

Choosing AVB Stream Sizes

... For the AVB device designer

2015 April 27

Jeff Koftinoff <jeff.koftinoff@gmail.com>

Table of Contents

CHOOSING AVB STREAM SIZES	1
Many Options	4
Trade-offs and optimizations	4
Maximize channel count limit per link / minimize stream count per device	5
Minimize channel count per stream / Maximize stream count per link	7
Maximize audio transport efficiency	8
Minimize the different stream sizes a device needs to support	8
Stream Dimension Calculations	10
AM824 Non Blocking Async packetization	10
48 kHz	10
96 kHz	11
192 kHz	12
AAF Streams, Quadlet aligned	13
48 kHz	13
96 kHz	14
192 kHz	15
Total Channels Per Network Link	16
48 kHz	16
96 kHz	17
48 kHz AM824 Max Channels Per Link	18
96 kHz AM824 Max Channels Per Link	20
192 kHz AM824 Max Channels Per Link	21
Channel Assignments	22
Conclusions	25

List of Tables

Table 1 - Stream and Channel Count AM824 Non Blocking, Async @ 48 kHz.....	10
Table 2 - Stream and Channel Count AM824 Non Blocking, Async @ 96 kHz.....	11
Table 3 – Stream and Channel Count AM824 Non Blocking, Async @ 192 kHz	12

Table 4 – Stream and Channel Count Quadlet aligned AAF @ 48 kHz.....	13
Table 5 - Stream and Channel Count Quadlet aligned AAF @ 96 kHz	14
Table 6 - Stream and Channel Count Quadlet aligned AAF @ 192 kHz.....	15
Table 7 – 48 kHz AM824 Non-Blocking Max Channels Per Link.....	18
Table 8 – 96 kHz AM824 Non-Blocking Max Channels Per Link.....	20
Table 9 – 192 kHz AM824 Non-Blocking Max Channels Per Link.....	21
Table 10 - Stereo Channel Layout.....	22
Table 11 - 5.1 Channel Layouts	23
Table 12 - 7.1 Channel Layouts	24
Table 13 – 8 Channels per Stream	25

List of Figures

Figure 1 – Total Channels per network link versus channels per stream @ 48 kHz.....	16
Figure 2 – Total Channels per network link versus channels per stream @ 96 kHz.....	17

Many Options

When designing an AVB Audio Device the designer has many stream format options available.

An AVB Stream can contain different media formats including encoding, sample rates and channel counts. Some formats that are supported include: integer 24 bit audio, integer 32 bit audio, floating point 32 bit audio, MIDI, SMPTE, MPEG video with audio, MJPEG video, SDI video with meta data, or control data.

The focus of this document will be 24 bit integer LPCM audio, 48 kHz and 96 kHz with various channel count options. The numbers in the tables in this document were calculated using the "AVB Bandwidth Calculator" available at <https://abc.statusbar.com/>. For AM824, "Non-blocking, Asynchronous Packetization" is used. For AAF, "32 bit audio (for quadlet aligned media) and default samples per frame" is used.

Trade-offs and optimizations

A device that sources 24 channels of audio onto the network, the designer has an option for 24 "single channel" streams, 12 "two channel" streams, 3 "eight channel" streams, or one "24 channel" stream.

The different options for transporting a set of audio channels over AVB streams each have different trade-offs, which are not necessarily obvious. Points to consider for optimization are to:

- Maximize channel count limit per link / minimize stream count per device
- Minimize channel count per stream / maximize stream count per link
- Maximize audio transport efficiency
- Minimize the different stream sizes a device needs to support

Maximize channel count limit per link / minimize stream count per device

For a system where many channels of audio need to be transported over a single link, the designer would typically need to maximize the channel count per stream to increase the transport efficiency and maximize the number of streams on the link to get as many channels as possible.

As the channel count per stream increases however it may decrease the number of streams that can fit; therefore the maximum channel count per stream is not necessarily the best option for maximum channel count per link.

Example:

On a gigabit link with 48 kHz streams one can fit 11 streams with 33 channels per stream, for a total of 363 audio channels.

With 48 kHz streams one can fit 12 streams with 32 channels per stream, for a total of 384 audio channels. Less channels per stream sometimes means more channels per network link.

Implementation of multiple stream sources and stream sinks in one device is expensive, either in CPU usage or FPGA real estate or both. In these cases it can be seen to be beneficial to reduce stream count and increase stream size to match.

When the designer maximizes the number of channels in a stream, the costs include:

- Additional requirements for Listeners to receive large channel count streams even when it only needs to listen to one audio channel, including dynamic audio mapping or channel selectors used to choose the appropriate audio channel.
- Requirements for Listeners to use gigabit Ethernet because the large channel count stream doesn't fit on a 100baseT link.
- Limitations of the number of streams that can fit on a network link, causing loss of flexibility of routing arbitrary channels to multiple Listeners elsewhere on the network.
- Requiring the grouped audio channels to come from a single Talker.
- Having streams that do not 'fit nicely' with the intended Listener channel count dimensions.

Minimize channel count per stream / Maximize stream count per link

For a system where many talkers need to provide individual channels to the network, or where maximum stream routing flexibility is needed, the designer would try to reduce the total number of audio channels per stream.

The benefits of using a small channel count stream include:

- Devices can provide network audio from a single microphone or stereo audio source without having to fill additional audio channels in a stream with duplicate channels or silence.
- Users can connect audio channels without having to consider selection of individual channels within a stream.
- Users can connect audio channels without quickly hitting bandwidth limitations of common network links
- Listeners can use 100baseT network ports

The costs of using a small channel count stream include:

- High channel count Talkers would be required to have a high number of stream sources, and implementation of a large number of stream sources in one device is expensive to implement in software or gate-ware.
- High channel count Listeners would be required to have a high number of stream sinks, and implementation of a a large number of stream sinks in one device is expensive to implement in software or gate-ware.
- The total number of streams that can be advertised on a network is limited - typically around 300 streams.
- The bandwidth efficiency when transporting small channel count streams is low; this greatly reduces the total number of channels that can be transported on one network link.

Maximize audio transport efficiency

Transport efficiency can be calculated by dividing the audio payload nominal bits per second by the total Ethernet frame transport bits per second.

Maximizing the audio transport efficiency can be accomplished by choosing a high sample rate and the most channels per stream possible for the desired link speed, however the combination of maximum values for sample rate and channels per stream are limited by the Ethernet maximum payload size.

The benefits for maximizing the audio transport efficiency are limited. Specifically, a stream is only forwarded down paths toward Listeners that are actively listening to the stream.

The costs for maximizing the audio transport efficiency include:

- Large channel counts in streams waste bandwidth going down the path to Listeners such as single channel speakers that are only interested in a single audio channel.
- Additional requirements for Listeners to receive large channel count streams even when it only needs to listen to one audio channel, including dynamic audio mapping or channel selectors used to choose the appropriate audio channel.
- Requirements for Listeners to use gigabit Ethernet because the large channel count stream doesn't fit on a 100baseT link.
- Limitations of the number of streams that can fit on a network link, causing loss of flexibility of routing arbitrary channels to multiple Listeners elsewhere on the network.
- Requiring the grouped audio channels to come from a single Talker.
- Having streams that do not fit nicely with the intended Listener channel count dimensions.
- Loss of efficiency when unpacking audio payloads which are unaligned to quadlet boundaries

Minimize the different stream sizes a device needs to support

Devices that support multiple channels and multiple sample rates may be more efficient to design with a single, common stream size across all combinations. The stream size that works with decent efficiency on all sample rates, AM824, and AAF, on Gigabit Ethernet as well as 100baseT is 8 channels

per stream. Table 13 – 8 Channels per Stream details the bandwidth, stream counts, and efficiency for AM824 and AAF stream formats at 48 kHz, 96 kHz, and 192 kHz.

Stream Dimension Calculations

AM824 Non Blocking Async packetization

48 kHz

Table 1 - Stream and Channel Count AM824 Non Blocking, Async @ 48 kHz shows the maximum stream count and channel count for various stream sizes for a single Ethernet link on both 100baseT and GigE network link speeds.

Table 1 - Stream and Channel Count AM824 Non Blocking, Async @ 48 kHz

Channels per Stream	SRP Bits/second per stream	Stream Count GigE	Stream Count 100baseT	Efficiency
1	6,592,000	113 (113 Ch.)	11 (11 Ch.)	24.49%
2	8,384,000	89 (178 Ch.)	8 (16 Ch.)	39.34%
8	19,136,000	39 (312 Ch.)	3 (24 Ch.)	72.18%
16	33,472,000	22 (352 Ch.)	2 (32 Ch.)	83.84%
24	47,808,000	15 (360 Ch.)	1 (24 Ch.)	88.62%
32	62,144,000	12 (384 Ch.)	1 (32 Ch.)	91.21%

96 kHz

Table 2 - Stream and Channel Count AM824 Non Blocking, Async @ 96 kHz

Channels per Stream	SRP Bits/second per stream	Stream Count GigE	Stream Count 100baseT	Efficiency
1	8,128,000	92 (92 Ch.)	9 (9 Ch.)	39.34%
2	11,456,000	65 (130 Ch.)	6 (12 Ch.)	56.47%
8	31,424,000	23 (184 Ch.)	2 (16 Ch.)	84.84%
16	58,048,000	12 (192 Ch.)	1 (16 Ch.)	91.21%
24	84,672,000	8 (192 Ch.)	0	93.96%
32	N/A	0	0	N/A

192 kHz

Table 3 – Stream and Channel Count AM824 Non Blocking, Async @ 192 kHz

Channels per Stream	SRP Bits/second per stream	Stream Count GigE	Stream Count 100baseT	Efficiency
1	11,200,000	66 (66 Ch.)	6 (6 Ch.)	56.47%
2	17,600,000	42 (84 Ch.)	2 (8 Ch.)	72.18%
8	56,000,000	13 (104 Ch.)	1 (8 Ch.)	91.21%
16	N/A	0	0	N/A
24	N/A	0	0	N/A
32	N/A	0	0	N/A

AAF Streams, Quadlet aligned

48 kHz

Table 4 – Stream and Channel Count Quadlet aligned AAF @ 48 kHz

Channels per Stream	SRP Bits/second per stream	Stream Count GigE	Stream Count 100baseT	Efficiency
1	5,824,000	128 (128 Ch.)	12 (12 Ch.)	26.67%
2	7,360,000	101 (202 Ch.)	10 (20 Ch.)	42.11%
8	16,576,000	45 (360 Ch.)	4 (32 Ch.)	74.42%
16	28,864,000	25 (400 Ch.)	2 (32 Ch.)	85.33%
24	41,152,000	18 (432 Ch.)	1 (24 Ch.)	89.72%
32	53,440,000	14 (448 Ch.)	1 (32 Ch.)	92.09%
48	78,016,000	9 (432 Ch.)	0	94.58%

96 kHz

Table 5 - Stream and Channel Count Quadlet aligned AAF @ 96 kHz

Channels per Stream	SRP Bits/second per stream	Stream Count GigE	Stream Count 100baseT	Efficiency
1	7,360,000	101 (101 Ch.)	10 (10 Ch.)	42.11%
2	10,432,000	71 (142 Ch.)	7 (14 Ch.)	59.26%
8	28,864,000	25 (200 Ch.)	2 (16 Ch.)	85.33%
16	53,440,000	14 (224 Ch.)	1 (16 Ch.)	92.09%
24	78,016,000	9 (216 Ch.)	0	94.58%
32	N/A	0	0	N/A

192 kHz

Table 6 - Stream and Channel Count Quadlet aligned AAF @ 192 kHz

Channels per Stream	SRP Bits/second per stream	Stream Count GigE	Stream Count 100baseT	Efficiency
1	10,432,000	71 (71 Ch.)	7 (7 Ch.)	59.26%
2	16,576,000	45 (90 Ch.)	4 (8 Ch.)	74.42%
8	53,440,000	14 (112 Ch.)	1 (8 Ch.)	92.09%
16	N/A	0	0	N/A

Total Channels Per Network Link

48 kHz

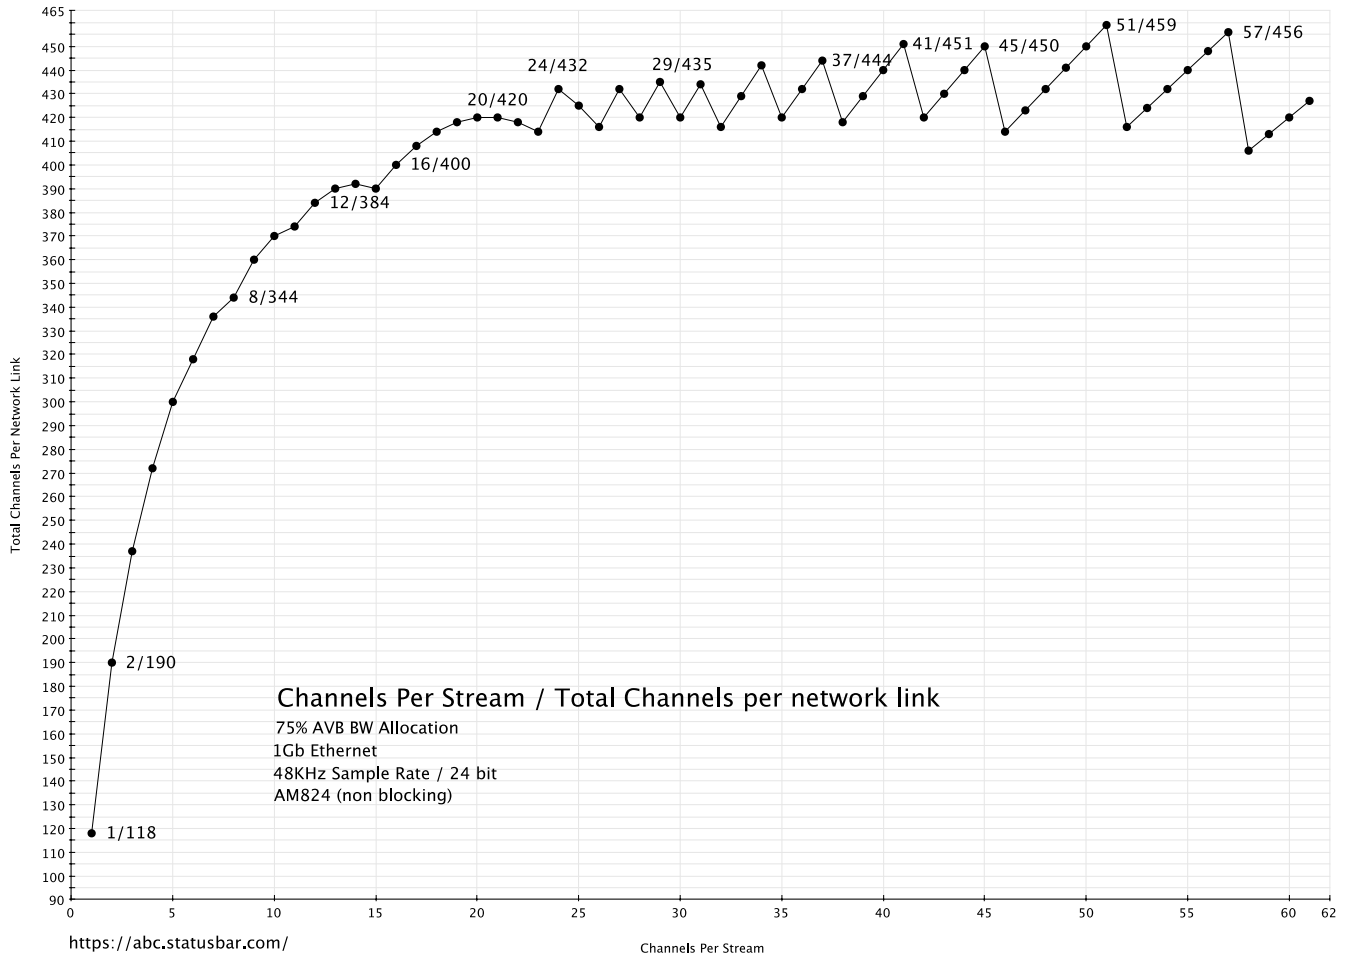


Figure 1 – Total Channels per network link versus channels per stream @ 48 kHz

Figure 1 – Total Channels per network link versus channels per stream @ 48 kHz shows the maximum number of channels possible on a single gigabit Ethernet link in one direction, using 48 kHz AM824 non-blocking synchronous packetization. It shows that 8 channels per stream provides a reasonable channel count and network link efficiency.

96 kHz

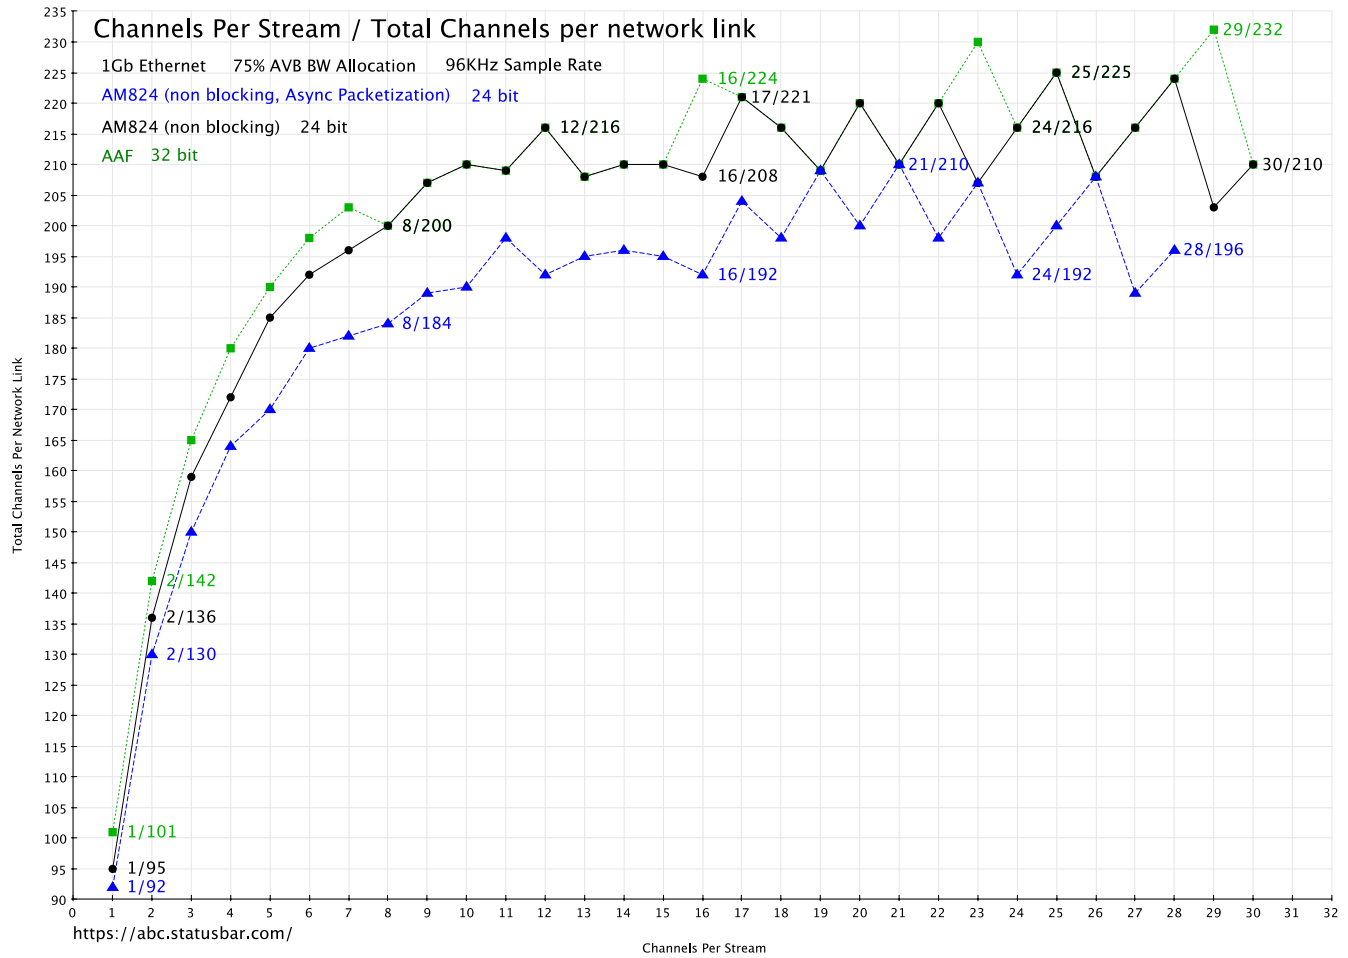


Figure 2 – Total Channels per network link versus channels per stream @ 96 kHz

Figure 2 – Total Channels per network link versus channels per stream @ 96 kHz shows the maximum number of channels possible on a single gigabit Ethernet link in one direction, using the following formats:

- AM824 non-blocking with asynchronous packetization
- AM824 non-blocking with synchronous packetization
- AAF with 24 bit audio or 32 bit audio, both with Quadlet alignment

This figure shows that 8 channels per stream provides a reasonable channel count and network link efficiency @ 96 kHz

48 kHz AM824 Max Channels Per Link

Table 7 – 48 kHz AM824 Non-Blocking Max Channels Per Link

Ch per Stream	Max Ch GigE	% diff from 8ch stream GigE	Max Ch 100base T	% diff from 8ch stream 100base T
1	118	-65.7%	11	-65.6%
2	190	-44.8%	18	-43.75%
8	344	0%	32	0%
16	400	+16%	32	0%
24	432	+25%	24	-25.0%
32	416	+20.9%	32	0%
48	432	+25.8%	N/A	N/A
60	420	+22.1%	N/A	N/A
61	427	+24.1%	N/A	N/A

Table 7 – 48 kHz AM824 Non-Blocking Max Channels Per Link shows that an 8 channel stream provides the high channel count for 100baseT network links with relatively high total channel count for GigE network links.

96 kHz AM824 Max Channels Per Link

Table 8 – 96 kHz AM824 Non-Blocking Max Channels Per Link

Ch per Stream	Max Ch GigE	% diff from 8ch stream GigE	Max Ch 100base T	% diff from 8ch stream 100base T
1	95	-52.5%	9	-43.75%
2	130	-35.0%	12	-25.0%
8	200	0%	16	0%
16	208	+4.0%	16	0%
24	216	+8.0%	N/A	N/A
25	225	+12.5%	N/A	N/A

Table 8 – 96 kHz AM824 Non-Blocking Max Channels Per Link shows that an 8 channel stream provides the highest channel count for both GigE and 100baseT network links with relatively high total channel count.

192 kHz AM824 Max Channels Per Link

Table 9 – 192 kHz AM824 Non-Blocking Max Channels Per Link

Ch per Stream	Max Ch GigE	% diff from 8ch stream GigE	Max Ch 100base T	% diff from 8ch stream 100base T
1	68	-34.6%	6	-25.0%
2	86	-17.3%	8	0%
8	104	0%	8	0%
12	108	+3.84%	N/A	N/A
13	104	0%	N/A	N/A
14	112	+7.69%	N/A	N/A
15	105	0.96%	N/A	N/A

Table 9 – 192 kHz AM824 Non-Blocking Max Channels Per Link Shows that the 8 channel stream is the only high channel stream count which also fits on a 100baseT network link and also provides a high total channel count for GigE network links.

Channel Assignments

IEEE 1722-rev1 Clause 8.4 defines channel assignments for common channel layouts. The audio channel assignments for stereo, 5.1 or 7.1 are mapped within an AVTP Audio stream according to Table 10 - Stereo Channel Layout, Table 11 - 5.1 Channel Layouts, or Table 12 - 7.1 Channel Layouts, respectively.

The channel ordering used in this standard is the same as used in SMPTE ST 2035:2009 Am1:2012, Audio Channel Assignments for Digital Television Recorders (DTRs) — Amendment 1. Other channel assignments can be configured by IEEE Std 1722.1-2013.

Devices which transmit audio based on content which can dynamically change between stereo, 5.1, or 7.1 formats shall always send an 8 channel stream and send silence on any unused audio channels.

Table 10 - Stereo Channel Layout

Channel	Description
0	Left
1	Right

Table 11 - 5.1 Channel Layouts

Channel	Description
0	Left
1	Right
2	Front Center
3	Low Frequency
4	Surround Left
5	Surround Right

Table 12 - 7.1 Channel Layouts

Channel	Description
0	Left
1	Right
2	Front Center
3	Low Frequency
4	Surround Left
5	Surround Right
6	Surround Back Left
7	Surround Back Right

Conclusions

Table 13 – 8 Channels per Stream

Stream Format	SRP Bits/second per stream	Stream Count GigE	Stream Count 100baseT	Efficiency
48 kHz AM824	19,136,000	39 (312 Ch.)	3 (24 Ch.)	72.18%
48 kHz AAF	16,576,000	45 (360 Ch.)	4 (32 Ch.)	74.42%
96 kHz AM824	31,424,000	23 (184 Ch.)	2 (16 Ch.)	84.84%
96 kHz AAF	28,864,000	25 (200 Ch.)	2 (16 Ch.)	85.33%
192 kHz AM824	56,000,000	13 (104 Ch.)	1 (8 Ch.)	91.21%
192 kHz AAF	53,440,000	14 (112 Ch.)	1 (8 Ch.)	92.09%

Table 13 – 8 Channels per Stream shows that the 8 channel stream is the one multi channel stream format which works with at reasonable to high efficiency with high total channel counts when used at 48 kHz, 96 kHz, and 192 kHz, as well as on 100baseT and GigE with all combinations.

The 8 channel stream, regardless of sample rate or packetization format provides the best trade-off for channel count, efficiency, and interoperability.

It is recommended that the AVB device designer use the following guidelines when deciding on the AVB Stream capabilities:

1. If your product is an AVB Talker for a single channel only, such as a single microphone, regardless of packetization format or sample rate, it is appropriate to use a 1 channel AVB Talker stream source.
2. If your product is an AVB Talker for two channels to 8 channels of audio, regardless of packetization format or sample rate, it is recommended to always use an 8 channel AVB Talker stream source. Use the channel assignments defined in the Channel Assignments section if necessary.
3. If your product is an AVB Talker of more than 8 channels of audio, regardless of packetization format or sample rate, it is recommended to always use 8 channel AVB Talker stream sources.
4. If your product is an AVB Listener, regardless of packetization format or sample rate it is recommended to use AVB Listener stream sinks which can sink both a single channel stream or an 8 channel stream.
5. If your product is an AVB Listener for more than 1 channel of audio, regardless of packetization format or sample rate, it is recommended to have a separate AVB Listener stream sink for each channel with each AVB Listener stream sink capable of listening to an 8 channel stream with the capability of selecting individual channels from each stream.