SOLUTION BRIEF

# What Just Happened™ (WJH) Telemetry

## Fully-integrated, Open and Actionable Network Visibility

With the increase in the number of mission-critical workloads running on denser and faster datacenter infrastructures, network downtimes have a bigger impact on business and revenue. According to an IHS report, the cost impact of downtime can range from $1M a year for a typical mid-size company to $60M for a large enterprise. The biggest challenge in reducing downtime is identifying the root-cause of problems. According to Microsoft and others, issue reassignments caused by incorrect root-cause have led to a ten-fold increase in time to resolution. As data center network speeds rapidly increase, so must the tools and mechanisms being used for network visibility.

### Challenges with Legacy Telemetry

Traditional network monitoring approaches like SNMP polling, sFlow sampling, and streaming telemetry, use a central collector for telemetry data storage and problem identification. As data speeds increase, these traditional approaches present a growing challenge to the transportation, processing, and storage of telemetry data. Even a small network can generate petabytes of data per day. Some networks use powerful server clusters to cope with collecting and analyzing the data. Even with this expensive infrastructure, it is hard to pinpoint the root-cause of problems without the appropriate network context.

The difficulty with commodity Ethernet switches is that they are unable to provide the critical telemetry data needed for root cause problem identification or issue resolution.  This forces operators to process massive amounts of telemetry data centrally, creating an artificial bottleneck to problem resolution.

### What Just Happened™ (WJH)

Mellanox Spectrum® Ethernet switches provide rich, contextual and actionable insights on a variety of topics including Layer-1 through Layer-4, ACLs and Buffer occupancy. A WJH agent collects events and visibility insights locally from Mellanox  Spectrum Ethernet switches, to provide instant answers to When, What, Why, Who and Where – critical questions in problem-solving. WJH is available on every Mellanox Spectrum platform running any Network Operating System including Onyx, Cumulus Linux, Linux Switch and Sonic. WJH can be integrated and extended with both third party and open-source tools.

With WJH,  traffic inspection, filtering, and issue identification are done by the switch platform where the network context is readily available. As a result, only issue-relevant data is streamed out. For example, when packets are dropped, WJH implements the appropriate drop counter as well as captures the relevant header data from the packet,  for more thorough analysis.

WJH also helps network operators by dramatically reducing mean time to innocence (MTTI) or issue resolution. Additionally, WJH provides insights to help improve resource utilization and capacity planning.

### Highlights

- Faster troubleshooting
- Switch-based accelerated telemetry
- Works in multi-vendor networks
- Isolates network issues from servers/ storage issues
- Easy to use
- No license required

*“Nimbix accelerates HPC, AI and machine/deep learning applications by providing purpose-built cloud computing technology that is optimized for these workloads,” said Rob Sherrard, Co-Founder & VP of Service Delivery at Nimbix Inc. “Mellanox’s switches are part of our best-in-class cloud computing infrastructure and we are excited to use WJH to increase the utilization of the data center fabric while improving the visibility of the overall network.”*

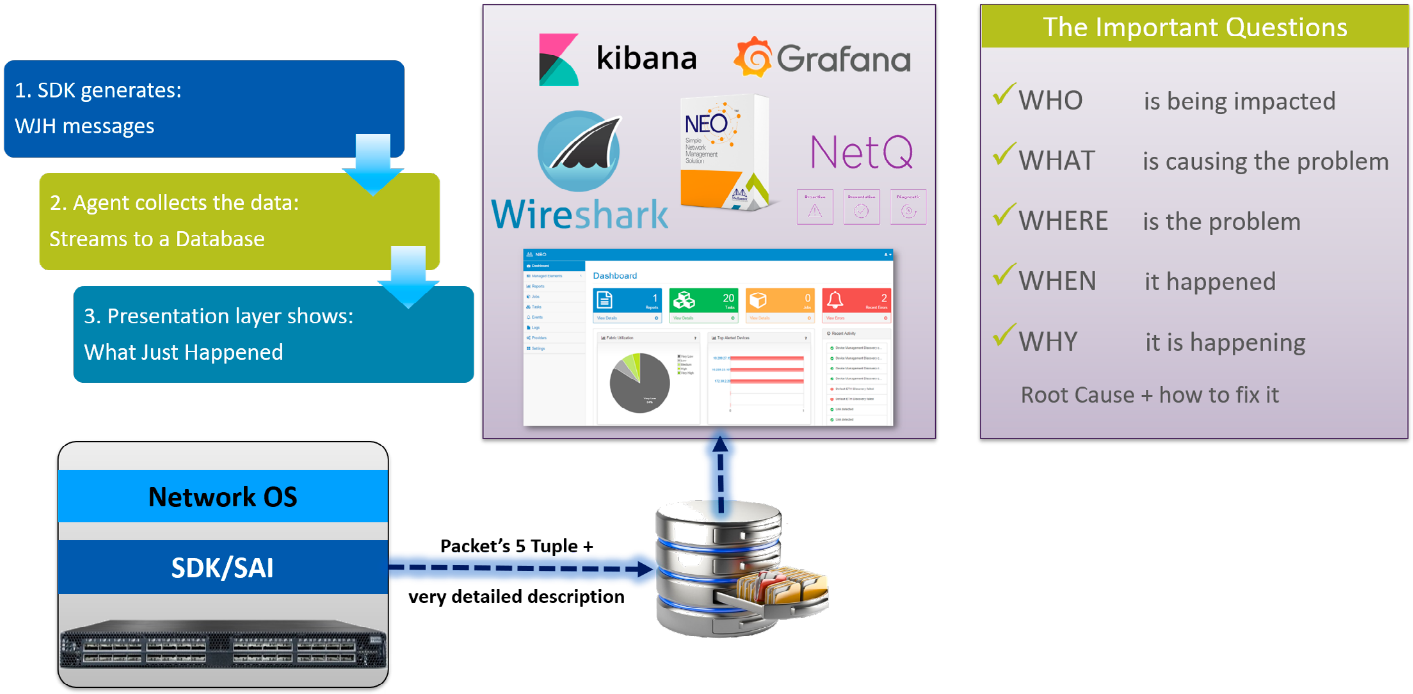

1. SDK generates: WJH messages

2. Agent collects the data: Streams to a Database

3. Presentation layer shows: What Just Happened

The Important Questions

| | |
|---|---|
| ✓WHO | is being impacted |
| ✓WHAT | is causing the problem |
| ✓WHERE | is the problem |
| ✓WHEN | it happened |
| ✓WHY | it is happening |

Root Cause + how to fix it

***Figure 1.*** *WJH - How does it work?*

## How Does it Work?

WJH is an advanced streaming telemetry technology that provides real-time visibility into network-based problems.  WJH goes beyond conventional telemetry solutions by providing actionable details on abnormal network behavior.  Traditional solutions try to extrapolate root causes of network issues by simply analyzing network counters and statistical packet sampling; WJH eliminates the guess-work from network troubleshooting.

The WJH solution leverages the unique hardware capabilities built into the Mellanox Spectrum and Mellanox Spectrum-2 Ethernet switch ASICs to inspect packets at multi-terabit speeds – faster than software or firmware-based solutions.  WJH inspects packets across all ports at line-rate, at speeds that overwhelm traditional Packet Inspection (PI) solutions.

WJH can be configured to monitor certain type of events in the switch and filter out others. Rich telemetry and contextual and actionable telemetry data related to the events of interest are collected from the switch platform. The collected information is locally accessible via a CLI. Detailed packet payload information can be captured in pcap files. The captured packet information can later be analyzed using tools such as Wireshark.

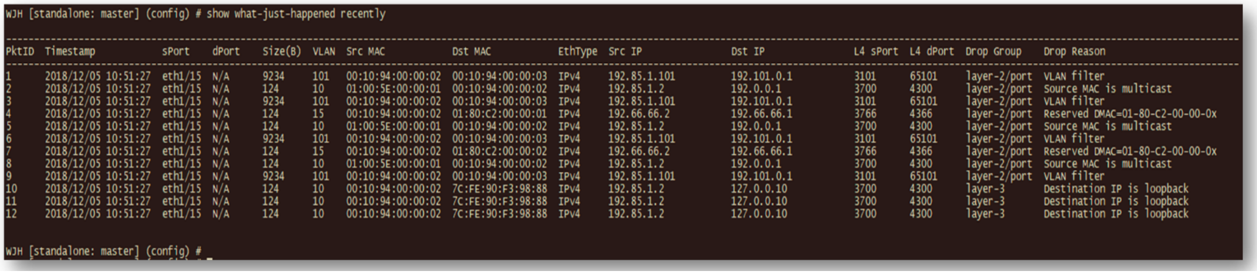

```
WJH [standalone: master] (config) # show what-just-happened recently

---------------------------------------------------------------------------------------------------------------------------------------------------------------------------
PktID  Timestamp            sPort    dPort     Size(B)  VLAN  Src MAC            Dst MAC            EthType  Src IP              Dst IP              L4 sPort  L4 dPort  Drop Group     Drop Reason
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------
1      2018/12/05 10:51:27  eth1/15  N/A       9234     101   00:10:94:00:00:02  00:10:94:00:00:03  IPv4     192.85.1.101        192.101.0.1         3101      65101     layer-2/port   VLAN filter
2      2018/12/05 10:51:27  eth1/15  N/A       124      10    01:00:5E:00:00:01  00:10:94:00:00:02  IPv4     192.85.1.2          192.0.0.1           3700      4300      layer-2/port   Source MAC is multicast
3      2018/12/05 10:51:27  eth1/15  N/A       9234     101   00:10:94:00:00:02  00:10:94:00:00:03  IPv4     192.85.1.101        192.101.0.1         3101      65101     layer-2/port   VLAN filter
4      2018/12/05 10:51:27  eth1/15  N/A       124      15    00:10:94:00:00:02  01:80:C2:00:00:01  IPv4     192.66.66.2         192.66.66.1         3766      4366      layer-2/port   Reserved DMAC=01-80-C2-00-00-0x
5      2018/12/05 10:51:27  eth1/15  N/A       124      10    01:00:5E:00:00:01  00:10:94:00:00:02  IPv4     192.85.1.2          192.0.0.1           3700      4300      layer-2/port   Source MAC is multicast
6      2018/12/05 10:51:27  eth1/15  N/A       9234     101   00:10:94:00:00:02  00:10:94:00:00:03  IPv4     192.85.1.101        192.101.0.1         3101      65101     layer-2/port   VLAN filter
7      2018/12/05 10:51:27  eth1/15  N/A       124      15    00:10:94:00:00:02  01:80:C2:00:00:02  IPv4     192.66.66.2         192.66.66.1         3766      4366      layer-2/port   Reserved DMAC=01-80-C2-00-00-0x
8      2018/12/05 10:51:27  eth1/15  N/A       124      10    01:00:5E:00:00:01  00:10:94:00:00:02  IPv4     192.85.1.2          192.0.0.1           3700      4300      layer-2/port   Source MAC is multicast
9      2018/12/05 10:51:27  eth1/15  N/A       9234     101   00:10:94:00:00:02  00:10:94:00:00:03  IPv4     192.85.1.101        192.101.0.1         3101      65101     layer-2/port   VLAN filter
10     2018/12/05 10:51:27  eth1/15  N/A       124      10    00:10:94:00:00:02  7C:FE:90:F3:98:88  IPv4     192.85.1.2          127.0.0.10          3700      4300      layer-3        Destination IP is loopback
11     2018/12/05 10:51:27  eth1/15  N/A       124      10    00:10:94:00:00:02  7C:FE:90:F3:98:88  IPv4     192.85.1.2          127.0.0.10          3700      4300      layer-3        Destination IP is loopback
12     2018/12/05 10:51:27  eth1/15  N/A       124      10    00:10:94:00:00:02  7C:FE:90:F3:98:88  IPv4     192.85.1.2          127.0.0.10          3700      4300      layer-3        Destination IP is loopback


WJH [standalone: master] (config) # _
```

***Figure 2.*** *WJH via CLI*

As an alternative to CLI, WJH functionality can be accessed via Mellanox NEO®. The filtering and packet data capture capabilities that are available via CLI are also available via NEO. Below is a screenshot of NEO and its WJH interface.

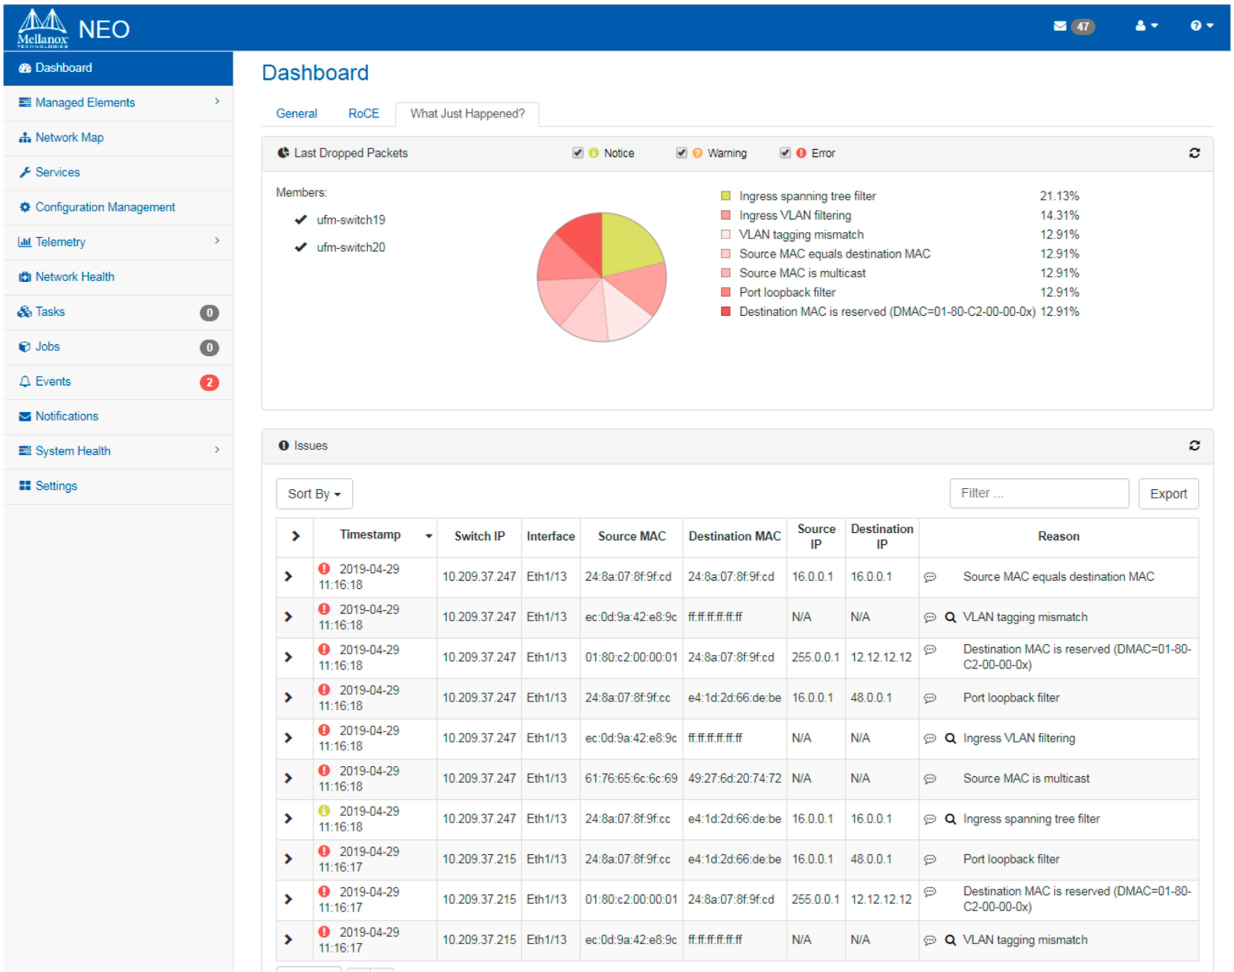

Members:

✓ ufm-switch19

✓ ufm-switch20

| | |
|---|---|
| 🟩 Ingress spanning tree filter | 21.13% |
| 🟥 Ingress VLAN filtering | 14.31% |
| 🟥 VLAN tagging mismatch | 12.91% |
| 🟥 Source MAC equals destination MAC | 12.91% |
| 🟥 Source MAC is multicast | 12.91% |
| 🟥 Port loopback filter | 12.91% |
| 🟥 Destination MAC is reserved (DMAC=01-80-C2-00-00-0x) | 12.91% |

| › | Timestamp | Switch IP | Interface | Source MAC | Destination MAC | Source IP | Destination IP | Reason |
|---|---|---|---|---|---|---|---|---|
| › | 2019-04-29 11:16:18 | 10.209.37.247 | Eth1/13 | 24:8a:07:8f:9f:cd | 24:8a:07:8f:9f:cd | 16.0.0.1 | 16.0.0.1 | Source MAC equals destination MAC |
| › | 2019-04-29 11:16:18 | 10.209.37.247 | Eth1/13 | ec:0d:9a:42:e8:9c | ff:ff:ff:ff:ff:ff | N/A | N/A | VLAN tagging mismatch |
| › | 2019-04-29 11:16:18 | 10.209.37.247 | Eth1/13 | 01:80:c2:00:00:01 | 24:8a:07:8f:9f:cd | 255.0.0.1 | 12.12.12.12 | Destination MAC is reserved (DMAC=01-80-C2-00-00-0x) |
| › | 2019-04-29 11:16:18 | 10.209.37.247 | Eth1/13 | 24:8a:07:8f:9f:cc | e4:1d:2d:66:de:be | 16.0.0.1 | 48.0.0.1 | Port loopback filter |
| › | 2019-04-29 11:16:18 | 10.209.37.247 | Eth1/13 | ec:0d:9a:42:e8:9c | ff:ff:ff:ff:ff:ff | N/A | N/A | Ingress VLAN filtering |
| › | 2019-04-29 11:16:18 | 10.209.37.247 | Eth1/13 | 61:76:65:6c:6c:69 | 49:27:6d:20:74:72 | N/A | N/A | Source MAC is multicast |
| › | 2019-04-29 11:16:18 | 10.209.37.247 | Eth1/13 | 24:8a:07:8f:9f:cc | e4:1d:2d:66:de:be | 16.0.0.1 | 16.0.0.1 | Ingress spanning tree filter |
| › | 2019-04-29 11:16:17 | 10.209.37.215 | Eth1/13 | 24:8a:07:8f:9f:cc | e4:1d:2d:66:de:be | 16.0.0.1 | 48.0.0.1 | Port loopback filter |
| › | 2019-04-29 11:16:17 | 10.209.37.215 | Eth1/13 | 01:80:c2:00:00:01 | 24:8a:07:8f:9f:cd | 255.0.0.1 | 12.12.12.12 | Destination MAC is reserved (DMAC=01-80-C2-00-00-0x) |
| › | 2019-04-29 11:16:17 | 10.209.37.215 | Eth1/13 | ec:0d:9a:42:e8:9c | ff:ff:ff:ff:ff:ff | N/A | N/A | VLAN tagging mismatch |

***Figure 3.*** *Mellanox NEO® interface for WJH*

WJH is an Open Ethernet solution that can be integrated into open source tools like Grafana, and Kibana, but it also works with turn-key data center-wide monitoring solutions like Cumulus NetQ. The data collected on the switch can be streamed out of the switch via gRPC in JSON format using a containerized agent (provided by Mellanox). The streamed telemetry data can be centrally stored in a time-series database such as InfluxDB. Visualization tools such as Grafana can be used to display the data from the database.

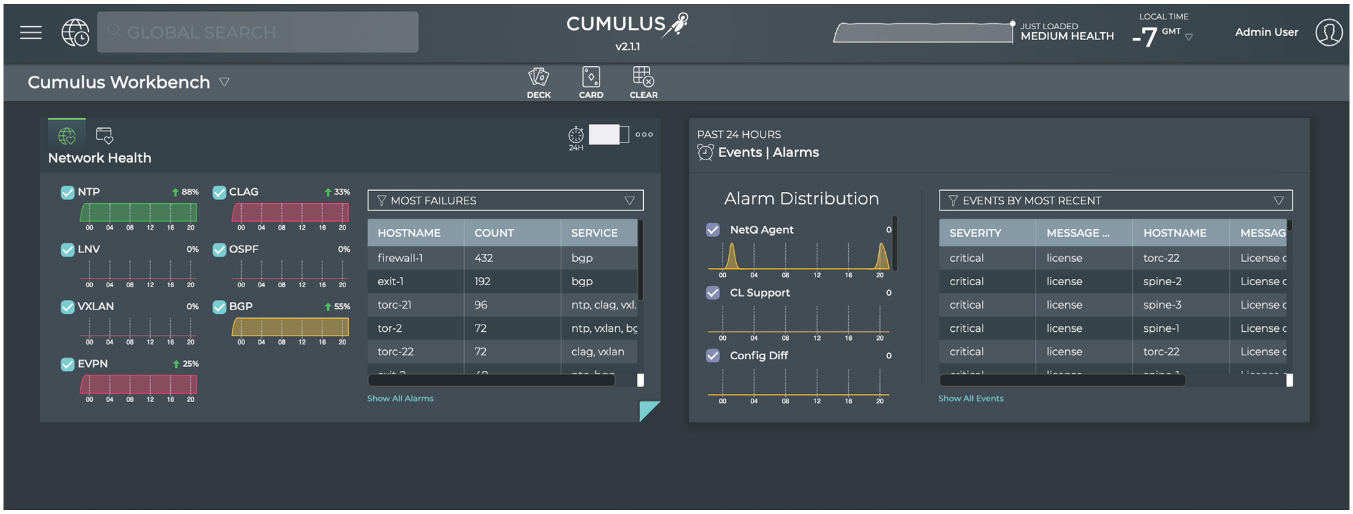

MOST FAILURES

| HOSTNAME | COUNT | SERVICE |
|---|---|---|
| firewall-1 | 432 | bgp |
| exit-1 | 192 | bgp |
| torc-21 | 96 | ntp, clag, vxl. |
| tor-2 | 72 | ntp, vxlan, bg |
| torc-22 | 72 | clag, vxlan |

Show All Alarms

EVENTS BY MOST RECENT

| SEVERITY | MESSAGE ... | HOSTNAME | MESSAG |
|---|---|---|---|
| critical | license | torc-22 | License c |
| critical | license | spine-2 | License c |
| critical | license | spine-3 | License c |
| critical | license | spine-1 | License c |
| critical | license | torc-22 | License c |

Show All Events

***Figure 4.*** *Cumulus NetQ dashboard*

WJH can also be integrated with third-party software such as Apstra and Cumulus NetQ to provide rich and comprehensive visibility into datacenter networks.

## Summary

Traditional network visibility methods fall short as datacenters adopt 100GbE and higher speeds. WJH provides an open, extensible and actionable telemetry that can help operators dramatically reduce Mean Time to Issue Resolution and improve infrastructure uptimes at even higher data rates. WJH leverages Mellanox Spectrum silicon level capabilities to provide granular, contextual and actionable telemetry information on datacenter networks. Mellanox Spectrum SDK is tightly integrated with WJH and provides an open and extensible API framework to build fully-integrated visibility frameworks using 3rd party as well as open-source tools.

WJH can directly be used on the switch platforms via CLI. Alternatively, the WJH functionality is also available through a Web UI or NEO. A containerized WJH streaming agent can be used to stream out WJH events out of the switch. Open-source time series databases such as InfluxDB and visualization tools such as Grafana can be used to visualize the data in a central dashboard.

350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com