

Leveraging Zabbix trend functions for network monitoring anomaly detection and enhanced visibility





Introducing the issue

- **Issue overview**
 - Lots of operational monitoring, almost no tactical vision
 - Lack of essential information for the network environment management
 - In general, the network behavior in relation to its traffic is unknown and network administrators are reactive professionals



Purpose of the implementation and expected results

- **What is the purpose of the implementation?**
 - Improving or generating a network tactical layer monitoring
 - Using Zabbix Trend Functions
- **What is the expected result?**
 - Some piece of information about network behavior

1

Contextualizing

2

Scenario

3

Ordinary data

4

What's new?

5

Baseline

6

Standard Deviation

7

Anomaly Rate

8

Conclusion





1

Contextualizing

Contextualizing

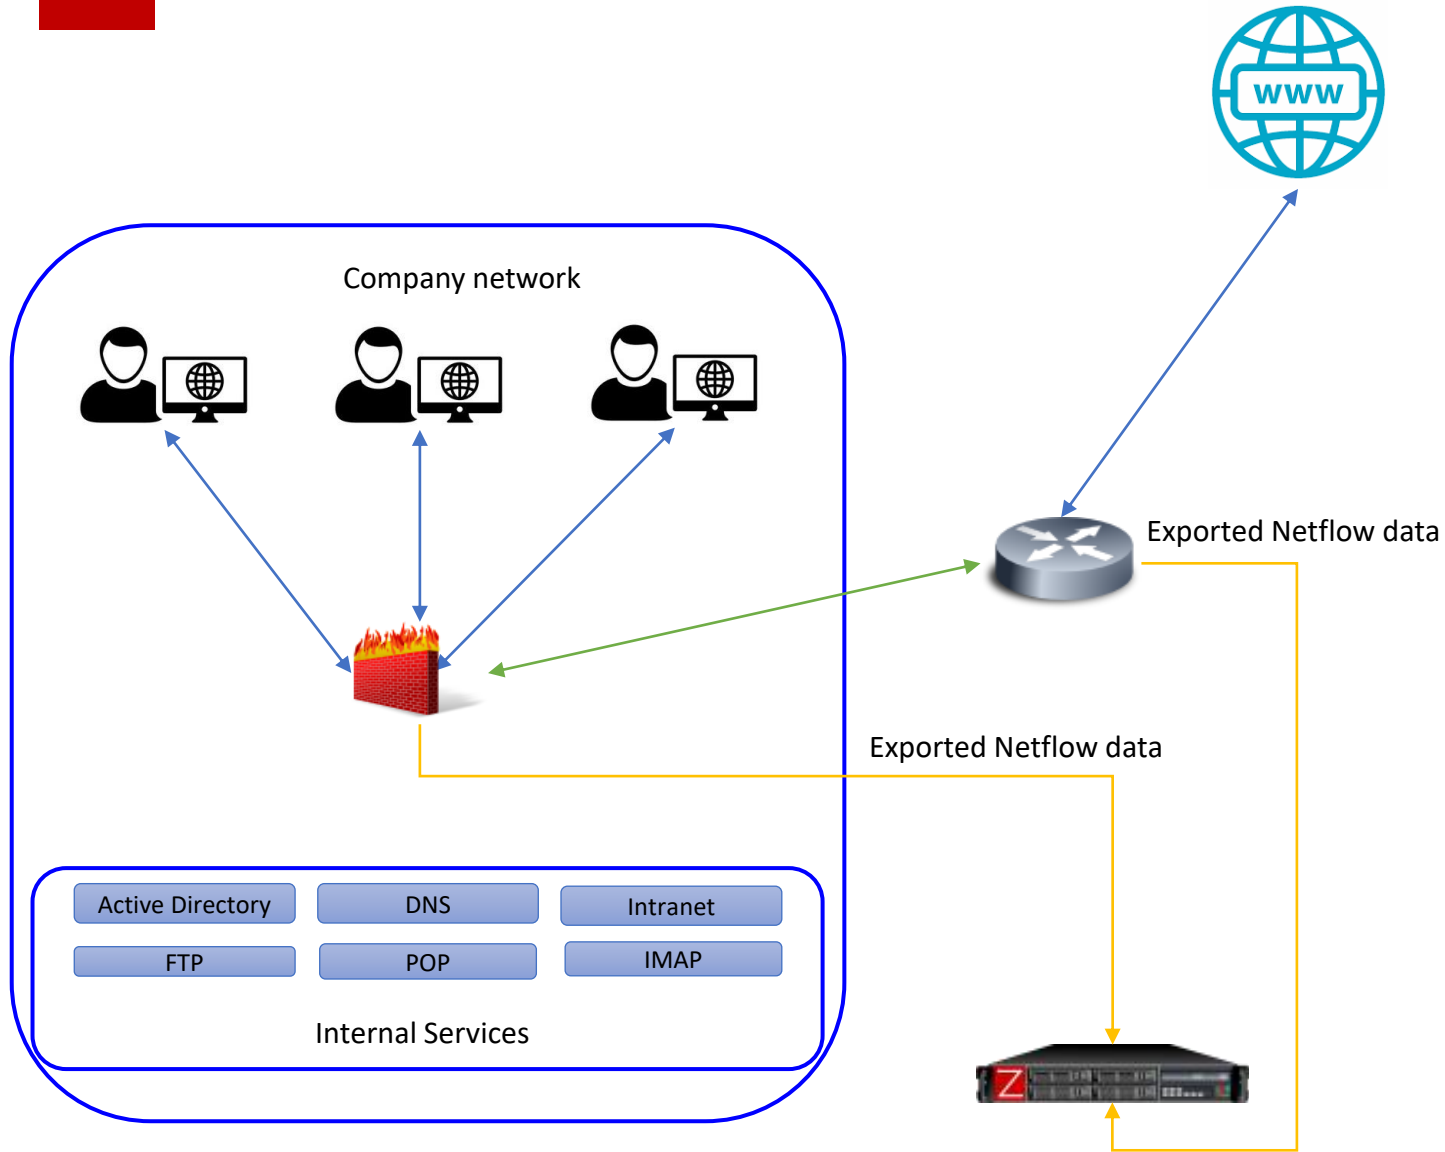
- **What do I really need to monitor?**
 - CPU, memory, disk, etc...
- **Talking about network traffic**
 - The 5-tuple concept and some additional information
- **Why and how Netflow?**
 - What's happening on the network?
- **How can Zabbix collect, store and analyse Netflow data?**
 - I can teach it and Zabbix can learn it (what if...?)
- **So, how can Zabbix help us understand the network behavior?**
 - Zabbix can bring to light some unknow behaviors



2

Scenario

Scenario



The 5-tuple:

- Source IP = x
- Source port = y
- Destination IP = z
- Destination port = w
- Protocol = p

Additional information

- Flows
- Packets
- Bytes



3

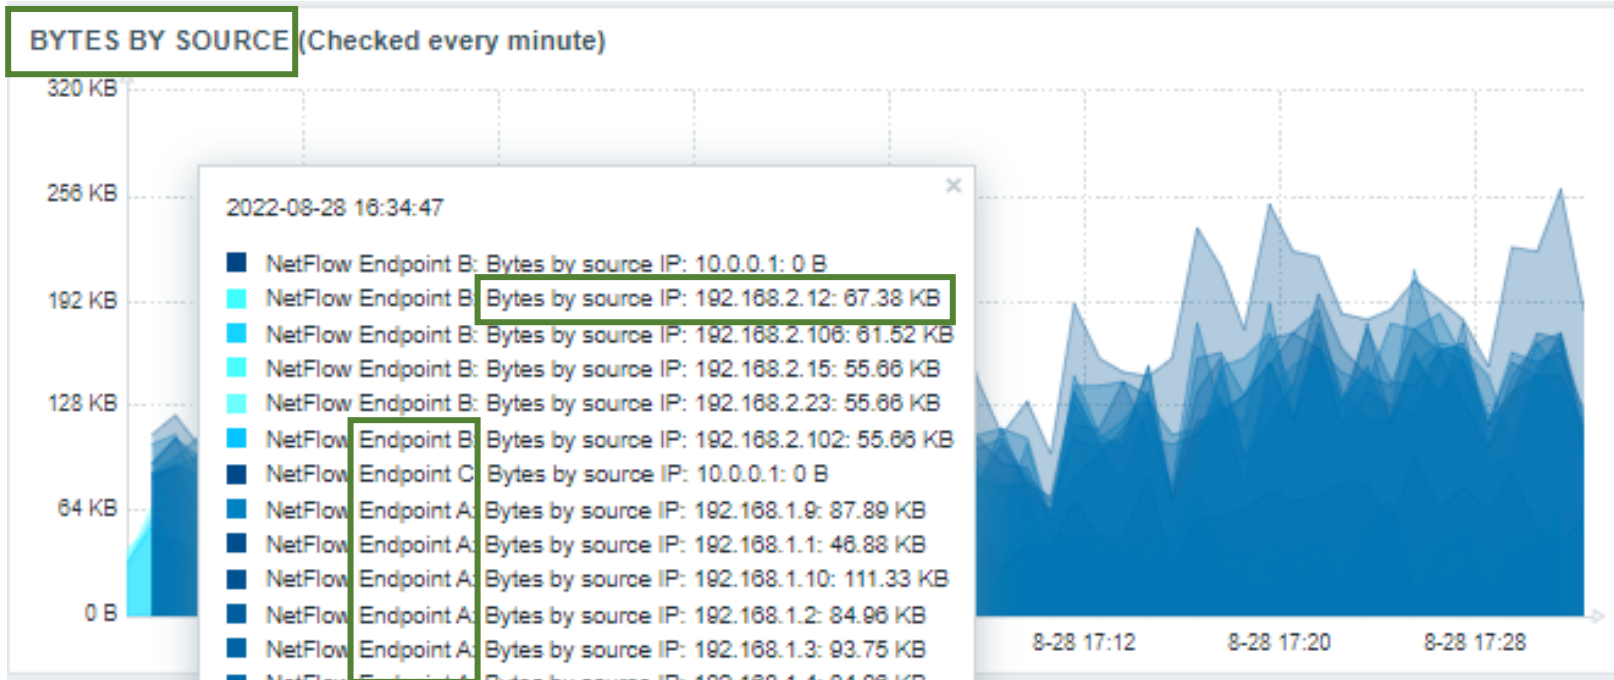
Ordinary data



Ordinary data

- **What do I know so far?**
- **Can I say my network behavior is OK with these data?**
- **What was expected and what is still not clear?**

Some prints – Part 1.1: Dashboard widget = Graph



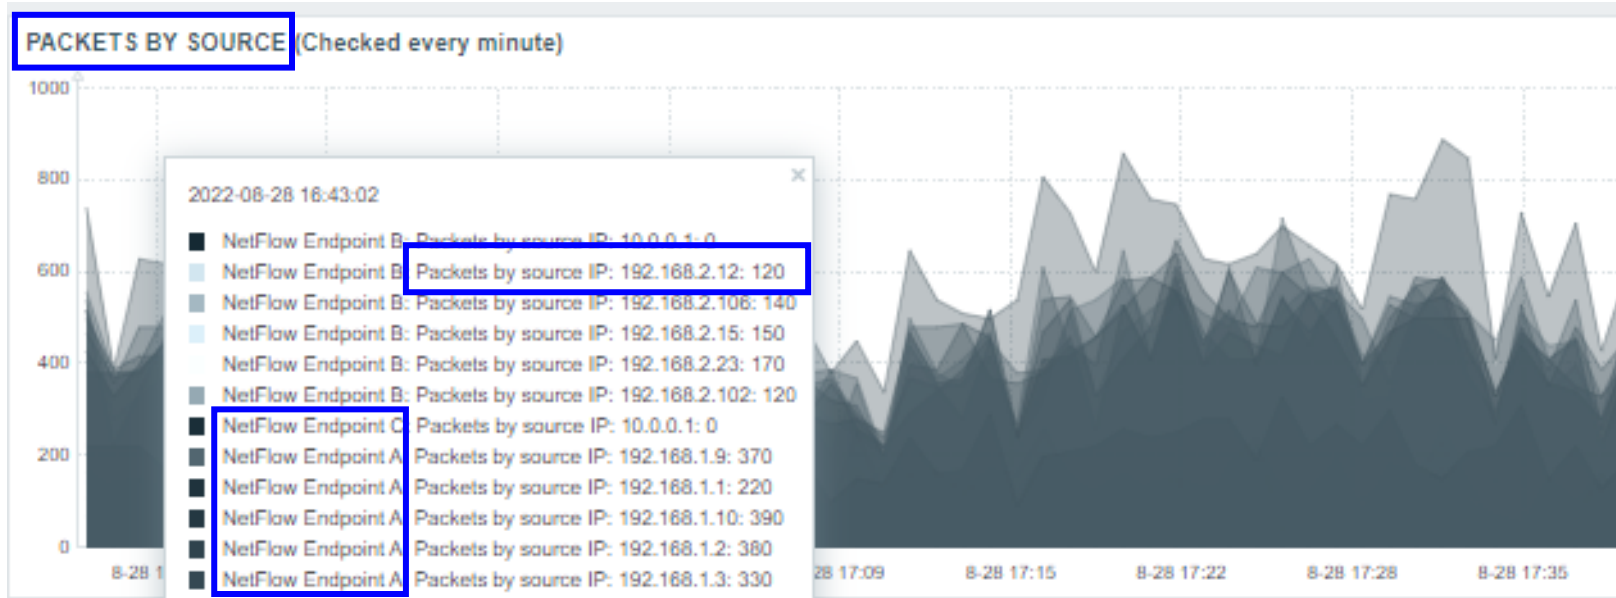
The 5-tuple:

- Source IP = x
- Source port = y
- Destination IP = z
- Destination port = w
- Protocol = p

Additional information

- Flows
- Packets
- Bytes

Some prints – Part 1.2: Dashboard widget = Graph



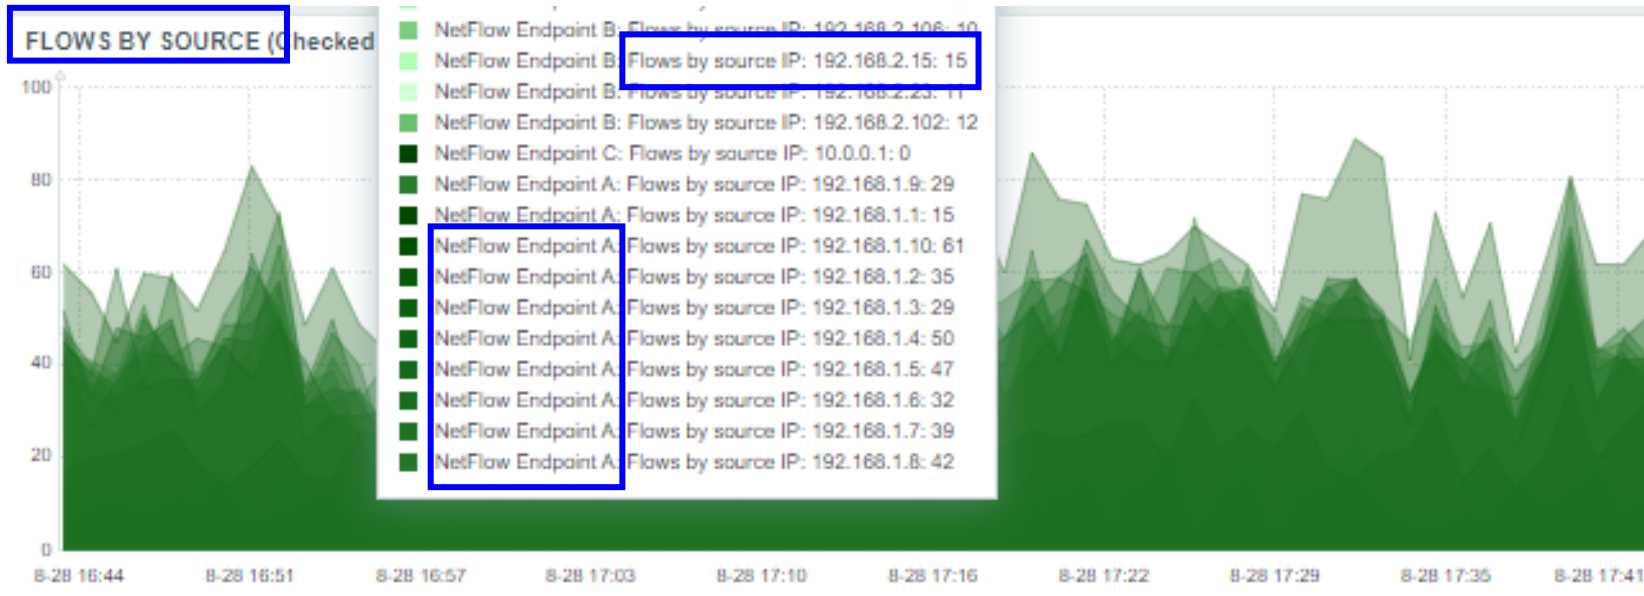
The 5-tuple:

- Source IP = x
- Source port = y
- Destination IP = z
- Destination port = w
- Protocol = p

Additional information

- Flows
- Packets
- Bytes

Some prints – Part 1.2: Dashboard widget = Graph



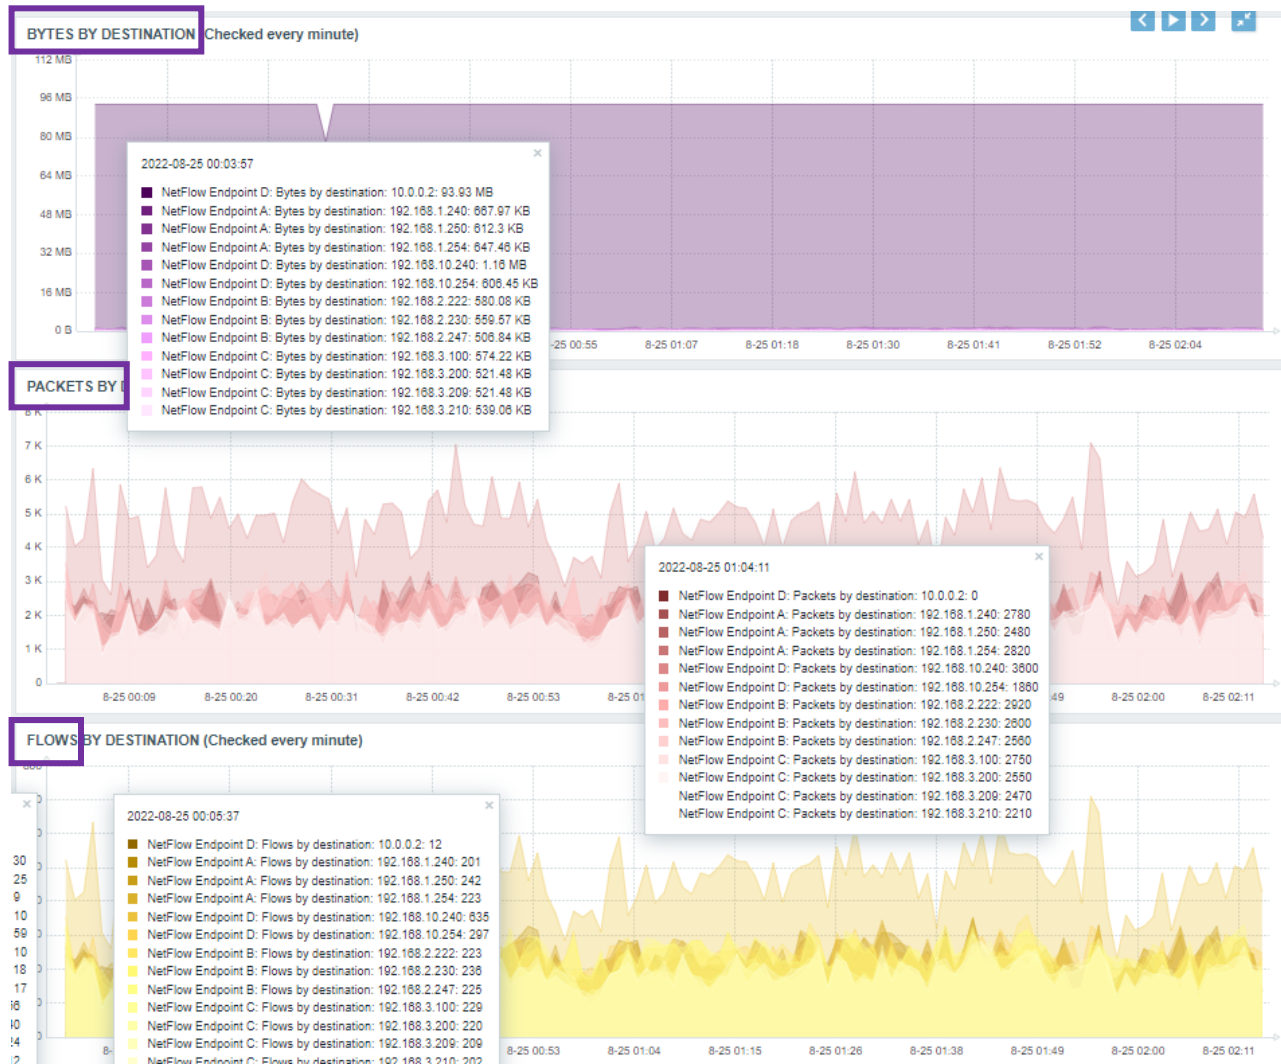
The 5-tuple:

- Source IP = x
- Source port = y
- Destination IP = z
- Destination port = w
- Protocol = p

Additional information

- Flows
- Packets
- Bytes

Some prints – Part 1.4: Dashboard widget = Graph



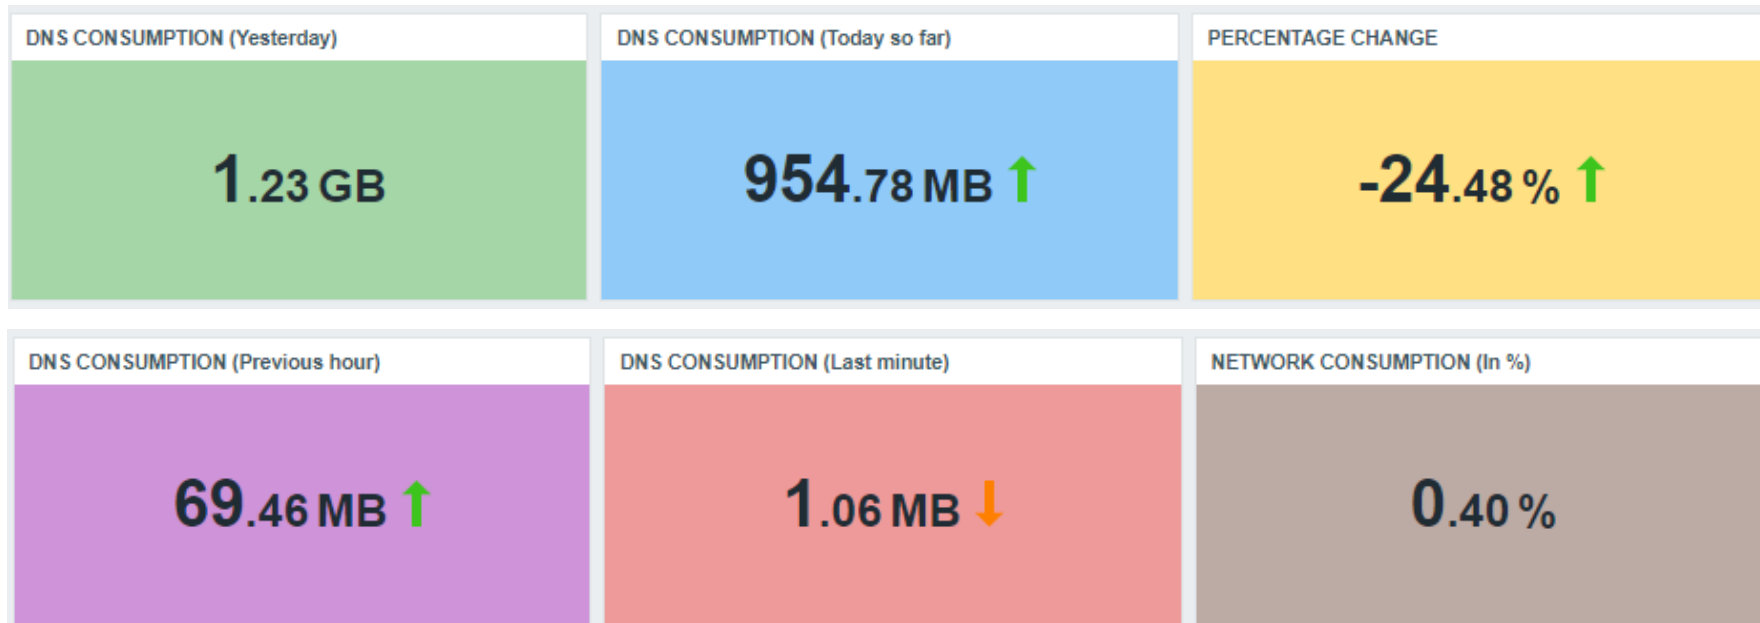
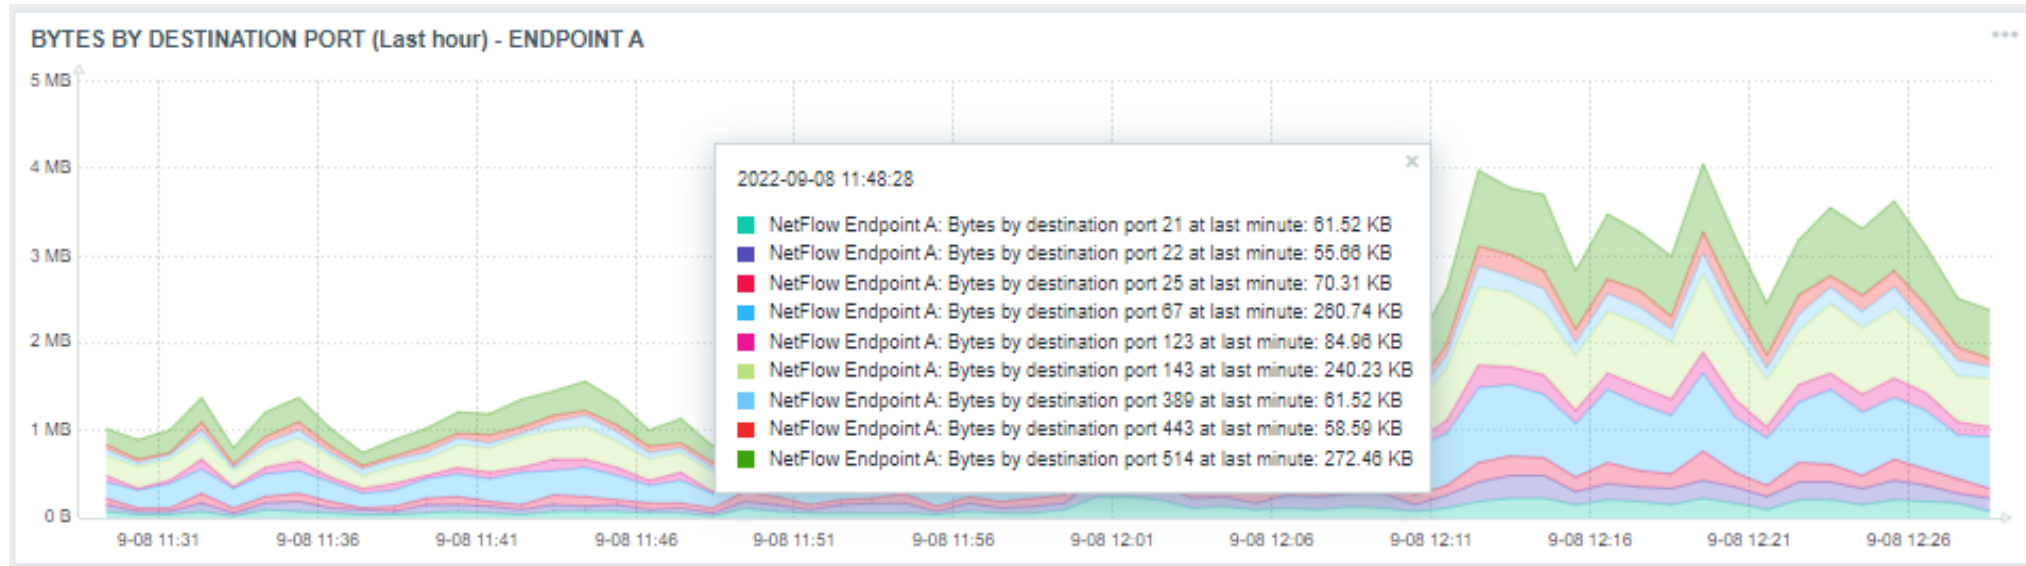
The 5-tuple:

- Source IP = x
- Source port = y
- Destination IP = z
- Destination port = w
- Protocol = p

Additional information

- Flows
- Packets
- Bytes

Some prints – Part 2: Dashboard widget = Top host + Graph





4


What's new?




Extraordinary data

- **What does extraordinary data means?**
- **What do you think about observability?**
- **Just monitor or search for something else?**

BYTES YESTERDAY 4.15 GB <small>2022-09-05 00:05:01</small>	BYTES TODAY SO FAR 2.66 GB ↑ <small>2022-09-05 18:06:28</small>	PERCENTAGE CHANGE -35.92 % ↑ <small>2022-09-05 18:05:46</small>	AVG / LAST HOUR > NOW <small>Using history values</small> 3.16 MB ↑ <small>2022-09-05 18:06:05</small>	AVG BYTES / PREVIOUS HOUR <small>Using trend values</small> 3.14 MB <small>2022-09-05 18:05:00</small>	STANDARD DEVIATION <small>Previous hour, 7 days, same timeframe</small> 0.99 STD DEV ↑ <small>2022-09-05 18:01:00</small>	ANOMALY RATE <small>Based on 24 hours, finding anomalies at last 2 hours, deviation 2</small> 0.13 RATE <small>2022-09-05 18:05:00</small>
--	---	---	--	---	--	---



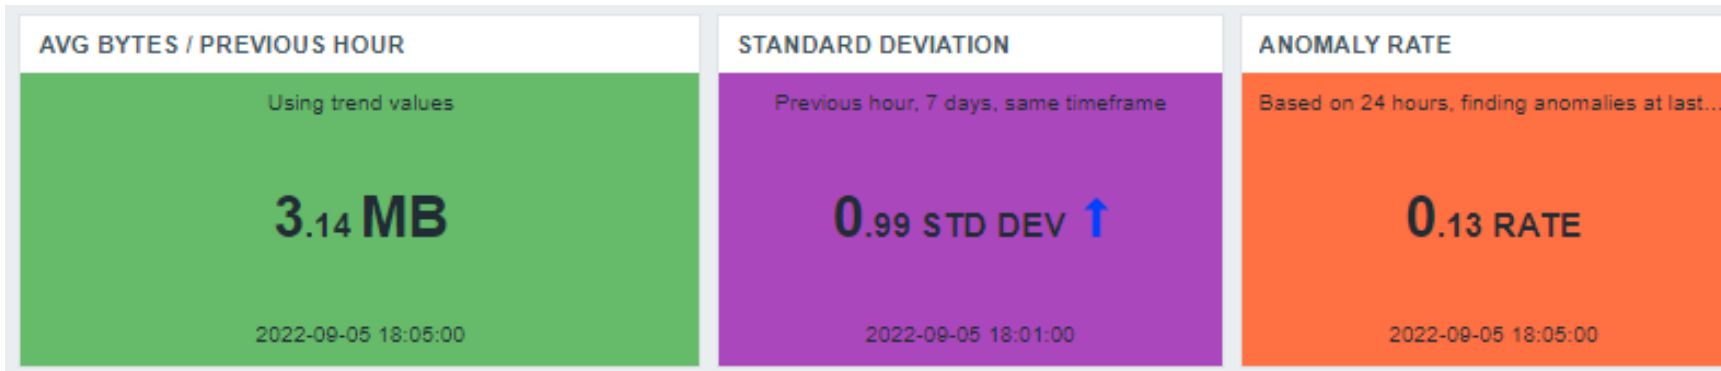
TCP BYTES YESTERDAY 2.76 GB <small>2022-09-05 00:05:01</small>	TCP BYTES TODAY SO FAR 2.06 GB ↑ <small>2022-09-05 18:06:36</small>	PERCENTAGE CHANGE -25.69 % ↑ <small>2022-09-05 18:06:30</small>	AVG / LAST HOUR > NOW <small>Using history values</small> 2.43 MB ↑ <small>2022-09-05 18:06:05</small>	AVG BYTES / PREVIOUS HOUR <small>Using trend values</small> 687.69 KB <small>2022-09-05 18:05:00</small>	STANDARD DEVIATION <small>Previous hour, 7 days, same timeframe</small> 1.39 DEV ↑ <small>2022-09-05 18:01:00</small>	ANOMALY RATE <small>Based on 24 hours, finding anomalies at last 12 hours, deviation 2</small> 0.13 RATE <small>2022-09-05 18:05:00</small>
UDP BYTES YESTERDAY 1.32 GB <small>2022-09-05 00:05:01</small>	UDP BYTES TODAY SO FAR 566.82 MB ↑ <small>2022-09-05 18:06:46</small>	PERCENTAGE CHANGE -58.12 % ↑ <small>2022-09-05 18:06:35</small>	AVG / LAST HOUR > NOW <small>Using history values</small> 697.65 KB ↑ <small>2022-09-05 18:06:05</small>	AVG BYTES / PREVIOUS HOUR <small>Using trend values</small> 687.69 KB <small>2022-09-05 18:05:00</small>	STANDARD DEVIATION <small>Previous hour, 7 days, same timeframe</small> 0.05 DEV ↓ <small>2022-09-05 18:01:00</small>	ANOMALY RATE <small>Based on 24 hours, finding anomalies at last 12 hours, deviation 2</small> 0.13 RATE <small>2022-09-05 18:05:00</small>
ICMP BYTES YESTERDAY 67.37 MB <small>2022-09-05 00:05:01</small>	ICMP BYTES TODAY SO FAR 50.83 MB ↑ <small>2022-09-05 18:06:30</small>	PERCENTAGE CHANGE -24.63 % ↑ <small>2022-09-05 18:06:30</small>	AVG / LAST HOUR > NOW <small>Using history values</small> 48.00 KB <small>2022-09-05 18:06:05</small>	AVG BYTES / PREVIOUS HOUR <small>Using trend values</small> 48.00 KB <small>2022-09-05 18:05:00</small>	STANDARD DEVIATION <small>Previous hour, 7 days, same timeframe</small> 0.16 DEV ↑ <small>2022-09-05 18:01:00</small>	ANOMALY RATE <small>Based on 24 hours, finding anomalies at last 12 hours, deviation 2</small> 0.08 RATE <small>2022-09-05 18:05:00</small>
OTHER BYTES YESTERDAY 0 B <small>2022-09-05 00:05:01</small>	OTHER BYTES TODAY SO FAR 0 B <small>2022-09-05 18:01:04</small>	PERCENTAGE CHANGE No data <small>2022-09-05 18:06:42</small>	AVG / LAST HOUR > NOW <small>Using history values</small> 0.00 B <small>2022-09-05 18:06:05</small>	AVG BYTES / PREVIOUS HOUR <small>Using trend values</small> 0.00 B <small>2022-09-05 18:05:00</small>	STANDARD DEVIATION <small>Previous hour, 7 days, same timeframe</small> 0.00 DEV <small>2022-09-05 18:01:00</small>	ANOMALY RATE <small>Based on 24 hours, finding anomalies at last 12 hours, deviation 2</small> 0.00 RATE <small>2022-09-05 18:05:00</small>



Comparing both

BYTES YESTERDAY	BYTES TODAY SO FAR	PERCENTAGE CHANGE
4.15 GB 2022-09-05 00:05:01	2.66 GB ↑ 2022-09-05 18:06:28	-35.92 % ↑ 2022-09-05 18:05:46

How far from the average?



Value for the last full hour

Deviations found



Key concept #1

- **What is Baseline?**
 - A way of analyzing past behavior so that I can predict what to receive



Key concept #2

- **What is Standard Deviation?**
 - A way of knowing how “far” values are from average



Key concept #3

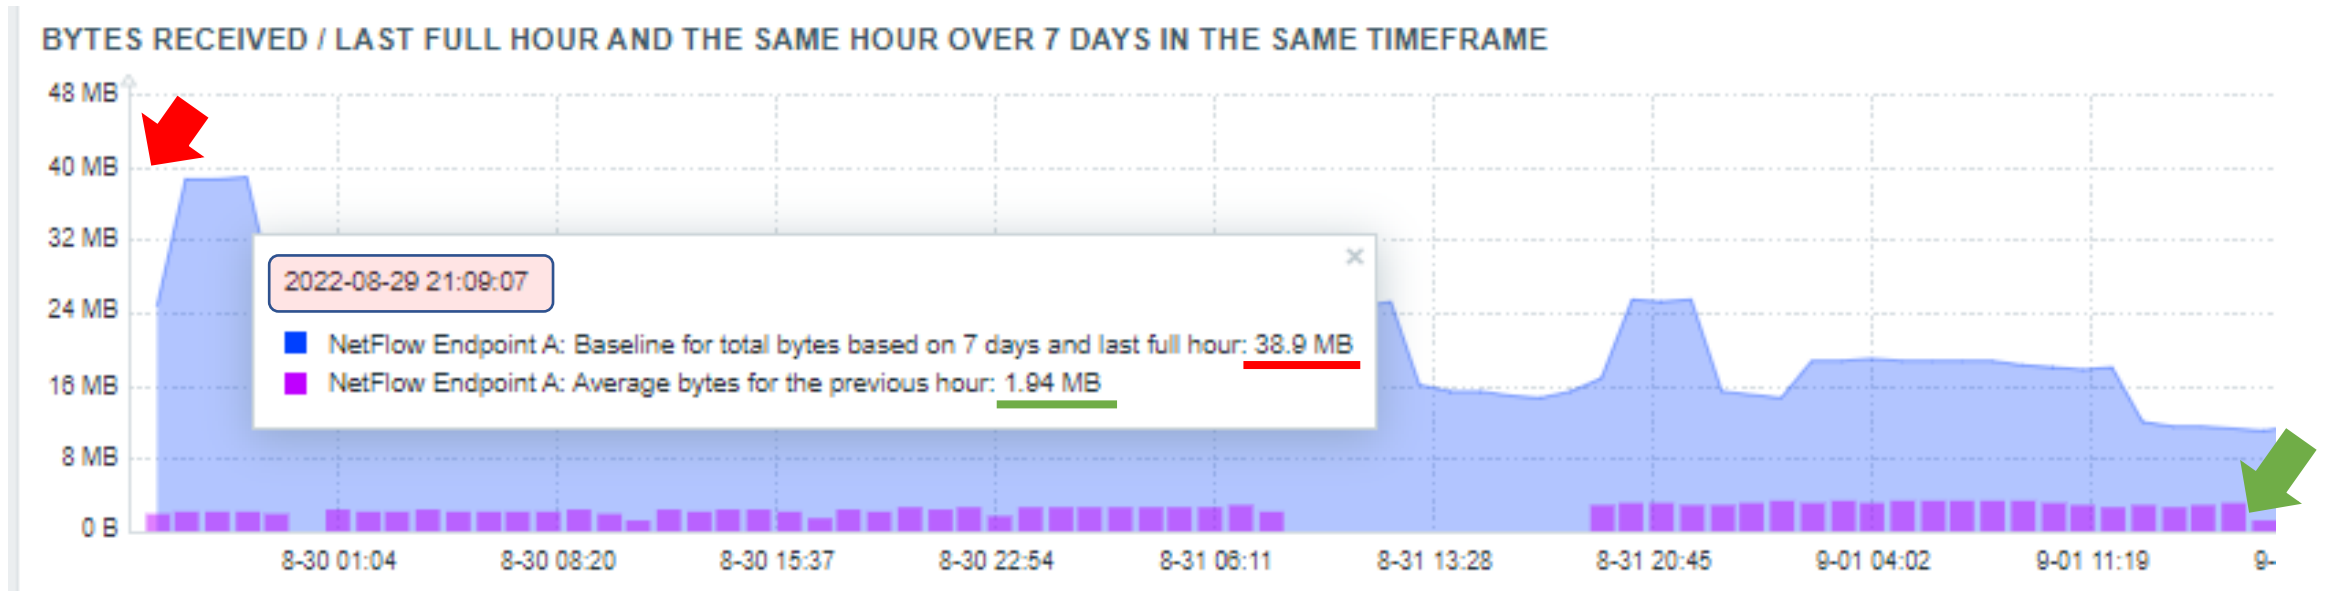
- **What is Anomaly Rate?**
 - it is an indicator based on the number of deviations reached in a detection period



5

Baseline

Baseline for total bytes received from a Netflow exporter (firewall, router, etc.)




`baselinewma(//total.bytes,1h:now/h,"d",7)`

`trendavg(//total.bytes,1h:now/h)`

Baseline for total bytes received from a Netflow exporter (firewall or router)

Same time, few days later



```
baselinewma(//total.bytes,1h:now/h,"d",7)
```

```
trendavg(//total.bytes,1h:now/h)
```


Graph explanation #Used baseline formula

Defining the data gathering period

Defining periodicity

Formula: `baselinewma(//total.bytes,1h:now/h,"d",7)`

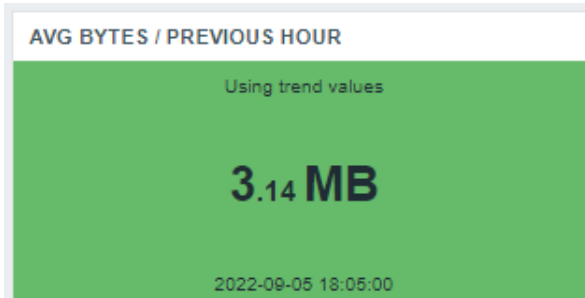
Seasonality



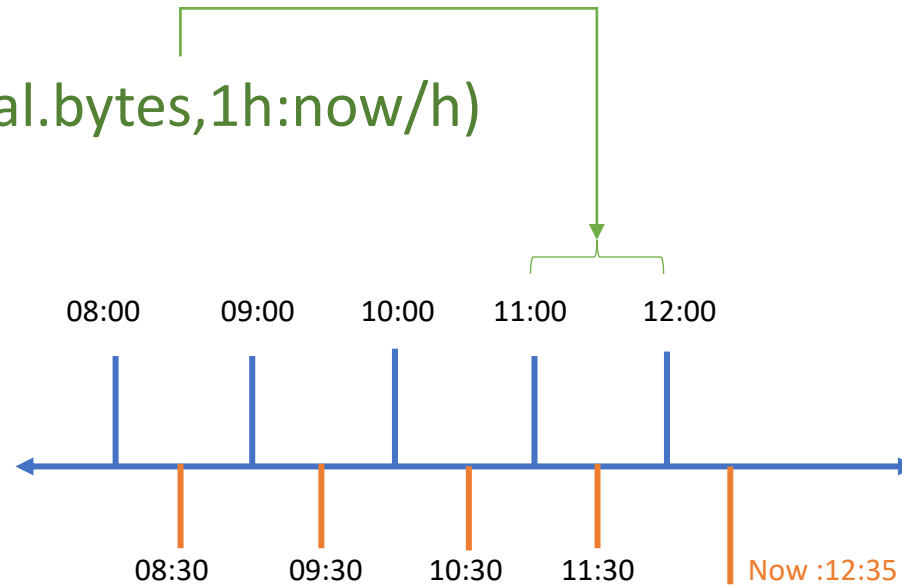
Based on received values in the last full hour: 11:00 – 11:59

“trendavg x avg”

Formula: `trendavg(//total.bytes,1h:now/h)`
The last full hour




Using trends



Formula: `avg(//total.bytes,1h)`
Last hour until now

Using history






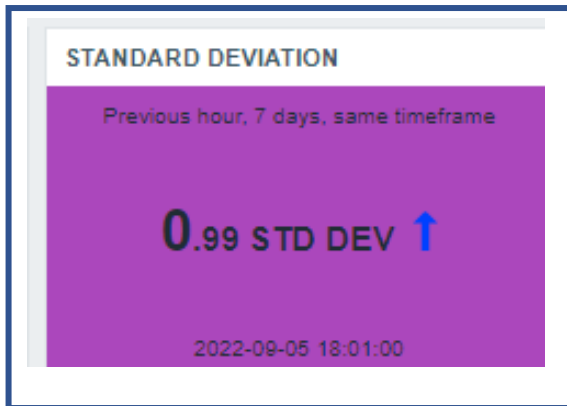
6

Standard Deviation

Standard Deviation



Standard Deviation



Formula: `baselinedev(//total.bytes,1h:now/h,"d",7)`

Standard Deviation

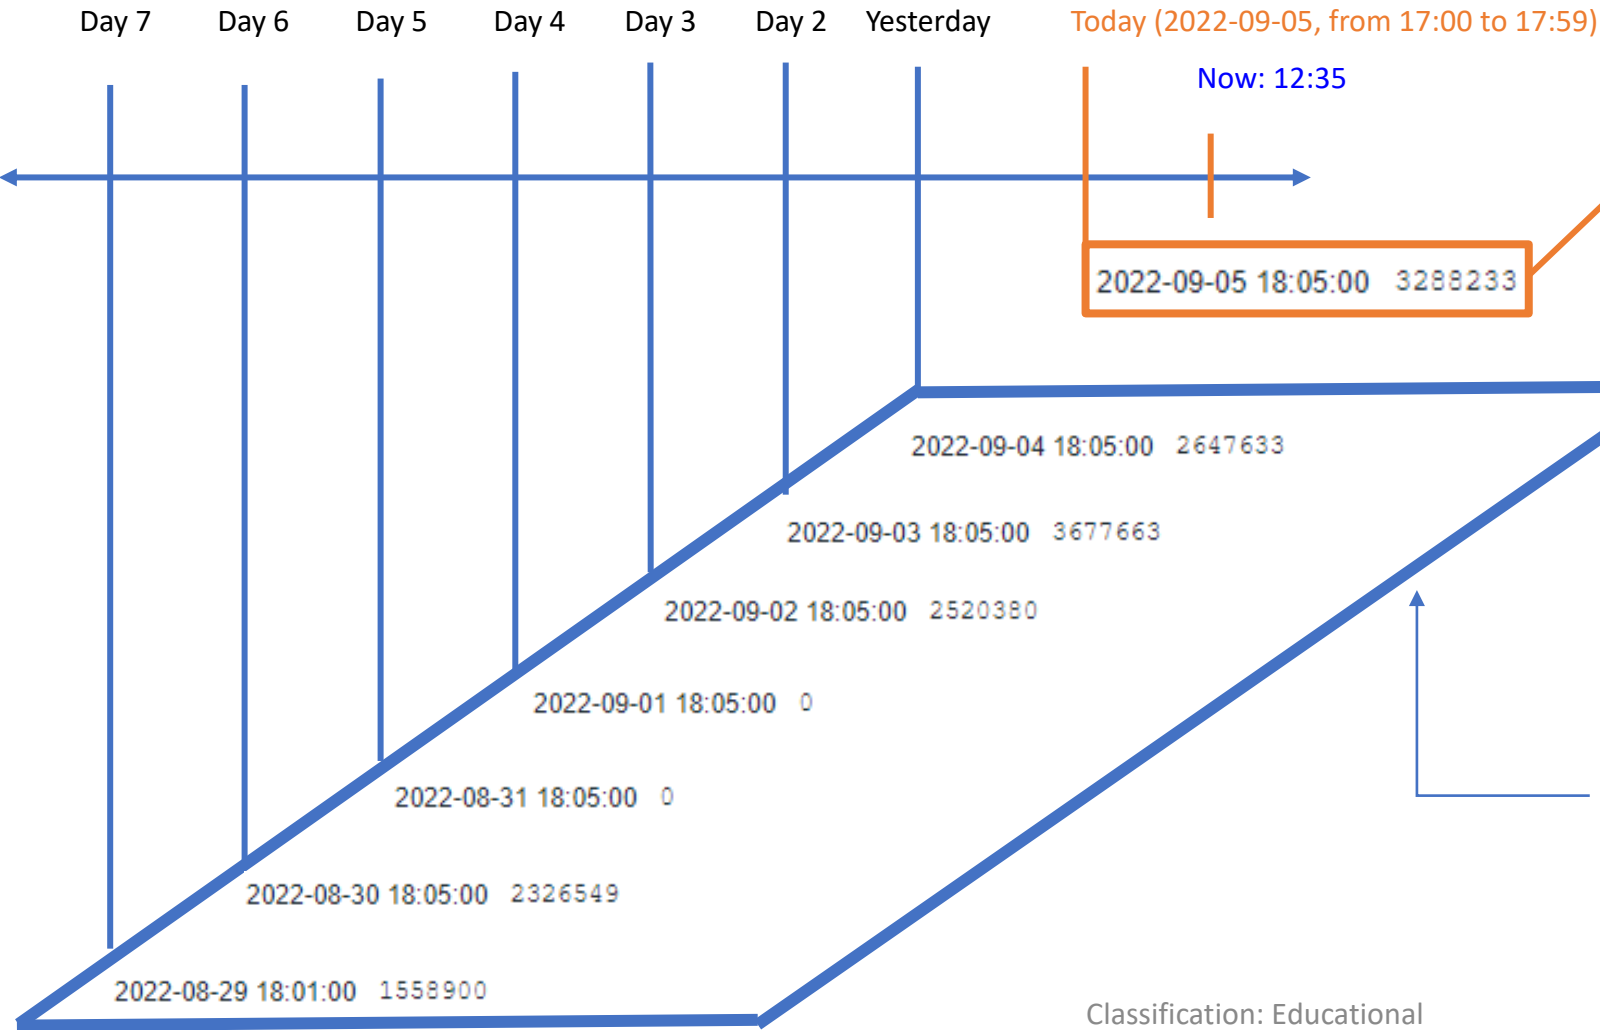
Formula: `baselinedev(//total.bytes,1h:now/h,"d",7)`

STANDARD DEVIATION

Previous hour, 7 days, same timeframe

0.99 STD DEV ↑

2022-09-05 18:01:00



`trendavg(//total.bytes,1h:now/h)`


Standard Deviation

STANDARD DEVIATION

Previous hour, 7 days, same timeframe

0.99 STD DEV ↑

2022-09-05 18:01:00






7

Anomaly Rate

Anomaly Rate



Formula: `trendstl(//trendavg.1,24h:now/d,24h,2h,2,"stddevpop")`



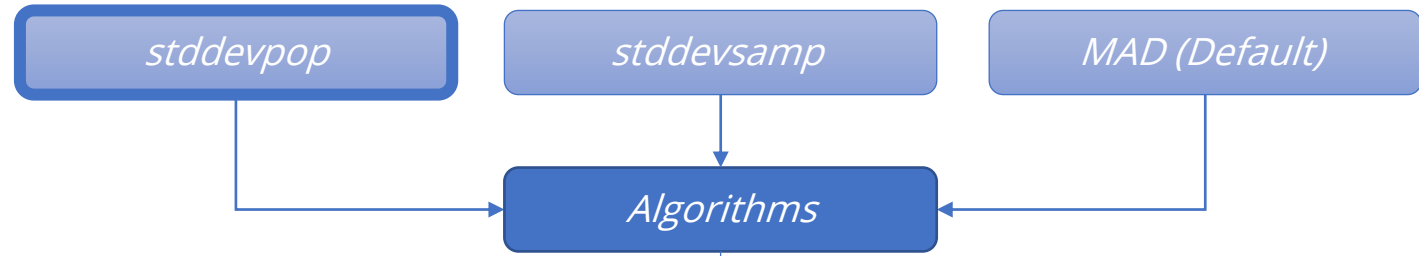
Anomaly Rate

ANOMALY RATE

Based on 24 hours, finding anomalies at last...

0.13 RATE

2022-09-05 18:05:00



Number of deviations to considering

Formula: `trendstl(//trendavg.1,24h:now/d,24h,2h,2,"stddevpop")`


Observability period

Expected periodicity

Looking for anomalies here

In the last 24 hours "Anomaly Rate" is 0.13

Anomaly Rate



Something is happening at about 6 PM every day and needs some investigation!



8

Conclusion



#Considerations

- A short time was used
- You can use graphs to show a Standard Deviation evolution
- In real life you must consider your business rules

Conclusion

- In general, our goal was achieved
 - Improving or generate network infrastructure visibility using some new Zabbix Trend Functions
 - Highlighting Baseline, Standard Deviation and Anomaly Rate ✓
 - Improving or generate a tactical layer monitoring in addition to operational layer monitoring ✓
 - Understanding our network behavior based on operational data received before by Zabbix (using Netflow integration) ✓



Conclusion


- **We have new challenges**
 - we need to create some triggers to be fired when a threshold is reached
 - we need to translate some business rules to Zabbix formulas and try to reflect the company reality
 - if you don't use Zabbix 6.+ yet, make a plan and start as soon as possible!

Conclusion

- **Operational layer monitoring is easy and necessary**
 - if you can monitor the tactical layer, it will be desirable
- **There is no limit to monitor when using Zabbix**
 - if out-of-the-box features are not enough, just extend Zabbix
- **More questions instead answers**
 - Before: Poor visibility
 - Now: Operational and tactical layer monitoring
 - Next: Investigation

Thank you!

 [instagram.com/uniredeinfo/](https://www.instagram.com/uniredeinfo/)

 [linkedin.com/company/unirede/](https://www.linkedin.com/company/unirede/)

 [facebook.com/unirede.net](https://www.facebook.com/unirede.net)

 [youtube.com/uniredeinfo](https://www.youtube.com/uniredeinfo)

 twitter.com/unirede

