FIRSTLOOK

RTInsights Research

Stratola
Say hello to strategy

Origin-to-Edge Observability in a multi-CDN world

Today, with customer or user experience being the #1 use case to address for practically all companies worldwide, delivering content to the end user promptly is critical. Organizations have been leveraging Content Delivery Networks (CDNs) to address this problem. Whether it be about improving website load times, reducing bandwidth costs, increasing content availability and redundancy, or even improving website security, CDNs can be a great solution. But, more often than not, organizations end up using multiple CDNs from different vendors. This also means they need a monitoring solution or a comprehensive observability platform for troubleshooting issues. The complexity of troubleshooting becomes multi-fold when they have multiple CDNs.

The Need for Comprehensive Observability

In a multi-CDN environment, traditional monitoring solutions often fall short. They may provide insights into individual CDNs but cannot offer a holistic view of the entire content delivery chain – from origin servers to the edge and finally to the end-user. This is where the concept of origin-to-edge observability becomes crucial.

Origin-to-edge observability refers to gaining visibility into every step of the content delivery process across all CDNs and infrastructure components. This level of insight is essential for several reasons -

- **Performance Optimization:** By understanding the performance characteristics of each CDN in different scenarios, organizations can make data-driven decisions about content routing and CDN selection.
- **Troubleshooting:** When issues arise, having end-to-end visibility allows teams to quickly identify the root cause, whether it's at the origin, within a specific CDN, or at the edge.
- **Capacity Planning:** Comprehensive observability provides insights into usage patterns and trends, enabling more accurate capacity planning and resource allocation.
- **Cost Management:** With visibility into the performance and usage of each CDN, organizations can optimize their CDN strategy to balance performance and cost-effectiveness.
- **SLA Management:** End-to-end observability allows organizations to verify that CDN providers are meeting their service level agreements (SLAs) and to hold them accountable when they fall short.

Challenges in Achieving Origin-to-Edge Observability

While the benefits of origin-to-edge observability are clear, achieving this level of insight in a multi-CDN environment presents several challenges -

- **Data Volume and Variety:** CDNs generate vast amounts of log data in various formats. Collecting, processing, and analyzing this data in real time requires robust data handling capabilities.
- **Data Correlation:** Linking log data from different CDNs and correlating it with origin server metrics and end-user experience data is a complex task that requires advanced analytics capabilities.
- **Real-Time Processing:** To be truly effective, observability solutions must process and analyze data in real time, allowing for immediate detection of issues and rapid response.
- **Scalability:** As traffic volumes fluctuate and new CDNs are added to the mix, observability solutions must be able to scale seamlessly to handle increased data loads.
- **Visualization and Reporting:** Presenting complex, multi-dimensional data in an easily understandable and actionable way is crucial for effective decision-making.

Organizations need a comprehensive solution that combines advanced data processing capabilities with powerful analytics and visualization tools to address these challenges and achieve true origin-to-edge observability in a multi-CDN world. Such a solution should include capabilities such as high-volume data ingestion, real-time data processing, cost-effective long-term data storage, advanced analytics, scalability and flexibility, and high-performance querying. As organizations grapple with these requirements, innovative solutions like Hydrolix are emerging to meet the unique challenges of origin-to-edge observability in a multi-CDN world.

About Hydrolix

Hydrolix is a cutting-edge streaming data lake optimized for log-intensive use cases, offering both real-time and historical analytics at scale. Designed as a cloud-native system, Hydrolix makes high-volume log use cases much more cost-effective, allowing organizations to keep all their log data readily available long-term without being forced to discard it or move it to less accessible tiered storage.

At its core, Hydrolix is built on several key tenets -

- Decoupled storage for cost-effective, long-term data retention
- Streaming real-time log processing at scale with transforms for data enrichment and normalization
- Advanced compression that reduces storage footprint by 20x-50x while optimizing query performance
- A fully indexed, partitioned datastore offering sub-second query latency even on trillion-row datasets

Hydrolix's architecture combines the structured data and high performance typically associated with traditional databases with the flexibility and cost benefits of a data lake. This unique approach provides "hot" storage performance at the cost of "cold" storage, enabling sub-second query performance regardless of whether you're analyzing data from a minute or a year ago.

Why Use Hydrolix

For organizations facing the challenges of origin-to-edge observability in a multi-CDN world, Hydrolix offers a powerful and cost-effective solution -

- **Scalability:** Hydrolix is designed for scale, making it ideal for enterprises ingesting and analyzing at least one terabyte of log data daily. Its architecture allows each primary subsystem—ingest, query, and storage—to scale independently.
- **Real-Time Processing:** With Hydrolix, you can ingest, transform, and analyze streaming data in real time, enabling immediate insights and rapid response to issues across your multi-CDN environment.
- **Cost-Effectiveness:** Hydrolix's innovative approach to data storage and processing, including advanced compression techniques, results in significant cost savings compared to traditional solutions.
- **Long-Term Data Retention:** Unlike many solutions that force you to discard or archive older data, Hydrolix allows you to keep all your log data readily available long-term, opening up new possibilities for historical analysis and trend identification.
- **High-Performance Querying:** Hydrolix offers sub-second query latency even on massive datasets, allowing you to quickly analyze performance across your entire content delivery ecosystem, regardless of the time range.
- **Flexibility:** As a cloud-native system, Hydrolix integrates seamlessly with existing cloud infrastructure and can adapt to changing requirements, providing maximum flexibility in managing performance and overall cloud spend.
- **Comprehensive Insights:** By making it feasible to retain and quickly analyze vast amounts of log data, Hydrolix enables deeper insights into your multi-CDN performance, opening up a wide range of use cases from business intelligence to monitoring.

Whether you're a direct enterprise customer or a B2B partner serving enterprises with log-intensive use cases, Hydrolix provides the tools needed to manage, analyze, and derive value from your log data at scale. For partners, Hydrolix offers the opportunity to quickly deliver white-label products to market, creating new revenue streams, offering new features, and saving on cloud costs.

By leveraging Hydrolix for origin-to-edge observability, organizations can gain unprecedented visibility into their multi-CDN environments, optimize performance, reduce costs, and ultimately deliver a superior user experience, all while maintaining the flexibility to adapt to future needs and challenges.

Dinesh Chandrasekhar has been a data management veteran, thought leader, and data practitioner for 30+ years. As the chief analyst and founder of Stratola, he speaks and writes on the latest topics in data-in-motion, real-time analytics, IoT, Observability, GenAI, and more. Follow his work at www.stratola.com.