version 1.0
- [4] Common Audio Profile, Version 1.0
- [5] Low Energy Audio – Context Types and Availability, <https://www.bluetooth.com/bluetooth-resources/low-energy-audio-context-types-and-availability>
- [6] Bluetooth Assigned Numbers, <https://www.bluetooth.com/specifications/assigned-numbers>
- [7] Bluetooth Core Specification, Version 5.4 or later