

AI-Powered Alert Analysis with ClickHouse Cloud Databases

From Alert Fatigue to Insightful Automation

Alkin Tezuysal - Director of Services

Boris Tyshkevich - Support Engineer

Altinity® is a Registered Trademark of Altinity, Inc. ClickHouse® is a registered trademark of ClickHouse, Inc.; Altinity is not affiliated with or associated with ClickHouse, Inc.

Alkin Tezuysal

Director of Services @AltinityDB
Open Source Database Evangelist



LinkedIn

linkedin.com/in/askdba/

X

[@ask_dba](https://twitter.com/ask_dba)

Boris Tyshkevich


Support Engineer @AltinityDB
ClickHouse Expert

Agenda

- Using AI for Alert Analysis
- Requirements and data
- Setting up the environment for generating smart dashboards
- Demo

The Goal

d





THE ALERT FATIGUE PROBLEM



Overwhelming alerts lead to missed critical signals.



High volume of alerts

Organizations receive thousands of alerts daily, often leading to alert fatigue among team members, as important notifications become lost in the noise. Drowning in irrelevant alerts reduces operational efficiency, hindering timely responses to genuine issues.



Critical alerts unnoticed

When teams are overwhelmed by alerts, they may overlook crucial information. This can result in missed opportunities for quick responses, increased risk factors, and ultimately, failure to act on significant issues that could affect overall performance.



Inefficient manual triage

Manual triage of alerts cannot scale effectively for large organizations. As alert volume increases, it becomes increasingly difficult for teams to prioritize and respond to alerts promptly, which can lead to burnout and lost productivity over time.



Impact on team morale

The constant influx of alerts can demoralize teams, leading to frustration and burnout. When employees feel overwhelmed, their engagement and efficiency decline, ultimately affecting the organization's ability to respond to critical incidents.



Need for smarter alerts


Instead of increasing alert volume, organizations should focus on creating smarter alerts that prioritize important information. This approach ensures that team members receive relevant data, allowing them to take immediate action on critical issues without being overwhelmed.



Leverage automation tools

Using automation tools can help filter alerts and prioritize them effectively. These tools can analyze patterns, escalating only high-priority alerts to human teams, thereby reducing noise and enabling better focus and more efficient responses.





WHY CLICKHOUSE FOR ALERT ANALYTICS?



High-performance querying

ClickHouse enables fast, real-time analysis of large datasets, making it ideal for alert analytics.

Comprehensive system logs

Includes detailed logs like `query_log` and `asynchronous_metrics_log` to enhance transparency and troubleshooting.




Scalable architecture

Easily scales with increasing alert volume, ensuring performance remains consistent as data grows.

Efficient storage model

ClickHouse uses a columnar storage model, optimizing data retrieval for analytics and alerting tasks.






Why Use LLMs for Data Analytics

Traditional BI platforms like Grafana, Tableau, Power BI, and Looker are great for dashboards and repeatable reports, but they're less useful when you want to *explore* data.

1. Rigid models —> You need pre-built dashboards before you can ask meaningful questions.
2. Slow iteration —> Every new hypothesis means redesigning charts, filters, or writing fresh SQL. That kills discovery speed.
3. Limited flexibility —> Dashboards are good for tracking known KPIs, but awkward for unexpected questions or new analyses.

New BI Architecture



Altinity MCP server

- Designed for ClickHouse®
- Fast go lang code
- Easy to install (no Python dependencies hell)
- All MCP access kinds (stdio, sse, http)
- execute_query tool as main instrument
- Dynamic tool creation (based on Parameterized View)
- Password and JWE auth
- Additional OpenAPI (REST) endpoint
- No OAuth support (yet, planned)
- OpenSource Apache 2.0 License

<https://github.com/Altinity/altinity-mcp>

Connecting to LLM

MCP HTTP URL - <https://mcp.demo.altinity.cloud/eyJ...kNw/http>

- Auth by JWE token with encrypted host/port/login/password
- Ready to paste into UI/Config of any modern chat application:
 - [Claude.ai](#)
 - ChatGPT
 - Claude code
 - Codex
 - Gemini CLI
 - Manus, LibreChat, ... (hundred of them)

Admin settings

Organization

Billing

Usage

Data and

Capabilities

Connectors

Claude

Desktop

Extensions

Connectors

Control which connectors your team members have access to.

Add custom connector BETA

Connect Claude to your data and tools. [Learn more about connectors](#) or get started with [pre-built ones](#).

Alerts

<https://mcp.demo.altinity.cloud/eyJ...kNw/http>

▼ Advanced settings

Only use connectors from developers you trust. Anthropic does not control which tools developers make available and cannot verify that they will work as intended or that they won't change.

Add

Cancel

Eating Our Own Dogfood

- Which data to explore for examples?
- Alerts Database
 - Not too big, but many people know the field and understand the problem
 - Could be anonymized for public demo
 - Analytic tasks require consistent data with clear business terms columns
 - Have structured columns (like status and priority)



Atlassian API



Table Schema

```
CREATE TABLE support.opsgenie_alerts
(
    `alert_id` String,
    `status` String,
    `acknowledged` UInt8,
    `is_seen` UInt8,
    `created_at` DateTime64(3),
    `updated_at` DateTime64(3),
    `priority` String,
    `owner` Nullable(String),
    `fire` UInt8,
    `alert` LowCardinality(String),
    `severity` LowCardinality(String),
    `cluster` String COMMENT 'clickhouse cluster, chi',
    `environment` LowCardinality(String) COMMENT 'client, tenant, customer'
)
```

Creating Agent

- Starting Chat each time from a “white list” explaining all requirements is impractical
- Better to give the model standard instructions and basic knowledge on a particular field
- An agent can be developed in any programming language using the LLM’s APIs
- Alternatively, we can build an agent inside the common Chat UI. For ChatGPT, it’s called GPTs; for Claude, it’s Projects.
- All agents’ instructions are placed into limited Context Window. Overloading the Agent’s prompt with too many details is a bad idea!
- Claude and ChatGPT has its own file storage for Agents with vector search (RAG). We can place any additional knowledge and refer to it in the prompt.
- Claude has new feature - Skills - additional instructions for the Agent explaining how to do something in a proper way. Skill’s prompt loaded automatically when a keyword or phrase is typed by a user in a normal chat conversation (they should have a proper description)

Set project instructions

Provide Claude with relevant instructions and information for chats within Alert Analyst. This will work alongside user preferences and the selected style in a chat.

You are an Alert Analyst, a specialized assistant who helps find problems in Clickhouse instances in the Cloud. Your job is to generate only safe, read-only SELECT queries to retrieve meaningful insights.

- All queries should operate exclusively on the table: support.opsgenie_alerts_v3.
- use column 'cluster' when asked about clickhouse clusters
- use column 'environment' as a client/organization/tenant/etc
- skip priority P3/P4) and all clusters and environments having dev and staging substrings in their name (until requested)
- P1: Critical, P2: High
- To simple questions, answer clearly and briefly.
- If source coverage is insufficient to answer confidently, state that clearly and specify the gap.
- State limitations or uncertainty when appropriate; do not speculate.
- Summarize retrieved facts and craft a focused, professional response.

Recommended Tools:



Alert Analyst



Share

Created by you · Shared with your org

Top Priority alerts by alert type and organization breakdown for the last 4 months



Sonnet 4.5 ▾



Search menu

Your c



Use style



chats are private until shared



Extended thinking



Research



Web search



Drive search



Gmail search

Connect



Calendar search

Connect

Top 1



alerts



Last m



altinity demo



Demo1: Top Priority alerts

Top Priority alerts by alert type and organization breakdown for the last 4 months

Top Organizations by Total Alert Volume:

1. quantum-labs-569 - 3,071 alerts (ClickHouseMaxPartCountForPartition)
2. prime-labs-156-stg - 2,357 alerts (ClickHouseDistributedFilesToInsertContinuouslyGrowing)
3. green-technologies-508-prod - 795 alerts (ClickHouseReplicasSumQueueSize)

Critical Patterns:

- "Flapping" alerts indicated by avg_fires_per_alert > 1:
 - vertex-works-331-eu: 11.3 fires/alert (MetricsExporterFetchErrors)
 - clear-partners-369: 6.0 fires/alert (MetricsExporterFetchErrors)
 - global-solutions-525-prod-us: 4.55 fires/alert (MaxPartCountForPartition)
- P1 Critical Alerts: 42 KubeletDown, 41 ZookeeperDigestMismatch, 39 BackupFailed incidents

Demo2: Chronic Problems

Which clusters have had the SAME alert firing for 30+ consecutive days? Not just open alerts, but closed and repeated during the next 24 hours. Analyze data for the last 6 months.

Most common chronic alerts:

- ClickHouseTooManyDetachedParts (35 clusters)
- ClickHouseBackupFailed (8 clusters)
- ClickHouseMaxPartCountForPartition (5 clusters)

Dashboards

- Text is great, but sometimes Visual is better
- LLM can create Artifact as HTML/JS code
- Random design and colors is not the corporates want to have
- Advanced prompting technology as Claude Skills can teach LLM to make stable design
- Embedding data into HTML will spend tokens
- Dashboard is a full fledged site and can load data directly from Clickhouse through Altinity-MCP



Alert Organization Breakdown

P1/P2 Priority Alerts Analysis by Organization & Alert Type
July 2025 - November 2025 | Production ClickHouse Clusters Only

TOTAL ORGANIZATIONS

89

With P1/P2 alerts

TOTAL ALERTS

11,847

Last 4 months

ALERT TYPES

45

Unique alert types

P1 CRITICAL

587

5% of total

FLAPPING ALERTS

28


Avg fires > 2.0

Key Insights

- quantum-labs-569/sirius-150 leads with 3,071 MaxPartCountForPartition alerts (26% of total volume)
- Flapping behavior detected: vertex-works-331-eu shows 11.3 avg fires per alert (MetricsExporterFetchErrors)
- Top 3 alert types: MaxPartCountForPartition (4,123), DistributedFilesToInsert (3,012), ReplicasSumQueueSize (2,156)
- P1 Critical Issues: KubeletDown (42), ZookeeperDigestMismatch (41), BackupFailed (39)
- 6 organizations account for 60% of all alert volume



Top 20 Organizations by Alert Volume




Alert Type Distribution



P1 vs P2 Priority Split



Top Flapping Alerts (High Fire Frequency)



Q&A

References

- [Altinity Cloud](#)
- [Altinity MCP Server](#)
- [Project \(Agent\) Prompt](#)
- [Transcript: Demo1 Top Priority alerts](#)
- [Transcript: Demo2: Chronic Problems](#)
- [Artifact: Static Dashboard](#)
- [Artifact: Skill-generated Dashboard](#)