100GE OVERVIEW AND DEPLOYMENTS

Tim Nagy

AusNOG-05 – 15 September 2011

AGENDA: 100GE OVERVIEW AND DEPLOYMENTS

100GE (and 40GE) Standards Overview

Optics/Electrical Overview and Update

Announced 100GE Deployments

Operational/deployment experience with 100GE in Australia

AGENDA: 100GE OVERVIEW AND DEPLOYMENTS

100GE (and 40GE) Standards Overview

Optics/Electrical Overview and Update

Announced 100GE Deployments

Operational/deployment experience with 100GE in Australia

The Ethernet Standard

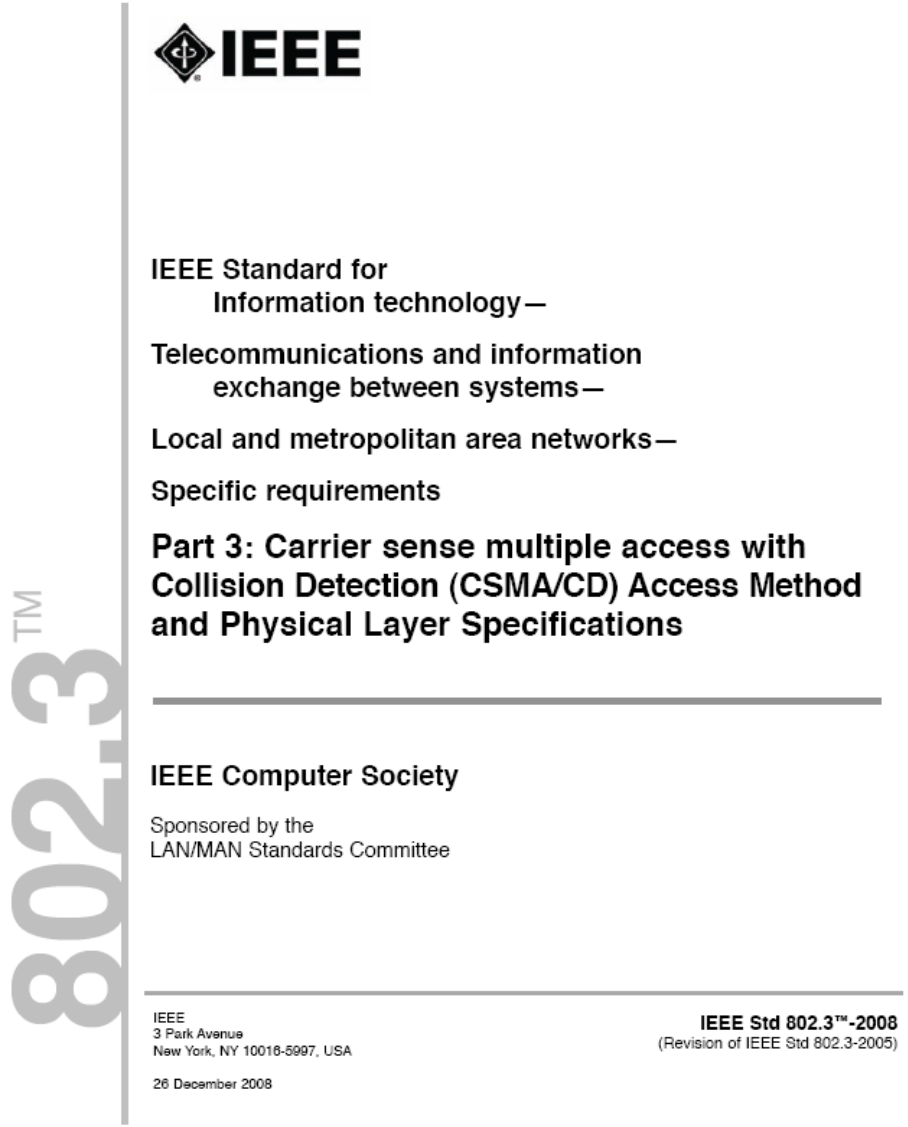

Free for individual download at
http://standards.ieee.org/getieee802/802.3.html

Prevailing Version
IEEE Std 802.3™-2008

2,977 pages

IEEE 802.3 Overview

Some of the major additions to IEEE Std 802.3 are identified in the marketplace with their project number:

- IEEE Std 802.3u™ added 100 Mb/s operation (also called Fast Ethernet)
- IEEE Std 802.3x™ specified full duplex operation and a flow control protocol
- IEEE Std 802.3z™ added 1000 Mb/s operation (also called Gigabit Ethernet)
- IEEE Std 802.3ae™ added 10 Gb/s operation (also called 10 Gigabit Ethernet)
- IEEE Std 802.3ah™ specified access network Ethernet (also called Ethernet in the First Mile)

These major additions are all now included in and are superseded by IEEE Std 802.3™-2008 and are not maintained as separate documents.

IEEE Std 802.3ba™ 40 Gb/s and 100 Gb/s Standard

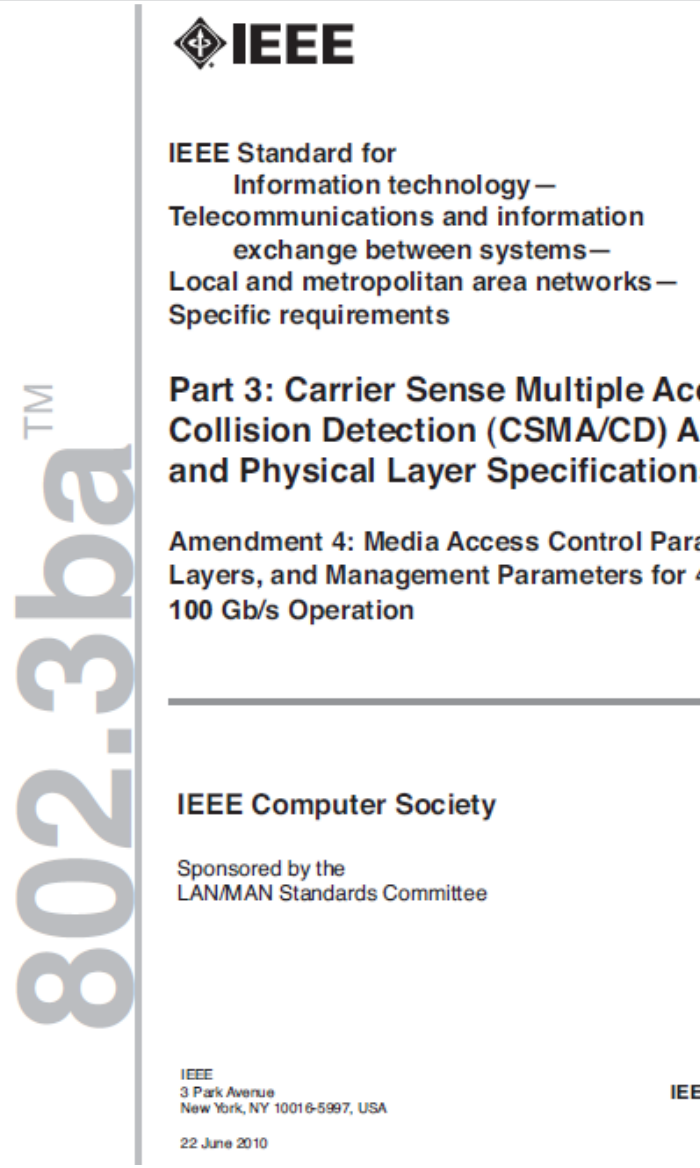

Free for individual download at
http://standards.ieee.org/getieee802/download/802.3ba-2010.pdf

Amendment to
IEEE Std 802.3™-2008

Published June 22, 2010

457 pages

WHAT DOES 100G REALLY MEAN?

Began broadcast
16 September 1955

460 Episodes
20,088 Days

@ ~20 Mbps (SD MPEG-2)
@ ~2 Mbps (SD MPEG-4)

~9.6 Gbps
= 40.1 Gbps

1 GbE
(1998)

10 GbE
(2002)

100 GbE
(2009)

AGENDA: 100GE OVERVIEW AND DEPLOYMENTS

100GE (and 40GE) Standards Overview

Optics/Electrical Overview and Update

Announced 100GE Deployments

Operational/deployment experience with 100GE in Australia

10G Module Review

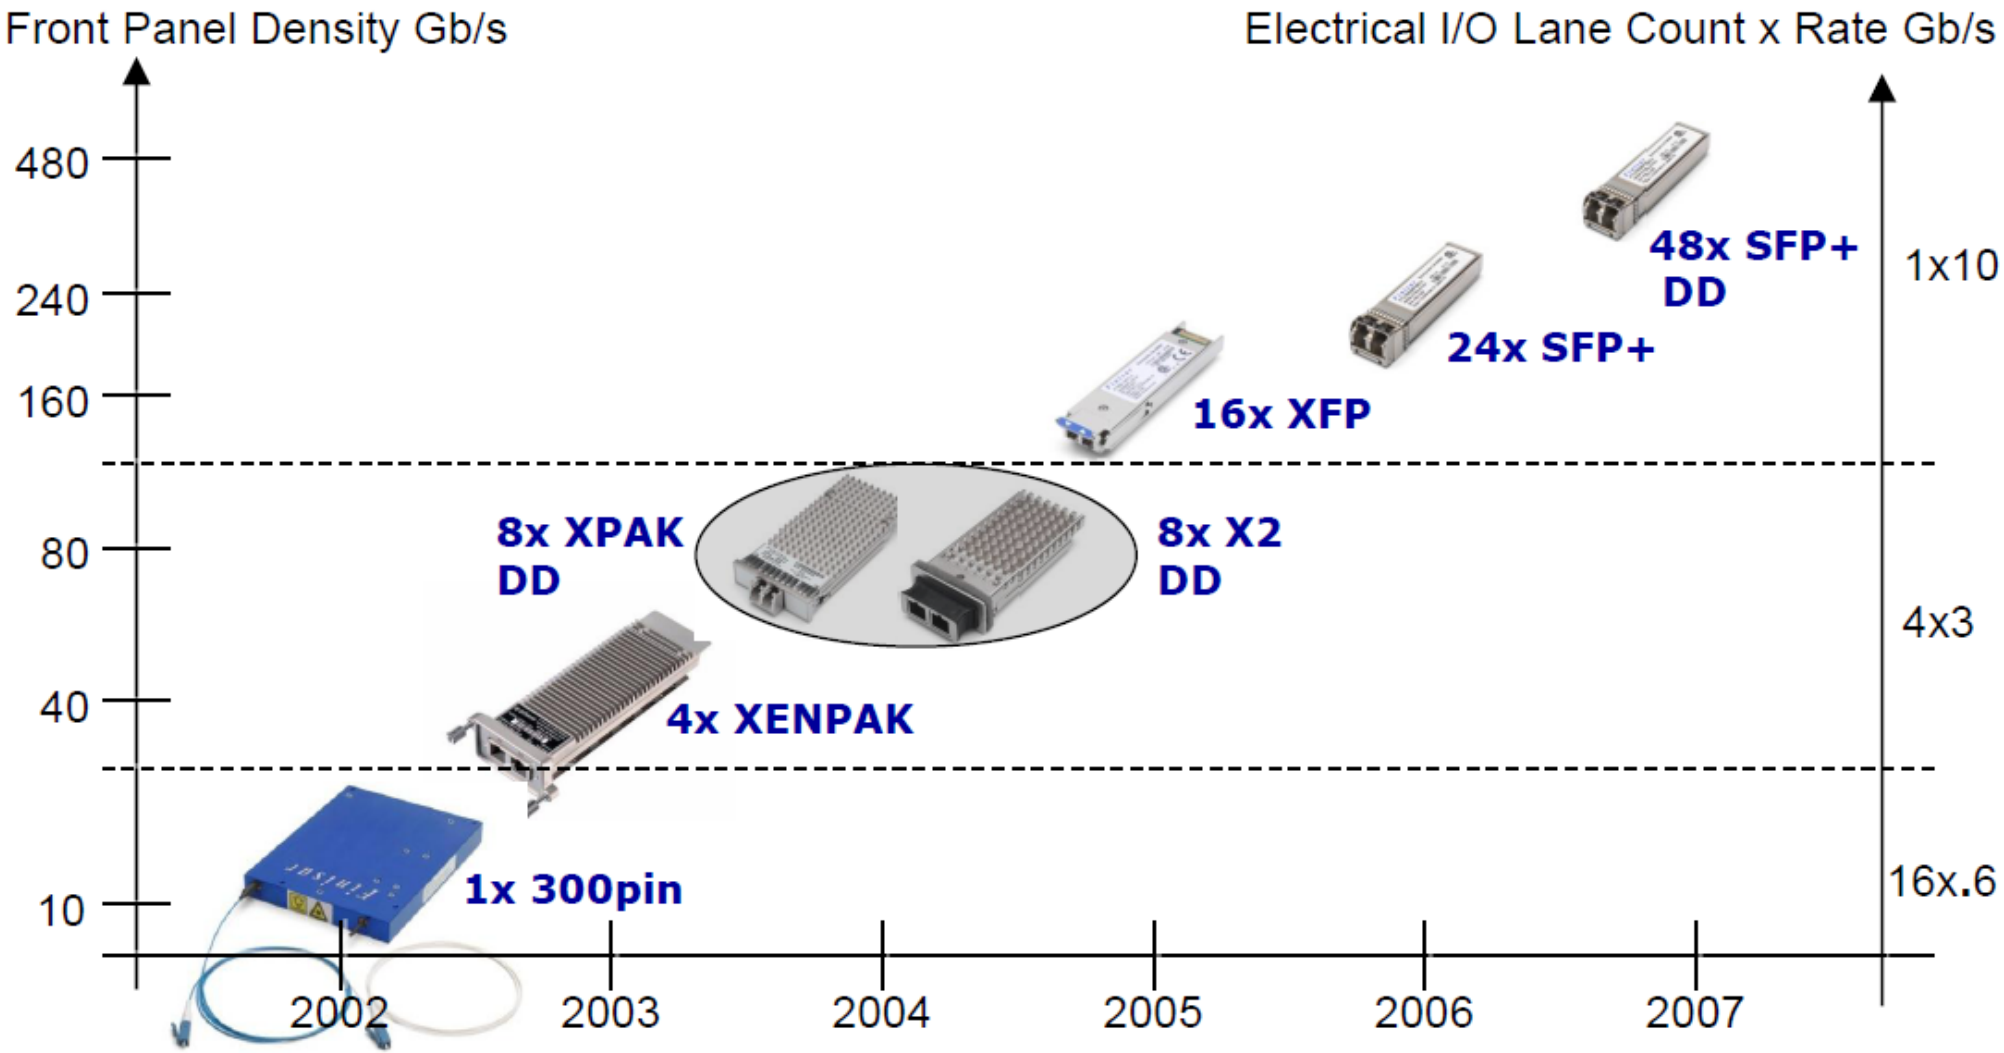

- Each new module enabled higher density, lower power, lower cost systems
- Only exception was duplication by competing X2 and XPAK modules

Source: www.cfp-msa.org/Documents/CFP_MSA_baseline_specifications.pdf

IEEE Std 802.3ba™ & 802.3bg™ Electrical and Optical Physical Medium Dependent (PMD) Sublayers

PMD Support	Nomenclature	
	40 Gbps	100 Gbps
Backplane (At least 1 m)	40GBASE-KR4	–No Objective–
Copper Cable Assembly (At least 7 m)	40GBASE-CR4	100GBASE-CR10
MMF (At least 100 m of OM3)	40GBASE-SR4	100GBASE-SR10
SMF (At least 2 km)	40GBASE-FR	–No Objective–
SMF (At least 10 km)	40GBASE-LR4	100GBASE-LR4
SMF (At least 40 km)	–No Objective–	100GBASE-ER4

100GE Module Options/Roadmap

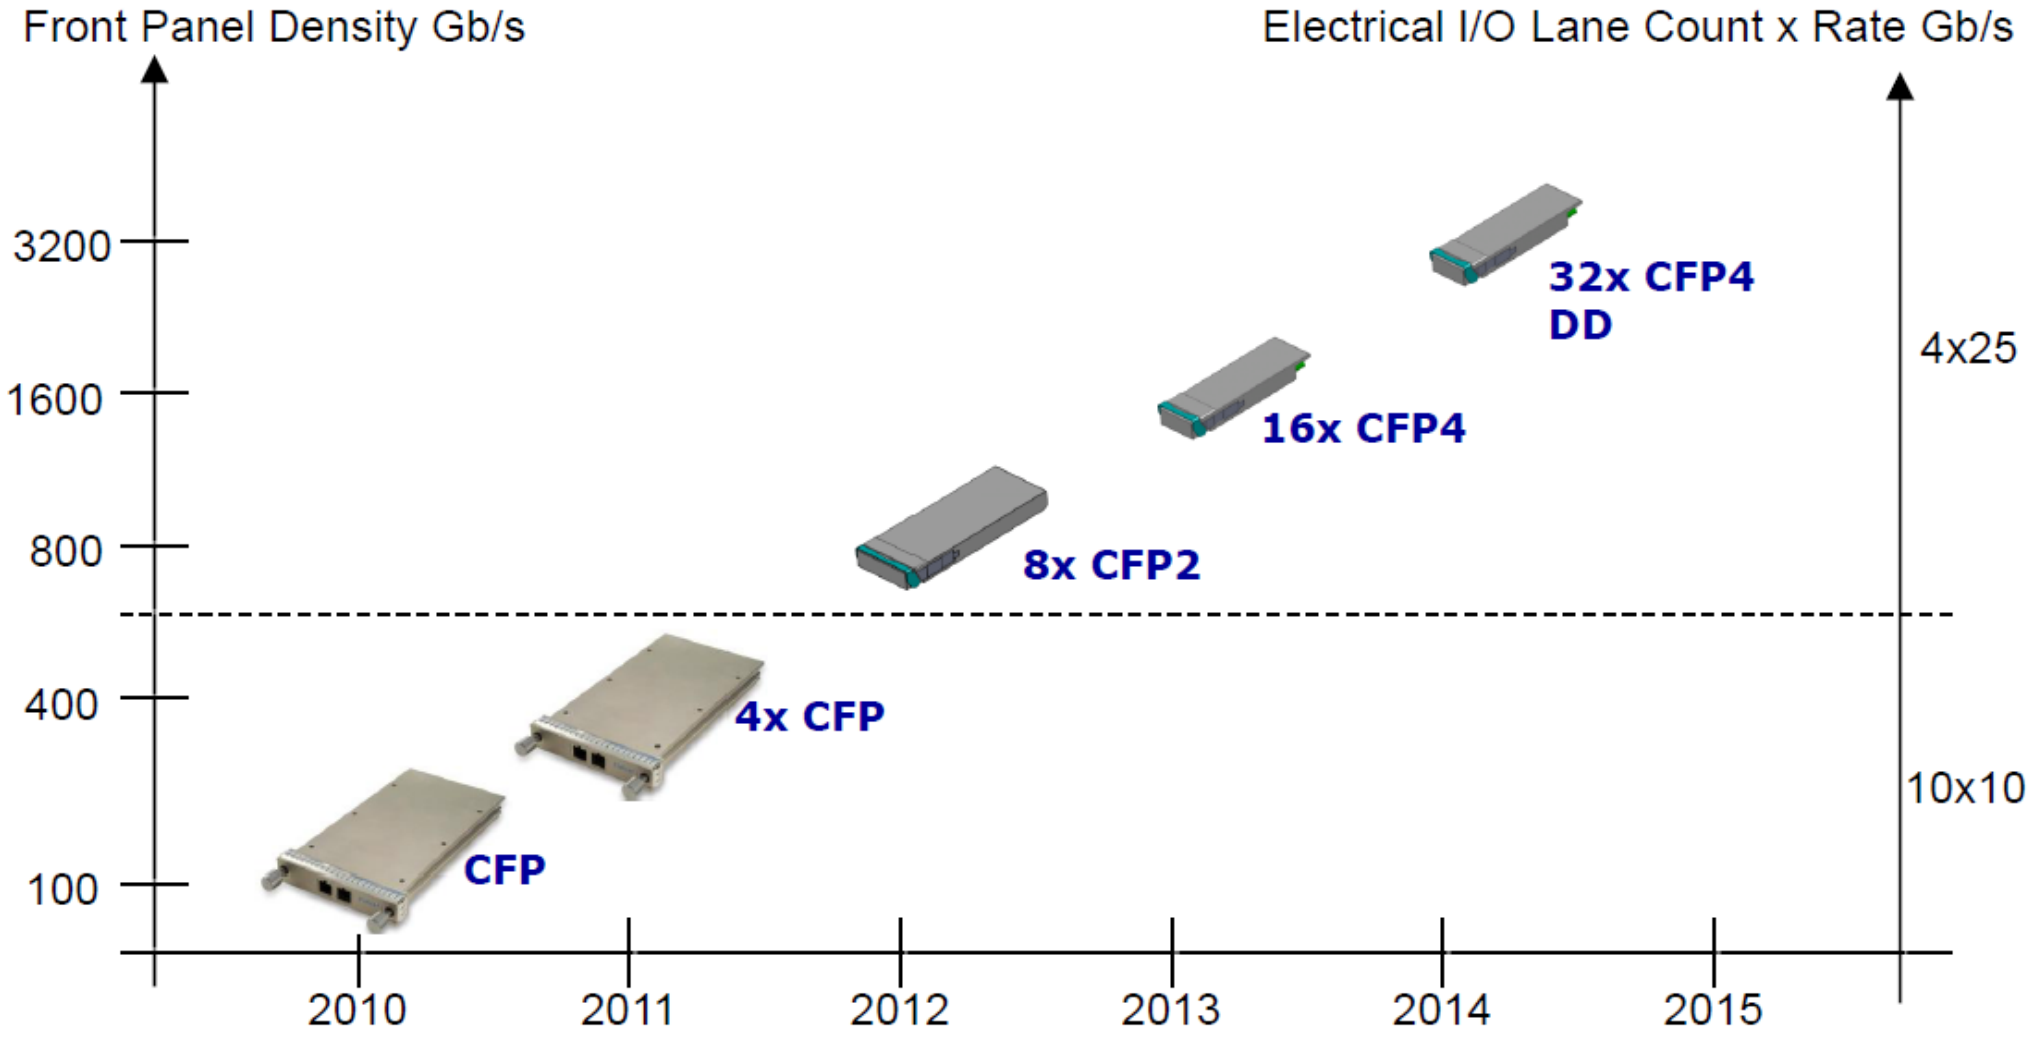

- Each new module enables higher density, lower power, lower cost systems
- Proposed XXP, extended CXP2 and QSFP2 modules duplicate this roadmap

Source: www.cfp-msa.org/Documents/CFP_MSA_baseline_specifications.pdf

PMD Support by Electrical Interface and Form Factor

PMD	Electrical Interface	Form Factor
40GBASE-CR4	XLPPI (4 x 10G)	QSFP
40GBASE-SR4	XLPPI (4 x 10G)	QSFP
	XLAUI (4 x 10G)	CFP
40GBASE-FR	XLAUI (4 x 10G)	CFP
40GBASE-LR4	XLPPI (4 x 10G)	QSFP
	XLAUI (4 x 10G)	CFP
100GBASE-CR10	CPPI (10 x 10G)	CXP
100GBASE-SR10	CPPI (10 x 10G)	CXP
	CAUI (10 x 10G)	CFP
10x10-2km	CAUI (10 x 10G)	CFP
100GBASE-LR4	CAUI (10 x 10G)	CFP
100GBASE-ER4	CAUI (10 x 10G)	CFP

10G – 120G Module Form Factor Comparison

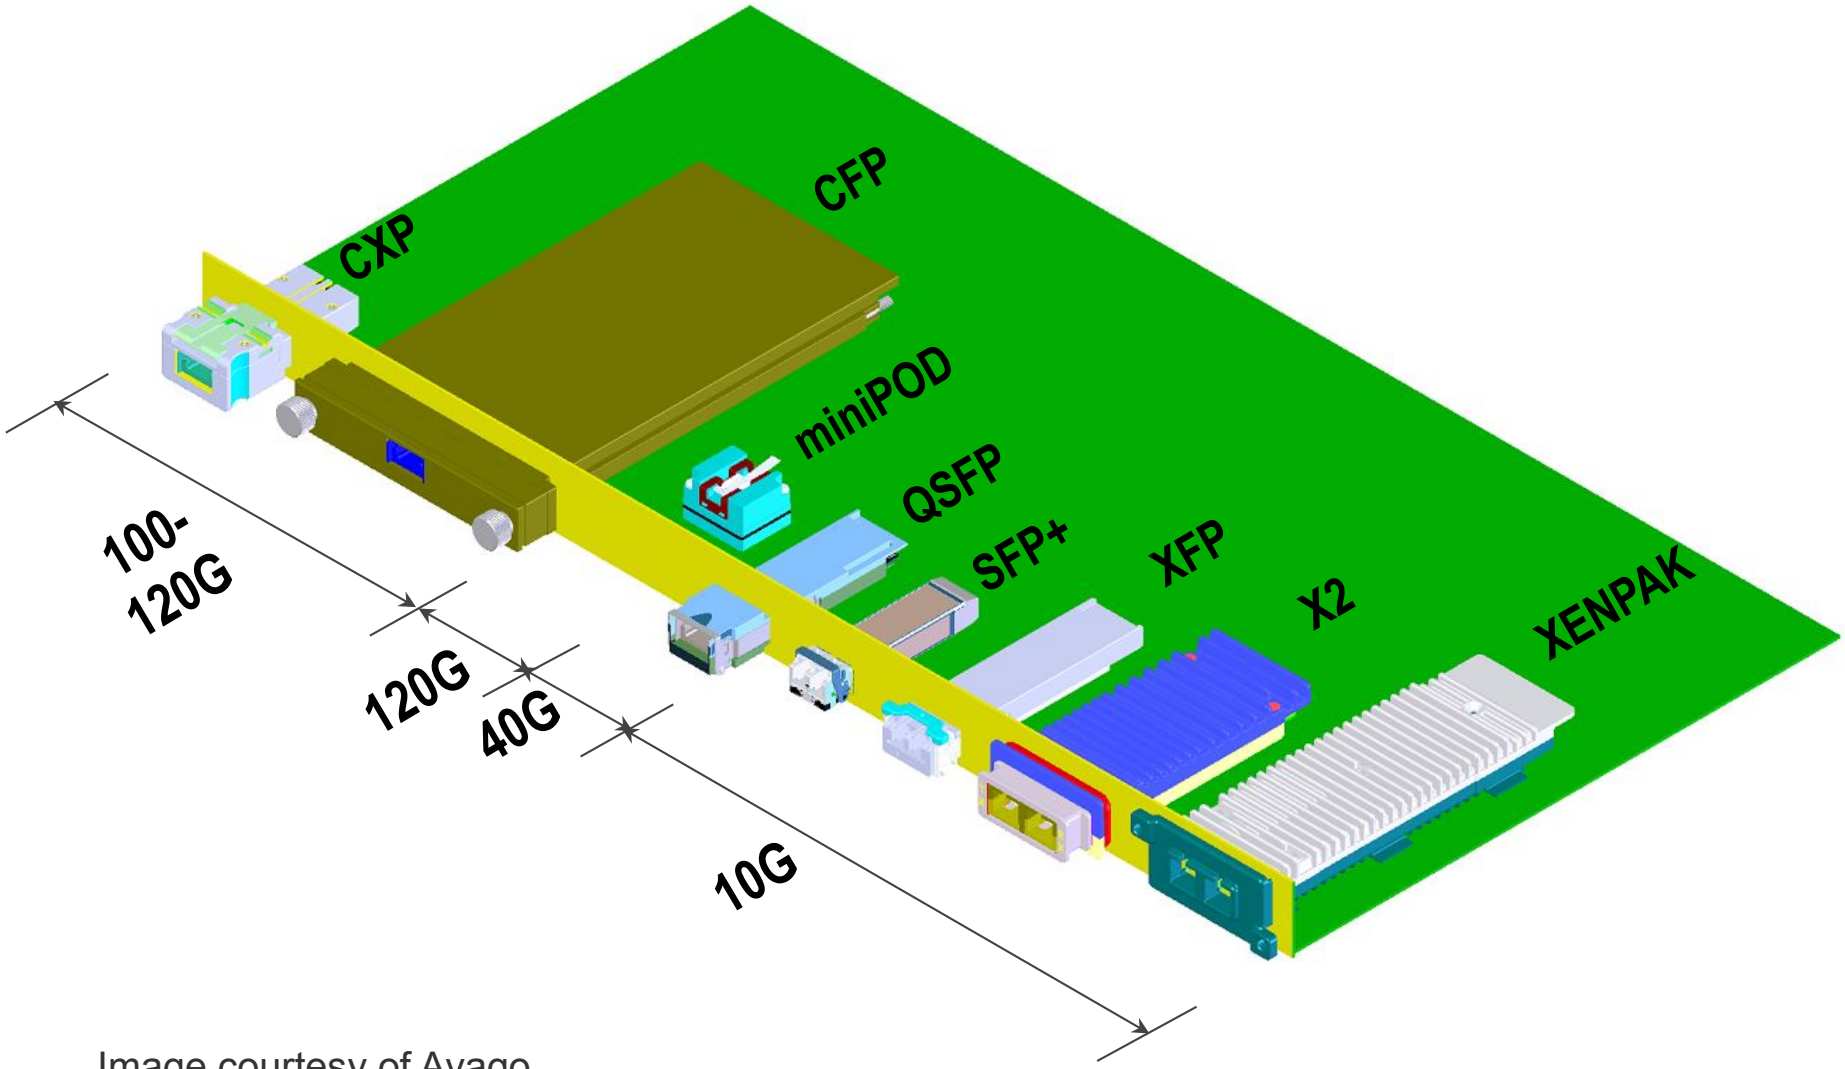

Image courtesy of Avago

CFP MSA

See http://www.cfp-msa.org/

- CFP MSA hardware specification, version 1.4
- CFP management specification, version 1.4

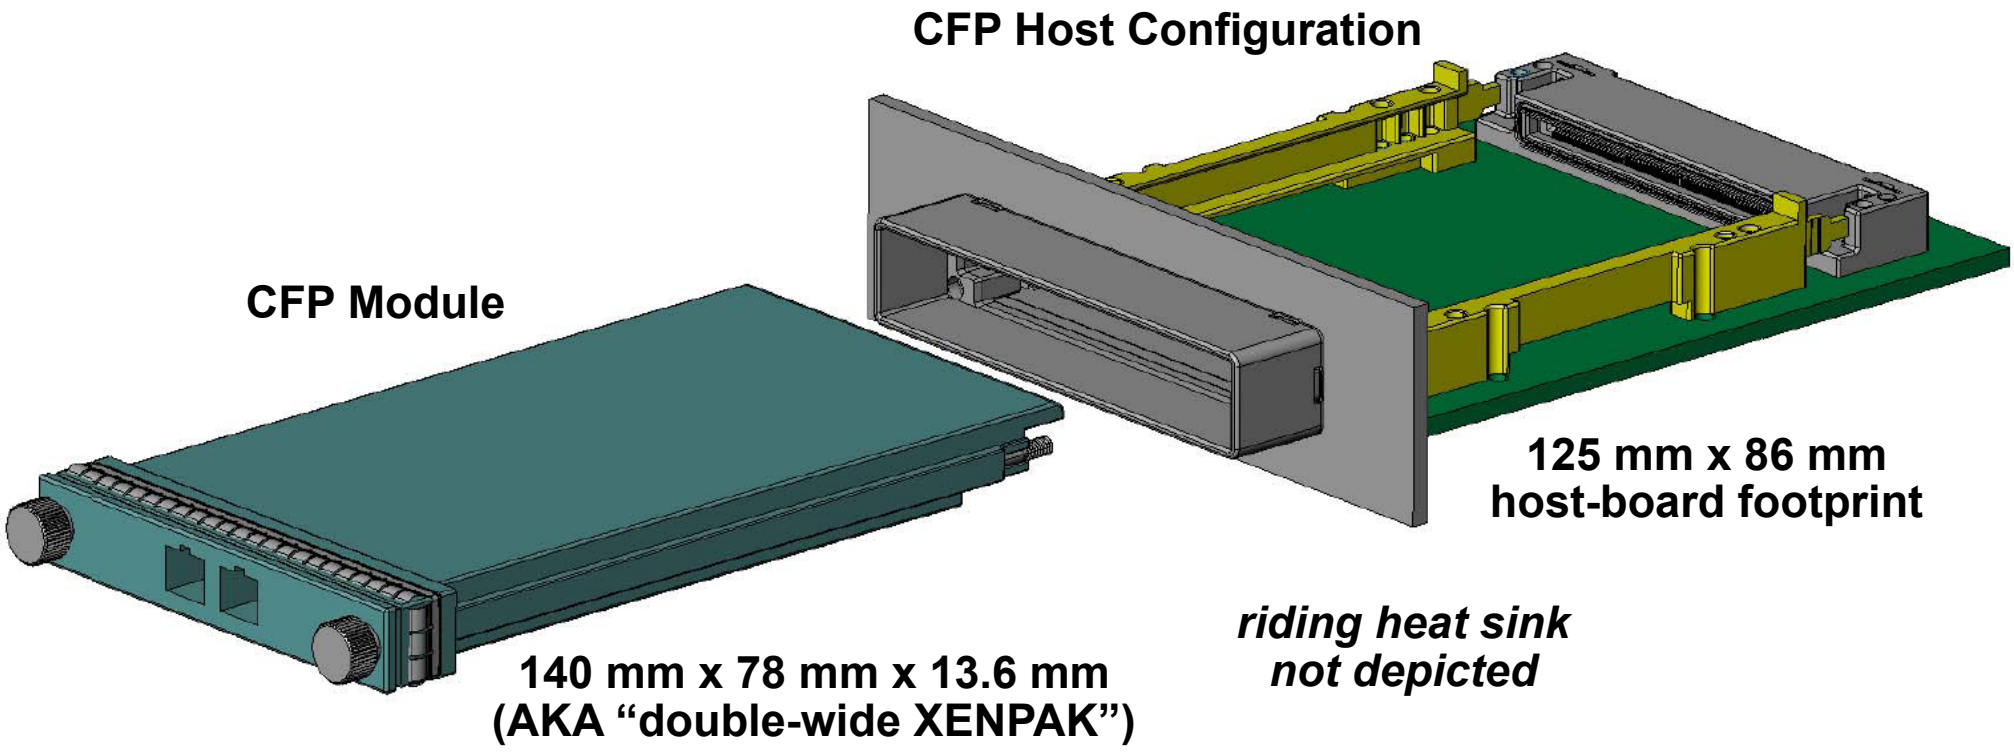

CFP-based 100GE

100GBASE-LR4 or –ER4

- pluggable CFP optics
- SMF SC-Duplex connector
- 10Km
- 40Km

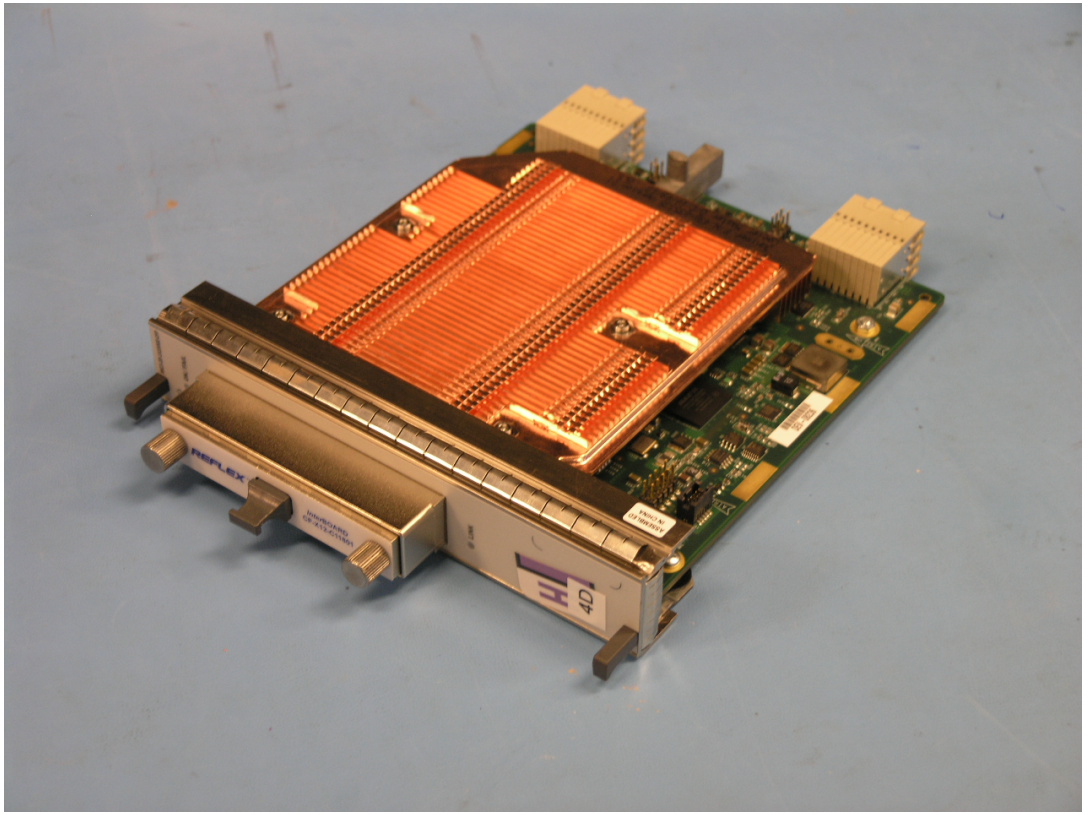

CXP-based 100GE

100GBASE-SR10

- 100m of MM Fiber
- 20 MM fibers in a “ribbon cable”

CXP Transceiver

- 850nm wavelength

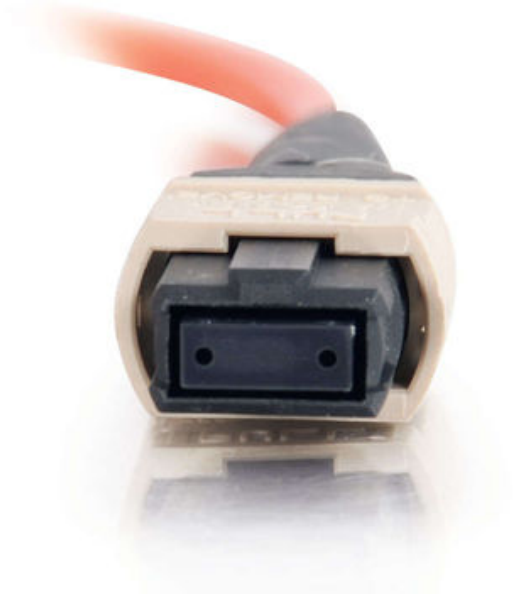

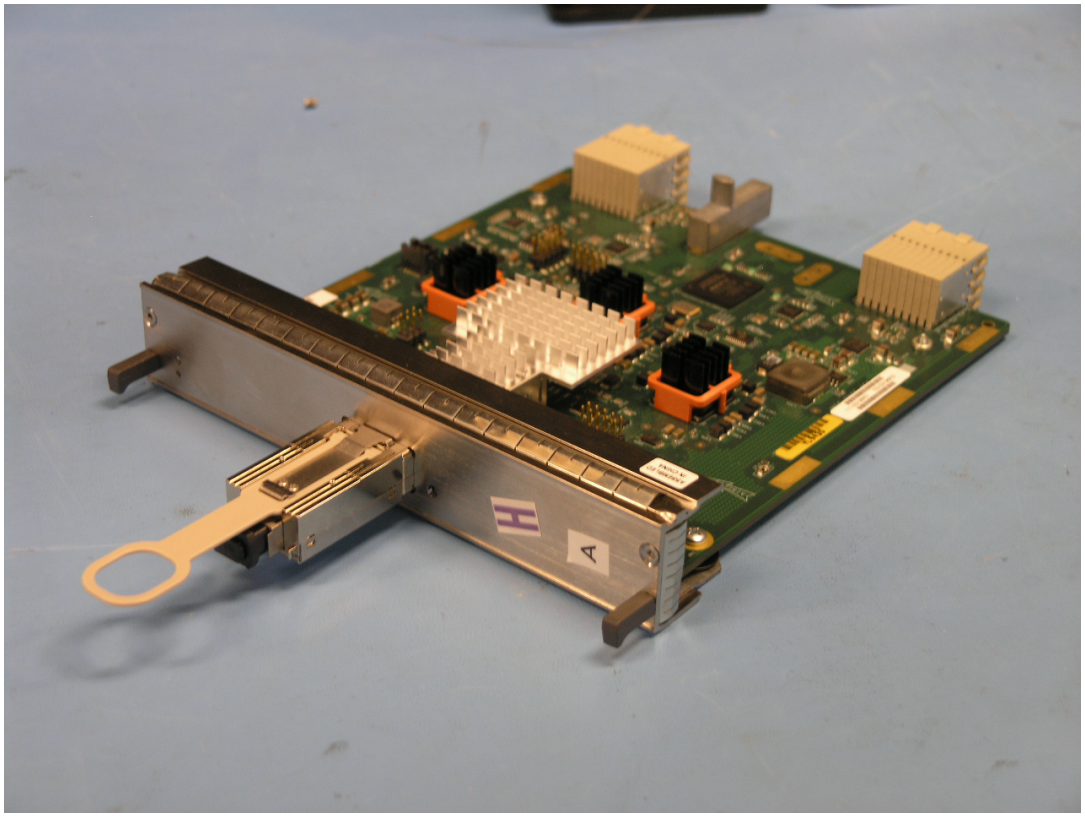

AGENDA: 100GE OVERVIEW AND DEPLOYMENTS

100GE (and 40GE) Standards Overview

Optics/Electrical Overview and Update

Announced 100GE Deployments

Operational/deployment experience with 100GE in Australia

SERVICE PROVIDERS WITH ANNOUNCED 100GE LIVE OR TRIAL NETWORKS

Amsterdam IX

AT&T

Comcast

iiNet

Internet2

JANET

Verizon

AGENDA: 100GE OVERVIEW AND DEPLOYMENTS

100GE (and 40GE) Standards Overview

Optics/Electrical Overview and Update

Announced 100GE Deployments

Operational/deployment experience with 100GE in Australia

100GE DEPLOYMENT SCENARIO

Two core routers with nx10GE connections for bandwidth; 100GE provided room for growth without adding connections to the LAG

Target was for just over 10km of fibre between two data centres in Sydney

Attenuation tests showed that the link met minimums for 100GBASE-LR4

Passive DWDM equipment in-line carrying multiple 10GE wavelengths, 100GE had to be carried over a single dark fibre pair in addition to the nx10GE wavelengths

100GE DEPLOYMENT STEPS

- Tested fibre attenuation, ready for deployment
- Passive DWDM filters specially made by the manufacturer to filter out the 4x25G wavelengths in the 100GE stream and present them on a single fibre to both routers
  - Ensured that nx10GE wavelengths were not overlapping the 4x25G wavelengths
  - Filters were passive and simply blocked all non-100GE wavelengths from coming through
- Equipment, especially the line card, is very heavy, so preferably use two people for installation
- Fibres plugged in at each site and links went active

CFP OPTICS AND LINE CARD

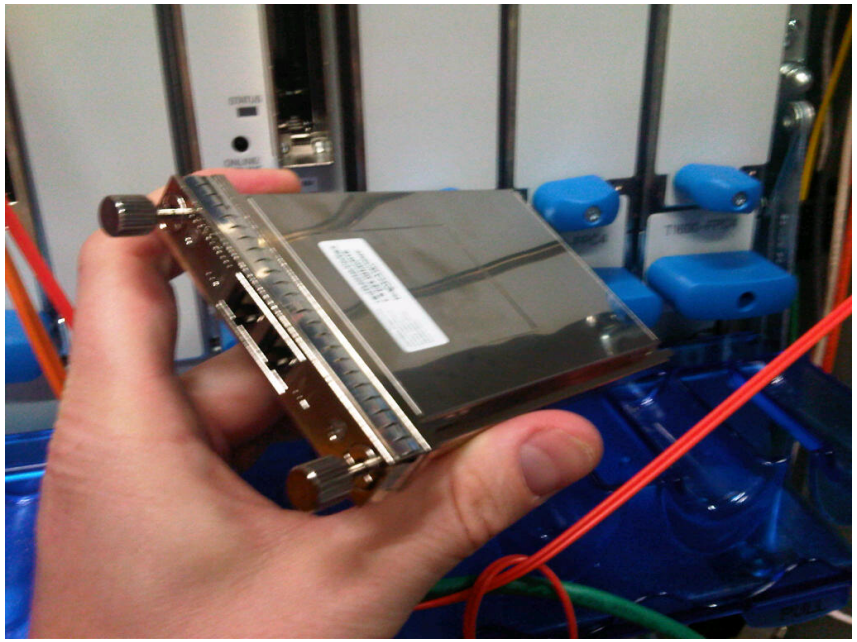

Weighs 35kg

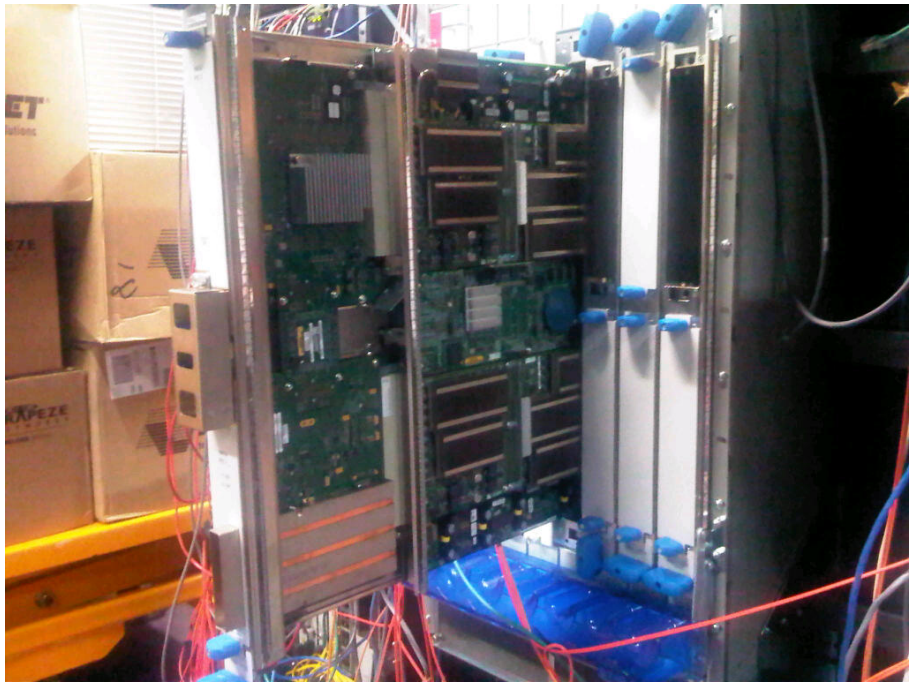

CONFIGURING THE INTERFACE

```
user@ROUTER> show chassis pic fpc-slot 7 pic-slot 1
FPC slot 7, PIC slot 1 information:
  Type                            100GE CFP
  ASIC type                       Brooklyn 100GE FPGA
  State                           Online
  PIC version                  1.3
  Uptime                          21 days, 5 hours, 54 minutes, 18 seconds

PIC port information:
                          Fiber                     Xcvr vendor
  Port  Cable type        type  Xcvr vendor         part number       Wavelength
  0     100GBASE LR4      SM    Opnext Inc.         TRC5E20ENF-SF150  1309 nm

user@ROUTER> show configuration interfaces et-7/0/0:0
gigether-options {
    802.3ad ae1;
}

user@ROUTER> show configuration interfaces et-7/0/0:1
gigether-options {
    802.3ad ae1;
}
```

CONFIGURING THE INTERFACE, CONT'D

```
user@ROUTER> show configuration interfaces ae1
traps;
mtu 9000;
aggregated-ether-options {
    minimum-links 2;
    link-speed 50g;
    lacp {
        active;
    }
}
unit 0 {
    family inet {
        no-redirects;
        address x.x.x.x/31;
    }
    family inet6 {
        address x:x:x:x::x/126;
    }
    family mpls;
}
```

UP AND RUNNING

Looking at the interface:

```
user@ROUTER> show interfaces ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 562
  Link-level type: Ethernet, MTU: 1614, Speed: 100Gbps, BPDU Error: None, MAC-
REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Disabled, Minimum links needed: 2, Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Current address: 28:c0:da:4e:07:f1, Hardware address: 28:c0:da:4e:07:f1
  Last flapped   : 2011-08-25 01:18:26 EST (1d 12:18 ago)
```

Note that first generation 100GE modules were made before ASICs could process 200 Gbps, so incoming traffic is split between two ASICs, hence the “ae” interface

CHECKING LIGHT LEVELS

```
user@ROUTER> show interfaces diagnostics optics et-7/0/0:0
Physical interface: et-7/0/0:0
    Module temperature                          :  47 degrees C / 117 degrees F
    Module voltage                              :  3.2510 V
…
 Lane 0
    Laser bias current                          :  44.908 mA
    Laser output power                          :  0.843 mW / -0.74 dBm
    Laser temperature                           :  52 degrees C / 126 degrees F
    Laser receiver power                        :  0.039 mW / -14.06 dBm
…
 Lane 1
    Laser bias current                          :  38.608 mA
    Laser output power                          :  1.028 mW / 0.12 dBm
    Laser temperature                           :  50 degrees C / 122 degrees F
    Laser receiver power                        :  0.037 mW / -14.32 dBm
…
 Lane 2
    Laser bias current                          :  49.597 mA
    Laser output power                          :  0.836 mW / -0.78 dBm
    Laser temperature                           :  58 degrees C / 136 degrees F
    Laser receiver power                        :  0.037 mW / -14.37 dBm
…
 Lane 3
    Laser bias current                          :  37.655 mA
    Laser output power                          :  1.052 mW / 0.22 dBm
    Laser temperature                           :  50 degrees C / 122 degrees F
    Laser receiver power                        :  0.044 mW / -13.53 dBm
```


everywhere