WHAT IS CLAIMED IS:

1. A computer-implemented method for accessing a media title for playback, the method comprising:

receiving from a client device a first request for a first range of data included in a first encoded version of the media title, wherein the first encoded version of the media title is stored across a first set of files; wherein the first range of data comprises a range of byte offsets or a range of time offsets;

determining, from a first file extent index, a plurality of files that are included in the first set of files across which the first encoded version of the media title is stored, and from which the first range of data is to be accessed, wherein, the first file extent index maps an identifier for the first encoded version of the media title to an identifier for each file included in the plurality of files and wherein the first file extent index maps each identifier for each file to a given range of data that is stored in the file;

retrieving the first range of data from the plurality of files, wherein at least a portion of the first range of data retrieved from each of the plurality of files falls within the given range of data that is stored in the file; and

transmitting the first range of data to the client device in a response to the first request.

2. The computer-implemented method of claim 1, further comprising:

receiving a second request for a second range of data that begins at a first time offset in the first encoded version of the media title; and

retrieving the second range of data based on the first file extent index and a first segment index for a first file included in the first set of files, wherein the first segment index maps one or more time offsets in the first encoded version to one or more offsets in the first file.

3. The computer-implemented method of claim 2, wherein retrieving the second range of data based on the first file extent index and the first segment index comprises:

determining a first offset in the first file to which the first time offset is mapped in the first segment index; and

retrieving, from the first file, at least a portion of the second range of data beginning at the first offset.

4. The computer-implemented method of claim 2, wherein retrieving the second range of data based on the first file extent index and the first segment index comprises:

determining, based on the first file extent index, that the second range of data begins in the first file and ends in a second file;

retrieving, from the first file, a first portion of the second range of data beginning at a first offset to which the first time offset is mapped in the first segment index and ending at a second offset to which a first identifier for the first file is mapped in the first file extent index; and

retrieving, from the second file, a second portion of the second range of data beginning at a third offset to which a second identifier for the second file is mapped in the first file extent index and ending at a fourth offset to which a second time offset delineating an end of the second range of data is mapped in a second segment index for the second file.

5. The computer-implemented method of claim 1, further comprising:

receiving a second request for a second range of data included in a second encoded version of the media title, wherein the second encoded version of the media title is stored across a second set of files that overlaps with the first set of files;

determining, based on a second file extent index, one or more additional files that are included in the second set of files and from which the second range of data is to be accessed, wherein, for each additional file included in the second set of files, the second file extent index maps an identifier for the additional file to a given range of data that is both stored in the additional file and is included in the second encoded version of the media title; and

retrieving the second range of data from the one or more additional files, wherein at least a portion of the second range of data retrieved from

each of the one or more additional files falls within the given range of data that is stored in the additional file.

6. The computer-implemented method of claim 1, wherein retrieving the first range of data from the plurality of files comprises:

converting a starting offset in the first range of data into a first offset in a first file included in the first set of files; and

retrieving a first portion of the first range of data from the first file, wherein the first portion begins at the first offset in the first file.

7. The computer-implemented method of claim 6, wherein retrieving the first range of data from the plurality of files further comprises retrieving a second portion of the first range of data from a second file, wherein the second portion ends at a second offset in the second file that corresponds to an end of the first range of data.

8. The computer-implemented method of claim 6, wherein converting the starting offset in the first range of data into the first offset in the first file comprises:

identifying the first file based on the starting offset and a number of bytes spanned by one or more ranges of data included in the first file extent index; and

calculating the first offset based on a difference between the starting offset and the number of bytes spanned by the one or more ranges of data.

9. The computer-implemented method of claim 1, wherein the first set of files includes a first file that includes a first portion of the media title in a first language and a second file that includes the first portion of the media title in a second language.

10. The computer-implemented method of claim 1, wherein the first file extent index includes a length of the first encoded version of the media title.

11. A non-transitory computer readable medium storing instructions that, when executed by a processor, cause the processor to perform the steps of:

receiving from a client device a first request for a first range of data included in a first encoded version of a media title, wherein the first encoded

version of the media title is stored across a first set of files; wherein the first range of data comprises a range of byte offsets or a range of time offsets;

determining, from a first file extent index and a plurality of file mappings between one or more file identifiers and one or more file locations, one or more files that are included in the first set of files across which the first encoded version of the media title is stored, and from which the first range of data is to be accessed, wherein, the first file extent index maps an identifier for the first encoded version of the media title to an identifier for each file included in the plurality of files and wherein the first file extent index maps each identifier for each file to a given range of data that is stored in the file;

retrieving the first range of data from the plurality of files, wherein at least a portion of the first range of data retrieved from each of the plurality of files falls within the given range of data that is stored in the file; and

transmitting the first range of data to the client device in a response to the first request.

12. The non-transitory computer readable medium of claim 11, wherein the instructions further cause the processor to perform the steps of:

receiving a second request for a second range of data that begins at a first time offset in the first encoded version of the media title; and

retrieving the second range of data based on the first file extent index and a first segment index for a first file included in the first set of files, wherein the first segment index maps one or more time offsets in the first encoded version to one or more offsets in the first file.

13. The non-transitory computer readable medium of claim 12, wherein retrieving the second range of data based on the first file extent index and the first segment index comprises:

determining a first offset in the first file to which the first time offset is mapped in the first segment index; and

retrieving, from the first file, at least a portion of the second range of data beginning at the first offset.

14. The non-transitory computer readable medium of claim 12, wherein the first file includes a first segment from which the at least a portion of the second range of data is retrieved and a second segment that stores content included in a second encoded version of the media title.

15. The non-transitory computer readable medium of claim 11, wherein the instructions further cause the processor to perform the steps of:

receiving a second request for a second range of data included in a second encoded version of the media title; and

retrieving the second range of data from one or more additional files included in a second set of files based on a second file extent index for the second encoded version of the media title, wherein the second set of files overlaps with the first set of files.

16. The non-transitory computer readable medium of claim 11, wherein retrieving the first range of data from the plurality of files comprises:

converting a starting offset in the first range of data into a first offset in a first file included in the first set of files;

retrieving a first portion of the first range of data from the first file, wherein the first portion begins at the first offset in the first file; and

retrieving a second portion of the first range of data from a second file that follows the first file in the first file extent index.

17. The non-transitory computer readable medium of claim 16, wherein converting the starting offset in the first range of data into the first offset in the first file comprises:

determining the first file based on a number of bytes spanned by one or more ranges of data included in the first file extent index; and

calculating the first offset based on a difference between the starting offset and the number of bytes spanned by the one or more ranges of data.

18. The non-transitory computer readable medium of claim 11, wherein the given range of data is specified in one or more byte offsets within the first file extent index.

19. The non-transitory computer readable medium of claim 11, wherein the first file extent index includes at least one of a length of the first encoded version of the media title, a modification time of the first encoded version of the media title, or a file extension.

20. A computer-implemented method for retrieving a media title for playback, the method comprising:

transmitting to a server machine a first request for a first range of data included in a first encoded version of the media title, wherein the first encoded version of the media title is stored across a first set of files, the first range of data resides in a plurality of files included in the first set of files across which the first encoded version of the media title is stored, and a first file extent index maps an identifier for the first encoded version of the media title to an identifier for each file included in the plurality of files and wherein the first file extent index maps each identifier for each file to a given range of data stored in the file; wherein the first range of data comprises a range of byte offsets or time offsets;

receiving the first range of data from the server machine in response to the first request; and

causing media content corresponding to the first range of data to be output for playback.

21. The computer-implemented method of claim 20, further comprising:

transmitting to the server machine a second request for a second range of data that begins at a first time offset in the first encoded version of the media title, wherein the first time offset is mapped to a first offset in a first file within a first segment index for the first file; and

receiving the second range of data from the server machine in response to the second request.

22. The computer-implemented method of claim 21, wherein at least a portion of the second range of data is retrieved from the first file beginning at the first offset.

23. The computer-implemented method of claim 20, further comprising:

transmitting to the server machine a second request for a second range of data included in a second encoded version of the media title, wherein the second encoded version of the media title is stored across a second set of files that overlaps with the first set of files, the second range of data resides in one or more additional files included in the second set of files, and, for each additional file included in the second set of files, a second file extent index maps an identifier for the additional file to a given range of data stored in the additional file; and

receiving the second range of data from the server machine in response to the second request.

24. The computer-implemented method of claim 20, wherein at least a portion of the first range of data is retrieved from a first file included in the first set of files beginning at a first offset corresponding to a starting offset in the first range of data.

25. The computer-implemented method of claim 20, wherein the first set of files includes a first file that includes an original version of a portion of the media title and a second file that includes an alternative version of the portion of the media title.

26. The computer-implemented method of claim 25, wherein the alternative version comprises at least one of an edit to the original version, a regional version of the portion, or a version that meets criteria for a content rating.