White Rabbit Switch advanced

White Rabbit
COLLABORATION

4 December 2024
CERN

Adam Wujek
White Rabbit Collaboration

Agenda

- Monitoring/Management
- Advanced CLI tools
- Configuration
- WRS: Interoperability with non-WR devices

Monitoring/Management

WRS: monitoring

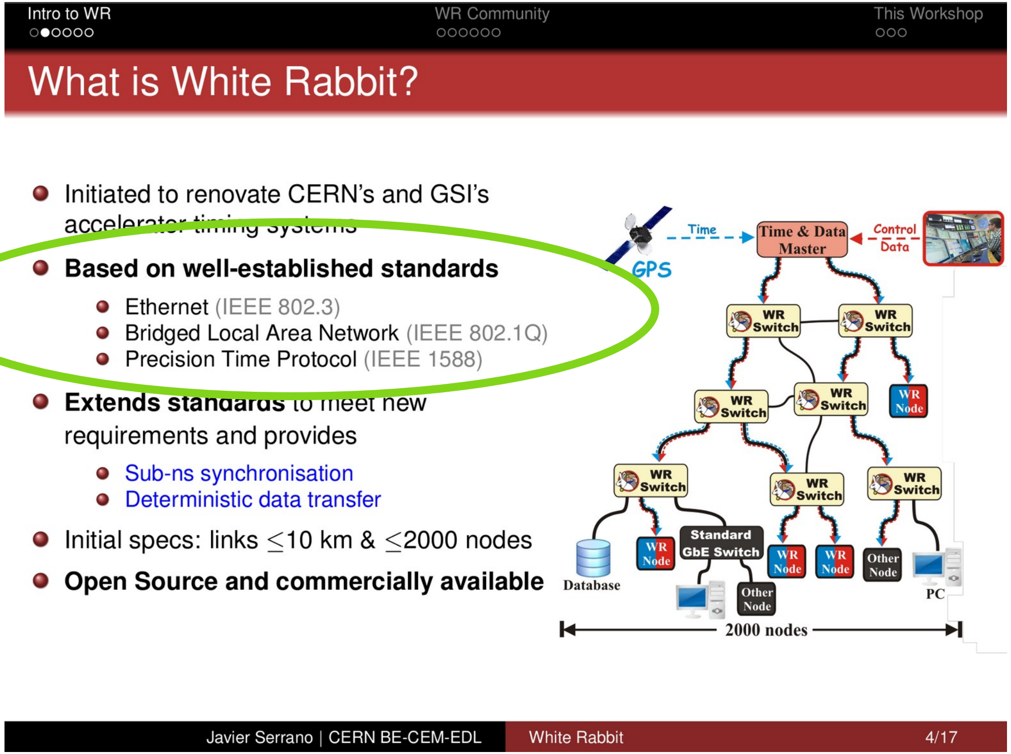

Intro to WR ○●○○○○ WR Community ○○○○○○ This Workshop ○○○

What is White Rabbit?

- Initiated to renovate CERN's and GSI's accelerator timing systems
- **Based on well-established standards**
  - Ethernet (IEEE 802.3)
  - Bridged Local Area Network (IEEE 802.1Q)
  - Precision Time Protocol (IEEE 1588)
- **Extends standards** to meet new requirements and provides
  - Sub-ns synchronisation
  - Deterministic data transfer
- Initial specs: links $\leq$10 km & $\leq$2000 nodes
- **Open Source and commercially available**

Javier Serrano | CERN BE-CEM-EDL White Rabbit 4/17

Let's use the power of standards!

LLDP

- Displays details of a device on the other side of a link

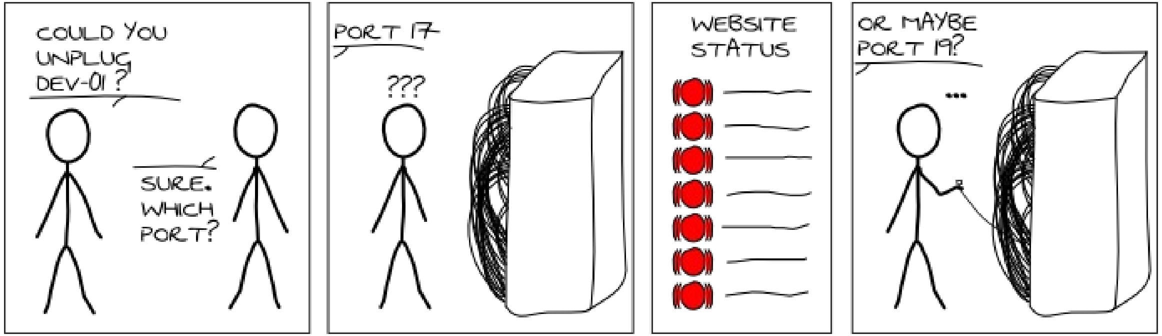

Source: https://lldpd.github.io/

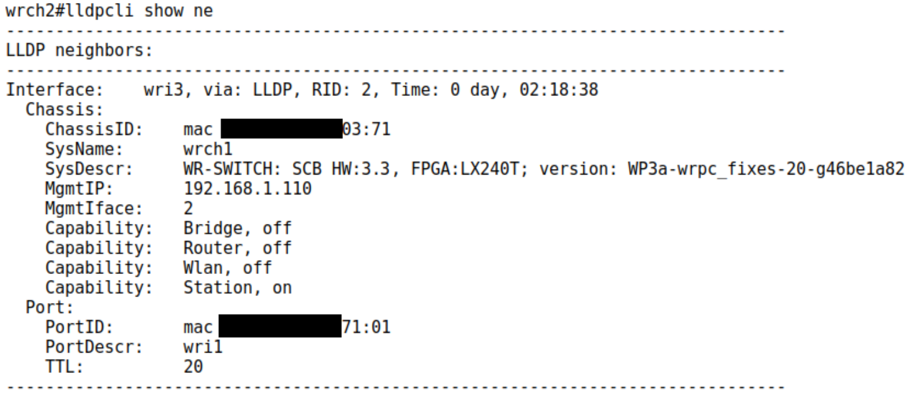

```
wrch2#lldpcli show ne
-------------------------------------------------------------------------------
LLDP neighbors:
-------------------------------------------------------------------------------
Interface:    wri3, via: LLDP, RID: 2, Time: 0 day, 02:18:38
  Chassis:
    ChassisID:    mac [redacted]03:71
    SysName:      wrch1
    SysDescr:     WR-SWITCH: SCB HW:3.3, FPGA:LX240T; version: WP3a-wrpc_fixes-20-g46be1a82
    MgmtIP:       192.168.1.110
    MgmtIface:    2
    Capability:   Bridge, off
    Capability:   Router, off
    Capability:   Wlan, off
    Capability:   Station, on
  Port:
    PortID:       mac [redacted]71:01
    PortDescr:    wri1
    TTL:          20
-------------------------------------------------------------------------------
```

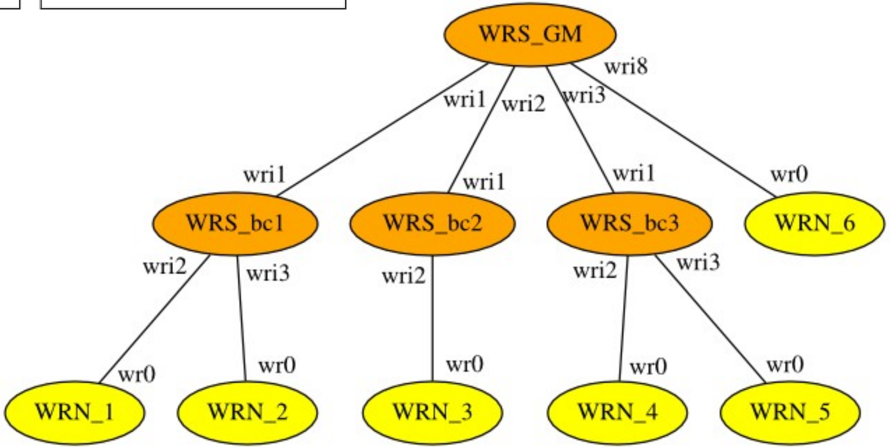

LLDP

- Real life example: GSI

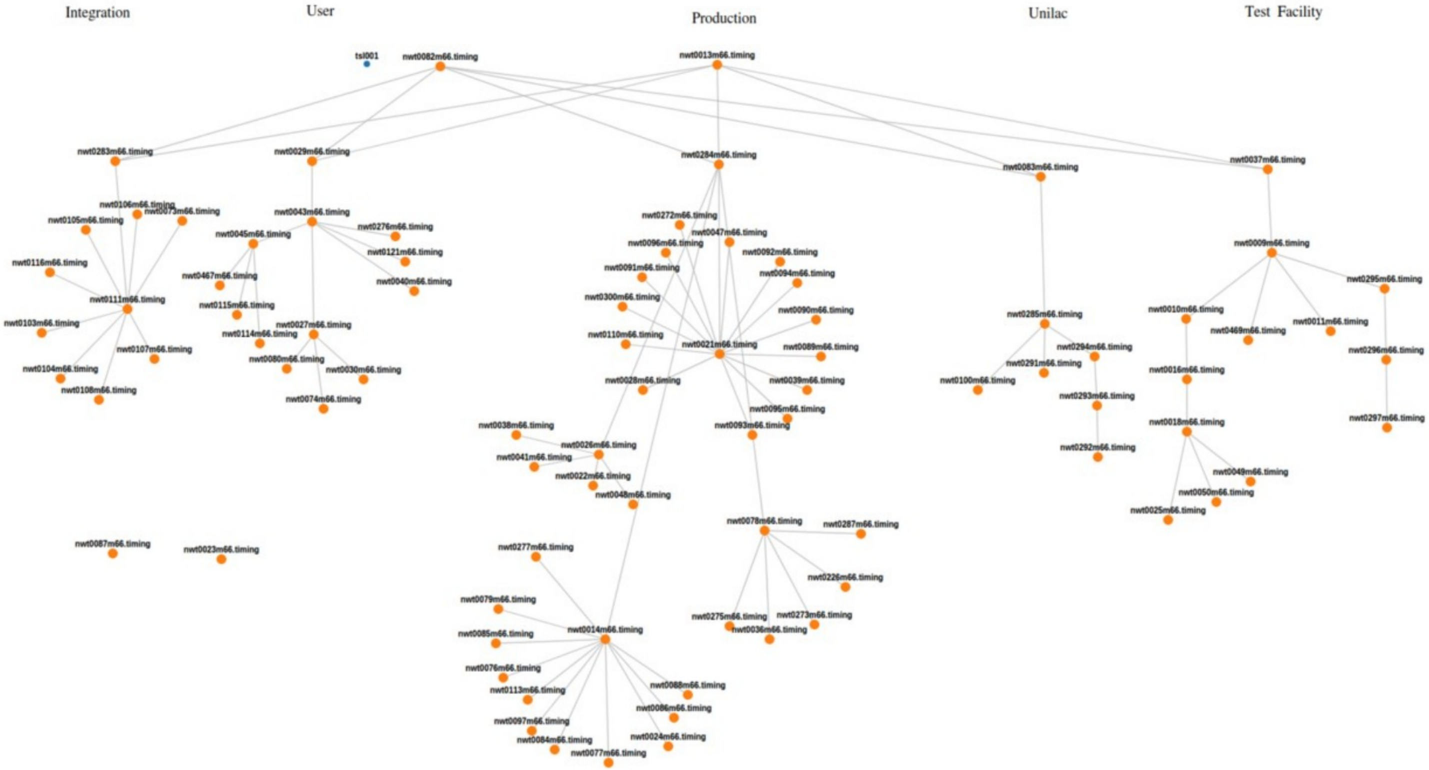

radius

- WRS can limit access to WR network
- Based on information on RADIUS server
- Uses VLANs to limit access

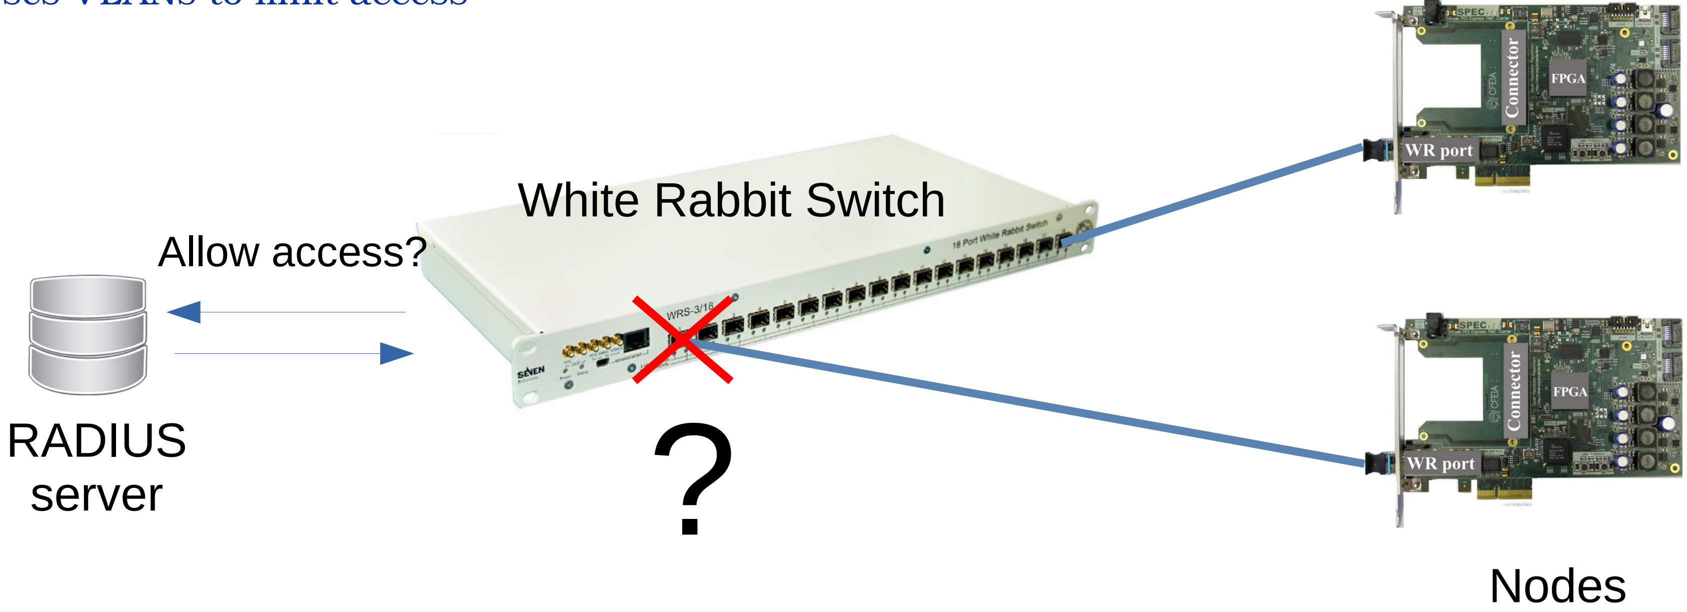

White Rabbit Switch

Allow access?

RADIUS server

?

Nodes

SNMP

- WR-SWITCH-MIB (in wr-switch-sw repo)
  - expert OIDs like:
    - Port status and configuration
    - Timing status and configuration
    - SFP monitoring (e.g., vendor, temperature)
  - status OIDs
    - Collective status of a subsystem
    - Errors/Warnings propagate upstream to wrsMainSystemStatus
    - Reason of errors is reported to syslog
  - More in “White Rabbit Switch: Failures and Diagnostics”
  - Most items cached with 5s timeout

```
$ snmpwalk -v 2c -c public -M +:. -m ALL .1.3.6.1.4.1.96.100
```

- Standard MIBs (subsets):
  - MIB-IP
  - Q-BRIDGE-MIB (VLANs)
  - BRIDGE-MIB (MAC routing tables)
  - Ongoing standardization of MIB for PTP
    - Implementation in the future
- Note: SNMP cannot be used to change the configuration of WRS

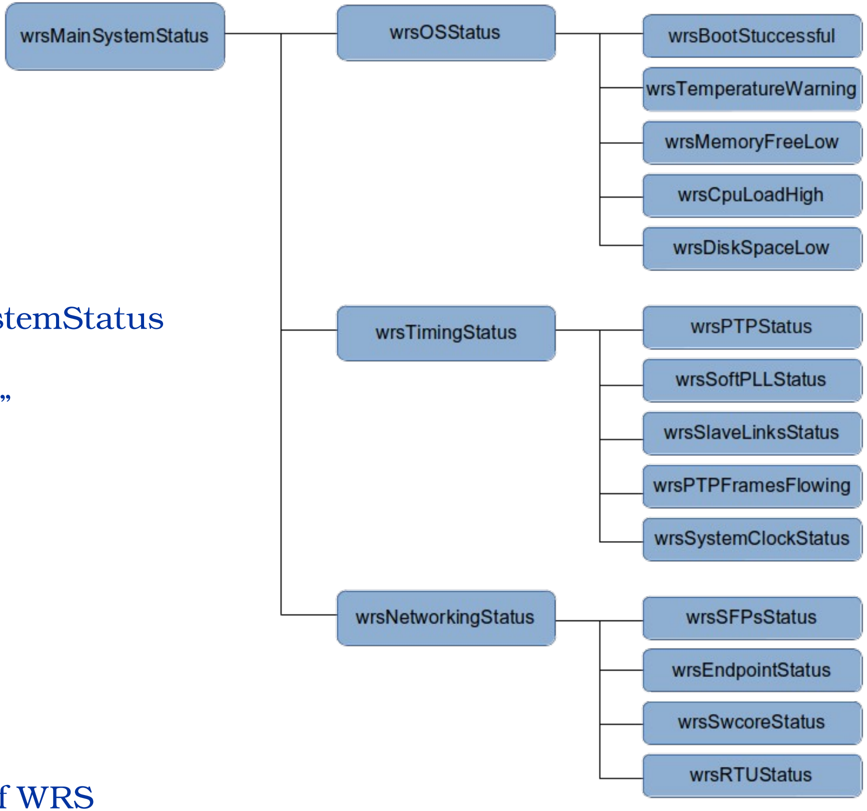

- wrsMainSystemStatus
  - wrsOSStatus
    - wrsBootSuccessful
    - wrsTemperatureWarning
    - wrsMemoryFreeLow
    - wrsCpuLoadHigh
    - wrsDiskSpaceLow
  - wrsTimingStatus
    - wrsPTPStatus
    - wrsSoftPLLStatus
    - wrsSlaveLinksStatus
    - wrsPTPFramesFlowing
    - wrsSystemClockStatus
  - wrsNetworkingStatus
    - wrsSFPsStatus
    - wrsEndpointStatus
    - wrsSwcoreStatus
    - wrsRTUStatus

SNMP: Failures and Diagnostics

- Document published with WR Switch firmware release
- Lists various errors reported by a switch
- Analyses problems that cause the error
- Proposes actions to mitigate problems
- Similar document exists for WRPC (node)

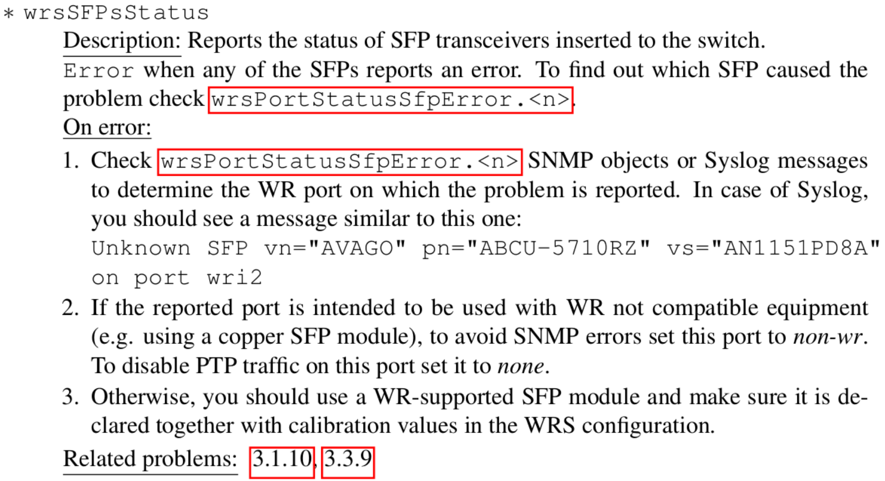

* wrsSFPsStatus

Description: Reports the status of SFP transceivers inserted to the switch. Error when any of the SFPs reports an error. To find out which SFP caused the problem check wrsPortStatusSfpError.<n>.

On error:

1. Check wrsPortStatusSfpError.<n> SNMP objects or Syslog messages to determine the WR port on which the problem is reported. In case of Syslog, you should see a message similar to this one:
   ```
   Unknown SFP vn="AVAGO" pn="ABCU-5710RZ" vs="AN1151PD8A"
   on port wri2
   ```
2. If the reported port is intended to be used with WR not compatible equipment (e.g. using a copper SFP module), to avoid SNMP errors set this port to *non-wr*. To disable PTP traffic on this port set it to *none*.
3. Otherwise, you should use a WR-supported SFP module and make sure it is declared together with calibration values in the WRS configuration.

Related problems: 3.1.10, 3.3.9

https://ohwr.org/project/wr-switch-sw/-/wikis/uploads/7b9d6bcb88a793067d9150b972c64e08/wrs_failures-v7.0.pdf

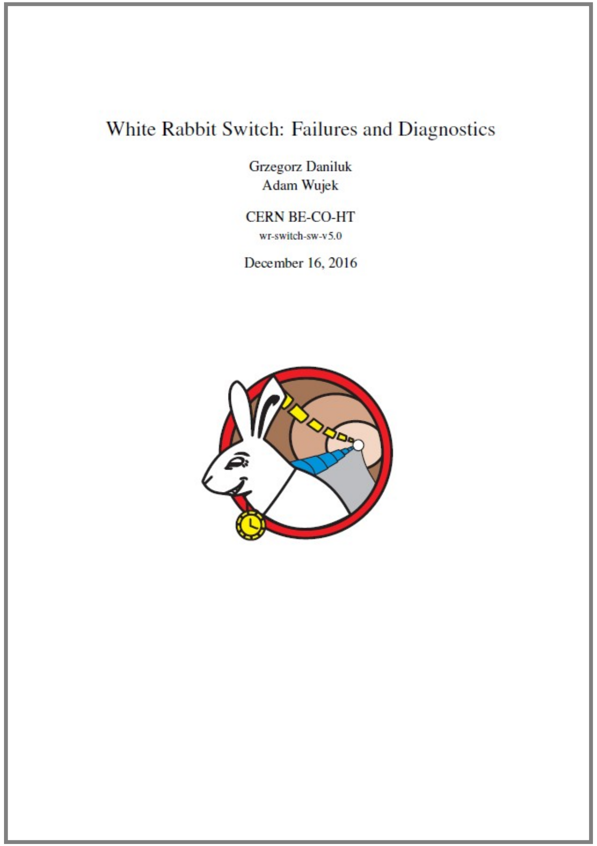

White Rabbit Switch: Failures and Diagnostics

Grzegorz Daniluk
Adam Wujek

CERN BE-CO-HT
wr-switch-sw-v5.0

December 16, 2016

SNMP: usage

Used already with different tools

- Grafana
- Net disco
- Nagios
- Zabbix

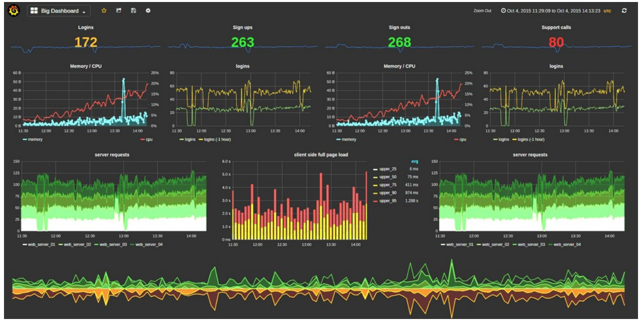

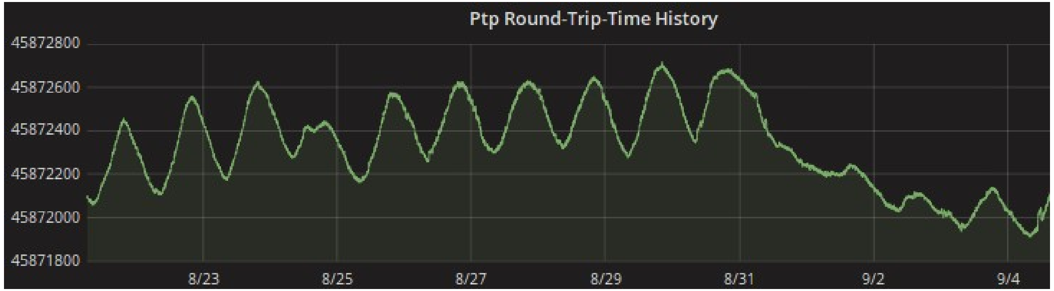

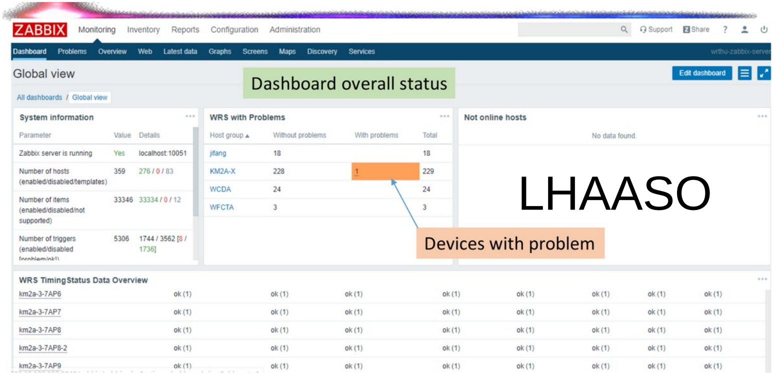

Global view

Dashboard overall status

System information

Parameter	Value	Details
Zabbix server is running	Yes	localhost:10051
Number of hosts (enabled/disabled/templates)	359	276 / 0 / 83
Number of items (enabled/disabled/not supported)	33346	33334 / 0 / 12
Number of triggers (enabled/disabled [problem/ok])	5306	1744 / 3562 [8 / 1736]

WRS with Problems

Host group ▲	Without problems	With problems	Total
jifang	18		18
KM2A-X	228	1	229
WCDA	24		24
WFCTA	3		3

Not online hosts

No data found.

LHAASO

Devices with problem

WRS TimingStatus Data Overview

km2a-3-7AP6	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)
km2a-3-7AP7	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)
km2a-3-7AP8	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)
km2a-3-7AP8-2	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)
km2a-3-7AP9	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)	ok (1)

Monitor: WRS sync status, packet flow of WRS port, RTT, temperature, etc..

Source: https://ohwr.org/project/white-rabbit/wikis/uploads/ea8d8eac37a97df1cf3cb097d081b8f4/WR_at_LHAASO

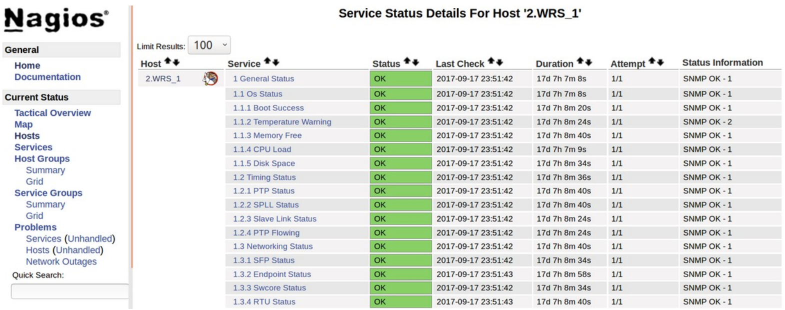

Service Status Details For Host '2.WRS_1'

Limit Results: 100

Host	Service	Status	Last Check	Duration	Attempt	Status Information
2.WRS_1	1 General Status	OK	2017-09-17 23:51:42	17d 7h 7m 8s	1/1	SNMP OK - 1
	1.1 Os Status	OK	2017-09-17 23:51:42	17d 7h 7m 8s	1/1	SNMP OK - 1
	1.1.1 Boot Success	OK	2017-09-17 23:51:42	17d 7h 8m 20s	1/1	SNMP OK - 1
	1.1.2 Temperature Warning	OK	2017-09-17 23:51:42	17d 7h 8m 24s	1/1	SNMP OK - 2
	1.1.3 Memory Free	OK	2017-09-17 23:51:42	17d 7h 8m 40s	1/1	SNMP OK - 1
	1.1.4 CPU Load	OK	2017-09-17 23:51:42	17d 7h 7m 9s	1/1	SNMP OK - 1
	1.1.5 Disk Space	OK	2017-09-17 23:51:42	17d 7h 8m 34s	1/1	SNMP OK - 1
	1.2 Timing Status	OK	2017-09-17 23:51:42	17d 7h 8m 36s	1/1	SNMP OK - 1
	1.2.1 PTP Status	OK	2017-09-17 23:51:42	17d 7h 8m 40s	1/1	SNMP OK - 1
	1.2.2 SPLL Status	OK	2017-09-17 23:51:42	17d 7h 8m 40s	1/1	SNMP OK - 1
	1.2.3 Slave Link Status	OK	2017-09-17 23:51:42	17d 7h 8m 24s	1/1	SNMP OK - 1
	1.2.4 PTP Flowing	OK	2017-09-17 23:51:42	17d 7h 8m 24s	1/1	SNMP OK - 1
	1.3 Networking Status	OK	2017-09-17 23:51:42	17d 7h 8m 40s	1/1	SNMP OK - 1
	1.3.1 SFP Status	OK	2017-09-17 23:51:42	17d 7h 8m 34s	1/1	SNMP OK - 1
	1.3.2 Endpoint Status	OK	2017-09-17 23:51:43	17d 7h 8m 58s	1/1	SNMP OK - 1
	1.3.3 Swcore Status	OK	2017-09-17 23:51:42	17d 7h 8m 34s	1/1	SNMP OK - 1
	1.3.4 RTU Status	OK	2017-09-17 23:51:43	17d 7h 8m 40s	1/1	SNMP OK - 1

Monitoring of processes

- Using monit
- Restart of a process if crashed
- Check every 10 seconds
- If restarted more than 5 times within 100s, restarts WRS
- Processes' restart counters available
  - Using wrs_dump_shmem for wrsw_hald, ppsi, wrsw_rtud
  - In files /tmp/start_cnt_* (e.g., start_cnt_lldpd) for lldpd, snmpd, sshd, syslogd, wrs_watchdog
- Cause of the last triggered WRS restart is in /tmp/monit_restart_reason
- Can be disabled in dot-config

Front panel LEDs

- Each LED can be off, green or orange color, or the combination of both giving yellow
- Two LEDs close to management port
  - Left (6): Power indicator
  - Right (5):
    - Off: Booting
    - Yellow: Before FPGA is programmed
    - Green: HAL starts successfully (desired state)
    - Orange: HAL crashed
- Two LEDs per port:
  - Left:
    - Blinking orange/off: LPDC calibration ongoing (during boot or fiber plugged)
    - Solid green: WR slave
    - Solid yellow : WR master
    - Solid orange: other mode (e.g., non-WR PTP)
  - Right:
    - Green: synced to master
    - Blinks orange/off: RX or TX of a frame
    - Blinks yellow/green: synced and RX or TX of a frame

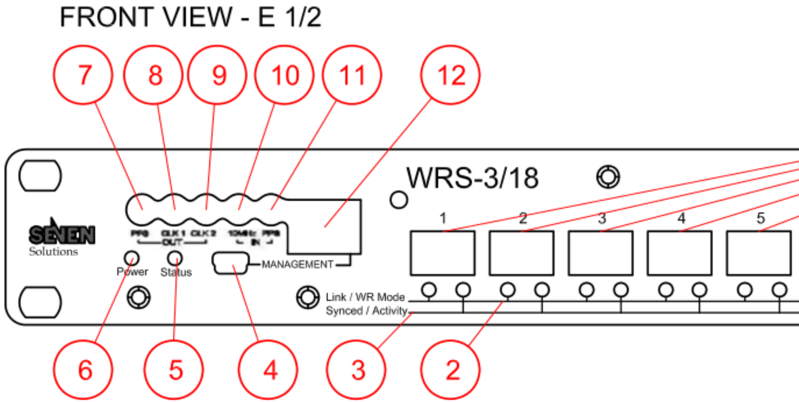

Other

- Syslog
  - If something is wrong “tool” of a second choice (after wr_mon)
  - Local and remote
    - Default location in /tmp/syslog
- Limit access to WRS
  - Set root password
    - store in dot-config as plain or md5sum
  - Disable root login
  - LDAP/kerberos
    - Use e-group for root access with sudo
- Preserve SSH host keys between firmware upgrades (new in v8.0)
- Preserve SSH authorized keys between firmware upgrades (new in v8.0)

Advanced CLI tools

wrs_dump_shmem

- Dump run-time internal data of
  - wrsw_hal (HAL)
  - ppsi
  - wrsw_rtud (RTU)
  - soft PLL stats
- wrsw_hal, ppsi, wrsw_rtu store internal data in shared memory in /dev/shm/wrs-shmem-*
- In format <name>: <value>
  - In most cases with interpretation of enumerations
- Useful when a requested information is not displayed by any other tool

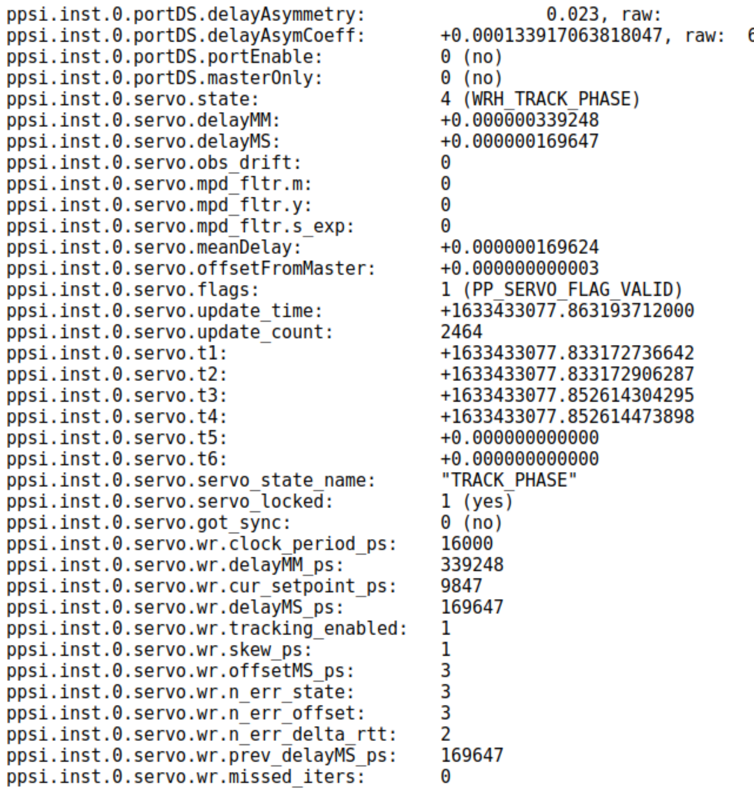

```
ppsi.inst.0.portDS.delayAsymmetry:                    0.023, raw:
ppsi.inst.0.portDS.delayAsymCoeff:        +0.00013391706381847, raw:  6
ppsi.inst.0.portDS.portEnable:            0 (no)
ppsi.inst.0.portDS.masterOnly:            0 (no)
ppsi.inst.0.servo.state:                  4 (WRH_TRACK_PHASE)
ppsi.inst.0.servo.delayMM:                +0.000000339248
ppsi.inst.0.servo.delayMS:                +0.000000169647
ppsi.inst.0.servo.obs_drift:              0
ppsi.inst.0.servo.mpd_fltr.m:             0
ppsi.inst.0.servo.mpd_fltr.y:             0
ppsi.inst.0.servo.mpd_fltr.s_exp:         0
ppsi.inst.0.servo.meanDelay:              +0.000000169624
ppsi.inst.0.servo.offsetFromMaster:       +0.000000000003
ppsi.inst.0.servo.flags:                  1 (PP_SERVO_FLAG_VALID)
ppsi.inst.0.servo.update_time:            +1633433077.863193712000
ppsi.inst.0.servo.update_count:           2464
ppsi.inst.0.servo.t1:                     +1633433077.833172736642
ppsi.inst.0.servo.t2:                     +1633433077.833172906287
ppsi.inst.0.servo.t3:                     +1633433077.852614304295
ppsi.inst.0.servo.t4:                     +1633433077.852614473898
ppsi.inst.0.servo.t5:                     +0.000000000000
ppsi.inst.0.servo.t6:                     +0.000000000000
ppsi.inst.0.servo.servo_state_name:       "TRACK_PHASE"
ppsi.inst.0.servo.servo_locked:           1 (yes)
ppsi.inst.0.servo.got_sync:               0 (no)
ppsi.inst.0.servo.wr.clock_period_ps:     16000
ppsi.inst.0.servo.wr.delayMM_ps:          339248
ppsi.inst.0.servo.wr.cur_setpoint_ps:     9847
ppsi.inst.0.servo.wr.delayMS_ps:          169647
ppsi.inst.0.servo.wr.tracking_enabled:    1
ppsi.inst.0.servo.wr.skew_ps:             1
ppsi.inst.0.servo.wr.offsetMS_ps:         3
ppsi.inst.0.servo.wr.n_err_state:         3
ppsi.inst.0.servo.wr.n_err_offset:        3
ppsi.inst.0.servo.wr.n_err_delta_rtt:     2
ppsi.inst.0.servo.wr.prev_delayMS_ps:     169647
ppsi.inst.0.servo.wr.missed_iters:        0
```

wrs_pstats

- Print frame counters for wr ports
- Can print different sets of counters (see help)

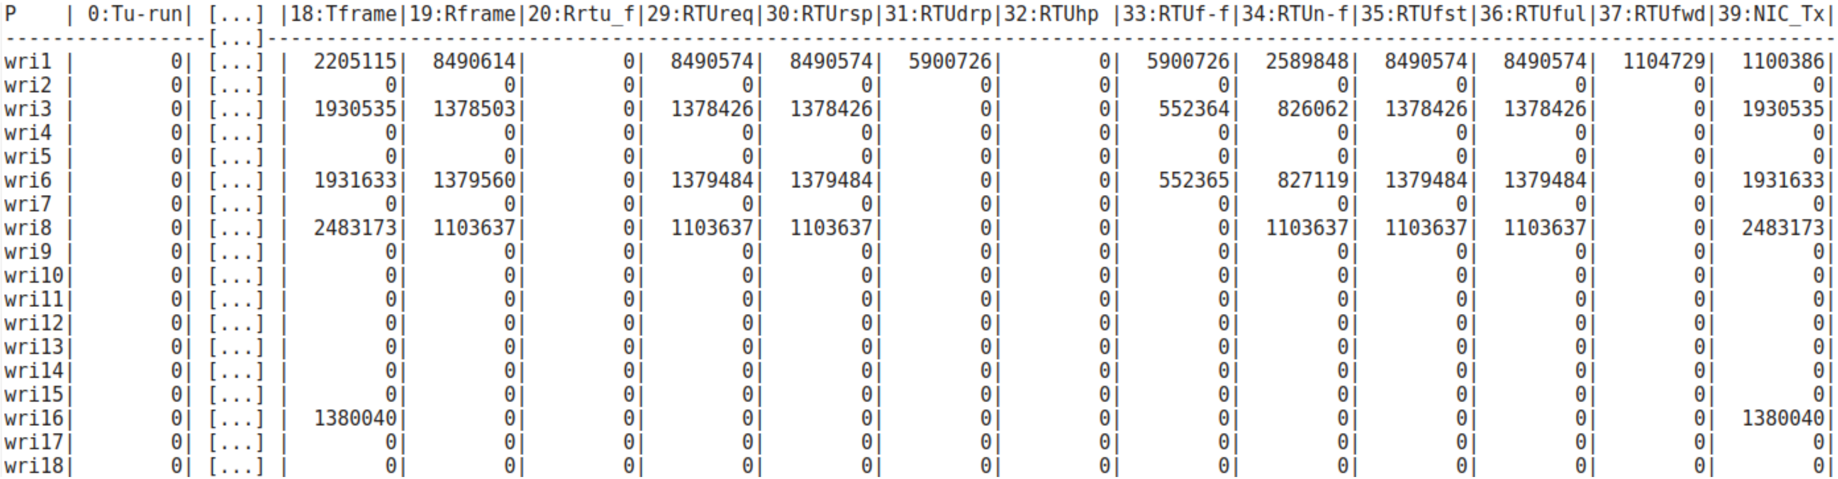

```
P    | 0:Tu-run| [...] |18:Tframe|19:Rframe|20:Rrtu_f|29:RTUreq|30:RTUrsp|31:RTUdrp|32:RTUhp |33:RTUf-f|34:RTUn-f|35:RTUfst|36:RTUful|37:RTUfwd|39:NIC_Tx|
-----------------[...]----------------------------------------------------------------------------------------------------------------------------------------------
wri1 |        0| [...] |  2205115|  8490614|        0|  8490574|  8490574|  5900726|        0|  5900726|  2589848|  8490574|  8490574|  1104729|  1100386|
wri2 |        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri3 |        0| [...] |  1930535|  1378503|        0|  1378426|  1378426|        0|        0|   552364|   826062|  1378426|  1378426|        0|  1930535|
wri4 |        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri5 |        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri6 |        0| [...] |  1931633|  1379560|        0|  1379484|  1379484|        0|        0|   552365|   827119|  1379484|  1379484|        0|  1931633|
wri7 |        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri8 |        0| [...] |  2483173|  1103637|        0|  1103637|  1103637|        0|        0|        0|  1103637|  1103637|  1103637|        0|  2483173|
wri9 |        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri10|        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri11|        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri12|        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri13|        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri14|        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri15|        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri16|        0| [...] |  1380040|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|  1380040|
wri17|        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
wri18|        0| [...] |        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|        0|
```

wrs_dump.sh

- Dump the current state of WRS
- Run over ssh from the host
- Can be used with old releases
- Script gets the following:
  - output of wrs_version
  - output of w command (logged users and uptime)
  - process list
  - output of wrs_dump_shmem
  - output of wrs_pstats
  - output of wr_mon
  - output of df
  - output of free
  - output of /proc/meminfo
  - output of ifconfig
  - tcpdump of up ports (on specified ports depending on the parameter)
  - output of PPSI's verbose messages (if selected by the parameter)
  - output of dmesg
  - output of wrs_vlans --list
  - output of wrs_vlans --plist
  - output of rtu_stat
  - output of wrs_sfp_dump -L -d -x
  - dot-config
  - shmem files
  - content of /tmp
- Gets the output of some commands twice

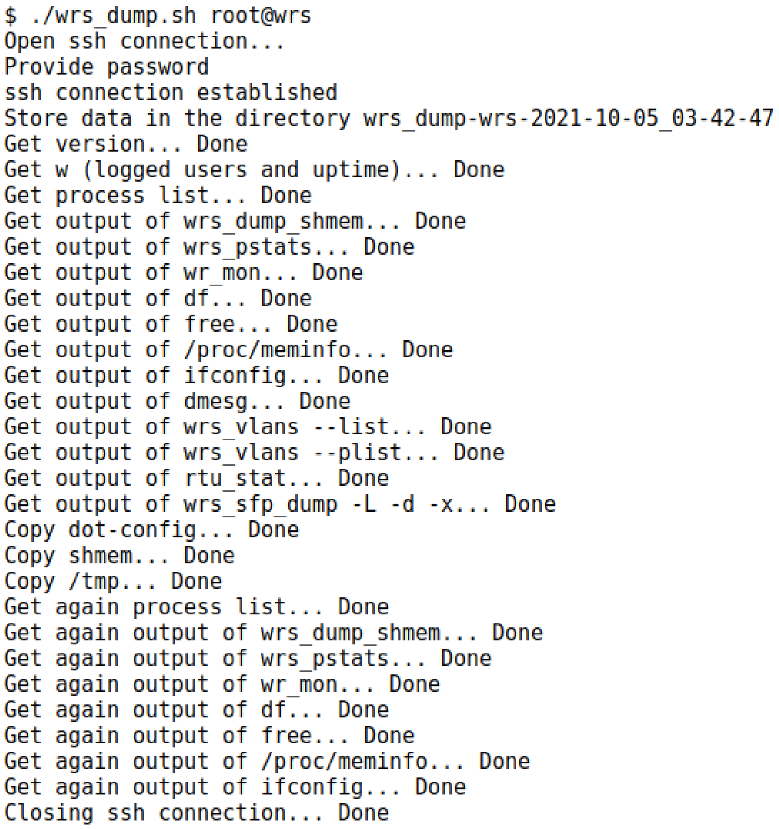

```
$ ./wrs_dump.sh root@wrs
Open ssh connection...
Provide password
ssh connection established
Store data in the directory wrs_dump-wrs-2021-10-05_03-42-47
Get version... Done
Get w (logged users and uptime)... Done
Get process list... Done
Get output of wrs_dump_shmem... Done
Get output of wrs_pstats... Done
Get output of wr_mon... Done
Get output of df... Done
Get output of free... Done
Get output of /proc/meminfo... Done
Get output of ifconfig... Done
Get output of dmesg... Done
Get output of wrs_vlans --list... Done
Get output of wrs_vlans --plist... Done
Get output of rtu_stat... Done
Get output of wrs_sfp_dump -L -d -x... Done
Copy dot-config... Done
Copy shmem... Done
Copy /tmp... Done
Get again process list... Done
Get again output of wrs_dump_shmem... Done
Get again output of wrs_pstats... Done
Get again output of wr_mon... Done
Get again output of df... Done
Get again output of free... Done
Get again output of /proc/meminfo... Done
Get again output of ifconfig... Done
Closing ssh connection... Done
```

ptpdump

- Frame level sniffer
- Dumps PTP frames on a given interface
- Useful in verification of the link status and configuration (PTP)
- In some scenarios can be used interchangeably with tcpdump

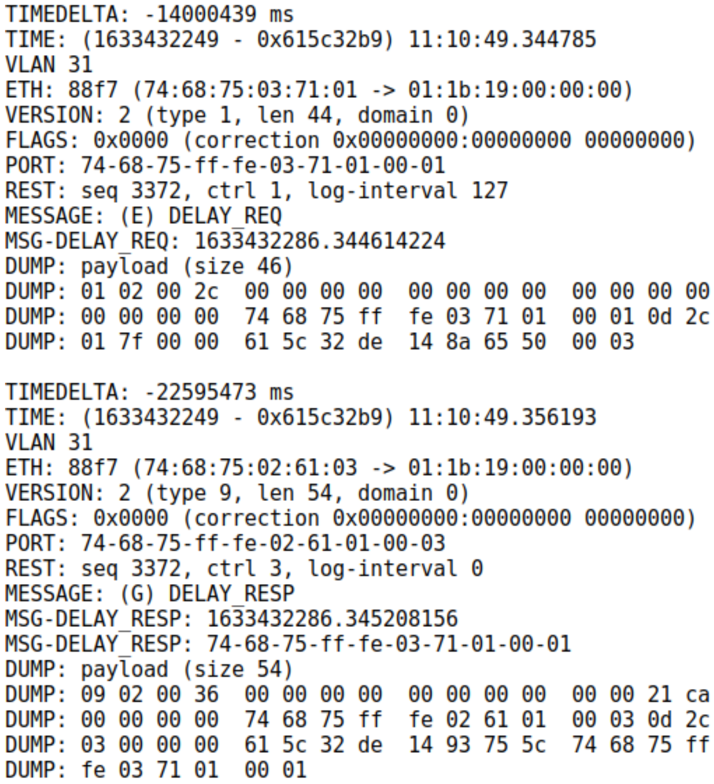

```
TIMEDELTA: -14000439 ms
TIME: (1633432249 - 0x615c32b9) 11:10:49.344785
VLAN 31
ETH: 88f7 (74:68:75:03:71:01 -> 01:1b:19:00:00:00)
VERSION: 2 (type 1, len 44, domain 0)
FLAGS: 0x0000 (correction 0x00000000:00000000 00000000)
PORT: 74-68-75-ff-fe-03-71-01-00-01
REST: seq 3372, ctrl 1, log-interval 127
MESSAGE: (E) DELAY_REQ
MSG-DELAY_REQ: 1633432286.344614224
DUMP: payload (size 46)
DUMP: 01 02 00 2c  00 00 00 00  00 00 00 00  00 00 00 00
DUMP: 00 00 00 00  74 68 75 ff  fe 03 71 01  00 01 0d 2c
DUMP: 01 7f 00 00  61 5c 32 de  14 8a 65 50  00 03

TIMEDELTA: -22595473 ms
TIME: (1633432249 - 0x615c32b9) 11:10:49.356193
VLAN 31
ETH: 88f7 (74:68:75:02:61:03 -> 01:1b:19:00:00:00)
VERSION: 2 (type 9, len 54, domain 0)
FLAGS: 0x0000 (correction 0x00000000:00000000 00000000)
PORT: 74-68-75-ff-fe-02-61-01-00-03
REST: seq 3372, ctrl 3, log-interval 0
MESSAGE: (G) DELAY_RESP
MSG-DELAY_RESP: 1633432286.345208156
MSG-DELAY_RESP: 74-68-75-ff-fe-03-71-01-00-01
DUMP: payload (size 54)
DUMP: 09 02 00 36  00 00 00 00  00 00 00 00  00 00 21 ca
DUMP: 00 00 00 00  74 68 75 ff  fe 02 61 01  00 03 0d 2c
DUMP: 03 00 00 00  61 5c 32 de  14 93 75 5c  74 68 75 ff
DUMP: fe 03 71 01  00 01
```

tcpdump & wireshark

- tcpdump - network sniffer
- wireshark - network sniffer with GUI
  - can dissect WR specific fields in PTP frames/packets since version 2.9
    - Version available in ubuntu 20.04 LTS
- tcpdump can be combined with wireshark to dissect frames in live over ssh:

  ```
  ssh root@wrs "tcpdump -i wri1 --immediate-mode --packet-buffered -w -" | wireshark -i -
  ```

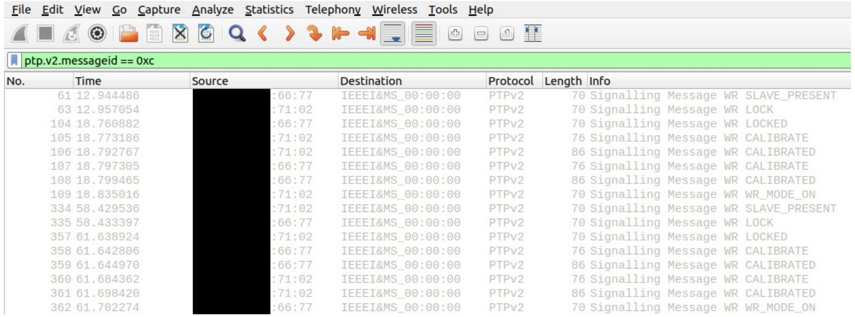

ptp.v2.messageid == 0xc

No.	Time	Source	Destination	Protocol	Length	Info
61	12.944486	[redacted]:66:77	IEEEI&MS_00:00:00	PTPv2	70	Signalling Message WR SLAVE_PRESENT
63	12.957054	[redacted]:71:02	IEEEI&MS_00:00:00	PTPv2	70	Signalling Message WR LOCK
104	18.760882	[redacted]:66:77	IEEEI&MS_00:00:00	PTPv2	70	Signalling Message WR LOCKED
105	18.773186	[redacted]:71:02	IEEEI&MS_00:00:00	PTPv2	76	Signalling Message WR CALIBRATE
106	18.792767	[redacted]:71:02	IEEEI&MS_00:00:00	PTPv2	86	Signalling Message WR CALIBRATED
107	18.797305	[redacted]:66:77	IEEEI&MS_00:00:00	PTPv2	76	Signalling Message WR CALIBRATE
108	18.799465	[redacted]:66:77	IEEEI&MS_00:00:00	PTPv2	86	Signalling Message WR CALIBRATED
109	18.835016	[redacted]:71:02	IEEEI&MS_00:00:00	PTPv2	70	Signalling Message WR WR_MODE_ON
334	58.429536	[redacted]:71:02	IEEEI&MS_00:00:00	PTPv2	70	Signalling Message WR SLAVE_PRESENT
335	58.433397	[redacted]:66:77	IEEEI&MS_00:00:00	PTPv2	70	Signalling Message WR LOCK
357	61.638924	[redacted]:71:02	IEEEI&MS_00:00:00	PTPv2	70	Signalling Message WR LOCKED
358	61.642806	[redacted]:66:77	IEEEI&MS_00:00:00	PTPv2	76	Signalling Message WR CALIBRATE
359	61.644970	[redacted]:66:77	IEEEI&MS_00:00:00	PTPv2	86	Signalling Message WR CALIBRATED
360	61.684362	[redacted]:71:02	IEEEI&MS_00:00:00	PTPv2	76	Signalling Message WR CALIBRATE
361	61.698420	[redacted]:71:02	IEEEI&MS_00:00:00	PTPv2	86	Signalling Message WR CALIBRATED
362	61.702274	[redacted]:66:77	IEEEI&MS_00:00:00	PTPv2	70	Signalling Message WR WR_MODE_ON

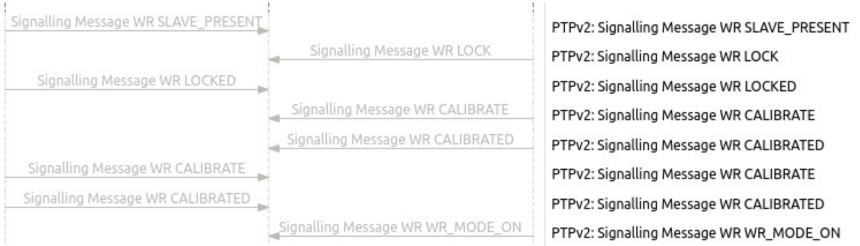
Signalling Message WR SLAVE_PRESENT
Signalling Message WR LOCK
Signalling Message WR LOCKED
Signalling Message WR CALIBRATE
Signalling Message WR CALIBRATED
Signalling Message WR CALIBRATE
Signalling Message WR CALIBRATED
Signalling Message WR WR_MODE_ON

PTPv2: Signalling Message WR SLAVE_PRESENT
PTPv2: Signalling Message WR LOCK
PTPv2: Signalling Message WR LOCKED
PTPv2: Signalling Message WR CALIBRATE
PTPv2: Signalling Message WR CALIBRATED
PTPv2: Signalling Message WR CALIBRATE
PTPv2: Signalling Message WR CALIBRATED
PTPv2: Signalling Message WR WR_MODE_ON

wr_vlans

Useful in configuration and verification of VLANs

- Print ports' VLAN configuration
  - wrs_vlans --plist
  - Mode of port (Trunk, Access etc.)
  - If tagging enabled, extra parameters: priority, PVID
- Print VLAN configuration
  - wrs_vlans --list
  - Used VID/FID
  - Association of ports to VLANs
- Configuration in runtime

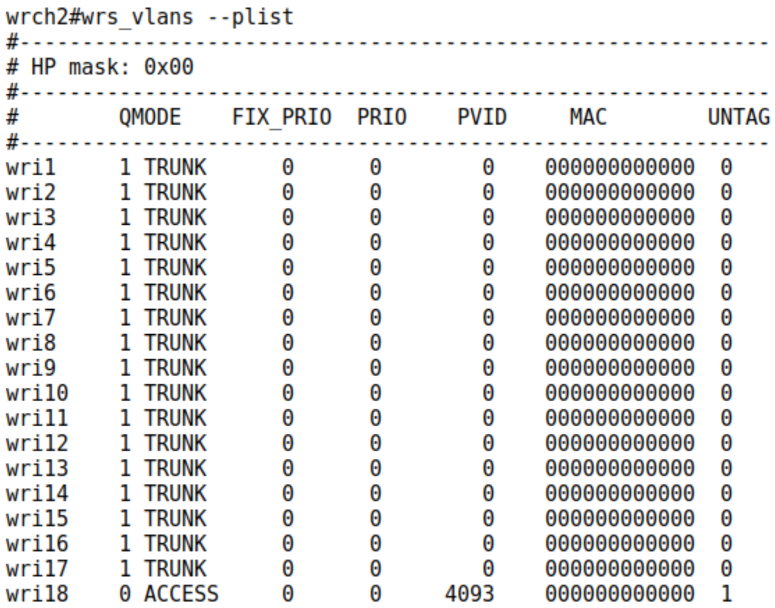

```
wrch2#wrs_vlans --plist
#------------------------------------------------------------
# HP mask: 0x00
#------------------------------------------------------------
#        QMODE     FIX_PRIO  PRIO     PVID      MAC          UNTAG
#------------------------------------------------------------
wri1     1 TRUNK       0       0         0     000000000000   0
wri2     1 TRUNK       0       0         0     000000000000   0
wri3     1 TRUNK       0       0         0     000000000000   0
wri4     1 TRUNK       0       0         0     000000000000   0
wri5     1 TRUNK       0       0         0     000000000000   0
wri6     1 TRUNK       0       0         0     000000000000   0
wri7     1 TRUNK       0       0         0     000000000000   0
wri8     1 TRUNK       0       0         0     000000000000   0
wri9     1 TRUNK       0       0         0     000000000000   0
wri10    1 TRUNK       0       0         0     000000000000   0
wri11    1 TRUNK       0       0         0     000000000000   0
wri12    1 TRUNK       0       0         0     000000000000   0
wri13    1 TRUNK       0       0         0     000000000000   0
wri14    1 TRUNK       0       0         0     000000000000   0
wri15    1 TRUNK       0       0         0     000000000000   0
wri16    1 TRUNK       0       0         0     000000000000   0
wri17    1 TRUNK       0       0         0     000000000000   0
wri18    0 ACCESS      0       0      4093     000000000000   1
```

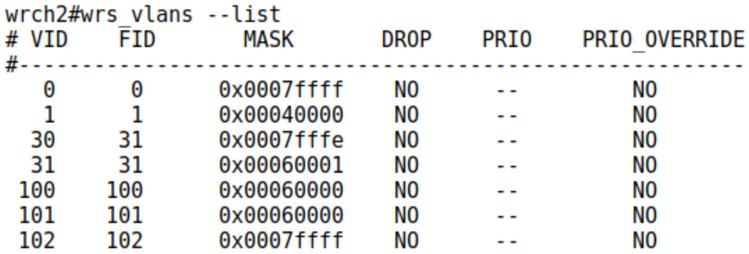

```
wrch2#wrs_vlans --list
# VID     FID        MASK        DROP    PRIO    PRIO_OVERRIDE
#-----------------------------------------------------------
    0       0     0x0007ffff     NO       --          NO
    1       1     0x00040000     NO       --          NO
   30      31     0x0007fffe     NO       --          NO
   31      31     0x00060001     NO       --          NO
  100     100     0x00060000     NO       --          NO
  101     101     0x00060000     NO       --          NO
  102     102     0x0007ffff     NO       --          NO
```

rtu_stat

- Print details about:
  - Switching tables
  - VLANs
  - Port mirroring details
- Can set the above parameters

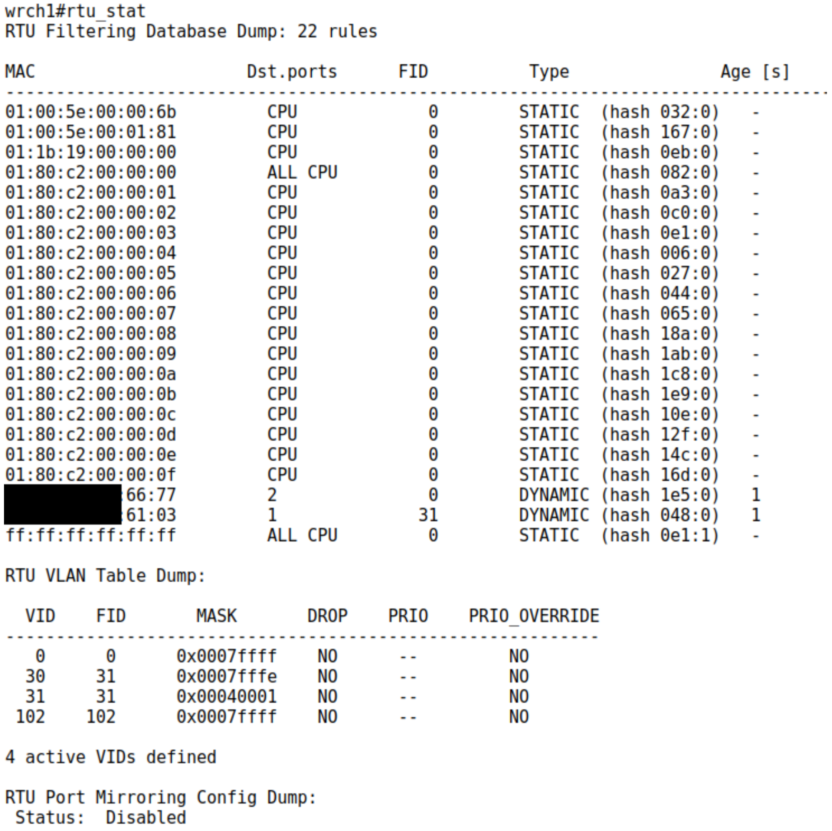

```
wrch1#rtu_stat
RTU Filtering Database Dump: 22 rules

MAC                       Dst.ports      FID         Type               Age [s]
---------------------------------------------------------------------------------
01:00:5e:00:00:6b           CPU            0          STATIC  (hash 032:0)   -
01:00:5e:00:01:81           CPU            0          STATIC  (hash 167:0)   -
01:1b:19:00:00:00           CPU            0          STATIC  (hash 0eb:0)   -
01:80:c2:00:00:00           ALL CPU        0          STATIC  (hash 082:0)   -
01:80:c2:00:00:01           CPU            0          STATIC  (hash 0a3:0)   -
01:80:c2:00:00:02           CPU            0          STATIC  (hash 0c0:0)   -
01:80:c2:00:00:03           CPU            0          STATIC  (hash 0e1:0)   -
01:80:c2:00:00:04           CPU            0          STATIC  (hash 006:0)   -
01:80:c2:00:00:05           CPU            0          STATIC  (hash 027:0)   -
01:80:c2:00:00:06           CPU            0          STATIC  (hash 044:0)   -
01:80:c2:00:00:07           CPU            0          STATIC  (hash 065:0)   -
01:80:c2:00:00:08           CPU            0          STATIC  (hash 18a:0)   -
01:80:c2:00:00:09           CPU            0          STATIC  (hash 1ab:0)   -
01:80:c2:00:00:0a           CPU            0          STATIC  (hash 1c8:0)   -
01:80:c2:00:00:0b           CPU            0          STATIC  (hash 1e9:0)   -
01:80:c2:00:00:0c           CPU            0          STATIC  (hash 10e:0)   -
01:80:c2:00:00:0d           CPU            0          STATIC  (hash 12f:0)   -
01:80:c2:00:00:0e           CPU            0          STATIC  (hash 14c:0)   -
01:80:c2:00:00:0f           CPU            0          STATIC  (hash 16d:0)   -
[redacted]  66:77           2              0          DYNAMIC (hash 1e5:0)   1
            61:03           1             31          DYNAMIC (hash 048:0)   1
ff:ff:ff:ff:ff:ff           ALL CPU        0          STATIC  (hash 0e1:1)   -

RTU VLAN Table Dump:

  VID     FID        MASK        DROP     PRIO    PRIO_OVERRIDE
-----------------------------------------------------------------
    0       0      0x0007ffff     NO       --          NO
   30      31      0x0007fffe     NO       --          NO
   31      31      0x00040001     NO       --          NO
  102     102      0x0007ffff     NO       --          NO

4 active VIDs defined

RTU Port Mirroring Config Dump:
 Status:  Disabled
```

Files in /tmp

Various files in /tmp:

- custom_boot_script_status
- dot-config_source
- dot-config_status
- hwinfo_read_status
- leapseconds_download_source
- leapseconds_check_status
- vlans_set_status
- wrs_auxclk_set_status
- wrs_throttling_set_status
- load_fpga_status
- load_lm32_status
- start_cnt_* (e.g., start_cnt_lldpd)
- syslog

Configuration

PTP devices roles

- Grandmaster (provider of time and frequency)
  - Synchronized to a primary reference time source
  - Needs source of frequency (via 1-PPS and 10MHz)
    - GPS
    - Atomic clock Rb, Cs
  - Needs source of time
    - NTP
    - NMEA (from v8.0)
    - IRIG-B (from v8.0)
- Boundary clock (distribution device)
  - Intended to be connected to another device with better clock
- Ordinary clock (end node)
  - Ordinary clock is implemented as Boundary clock with no slaves connected
- In WR additionally:
  - Arbitrary Grand-Master
    - Like GM, but without source of time
  - Free-Running master
    - Like GM, no external reference nor source of time, using local oscillator
  - Custom
    - For testing, without constraints on parameters

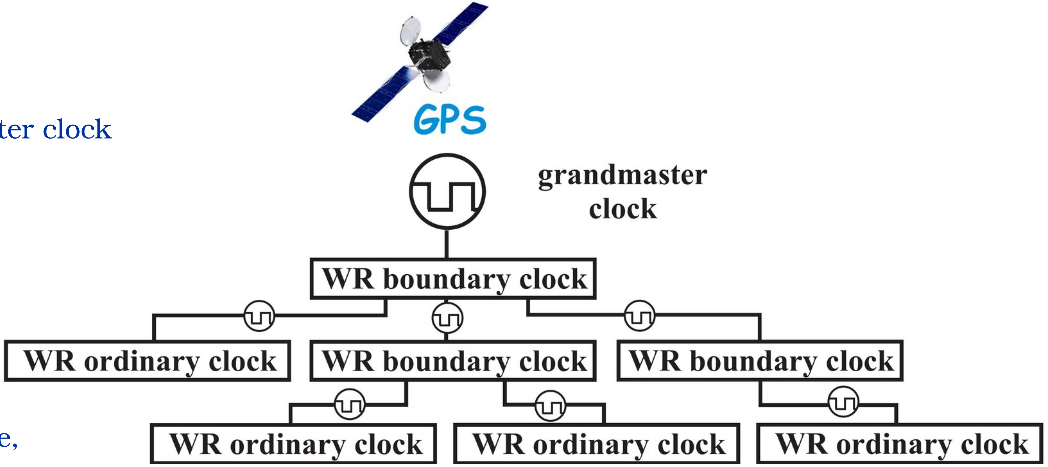

Assignment of port roles

PTP standard allows two types of port roles assignment (e.g., master/slave)

- Static (external port configuration)
- Dynamic (based on Best Master Clock Algorithm, BMCA)
  - A port becomes slave, that has a peer with best metrics
  - If local clock has even better metrics, local device becomes timing source for other devices (e.g., grandmaster)
  - Decision is based on predefined metrics (lower = better) like:
    - priority1
    - clockClass
    - accuracy
    - variance
    - priority2
    - clockIdentity
  - Different profiles may implement different BMCA
  - As today, the only way to provide timing redundancy
    - Takes 10's of seconds to minutes to propagate changes over the network
    - Avoid data loops with VLANs, no Spanning Tree Protocol implemented

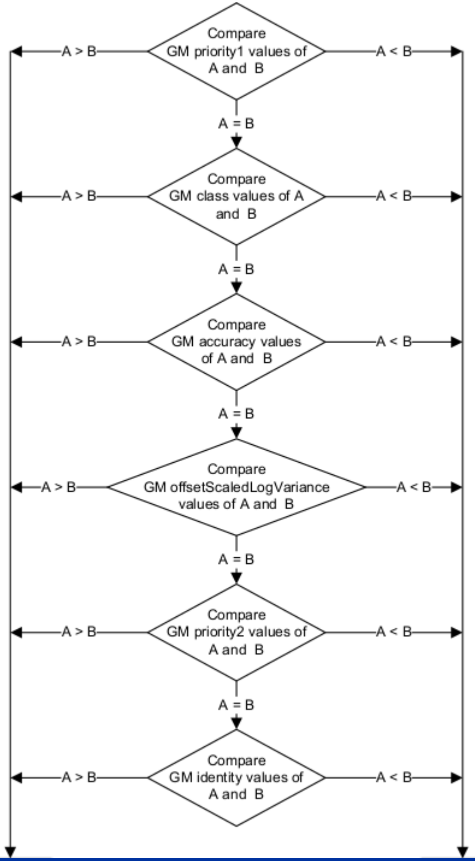
Compare GM priority1 values of A and B
A > B
A < B
A = B
Compare GM class values of A and B
A > B
A < B
A = B
Compare GM accuracy values of A and B
A > B
A < B
A = B
Compare GM offsetScaledLogVariance values of A and B
A > B
A < B
A = B
Compare GM priority2 values of A and B
A > B
A < B
A = B
Compare GM identity values of A and B
A > B
A < B

Profiles

- A PTP Profile defines a set of default and allowed values of parameters
- In WRS setting of profile is split to:
  - Global profile for parameters that are common or independent to all ports
    - Parameters of local clock (e.g., priority1, clockClass)
    - In WRS can be configured as:
      - Default PTP Profile
      - White Rabbit / High Accuracy (allows to use White Rabbit or L1Sync extension)
      - Custom (does not put restrictions on parameters values)
  - Port profile for parameters that are specific to a port
    - Extension (White Rabbit or L1Sync)
    - PTP messages rates
    - In WRS can be configured as:
      - Keep global (use the same profile as Global Profile)
      - Default PTP Profile
      - White Rabbit / High Accuracy (allows to use White Rabbit or L1Sync extension)
      - Custom (does not put restrictions on parameters values)

PTP Extensions

- A PTP extension refers to an enhancement or addition to the core PTP protocol
- WR switch's ports can be configured to use as extension:
  - None
  - White Rabbit
  - L1Sync
  - Autonegotiation between White Rabbit and L1Sync

- Autonegotiation
  - Allows to autodetect peer's extension
  - Detection is done with priority:
    - L1Sync
    - White Rabbit
    - None (standard PTP)

SFP database matching

- Why needed? To get relative calibration values of SFPs RX and TX
  - For absolute calibration, values will be read from SFPs EEPROM (see SFF-8472, rev 12.4.2)
- Entries in dot-config like:
  CONFIG_SFP00_PARAMS="vn=Axcen Photonics,pn=AXGE-3454-0531,vs=AX12390009629,
  tx=0,rx=0,wl_txrx=1490+1310"
  - vn – vendor name
  - pn – product name
  - vs – vendor serial
  - tx,rx – relative TX, RX delays (they may be negative!)
  - wl_txrx – TX and RX wavelengths in nm
- VN, PN, VS, TX wavelength are read from SFP’s EEPROM
- Matching order:
  - TX wavelength, PN, VN, VS
  - TX wavelength, PN, VN
  - TX wavelength, PN
- After match, deltas are taken from tx= and rx=

Fiber database matching

- Why needed?
  - To get delayCoefficient (alpha) of a fiber
  - Fiber has no memory that can store this value
- Entries in dot-config like:
  CONFIG_PORT01_FIBER=0
  CONFIG_FIBER00_PARAMS="alpha_1310_1490=0.0002743"
- Port has to define which fiber type is used (marked with gray)
- For matching alpha_XX_YY, values of XX and YY are taken from wl_txrx parameter of SFP DB entry. Like:
- CONFIG_SFP00_PARAMS="vn=Axcen Photonics,pn=AXGE-3454-0531,vs=AX12390009629, tx=0,rx=0,wl_txrx=1490+1310"
- If alpha_XX_YY is not matched, swapped XX and YY is tried (alpha_YY_XX)

Note alpha in the opposite direction is not -alpha, but: $\frac{1}{1+alpha}-1$ which is equivalent to: $-\frac{alpha}{1+alpha}$

Dot-config overwrite

- New in v8.0
- Local subset of dot-config that overwrites values from main dot-config at boot-up
- Intention is to store e.g., management port configuration (IP, etc.), calibration values for HW or SPLL settings
- Kept between firmware upgrades
- Intention that it is small
- Option in main dot-config, enabled by default

Interoperability with non-WR devices

Interoperability with non-WR devices

- WRS can work with other PTP devices
- Recommended settings:
  - Default profile
  - PTP over Ethernet
- Some devices requires different settings
- Successfully working with:
  - Oscilloquartz (standard profile with UDP)
  - Juniper (AES-67 profile)
  - Meinberg (profiles: standard, AES-67, ITU-T G.8275.1)
  - Microchip (profiles: standard, ITU-T G.8275.1)
  - Raspberry CM4 with ptp4l

Interoperability with non-WR devices: Oscilloquartz

- WRS:
  - set PTP version 2.0
    - enable: "Global PTP configuration (profile, timing mode, BMCA and external port config, ...)" -> "Allow the overwrite of the default PTP version per instance"
    - for a given port: "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "Force PTP version" and select "PTP version (v2.0 (IEEE1588-2008))"
  - set to use PTP over L3 for a port
    - "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "Network protocol" -> UDP/IPv4
  - set IP address for a given port
    - (from WRS firmware v8.0): "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "IP address"
    - (prior to WRS firmware v8.0) manually with e.g., a command: ifconfig wri5 192.168.255.100
  - select default PTP profile
    - "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "Network protocol" -> "PTP Profile (Default (IEEE 1588))"
- Oscilloquartz:
  - select profile “IEEE1588-2008 annex J.3”
  - set IP address on a port
  - set properly syncE (not to expect ESMC)

Interoperability with non-WR devices: Juniper

- WRS:
  - set PTP version 2.0
    - enable: "Global PTP configuration (profile, timing mode, BMCA and external port config, ...)" -> "Allow the overwrite of the default PTP version per instance"
    - for a given port: "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "Force PTP version" and select "PTP version (v2.0 (IEEE1588-2008))"
  - set to use PTP over L3 for a port
    - "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "Network protocol" -> UDP/IPv4
  - set IP address for a given port
    - (from WRS firmware v8.0): "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "IP address"
    - (prior to WRS firmware v8.0) manually with e.g., a command: ifconfig wri5 192.168.255.100
  - select default PTP profile
    - "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "Network protocol" -> "PTP Profile (Default (IEEE 1588))"
  - adjust message rates
    - enable: "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "Overwrite default minDelayRequestInterval"
    - set: "Port PTP/Timing configuration" "PORT XX" -> "Instance 1" -> "minDelayRequestInterval"
- Juniper:
  - select profile AES-67
  - set IP address on a port

Questions?

More info

WRS: User manual:

https://ohwr.org/project/wr-switch-sw/-/wikis/uploads/d1f78666704fb292982453e1429b9f10/wrs-user-manual-v7.0.pdf

WRS: Developer manual:

https://ohwr.org/project/wr-switch-sw/-/wikis/uploads/1a5a73c1528ccfe7e739e0dfc8e0ecd1/wrs-developer-manual-v7.0.pdf

WRS: Failures and Diagnostics:

https://ohwr.org/project/wr-switch-sw/-/wikis/uploads/7b9d6bcb88a793067d9150b972c64e08/wrs_failures-v7.0.pdf

WRS: Radius Vlan:

https://ohwr.org/project/wr-switch-sw/-/wikis/uploads/5f86a996d29a2fb21a389c27da7781db/wrs-radiusvlan-v6.1.pdf