initMAX

ZABBIX PREMIUM PARTNER

ZABBIX CERTIFIED TRAINER

Webinar

Zabbix and advanced database monitoring

all your microphones are muted
ask your questions in Q&A, not in the Chat
use Chat for discussion, networking or applause

1

Why Database Monitoring

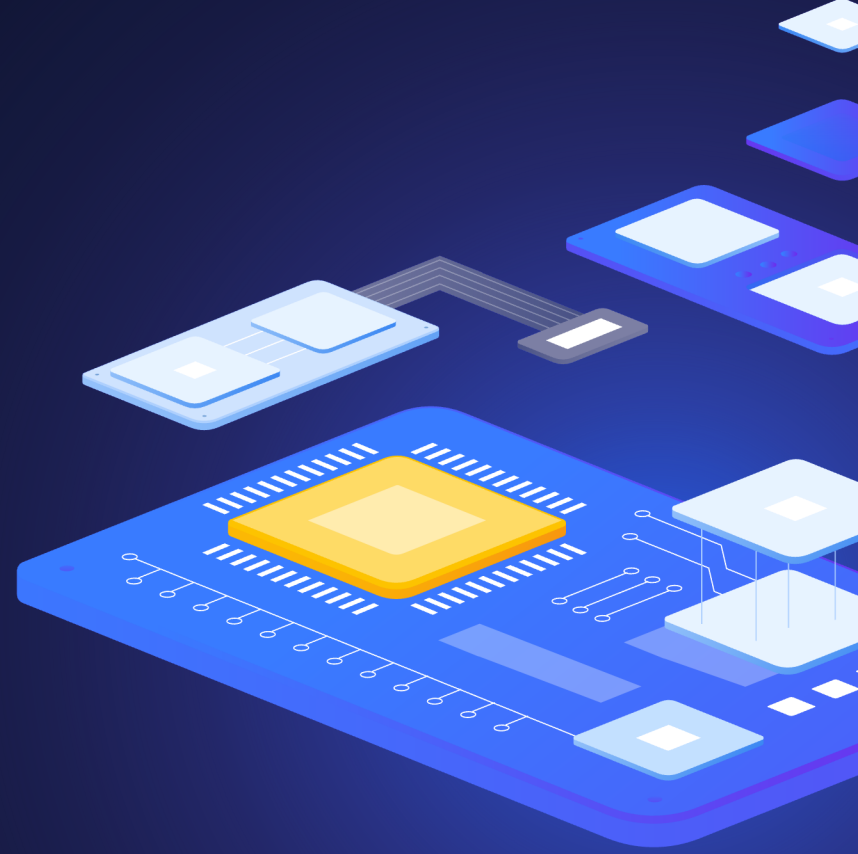

initMAX

Why Database Monitoring

- **Business continuity**
  - Proactive monitoring helps prevent downtime and data loss
- **Critical infrastructure component**
  - Databases form the backbone of most business applications
- **Bottleneck identification and performance impact**
  - Database issues directly affect application performance and user experience
- **Resource utilization and capacity planning**
  - Databases consume significant system resources (CPU, memory, disk I/O)

initMAX

Why Database Monitoring

initMAX

Monitoring Best Practices

- **Database credentials are sensitive**
  - Always use least privilege principle for your credentials.

- **Think about best way how to use authentication**
  - Configuration Files
  - Regular passwords in Frontend
  - User macros / secret macros

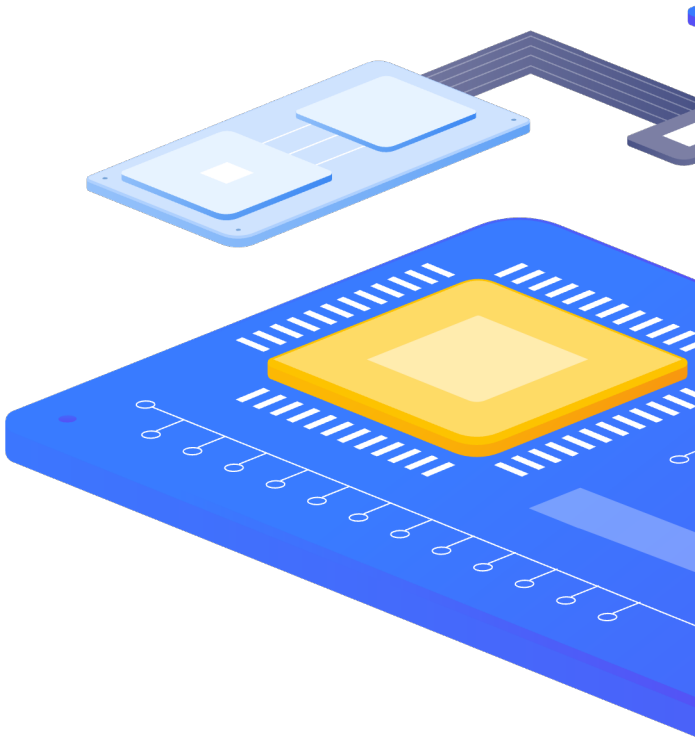

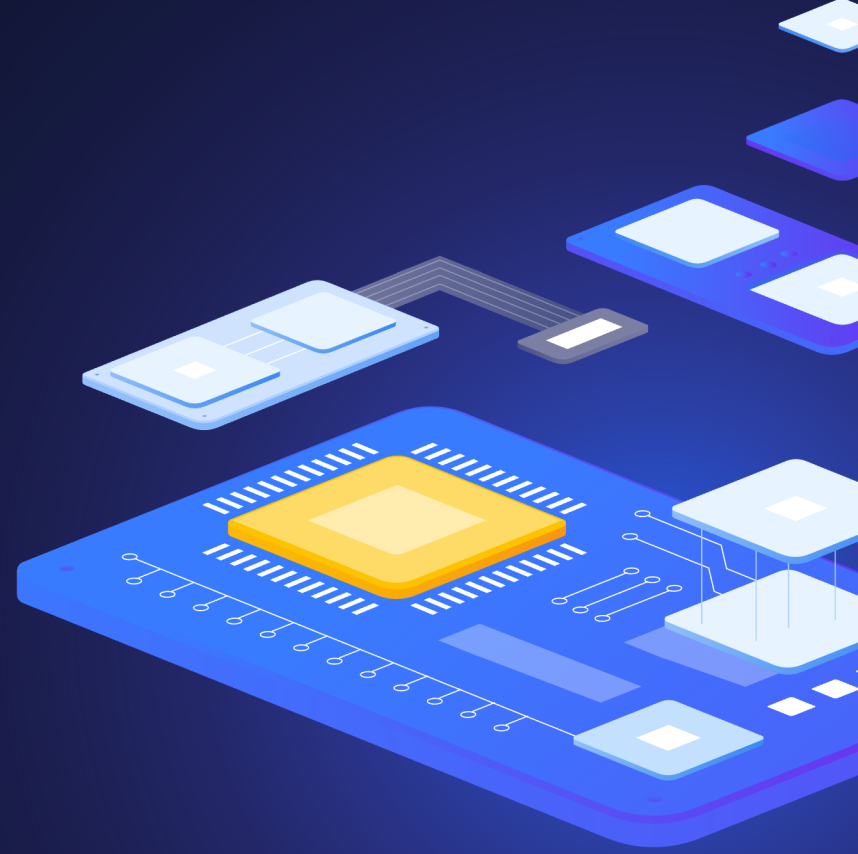

2

Monitoring Methods and Technologies

initMAX

Monitoring Methods and Technologies

- **Agent-based monitoring**
  - Agent2 plugins for major databases
  - Custom scripts and user parameters
- **Direct database querying**
  - ODBC for SQL-based monitoring
  - Native database connections
- **Application interfaces**
  - JMX for Java-based databases
  - HTTP API monitoring
- **Network-level monitoring**
  - TCP/IP checks for availability

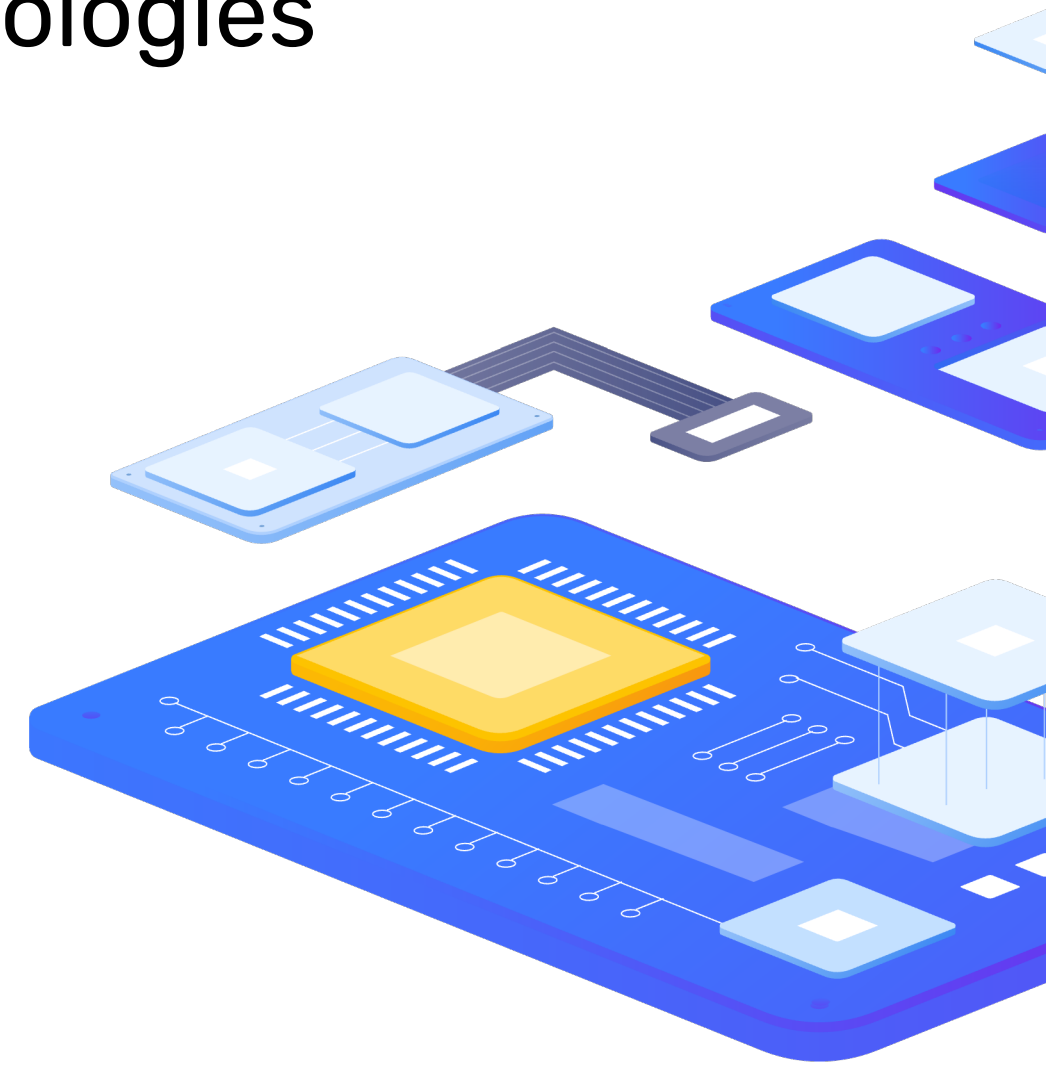

initMAX

Agent2 Plugins

- Plugins for Zabbix Agent 2 database monitoring
  - Some of them are included in installation package

- **Available plugins**
  - PostgreSQL (requires separate installation)
  - MySQL/MariaDB/Percona
  - MSSQL (requires separate installation)
  - Oracle
  - MongoDB (requires separate installation)
  - Redis

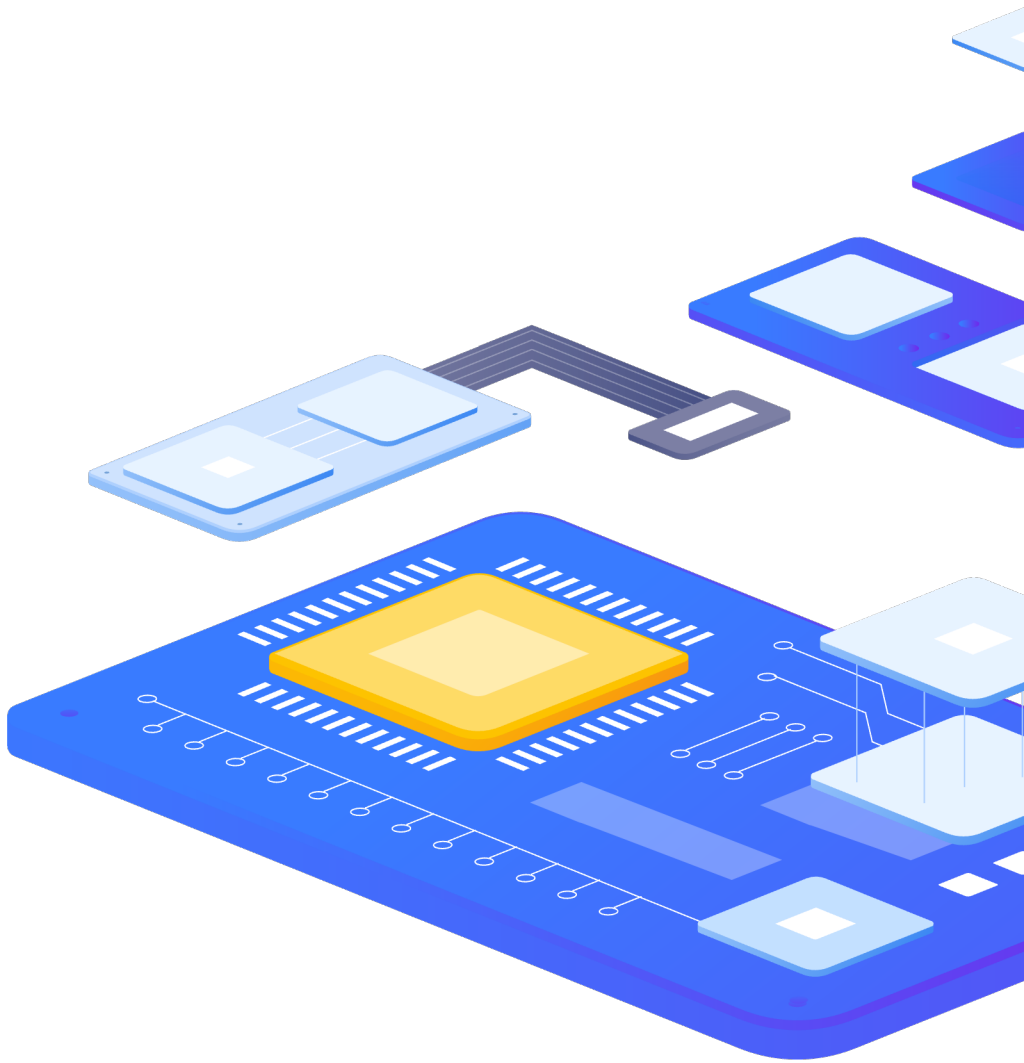

initMAX

Agent2 Plugins - installation

- **Built-in plugins (MySQL, Oracle, Redis)**
  - No additional installation needed

- **External plugins (PostgreSQL, MSSQL, MongoDB)**
  - Linux: Install via package manager or build from source

```
dnf install zabbix-agent2-plugin-mongodb.x86_64 \
        zabbix-agent2-plugin-mssql.x86_64 \
        zabbix-agent2-plugin-postgresql.x86_64
```

  - Windows: Download corresponding plugin package
    - https://cdn.zabbix.com/zabbix/binaries/stable/7.4/latest/
      - zabbix_agent2_plugins-7.4-latest-windows-amd64.msi
      - zabbix_agent2_plugins-7.4-latest-windows-amd64.zip

initMAX

PostgreSQL Installation (Alma Linux 9)

- Add PostgreSQL Repository

```
sudo dnf install -y https://download.postgresql.org/pub/repos/yum/reporpms/EL-9-x86_64/pgdg-redhat-repo-latest.noarch.rpm
```

- Disable built-in PostgreSQL Module and Install Server

```
sudo dnf -y module disable postgresql
# Example for PostgreSQL 18
sudo dnf install -y postgresql18-server postgresql18-contrib
```

- Initialize PostgreSQL Database

```
sudo /usr/pgsql-18/bin/postgresql-18-setup initdb
```

initMAX

PostgreSQL Setup

- Create new user for monitoring purposes

```
sudo -u postgres psql
CREATE USER my_new_user WITH PASSWORD 'my_secure_password';
```

- Create database and grant privileges

```
CREATE DATABASE my_database;
GRANT ALL PRIVILEGES ON DATABASE my_database TO my_new_user;
```

- Grant Schema Permissions

```
\c my_database
GRANT ALL ON SCHEMA public TO my_new_user;
```

initMAX

PostgreSQL Setup

- Create Example Table

```
CREATE TABLE employees (
    id SERIAL PRIMARY KEY,
    name VARCHAR(100) NOT NULL,
    department VARCHAR(50),
    joined_date DATE DEFAULT CURRENT_DATE
);
```

- Add Sample Data

```
INSERT INTO employees (name, department, joined_date) VALUES
('Dmitry', 'IT', '2026-01-15'),
('Elena', 'Marketing', '2026-02-01'),
('Alex', 'Sales', '2026-02-10'),
('Sarah', 'IT', '2026-02-15');
```

initMAX

Agent2 Plugins – configuration

- **Configuration**
  - Check Include directive in zabbix_agent2.conf

```
#Linux
Include=/etc/zabbix/zabbix_agent2.d/plugins.d/*.conf
#Windows
Include=.\zabbix_agent2.d\plugins.d\*.conf
```

  - Set connection parameters
    - either in individual plugin config files or directly in zabbix_agent2.conf
    - or configure in frontend macro and pass with item key call
  - Set custom SQL queries path

```
Plugins.<PluginName>.CustomQueriesPath=<path_to_directory>
```

  - Restart agent after configuration changes

initMAX

Agent2 Plugins – custom SQL queries

- **Create custom queries**
  - create .sql files in the custom queries path
  - insert SQL query content into file
  - can use parameters using "$1" syntax

```
#chunks_size.sql
SELECT SUM(total_bytes) FROM chunks_detailed_size($1);
```

  - must restart agent after adding new .sql files

- **Using custom queries**
  - Item key

```
<plugin>.custom.query[uri,<username>,<password>,queryName,<args...>]
```

initMAX

Agent2 Plugins – custom SQL queries

- **Set required authentication parameters**
  - Don't forget to use Secret text

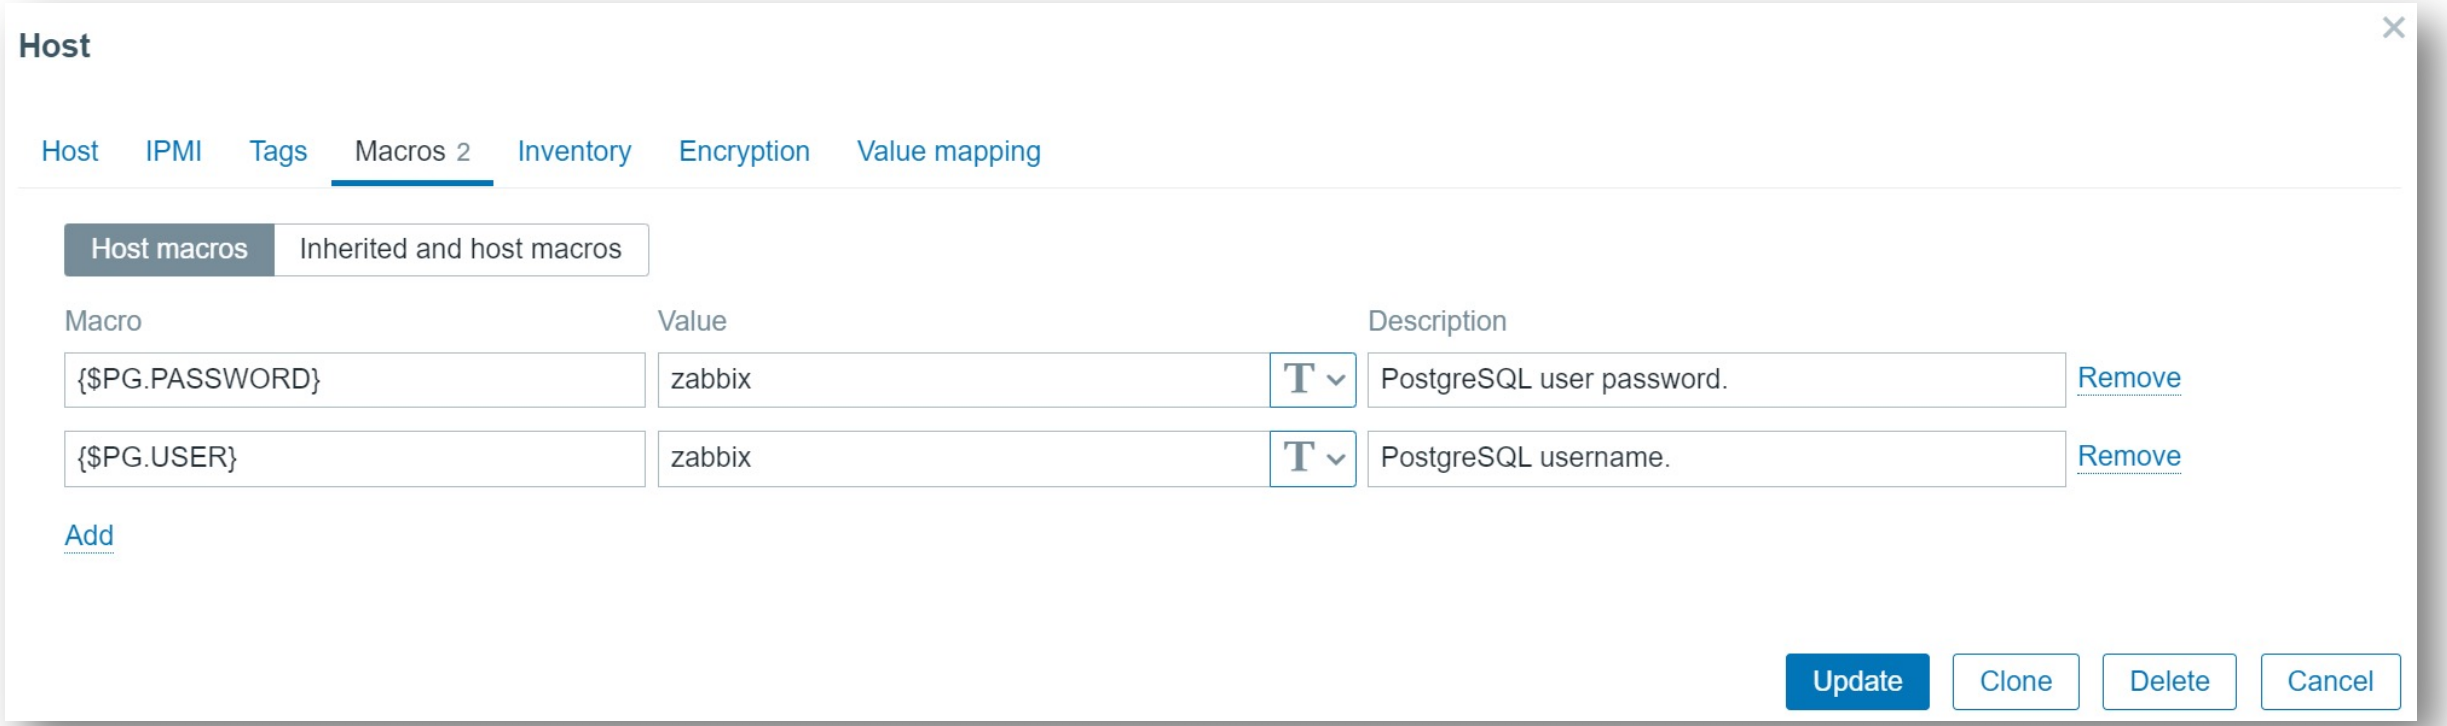

initMAX

ODBC Monitoring

- **Direct SQL-based monitoring**
  - Open Database Connectivity (ODBC) standard
  - C language middleware API for accessing database management systems (DBMS)
  - Execute custom queries for precise metrics
- **Wide database compatibility**
  - Zabbix may query any database, which is supported by ODBC
  - MySQL, PostgreSQL, Microsoft SQL Server, Oracle, and others
- **Implementation requirements**
  - ODBC drivers on Zabbix server/proxy
  - DSN configuration and proper permissions
  - Item keys using db.odbc.select[] syntax

initMAX

ODBC Monitoring - installation

- First install unixODBC, then specific DB driver

```
dnf install unixODBC unixODBC-devel postgresql-odbc
```

- Configure installed ODBC drivers in /etc/odbcinst.ini

```
[PostgreSQL_driver]
Description = PostgreSQL ODBC driver
Driver = /usr/lib/odbc/psqlodbca.so
Setup = /usr/lib/odbc/libodbcpsqlS.so
```

- Configure DSN in /etc/odbc.ini

```
[PostgreSQL-Zabbix]
Description = PostgreSQL connection
Driver = PostgreSQL_driver
Database = Zabbix
Servername = localhost
UserName = Zabbix
Password = password
Port = 5432
```

initMAX

ODBC Monitoring - installation

- Make sure that Driver name actually matches
- Different Drivers behave differently
  - MSSQL Does not support Username/Password in Config File

- Verify ODBC Configuration with **isql**

```
[root@vbox ~]# isql PostgreSQL-Zabbix
+---------------------------------------+
| Connected!                            |
|                                       |
| sql-statement                         |
| help [tablename]                      |
| quit                                  |
|                                       |
+---------------------------------------+
```

initMAX

ODBC Monitoring – usage example

- Item type
  - Database monitor
- Item key

```
#db.odbc.select[<unique short description>,<dsn>,<connection string>]
db.odbc.select[pgsql_active,PostgreSQL-Zabbix]
```

- SQL query

```
SELECT count(*) FROM pg_stat_activity;
```

- Detailed guide
  https://www.zabbix.com/documentation/current/en/manual/config/items/itemtypes/odbc_checks

initMAX

ODBC Monitoring – Supported Keys

- Simple Select Key

```
db.odbc.select[] - this item returns one value (the first column of the first row of the
SQL query result);
```

- Key for JSON Data

```
db.odbc.get[] - this item returns multiple rows/columns in JSON format;
```

- Key For Low Level Discovery Data

```
db.odbc.discovery[] - this item returns low-level discovery data.
```

initMAX

Alternative Monitoring Methods

- **JMX Agent**
  - For Java-based DBs (Cassandra, Neo4j)
  - Requires Zabbix Java Gateway
  - Items using jmx[] key format
- **HTTP Agent**
  - REST API monitoring (Elasticsearch, CouchDB)
  - No agent needed, monitored by server/proxy
  - JSON/XML data processing
- **Simple Checks**
  - Basic availability monitoring
  - Standard port checks
    - MySQL (3306), PostgreSQL (5432) › MS SQL (1433), Oracle (1521), MongoDB (27017)
  - Item keys: net.tcp.service[], net.tcp.port[]

initMAX

Custom Monitoring Approaches

- **User Parameters**
  - Defined in agent configuration files
  - Agent restart required after changes
- **External Checks**
  - Scripts executed by Zabbix server/proxy
  - No agent required on monitored host

- **Example: SQLite monitoring**
  - File-based database without client-server architecture
  - Advanced approach - create shell script in /etc/zabbix/scripts/
  - Make scripts executable (chmod +x)
  - Reference in UserParameters

```
UserParameter=sqlite.stats[*],/etc/zabbix/scripts/sqlite_stats.sh "$1"
```

3

Native Database Monitoring in Zabbix

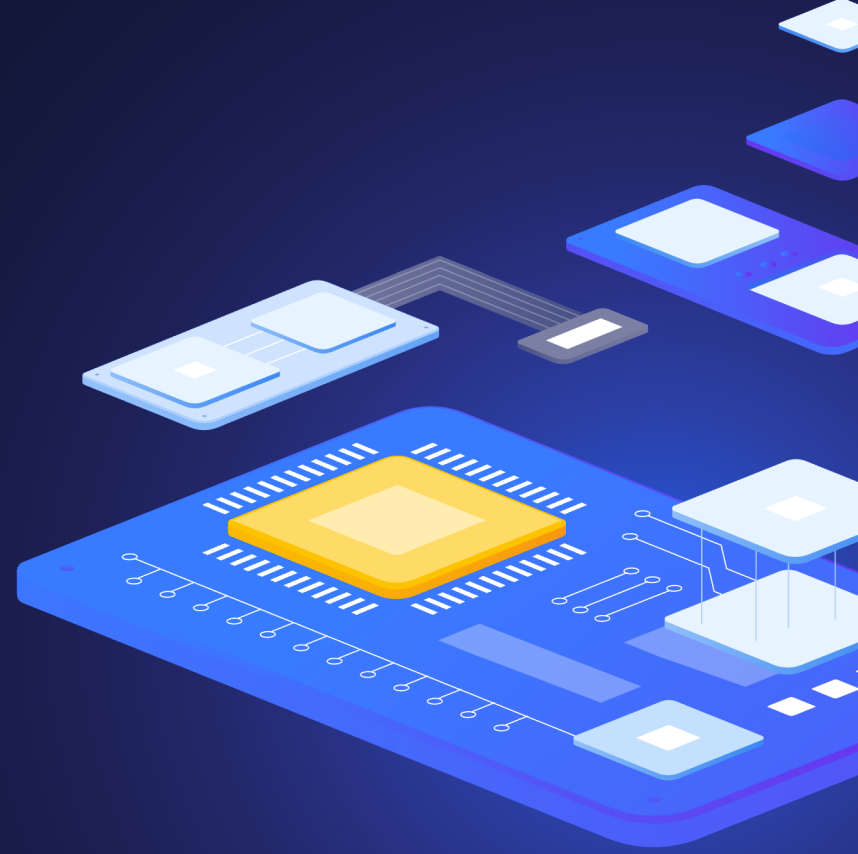

initMAX

Native Database Monitoring in Zabbix

- **Zabbix database integrations**
  - https://www.zabbix.com/integrations?cat=databases

- **PostgreSQL (+TimescaleDB)**
  - Agent2 plugin (requires installation)
    - PostgreSQL by Zabbix agent 2 (active)
  - Agent (requires user parameters)
    - PostgreSQL by Zabbix agent
  - ODBC monitoring
    - PostgreSQL by ODBC
  - Metrics: Discovers databases, monitors size, replication, checkpoints, WAL logs, etc.

initMAX

Native Database Monitoring in Zabbix

- **MySQL/MariaDB/Percona**
  - Agent2 plugin (built-in)
    - MySQL by Zabbix agent 2
  - Zabbix Agent (agentd)
    - MySQL by Zabbix agent
  - ODBC monitoring
    - MySQL by ODBC
  - Metrics: Connections, queries/sec, buffer usage, InnoDB metrics, etc.

initMAX

Native Database Monitoring in Zabbix

- **Oracle Database**
  - Agent2 plugin (built-in), ODBC monitoring
  - Tablespace usage, sessions, cache ratios, locks, etc.
- **Microsoft SQL Server**
  - Agent2 plugin (requires installation), ODBC monitoring
  - Buffer cache, page life expectancy, lock waits
- **MongoDB**
  - Agent2 plugin (requires installation), HTTP Agent
  - Operations count, connections, memory usage, replication
- **Redis**
  - Agent2 plugin (built-in)
  - Commands processed, memory usage, connections, persistence, etc.

4

Case study

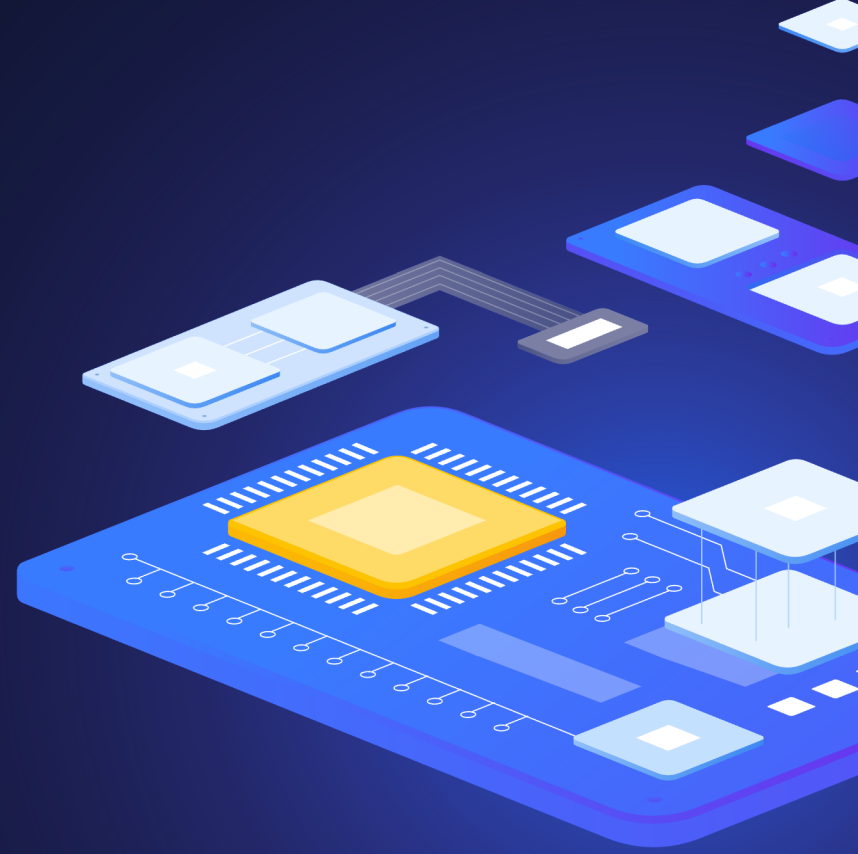

initMAX

Case Study - PostgreSQL Patroni cluster

- **Infrastructure**
  - PostgreSQL on Patroni (2-node + witness)
  - High-availability cluster with etcd
  - F5 load balancer
    - Patroni API integration
    - Automated failover testing
- **Scale**
  - ~6.7K hosts monitored (6.6K enabled)
  - ~1.2M items, ~530K triggers
  - ~15K values per second throughput

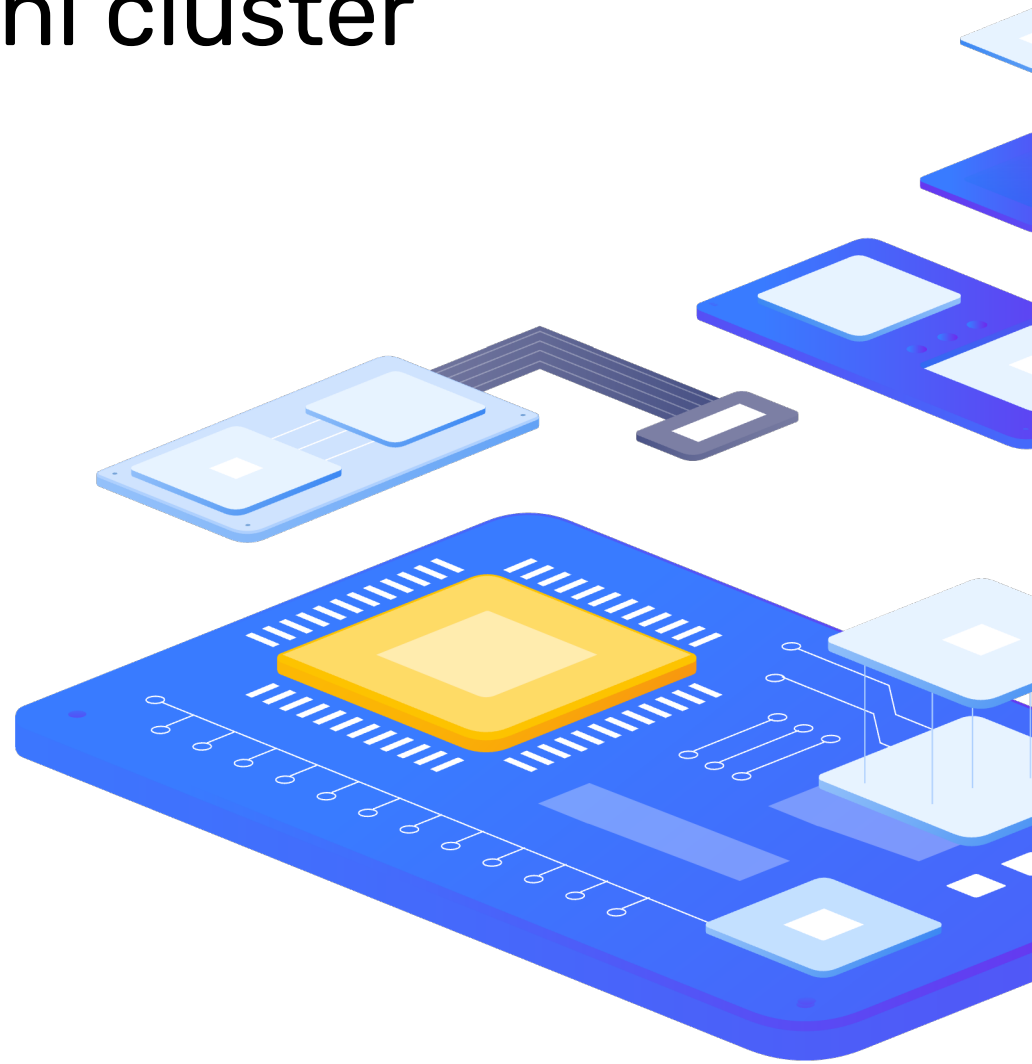

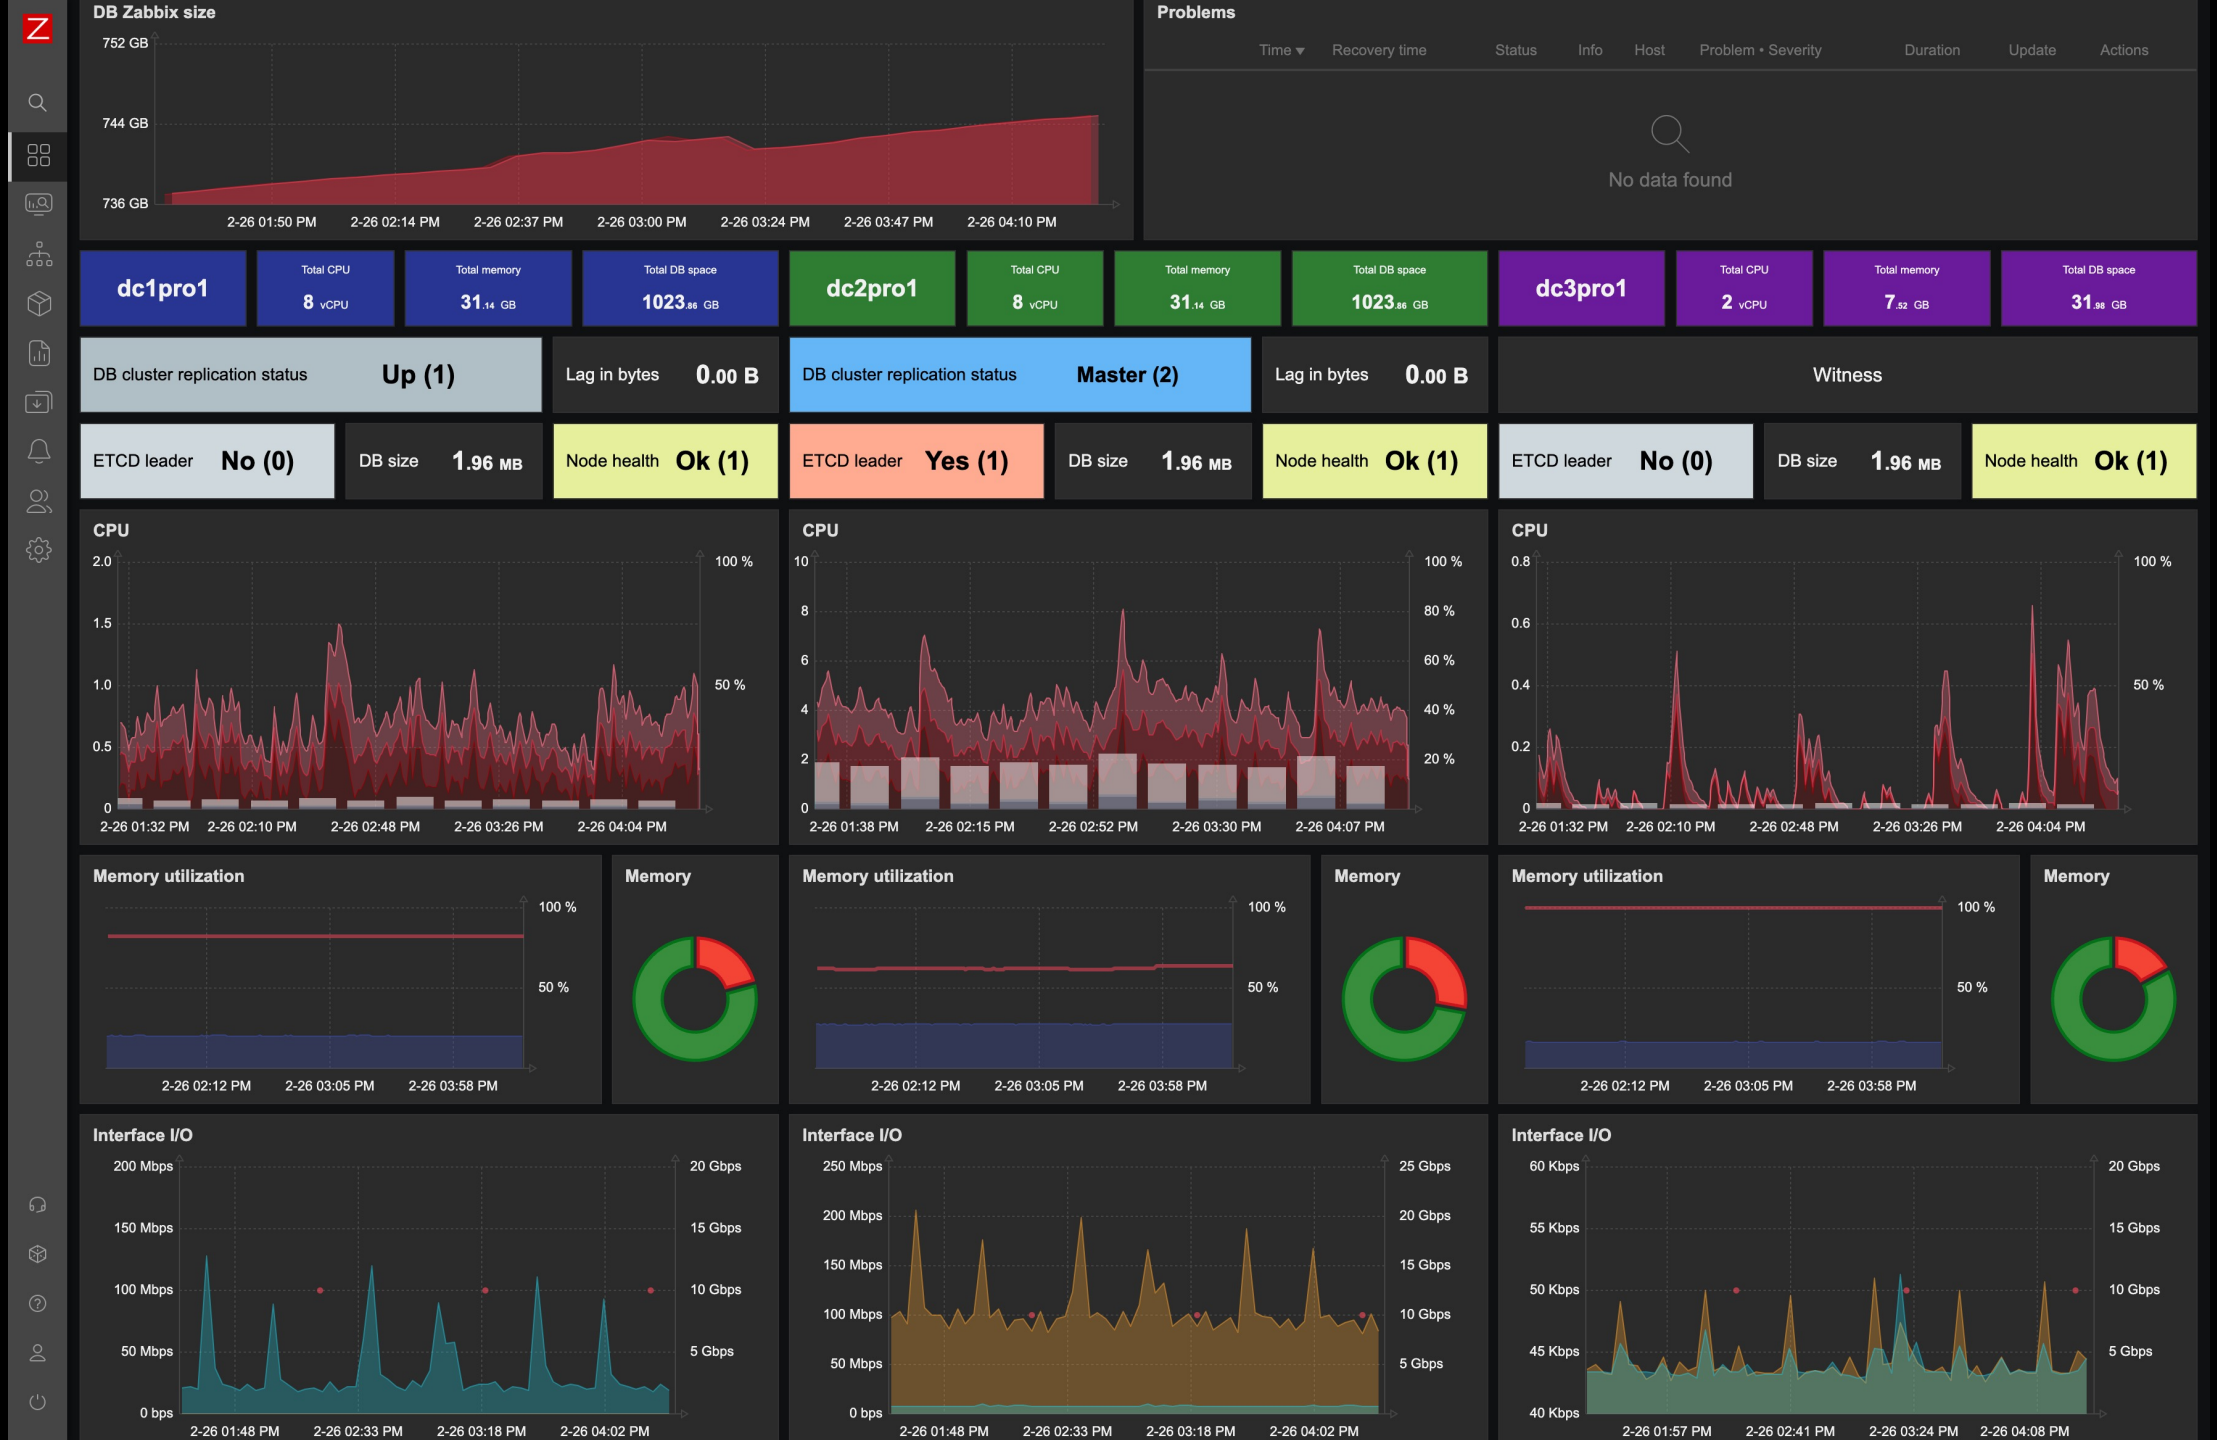
DB Zabbix size
752 GB
744 GB
736 GB
2-26 01:50 PM
2-26 02:14 PM
2-26 02:37 PM
2-26 03:00 PM
2-26 03:24 PM
2-26 03:47 PM
2-26 04:10 PM
Problems
Time
Recovery time
Status
Info
Host
Problem • Severity
Duration
Update
Actions
No data found
dc1pro1
Total CPU
8 vCPU
Total memory
31.14 GB
Total DB space
1023.86 GB
dc2pro1
Total CPU
8 vCPU
Total memory
31.14 GB
Total DB space
1023.86 GB
dc3pro1
Total CPU
2 vCPU
Total memory
7.52 GB
Total DB space
31.98 GB
DB cluster replication status
Up (1)
Lag in bytes
0.00 B
DB cluster replication status
Master (2)
Lag in bytes
0.00 B
Witness
ETCD leader
No (0)
DB size
1.96 MB
Node health
Ok (1)
ETCD leader
Yes (1)
DB size
1.96 MB
Node health
Ok (1)
ETCD leader
No (0)
DB size
1.96 MB
Node health
Ok (1)
CPU
2.0
1.5
1.0
0.5
0
100 %
50 %
2-26 01:32 PM
2-26 02:10 PM
2-26 02:48 PM
2-26 03:26 PM
2-26 04:04 PM
CPU
10
8
6
4
2
0
100 %
80 %
60 %
40 %
20 %
2-26 01:38 PM
2-26 02:15 PM
2-26 02:52 PM
2-26 03:30 PM
2-26 04:07 PM
CPU
0.8
0.6
0.4
0.2
0
100 %
50 %
2-26 01:32 PM
2-26 02:10 PM
2-26 02:48 PM
2-26 03:26 PM
2-26 04:04 PM
Memory utilization
100 %
50 %
2-26 02:12 PM
2-26 03:05 PM
2-26 03:58 PM
Memory
Memory utilization
100 %
50 %
2-26 02:12 PM
2-26 03:05 PM
2-26 03:58 PM
Memory
Memory utilization
100 %
50 %
2-26 02:12 PM
2-26 03:05 PM
2-26 03:58 PM
Memory
Interface I/O
200 Mbps
150 Mbps
100 Mbps
50 Mbps
0 bps
20 Gbps
15 Gbps
10 Gbps
5 Gbps
2-26 01:48 PM
2-26 02:33 PM
2-26 03:18 PM
2-26 04:02 PM
Interface I/O
250 Mbps
200 Mbps
150 Mbps
100 Mbps
50 Mbps
0 bps
25 Gbps
20 Gbps
15 Gbps
10 Gbps
5 Gbps
2-26 01:48 PM
2-26 02:33 PM
2-26 03:18 PM
2-26 04:02 PM
Interface I/O
60 Kbps
55 Kbps
50 Kbps
45 Kbps
40 Kbps
20 Gbps
15 Gbps
10 Gbps
5 Gbps
2-26 01:57 PM
2-26 02:41 PM
2-26 03:24 PM
2-26 04:08 PM

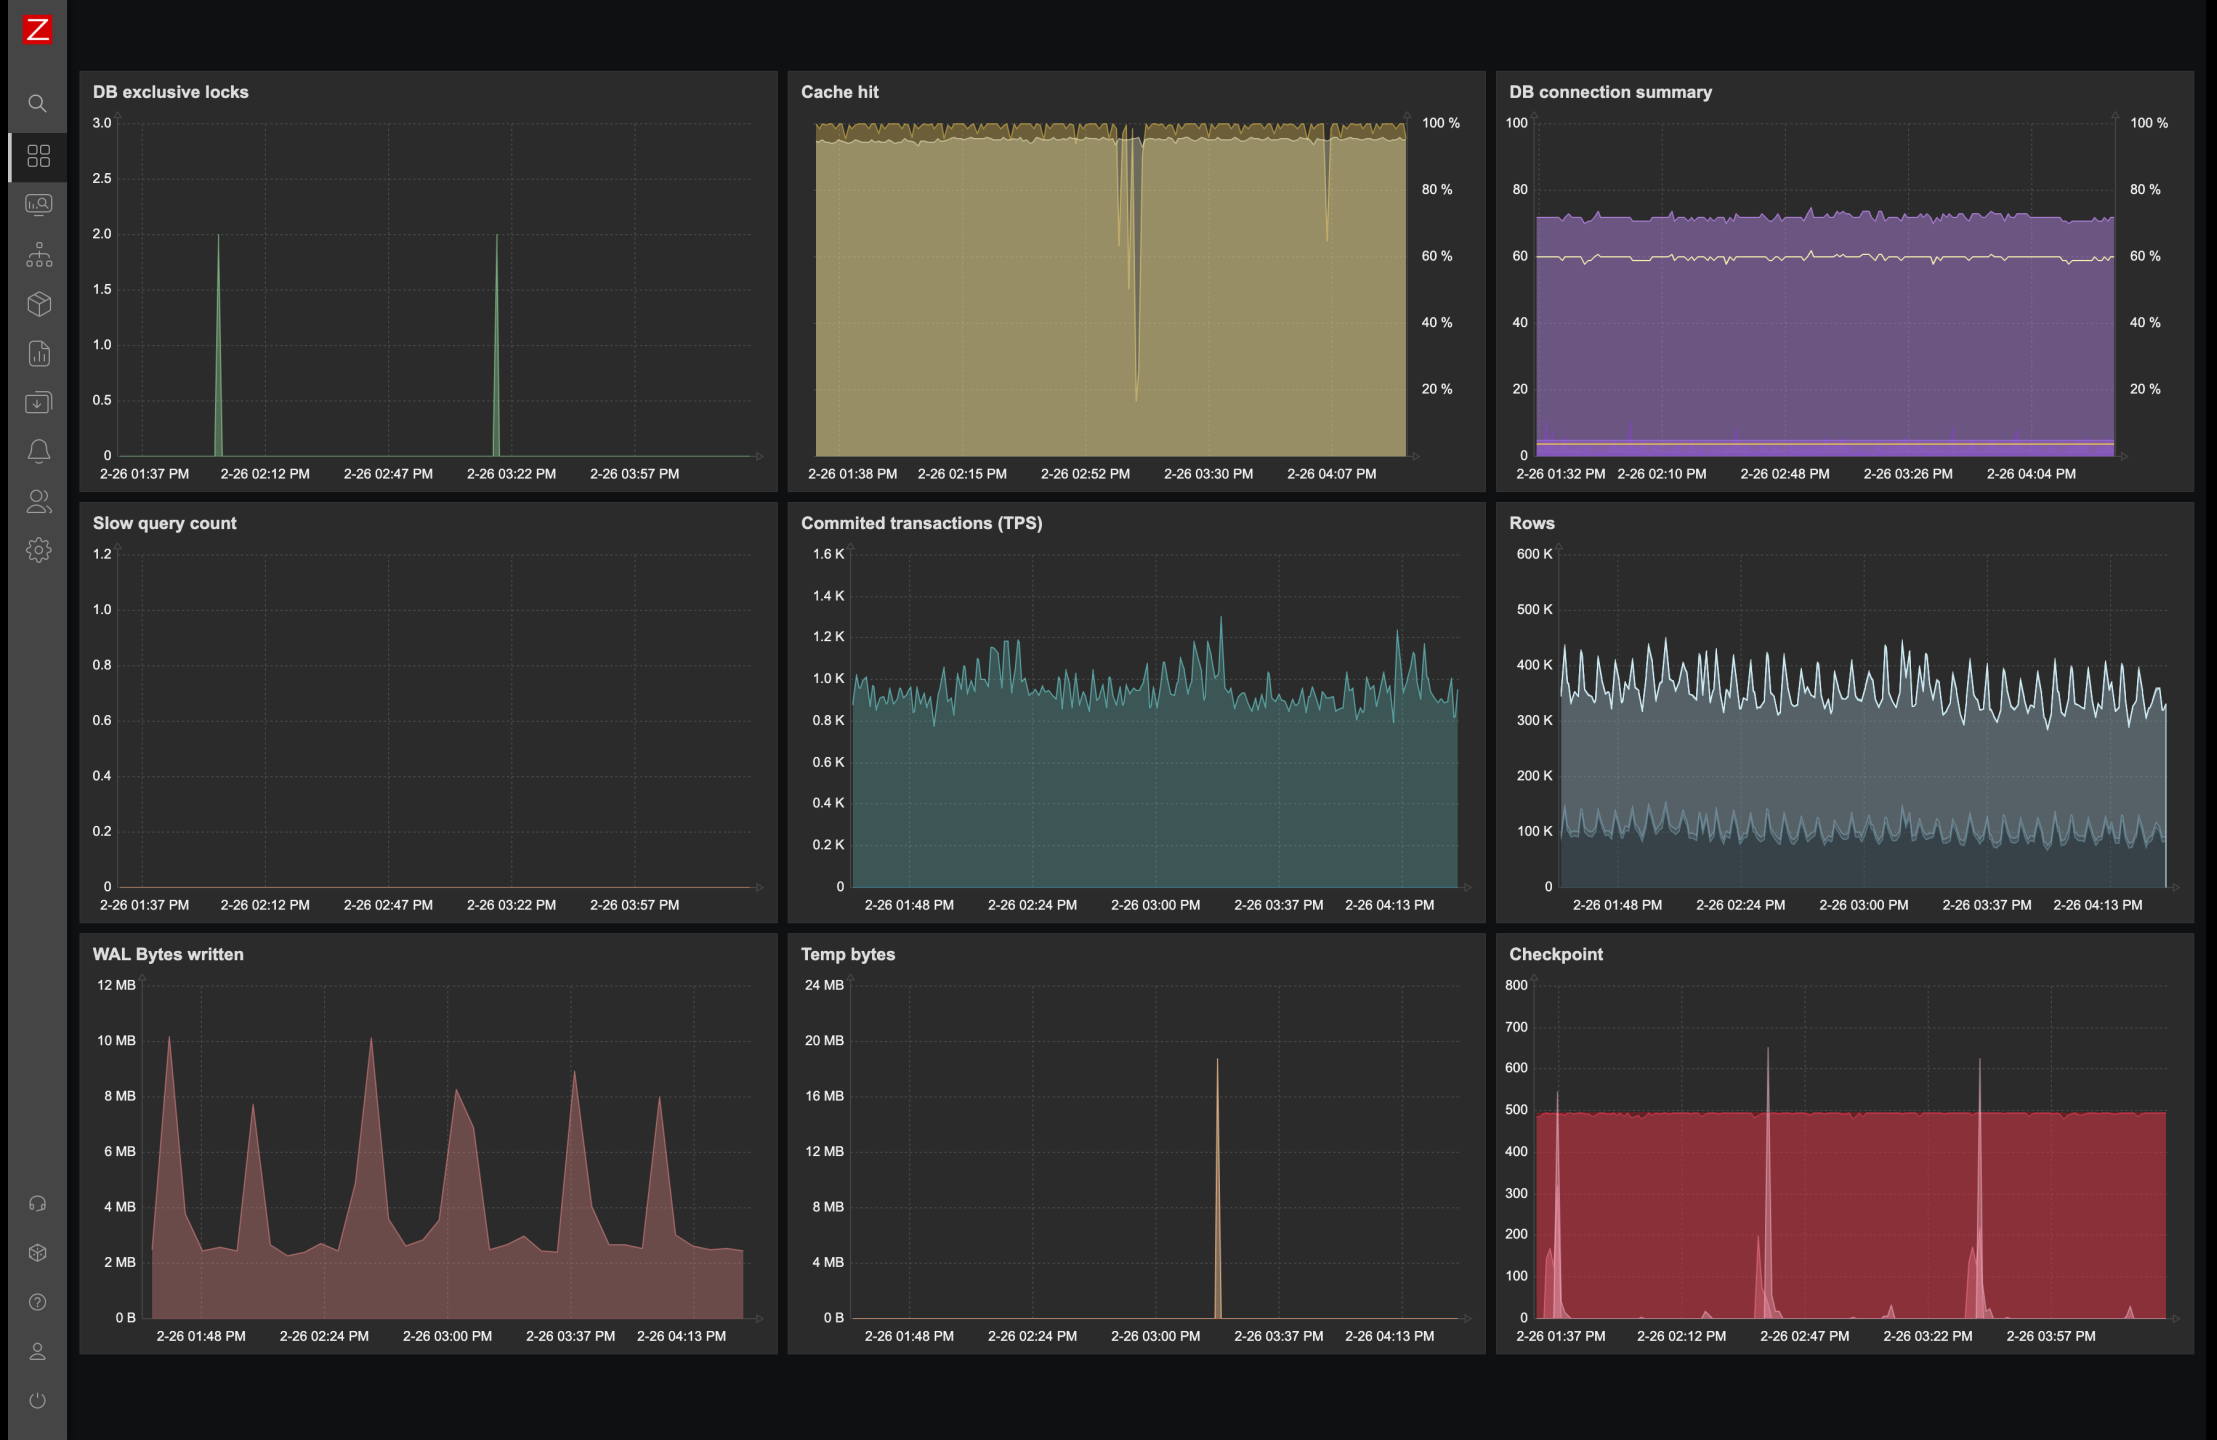
DB exclusive locks
3.0
2.5
2.0
1.5
1.0
0.5
0
2-26 01:37 PM
2-26 02:12 PM
2-26 02:47 PM
2-26 03:22 PM
2-26 03:57 PM
Cache hit
100 %
80 %
60 %
40 %
20 %
2-26 01:38 PM
2-26 02:15 PM
2-26 02:52 PM
2-26 03:30 PM
2-26 04:07 PM
DB connection summary
100
80
60
40
20
0
100 %
80 %
60 %
40 %
20 %
2-26 01:32 PM
2-26 02:10 PM
2-26 02:48 PM
2-26 03:26 PM
2-26 04:04 PM
Slow query count
1.2
1.0
0.8
0.6
0.4
0.2
0
2-26 01:37 PM
2-26 02:12 PM
2-26 02:47 PM
2-26 03:22 PM
2-26 03:57 PM
Commited transactions (TPS)
1.6 K
1.4 K
1.2 K
1.0 K
0.8 K
0.6 K
0.4 K
0.2 K
0
2-26 01:48 PM
2-26 02:24 PM
2-26 03:00 PM
2-26 03:37 PM
2-26 04:13 PM
Rows
600 K
500 K
400 K
300 K
200 K
100 K
0
2-26 01:48 PM
2-26 02:24 PM
2-26 03:00 PM
2-26 03:37 PM
2-26 04:13 PM
WAL Bytes written
12 MB
10 MB
8 MB
6 MB
4 MB
2 MB
0 B
2-26 01:48 PM
2-26 02:24 PM
2-26 03:00 PM
2-26 03:37 PM
2-26 04:13 PM
Temp bytes
24 MB
20 MB
16 MB
12 MB
8 MB
4 MB
0 B
2-26 01:48 PM
2-26 02:24 PM
2-26 03:00 PM
2-26 03:37 PM
2-26 04:13 PM
Checkpoint
800
700
600
500
400
300
200
100
0
2-26 01:37 PM
2-26 02:12 PM
2-26 02:47 PM
2-26 03:22 PM
2-26 03:57 PM

DB exclusive locks

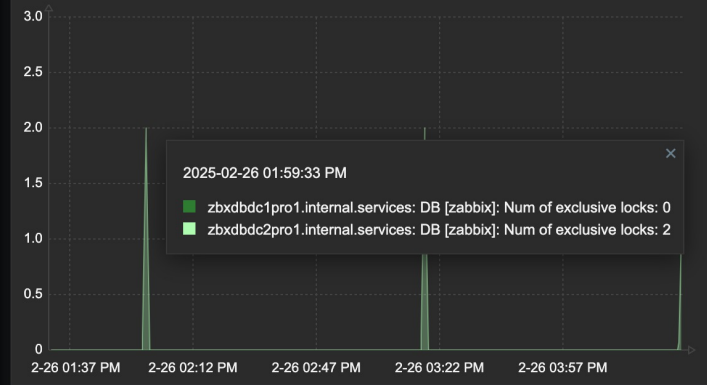

Cache hit

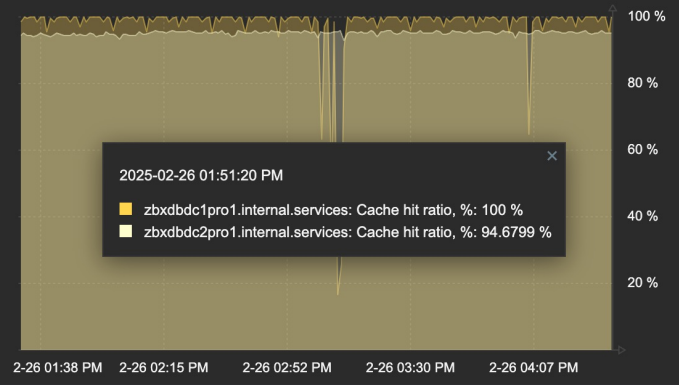

DB connection summary

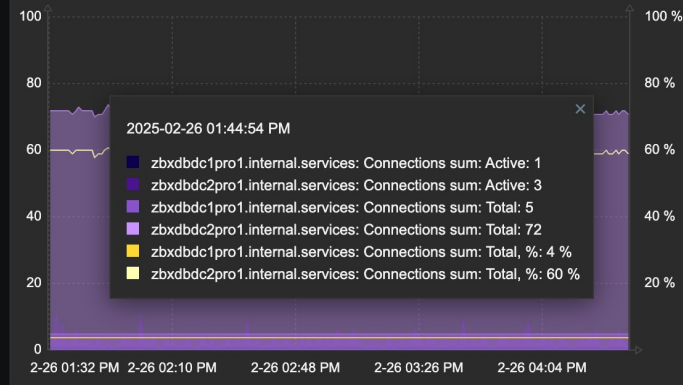

Slow query count

Commited transactions (TPS)

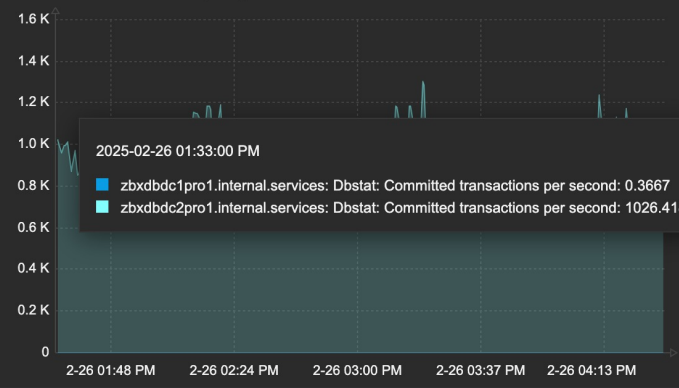

Rows

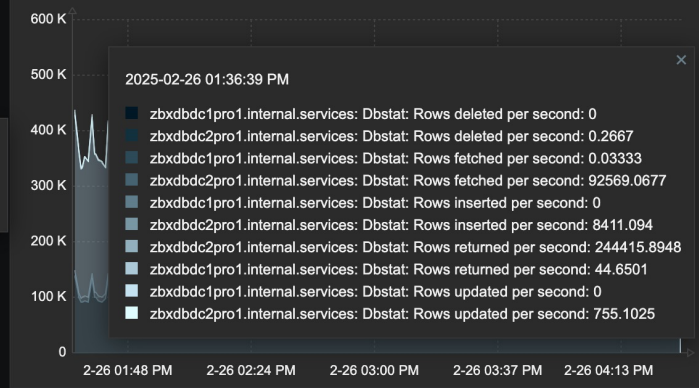

WAL Bytes written

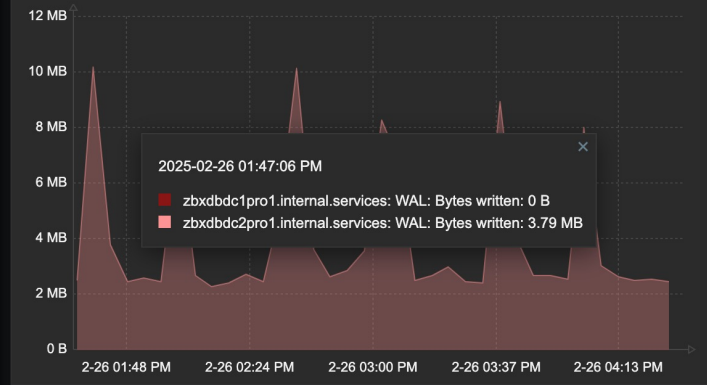

Temp bytes

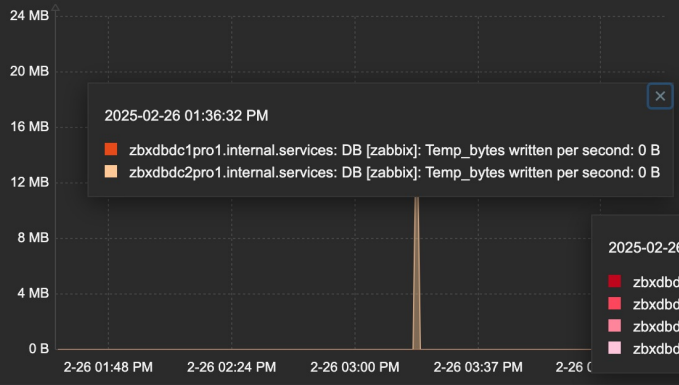

Checkpoint

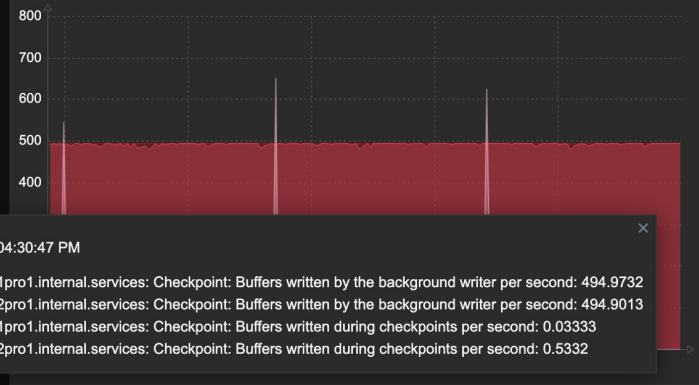

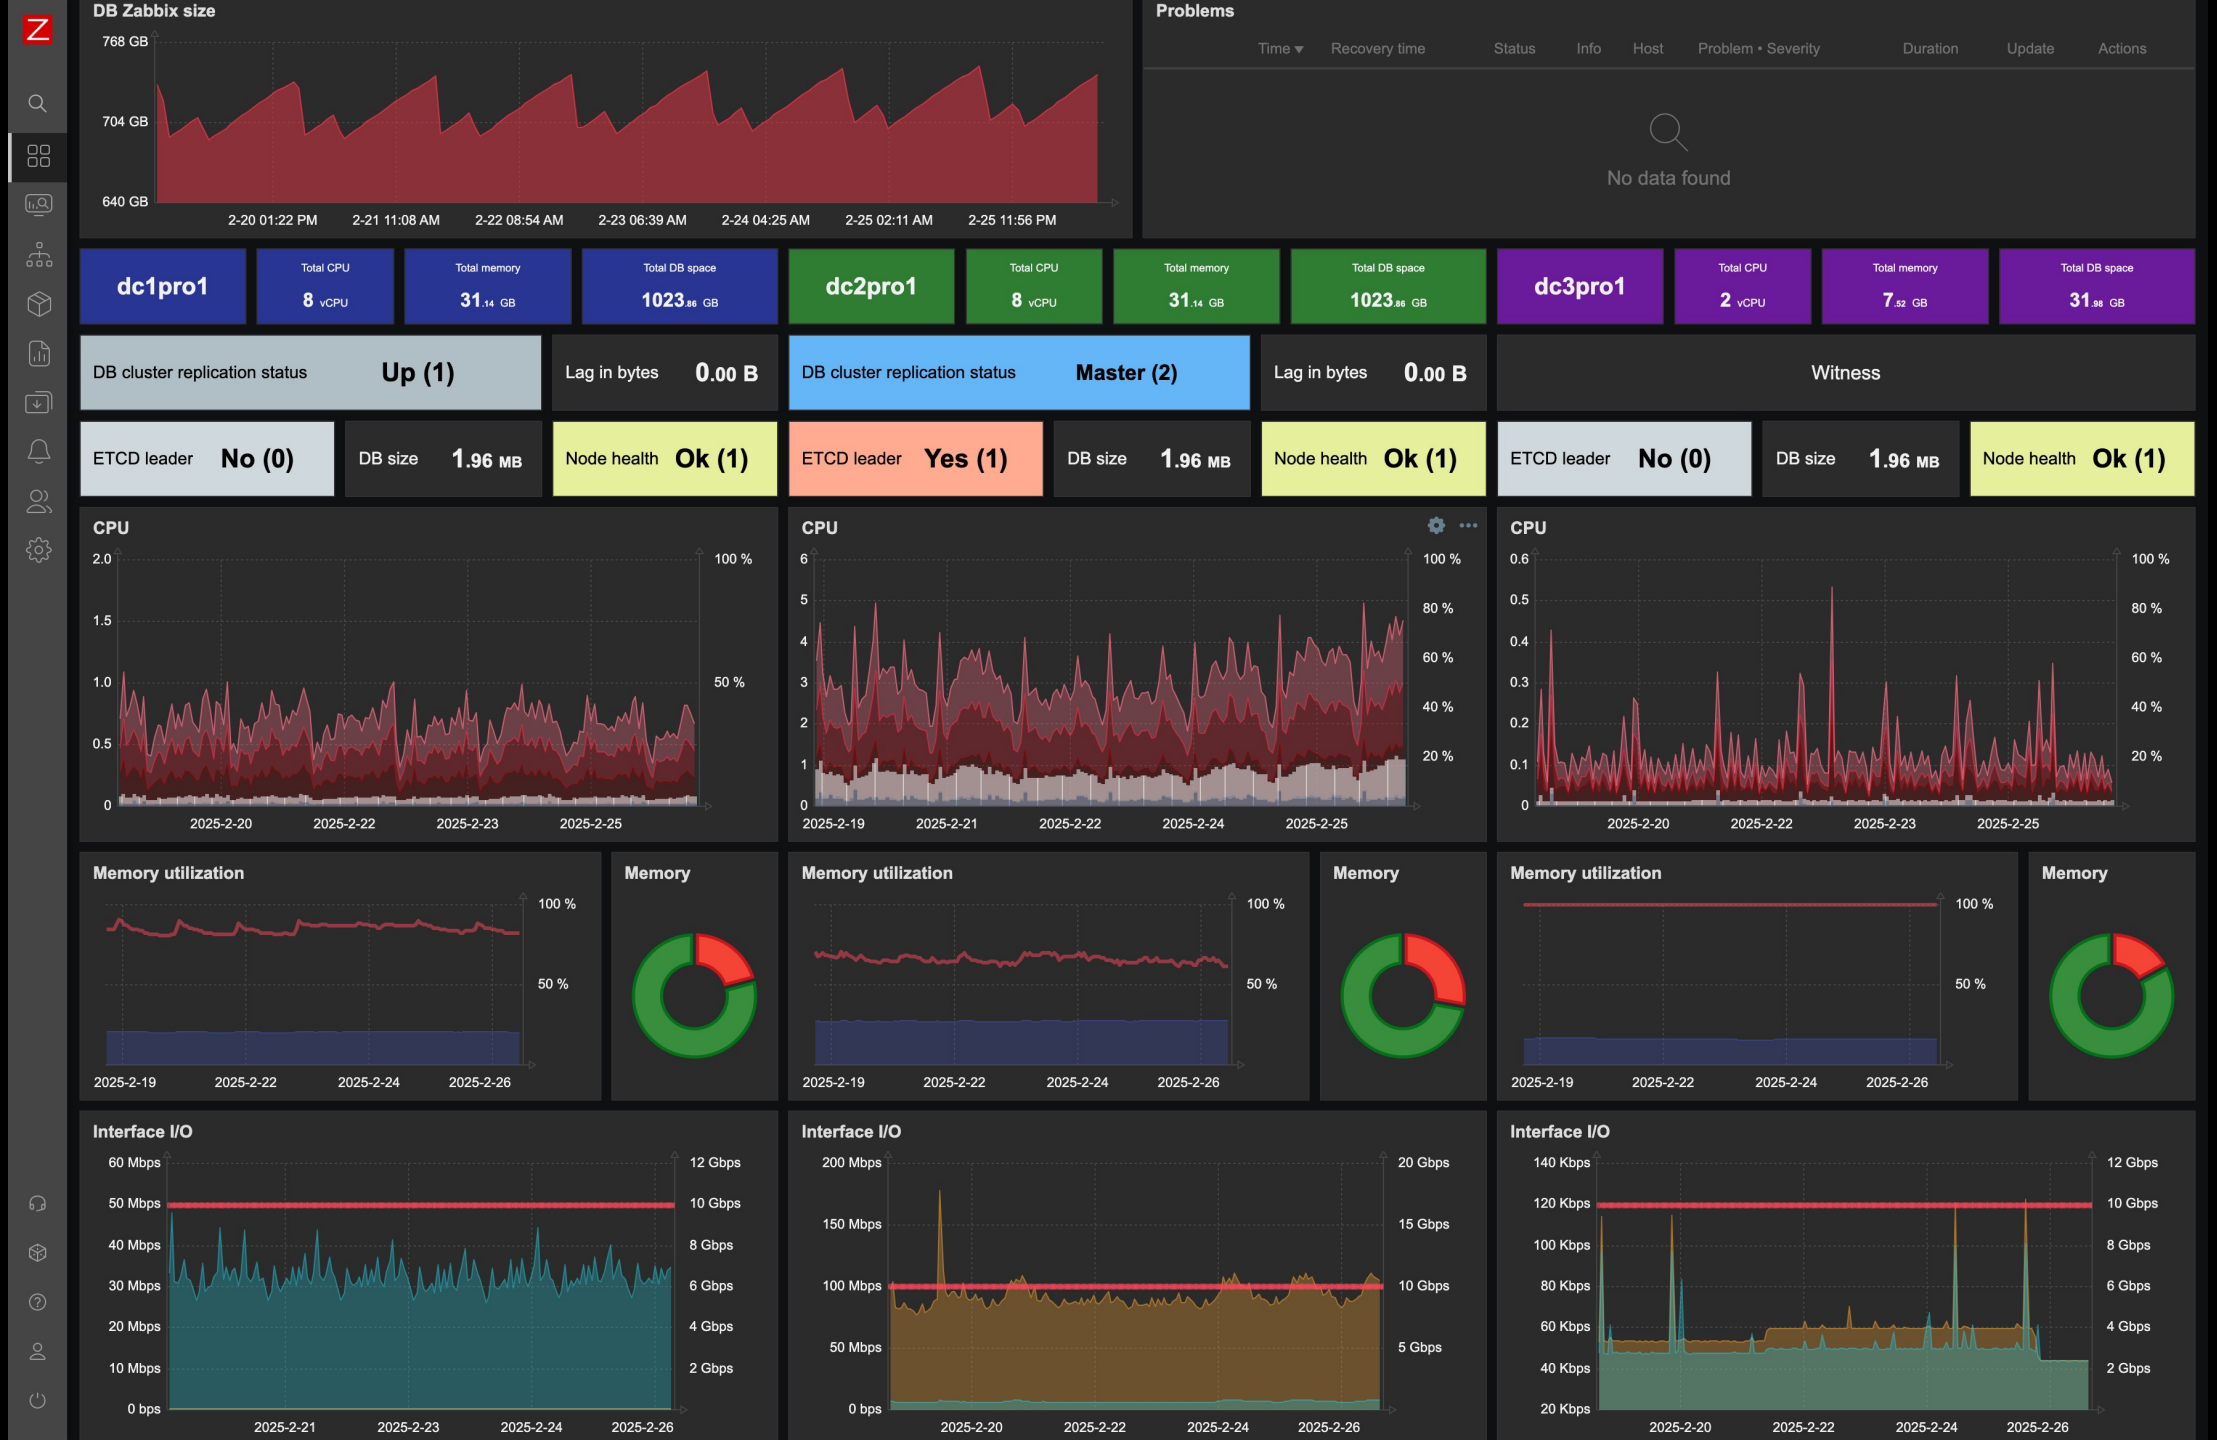
DB Zabbix size
768 GB
704 GB
640 GB
2-20 01:22 PM
2-21 11:08 AM
2-22 08:54 AM
2-23 06:39 AM
2-24 04:25 AM
2-25 02:11 AM
2-25 11:56 PM
Problems
Time
Recovery time
Status
Info
Host
Problem • Severity
Duration
Update
Actions
No data found
dc1pro1
Total CPU
8 vCPU
Total memory
31.14 GB
Total DB space
1023.86 GB
dc2pro1
Total CPU
8 vCPU
Total memory
31.14 GB
Total DB space
1023.86 GB
dc3pro1
Total CPU
2 vCPU
Total memory
7.52 GB
Total DB space
31.98 GB
DB cluster replication status
Up (1)
Lag in bytes
0.00 B
DB cluster replication status
Master (2)
Lag in bytes
0.00 B
Witness
ETCD leader
No (0)
DB size
1.96 MB
Node health
Ok (1)
ETCD leader
Yes (1)
DB size
1.96 MB
Node health
Ok (1)
ETCD leader
No (0)
DB size
1.96 MB
Node health
Ok (1)
CPU
2.0
1.5
1.0
0.5
0
100 %
50 %
2025-2-20
2025-2-22
2025-2-23
2025-2-25
CPU
6
5
4
3
2
1
0
100 %
80 %
60 %
40 %
20 %
2025-2-19
2025-2-21
2025-2-22
2025-2-24
2025-2-25
CPU
0.6
0.5
0.4
0.3
0.2
0.1
0
100 %
80 %
60 %
40 %
20 %
2025-2-20
2025-2-22
2025-2-23
2025-2-25
Memory utilization
100 %
50 %
2025-2-19
2025-2-22
2025-2-24
2025-2-26
Memory
Memory utilization
100 %
50 %
2025-2-19
2025-2-22
2025-2-24
2025-2-26
Memory
Memory utilization
100 %
50 %
2025-2-19
2025-2-22
2025-2-24
2025-2-26
Memory
Interface I/O
60 Mbps
50 Mbps
40 Mbps
30 Mbps
20 Mbps
10 Mbps
0 bps
12 Gbps
10 Gbps
8 Gbps
6 Gbps
4 Gbps
2 Gbps
2025-2-21
2025-2-23
2025-2-24
2025-2-26
Interface I/O
200 Mbps
150 Mbps
100 Mbps
50 Mbps
0 bps
20 Gbps
15 Gbps
10 Gbps
5 Gbps
2025-2-20
2025-2-22
2025-2-24
2025-2-26
Interface I/O
140 Kbps
120 Kbps
100 Kbps
80 Kbps
60 Kbps
40 Kbps
20 Kbps
12 Gbps
10 Gbps
8 Gbps
6 Gbps
4 Gbps
2 Gbps
2025-2-20
2025-2-22
2025-2-24
2025-2-26

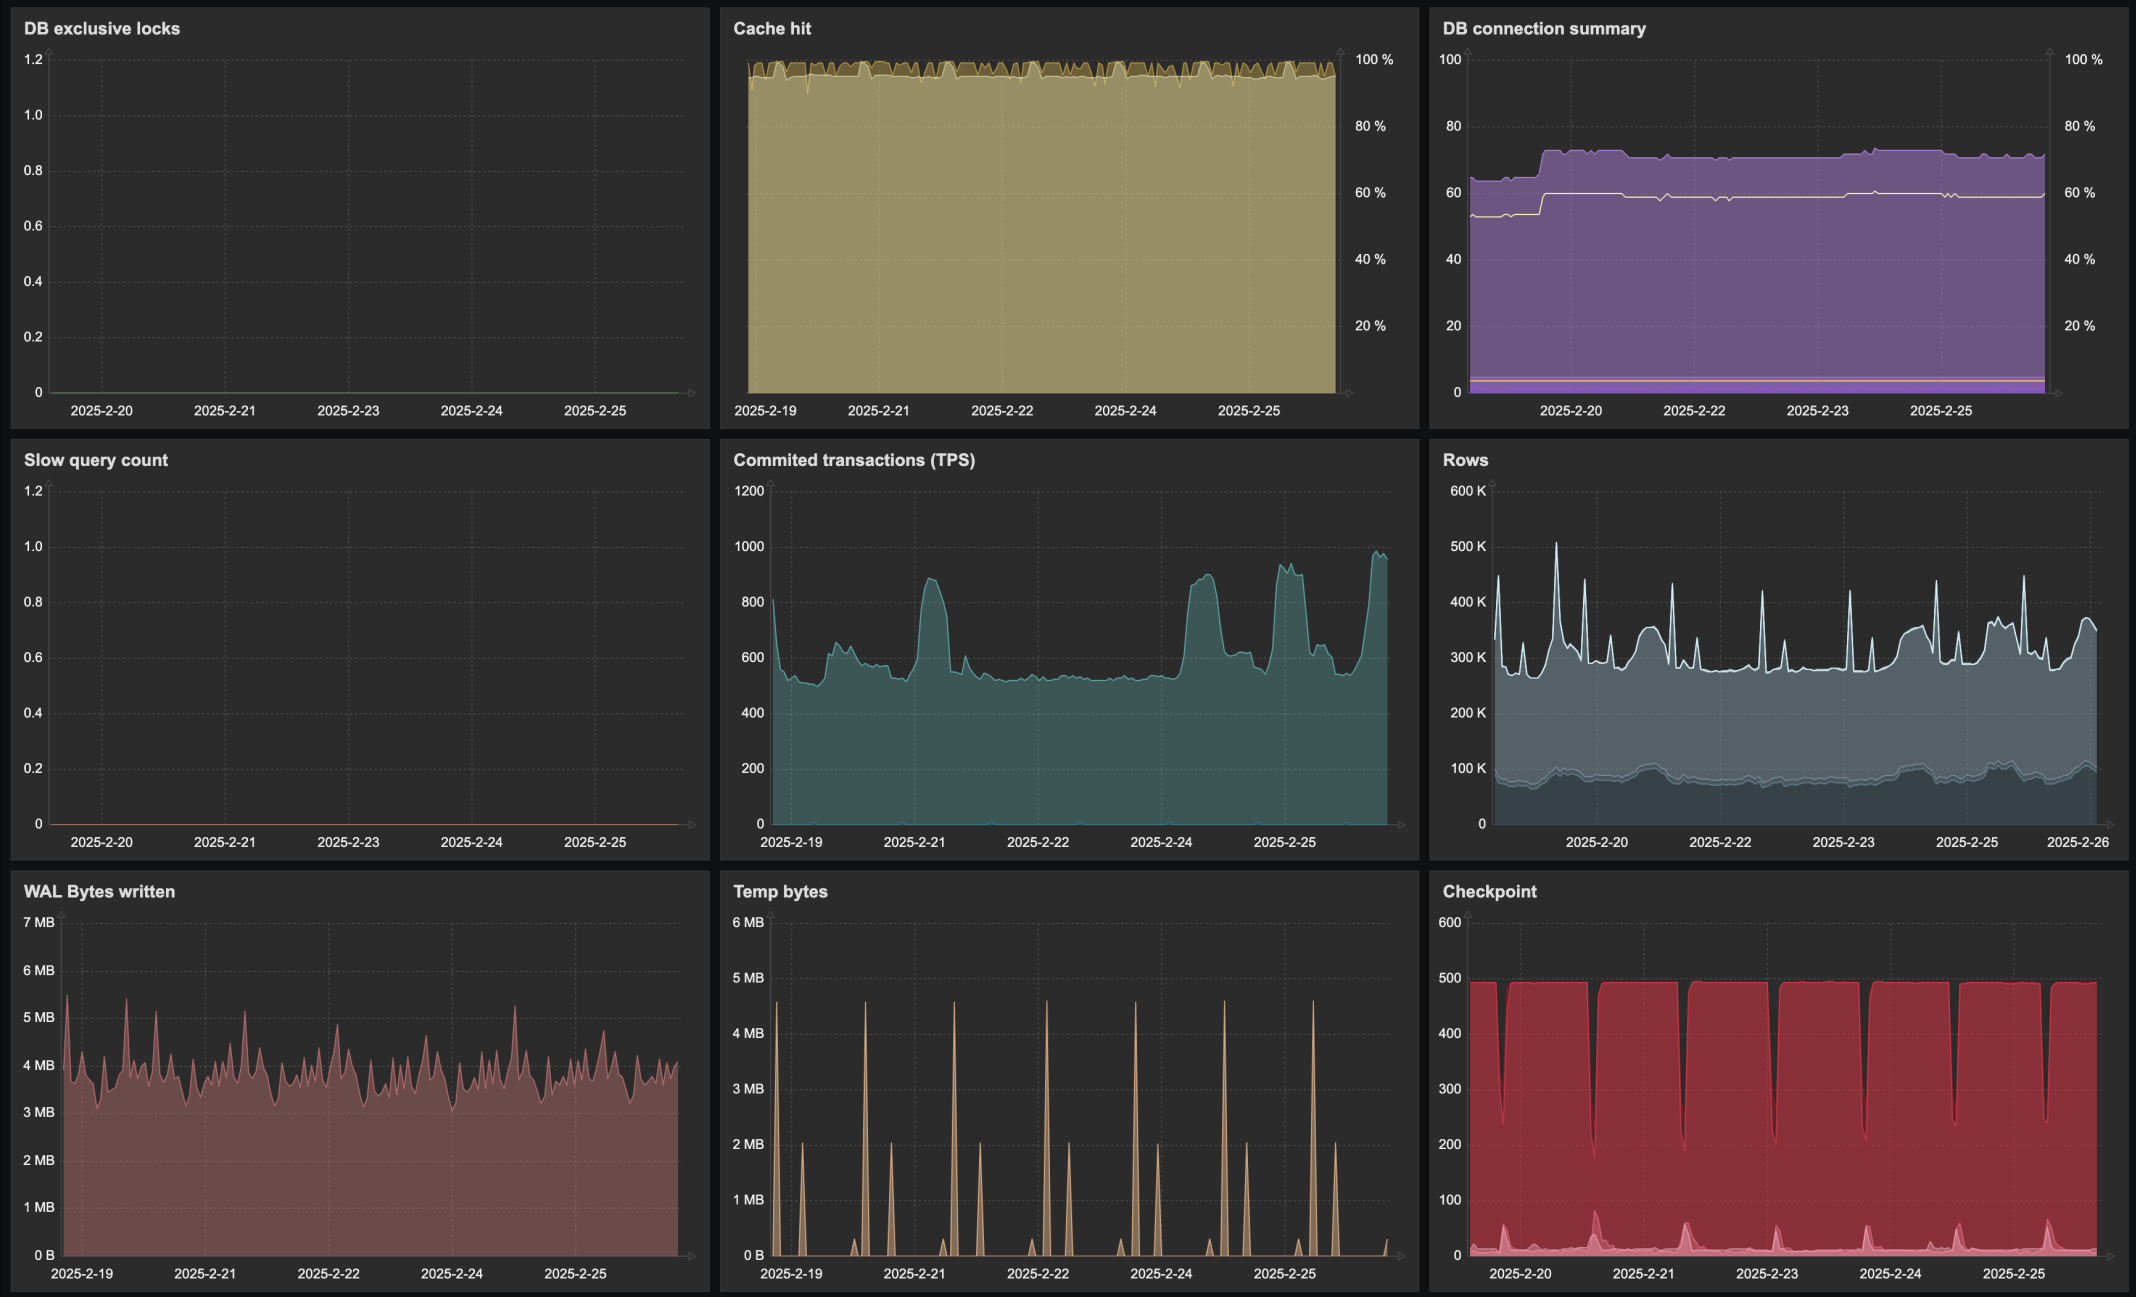
DB exclusive locks
1.2
1.0
0.8
0.6
0.4
0.2
0
2025-2-20
2025-2-21
2025-2-23
2025-2-24
2025-2-25
Cache hit
100 %
80 %
60 %
40 %
20 %
2025-2-19
2025-2-21
2025-2-22
2025-2-24
2025-2-25
DB connection summary
100
80
60
40
20
0
100 %
80 %
60 %
40 %
20 %
2025-2-20
2025-2-22
2025-2-23
2025-2-25
Slow query count
1.2
1.0
0.8
0.6
0.4
0.2
0
2025-2-20
2025-2-21
2025-2-23
2025-2-24
2025-2-25
Commited transactions (TPS)
1200
1000
800
600
400
200
0
2025-2-19
2025-2-21
2025-2-22
2025-2-24
2025-2-25
Rows
600 K
500 K
400 K
300 K
200 K
100 K
0
2025-2-20
2025-2-22
2025-2-23
2025-2-25
2025-2-26
WAL Bytes written
7 MB
6 MB
5 MB
4 MB
3 MB
2 MB
1 MB
0 B
2025-2-19
2025-2-21
2025-2-22
2025-2-24
2025-2-25
Temp bytes
6 MB
5 MB
4 MB
3 MB
2 MB
1 MB
0 B
2025-2-19
2025-2-21
2025-2-22
2025-2-24
2025-2-25
Checkpoint
600
500
400
300
200
100
0
2025-2-20
2025-2-21
2025-2-23
2025-2-24
2025-2-25

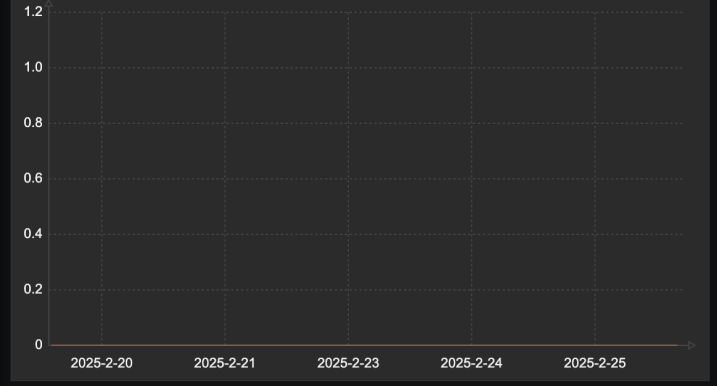

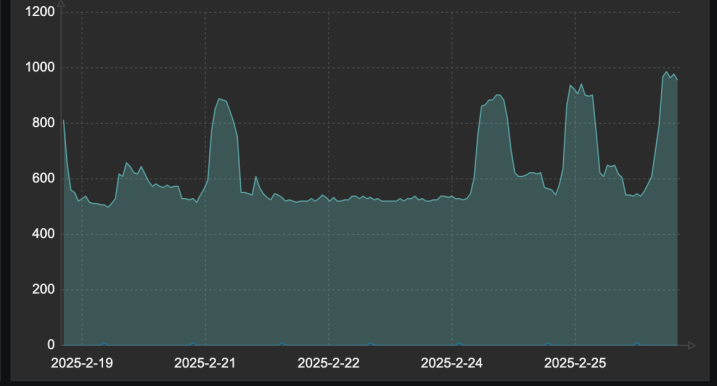

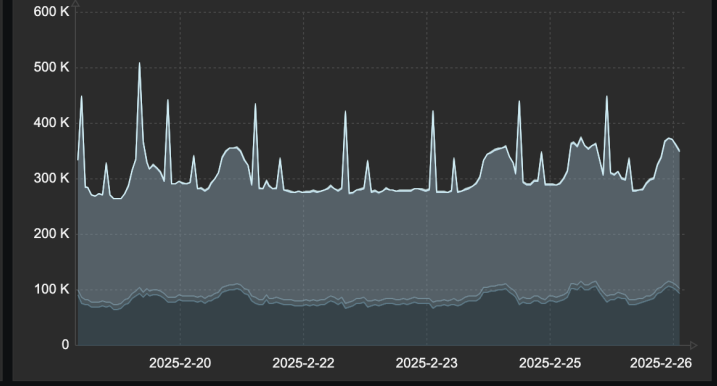

WAL Bytes written

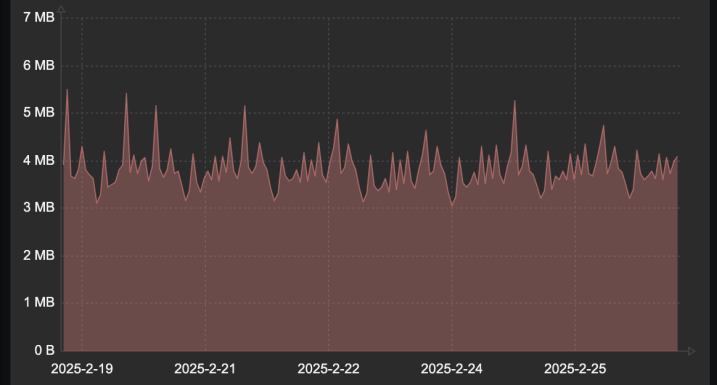

Temp bytes

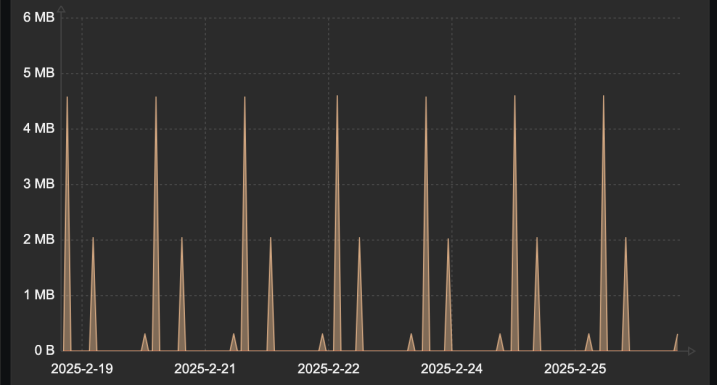

Checkpoint

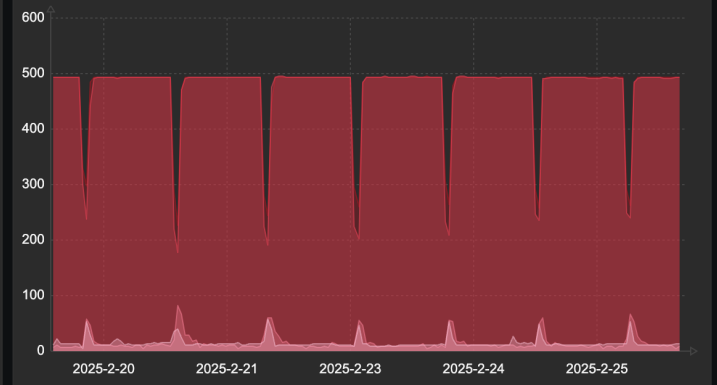

Total Chunks size per Hypertable

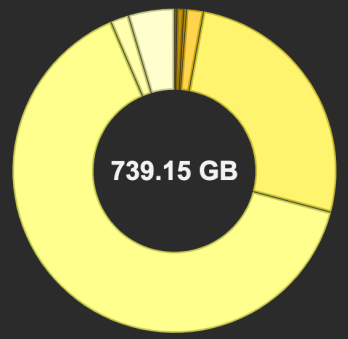

Total Chunks size per Hypertable in time

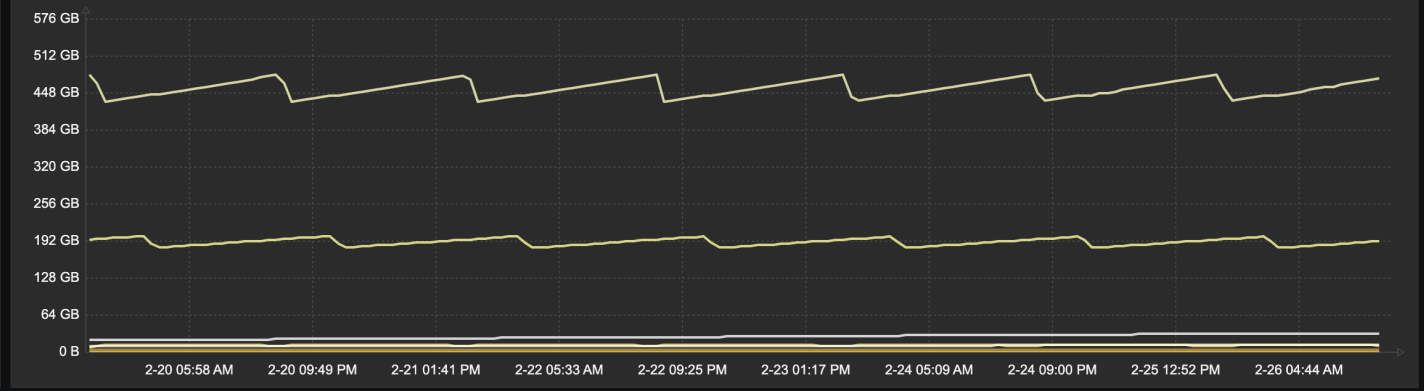

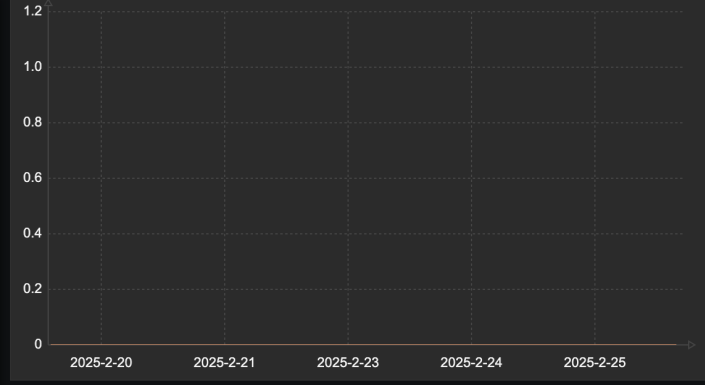

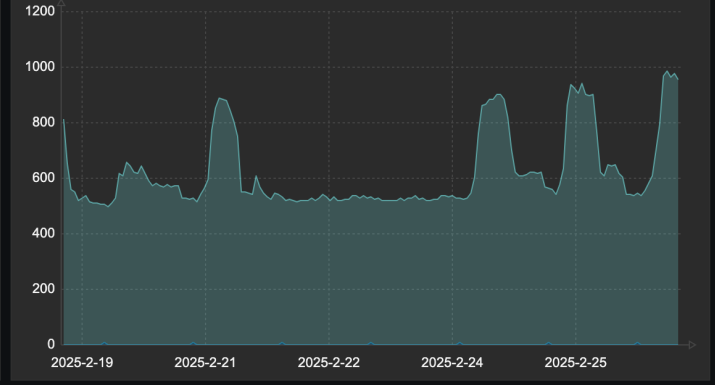

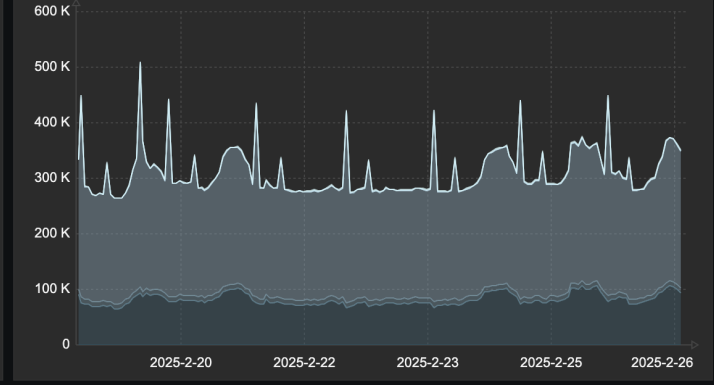

WAL Bytes written

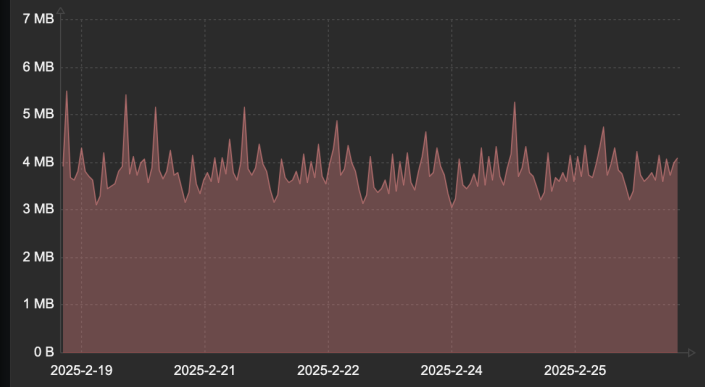

Temp bytes

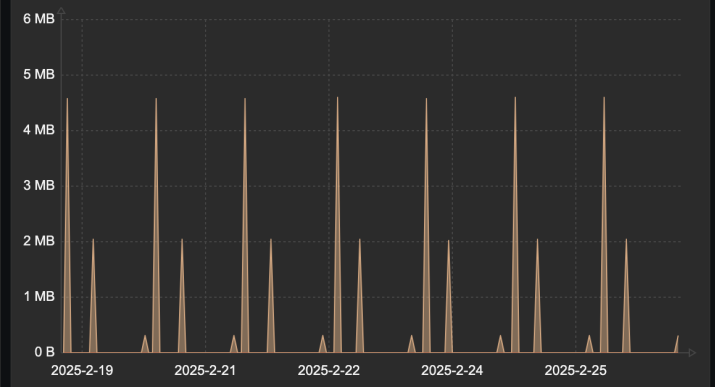

Checkpoint

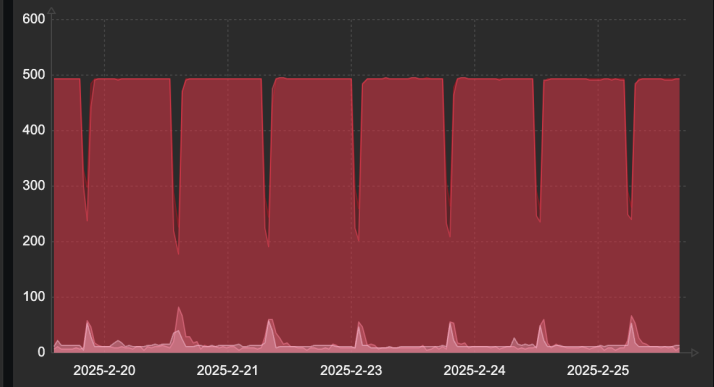

Total Chunks size per Hypertable

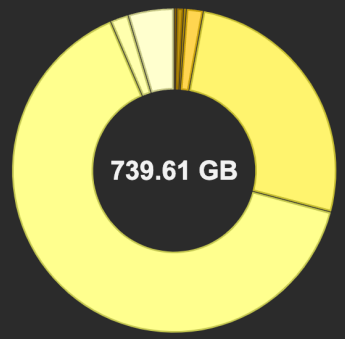

Total Chunks size per Hypertable in time

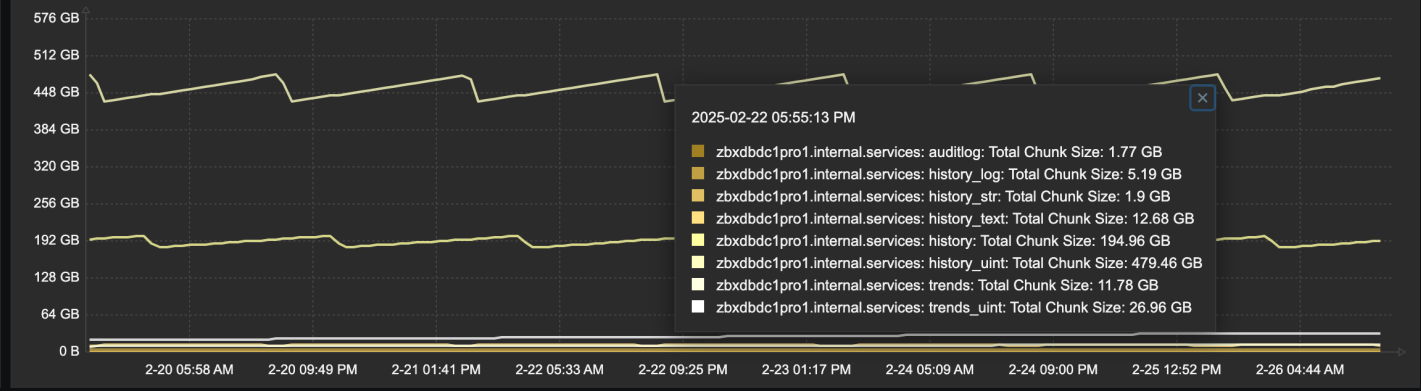

5

Monitoring Your Monitoring

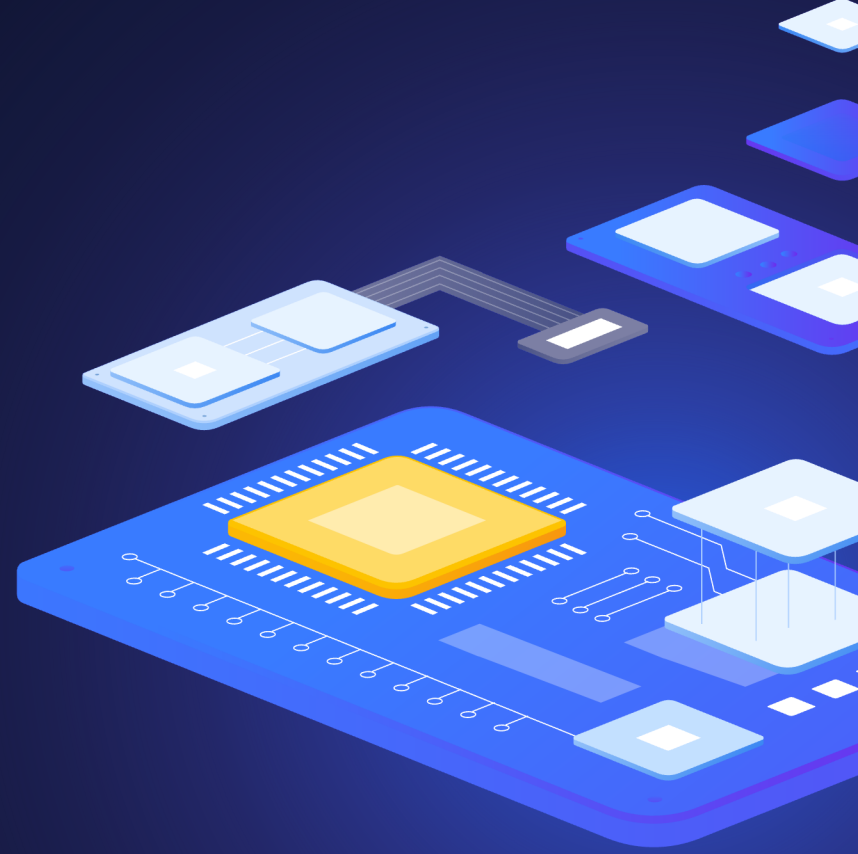

initMAX

Monitoring Your Monitoring

- **Why Monitor Zabbix Database Externally**
  - If Zabbix runs into issues, it may become blind to its own problems.
  - External monitoring prevents unnoticed failures and ensures reliability.

- **Recommended Monitoring Tool: pgwatch**
  - Designed specifically for PostgreSQL database monitoring.
  - Reads real-time metrics directly from the database.
  - Runs efficiently in containers and integrates with Grafana.
  - Provides a live view of database performance independent of Zabbix.

- https://pgwatch.com/
- https://github.com/cybertec-postgresql/pgwatch

initMAX

Monitoring Your Monitoring

- **Key Metrics to Track**
  - Database performance and query response times.
  - Storage usage and capacity trends.
  - Connection pool utilization and limits.
  - Replication status and potential lag issues.

- **Monitoring Strategies**
  - Deploy separate monitoring instances to avoid single points of failure.
  - Enable cross-cluster monitoring for distributed setups.
  - Implement alternative alerting paths independent of Zabbix.

initMAX

Tips and tricks

From time to time check our wiki or social networks

- Throttling, PostgreSQL tuning, …
- https://www.initmax.com/wiki/zabbix-java-gateway-installation-with-tomcat-monitoring/
- https://www.initmax.com/wiki/zabbix-7-0-instructions-for-installation-in-5-minutes/
- https://www.initmax.com/wiki/zabbix-7-0-and-increasing-system-limits/
- https://www.initmax.com/wiki/zabbix-migration-from-mysql-to-postgresql/
- https://www.initmax.com/wiki/how-to-set-up-snmp-trap-in-zabbix/
- https://www.initmax.com/wiki/microsoft-teams-integration-in-five-steps/
- https://www.initmax.com/wiki/reporting-in-zabbix-7-0/

- GitHub (Templates, widgets, modules, …)
- https://www.github.com/initMAX

initMAX

initMAX E-Shop

- Custom visualization widgets
- UX Improvement Modules
- AI Integration with Zabbix
- Both **FREE** and PRO Versions

- initMAX eshop link
  https://www.initmax.com/eshop/

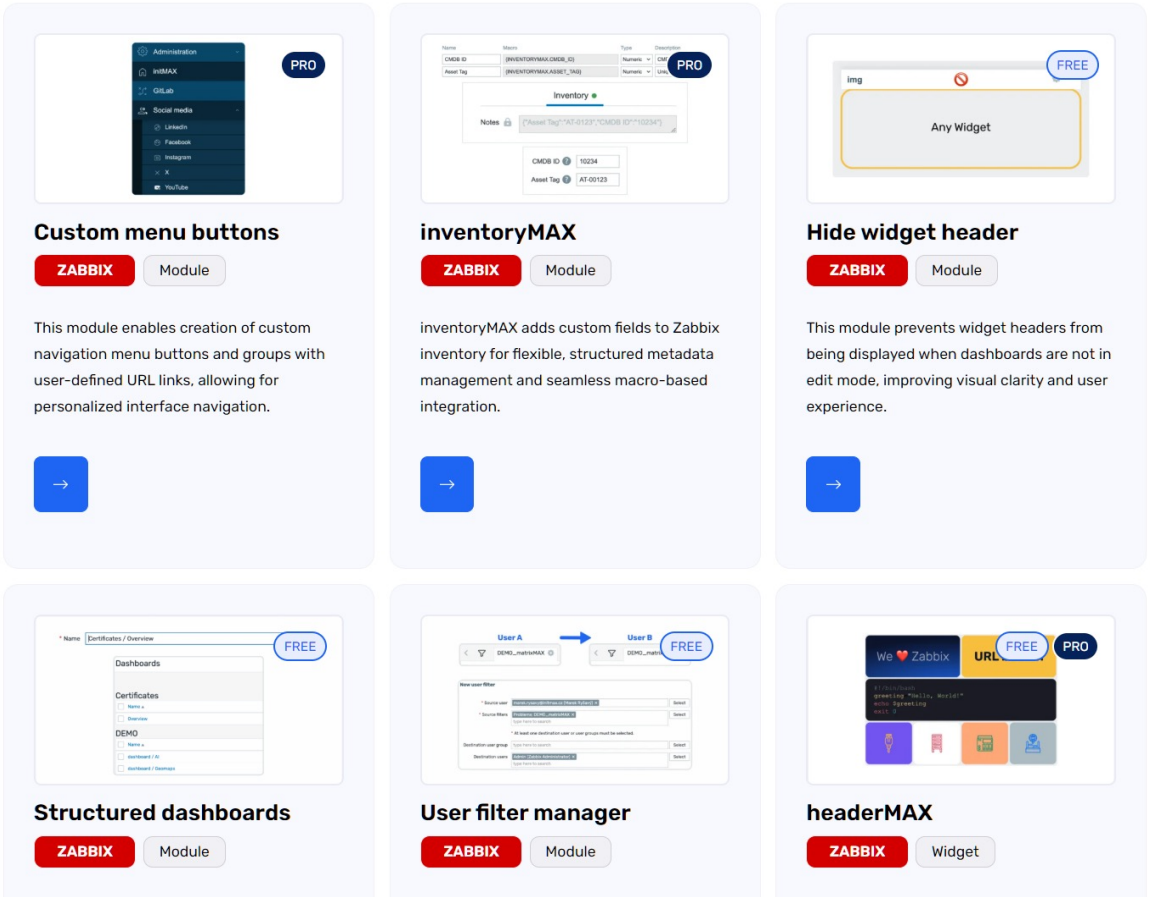

- Video explanation of widgets and modules
  https://www.youtube.com/watch?v=fpW6TR7DQdU&list=PLF7Hh_ikyQDpHiHhXwLtw57OCDF9jn7zL

Zabbix and advanced database monitoring

initMAX

dmitrylambert.com

- FREE Template builder out of JSON
- Tag support
- Low Level Discovery support
- Zabbix 7.0 / 7.2 / 7.4 version support
- Extract Simple JSONPATH
- Download ready to apply Template

- www.dmitrylambert.com

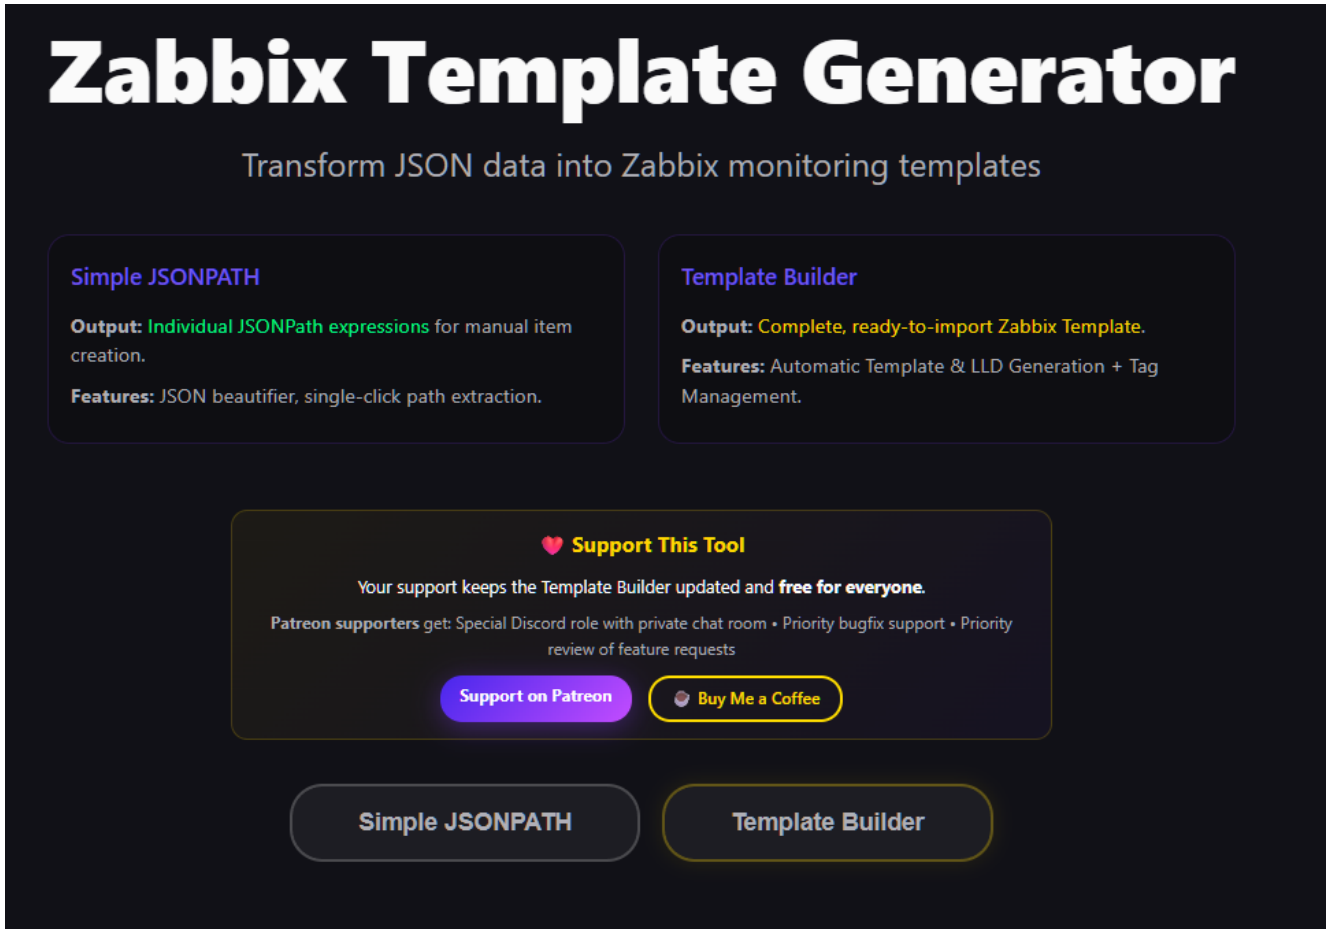

Questions?

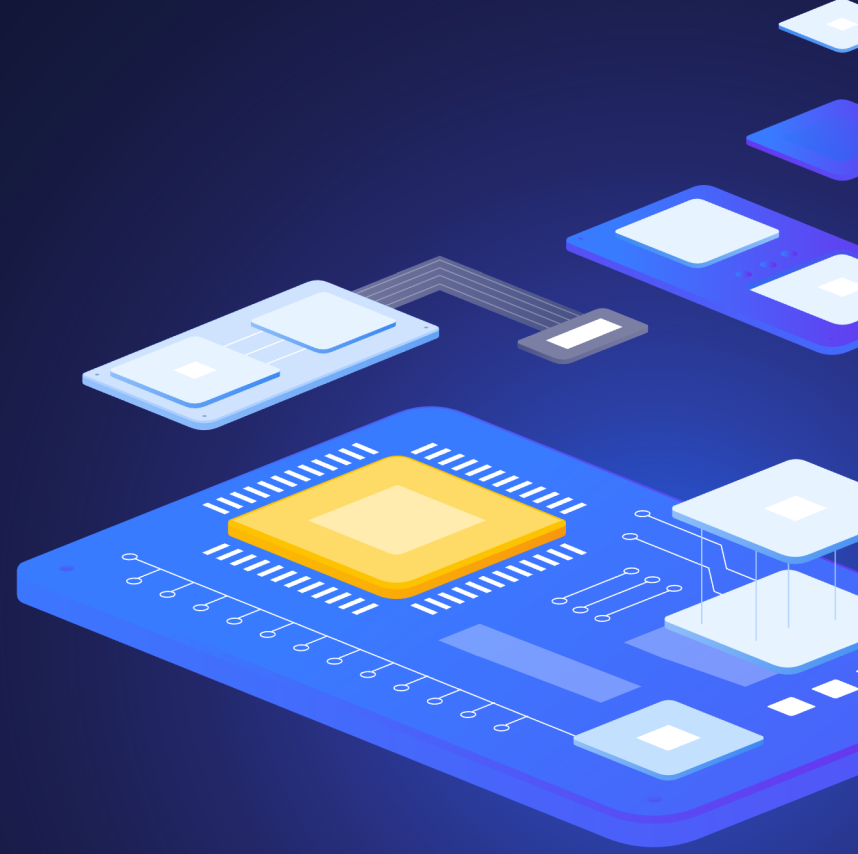

Zabbix and advanced database monitoring

initMAX

Contact us:

Phone:	+420 800 244 442
Web:	https://www.initmax.com
Email:	tomas.hermanek@initmax.com
LinkedIn:	https://www.linkedin.com/company/initmax
Twitter:	https://twitter.com/initmax
Tomáš Heřmánek:	+420 732 447 184