

# Enterprise-Class 24\*7 Consultative Support and Managed Services for ClickHouse by ChistaDATA Inc.

---

We are a full-stack ClickHouse infrastructure operations Consulting, Support and Managed Services provider with core expertise in performance, scalability and data SRE. Based out of California, Our consulting and support engineering team operates out of San Francisco, Vancouver, London, Germany, Russia, Ukraine, Australia, Singapore and India to deliver 24\*7 enterprise-class consultative support and managed services. We operate very closely with some of the largest and planet-scale internet properties like PayPal, Garmin, Honda cars IoT project, Viacom, National Geographic, Nike, Morgan Stanley, American Express Travel, VISA, Netflix, PRADA, Blue Dart, Carlsberg, Sony, Unilever etc.

## 👉 What is ClickHouse?

ClickHouse is an open source columnar database management system built for web-scale, real-time streaming data analytics using SQL queries. ClickHouse is capable of delivering optimal, horizontally scalable, fault-tolerant and highly available data analytics solutions for planet-scale internet / mobile properties and the Internet of Things (IoT). Modern hardware efficient ClickHouse columnar storage format allows fitting more hot data in RAM, which leads to shorter response times. ChistaDATA provides ClickHouse Consulting and Support (24\*7) to deliver optimal, scalable and highly available web-scale data analytics platforms.

## 👉 Quick facts on ClickHouse

- Built for Massively Parallel Processing Systems, Large/complex queries can be run in parallel with minimal or no effort, The modern hardware infrastructure ready!
- Data compression – ClickHouse supports data compression and this improves query performance.
- Horizontally scalable columnar database system – ClickHouse is built for web-scale data analytics, Data can be replicated across several ClickHouse Shards. ClickHouse has a distributed database analytics-ready columnar database system.
- In ClickHouse, data is not just stored by columns but is also processed by vectors to achieve high CPU performance.
- Web-Scale data analytics-ready – Primary keys are allowed, The data extraction for specific clients through Metrika counter over a specific time range makes low latency query analytics possible.

- Flexible aggregation – Aggregate functions for partial data with approximated calculation (minimal data retrieval option). Random keys aggregation instead of all keys for higher accuracy using minimal resources.
- Maximum availability and self-healing – Asynchronous multi-master replication with auto-failover capabilities.
- SQL based – ClickHouse supports SQL, JOINS, subqueries including FROM, IN, JOIN clauses; and scalar subqueries are allowed. Correlated subqueries are not allowed.

## ☛ Why do we recommend ClickHouse over many other columnar database systems ?

- **Compact data storage** – Ten billion UInt8-type values should exactly consume 10GB uncompressed to efficiently use the available CPU. 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.
- **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.
- **Massively Parallel Processing** – ClickHouse is capable of Massively Parallel Processing very large/complex SQL(s) optimally and cost-efficiently
- **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.
- **ClickHouse support Primary Key** – ClickHouse permits real-time data updates with a primary key (there will be no locking when adding data). Data is sorted incrementally using the merge tree to perform queries on the range of primary key values.
- **Built for statistical analysis and supporting partial aggregation** – ClickHouse is a statistical query analysis-ready columnar database store supporting aggregate functions for approximated calculation of the number of various values, medians, and

quantiles. ClickHouse supports aggregation for a limited number of random keys, instead of for all the keys. You can query on a part (sample) of data and generate approximate results reducing disk I/O operations considerably.

- **Supports SQL** – ClickHouse supports SQL, Subqueries are supported in FROM, IN, and JOIN clauses, as well as scalar subqueries. Dependent subqueries are not supported.
- **Supports data replication** – ClickHouse supports asynchronous multi-master and master-slave replication.

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

- **Consulting** – We are experts in building optimal, scalable (horizontally and vertically), highly available and fault-tolerant ClickHouse powered streaming data analytics platforms for planet-scale internet / mobile properties and the Internet of Things (IoT). Our elite-class consultants work very closely with your business and technology teams to build custom columnar database analytics solutions using ClickHouse.
- **Database Architect services** – We architect, engineer and deploy data analytics platforms for you. We will take care of your data analytics ecosystem so that you can focus on business.
- **ClickHouse Enterprise Support** – We have 24\*7 enterprise-class support available for ClickHouse, Our support team will review and deliver guidance for your data analytics platforms architecture, SQL engineering, performance optimization, scalability, high availability and reliability.
- Pay only for hours we have worked for you, This makes us affordable for startups and large corporations equally.
- ClickHouse Training

## 👁️ **ClickHouse Consulting Plans (we do both on-site and remote ClickHouse consulting) from ChistaDATA Inc.**

If you are building a web-scale columnar database systems analytics and your business demands on-site ClickHouse consultants, We are available on short notice. We work very closely with your team on-site guiding them both strategically and technically on building optimal, scalable and highly available ClickHouse database infrastructure operations.

<b>On-Site ClickHouse Consulting</b>	<b>Rate ( plus GST / Goods and Services Tax where relevant )</b>
Per-Diem	US \$350 / hour

We can do almost everything remote on ClickHouse, This includes performance, scalability and high availability. Our technical account manager will be working very closely with your team to understand the goals and build short/long-term deliverables managing ChistaDATA ClickHouse Consultants.

<b>Remote ClickHouse Consulting</b>	<b>Rate ( plus GST / Goods and Services Tax where relevant )</b>
Per-Diem	US \$250 / hour

## ☛ ChistaDATA Subscription Plans for Startups and SMEs (Small and Medium Enterprises)

If you are a startup, We have flexible ClickHouse Managed Services options available:

Avg. Hours / Month	Quarterly ( plus GST / Goods and Services Tax where relevant )	Six-Monthly ( plus GST / Goods and Services Tax where relevant )	Annually ( plus GST / Goods and Services Tax where relevant )
4	US \$2,100.00	US \$4,200.00	US \$8,400.00
8	US \$3,360.00	US \$6,720.00	US \$13,440.00
12	US \$3,780.00	US \$7,560.00	US \$15,120.00
16	US \$4,200.00	US \$8,400.00	US \$16,800.00
20	US \$4,900.00	US \$9,800.00	US \$19,600.00
24	US \$7,000.00	US \$14,000.00	US \$24,500.00
28	US \$9,100.00	US \$18,200.00	US \$28,000.00
32	US \$10,500.00	US \$21,000.00	US \$31,500.00
36	US \$14,000.00	US \$28,000.00	US \$42,000.00
40	US \$17,500.00	US \$34,500.00	US \$49,000.00

## ☛ ClickHouse Enterprise Support (24\*7)

You get access to our seasoned ClickHouse support team 24\*7 for a fraction of the cost of hiring a full-time Sr. level ClickHouse consultant. We will help you in building a planet-scale data analytics platform using ClickHouse which is optimal, scalable and highly available.

ChistaDATA ClickHouse Enterprise Support	Rate ( plus GST / Goods and Services Tax where relevant )
Unlimited ClickHouse Instances	US \$25,000 / Year

- **Enterprise-Class ClickHouse Support**

- Technical Account Manager to clearly understand your business goals and orchestrate our support operations.
- 30 Minute Response Time on Severity 1 (Urgent) Issues.
- 10 Named Customer Contacts.
- Support channels – Phone, Email, Slack, Skype, Google Hangouts and Phone.
- **Technical support — 30 minute response time (S1)**
  - Support -levels – We have a very well defined support infrastructure operations function:
    - **Severity 1**– Immediate attention is needed, The customer’s business is severely impacted and database infrastructure is unavailable.
      - **Response time (SLA) – 30 minutes.**
    - **Severity 2**– Customer database infrastructure is available (up and running) but performance/scalability issues are directly impacting business.
      - **Response time (SLA) – 12 hours.**
    - **Severity 3**– Low impact situation, Customer business and production infrastructure is functioning normally, but the problem is impacting the development ecosystems, also causing a delay in production deployment.
      - **Response time (SLA) – 24 hours.**
    - **Severity 4**– Low to no impact situation, It is more about knowing the features and capability of components before considering the adoption.
      - **Response time (SLA) – 48 hours.**

- **DBA Support**

- Recommendations for database architecture and design.
- Recommendations for optimal SQL engineering.
- Recommendations for ClickHouse Performance optimization and tuning.
- Recommendation for index design, optimization and usage.
- Recommendations for ClickHouse backup and disaster recovery.
- Recommendations for ClickHouse high availability and auto-failover.
- Recommendations for ClickHouse data archiving and partitioning.
- Recommendations for ClickHouse maintenance operations.