

Building Open Source Database Infrastructure for OLTP and OLAP

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

Enterprise Database Infrastructure

- **Online Transaction Processing Systems (OLTP)**
 - Real-time and concurrent transaction processing systems
 - Rapid Response(ultra-low latency) Transaction Computing Systems
 - Frequent UPDATES and DELETES
 - Not storage intensive
 - Most of the queries just search for a single record or few
- **Column Stores for Data Analytics and Statistical Computing**
 - No UPDATES and significantly READ intensive (mostly SORT/SEARCH intensive)
 - Stores large amount of Historical Data – Significantly high storage investment
 - More complex query structure analyzing a larger set of rows
 - Data is stored in Columnar Format- Data of a specific set of columns are homogeneous. This makes data compression more efficient and better query performance due to lower disk I/O.

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

Complexities in Database Infrastructure Operations – Applicable for both OLTP & OLAP ecosystem

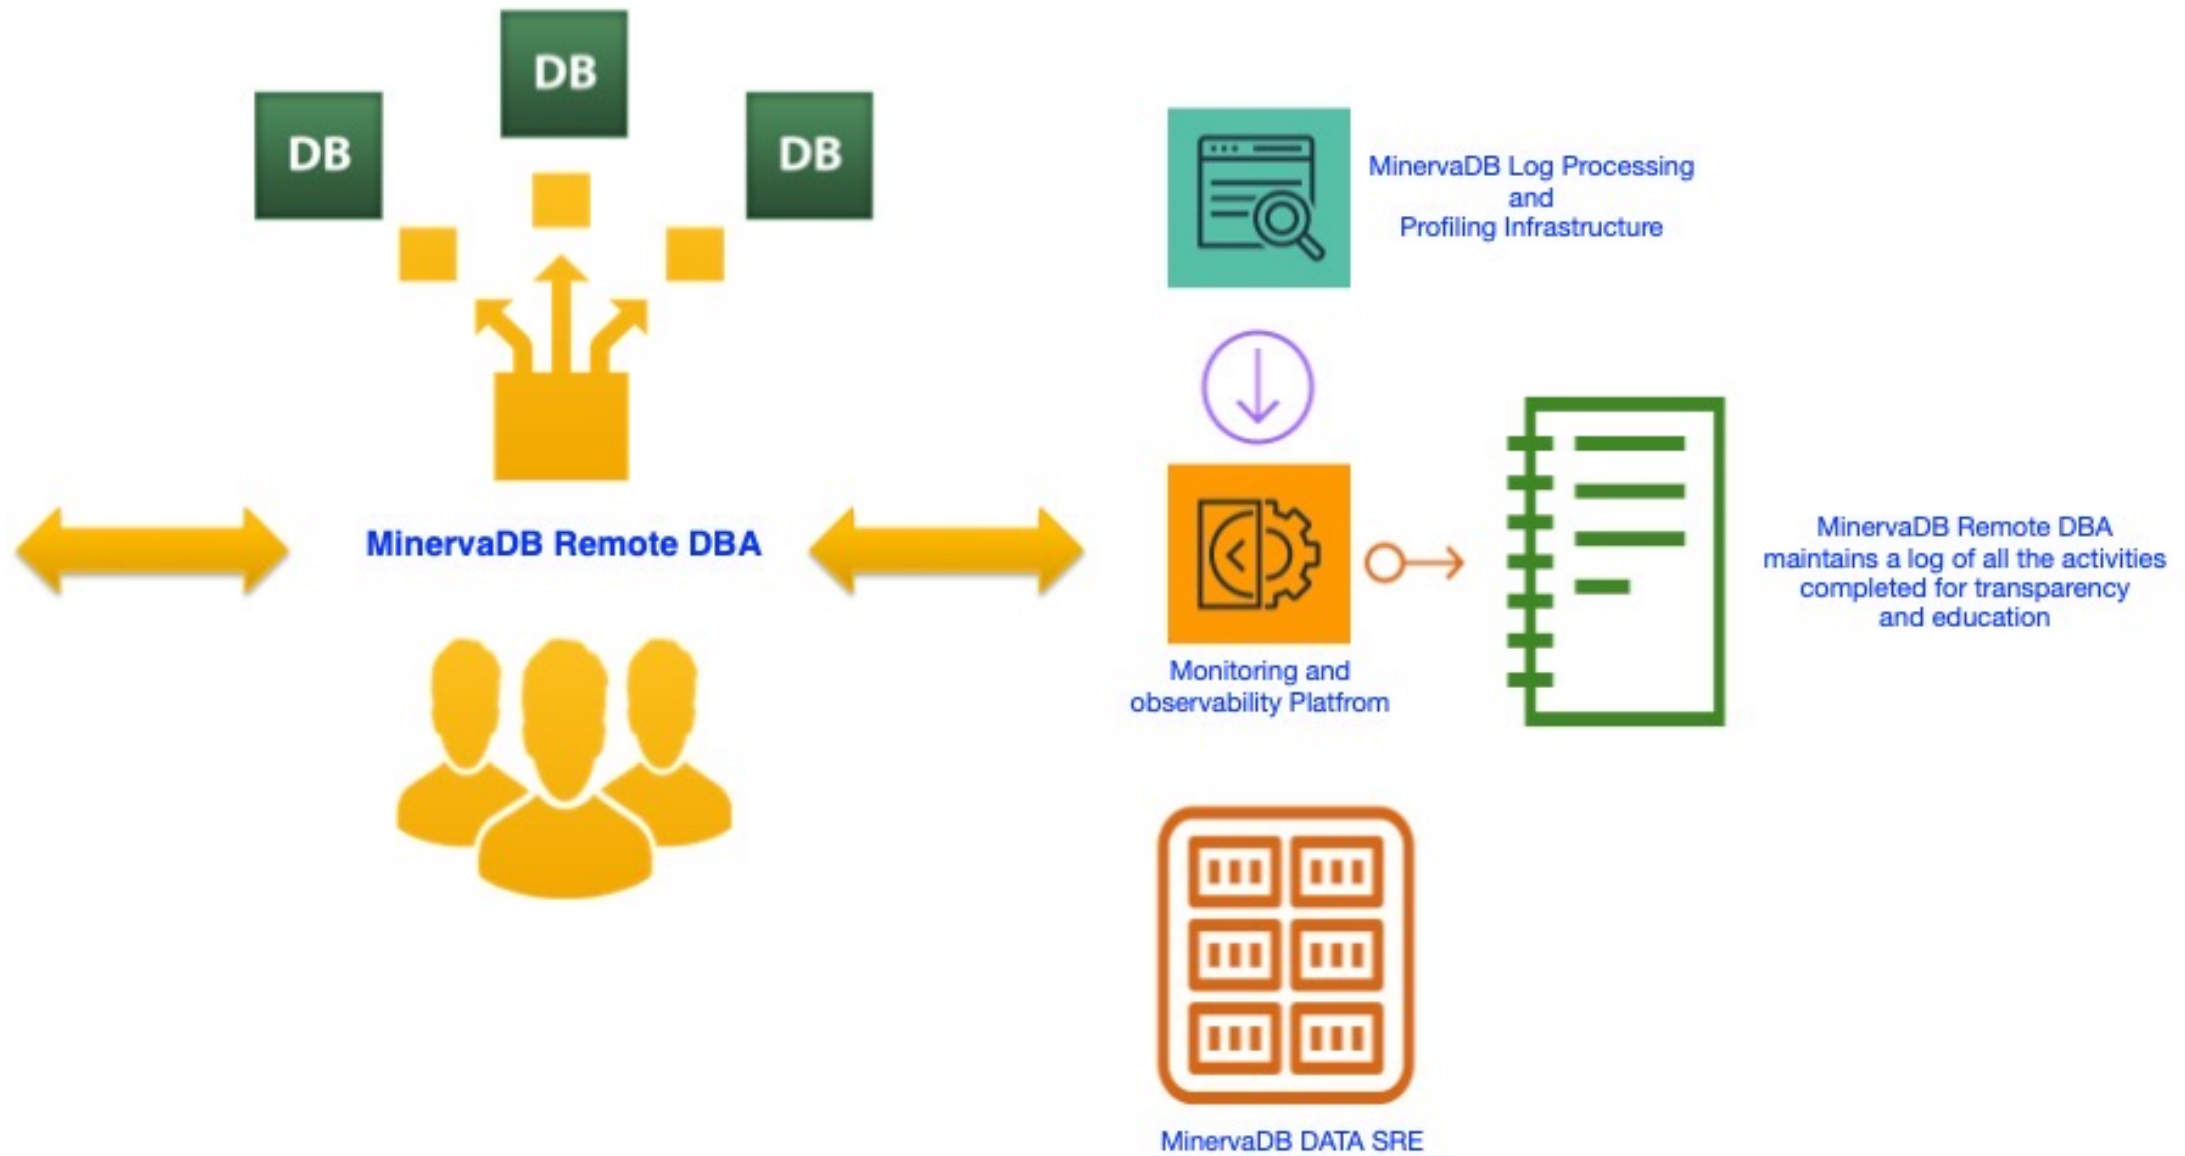
| | | |
|-------------------|---------------|-----------------------------|
| Performance | Scalability | Data SRE |
| Automation/DevOps | Monitoring | Scale-out/Replication |
| Audit Logs | Data Security | Self-Healing/Fault-Tolerant |

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

Partner for Open Source Database Systems Engineering – PostgreSQL, MySQL and ClickHouse

- Globally distributed elite-class team of Database Infrastructure Engineering and Operations Experts available 24*7
- Custom Engineering for PostgreSQL, MySQL and ClickHouse
 - Customized Database Infrastructure for Performance, Scalability, Reliability and High Availability
 - Bug fixing and Patching for PostgreSQL, MySQL and ClickHouse
- Full-Service and Platform Neutral Open Source Database Systems Team:
 - Architecture Engineering – Building Planet-Scale Data Infrastructure to Scale
 - Lean Database Infrastructure Operations – Cost-Efficient and Responsive

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP



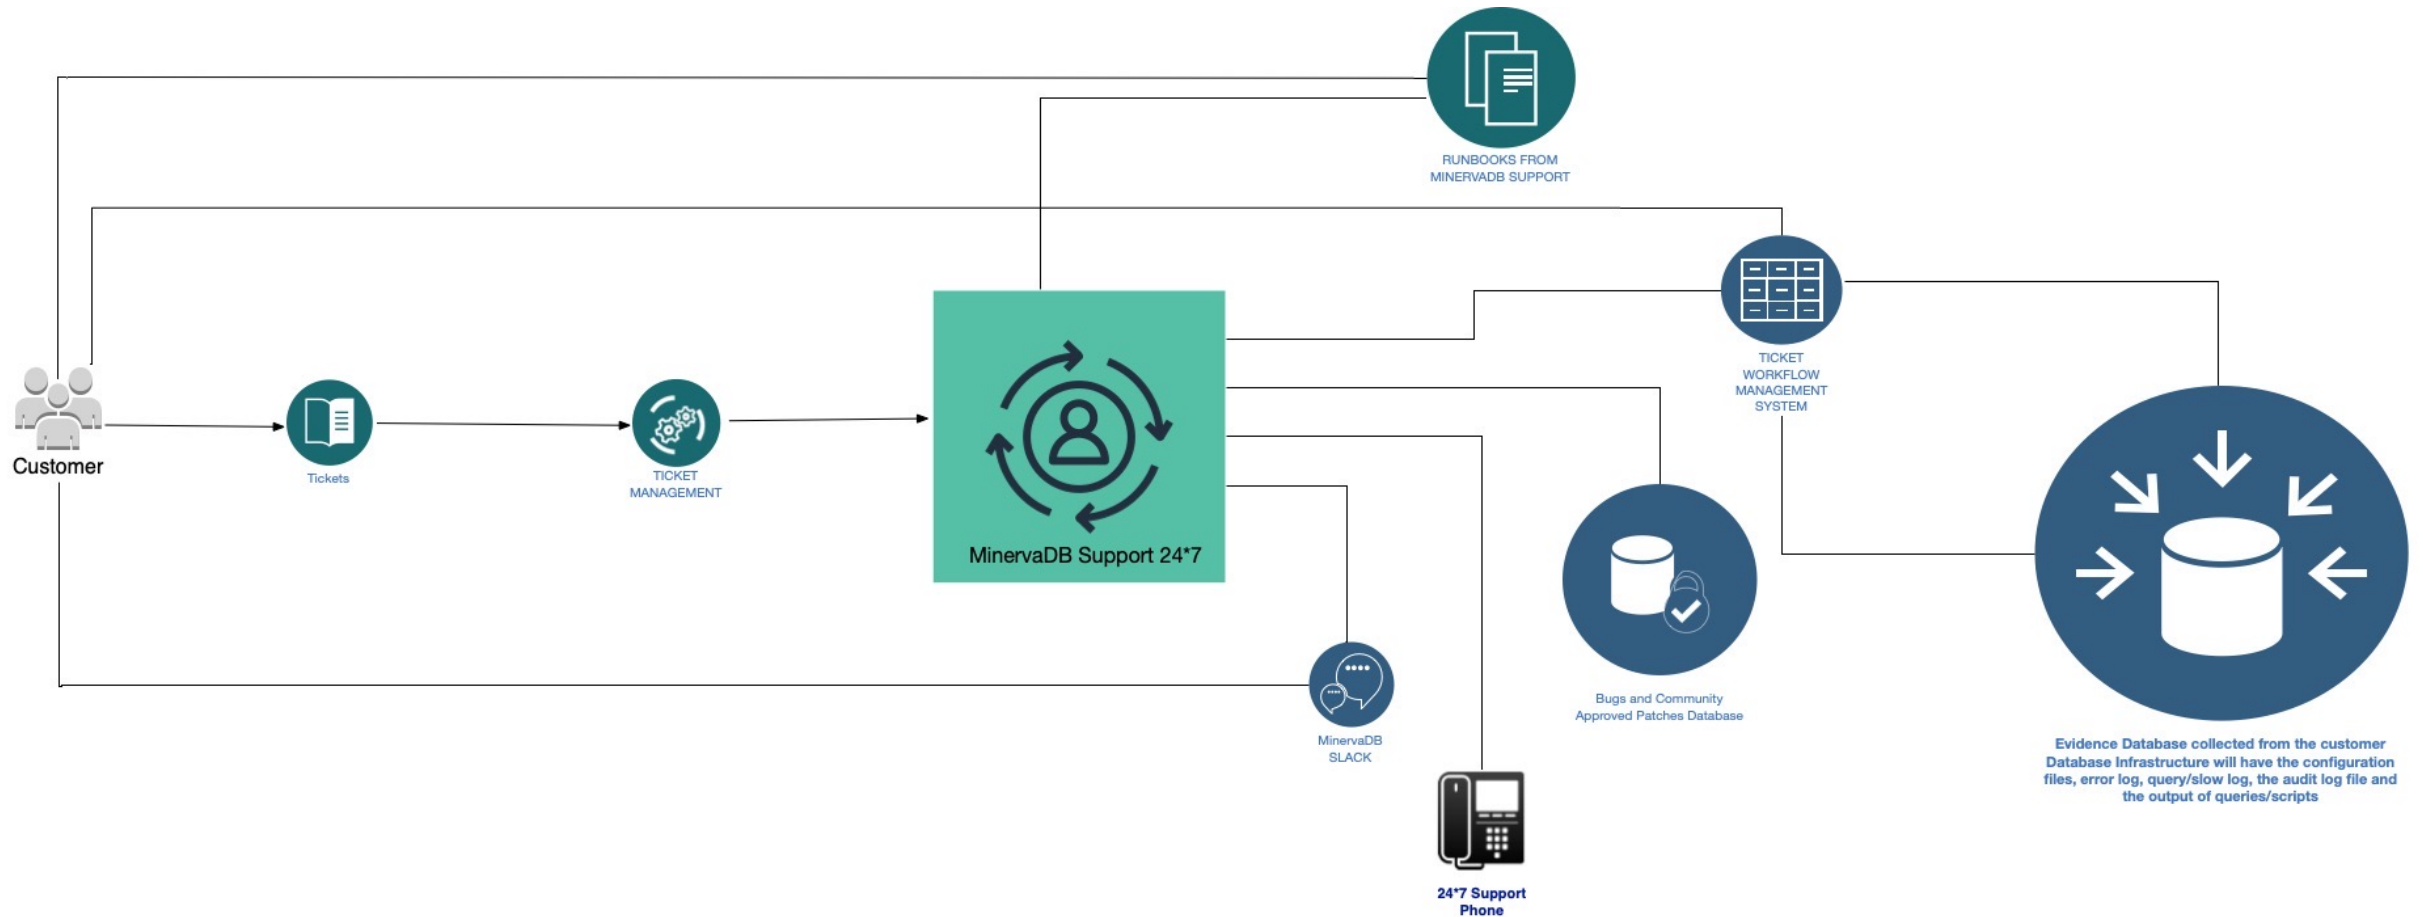
Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

Consultative Support (24*7) for Open Source Database Systems Infrastructure Operations

- Recommendations for optimal MySQL, PostgreSQL and ClickHouse installation and configuration – Do it right the first time !!
- Bug fixing (Community Fixes) and patching:
 - MySQL
 - PostgreSQL
 - ClickHouse
- You can report unlimited support incidents from 12 logins.
- Recommendations for custom installation, configuration, schema development, SQL engineering, optimal indexing, performance tuning, building fault tolerant & self healing replication solutions and database security.

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

Consultative Support (24*7) for OLTP and OLAP



Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

PostgreSQL (24*7)Consultative Support

- Recommendations for optimal PostgreSQL installation and configuration – Do it right the first time !!
- PostgreSQL Bug fixing (Community Fixes) and patching.
 - *Community Fixes mean that where possible, we will report bugs and any MinervaDB-created bug fixes to the upstream software vendor or open source project.*
- **Proactive PostgreSQL support – Recommendations for custom PostgreSQL installation, configuration, schema development, SQL engineering, optimal indexing, performance tuning, building fault-tolerant & self-healing replication solutions and database security.**
- Recommendations for optimal schema design, indexing and SQL engineering – Building high-performance PostgreSQL applications.
- Recommendation for capacity planning and sizing – We guide you in optimally sizing CPU, memory, disk and network for application performance, scalability and reliability.
- Recommendations and guidance for high database availability and disaster recovery – We support customers in building highly available and reliable applications.
- Recommendations and guidance for database sharding and clustering solutions – Horizontal scalability and reliable web-scale database infrastructure operations.
- Recommendations and guidance for scale-out and replication solutions
- Recommendations and guidance for database systems upgrades and migration – Optimal PostgreSQL upgrades and migration in zero downtime.

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

MySQL (24*7)Consultative Support

- MySQL Bug fixing (Community Fixes) and patching.
 - *Community Fixes mean that where possible, we will report bugs and any MinervaDB-created bug fixes to the upstream software vendor or open source project.*
- **Proactive MySQL support – Recommendations for custom MySQL installation, configuration, schema development, SQL engineering, optimal indexing, performance tuning, building fault-tolerant & self-healing replication solutions and database security.**
- Recommendations for optimal schema design, indexing and SQL engineering – Building high-performance MySQL applications.
- Recommendation for capacity planning and sizing – We guide you in optimally sizing CPU, memory, disk and network for application performance, scalability and reliability.
- Recommendations and guidance for high database availability and disaster recovery – We support customers in building highly available and reliable applications.
- Recommendations and guidance for database sharding and clustering solutions – Horizontal scalability and reliable web-scale database infrastructure operations.
- Recommendations and guidance for scale-out and replication solutions – We support both synchronous and asynchronous replication.

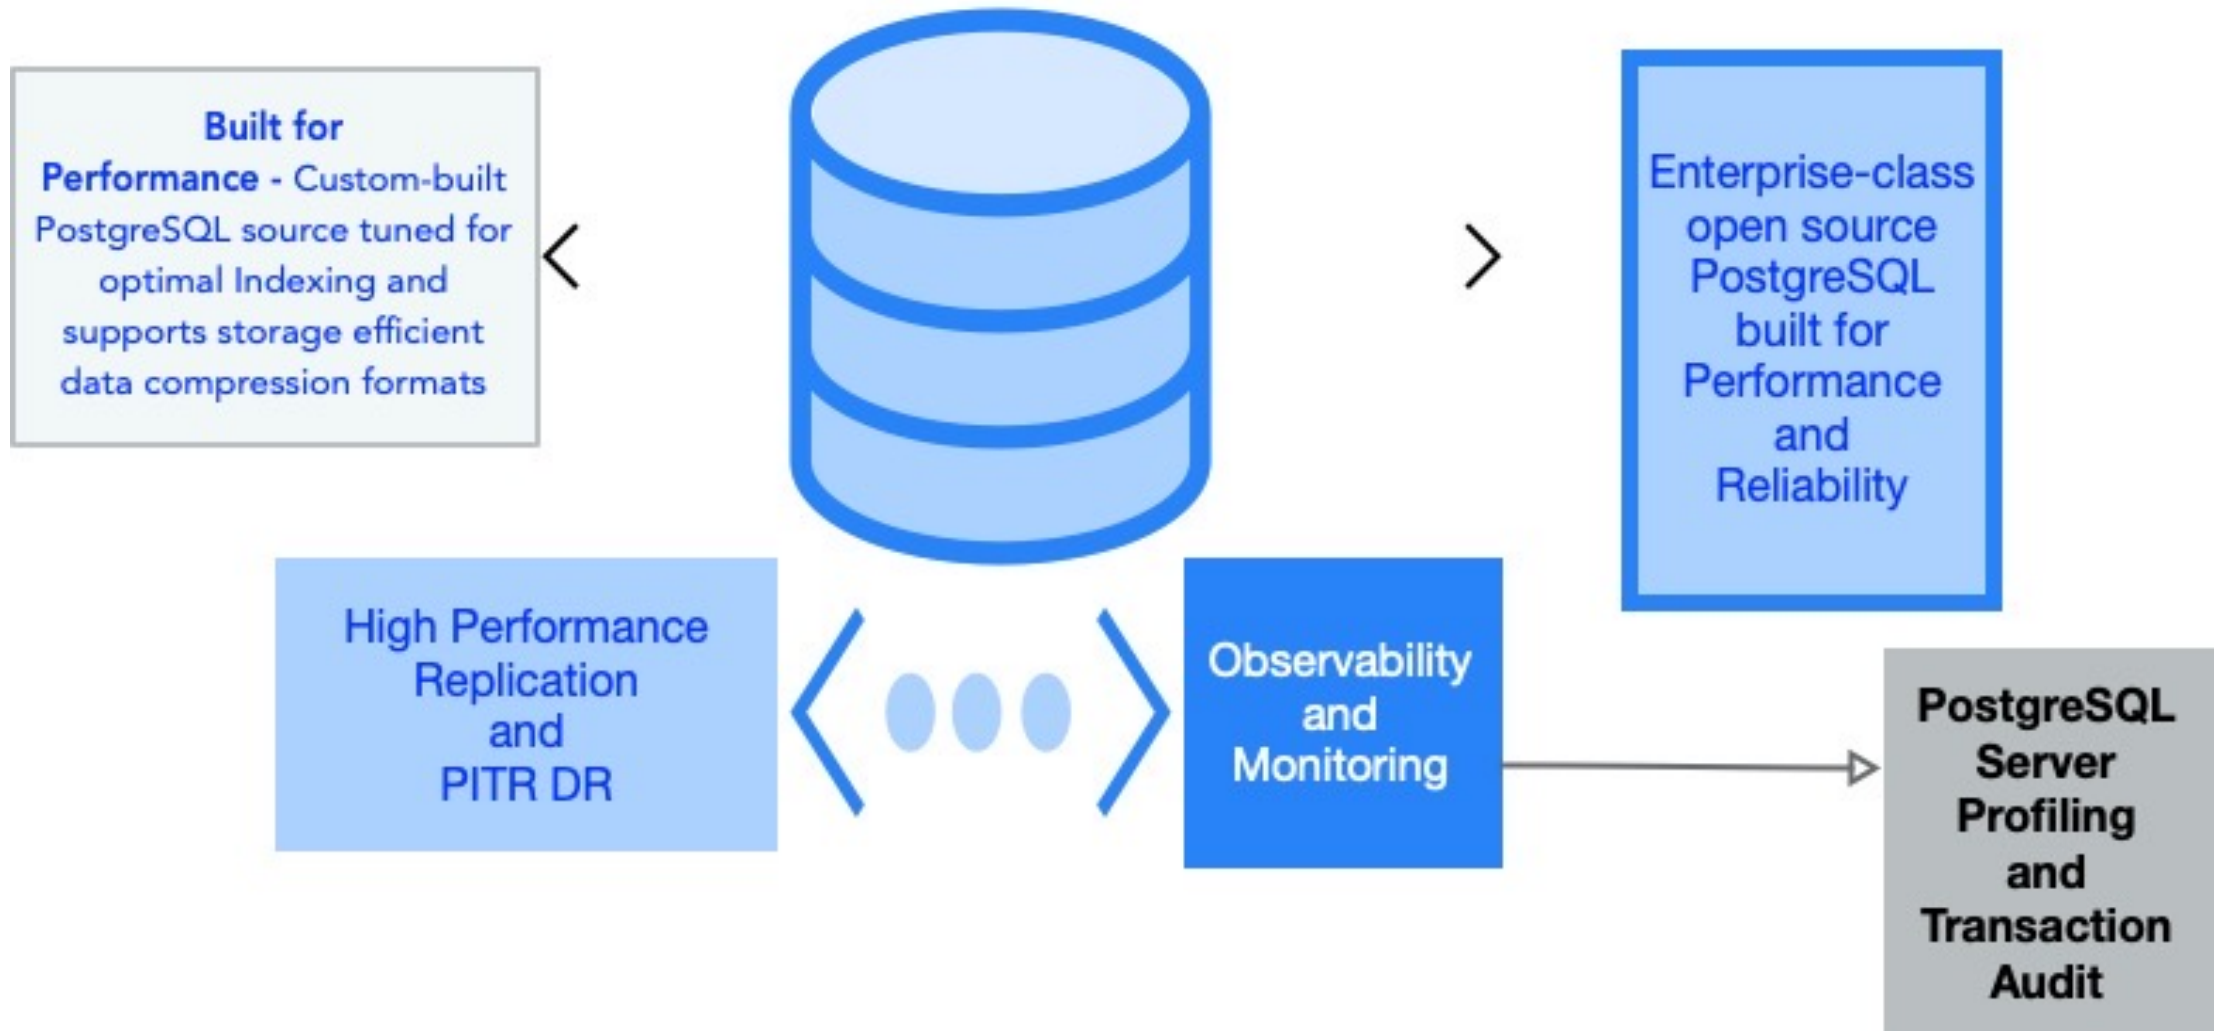
Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

Several Reasons to Engage with a Forever Open Source Database Systems Company

- Open Source and 100% GPL forever – No lock-in and available on Cloud/DBaaS (Database as a Service)
- A globally distributed and highly skilled team of Database Engineers with a deep understanding of optimizers, resource distribution/utilization algorithms and file systems engineering
- Lean Team Culture – Startup-friendly and specialists in DevOps. Extreme automation capabilities for scale and reliability
- Based out of San Francisco Bay Area. But, we have global teams operating from 11 cities worldwide to deliver 24*7 Consultative Support and Managed Services.

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

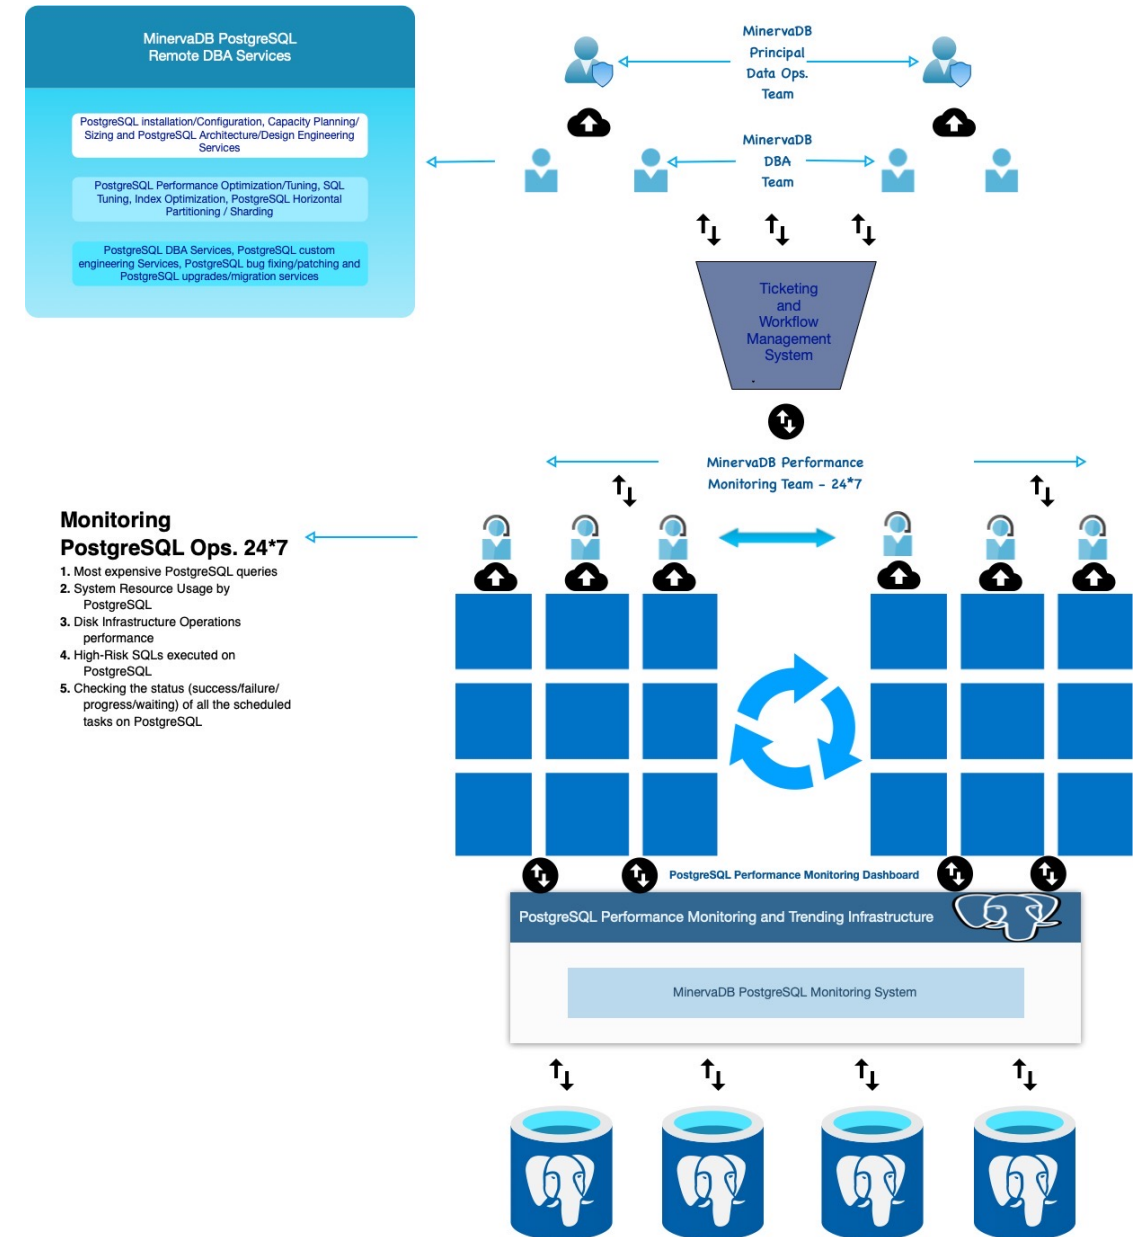
MinervaDB Server for PostgreSQL



Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

PostgreSQL Managed Services

- PostgreSQL Installation and Configuration
- PostgreSQL Performance Benchmarking
- PostgreSQL Capacity Planning/Sizing
- PostgreSQL Performance Audit / Health-Check / Diagnostics
- PostgreSQL Server Tuning
- PostgreSQL Query Optimization
- PostgreSQL Index Optimization
- PostgreSQL Disk Optimization
- PostgreSQL Scale-out/Replication Solutions
- PostgreSQL Data SRE Services
- PostgreSQL Custom Software Engineering / Feature Enhancement
- PostgreSQL Bug Fixing/Patching
- PostgreSQL Upgrades and Migration
- Building Secured Database Infrastructure on PostgreSQL



Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

MinervaDB Server for MySQL

MinervaDB Server for MySQL is an Enterprise-Grade MySQL infrastructure built from MySQL source (100% compatible with MySQL GA) addressing web-scale performance, scalability, high availability, data reliability and data security. Our MySQL engineers are highly experienced in optimal server architecture/configuration, storage engine optimization, InnoDB, RocksDB custom feature development, horizontal partitioning/sharding, transaction audit, data encryption, observability/monitoring and bug fixing/patching.

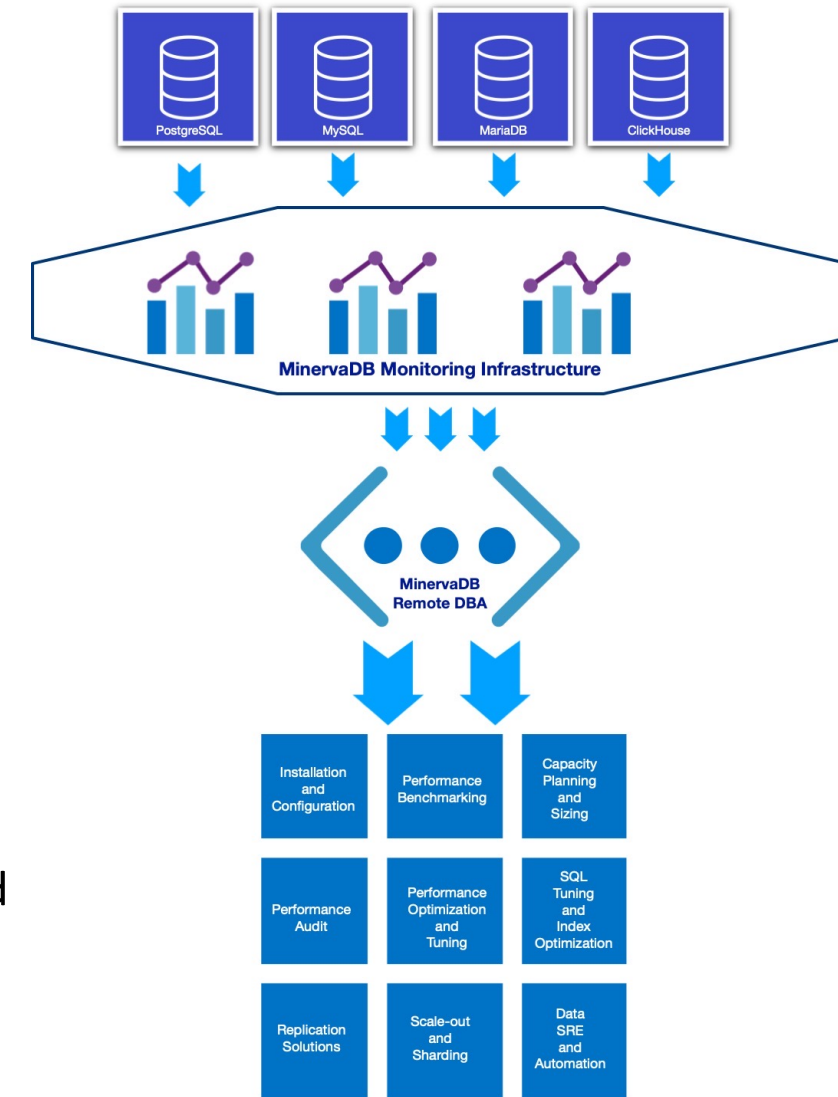


MySQL Server custom-built for performance, Scalability and Reliability

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

Managed Services for MySQL, PostgreSQL, MariaDB and ClickHouse

- Installation and configuration – Custom installation and configuration for MySQL, MariaDB, Percona Server and PostgreSQL.
- Optimal schema design and SQL engineering – Design and build optimal schema and SQL for performance and scalability.
- Optimal indexing – Design optimal indexes and re-write / tune SQL to use indexes efficiently.
- Capacity planning and sizing – Proactively sizing CPU, memory, storage and network for future growth.
- Performance health check and diagnostics – Proactive performance audit and health check to avoid potential performance bottleneck.
- Performance optimization and tuning – Tuning hardware, Linux and MySQL for performance, scalability and reliability.
- SQL optimization – Optimize SQL on response time, execution plan efficiency and index usage.
- Disk I/O tuning – Tuning disk for transaction performance, faster recovery and reliability.
- Scale-out and replication – Synchronous and asynchronous replication solution addressing performance, scalability and reliability.



Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

ChistaDATA Server for ClickHouse – Open Source Column Store for Data Analytics

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

ColumnStore and Row-Based Database Managed System – Why it's better to use ColumnStores for SORT/SEARCH intensive Analytics Operations

- **Compact data storage** – Ten billion UInt8-type values should exactly consume 10GB uncompressed to use the available CPU efficiently. Optimal storage, even when uncompressed, benefits performance and resource management. ClickHouse is built to store data efficiently without any garbage.
- **CPU efficient** – Whenever possible, ClickHouse operations are dispatched on arrays rather than on individual values. This is called “vectorized query execution,” and it helps lower the cost of actual data processing.
- **Massively Parallel Processing** – ClickHouse is capable of Massively Parallel Processing very large/complex SQL(s) optimally and cost-efficiently

ChistaDATA Server for ClickHouse - ColumnStore for SORT/SEARCH intensive query operations with vectorized query execution capabilities

Row-Oriented Tables - Table records stored in a sequence of rows.

| Customer_ID | Customer_Plan | Subscription_Details | Events | Event_Value | Event_Category |
|-------------|---------------|----------------------|---------------|-------------|----------------|
| 2355216 | ABS—GAM | MEDIA & GAMING | TEXR-GAME | \$4.90 | SOCIAL-GAMING |
| 2346221 | EDU-INFOT | BIO-LIFE | CUS-PROT-SYN | \$8.25 | MEDIA-INFOTAIN |
| 66231125 | IPTV-SPORT | SPORT-MEDIA | EURO-CLB-SOCC | \$6.49 | MEDIA-SPORT |

Column-Oriented Tables - Each column is stored in a separate file and data from the same column are stored together

| Customer_ID | 2355216 | 2346221 | 66231125 |
|----------------------|----------------|----------------|---------------|
| Customer_Plan | ABS—GAM | EDU-INFOT | IPTV-SPORT |
| Subscription_Details | MEDIA & GAMING | BIO-LIFE | SPORT-MEDIA |
| Events | TEXR-GAME | CUS-PROT-SYN | EURO-CLB-SOCC |
| Event_Value | \$4.90 | \$8.25 | \$6.49 |
| Event_Category | SOCIAL-GAMING | MEDIA-INFOTAIN | MEDIA-SPORT |

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

ChistaDATA Server for ClickHouse

- **Can store data in disk** – The columnar database systems like SAP HANA and Google PowerDrill can only work in the RAM.
- **Data compression** – ClickHouse supports two kinds of compression LZ4 and ZSTD. LZ4 is faster than ZSTD, but the compression ratio is smaller. ZSTD is faster and compresses better than traditional Zlib but slower than LZ4. We recommend customers LZ4 when I/O is fast enough, so decompression speed will become a bottleneck. When using super ultra-fast disk subsystems, you have an option to specify “none” compression. ZSTD is recommended when I/O is the bottleneck in queries with large range scans.
- **Can store data in disk** – The columnar database systems like SAP HANA and Google PowerDrill can only work in the RAM.
- **Built for web-scale data analytics** – ClickHouse supports sharding and distributed processing. This makes ClickHouse the most preferred columnar database system for web-scale. Each shard in ClickHouse can be a group of replicas addressing maximum reliability and fault tolerance.



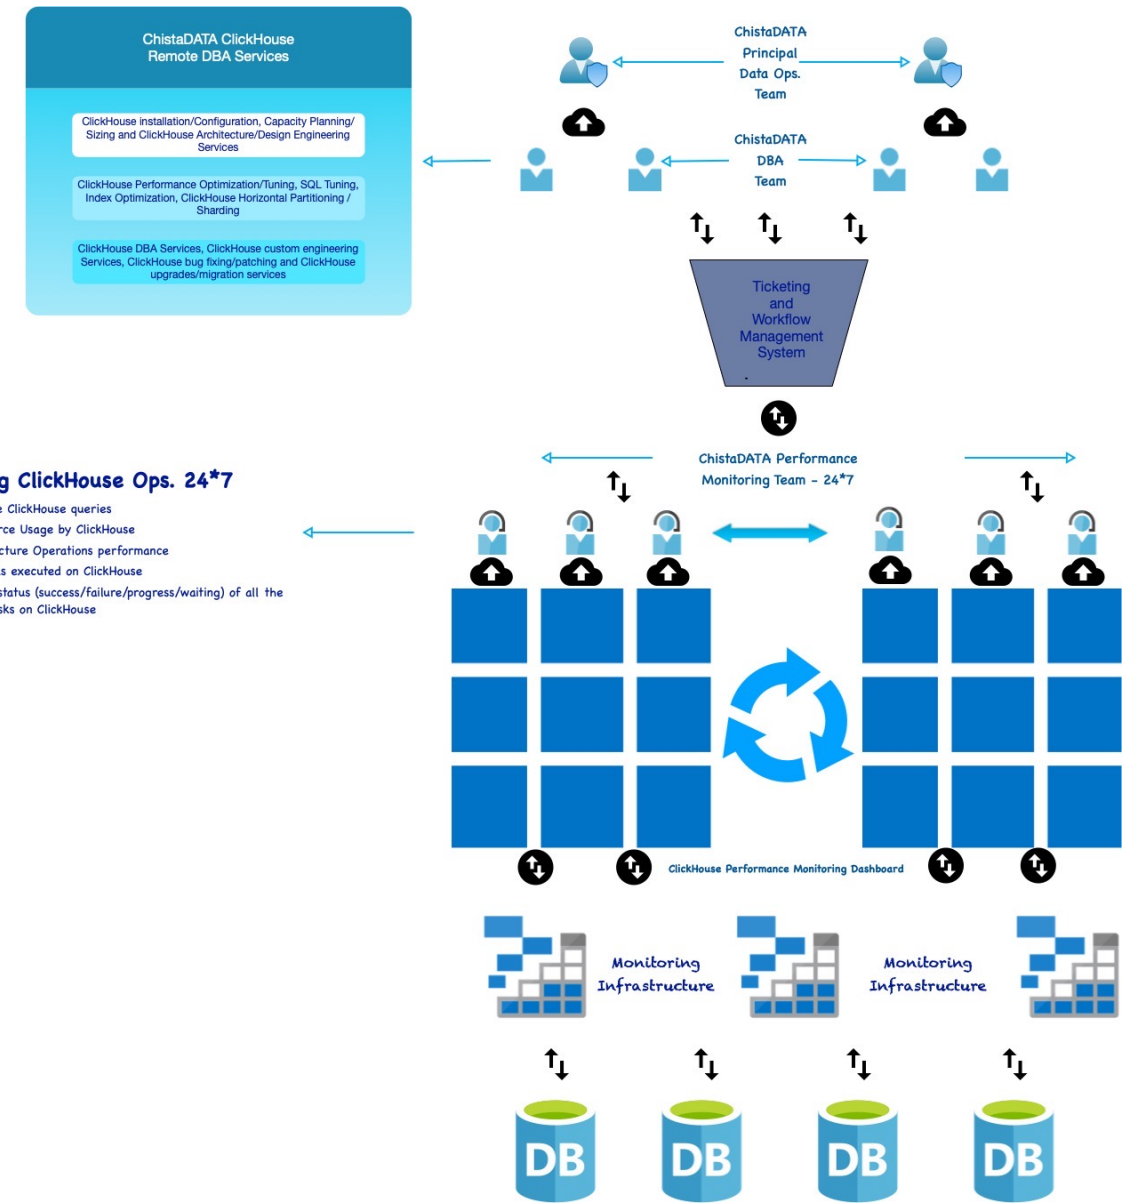
Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

How ChistaDATA can help you in building a web-scale real-time streaming data analytics using ClickHouse?

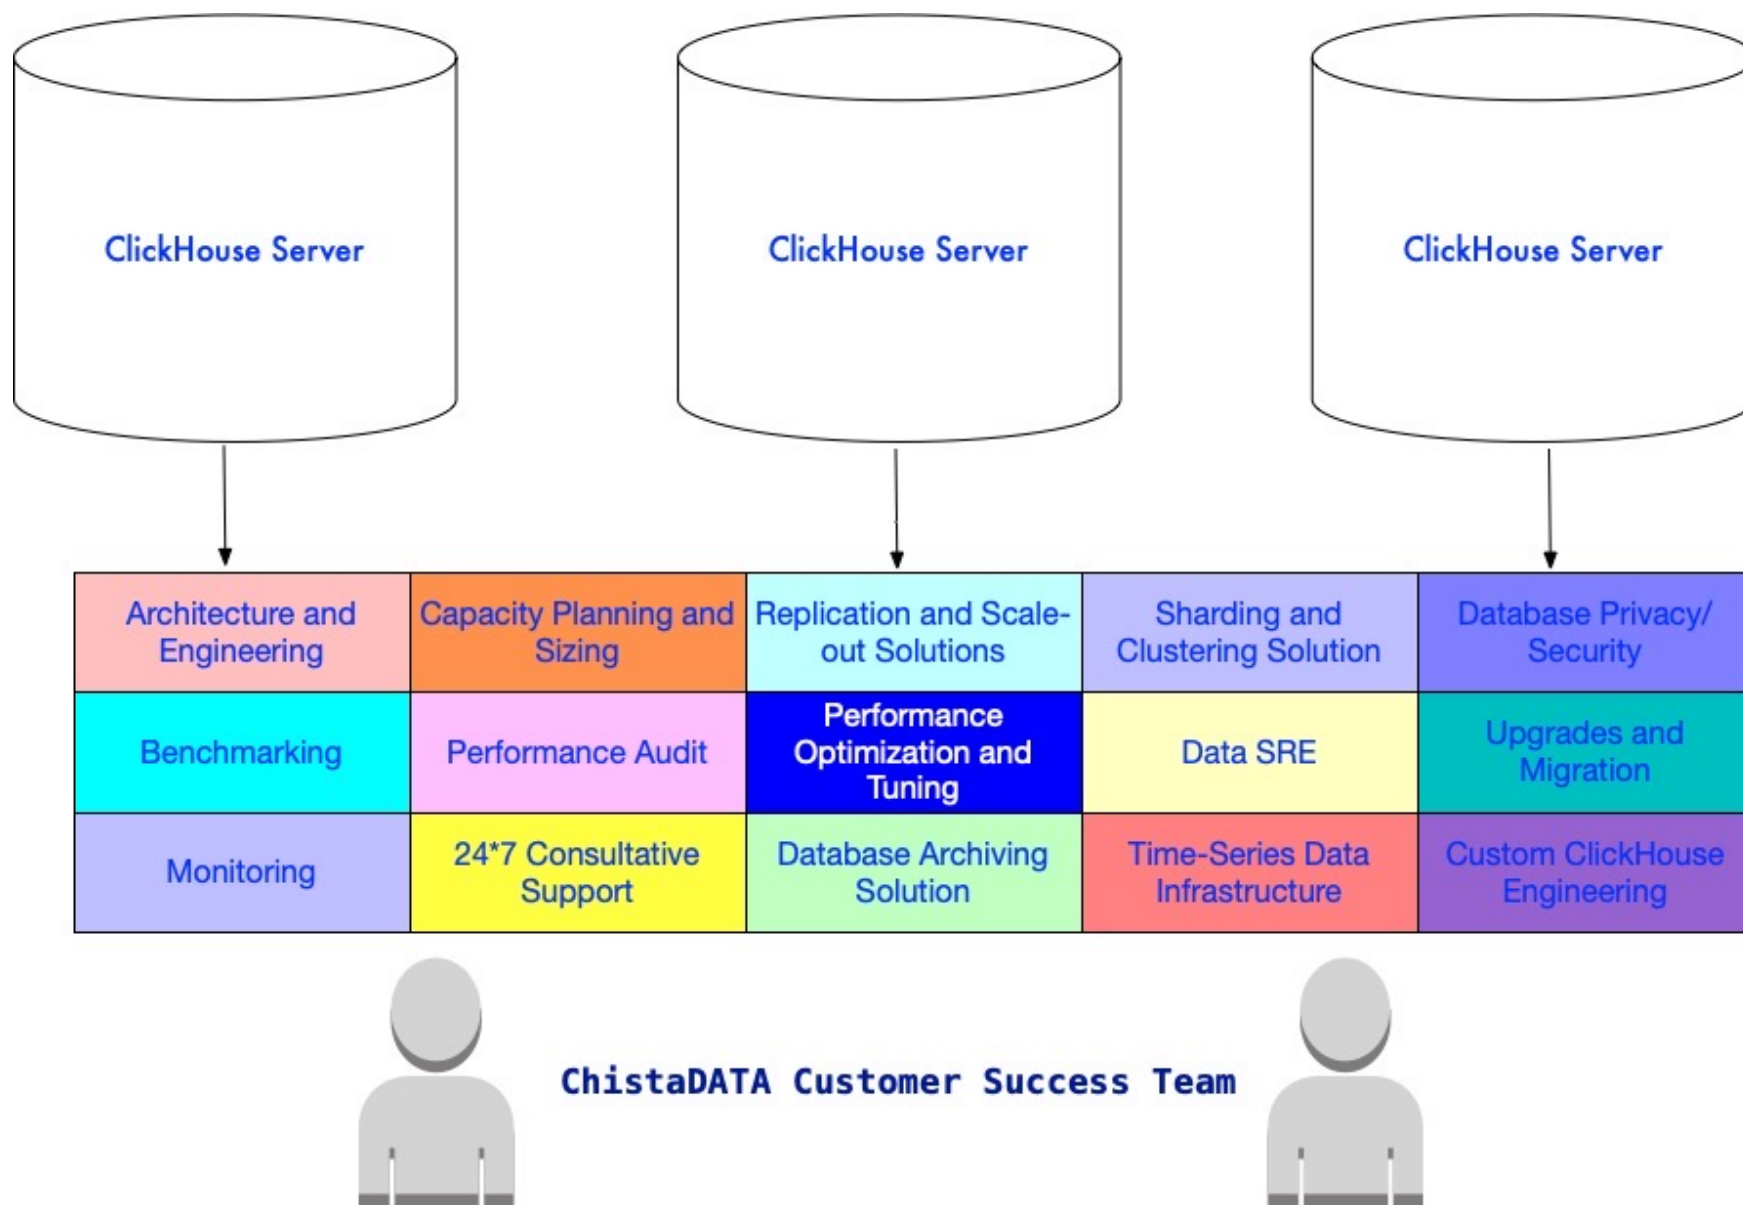
- ChistaDATA Server for ClickHouse (and all tools essential for Data Ops. @ Scale) will be Open Source (100% GPL forever) and free. We are committed to helping corporations build Open Source ColumnStore for high-performance Data Analytics.
- Global Team available 24*7 for ClickHouse Consultative Support and Managed Services.
- Our team has built and managed Data Ops. Infrastructure of some of the largest internet properties. We know very well the best practices for building optimal, scalable, highly reliable and secured Database Infrastructure @ scale.
- Lean Team Culture – Startup-friendly and specialists in DevOps. and Automation for Database Systems Maintenance Operations.

Monitoring ClickHouse Ops. 24*7

1. Most expensive ClickHouse queries
2. System Resource Usage by ClickHouse
3. Disk Infrastructure Operations performance
4. High-Risk SQLs executed on ClickHouse
5. Checking the status (success/failure/progress/waiting) of all the scheduled tasks on ClickHouse

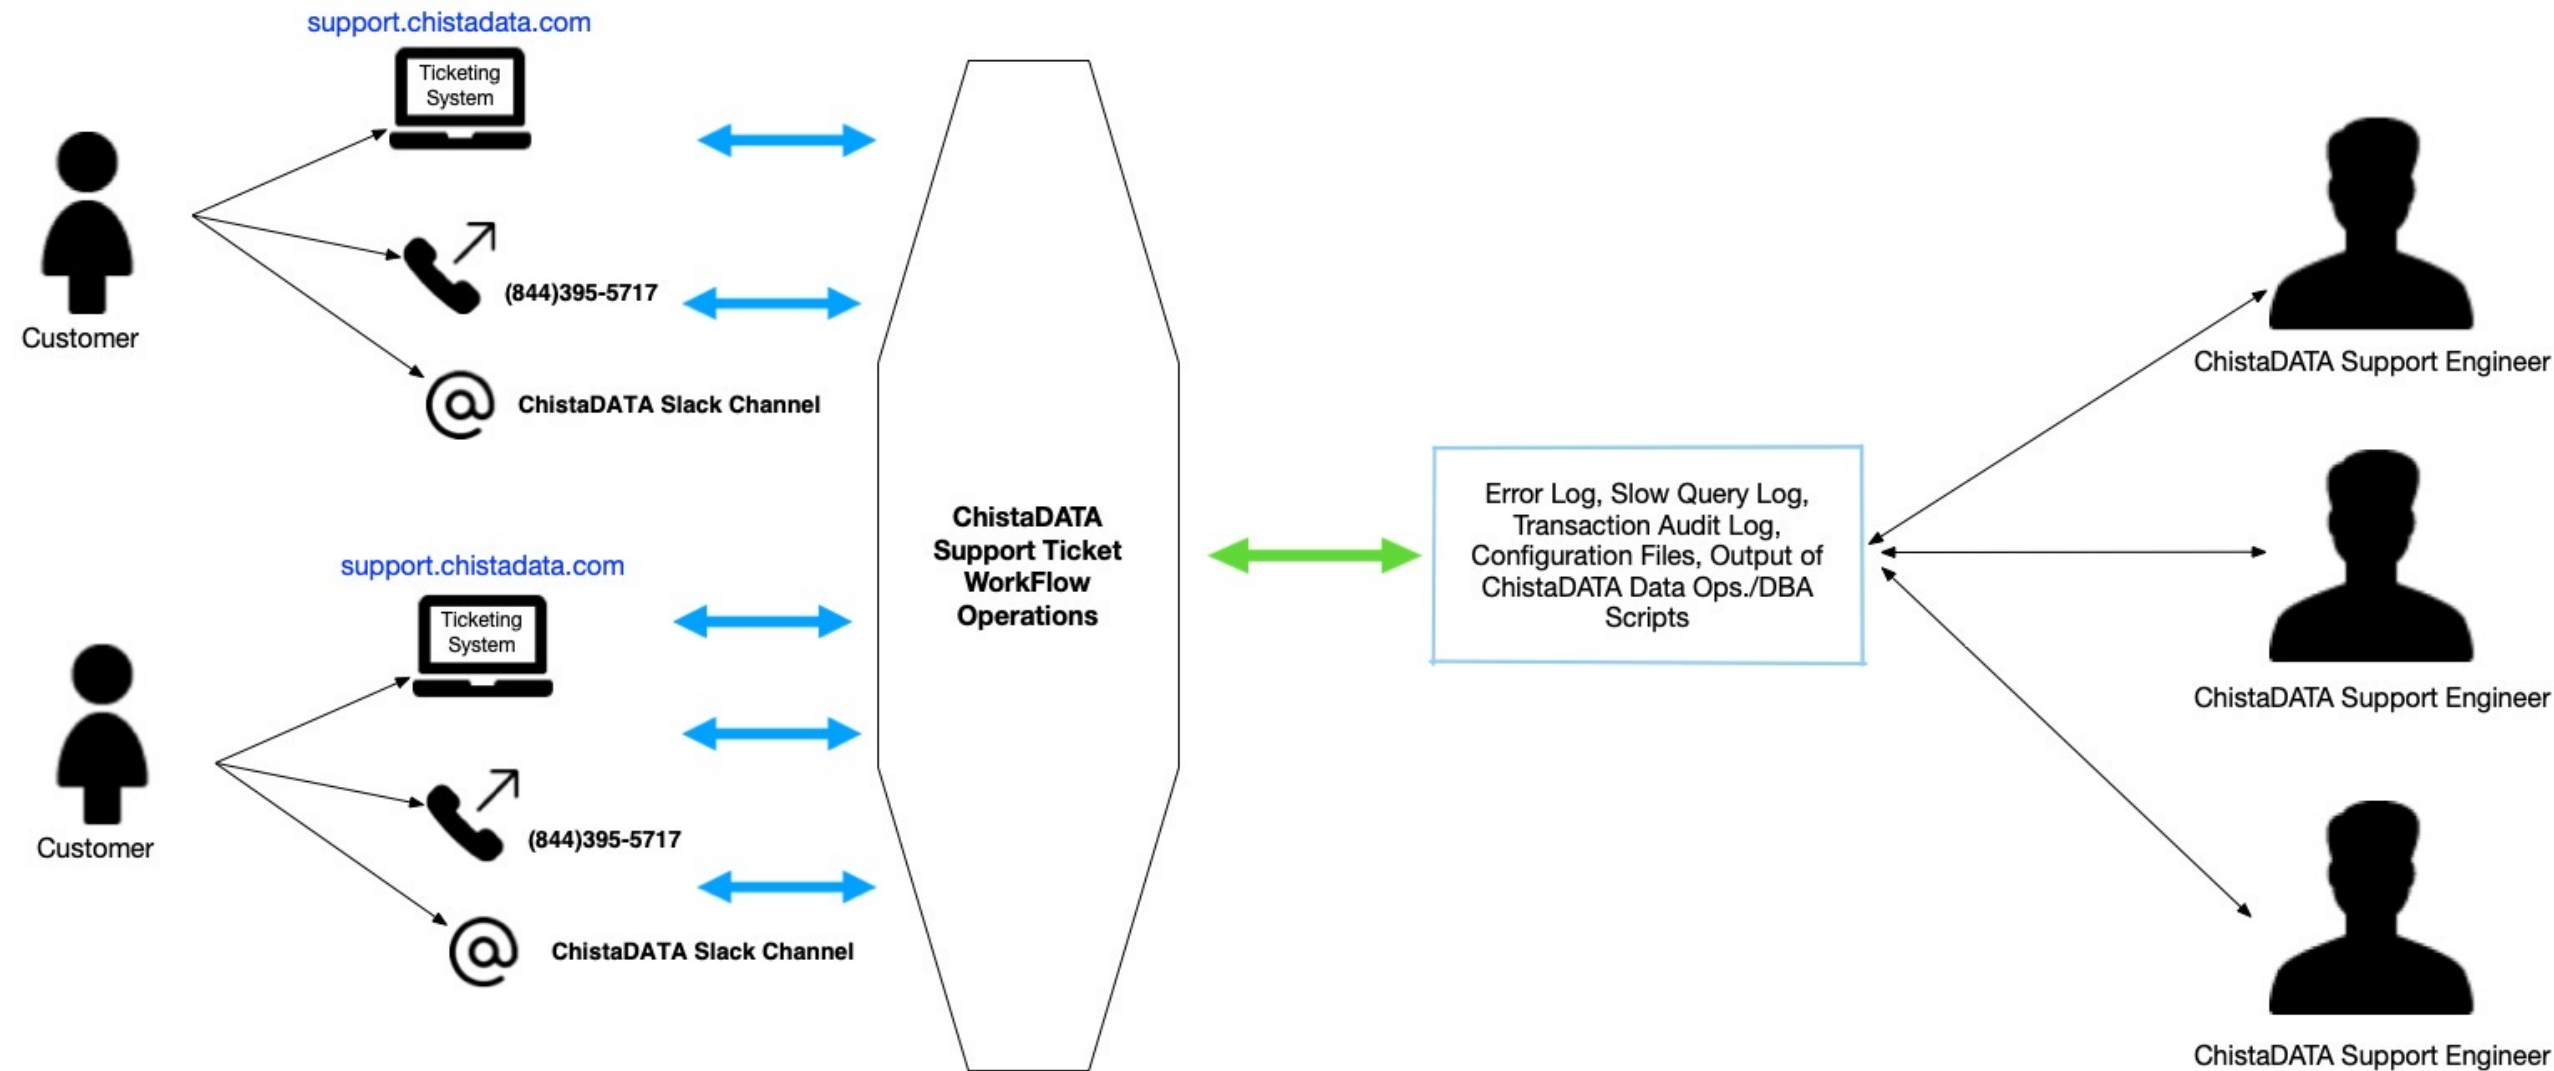


Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP




Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

Consultative Support (24*7) for ClickHouse – ChistaDATA



Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

ClickHouse Performance Troubleshooting



Monitoring ClickHouse Performance

Most Expensive
Linux Process

Disk IO Statistics

Current RAM
Distribution/
Consumption

Historical
Performance
Statistics

Response Time Per
Query in ClickHouse

Load Generated on
ClickHouse by Active
Sessions / Process

Physical Reads (Total
Bytes Per Sec) on
ClickHouse Server

Physical Writes (Total
Bytes Per Sec) on
ClickHouse Server

Database Growth
Rate and
Current DB Size

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

A partial list of customers (from approximately 500 globally)

- Fox Media
- Twitch
- Home Depot
- Comcast Media Networks
- Viacom
- APC
- Prada
- Shutterfly
- Instacart
- Sabre Corporation
- Verisign
- Tiffany & Co
- Strabag
- Danone
- Asahi
- Castel Group
- Shell Petroleum
- GoDaddy
- Grab
- DataCloud
- Garmin
- ClassPlus
- MPL
- WrestleMania
- Kirin
- Blue Dart
- LibLynx
- SriLankan Airlines
- Callaway Golf
- YONEX
- BharatPe
- Dean Food Group
- Flip.id
- Fitbit
- Sprint Global
- Sony
- Chevrolet
- BankBazaar
- MIX
- AOL
- eBay
- Forbes.com
- National Geographic
- Apigee
- PayPal
- Yahoo
- The New York Times Company
- Priceline.com
- Freshdesk
- OLA
- Flipkart
- Paytm
- PetSmart
- Orange Communications
- ESPN
- Go-Jek
- Thomson Reuters Corporation

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP

MinervaDB Inc. contacts – Sales & General Inquiries

| Business Function | Contact |
|---|--|
| CONTACT GLOBAL SALES (24*7) | (844) 588-7287 |
| MinervaDB Inc. FAX | +1 (209) 314-2364 |
| MinervaDB Inc. Email - General / Sales / Consulting / Support | contact(at)minervadb.com |
| CORPORATE ADDRESS: CALIFORNIA | MinervaDB Inc., PO Box 2093 PHILADELPHIA PIKE #3339 CLAYMONT, DE 19703 |
| CORPORATE ADDRESS: CALIFORNIA | MinervaDB Inc., 340 S LEMON AVE #9718 WALNUT 91789 CA, US |
| CORPORATE ADDRESS: HOUSTON | MinervaDB Inc., 1321 Upland Dr. PMB 19322, Houston, TX, 77043, US |

Enterprise-class Open Source Database Infrastructure Operations for OLTP and OLAP