Das kannste schon so machen, ist dann aber... [*]

TALES OF DAILY CHECK_MK USAGE

[*] Lost in translation

Network – All Interfaces? Not really!?

- All Interfaces deserve to be monitored
- Customers often say: „No, I‘m not interested in monitoring all interfaces“ - Yes! They just don‘t know their errors yet.
- Pro-Tip: No errors in a LAN are acceptable! And they can be monitored and addressed in some simple steps.

Network

- But my network doesn‘t have errors!

Don‘t be too sure

Network – How to monitor?

- By „Description“ or by „Alias“ - depends on vendor
- All states!

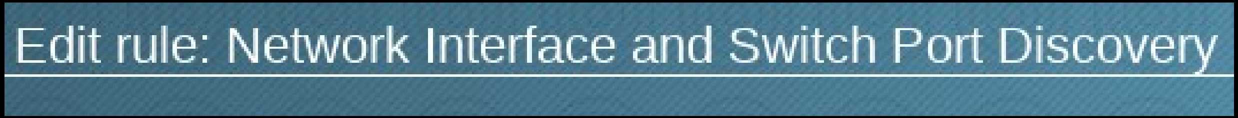

Edit rule: Network Interface and Switch Port Discovery

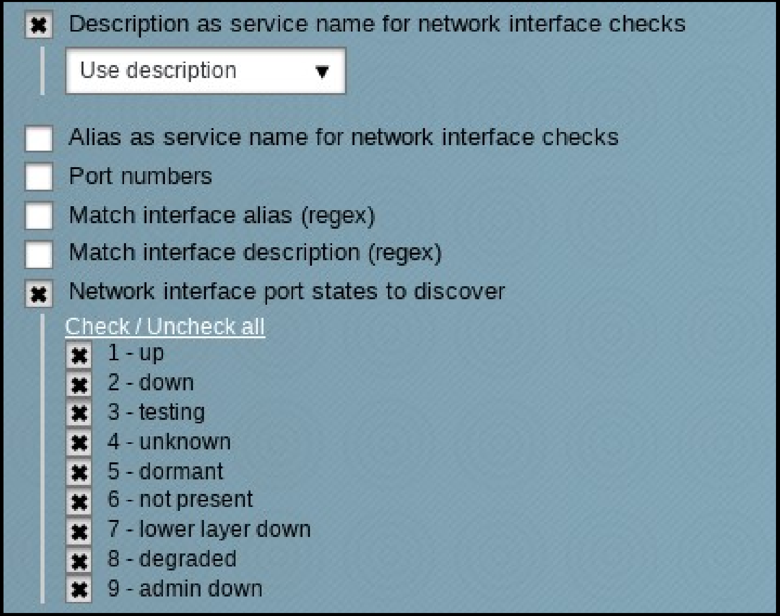

☒ Description as service name for network interface checks

Use description ▼

□ Alias as service name for network interface checks

□ Port numbers

□ Match interface alias (regex)

□ Match interface description (regex)

☒ Network interface port states to discover

Check / Uncheck all

☒ 1 - up
☒ 2 - down
☒ 3 - testing
☒ 4 - unknown
☒ 5 - dormant
☒ 6 - not present
☒ 7 - lower layer down
☒ 8 - degraded
☒ 9 - admin down

Network – How to monitor?

- All types!

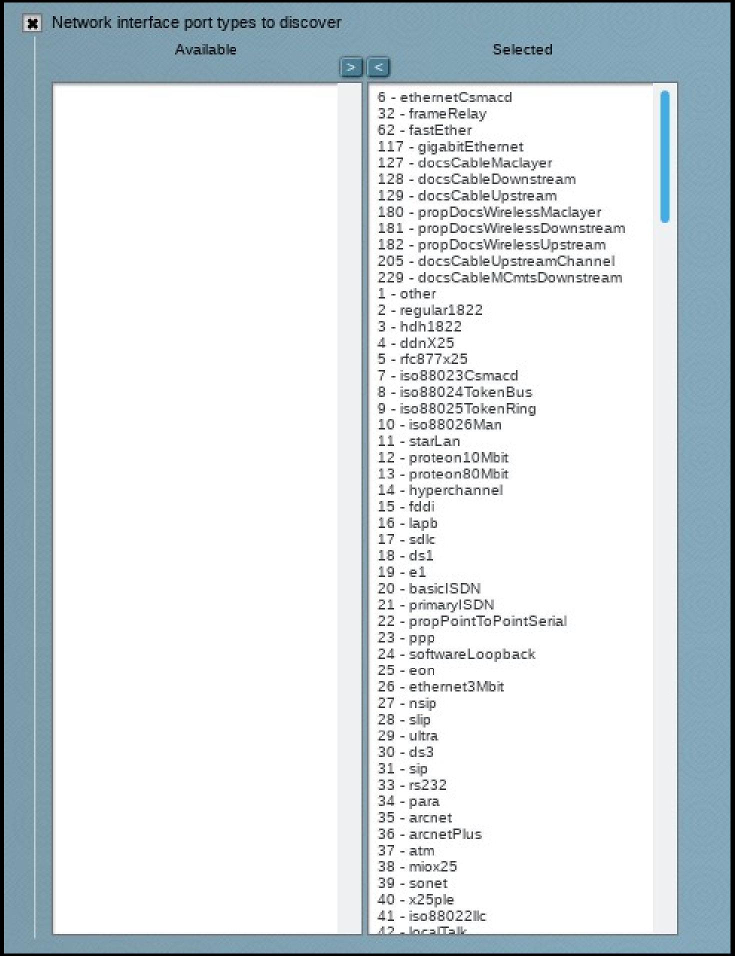

☒ Network interface port types to discover

6 - ethernetCsmacd
32 - frameRelay
62 - fastEther
117 - gigabitEthernet
127 - docsCableMaclayer
128 - docsCableDownstream
129 - docsCableUpstream
180 - propDocsWirelessMaclayer
181 - propDocsWirelessDownstream
182 - propDocsWirelessUpstream
205 - docsCableUpstreamChannel
229 - docsCableMCmtsDownstream
1 - other
2 - regular1822
3 - hdh1822
4 - ddnX25
5 - rfc877x25
7 - iso88023Csmacd
8 - iso88024TokenBus
9 - iso88025TokenRing
10 - iso88026Man
11 - starLan
12 - proteon10Mbit
13 - proteon80Mbit
14 - hyperchannel
15 - fddi
16 - lapb
17 - sdlc
18 - ds1
19 - e1
20 - basicISDN
21 - primaryISDN
22 - propPointToPointSerial
23 - ppp
24 - softwareLoopback
25 - eon
26 - ethernet3Mbit
27 - nsip
28 - slip
29 - ultra
30 - ds3
31 - sip
33 - rs232
34 - para
35 - arcnet
36 - arcnetPlus
37 - atm
38 - miox25
39 - sonet
40 - x25ple
41 - iso88022llc
42 - localTalk

Network – OK fine, but what about the access ports...?

- Access ports are allowed to change state (up/down)
- Access ports are allowed to change speed as well, e.g. 1Gbit/s while powered on, 10Mbit/s in WOL mode

Network – OK fine, but what about the access ports...?

- But access ports are not allowed to have errors!
- How to distinguish access from non-access ports in monitoring?
- Port naming schemes and appropriate rules! Remember: „by Description/by Alias“

Network – Distinguish access from non-access

- Name important interfaces in your switches/routers, e.g.:
  - Uplink core X
  - Uplink MPLS 10 Mbit
  - ap23
  - esx-vmnic0
- Don‘t name your access ports, unless you are able and willingly to handle all changes
- It doesn‘t matter what‘s your naming scheme, but let it be consistent → efficient Check_MK ruleset!

Network – Distinguish access from non-access - Rules

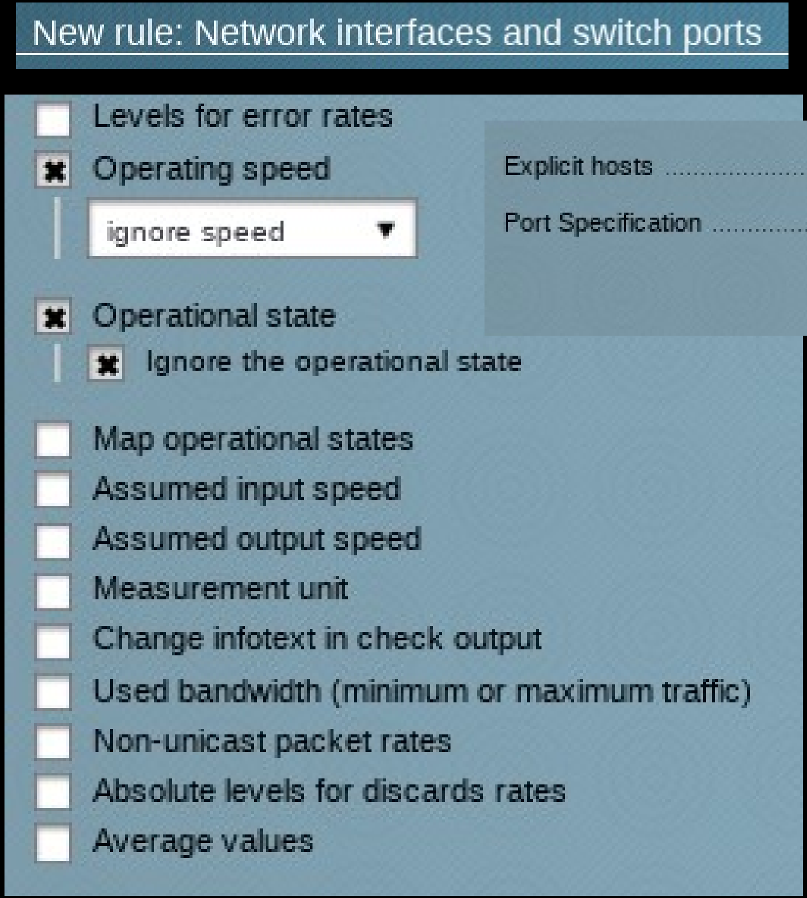

New rule: Network interfaces and switch ports

- □ Levels for error rates
- ☒ Operating speed
  - ignore speed ▼
- ☒ Operational state
  - ☒ Ignore the operational state
- □ Map operational states
- □ Assumed input speed
- □ Assumed output speed
- □ Measurement unit
- □ Change infotext in check output
- □ Used bandwidth (minimum or maximum traffic)
- □ Non-unicast packet rates
- □ Absolute levels for discards rates
- □ Average values

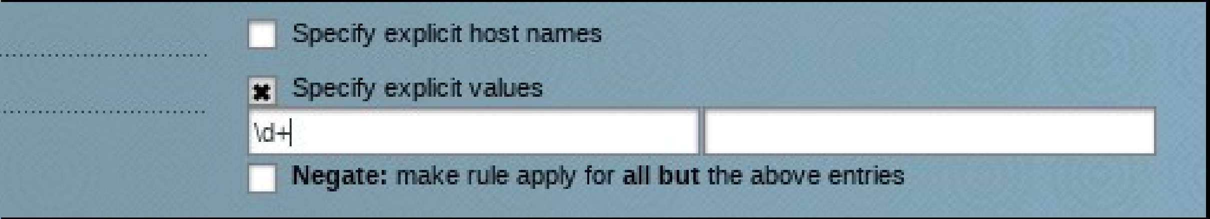

Explicit hosts .... □ Specify explicit host names

Port Specification .... ☒ Specify explicit values

\d+

□ **Negate:** make rule apply for **all but** the above entries

\d+
Regex for several digits, translates to
e.g. „Interface 000034“ aka
Port doesn‘t have a name

Network – another takeaway

- Note: The default levels of 0.01%/0.1% (WARN/CRIT) for interface errors make sense! Don‘t change them! Never*
- * OK sometimes as some WLAN-vendors pass errors of the radio interface to the counters...

Network – why?

- If you allow Check_MK to search, you‘ll find:
  - broken patch- and installation cables
  - dirty fiber optics
  - configuration error: one sided, deactivated Auto-Negotiation, very common → It‘s a protocol and not electro-magic, administrators tend not to know that :-(
  → Duplex Mismatch!!!11!

Network – why?

- If you allow Check_MK to search, you‘ll find:
  - Configuration error: Trunks LACP vs. static

  - broken firmware (Printer, yes, really)
  - overloaded embedded systems/IP stacks

Network – why?

- If you allow Check_MK to search, you‘ll find:
  <add your network phenomena here>

Network – real life example (1)

- Errors on an uplink interface in a metro network → single mode fiber optic
- Analysis: CRC error, on one side of the line, receiving direction (of course)
- Impact: iperf measurements far away from GBit/s

Network – real life example (1)

- Try and error: change of patch cables, CRC errors disappear immediately, iperf reaches up to 1 Gbit/s
- Take away: Clean your optics/cables. Every time before you plug in. Always. Don‘t discuss. Simply do it! Yes, also with fibre channel!

Network – real life example (2)

- Check_MK rollout without support of the local admin
- Massive errors on one network interface – OK, let me look later…
- Later: Oh, another interface error, this time on a server – is there a connection?

Network – real life example (2)

- Configuration of port names and also trunk/channel names (consistent naming schemes are nice little things) → all related ports are one below the other in the view
- Ohh, whats that?
  Trunk member one with 1Gbit/s
  Trunk member two with 100Mbit/s

Network – real life example (2)

- Check of configuration: OK
- Plug out, plug in 100Mbit/s – mhh
- Plug out, plug in 100Mbit/s – grrr
- Change of patch cable – yay. 1Gbit, errors disappear. Magic.

Network – real life example (2)

- Even later: Accounting lady comes to the admin:
- Lady: „Hans*, what did you do?”
- Hans: „Ehm, ehh, nothing? Why?“
- Lady: „SAP is suddenly lightning fast!“

* name changed, but known to the author

Network – real life example (2)

- Conclusions:
  Check your 5$ patch cables before:
  - you make the SAP consultant rich
  - you upgrade RAM/CPU or even the whole server
- Check_MK does that for you. Automagically. Reliable. If you allow it to do

Network – real life example (3)

- Switch interface between firewall / MPLS router has errors
- Analysis: Collisions, 10Mbit/s half-duplex
- Question to the customer: „Didn‘t you say you have a 34Mbit/s line?“ - „Yes, we bought an upgrade from 10 Mbit/s 1 year ago...“
- Calling the telco… suddenly autonegotiation is on…
- So they increased traffic shaping in the backend but forgot to set the interface from 10/full to 100/full

Network – real life example (3)

- Take away (mostly for the German audience): Deutsche Telekom doesn‘t like autonegotiation. Almost always turned off on business routers.
- And they don‘t tell the customer. Conspiration theory: Saves bandwidth
- Colt behaves similar
- Result: I can find such an error in ~50% of all Check_MK roll-outs in Germany

Network – further real life examples

- Duplex mismatch on an 10Mbit/s “Ethernet Connect” line to a remote office –> less than 1Mbit/s throughput
- „Ethernet Connect“ is a product of which telco? <You name it>

Network – further real life examples

- Error on a Cisco Switch, all links are affected by collisions
- All ports are 100FX optical lines and set to half-duplex.
- Ehm, you can do that, but probably you shouldn‘t. Or why defines FX standard separate send and receive fibers?

Network - duplex-mismatch-take-away

- Duplex mismatches are common.
- Admins often do not detect it by manual checking and underestimate the problem:

  A duplex mismatch degrades a 10Mbit/s line to something around 200Kbit/s

Network – real life example (4)

- Company with 20 locations Europe-wide. Low bandwidth MPLS connections → „country-locations“
- ISP doesn‘t grant SNMP access, says: „Our monitoring says, that you need more bandwidth – please insert coin“
- Conspiration theory: Companies have captalistic motivations

Network – real life example (4)

- Solution: Naming scheme for the Switch to MPLS Router interfaces, e.g.
  MPLS 2Mbit
  ADSL 16Mbit

  etc.
- Interface rules matching that name, including speed, upper limits, measurement unit, averaging

Network – real life example (4)

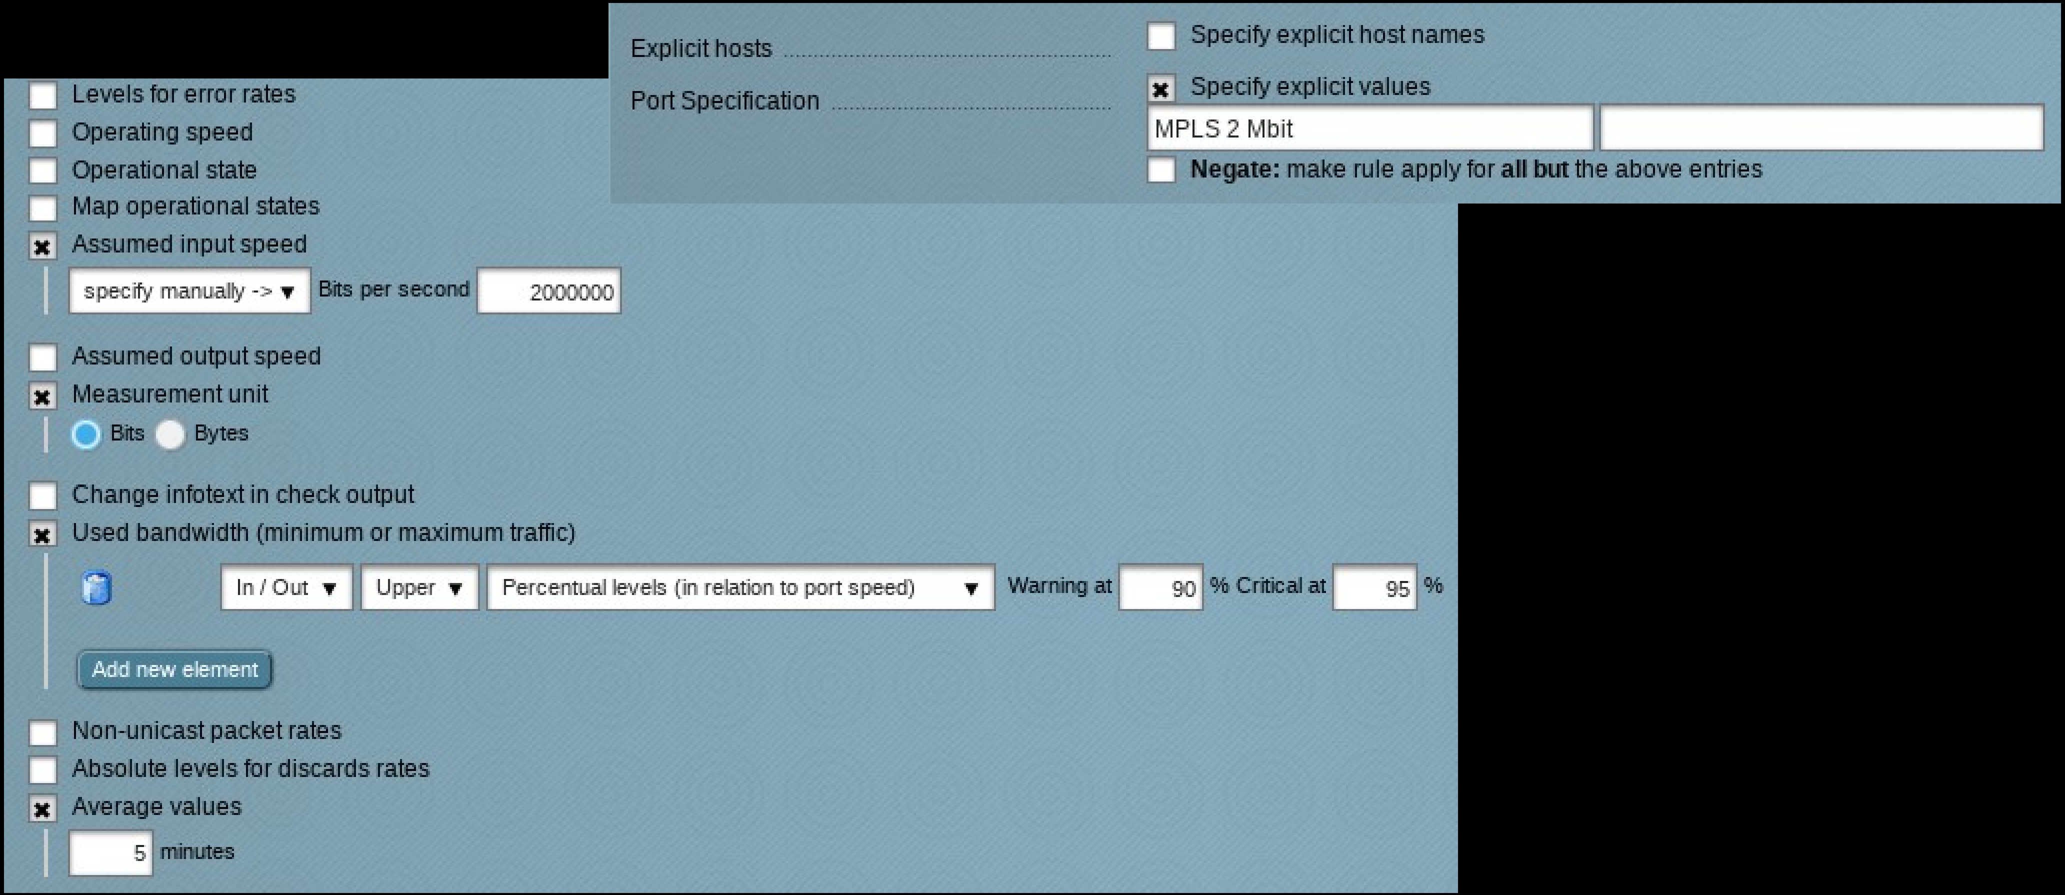

Explicit hosts ☐ Specify explicit host names

Port Specification ☒ Specify explicit values

MPLS 2 Mbit

☐ **Negate:** make rule apply for **all but** the above entries

☐ Levels for error rates

☐ Operating speed

☐ Operational state

☐ Map operational states

☒ Assumed input speed

specify manually -> ▼ Bits per second 2000000

☐ Assumed output speed

☒ Measurement unit

● Bits ○ Bytes

☐ Change infotext in check output

☒ Used bandwidth (minimum or maximum traffic)

In / Out ▼ Upper ▼ Percentual levels (in relation to port speed) ▼ Warning at 90 % Critical at 95 %

☐ Non-unicast packet rates

☐ Absolute levels for discards rates

☒ Average values

5 minutes

Network – real life example (4)

- Two weeks later: Alarm! More than 95% bandwidth usage since 1 hour!
- Analysis of the flows with NTOP-ng: ssh traffic from an IPSec peer
- Students of the partner university didn‘t know the limitation of 2 Mbit/s of that location
- <scp -l limit> is your friend

Network – real life example (5)

- WLAN configuration gets updated, radios are now allowed to use N standard, up to 300 Mbit/s
- Uhmm, the LAN interfaces of the access points are connected to 100Mbit/s switch interfaces, what could possibly go wrong…?

Network – real life example (5)

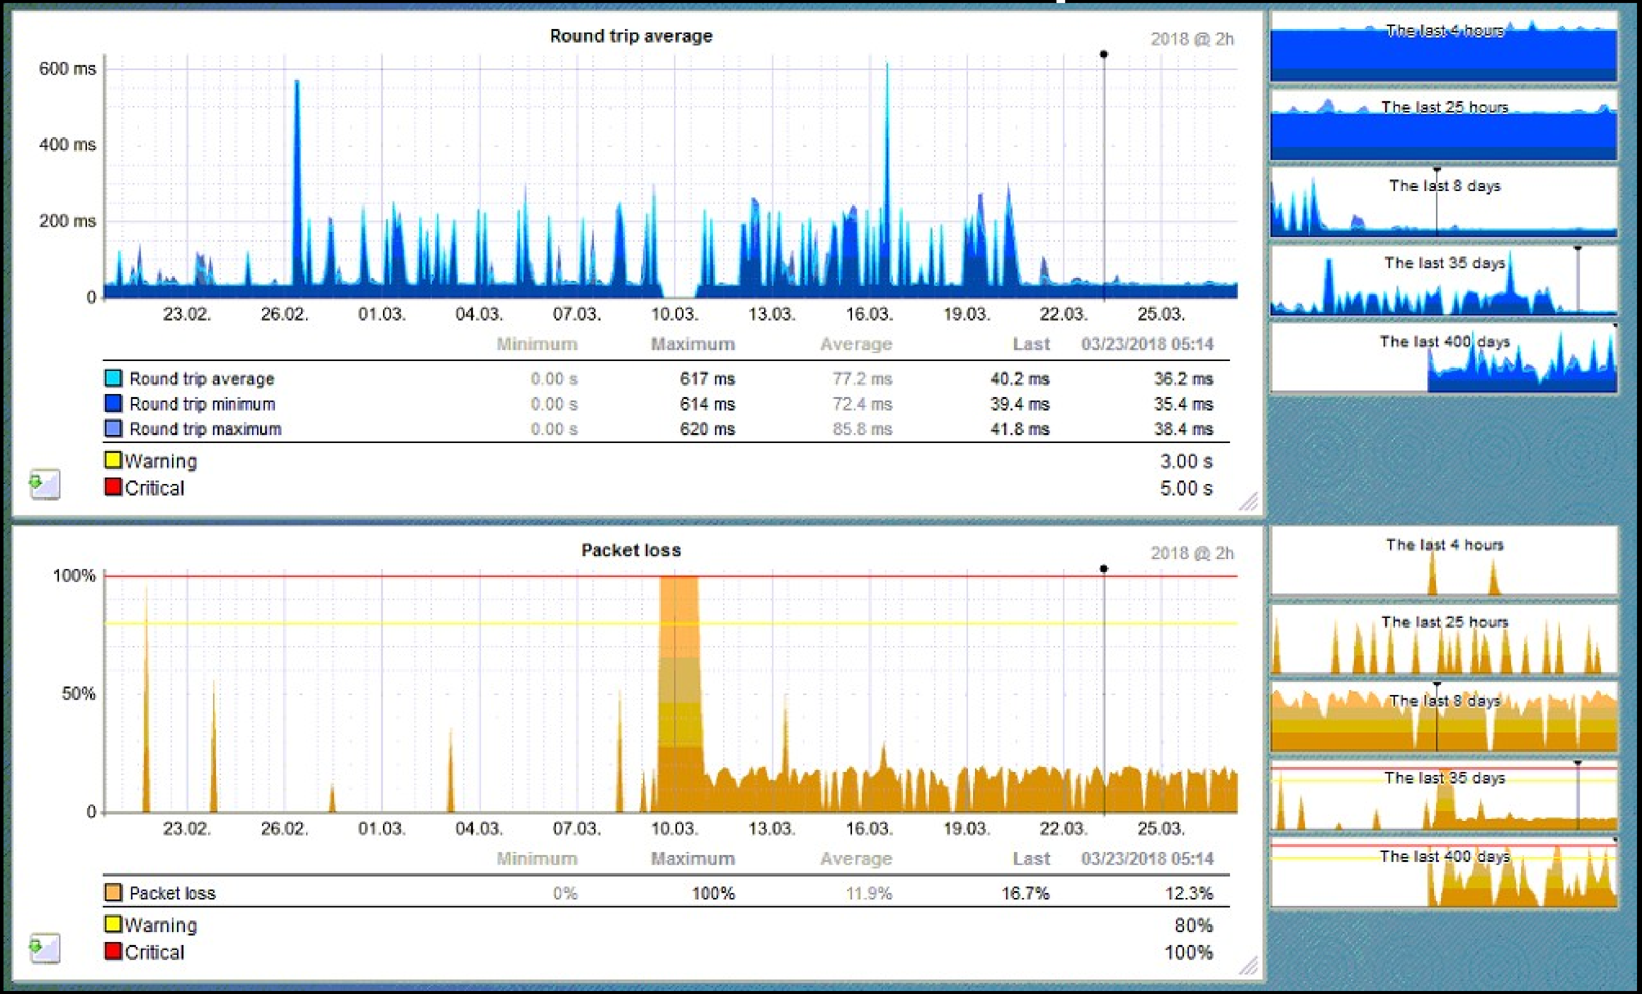

	Minimum	Maximum	Average	Last	03/23/2018 05:14
Round trip average	0.00 s	617 ms	77.2 ms	40.2 ms	36.2 ms
Round trip minimum	0.00 s	614 ms	72.4 ms	39.4 ms	35.4 ms
Round trip maximum	0.00 s	620 ms	85.8 ms	41.8 ms	38.4 ms
Warning					3.00 s
Critical					5.00 s

	Minimum	Maximum	Average	Last	03/23/2018 05:14
Packet loss	0%	100%	11.9%	16.7%	12.3%
Warning					80%
Critical					100%

Network – further real life examples

- Errors on all switch interfaces with connected UTAX printers.(re-branded Kyocera printers)
- Reason: unclear
- Solution: replacing UTAX firmware with the original Kyocera firmware

Network – further real life examples

- Packet-loss, timeouts, slow printouts
- Analysis: ~2Mbit/s basic load on all switchports!?
- Wireshark: Broadcast, Multicast caused by > 13.000 MAC addresses in one VLAN!
- Not so optimal: Embeded TCP/IP stack has to check all multicast packets before it can decide to drop.
- Even 1 core of a 8 core Xeon CPU was 100% busy: avahi-daemon handling multicast requests
- Solution: hang the DJ, VLAN segmentation is your friend

Network – further real life examples

- 10Mbit/s WAN: Bandwidth/packet loss OK, but latency up to DNS timeouts
- Analysis with the Check_MK metrics – packet rate: 10thousands of packets! Small sized as used bandwidth is low.
- Further analysis: POP3 mail fetcher goes wild, as 100MB/Mail are allowed on ISP side, but just 10MB on Exchange side, mail gets refused. Result: Interprets it as network error and reduces packet size.
- Quality programming meets quality administration

Network – further real life examples

- Periodic errors on all switches.
- Analysis: Giants. But only where the VLAN named “WLAN-Mgmt“ is connected.
- Reason: Aruba access points configured to use „client data tunnel“ - needs jumbo frame support in that VLAN.
- Read the specs. Or use Check_MK.

Network – further real life examples

- Switch interface counters stop working (seen on some HP Procurve firmwares)
- Without check_mk: show interfaces - „Wow looks great, no errors!“
- show interfaces is a point in time view, without any time correlation
- Better monitor your interfaces. Yes. All. Do it. Doesn’t hurt.

Network – further real life examples

- Cisco core switch, configured as „virtual stackwise“
- Periodic errors on all connected edge/access switches:

Network – further real life examples

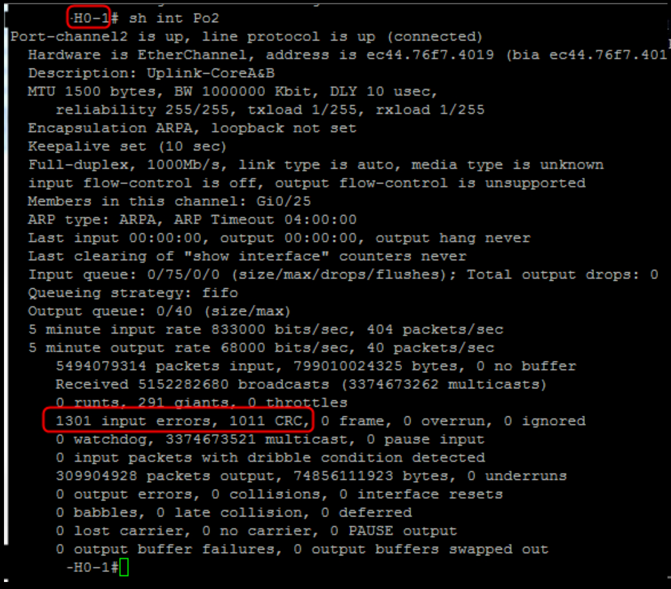

```
      -H0-1# sh int Po2
Port-channel2 is up, line protocol is up (connected)
  Hardware is EtherChannel, address is ec44.76f7.4019 (bia ec44.76f7.401
  Description: Uplink-CoreA&B
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, link type is auto, media type is unknown
  input flow-control is off, output flow-control is unsupported
  Members in this channel: Gi0/25
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:00, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 833000 bits/sec, 404 packets/sec
  5 minute output rate 68000 bits/sec, 40 packets/sec
     5494079314 packets input, 799010024325 bytes, 0 no buffer
     Received 5152282680 broadcasts (3374673262 multicasts)
     0 runts, 291 giants, 0 throttles
     1301 input errors, 1011 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 3374673521 multicast, 0 pause input
     0 input packets with dribble condition detected
     309904928 packets output, 74856111923 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 PAUSE output
     0 output buffer failures, 0 output buffers swapped out
      -H0-1#
```

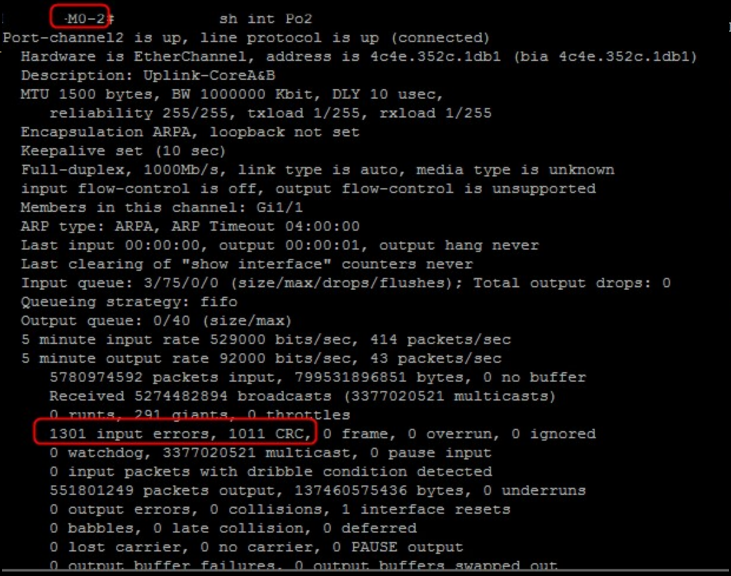

```
      -M0-2#            sh int Po2
Port-channel2 is up, line protocol is up (connected)
  Hardware is EtherChannel, address is 4c4e.352c.1db1 (bia 4c4e.352c.1db1)
  Description: Uplink-CoreA&B
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, link type is auto, media type is unknown
  input flow-control is off, output flow-control is unsupported
  Members in this channel: Gi1/1
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 3/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 529000 bits/sec, 414 packets/sec
  5 minute output rate 92000 bits/sec, 43 packets/sec
     5780974592 packets input, 799531896851 bytes, 0 no buffer
     Received 5274482894 broadcasts (3377020521 multicasts)
     0 runts, 291 giants, 0 throttles
     1301 input errors, 1011 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 3377020521 multicast, 0 pause input
     0 input packets with dribble condition detected
     551801249 packets output, 137460575436 bytes, 0 underruns
     0 output errors, 0 collisions, 1 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 PAUSE output
     0 output buffer failures, 0 output buffers swapped out
```

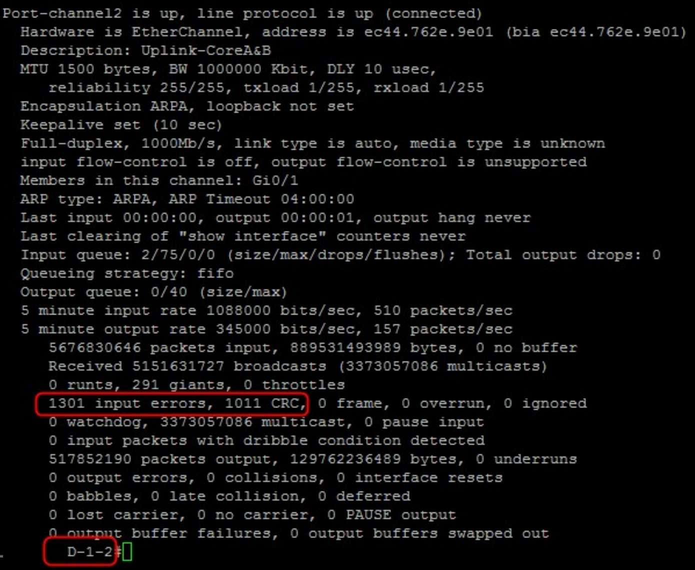

```
Port-channel2 is up, line protocol is up (connected)
  Hardware is EtherChannel, address is ec44.762e.9e01 (bia ec44.762e.9e01)
  Description: Uplink-CoreA&B
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, link type is auto, media type is unknown
  input flow-control is off, output flow-control is unsupported
  Members in this channel: Gi0/1
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 2/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 1088000 bits/sec, 510 packets/sec
  5 minute output rate 345000 bits/sec, 157 packets/sec
     5676830646 packets input, 889531493989 bytes, 0 no buffer
     Received 5151631727 broadcasts (3373057086 multicasts)
     0 runts, 291 giants, 0 throttles
     1301 input errors, 1011 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 3373057086 multicast, 0 pause input
     0 input packets with dribble condition detected
     517852190 packets output, 129762236489 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 PAUSE output
     0 output buffer failures, 0 output buffers swapped out
       D-1-2#
```

- Exact same error count on all connected switches
- CRC Errors, aka broken packets.
- You had one job…

Whats good for a network...

- ...can‘t be bad for Fibre Channel
- Works the same: Port names, error rates and so on
- A dirty optic / cable is even more thrilling than in IP networks (OK, OK packet loss in iSCSI networks is also a nightmare)

Beyond the network

- Some other Check_MK best practice tips
- ..and stories

Server - CPU

- Create this rule, always. Don‘t discuss. Except when your customers is mining bitcoins or so.

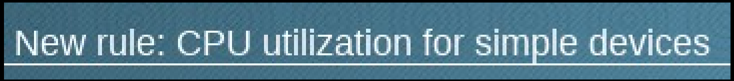

New rule: CPU utilization for simple devices

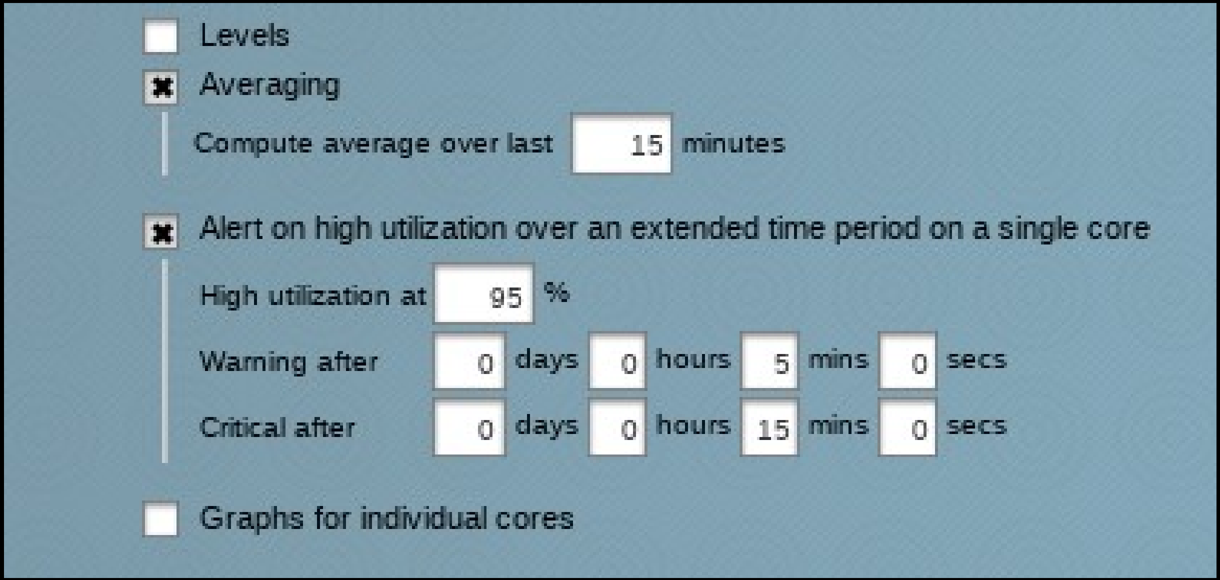

☐ Levels

☒ Averaging

Compute average over last 15 minutes

☒ Alert on high utilization over an extended time period on a single core

High utilization at 95 %

Warning after 0 days 0 hours 5 mins 0 secs

Critical after 0 days 0 hours 15 mins 0 secs

☐ Graphs for individual cores

Server -CPU – real life example (1)

- Continuous high CPU load on one core (Domain controller)
- Customer: „Yes, I know, I need to replace the hardware“
- Some analysis later: Backup Exec going crazy, wants to write a log to c:\program files (x86)\xxx and dosen‘t have permission → known bug
- Bugfix installed, CPU down to almost 0, ~80 Watt less power consumption according to ILO monitoring.

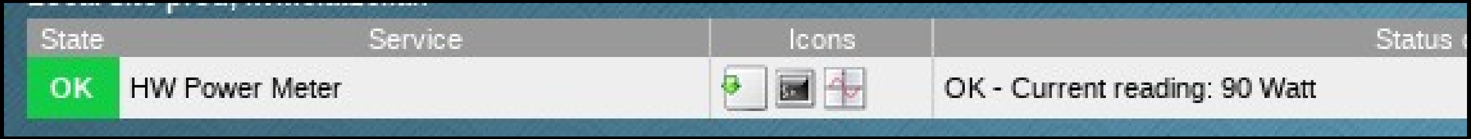

State	Service	Icons	Status
OK	HW Power Meter		OK - Current reading: 90 Watt

Server -CPU – real life example (2)

- Citrix Logon-times far beyond one minute
- Customer: „We have a network issue!“
- Analysis: No network errors at all in Check_MK. But CPU of file server goes high. Peridically, especially in the morning!
- Maybe just one CPU for a file server is a bad idea?
- 4 CPUs and the „network problem“ is gone

ESX Server - Snapshots

- Technical backgrounds of snapshots often unknown by customers
- Old, forgotten snapshots are evil, as snapshots are redo logs that have to be replayed/commited during delete
- Consumes up to the same space as the configured VMDK. Dangerous on almost full LUNs!
- I/O fun for the storage backend during deletion
- Some snapshot based backup solutions tend to „forget“ snapshots after failed/crashed jobs

ESX Server - Snapshots

- So don‘t discuss. Set the following rule. Always.
- If you need to restore a snapshot older than 2 days or so, you are most commonly already in trouble

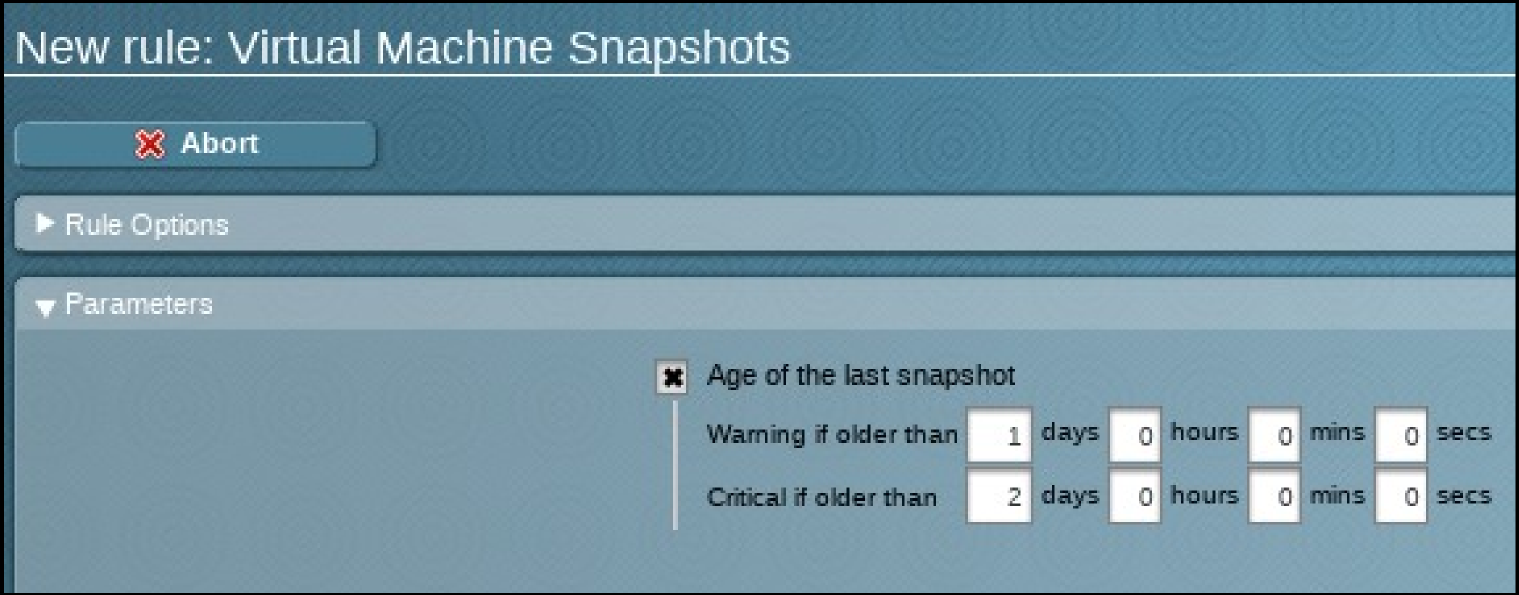

New rule: Virtual Machine Snapshots

☒ Age of the last snapshot

Warning if older than 1 days 0 hours 0 mins 0 secs

Critical if older than 2 days 0 hours 0 mins 0 secs

ESX Server – Snapshots – real life example

- Check_MK roll-out in a relatively new, big, complicated call center installation
- Almost all VMs have snapshots > 200 days
- Storage space is up to be exhausted!
- Check_MK is detecting over-provisioning!

ESX Server – Snapshots – real life example

- Monitoring admin talks with several(!) field engineers:
  “Yeah, after I finished my setup I did a snapshot. So if one of my colleagues breaks my machine one day I can go back to that point.“[sic]

Processes and Services

- Not only for monitoring, but also to control deployments
- „Is the <AV agent/Backup agent/whatever> everywhere installed?“
- „Yes sure, we don‘t need that rule:“

Processes and Services

- → Manual Checks / Windows Services

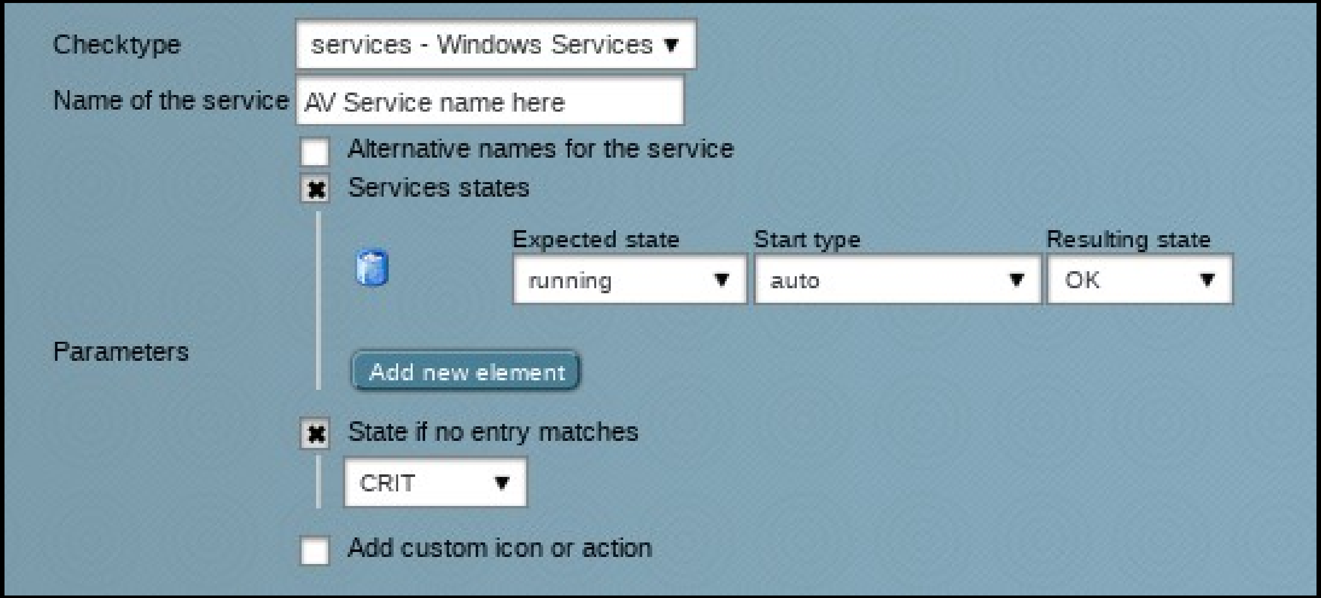

Checktype: services - Windows Services ▾

Name of the service: AV Service name here

Parameters

☐ Alternative names for the service

☒ Services states

Expected state	Start type	Resulting state
running ▾	auto ▾	OK ▾

☒ State if no entry matches

CRIT ▾

☐ Add custom icon or action

- „Uhm, sorry, I forgot that server…“

Processes and Services

- Similar rules can be used to e.g.:
  - Teamviewer service should be installed but not running
  - Monitor all services that are non standard → Discover * auto/running, disable all standard Windows Services by „Disabled Services“ rule

Questions?

Contact:

alexander.wilms[at]bechtle.com