

Experience with Digital Optical Monitoring in NRENs



Marian Ďurkovič
md@bts.sk



What is DOM ?

- Ability to monitor certain parameters on optical transceivers without service interruption
- Data accessible via CLI and SNMP

```
#show interfaces transceiver
```

Port	Temperature (Celsius)	Voltage (Volts)	Current (mA)	Optical Tx Power (dBm)	Optical Rx Power (dBm)
Te1/1	30.5	0.00	32.9	-1.6	-4.7
Te1/2	29.9	0.00	60.8	2.0	-6.8
Te1/3	28.8	0.00	7.2	-2.9	-2.9
Te1/4	30.5	0.00	101.8	1.3	-21.8
Gi4/12	39.1	3.30	35.6	2.8	-1.2
Gi5/1	29.9	3.29	24.7	3.1	-12.6
Gi5/2	30.0	3.32	20.6	-5.5	-9.0

- Allows much better management of optical links

Ways to manage an optical link

A: monitoring by enduser

customer calls if his service is down or severely non-performant

+ cheap

- service is already degraded or unusable when we notice
- unsuitable for detecting failing backup links



B: close monitoring of interface statistics

input-errors indicate degraded links

- + fix problems before users get aware
- + frequent checking of interface statistics can be done by scripts
- once crc errors occur, link quality is already degraded, noticable performance problems likely
- unsuitable for detecting failing backup links

```
swiMY2>sh int g0/1
GigabitEthernet0/1 is up, line protocol is up
Hardware is BCM1250 Internal MAC, address is 0005.dc4b.001b (bia 0005.dc4b.001b)
Description: bidir CWDM-1550 to swiEL2
...
5 minute input rate 3009000 bits/sec, 448 packets/sec
5 minute output rate 1872000 bits/sec, 646 packets/sec
1859916423 packets input, 508652634 bytes, 543 no buffer
Received 15553493 broadcasts, 0 runs, 0 giants, 0 throttles
51 input errors, 51 CRC, 0 frame, 0 overrun, 0 ignored
0 watchdog, 15553336 multicast, 0 pause input
0 input packets with dribble condition detected
2004072000 packets output, 201000988 bytes, 0 underruns
```

```
From: NOC Account <noc@switch.ch>
Subject: New input errors on link MY2_G0/3
Date: 12. August 2007 19:25:03 GMT+02:00
To: NOC SWITCH <noc@switch.ch>
```

```
2585 interfaces
sorted by input errors and output drops
tunnel interfaces skipped
only interfaces which matches: ^GI^T
excluded links: swi6netCE1_g0/0, cs3_g7/4
print if in-errors >= 1
compare result with /home/noc/lan-mon/lan/check-linkerrors.lastcheck
2007-08-12 19:20 00:20 MY2 G0/3 in-error: 7 ie/s: 0.0 out-drops: 0 od/s: 0.0
```

```
All links with input errors >= 1:
2585 interfaces
sorted by input errors and output drops
tunnel interfaces skipped
only interfaces which matches: ^GI^T
excluded links: swi6netCE1_g0/0, cs3_g7/4
print if in-errors >= 1
2007-08-12 19:20 00:20 MY2 G0/3 in-error: 7 ie/s: 0.0 out-drops: 0 od/s: 0.0
```

```
interface descriptions:
MY2_G0/3 bidir CWDM-1530 to swiBR2
```

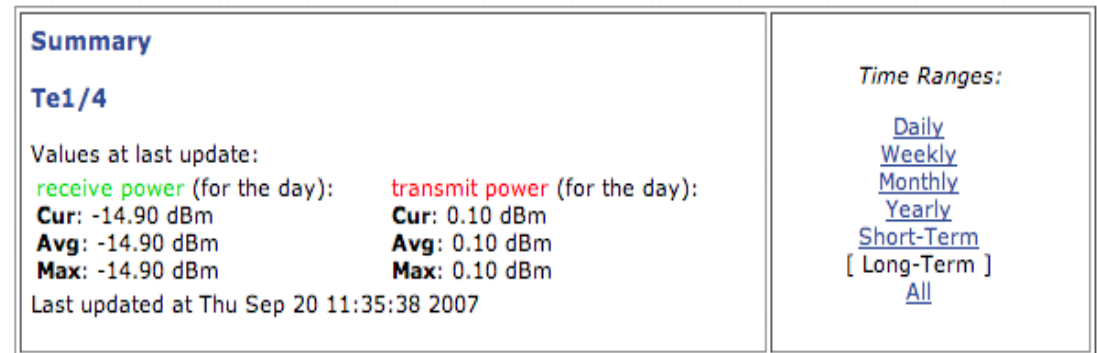
Ways to manage an optical link (2)

C: proactive monitoring on the optical level


use changing optical power levels as early indicators - before errors pop up

- + get aware of degrading fiber paths
- + get aware of transceiver problems
- + maintenance: test outcome of splicing/patching work
 - from remote
 - non-intrusively
- + works on unloaded backup links as well
- requires additional hardware & software capabilities

Graphs for swiEL2 - Te1/4



Yearly graph



300km 10GE link with DWDM XENPAKs Lausanne - Manno



DOM for local links?

- Sometimes considered unnecessary
- But technology is advancing rapidly
 - for 100m copper links, we now have:

```
#show cable-diagnostics tdr interface g9/5
```

Interface	Speed	Pair	Cable length	Distance to fault	Channel	Pair status
Gi9/5	1000	1-2	40 +/- 6 m	N/A	Pair B	Terminated
		3-4	40 +/- 6 m	N/A	Pair A	Terminated
		5-6	38 +/- 6 m	N/A	Pair D	Terminated
		7-8	40 +/- 6 m	N/A	Pair C	Terminated


- DOM needed to achieve decent L1 manageability for fibre links
 - LAN core / trunks



DOM Mechanics


SFF8472 - Diagnostic Monitoring Interface for Optical Transceivers

- Extension of serial ID interface for GBICs/SFPs
- Accessible the same way (2-wire I²C interface)
- Slightly modified for 10GE pluggables
 - derivatives of SFF8472 included in XENPAK/XFP MSAs
- DOM requires:
 - support in transceivers (sensors, data export)
 - support in IOS (data processing & presentation)
 - linecards not restricting DOM




Support in IOS

- First implemented on Cat4500
 - 12.1(20)EW, October 2003
- Support on Cat3750/3560/2970/2960
 - 12.2(20)SE, May 2004
 - Cat3750E/3650E: SFPs now, X2s in 12.2(44)SE
- Unofficial support on Cat6500
 - 12.2(17d)SXB2, July 2004
- Officially announced in April 2005
 - 12.2(18)SXE
- DOM available on all modern Catalyst switches



10GE Transceivers (1)


- XENPAKs
 - MSA rev 2.1 (Feb 2002) – no DOM defined
 - MSA rev 3.0 (Sep 2002) – DOM support added
 - LW/ZR/DWDM XENPAKs always MSA 3.0 with DOM
 - but for SR/LR/ER DOM support was unsure
 - customers received MSA 2.1 XENPAKs even in Sept 2005
- X2s - fine
 - based on XENPAK MSA 3.0
- XFPs - fine
 - DOM support in XFP MSA since the first draft



10GE Transceivers (2)


Situation almost fixed in 2007:

- XENPAK ER
 - EOS/EOL announced in March 2007
 - replaced by ER+ with DOM (Cisco P/N: 10-1888-04)
- XENPAK LR
 - EOS/EOL announced in June 2007
 - replaced by LR+ with DOM (Cisco P/N: 10-1838-04)
- XENPAK SR
 - no EOS/EOL announcement, no SR+ yet
 - DOM-capable SRs available from multiple vendors



GE Transceivers (1)

- GBICs
 - legacy form factor, no development anymore
 - DOM support in CWDM/DWDM versions only
- SFPs
 - all versions available with DOM support
 - but DOM-capable SFP-GE-S -L,-Z only for SPAs/ES20
 - GLC-SX-MM, GLC-LH-SM, GLC-ZX-SM (no DOM) for other linecards and/or other Catalyst switches
 - it means different SFPs are used even in the same chassis
 - unnecessarily complicated, sparing is a nightmare
 - solution quite simple: EOS/EOL for GLC-* versions
 - replacement already available: SFP-* versions



GE Transceivers (2)


- A closer look at some GLC-LH-SM SFPs:

Vendor Name	: FINISAR CORP.	Vendor Name	: CISCO-FINISAR
Vendor PN	: FTRJ-1319-7D-CSC	Vendor PN	: FTRJ1319P1BTL-C7
Date code	: 030508	Date code	: 060812

- From Finisar's application note AN-2030:

Digital diagnostic features are implemented in all Finisar SFP transceivers that contain a "D" in the part number suffix (for example, FTRJ -1319-7D-2.5), as well as DWDM and CWDM GBICs. All next generation Finisar SFPs utilizing the new part numbering scheme (e.g. FTRJ1621P1BCL) also have the same diagnostic capability.

- Those SFPs do have DOM hardware
 - some vendors say it's not feasible to produce two different HW versions
 - they just use different EEPROM image without DOM data
 - another strong argument to discontinue GLC-* versions



Back to IOS Support

- Upto 12.2(18)SXF10, DOM implementation was strictly based on SFF8472
 - all compliant transceivers worked fine
 - successful operation in NRENs for 2+ years
- Changes implemented in recent IOS versions
 - 12.2(33)SXH, 12.2(33)SRA4, 12.2(33)SRB
 - DOM stopped working for most of our transceivers
 - even those that work, don't work 100% correctly



Diversions from SFF8472

- DOM capability byte in transceiver's EEPROM is not honoured
 - IOS checks Cisco's own DOM-support list
 - DOM disabled if transceiver not listed there
- Alarm & warning thresholds are overridden with static values from IOS
 - threshold values from transceiver manufacturer no longer accessible
- Apparently an attempt to fix problems with bad transceivers – but too many unwanted side-effects

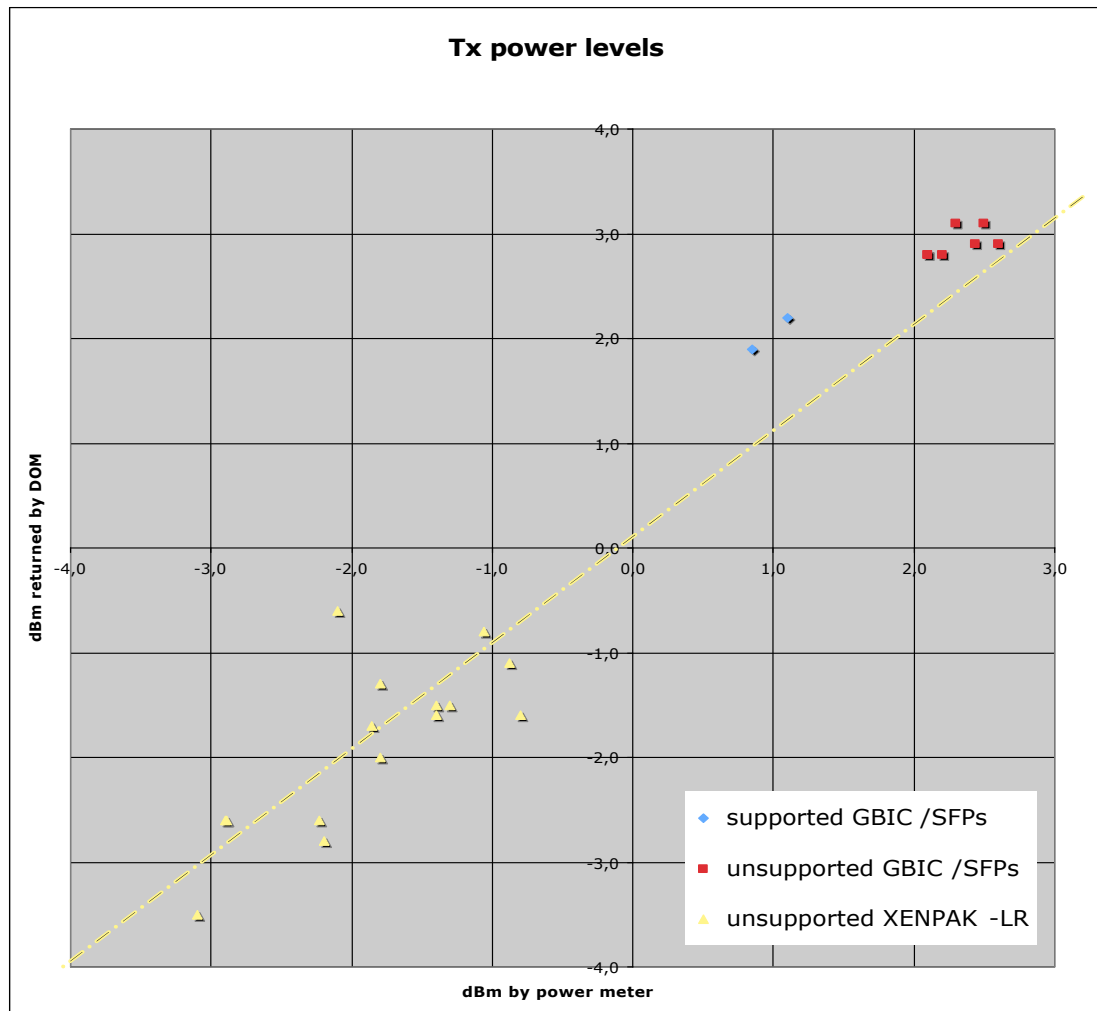


DOM-support list (1)

- LR+/ER+ XENPAKs with official DOM support only introduced in 2007, SR+ still not announced
 - DOM disabled on “old” SR/LR/ER XENPAKs even though it was working fine for years
 - also disabled on CWDM GBICs, some SFPs etc.
 - NRENs seriously hit by this – as usual, we are the early adopters and our XENPAKs are the old versions
- Introduces problem with all future transceivers
 - 10GBase-LRM not on the list as well
- Let’s verify accuracy of DOM data

Precision of DOM values - Tx

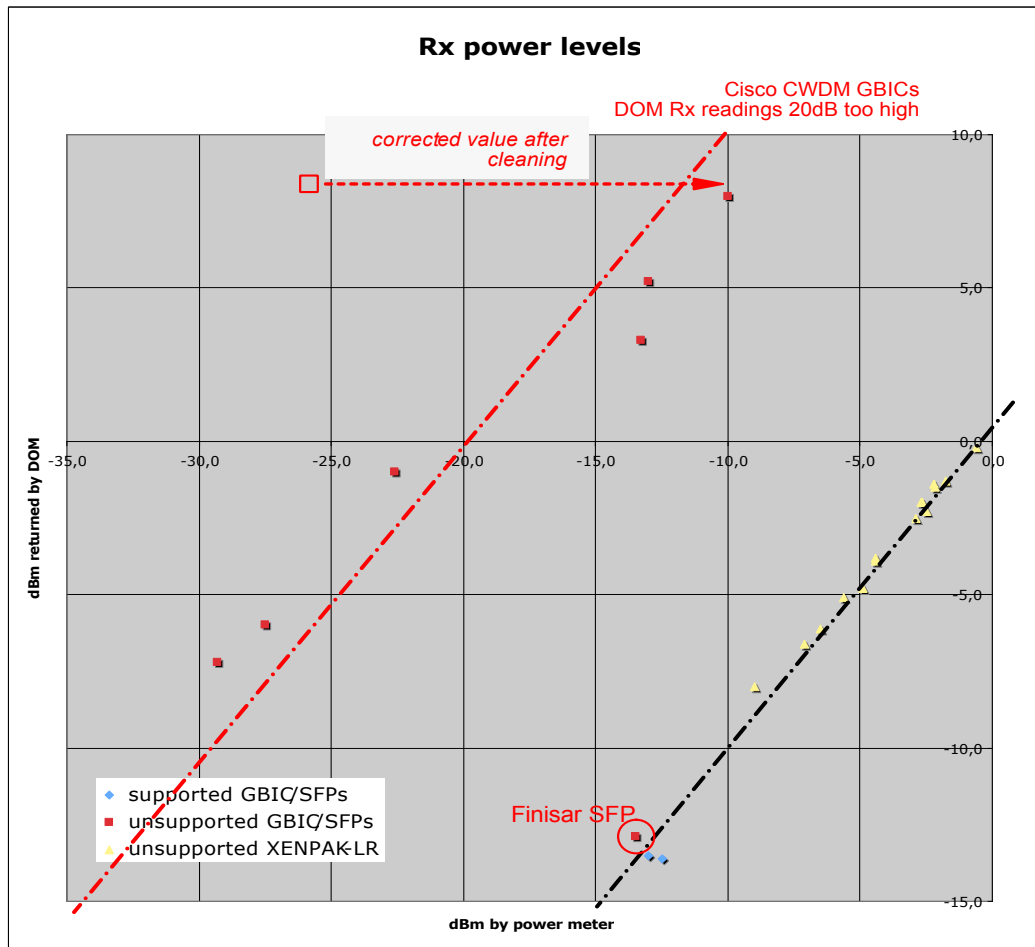
measured around Zurich on SWITCHlan backbone and between SWITCH's datacenters



- Tx readings of all plugins were reasonable
- accuracy mostly within +/- 1dBm limits
- two of the tested XENPAK-LR did not return any DOM values

Precision of DOM values - Rx

measured around Zurich on SWITCHlan backbone and between SWITCH's datacenters




- CWDM SFP: correct values, even from Finisar plugin
- CWDM GBIC: Rx readings are 20dB too high, regardless of IOS release
- DWDM GBIC: correct values
- XENPAK-LR: very precise, but two tested modules did not report any DOM values




DOM-support list (2)

- All tested transceivers return reasonable data
- CWDM GBICs are pre-SFF8472 devices
 - they report RX power in units of 1 nW
 - SFF8472 specifies 0.1 μ W units, i.e. 100-times more
 - need to subtract 20 dB from computed RX values
 - easy to adjust this in IOS
- Missing DOM on old transceivers bugfixed recently
- Alternative solution: introduce interface-level CLI to disable DOM if needed
 - in sync with “speed nonegotiate”, “no cdp enable” etc.




Static Thresholds (1)

- Some transceivers have incorrectly programmed alarm & warning thresholds in EEPROM
 - they might trigger false alarms / SNMP traps
- However, this problem is not fully fixable in IOS
 - several transceiver versions from several manufacturers are grouped under one name – e.g. XENPAK-LR
 - each of them might use different components with different threshold limits
 - it's therefore not possible to define universal limits
- Transceiver manufacturers need to fix this
 - they have done already, so recent & future pluggables should be OK



Static Thresholds (2)

- Any static threshold implemented in IOS means
 - at best only partial fix for broken transceiver
 - regression for all compliant transceivers
- Current implementation
 - no thresholds (thus no alarms) for laser current
 - “average” values for other parameters, some inaccurate
 - specific components not taken into account
- Some modifications needed
 - check validity of EEPROM thresholds first
 - if they look reasonable, prefer & use them
 - fallback to static thresholds otherwise
 - or add user-configurable thresholds from CatOS 8.6 ?



Linecard problems (1)

- DOM works fine on most linecards
 - even on really old ones like WS-X6516-GBIC
 - but not on WS-X6724-SFP / WS-X6748-SFP
- This is major problem for GE connections
 - WS-X6516-GBIC works, but only few GBICs with DOM
 - SFPs with DOM available, but linecards don't work
- Different answers from different people
 - hardware limitation, no workaround
 - possibly fixable in linecard firmware
 - further investigation needed




Linecard problems (2)

- The symptoms are really strange:
 - “show idprom int gig x/y” works
 - this means 2-wire I²C interface must be functional
 - DOM is accessible the same way
 - IOS tries to read it, but receives error from linecard firmware
- If no fix/workaround possible, please consider new HW revision of SFP linecards soon
 - otherwise there’s no DOM on Cat6500 for GE links



Summary

- DOM very useful technology worth developing
- Well supported by transceiver manufacturers
- Need to sort out a few things to be complete
 - EOS/EOL for DOM-incapable transceivers
 - fix DOM support on SFP linecards
 - fix a few IOS problems
- Please consider DOM in every future design
- As usual, NRENs will be glad to help



Acknowledgements

- ACONET: Christian Panigl, Harald Michl
- CESNET: Václav Novák, Josef Verich
- SWITCH: Felix Kugler, Simon Leinen

Questions ?