Zabbix at Scale
-
Global Problem View for Multi-Instance Environments

IntelliTrend GmbH

Contact: Wolfgang Alper

www.intellitrend.de

wolfgang.alper@intellitrend.de

ZABBIX PREMIUM PARTNER

Zabbix at Scale

-

Global Problem View for Multi-Instance Environments

With special thanks to

Deutsche Telekom Technik GmbH

for a great project

ZABBIX
PREMIUM PARTNER

IntelliTrend
IT-Services

When one Zabbix server is not enough

Where it all began

When one Zabbix server is not enough

It started with an overloaded Zabbix Server ...

- Number of Hosts increased
- Number of Items increased
- Number of Triggers increased

This led to a significant increase in utilization, especially for NVPS = **n**ew **v**alues **p**er **s**econd.

Performance Tuning

First step was to analyze and improve the performance of the Zabbix System

- Use proxies for data collection (and preprocessing)
- Use larger time periods in items
- Clean up templates of outdated items
- Use improved items if available
- Use Zabbix Agent active where possible
- Check Zabbix server configuration (cache, poller, etc)
- Check Zabbix proxy configuration (cache, poller, etc)

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

Performance Tuning

Next step was to analyze and improve the performance of the Database System

- Check Database version
- Check Database server settings (huge topic…)
- Check Database server hardware (cpu, memory, storage, disk i/o, volume separation for db and log …)
- Check Database server metrics for changes (reads, writes …)
- Use Partitioning

ZABBIX PREMIUM PARTNER

When one Zabbix server is not enough

The reality of “mega-server”

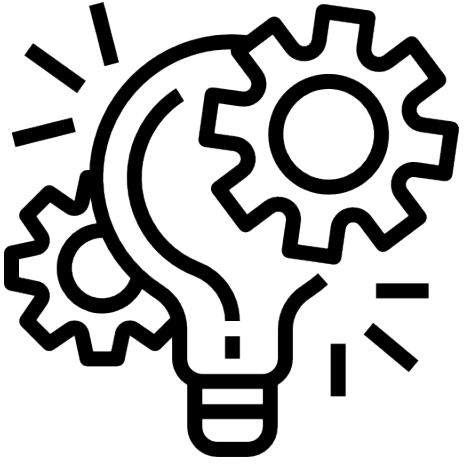

ZABBIX
PREMIUM PARTNER

IntelliTrend
IT-Services

There is always a limit

- No matter how much you tune a single box, CPU, RAM, disk-I/O or network will saturate eventually.
- A “mega-server” upgrade is a big-bang, high-risk investment (long lead-times, large CapEx).
- When the ceiling is hit you must re-architect – not just add more resources.

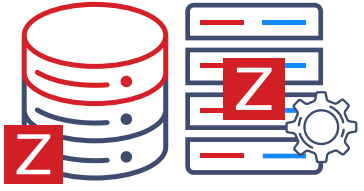

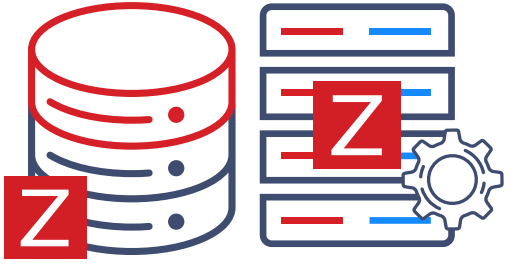

ZABBIX
PREMIUM PARTNER

IntelliTrend
IT-Services

Horizontal scaling = predictable and cost-effective

- **Predictable growth:** Add another identical node whenever CPU, RAM or I/O start to saturate.
- **Cloud-friendly:** Commodity VMs, containers or bare-metal instances are the default offering in every public-cloud catalog.
- **Cost-effective:** Small, off-the-shelf servers (or modest cloud instances) cost far less per unit of compute than a custom high-end appliance. So **capital-expenditure** (CapEx) **stays low** and **operational-expenditure** (OpEx) **can be optimised** by right-sizing each node.
- **Higher availability:** Failure of a single node only removes a slice of capacity; the rest of the cluster continues to operate.
- **Simpler (hardware) lifecycle:** Replace individual nodes without downtime for the whole service.

ZABBIX PREMIUM PARTNER

Scaling using multiple Zabbix Server

Data Collection: To scale use multiple Zabbix Proxies

Problem Detection and Alerting: To scale use multiple Zabbix Servers

Problem View: Each Zabbix Instance has it’s own problem view

Central problem **view** gets **lost.** Ops-team needs to check all instances.

ZABBIX PREMIUM PARTNER

When one Zabbix server is not enough

The Global Problem View Server idea

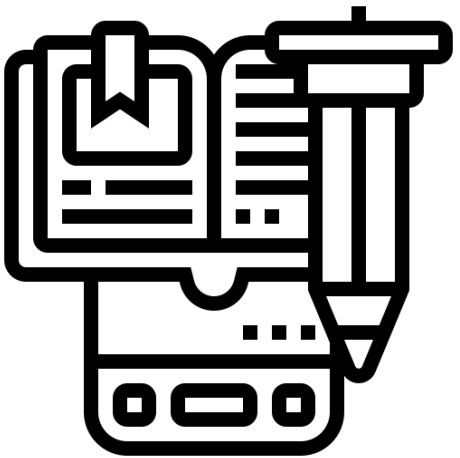

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

The GPV-Server idea

What if there were a **central server** that would **unify** the view of **problems** across multiple Zabbix servers?

Aggregated problems
across all Zabbix instances

The GPV-Server idea

How it should work

- **Pull** open problems from any number of Zabbix servers via the Zabbix API (read-only token).
- **Normalise** host names, tags, and severity into a common schema.
- **Store** the unified list in its own tiny DB.
- **Cache** metadata for quick access.
- **Provide** an API that shows all problems together.
- **Integrate** into the Zabbix frontend using the module API.
- **Support** Zabbix HA setups (Server and Frontend).

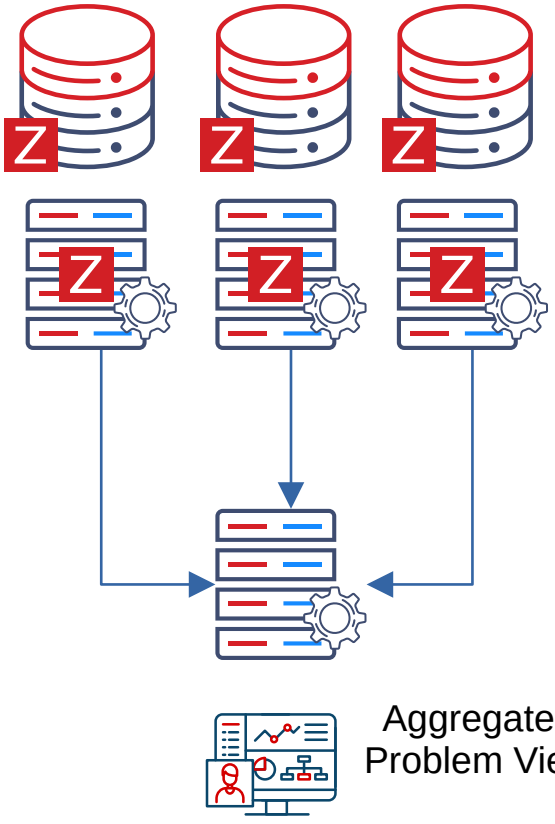

The GPV-Server idea

From the user's perspective

- One or more **Master Zabbix Server** that show problems from related Zabbix instances.
- **Filter settings** in Problem View to select specific Zabbix instances if needed.
- **Support** for Zabbix Server HA and Zabbix Frontend HA.
- Global-Problem-View **Dashboard Widget** that allows to create specific Dashboards by Instance or as summary.
- **Add** a new Zabbix Server by just adding a new endpoint.
- **Jump to** Source Instance by clicking on the problem in the global problem view.

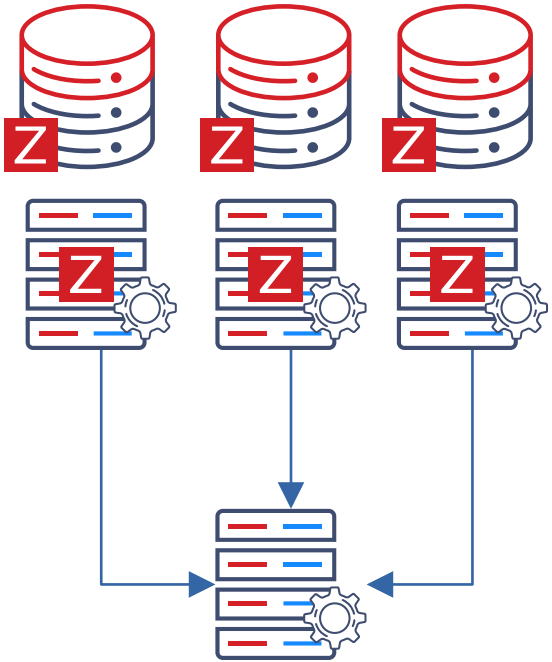

When one Zabbix server is not enough

GPV-Server architecture

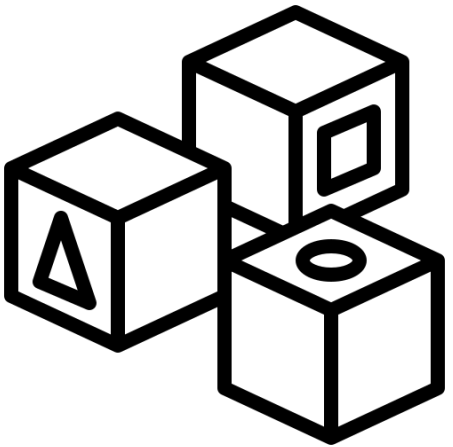

ZABBIX
PREMIUM PARTNER

IntelliTrend
IT-Services

GPV-Server architecture

Design considerations

The GPV-Server **cannot** simply perform parallel request using the Zabbix API to call Zabbix instances when a client requests a list of **current problems**. Why:

- The GPV-Server needs to perform **multiple API calls** per Zabbix instances to produce an open problem list with all the metadata required by the Zabbix frontend: problem.get, trigger.get, event.get, mediatype.get, user.get and maintenance.get
- Some of these API calls have **sequential dependencies**. They cannot be executed in parallel, because their parameters depend on a previous API call.
- This means the GPV-Server would perform at least **6 API calls** per Zabbix instance. Example with 10 Zabbix instances: 6 x 10 = 60 API calls just to load the list of problems. And - the user would have to **wait** for the **slowest** servers to complete or timeout.

This is why the design adds a db **cache** for open/closed problems on the GPV-Server.

GPV-Server architecture

Replication process

- **Replication Manager** creates and manages **Syncers**.
- **One syncer process** per Zabbix instance.
- Each syncer process has an independent asynchronous **replication loop** and a separate independent **ping loop**.
- The replication loop **replicates events** and all required metadata.
- The ping loop **checks connectivity** with a single frontend or all HA frontends for the syncer's target.

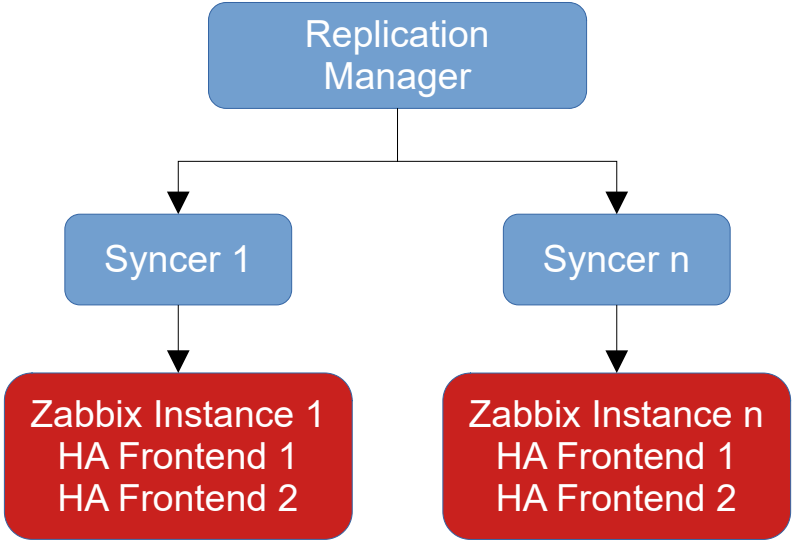

GPV-Server architecture

Optimization when using HA Frontends

- To increase the efficiency of the replication we can take advantage of HA frontends by doing Round-Robin style load balancing.
- Example: All events with an eventid between 0 and 150 to be fetched using a parallel loading strategy using all HA frontend nodes.
- Of cause this is only possible, if there are multiple frontend nodes.

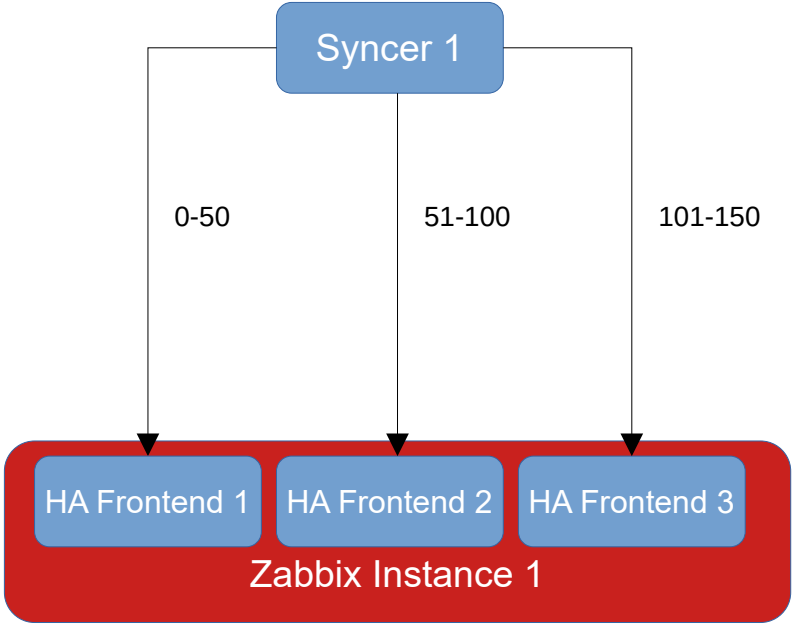

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

GPV-Server architecture

Overload handling

- If a **syncer** process **gets overloaded** (new events open faster than we are able to replicate them) we handle this similar to a WAL based replication of a database.
- We simply accept that we fall behind and **keep replicating** at a **steady maximum** replication rate until we eventually catch back up.
- This approach leads to **predictable behaviour** and avoids unexpected load spikes on the GPV server.

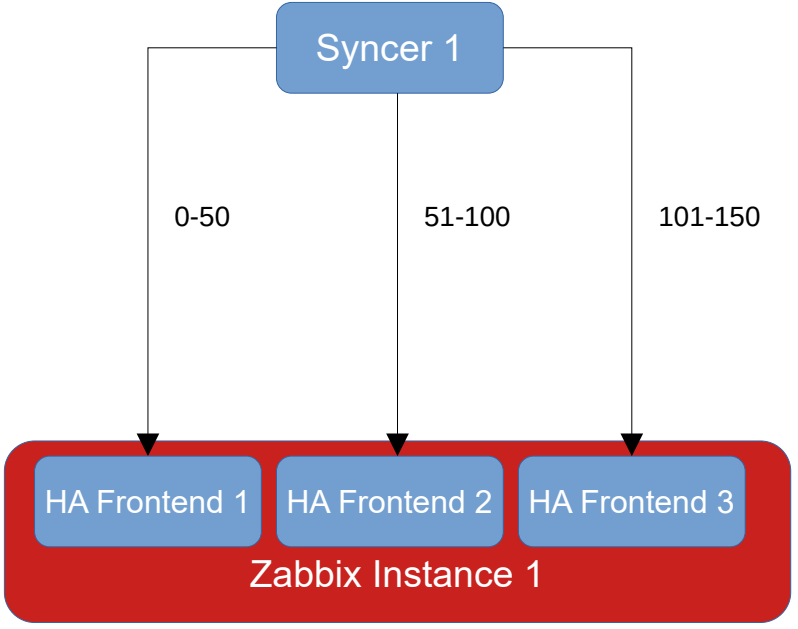

GPV-Server architecture

GPV-Server HA setup

- The GPV-Server can be deployed in an active/active **HA setup**.
- There is no synchronization needed, because the GPV-Server uses its database only as a **cache**, similar to a Zabbix-Proxy.
- The GPV frontend modules support configuring multiple GPV-Server nodes and switches between them automatically based on their availability and replication status.
- This makes it easy to deploy the GPV-Server in **cloud environments** like Kubernetes.

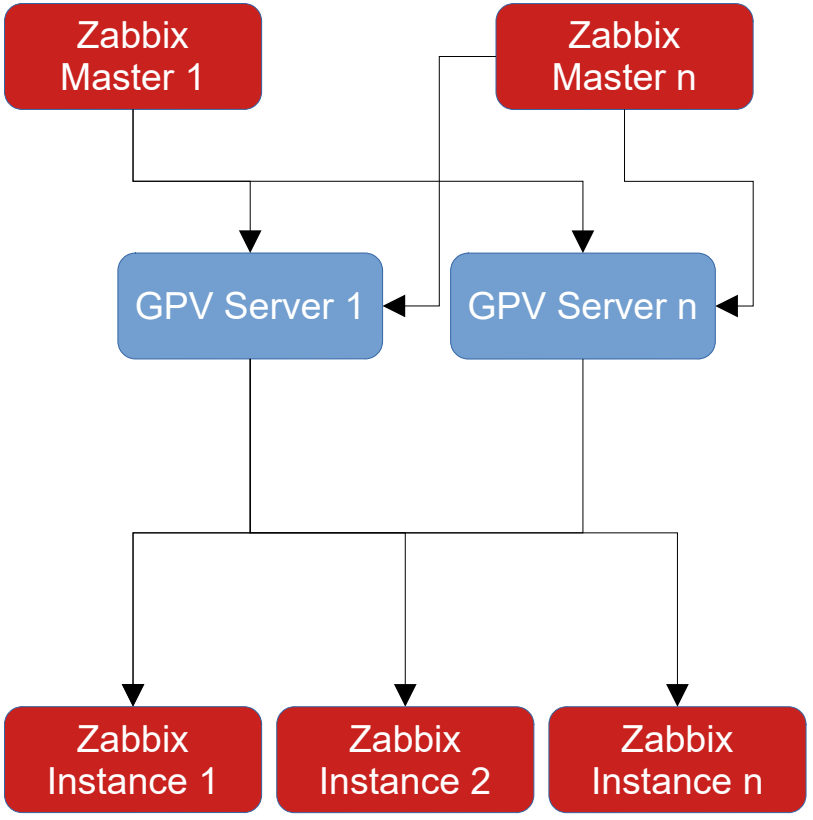

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

GPV-Server architecture

Overview of Zabbix instances states

A Zabbix instance can be in one of the following states from the perspective of the GPV-Server:

- INITIALIZING: The instance is currently being initialized and is not yet operational.
- ACTIVE: The instance is fully replicated and there are no problems.
- DEGRADED: The instance is replicated and its status can be used, but either it has increased replication delay or one of its HA frontends is unavailable.
- FAILURE: Replication for the instance has failed and should no longer be used.

Configured instances

ID	Instance name	Instance state	Zabbix version	Connection state	Latency (min/avg/max)	Last contacted	Replication state	Replicated until	Average replication duration
1	Zabbix Instance 01	Active	7.0.18	Online (2 / 2 nodes)	22 ms / 25.2 ms / 27 ms	2s ago	Idle	9s ago	264.1ms
2	Zabbix Instance 02	Degraded	7.0.18	Degraded (1 / 2 nodes)	15 ms / 17 ms / 19 ms	2s ago	Idle	7s ago	270.4ms
3	Zabbix Instance 03	Active	7.0.18	Online (2 / 2 nodes)	23 ms / 24.7 ms / 27 ms	2s ago	Idle	7s ago	271.1ms

GPV-Server architecture

Overview of GPV-Server states

A GPV Server can have one of the following states:

- INITIALIZING: One or more instances on the GPV server are currently being initialized. Another GPV server with ACTIVE status should be preferred, if available.
- ACTIVE: All instances have ACTIVE status. The server is operating normally.
- DEGRADED: One or more instances have either DEGRADED or FAILURE status. At least one instance has a status other than FAILURE.
- FAILURE: Replication for the instance has failed and should no longer be used.

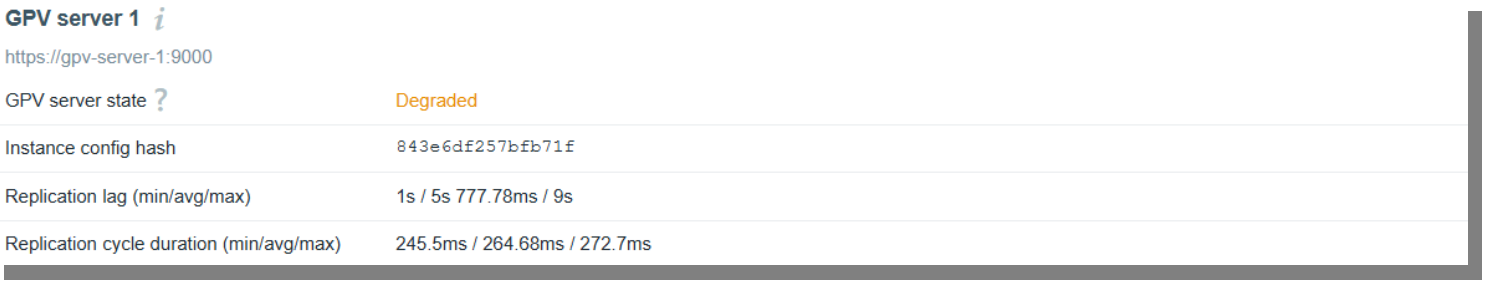

When one Zabbix server is not enough

Zabbix redundant HA setups

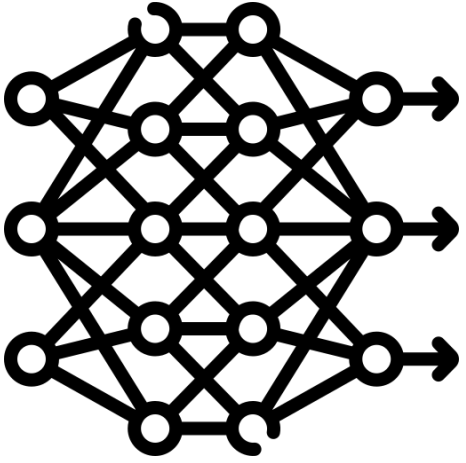

ZABBIX
PREMIUM PARTNER

IntelliTrend
IT-Services

Overview - Zabbix HA and Geo redundancy

High Availability and (Geo) redundancy can be implemented in different areas of the Zabbix software stack:

- **HA Zabbix Database Service**:
  Database Cluster with a minimum of 3 Nodes depending on setup
- **HA Zabbix Server**:
  Zabbix Server Cluster (>= V6.0) with a minimum of 2 Nodes
- **HA Zabbix Frontend**:
  Multiple Webserver running the Zabbix Frontend
- **HA Zabbix Proxies**:
  Zabbix Proxy Groups (>= V7.0) with a minimum of 2 Zabbix Proxies per group

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

Example of a HA setup for one Zabbix instance

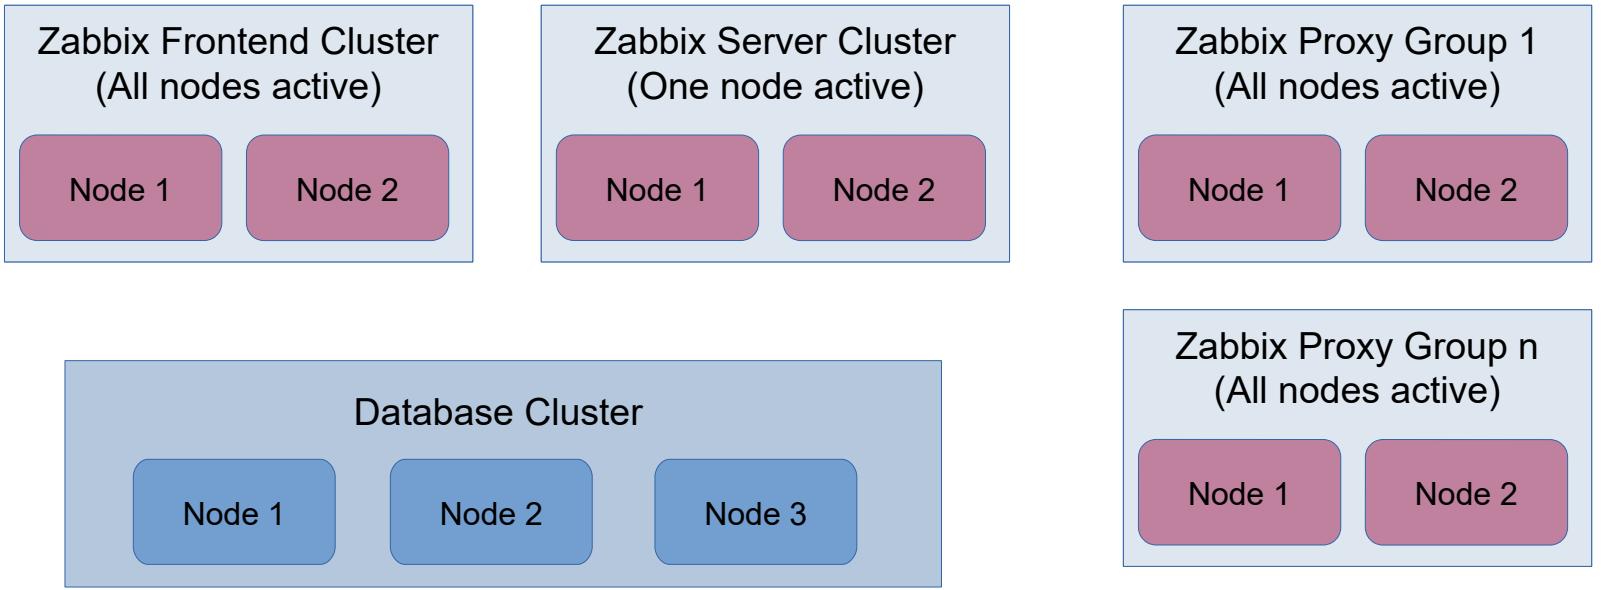

Global problem view across multiple HA instances

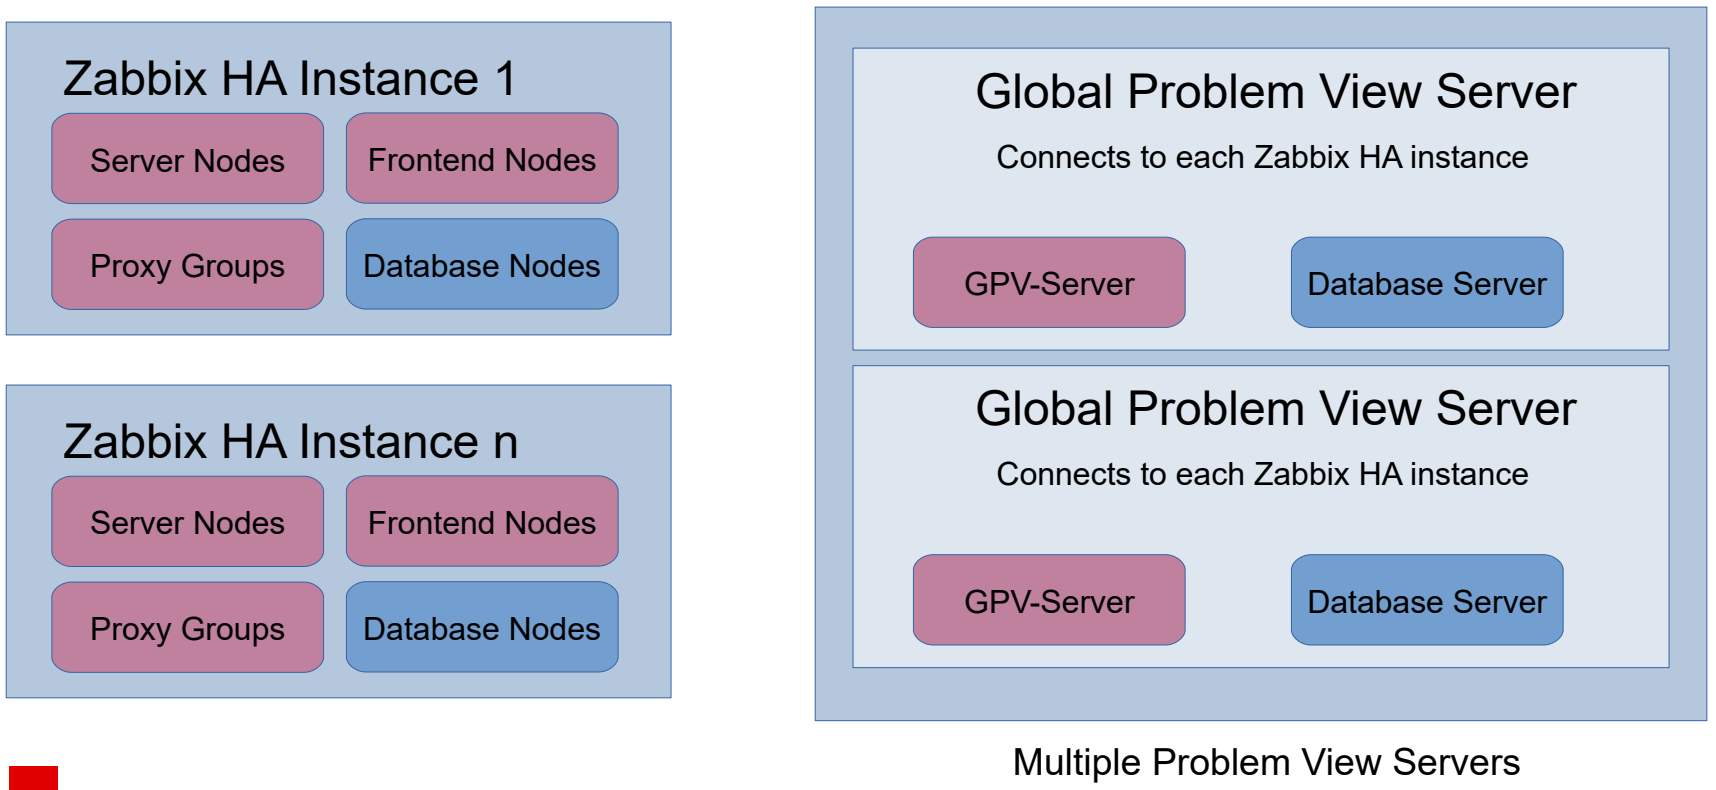

Multiple Problem View Servers

When one Zabbix server is not enough

GPV-Server implementation

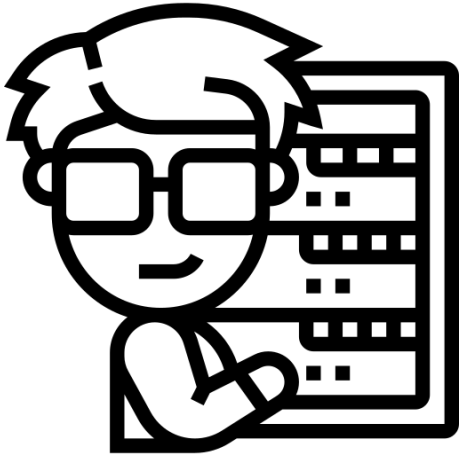

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

GPV-Server Configuration

GPV-Server Zabbix instance configuration

```
instances:
  - name: Zabbix Instance 01
    id: 1
    token: 9d4axxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx

    apis:
      - https://zabbix7-web-01/api_jsonrpc.php
    frontend: https://zabbix7-web-01
  - name: Zabbix Instance 02
    id: 2
    token: 4ed8xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
    apis:
      - https://zabbix7-web-02/api_jsonrpc.php
    frontend: https://zabbix7-web-02
```

GPV-Server Frontend Modules

There are two modules that can be installed on any Zabbix frontend, that should act as a master server:

- **Global problem view**: Adds a new top level menu “Global problem view” with two submenus to access the global problems and also the state on instances and GPV-Server.

- **Global problem widget**: Adds a new dashboard widget that can be used to show all problems in a dashboard, or filter by a given instance.

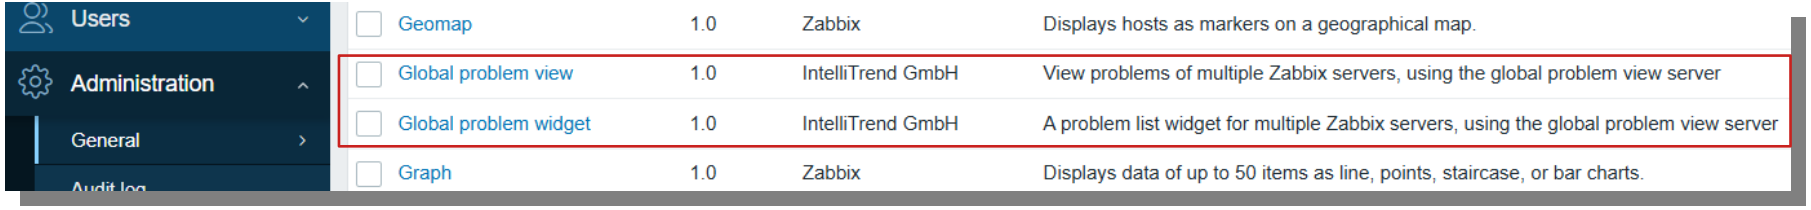

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

GPV-Server Frontend Module – Problem overview

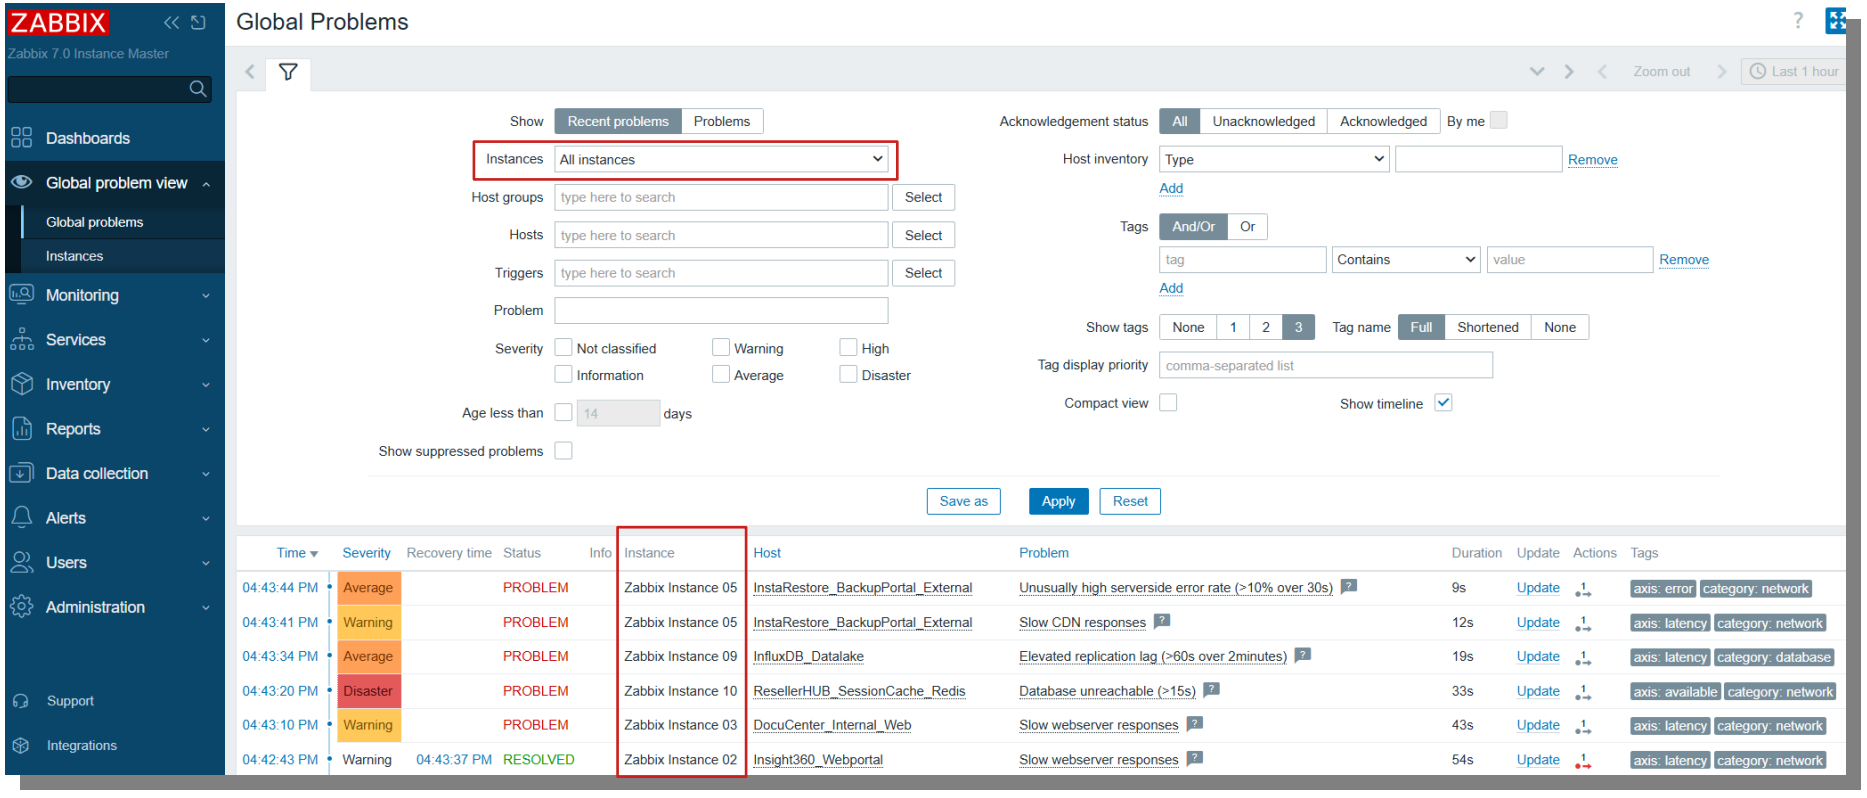

The global problem view page mimics the orginal problem page and includes an instance selector.

GPV-Server Frontend Module – Problem details

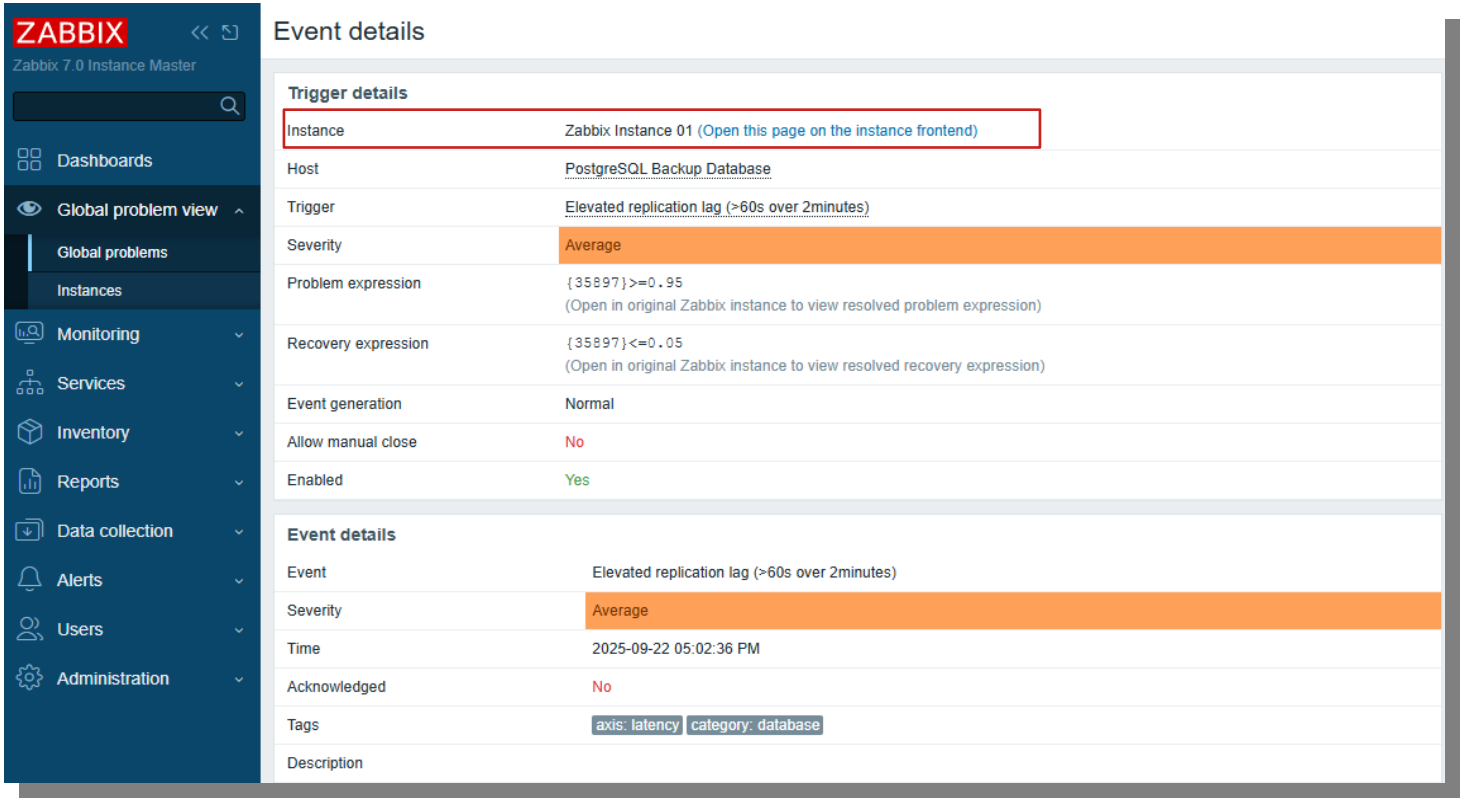

The global problem details page mimics the orginal page and includes the instance information with a link to the instance

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

GPV-Server Dashboard Widget

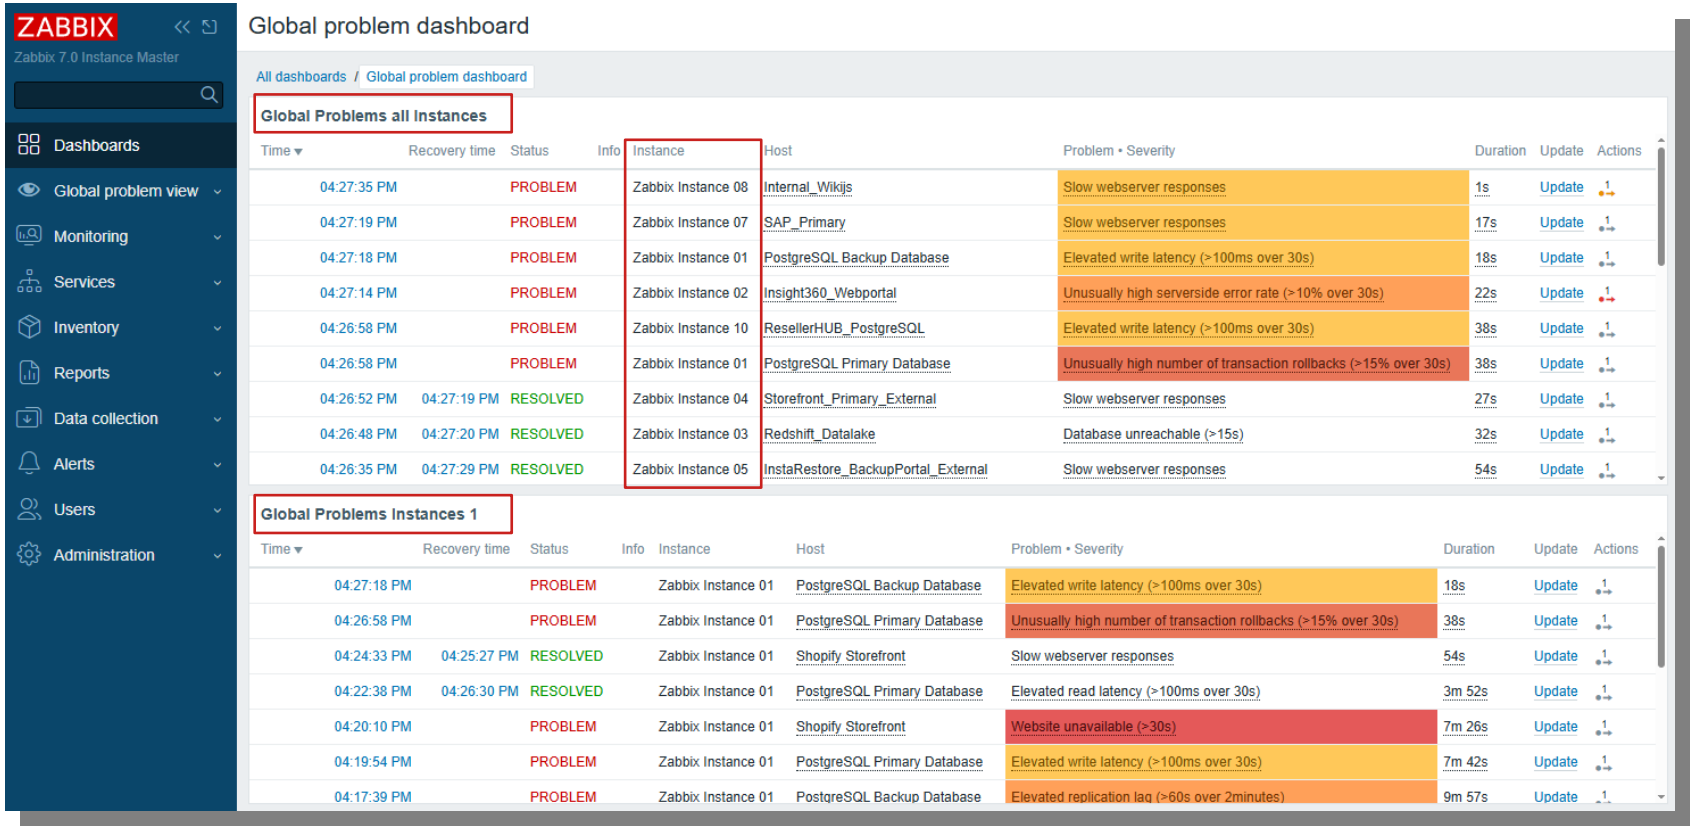

GPV-Server dashboard widget allows to create Zabbix Instance specific problem dashboards. Can be combined with any other widget.

GPV-Server Dashboard Widget Settings

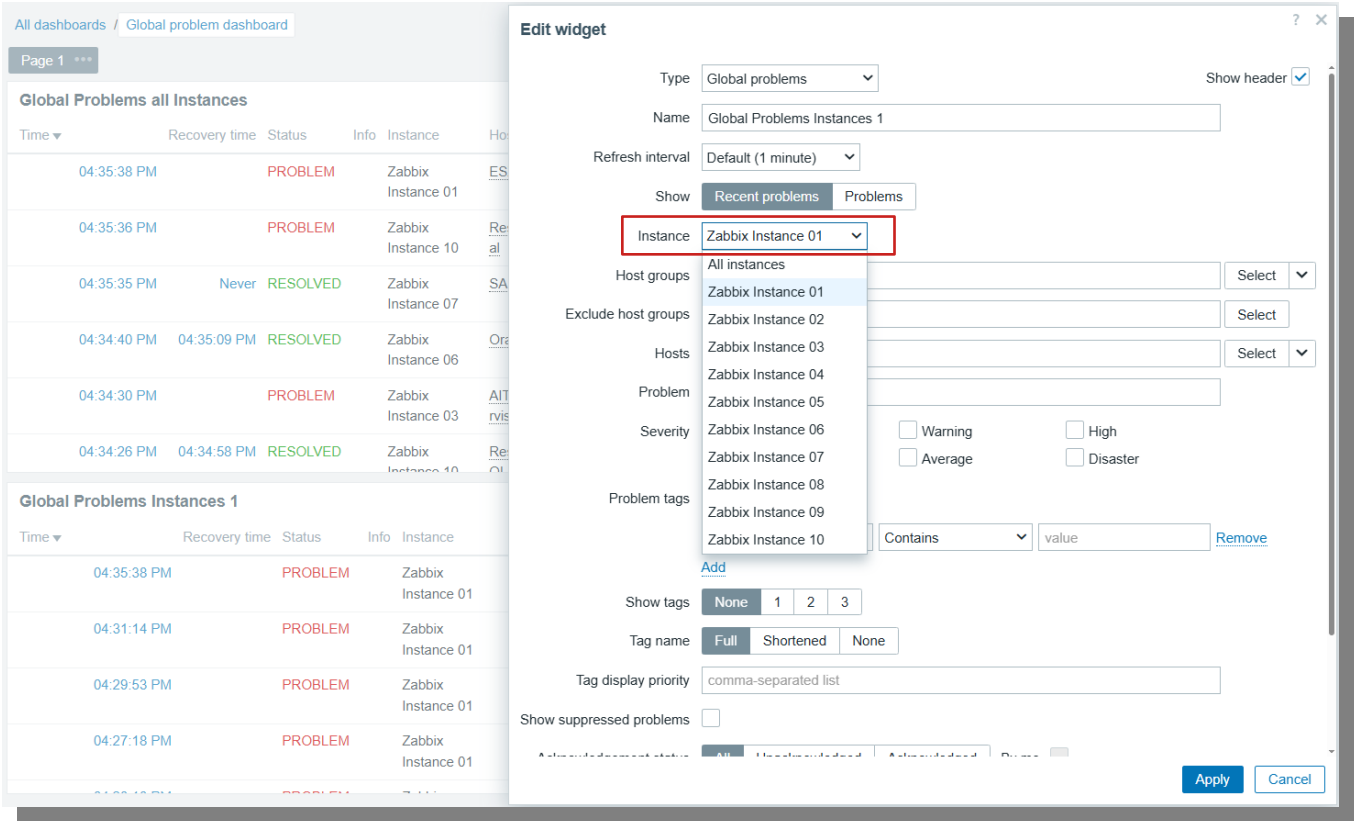

GPV-Server dashboard widget supports many configuration options including an instance selector

GPV-Server Instance overview

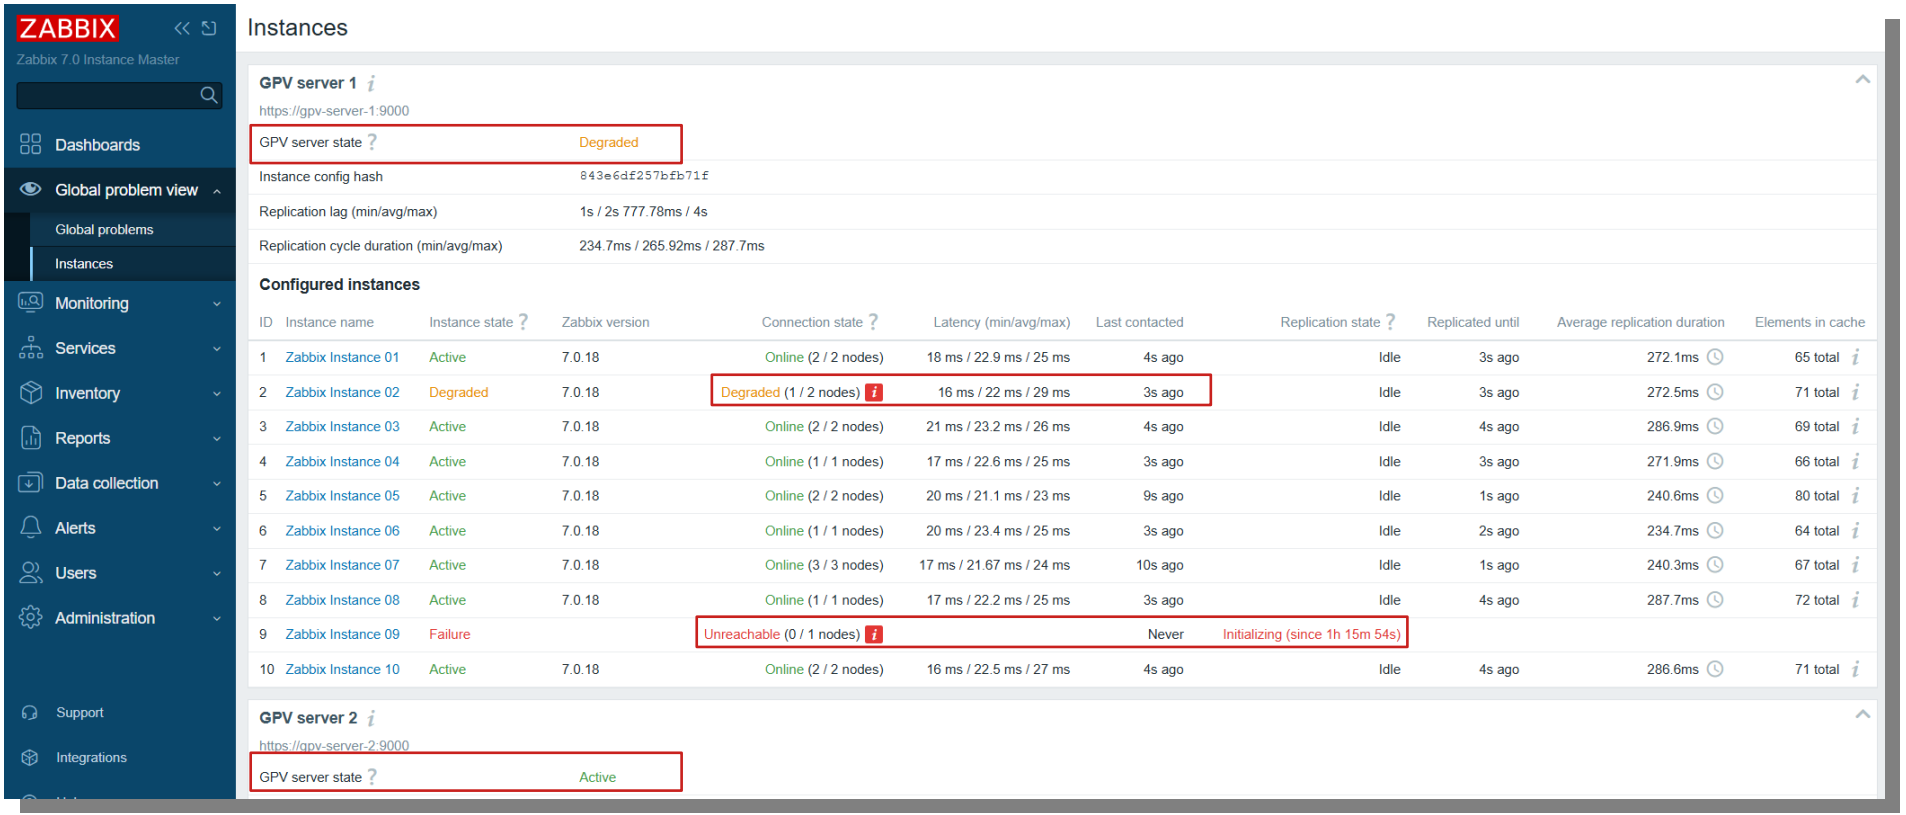

Instance page shows status of each Zabbix instance and each GPV-Server in realtime

GPV-Server Monitoring API

```
{
   "jsonrpc": "2.0",
   "method": "status.get",
   "id": 1,
   "params":
      {"selectInstances": true}
}
```


```
{
  "result": {
    "gpv_server_state": "active",
    "error": "",
    "instance_config_hash": "8486665c2ebd3dcf",
    "instances": [
      {
        "name": "Zabbix Instance 01",
        "instance_state": "active",
        "loop_state": "cooldown",
        "frontend_url": "https://zabbix7-web-01",
        "connection_state": "connected",
        "connection_error": "",
        "zabbix_version": "7.0.18",
        "api_urls": [
          "https://zabbix7-web-01/api_jsonrpc.php"
        ],
        "recent_loop_durations": [
              225, 262, 255, 265, 244, 266, 260, 261, 282
        ],
        "id": 1,
        "init_started_at": 1758468831,
        "replicated_until": 1758526497,
        "replicated_until_event_id": 42169,
        "replication_lag": 9,
        "last_loop_duration": 282,
        "avg_loop_duration": 257.1,
        "last_seen": 1758526504,
        "avg_node_latency": 23.4,
        "best_node_latency": 23,
        "worst_node_latency": 25,
        "nodes_connected": 2
...
      }
```


HTTP JSON-RPC Endpoint provides metrics about the GPV-Server

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

When one Zabbix server is not enough

Single pane of glass

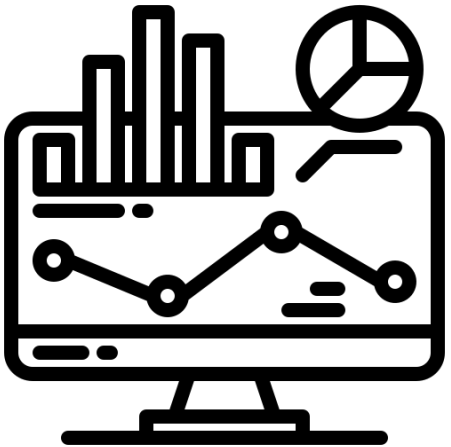

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

Single pane of glass

A Global View Server can serve as the basis for a central overview for **IT System Integrators** and **MSPs** ,who use many different individual Zabbix Servers for their customers.

Those setups will also include:

- **SSO**: Single Sign-On to simply access to all instances. To support OAuth2.0 we had to develop a **"Zabbix OAuth to SAML Proxy"** that automatically authenticates users with a valid OAuth session by acting as a SAML identity provider for Zabbix.
- **Instance Monitoring**: Health-Check of each Zabbix instance.
- **Service Monitoring**: Health-Check of an entire Zabbix service across many instances.

ZABBIX PREMIUM PARTNER

IntelliTrend IT-Services

Zabbix as a Service on the Master

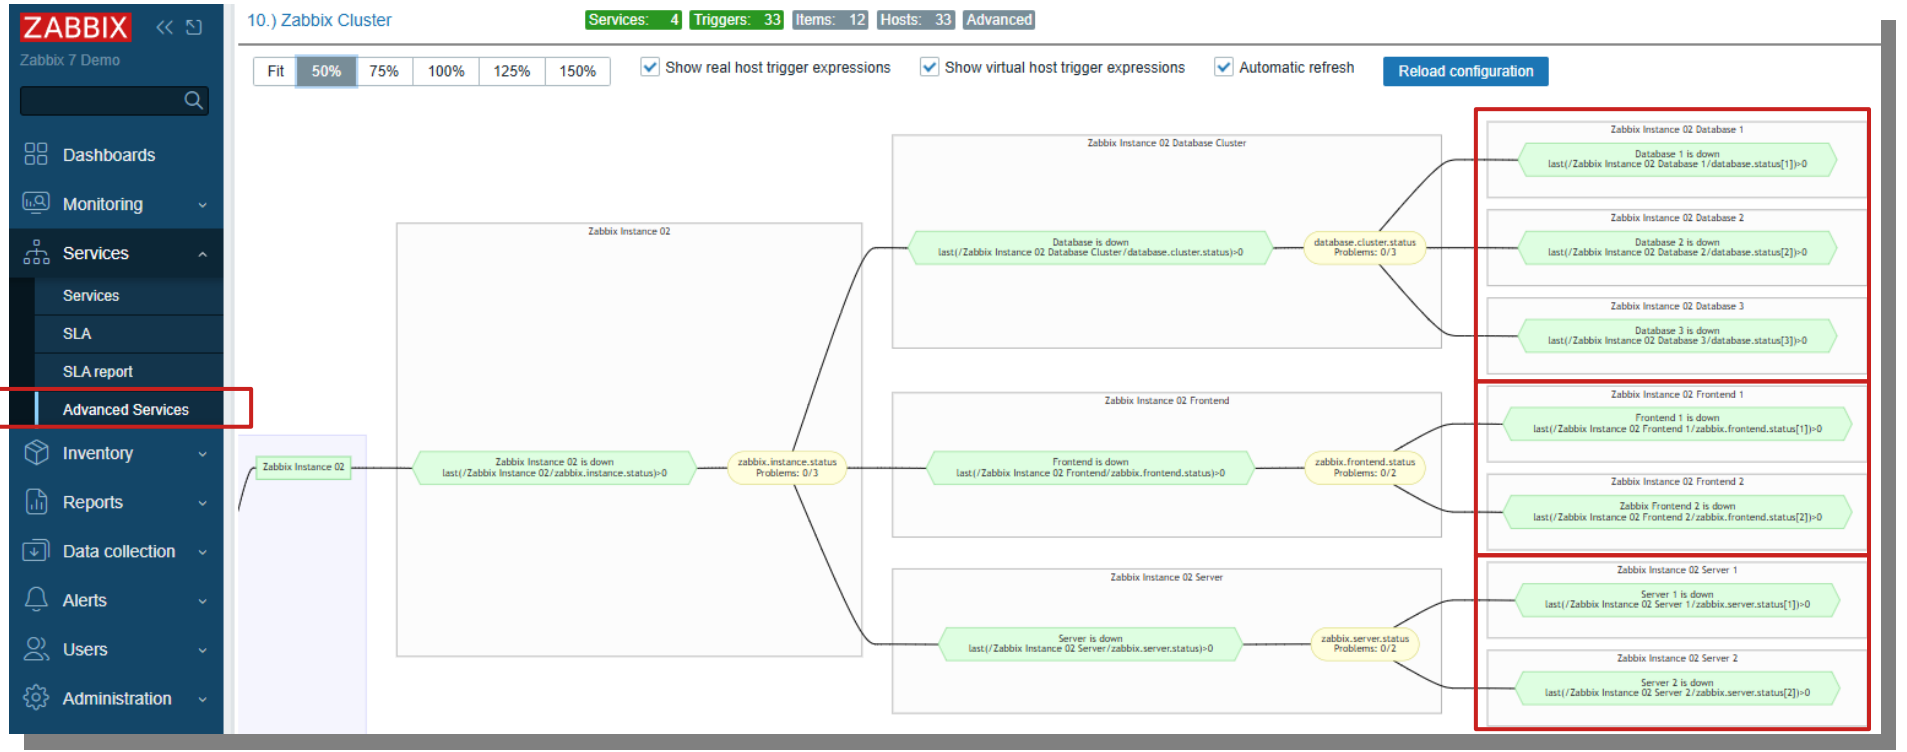

Zabbix as a Service on the master including all instances using the Advanced Service module

Zabbix Service - Anomaly Detection using AI

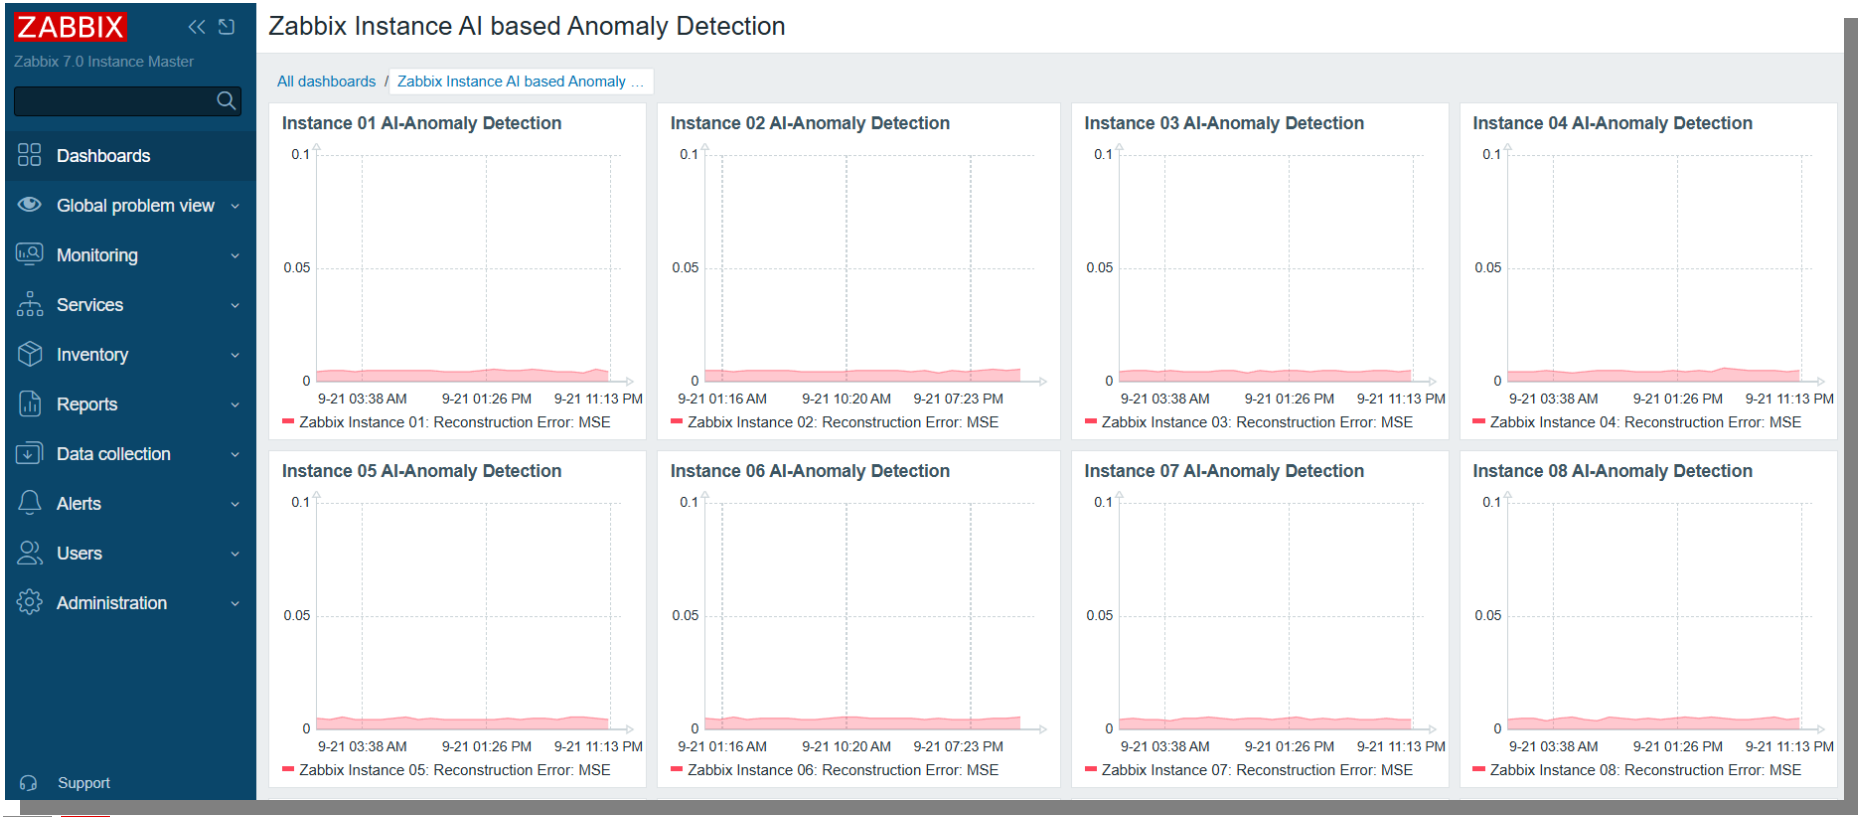

Check for anomalies on each Zabbix instance via the master using individual trained models.

Dashboard - Utilization of Zabbix instances

Zabbix Instances

Host	Address	Server version	Status	Required performance	Config cache	Value cache	History index cache	History write cache	Trend write cache
Zabbix Instance 01	10.0.120.1:10051	7.0.18	Up (1)	7481.1581 NVPS	1.0719 %	5.099 %	0.4515 %	0 %	2.9154 %
Zabbix Instance 02	10.0.120.2:10051	7.0.18	Up (1)	5350.1731 NVPS	3.1955 %	2.0749 %	13.9879 %	0.04994 %	0.0036 %
Zabbix Instance 03	10.0.120.3:10051	7.0.18	Up (1)	3362.1242 NVPS	2.5922 %	2.1029 %	0.4517 %	0.00639 %	0.0288 %
Zabbix Instance 04	10.0.120.4:10051	7.0.18	Up (1)	3945.6279 NVPS	1.2935 %	32.2604 %	1.9426 %	0.01783 %	23.8251 %
Zabbix Instance 05	10.0.120.5:10051	7.0.18	Up (1)	5431.3414 NVPS	0.3193 %	7.0062 %	0.0608 %	0 %	12.7068 %
Zabbix Instance 06	10.0.120.6:10051	7.0.18	Up (1)	4964.5955 NVPS	2.193 %	0.4568 %	0.06178 %	0.000082 %	0.711 %
Zabbix Instance 07	10.0.120.7:10051	7.0.18	Up (1)	7012.5817 NVPS	0.3698 %	0.1937 %	0.7519 %	0.01135 %	0.0288 %
Zabbix Instance 08	10.0.120.8:10051	7.0.18	Up (1)	6502.3986 NVPS	1.1294 %	0.6217 %	0.6142 %	0 %	4.2791 %
Zabbix Instance 09	10.0.120.9:10051	7.0.18	Up (1)	7528.3964 NVPS	15.2103 %	5.5564 %	0.3453 %	0.5937 %	0.762 %
Zabbix Instance 10	10.0.120.10:10051	7.0.18	Up (1)	3477.6317 NVPS	0.05022 %	15.0332 %	0.4471 %	0 %	0.8389 %

Designated Master Server provides central overview

Zabbix meets AI

Summary

Zabbix at Scale – Summary

- Any “mega-server” can exhaust CPU, RAM, I/O or network – it can’t be extended indefinitely.
- Distributing the load across multiple Zabbix Server is an option.
- Multiple Zabbix Servers are aggregated into one Global Problem.
- This allows to scale one Zabbix Server into many Zabbix instances, but keeps a single pane of glass.
- The concept of a global view can also be used by MSP’s or IT System Integrators, where many customers use numerous Zabbix Servers.
- SSO (OAuth 2.0 / SAML) lets users authenticate once and instantly jump into the problem’s source.
- The server designated as the Zabbix Master Server which provides the global view, can also be used to monitor the individual Zabbix Instances.

Zabbix at Scale
-
Global Problem View for Multi-Instance Environments

Thank You!

IntelliTrend GmbH

Contact: Wolfgang Alper

www.intellitrend.de

wolfgang.alper@intellitrend.de

ZABBIX PREMIUM PARTNER