



OCTOBER 8 • 10, 2025
RIGA • LATVIA

Multi-Tenancy Zabbix Metrics Processing and Streaming

with Kafka, KSQL, and Grafana



About this talk

Real-world use case

The story of a solution we built and keep developing for one of our customers that had a special set of challenges and requirements.

Example of integration

Especially in larger environments, even the best of all monitoring utilities have to play the role of being just a little part of something.

Integration

A showcase of how data-driven observability can be built as built as a modern, microservice-based, extendible platform. platform.

In...Q...what?

Inaugurating tomorrow's IT today

Trusted Zabbix Partner & Trainer

Deep expertise from countless projects — we know Zabbix inside-out and help customers get the most out of it.

Observability for Modern IT

Modern applications have observability built in. It's the key to assuring quality, guiding development, and keeping full control of infrastructure and apps.

Automation & Cloud-Native Practices

No manual work, no fragile setups – everything must be automatable and declarative. We design microservices and integrate them with leading open-source tech.

We're a specialized IT consulting company focused on IT Automation, Monitoring, Observability and Container Platforms.

We help our customers adopt and develop custom solutions, built on open-source, cloud native technologies.



Christian Anton

Owner | Consultant | Trainer

ZABBIX
CERTIFIED PARTNER

ZABBIX
CERTIFIED TRAINER

The Customer and Their Requirements

Health Care Organization

Consumes IT services from a service provider
Also runs some of their own infrastructure

Metrics View

Organization-specific metrics needed to be made available to the customer

Extendable and Multi-tenant

The solution should be capable of onboarding additional service provider's monitoring as customer-owned equipment

Shared Service Zabbix

Zabbix is used by the Service Provider to monitor all kinds of IT equipment

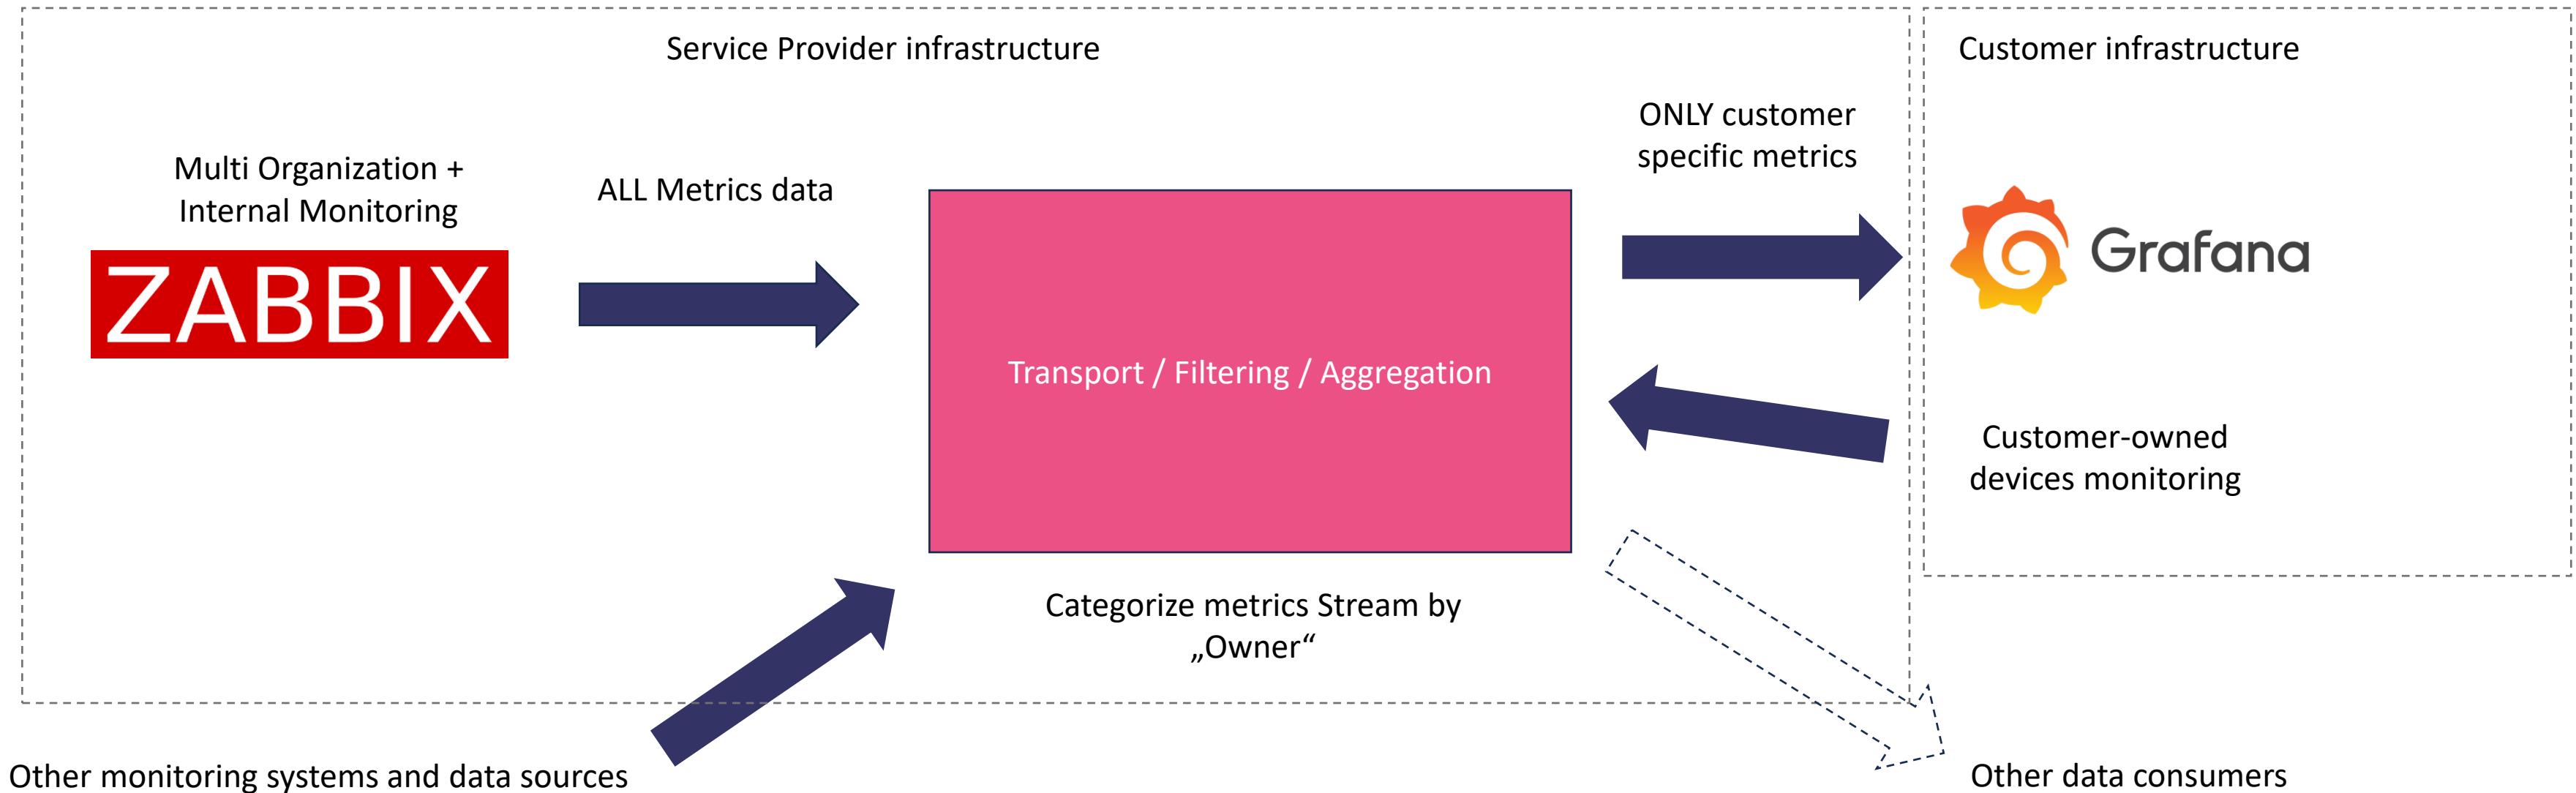
Visualization layer

Grafana was chosen as the desired platform for visualization and dashboards

Platform of choice

Stack should run in Red Hat Openshift, automated with ArgoCD

Basic Architecture



Turning the Idea into Architecture



Stream Processing

Apache Kafka – a robust and high performing data streaming platform with a large ecosystem



Target Format

Use Prometheus metrics format at the destination end - best to consume for Grafana



STRIMZI

Kafka Deployment

Strimzi operator offers fully declarative management of Kafka Instances in Kubernetes + Openshift



Storage

Thanos is a high-performance long-term metrics store – S3 buckets downsampling, housekeeping

Getting Data into the Stack

Time-Series Data

High Volume (~6k NVPS)

Timestamp + Item-ID + Value

Datadog Vector

Future: Native HTTP Exporter + Kafka Bridge

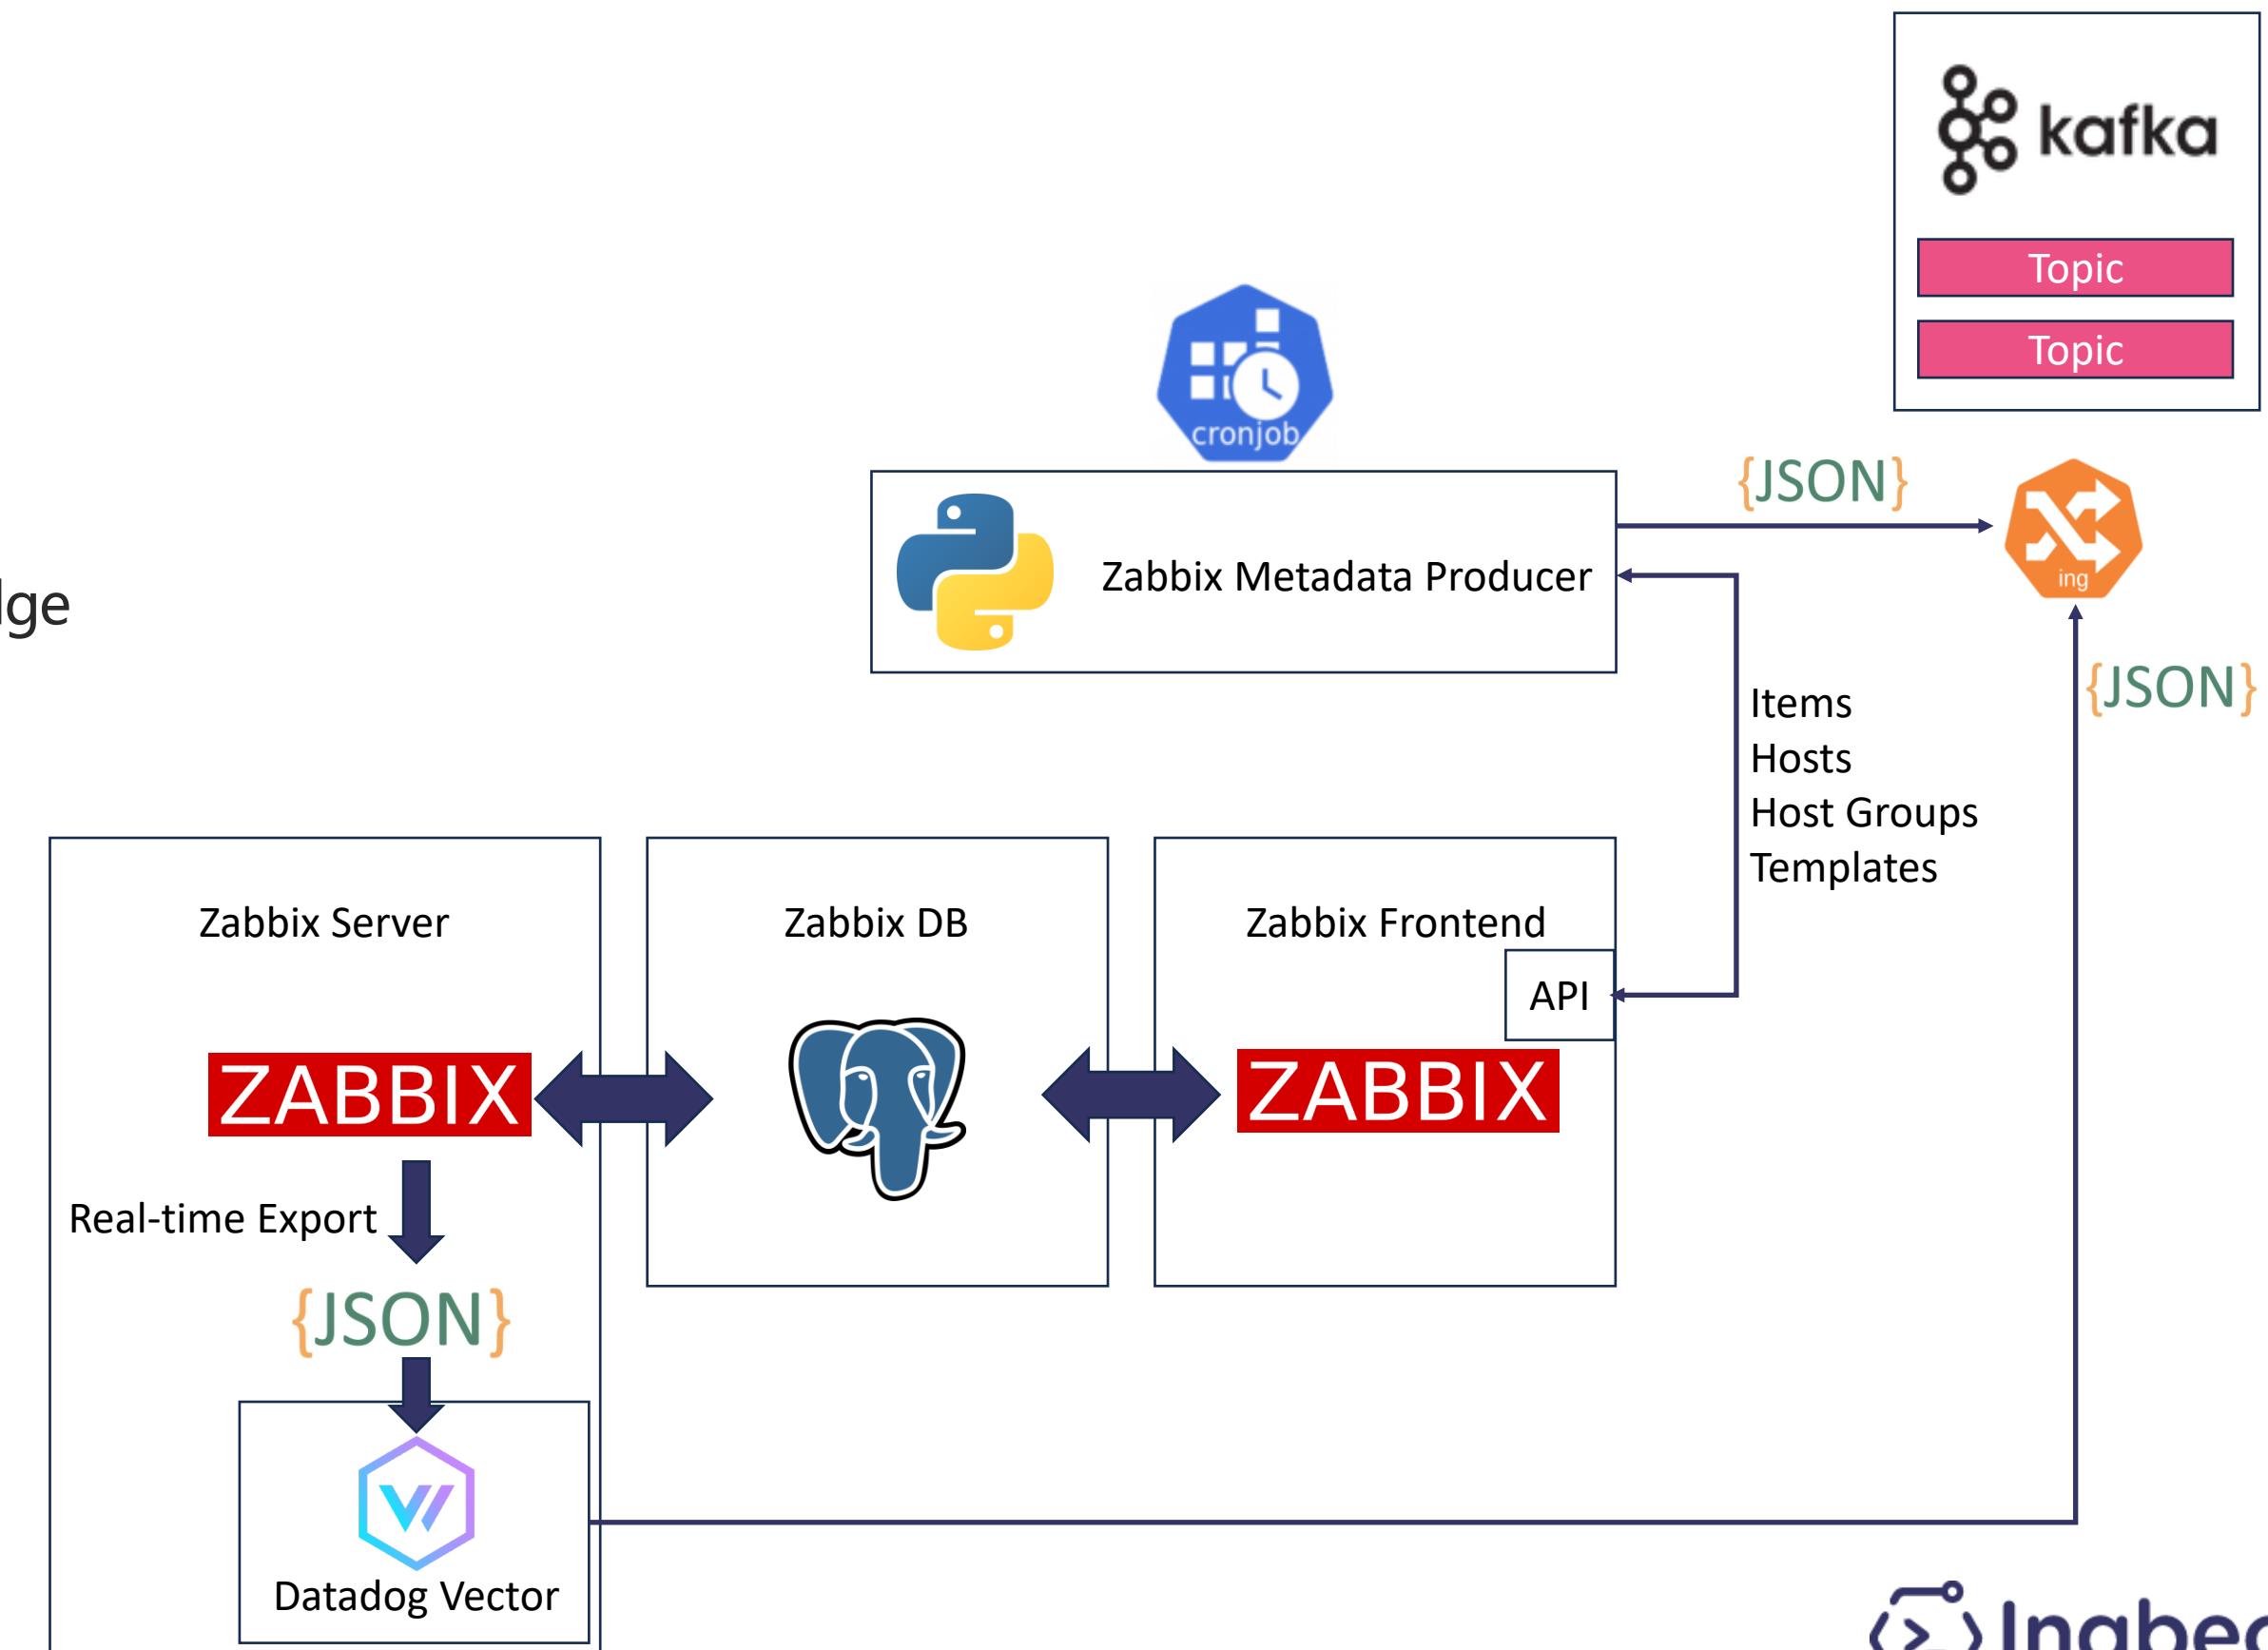
Metadata Producer

Python program, runs as Cron job

Mapping of Item-ID with Item Name, Key, Host Name, Group, Template, Tags, ...

Categorization Ruleset

- Tag selection
- Regex on Item Keys, Names
- Host Names and Host Groups



Overview **Messages** Consumers Settings Statistics

Seek Type Partitions Key Serde Value

Offset Offset All items are selected. String S

Search: 16017036 + Add Filters

	Offset	Partition	Timestamp	Key	Preview
+	54214519261	2	4.10.2025, 12:54:09	16017036	
-	54214407962	2	4.10.2025, 12:53:09	16017036	

Key Value Headers

```
{  
    "clock": 1759575189,  
    "itemid": 16017036,  
    "ts": "2025-10-04T10:53:09Z",  
    "value": 10.796453  
}
```

Item Value Message

Item Tags 1 Preprocessing 1

Parent items Linux by Zabbix agent

* Name CPU utilization

Type Dependent item

* Key system.cpu.util

Type of information Numeric (float)

* Master item Test Host: CPU idle time X

Units %

* History Do not store Store up to 31d

* Trends Do not store Store up to 365d

Value mapping Select

Populates host inventory field -None-

Description CPU utilization expressed in %.

Enabled

Latest data

Update **Clone** **Execute now** **Test** **Clear history and trends**

```

501 - name: Test Host
502   attribute_name: _customer
503   attribute_value: 
504   filters:
505     - item_attr: host_host
506       regex: testhost
507     - item_attr: hostgroups
508       regex: testgroup
509     - item_attr: key
510       regex: ^icmpping|system.cpu.util.*|agent.ping|vm.memory.utilization|vm.memory.si

```

The screenshot shows the Zabbix Item configuration interface. On the left, the 'Item' configuration page is displayed with fields for Name, Type, Key, and Master item. It also includes sections for History, Trends, Value mapping, and Description. Below the configuration are buttons for Update, Clone, Execute now, Test, and Clear history and trends. On the right, a table lists three data entries from the history and trends section. The table has columns for '+' (add), itemid, hostid, delay, key, name, type, value_type, value_type_num, master_itemid, flags, history, templateid, status, trends, units, valuemapid, state, error, params, and description. Below the table is a JSON preview of the item configuration object, which includes all the fields from the configuration page and the assignment filter (host/host: testhost).

	itemid	hostid	delay	key	name	type	value_type	value_type_num	master_itemid	flags	history	templateid	status	trends	units	valuemapid	state	error	params	description	host_host	host_name	hostgroups	host_templates	lastclock	tags	name_resolved	_customer
+	103662023	2	4.10.2025, 12:35:32	system.cpu.util	CPU utilization	Dependent item	numeric float	0	16017009	plain	31d	15501385	enabled	365d	%	0	normal			CPU utilization expressed in %.	testhost	Test Host	Linux, Linux VM, testgroup, zx - Zabbix	Linux by Zabbix agent, Zabbix server health	1759574709	component: cpu	CPU utilization	
+	103662240	2	4.10.2025, 12:42:28	system.cpu.util	CPU utilization	Dependent item	numeric float	0	16017009	plain	365d	15501385	enabled	365d	%	0	normal			CPU utilization expressed in %.	testhost	Test Host	Linux, Linux VM, testgroup, zx - Zabbix	Linux by Zabbix agent, Zabbix server health	1759574709	component: cpu	CPU utilization	
-	103662457	2	4.10.2025, 12:45:31	system.cpu.util	CPU utilization	Dependent item	numeric float	0	16017009	plain	365d	15501385	enabled	365d	%	0	normal			CPU utilization expressed in %.	testhost	Test Host	Linux, Linux VM, testgroup, zx - Zabbix	Linux by Zabbix agent, Zabbix server health	1759574709	component: cpu	CPU utilization	

```
{
  "itemid": "16017036",
  "hostid": "29931",
  "delay": "0",
  "key": "system.cpu.util",
  "name": "CPU utilization",
  "type": "Dependent item",
  "value_type": "numeric float",
  "value_type_num": "0",
  "master_itemid": "16017009",
  "flags": "plain",
  "history": "31d",
  "templateid": "15501385",
  "status": "enabled",
  "trends": "365d",
  "units": "%",
  "valuemapid": "0",
  "state": "normal",
  "error": "",
  "params": "",
  "description": "CPU utilization expressed in %.",
  "host_host": "testhost",
  "host_name": "Test Host",
  "hostgroups": "Linux, Linux VM, testgroup, zx - Zabbix",
  "host_templates": "Linux by Zabbix agent, Zabbix server health",
  "lastclock": "1759574709",
  "tags": "component: cpu",
  "name_resolved": "CPU utilization",
  "_customer": ""
}
```

Item configuration, assignment filter, metadata message

Processing Data

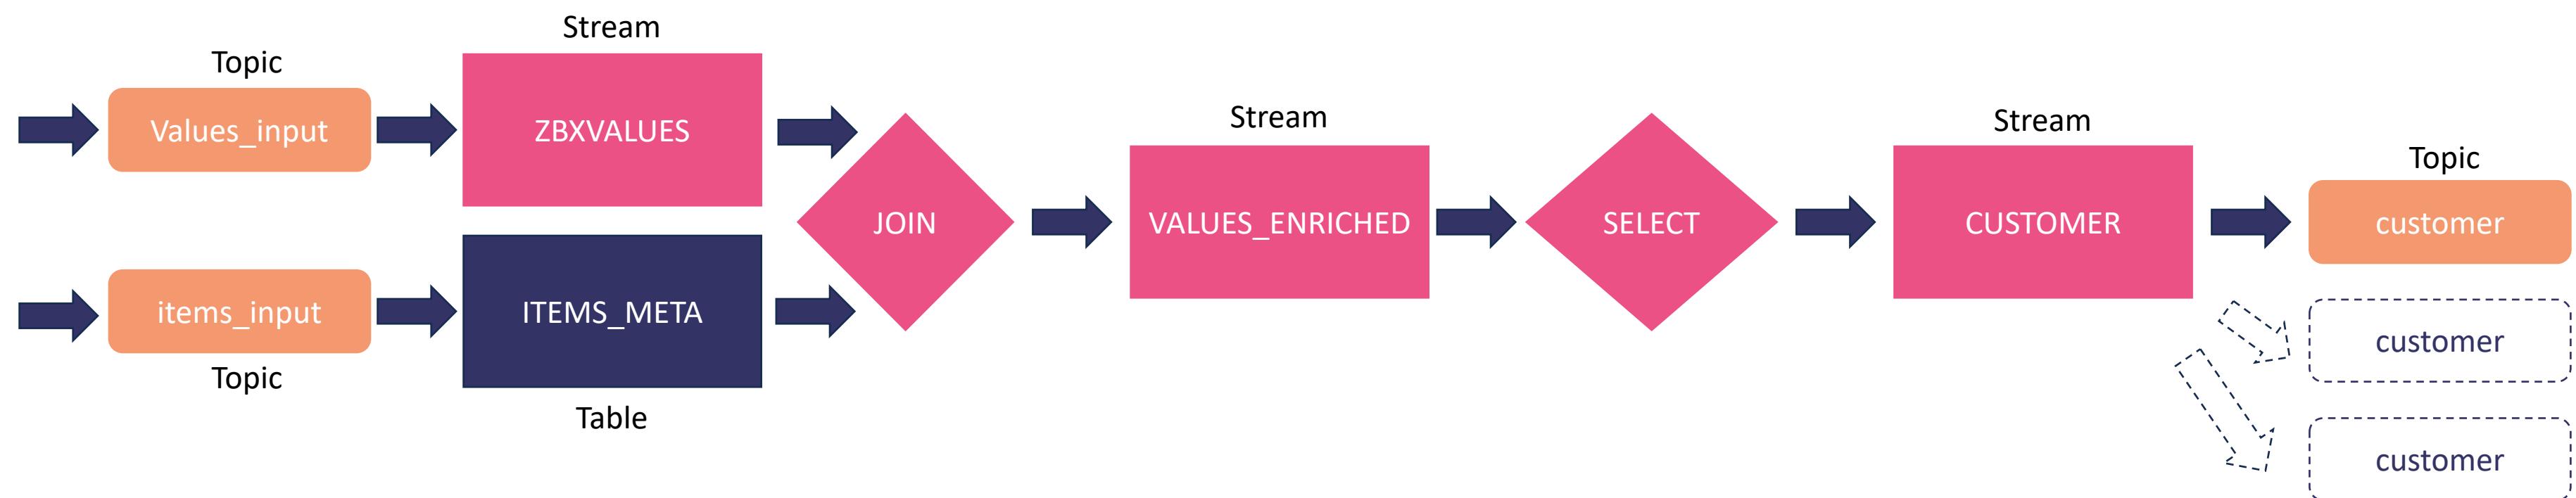
ksqldb

SQL-like interface on top of Kafka

Generates Kafka Streams Applications internally

Real-Time Stream Processing

No JAVA Code



```
1 sqlQueries: |
2   CREATE STREAM ZBXVALUES (itemid STRING KEY, clock INT, value STRING) WITH (KAFKA_TOPIC='values_input', VALUE_FORMAT='json');
3
4   CREATE TABLE ITEMS_META (itemid STRING PRIMARY KEY, master_itemid STRING, delay VARCHAR, name VARCHAR, key VARCHAR, type VARCHAR, hostid INT, host_name
5   VARCHAR, host_host VARCHAR, hostgroups VARCHAR, host_templates VARCHAR, units VARCHAR, description VARCHAR, tags VARCHAR, _customer VARCHAR) WITH
6   (KAFKA_TOPIC='items_input', VALUE_FORMAT='json');
7
8   CREATE STREAM VALUES_ENRICHED WITH (VALUE_FORMAT='json', PARTITIONS=3) AS SELECT i.itemid AS itemid, i.master_itemid, v.clock, v.value, i.delay, i.name,
9   i.key, i.host_name, i.host_host, i.hostgroups, i.host_templates, i.units, i.tags, i.type, i._customer, 'zabbix' AS zabbix_instance FROM ZBXVALUES v LEFT
10  JOIN ITEMS_META i ON v.itemid = i.itemid;
11
12  CREATE STREAM CUSTOMER WITH (KAFKA_TOPIC='customer', VALUE_FORMAT='json', PARTITIONS=3) AS SELECT itemid AS itemid_key, AS_VALUE(itemid) AS itemid,
13  master_itemid, clock, value, delay, name, key, host_name, host_host, hostgroups, host_templates, units, tags, type, zabbix_instance FROM VALUES_ENRICHED
14  WHERE _customer = 'ourcustomer';
```

KsqlDB queries

Overview **Messages** Consumers Settings Statistics

Seek Type Partitions Key Serde Value Serde

Offset Offset All items are selected. String String Clear all Cancel

16017036 + Add Filters

Polling partitions: [0, 1, 2] 57668 ms 3 GB 30278

Offset	Partition	Timestamp	Key	Preview	Value	Preview
+ 917112818	1	4.10.2025, 13:23:09	16017036		{"ITEMID": "16017036", "MASTER_ITEMID": "160170...	
- 917109966	1	4.10.2025, 13:22:09	16017036		{"ITEMID": "16017036", "MASTER_ITEMID": "160170...	

Key Value Headers

Key	Value	Headers	Timestamp	Timestamp type
{ "ITEMID": "16017036", "MASTER_ITEMID": "16017009", "CLOCK": 1759576929, "VALUE": "11.081357999999994", "DELAY": "0", "NAME": "CPU utilization", "KEY": "system.cpu.util", "HOST_NAME": "Test Host", "HOST_HOST": "testhost", "HOSTGROUPS": "Linux, Linux VM, testgroup, zx - Zabbix", "HOST_TEMPLATES": "Linux by Zabbix agent, Zabbix server health", "UNITS": "%", "TAGS": "component: cpu", "TYPE": "Dependent item", "ZABBIX_INSTANCE": "..." }			4.10.2025, 13:22:09	CREATE_TIME
			Key Serde	String Size: 8 Bytes
			Value Serde	String Size: 414 Bytes

Combined Kafka message

Getting Data Out

Metrics Receiver

Python program, running as a service

Consume data from Kafka topic

Conversion into Prometheus metrics format

Meta information as labels

Push data using Prometheus remote_write protocol

Thanos

S3 bucket (on prem) as data storage

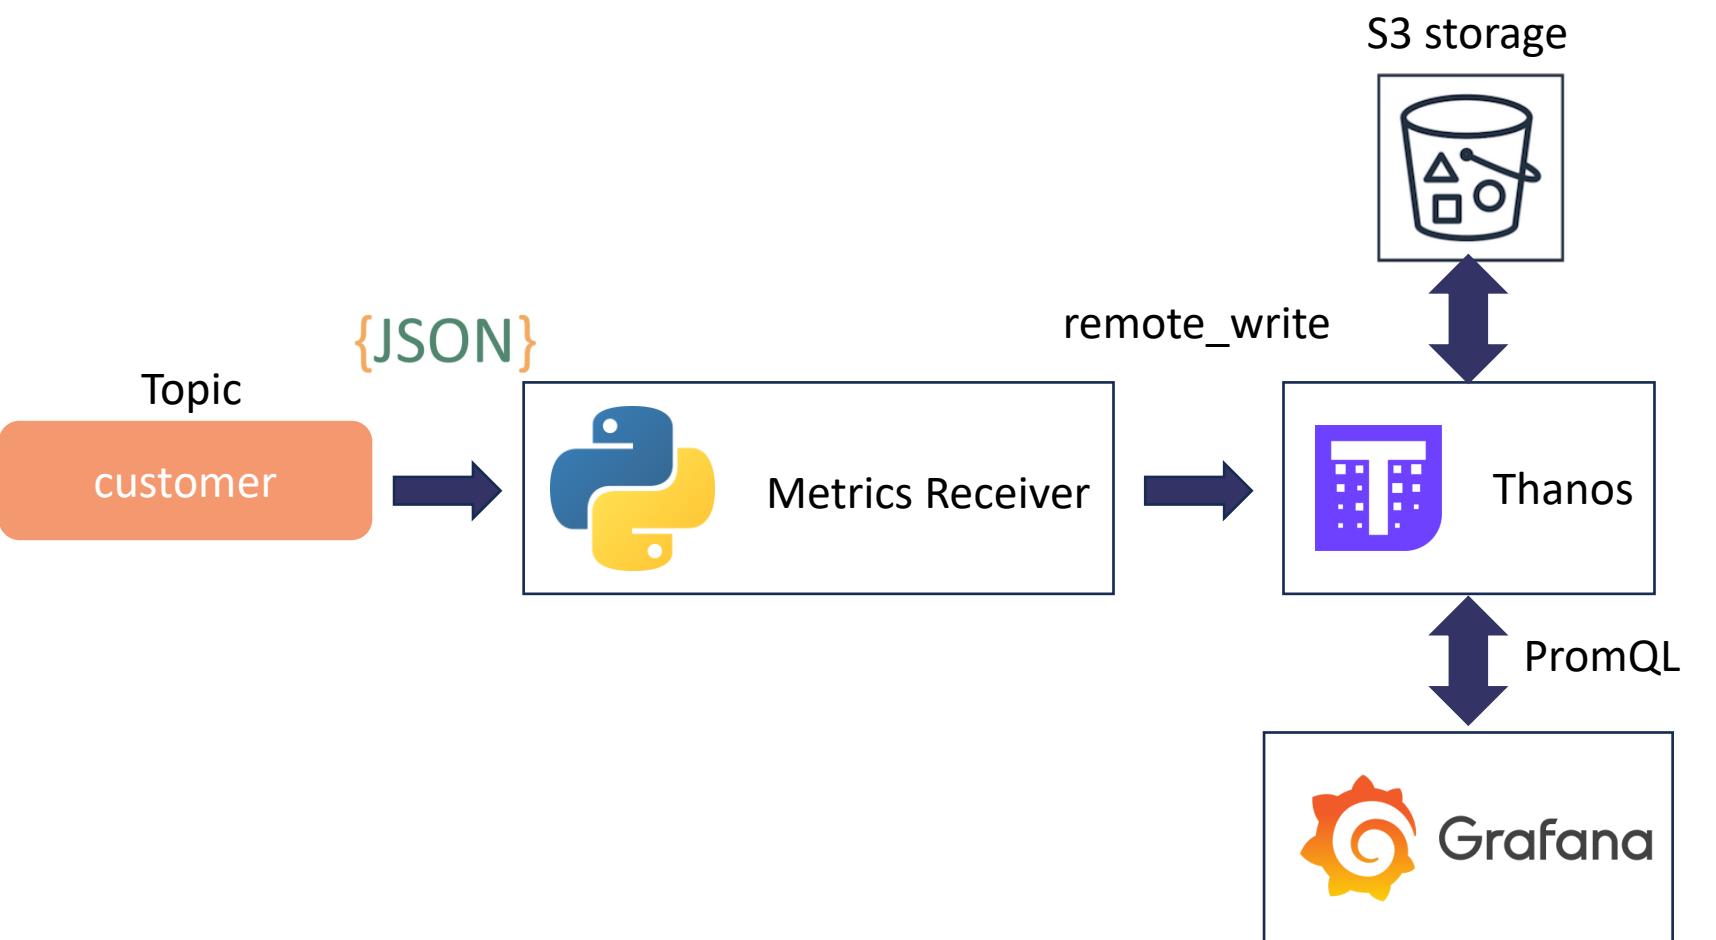
Supports downsampling, "housekeeping"

Provides Query API for Grafana

Grafana

Presentation layer

Rich dashboarding capabilities using PromQL





A single metric in Zabbix and Thanos

There's so much more...

Redis exporter

Generation of Prometheus metrics out of Redis HASH objects

Redis feeders

Query inventory data of Servers, Databases, Network Devices, Printers... from a CMDB and generate Hash entries in the Redis DB.

Also generate entries of “Locations” with Geo coordinates, addresses ...

Messaging transport bridge

A little microservice implementing a “Subscriber Model” so that users can just invite a Bot that emits specific alerts to a personal or group chat

Conclusion

Successfully

built an enterprise-ready, cloud-native observability pipeline with Zabbix as the central data generator

Full Multi-Tenancy

both for additional data sources (e.g. Zabbix, Prometheus) and for customer-specific data isolation

High-Performance

filtering and enrichment of Zabbix monitoring data to enable secure, scalable multi-tenant usage

Reusable

modular architecture that can be adapted or extended as a blueprint for other large-scale observability platforms

Thank you! ! !

christian.anton@inqbeo.de