

# Zabbix meets AI

-

## Leverage machine learning for detecting anomalies



IntelliTrend GmbH

 [www.intellitrend.de](http://www.intellitrend.de)



**ZABBIX**  
PREMIUM PARTNER



Contact: Wolfgang Alper

[wolfgang.alper@intellitrend.de](mailto:wolfgang.alper@intellitrend.de)

# Zabbix meets AI

First of all  
This talk is not about ChatGPT ;-)

Nevertheless – ChatGPT is a great product

# Zabbix meets AI

Time travel

-

Back to the year 2020

# Zabbix meets AI

It all started in May 2020

- One of our customers operates webshops
- These webshops are monitored by Zabbix
- On May 19th 2020 our customer found out that one of the webshops had been hacked
- Zabbix did not alert the hack
- The customer was not happy about this...



# Zabbix meets AI

## Customer:

"Why did the network monitoring not alert that the webshop was hacked?"

## We:

"Because Zabbix was not configured for this specific scenario."



# Zabbix meets AI

## Customer:

"If a doctor takes a blood screen, he can determine whether something is wrong, why can't Zabbix network monitoring do that?"

## We:


"Again, because Zabbix was not configured for this specific scenario."

"But wait, **maybe** there is a way for Zabbix to learn about the special features of a particular environment and then detect that something is `different`."




# Zabbix meets AI

## The idea of anomaly detection




# Zabbix meets AI – The idea of anomaly detection


## Network In




## Network Out




## CPU Load




## Memory Used



## Processes Count



## Free Disk Space



Simple “Blood screen” of a system



# Zabbix meets AI – The idea of anomaly detection

Place individual metrics in a system-specific **context** to detect **anomalies**

- Instead of having **one trigger** that looks at **one metric** at a time, have a system that looks at **multiple metrics** at once
- Instead of using **simple** trigger functions or **simple** aggregate functions, look at the data **as a whole** over a period of time
- Instead of using triggers with **static conditions**, let the system “learn” the specific characteristics over time with **variable conditions**



# Zabbix meets AI – The idea of anomaly detection


## Examples

- Monitoring server rooms:  
Do not only pay attention to the temperature, but put it in context with the power consumption of the systems, the air conditioning system etc.
- Server utilization:  
Do not only pay attention to the CPU utilization, but put it in a context with the number of users, the memory utilization, the network traffic etc.




# Zabbix meets AI

## Machine learning and neural networks



# Machine learning and neural networks

- **Artificial intelligence** (AI) is the overarching term
- **Machine learning** (ML) is a sub-area of AI
- **Neural network** (NN) is a sub-area of ML and also forms the backbone of deep learning algorithms (sometimes called ANN – Artificial neural network)
- **Deep learning** (DL) is a sub-area of ML where the number of node layers or depth is higher than in a single neural network (sometimes called DNN – Deep neural network)



**Machine learning** is not necessarily based on **NN** or **DL**

# Machine learning and neural networks

Example of trigger function in Zabbix using machine learning:

```
trendstl (/host/key,eval period:time shift,detection period,season,  
<deviations>,<devalg>,<s window>)
```

Returns the rate of anomalies during the detection period

**eval period** - the time period that must be decomposed

**detection period** - the time period before the end of eval period

**season** - the shortest time period where a repeating pattern is expected

**deviations** - the number of deviations to count as anomaly

**devalg** - the deviation algorithm, can be stddevpop, stddevsamp or mad


**s window** - the span of the loss window for seasonal extraction

These functions are designed to work with time series data of a **single metric**,  
therefore there is **no context** to other **related metrics**



# Zabbix meets AI

Simple Data vs. Time series data  
Univariate vs. Multivariate data



# Machine learning and neural networks

Difference between **univariate** and **multivariate** data / **simple** and **time-series** data

- Univariate simple data: **One value**  
Example: Temperature: 23 C
- Univariate time series data: **List of values**  
Example: 23 C, 24 C, 25 C

- Multivariate simple data: **Multiple metrics, one value per metric**  
Example: Temperature: 23 C, Humidity: 56 %, Air Pressure: 998 mbar
- Multivariate time series data: **Multiple metrics, list of value sets**  
Example: 23 C/56 %/998 mbar, 24 C/55 %/1002 mbar, 25 C/54 %/1004 mbar

We want **multivariate time-series data** because we need **contextual relationships** and **temporal context**

# Machine learning and neural networks

## Univariate times-series data

```
[{"temp":23, "time":1723988060}, {"temp":24, "time":1723988120}, {"temp":25, "time":1723988180}]
```

```
[{"apress":998, "time":1723988060}, {"apress":1002, "time":1723988120}, {"apress":1004, "time":1723988180}]
```

```
[{"humidity":56, "time":1723988060}, {"humidity":55, "time":1723988120}, {"humidity":54, "time":1723988180}]
```


## Multivariate times-series data

```
[{"temp":23, "hum":56, "apress":998, "time":1723988060}, {"temp":24, "hum":55, "apress":1002, "time":1723988120}, {"temp":25, "hum":54, "apress":1004, "time":1723988180}]
```





# Zabbix meets AI

## Anatomy of Neural networks



# Machine learning and neural networks

Each line represents a connection to an artificial neuron with individual “Weights”

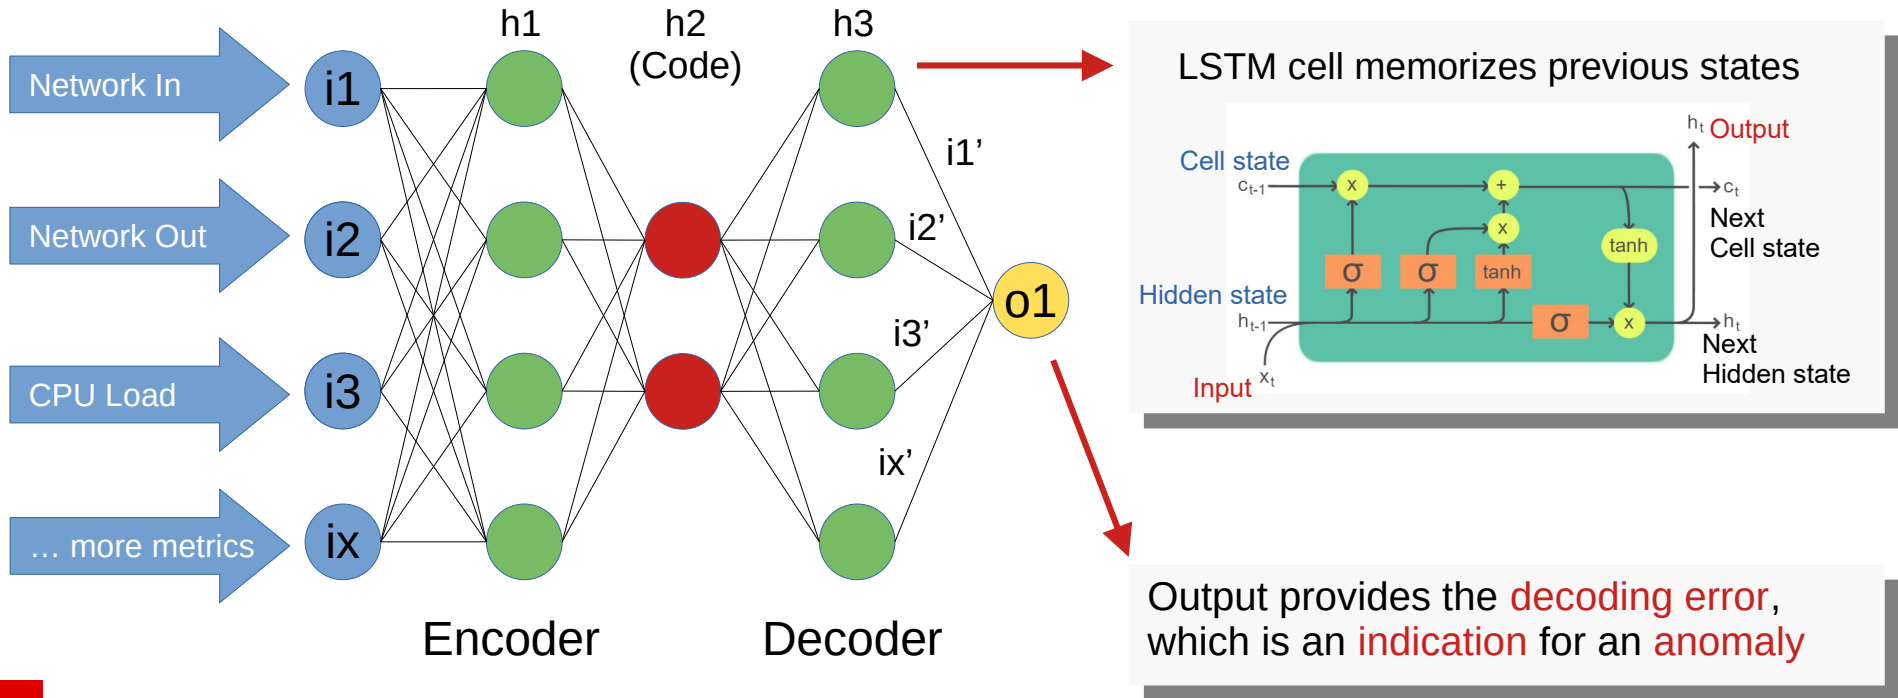


Simple example with 46 trainable weights  
(without biases)  
Each line represents a weight

## Basic design of neural networks

# Deep neural networks - Autoencoder

Deep neural network to detect anomalies in multivariate time-series data using LSTM cells



ZABBIX  
PREMIUM PARTNER

Autoencoders have an **encoder** and **decoder**, acting as a **data compressor**

# Deep neural networks - Autoencoder

With an Autoencoder using Long Short-Term Memory cells (LSTM) there is support for:

- ✓ Multivariate data
- ✓ Times-series data
- ✓ Temporal context
- ✓ Contextual relationships


What about training?

- To train a model, training data is required
- Usually **supervised learning** requires **labeled data**, which is not always easy to get
- However, the **Autoencoder** architecture allows **unsupervised training** with **unlabeled data** which is great for this use-case



# Zabbix meets AI

## Preparing the data for training, testing and inference



# Preparing the data – Training and inference

- **Training:** Train a model with training data so it “learns” features
- **Inference:** Use a trained model with production data

How to prepare the data for training, testing and inference?


- ▶ **Extraction:** Data needs to be extracted from **Zabbix**, using the **Zabbix-API** or **Zabbix database**
- ▶ **Data Enrichment:** To provide additional context, certain metrics, such as timestamps, need to be pre-processed by **creating additional inputs** like hour/minute, day of the week, month of the year
- ▶ **Normalization:** Neural networks prefer a min-max **normalization** of the values in the datasets in the range of **[0, 1]** or **[-1, 1]**
- ▶ **Windowing:** Divide multivariate time-series data into a **sliding window** of data




# Preparing the data – Extraction

Use of **tags** to identify the **items** that should be used for **training** and later for **inference**

<input type="checkbox"/> Host	Name ▲	Last check	Last value	Change	Tags
<input type="checkbox"/> IMS-Smart IntelliTrend...	Air pressure [Pa] ?	25s	100.81 KPa	+2.6562 Pa	df-ai: df-ai Application: Sensor d...
<input type="checkbox"/> IMS-Smart IntelliTrend...	Humidity [%] ?	25s	39.7236 %	-3.2168 %	df-ai: df-ai Application: Sensor d...
<input type="checkbox"/> IMS-Smart IntelliTrend...	Temperature [C] ?	25s	26.49 C		df-ai: df-ai Application: Sensor d...
<input type="checkbox"/> IMS-Smart IntelliTrend...	Total volatile organic compound TVOC [ppb] ?	25s	39 ppb	+6 ppb	df-ai: df-ai Application: Sensor d...



```
{
  "jsonrpc": "2.0",
  "method": "history.get",
  "params": {
    "output": "extend",
    "history": 0,
    "itemids": "912257",
    "sortfield": "clock",
    "sortorder": "DESC",
    "limit": 10
  },
  "id": 1
}
```



```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "itemid": "912257",
      "clock": "1724588945",
      "value": "26.31999969",
      "ns": "227934509"
    },
    {
      "itemid": "912257",
      "clock": "1724588885",
      "value": "26.32999992",
      "ns": "87916196"
    }
  ]
}
```


Using the **Zabbix-API** to get **univariate time-series** data for a given item

# Preparing the data – Data Enrichment

```
{  
  "itemid": "912257",  
  "clock": "1724588945",  
  "value": "26.31999969",  
  "ns": "227934509"  
}
```

- Timestamps **always increase** their **value** per second (monotonic increasing function)
- They need to be **transformed** into something that **changes periodically** over time
- Both **sine** and **cosine** are required to **differentiate** between positions with the same **values**

```
# Timestamp conversion to single features  
day = 24*60*60  
year = (365.2425)*day  
ts['day_sin'] = sin(timestamp * (2*pi/day))  
ts['day_cos'] = cos(timestamp * (2*pi/day))  
ts['year_sin'] = sin(timestamp * (2*pi/year))  
ts['year_cos'] = cos(timestamp * (2*pi/year))
```

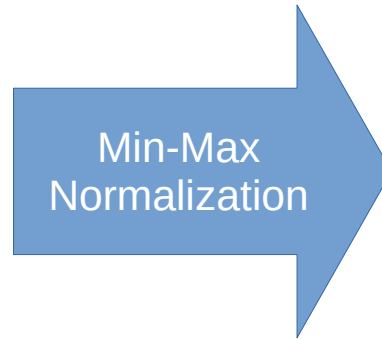


Adding features from timestamps so that the network **learns** the **periodicity** in the data



# Preparing the data – Normalization

```
{
  "jsonrpc": "2.0",
  "result": [
    {
      "itemid": "912257",
      "clock": "1724588945",
      "value": "26.31999969",
      "ns": "227934509"
    },
    {
      "itemid": "912257",
      "clock": "1724588885",
      "value": "26.32999992",
      "ns": "87916196"
    }
  ],
}
```



```
[
  {
    "itemid": "912257",
    "clock": "1724588945",
    "value": "0.7576",
  },
  {
    "itemid": "912257",
    "clock": "1724588885",
    "value": "0.8245",
  },
]
```

## Challenges:

- With large datasets, min-max normalization can take some time and resources
- Detecting outliers is difficult, because they could be valid data points in the dataset

# Preparing the data – Windowing

```
...  
...  
{"temp":0.75, "hum":0.56, "apress":0.79, "ts_ds":0.50, ts_dc:0.34},  
{"temp":0.82, "hum":0.55, "apress":0.80, "ts_ds":0.64, ts_dc:0.77},  
{"temp":0.84, "hum":0.54, "apress":0.86, "ts_ds":0.77, ts_dc:0.64},  
{"temp":0.83, "hum":0.53, "apress":0.88, "ts_ds":0.87, ts_dc:0.50},  
...  
...
```

```
...  
...  
{"temp":0.75, "hum":0.56, "apress":0.79, "ts_ds":0.50, ts_dc:0.34},  
{"temp":0.82, "hum":0.55, "apress":0.80, "ts_ds":0.64, ts_dc:0.77},  
{"temp":0.84, "hum":0.54, "apress":0.86, "ts_ds":0.77, ts_dc:0.64},  
{"temp":0.83, "hum":0.53, "apress":0.88, "ts_ds":0.87, ts_dc:0.50},  
...  
...
```

0.75	0.82	0.84
0.56	0.55	0.54
0.79	0.80	0.86
0.50	0.64	0.77
0.34	0.77	0.64

0.82	0.84	0.83
0.55	0.54	0.53
0.80	0.86	0.88
0.64	0.77	0.87
0.77	0.64	0.50

Prepare an overlapping window with multivariate time-series data with a constant size



ZABBIX  
PREMIUM PARTNER


# Preparing the data – Final windows of data

Batch of windowed data

0.82	0.84	0.83	0.75	0.82	0.84
0.55	0.54	0.53	0.56	0.55	0.54
0.80	0.86	0.88	0.79	0.80	0.86
0.64	0.77	0.87	0.50	0.64	0.77
0.77	0.64	0.50	0.34	0.77	0.64

Windowed data

Windowed data




Transfer of **batches** of **windowed** time series **data** to the deep learning model



ZABBIX  
PREMIUM PARTNER

# Zabbix meets AI

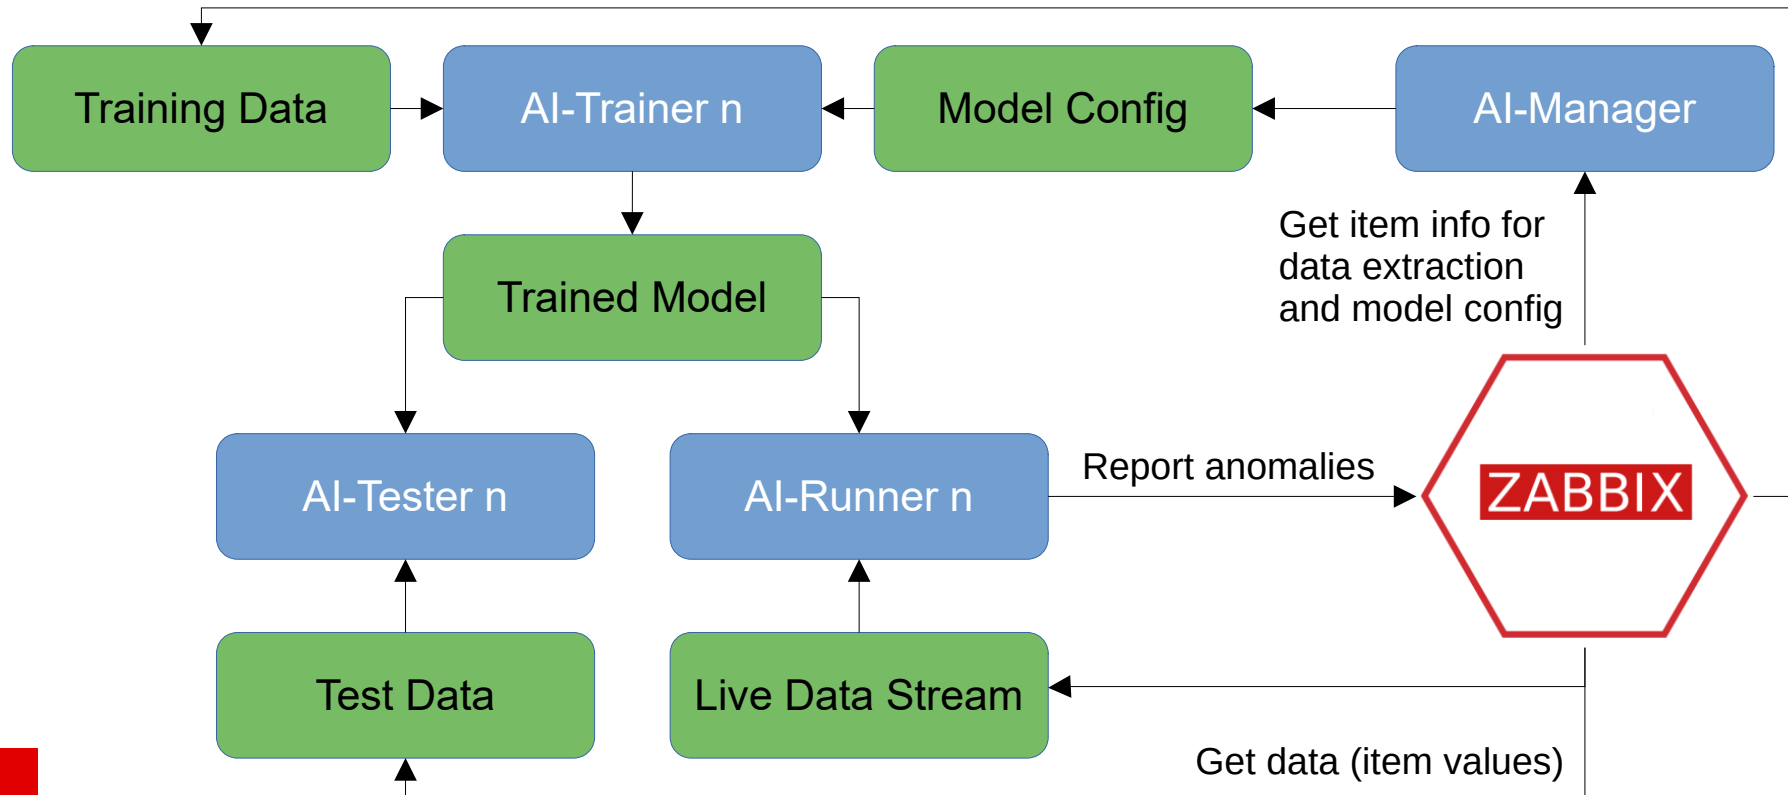
Building blocks  
to integrate **AI** into **Zabbix**



# Zabbix meets AI – Building blocks

- **Create a model configuration:** Select **model type** and define model **hyperparameters**
- **Extract training and test data:** Select data from **Zabbix history/trends** for specific items
- **Create and train the model:** Model **inputs** will be **adjusted** to **match items** from Zabbix
- **Test the model:** Test the model based on **existing historical data** from Zabbix
- **Use the model with live data:** Continuously **extract** and **prepare** the **values** for selected items and pass them to the model for **evaluation**
- **Pass the result of the evaluation back to Zabbix:** The **anomaly probability** should be send to **Zabbix** as an **item**, so mechanisms like trigger, actions, alerts etc. could be used
- **Support multiple models:** To detect anomalies across multiple **different systems**, **multiple models** must be usable at the **same time**

# Zabbix meets AI – Building blocks



Number of n AI Trainer/Tester/Runner need to scale horizontally like Zabbix Proxies


# Zabbix meets AI

Do you remember  
the  
time travel?


# Zabbix meets AI – Data from the time travel

Some metrics of the system in question ...


AI-Webshop: Context switches




AI-Webshop: CPU usage all




AI-Webshop: Interrupts per second



AI-Webshop: Network incoming bandwidth



AI-Webshop: Memory free



It is **difficult** for a human being to **recognize anomalies** in this type of data




**ZABBIX**  
PREMIUM PARTNER



# Zabbix meets AI

## Creating the model



# Zabbix meets AI – Creating a model

Item selection for the AI model is based on **tags** and **hosts** in Zabbix

<input type="checkbox"/>	Name ▲	Triggers	Key	Interval	History	Trends	Type	Status	Tags	Info
<input type="checkbox"/>	... <a href="#">Template AI-Webshop: Context switches</a>		contextSwitches	90d	365d		Zabbix Agent	<a href="#">Enabled</a>	df-ai	
<input type="checkbox"/>	... <a href="#">Template AI-Webshop: CPU usage all</a>		cpuUsageAll	90d	365d		Zabbix Agent	<a href="#">Enabled</a>	df-ai	
<input type="checkbox"/>	... <a href="#">Template AI-Webshop: Interrupts per second</a>		interruptsPerSec	90d	365d		Zabbix Agent	<a href="#">Enabled</a>	df-ai	
<input type="checkbox"/>	... <a href="#">Template AI-Webshop: Memory free</a>		memoryPFree	90d	365d		Zabbix Agent	<a href="#">Enabled</a>		
<input type="checkbox"/>	... <a href="#">Model One: Reconstruction Error [MAE]</a>		m1.loss.mae	90d	365d		Zabbix trapper	<a href="#">Enabled</a>		
<input type="checkbox"/>	... <a href="#">Model One: Reconstruction Error [MAPE]</a>		m1.loss.mape	90d	365d		Zabbix trapper	<a href="#">Enabled</a>		
<input type="checkbox"/>	... <a href="#">Model One: Reconstruction Error [MSE]</a>		m1.loss.mse	90d	365d		Zabbix trapper	<a href="#">Enabled</a>		
<input type="checkbox"/>	... <a href="#">Model One: Reconstruction Error [MSLE]</a>		m1.loss.msle	90d	365d		Zabbix trapper	<a href="#">Enabled</a>		
<input type="checkbox"/>	... <a href="#">Model Two: Reconstruction Error [MAE]</a>		m2.loss.mae	90d	365d		Zabbix trapper	<a href="#">Enabled</a>		
<input type="checkbox"/>	... <a href="#">Model Two: Reconstruction Error [MAPE]</a>		m2.loss.mape	90d	365d		Zabbix trapper	<a href="#">Enabled</a>		
<input type="checkbox"/>	... <a href="#">Model Two: Reconstruction Error [MSE]</a>		m2.loss.mse	90d	365d		Zabbix trapper	<a href="#">Enabled</a>		
<input type="checkbox"/>	... <a href="#">Model Two: Reconstruction Error [MSLE]</a>		m2.loss.msle	90d	365d		Zabbix trapper	<a href="#">Enabled</a>		
<input type="checkbox"/>	... <a href="#">Template AI-Webshop: Network incoming bandwidth</a>		ifInBps	90d	365d		Zabbix Agent	<a href="#">Enabled</a>	df-ai	
<input type="checkbox"/>	... <a href="#">Template AI-Webshop: Process count</a>		processCount	90d	365d		Zabbix Agent	<a href="#">Enabled</a>	df-ai	

Reconstruction error is sent by the AI-Runner as an **indication** of an **anomaly** per model

# Zabbix meets AI – Creating a model

Model configuration with preferences using DataForge with Zabbix

**DATAFORGE**

< Model config details

**Architecture**

Architecture LSTM Autoencoder ▾

Latent space size factor  
1

**Options**

Window size  
8

Time step  
60

Epochs  
15

Use outlier clipping

Outlier threshold  
3

Use dropout

Activation ReLU ▾

Loss MSE ▾

Optimizer Adam ▾

Scaler MIN\_MAX ▾

Model depth

Model width  
4

Model slope  
0.8


## Global model parameters:

- Architecture
- Windows Size
- Epochs
- Outlier clipping
- ...
- other model hyperparameters

**Expected weights count:**  
35.000 – 40.000 Parameter

# Zabbix meets AI

## Training and testing the model



# Zabbix meets AI – Training a model

## Create training or test data

The screenshot shows the 'DATAFORGE' interface with a sidebar on the left containing navigation options like Dashboard, Zabbix Client, Open problems, Closed problems, Triggers, Host groups, Hosts, Items, Graphs, Services, Favorites, Self Provisioning, Maintenances, Alerts, AI, Training Data, and Models. The main content area is titled '< Training data details' and contains the following sections:

- General**: Name 'Webshop Attack Dataset', Description 'Datasets related to the attack on the customers Webshop'.
- Training data configuration**: A red-bordered box containing:
  - Hosts: 'Webshop X' (with a close button).
  - Schedule mode: 'Manual' (dropdown).
  - From: 'Aug 11, 2024 24:00'.
  - To: 'Aug 28, 2024 24:00'.
  - Storage options: Storage duration '90 Days 0 Hours'.
- Training data**: A message stating 'There have been no training data extracted for this configuration yet. [Click here to start an training data extraction job.](#)' with an 'EXTRACT DATA' button.

### Training and test datasets

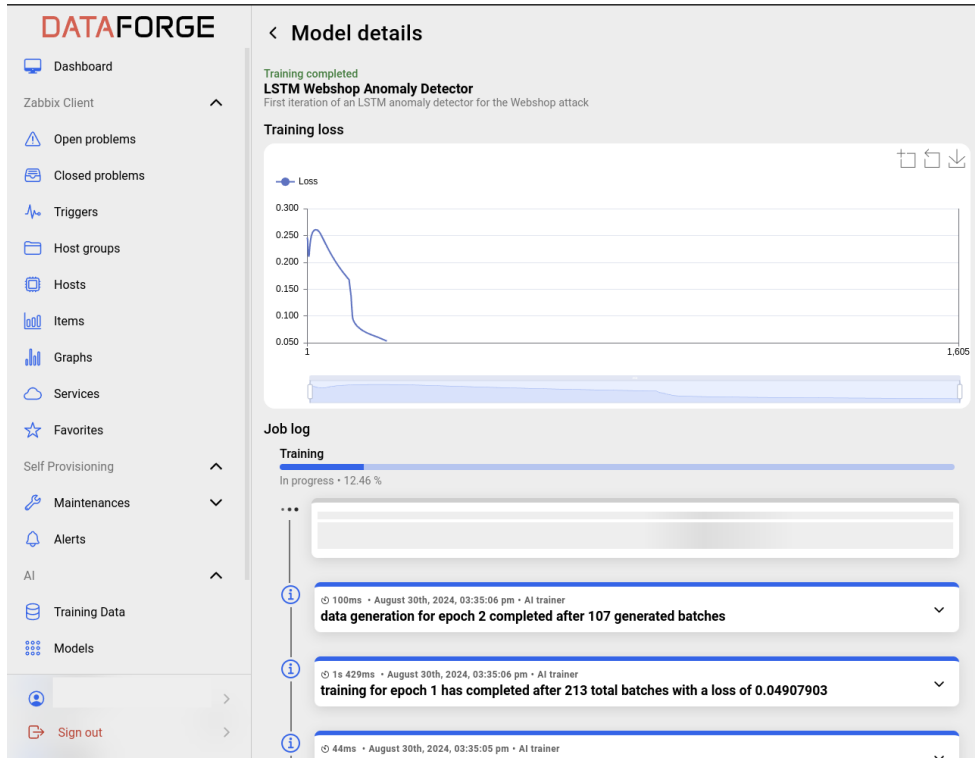
- **One** or **multiple** hosts
- When to create/update data
- Time period
- Data storage period

The data is extracted from **items** that have a corresponding **tag** in Zabbix



# Zabbix meets AI – Training a model

Training with created training data set and model configuration



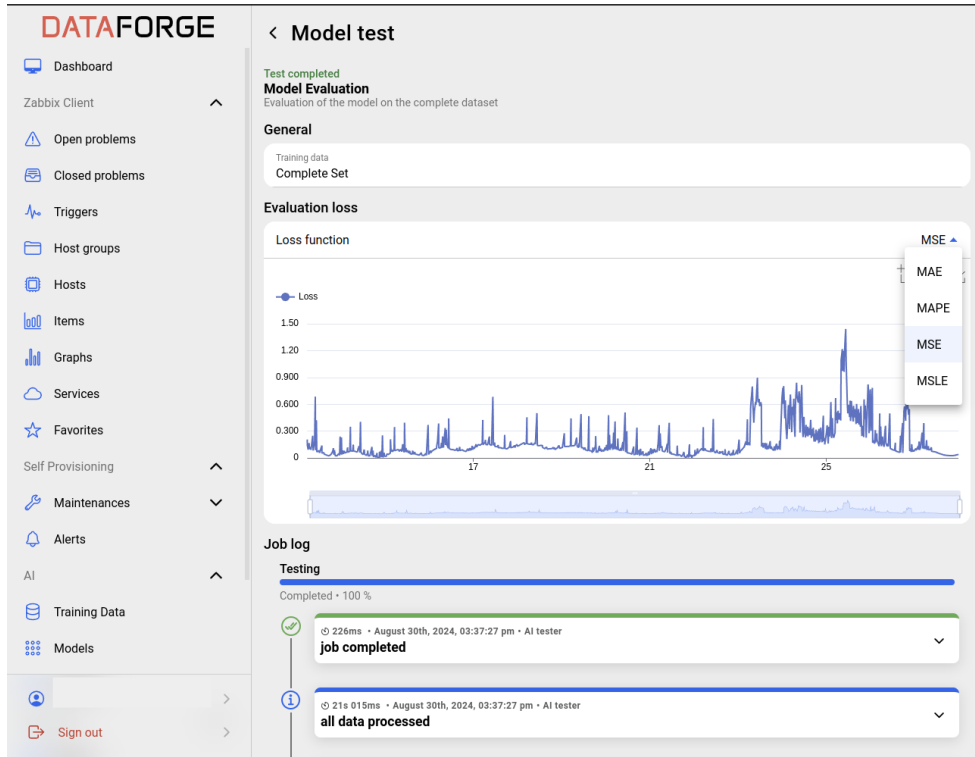
- Training loss graph and progress is streamed in realtime
- Loss graphs gives an estimation of the quality of the model



**ZABBIX**  
PREMIUM PARTNER

# Zabbix meets AI – Test a model

Testing a model is similar to training, except that it uses a different dataset




- **Loss graphs** shows the **performance** on historical data selected from Zabbix
- Different **loss functions** are available to calculate the value of the **anomaly**
- The **output** of these functions is available in **Zabbix** as an item and can be used with a **trigger**
- **Testing** a model with real data is a good way to test the **performance** before going live



# Zabbix meets AI

## Using the model - Inference





# Zabbix meets AI – Using a model in production

**DATAFORGE**

**< Inference details**

**General**

Name  
Webshop Anomaly Detector Inference

Description  
Runs the LSTM based anomaly detection model on live data from the Webshop

Switch to latest model

Host name  
Webshop

Model  
LSTM Webshop Anomaly Detector (#2) ▾

AI runner  
DataForge AI Runner ▾

Loss functions  
MSE, MAE, MAPE, MSLE ▾


- Set a **name** to **identify** the model
- Select whether the model **update** should be provisioned **automatically**
- Select the **host(s)** the model should be used with
- Select the **AI Runner** instance to be used
- Select the **loss functions** that are to be used to calculate the **anomalies** and sent to **Zabbix**




# Zabbix meets AI – Using a model in production

Model **output** with various loss functions as items in **Zabbix** based on the **webshop** data  
**Reconstruction errors** represent the probabilities for **anomalies** based on their **loss function**


AI-Webshop: Reconstruction Error [MAE]




AI-Webshop: Reconstruction Error [MSE]



AI-Webshop: Reconstruction Error [MAPE]



AI-Webshop: Reconstruction Error [MSLE]



Example trigger expressions in Zabbix to alert an anomaly:

```
avg (/AI-Webshop/1.loss.mae,600) >={ $MIN_ANOMALY_MAE }  
last (/AI-Webshop/m1.loss.mape) >={ $MIN_ANOMALY_MAPE }  
min (/AI-Webshop/1.loss.mse,300) >={ $MIN_ANOMALY_MSE }  
min (/AI-Webshop/1.loss.msle,300) >={ $MIN_ANOMALY_MSLE }
```

Available loss functions to indicate an anomaly:

MAE	- Mean Absolute Error
MAPE	- Mean Absolute Percentage Error
MSE	- Mean Squared Error
MSLE	- Mean Squared Logarithmic Error



**ZABBIX**  
PREMIUM PARTNER

# Zabbix meets AI

## Summary



# Zabbix meets AI – Summary

- ▶ Deep neural networks are a great way to detect **anomalies** in **time-series data** due to their ability to maintain **temporal context**
- ▶ They excel at analyzing **multiple metrics** simultaneously, enabling the creation of **contextual relationships** between data from **various sources**
- ▶ **Zabbix's flexible architecture** and **robust API** make it an ideal platform for integrating AI
- ▶ By combining metrics from **multiple hosts** into a **single model**, anomaly detection becomes even more powerful, particularly when monitoring **entire services**
- ▶ The **results** of these **AI-driven detections** can be seamlessly fed back into **Zabbix** as **items**, allowing to leverage any of Zabbix's **trigger functions** to generate **alerts**
- ▶ **Zabbix's versatility in metric collection** extends beyond just computer and network metrics, making AI-based **anomaly detection** applicable to a **wide range of domains**

Whether it's monitoring sales and visitor traffic for an online shop, or tracking the power generation of a solar park, **any metric** that can be collected **in Zabbix** can be utilized for **anomaly detection**



# Zabbix meets AI

-

## Leverage machine learning for detecting anomalies



IntelliTrend GmbH

# Thank You!



Contact: Wolfgang Alper

[wolfgang.alper@intellitrend.de](mailto:wolfgang.alper@intellitrend.de)

 [www.intellitrend.de](http://www.intellitrend.de)



**ZABBIX**  
PREMIUM PARTNER