

THE BER OBJECTIVE FOR 400GE

David Ofelt – Juniper Networks

2013-05-17



I'M HERE FOR YOU...



HISTORY

BER target is a standard objective in IEEE 802.3

Most/all non-twisted pair PMDs from 1G to 100G have:

Support a BER better than or equal to 10^{-12} at the MAC/PLS service interface.

THE ISSUE

100x increase in bit rate with no improvement in specified BER

Historically there has been enough margin to make actual BER experienced be effectively zero.

Technology building blocks are not scaling at the rate we are scaling interface performance

- We are harvesting margin
- Engineering links from silicon to silicon

We are in danger of achieving our minimum BER target

IS 10^{-12} A PROBLEM?

At 100Gb/s a BER of 10^{-12} results in an error every 10s

Large boxes will get multiple errors per second

Whether this is a problem depends on how you view the topic...

BIT-CENTRIC VIEW

One view of the BER is that it states the reliability of a single bit

Rate doesn't matter

Interface count doesn't matter

Each bit has a 10^{-12} chance of being wrong

- Doesn't matter if that bit is on a 1GE or a 100GE
- Doesn't matter if that bit is arriving on a small box or a jumbo box

SYSTEM-CENTRIC VIEW

Customers look at all the errors on a machine

- Some ask us to explain every single bit error they see
- It doesn't matter the machine has a lot of interfaces
- It doesn't matter what speed the interfaces are

Frustrating, but understandable

In this view the BER target should get better with

- Interface speed
- Machine size

Even ignoring the machine size- the focus will be on the port

Irony is that small cheap slow boxes can easily look more reliable than huge, expensive, fast boxes.

OTHER METRICS

BER is the metric we use in our IEEE 802.3 objectives

- Well understood
- Commonly used to monitor health of links

Frame Error Rate (FER) also specified

- Some recent analysis (not comprehensive list):
 - IEEE 802.3bj
 - brown_3bj_02_0912.pdf
 - cideciyan_3bj_01a_0912.pdf
 - IEEE 802.3bm
 - anslow_03_0113_optx.pdf
 - brown_3bj_02_513.pdf
- Somewhat awkward due to packet-size dependency
- Nice because it maps to packet-level reliability

Should we continue to use BER as primary objective?

FEC

FEC complicates things since it fixes errors

- There is a pre-FEC BER and a post-FEC BER
- Only the post-FEC BER affects the user
- The pre-FEC BER indicates how hard we need to work on the channel

FEC can allow us to both:

- Make the channel/module/testing easier
- Provide a more reliable interface to the user

SUMMARY

We need to discuss the BER objective for 400GE

- Hopefully we won't cut and paste the previous objective without careful discussion

If we are going to see non-zero BER for modern interfaces, we need to start educating the industry on what to expect.