hydrolix

Beyond Silos

UNIFYING MULTI-SERVICE MONITORING IN REAL TIME

FIELD GUIDE | Multi-Service Monitoring

CONTENTS

EXECUTIVE SUMMARY 2

Beyond Silos: Unifying Multi-Service Monitoring in Real Time

CHAPTER 1 5

The Multi-Service Ecosystem: A Hidden Monitoring Challenge

CHAPTER 2 8

One CDN, Many Headaches: The High Cardinality Problem

CHAPTER 3 11

Multi-CDN Monitoring: Once A Near-Impossibility, Now a Reality

CHAPTER 4 14

Beyond CDN: Stitching Together a Holistic View

CHAPTER 5 18

The Right Approach: Key Technical Ingredients for Multi-Service Observability

CHAPTER 6 22

Outcomes and Business Impact: Why It Matters for CTOs

CONCLUSION 26

Reimagining Observability: Seeing the Whole Chessboard

© 2025 Hydrolix, Inc. All rights reserved

EXECUTIVE SUMMARY

Beyond Silos: Unifying Multi-Service Monitoring in Real Time

Modern enterprises run on countless third-party services, from content delivery networks (CDNs) to cloud databases and SaaS applications. Yet many organizations still monitor these services in isolation, if at all. This fragmented approach leaves blind spots.

Consider, for example, the case of a travel website and its Black Friday experience. The company launches new product pages to offer Black Friday specials and discounted travel packages for the holidays, customized for visitor preferences. However, the company fails to warm up its CDN. Cache rates spike, resulting in an increase in the load on the origin server. Then, because of the dynamic nature of the pages, database performance degrades, causing product pages to time out. It's the makings of a Black Friday disaster that could result in immeasurable loss of revenue and damage to the brand's reputation

Each of the services involved in this scenario—the CDN, the origin ecommerce app, and the database—generates high volumes of high-cardinality, timestamped logs that could be used to identify the problem, but if the log data from these multiple sources are siloed or sampled, it can take hours to determine the issue. If, on the other hand, you have a system that correlates the data by time and page or item SKU, the problem—cache hitv rates—becomes immediately apparent.

© 2025 Hydrolix, Inc. All rights reserved

Too often, issues that span multiple services go undiagnosed until they impact customers. It's time to challenge the conventional thinking that believes you *can't* monitor everything as one cohesive system.

Multi-CDN monitoring offers a proof point that even the hardest visibility problems can be solved. If we can successfully monitor multiple CDN providers in real time—a task once deemed nearly impossible—what else might be possible? This eBook explores how multi-CDN observability was achieved and why it signals a larger shift toward holistic, real-time monitoring across *all* the services your business depends on. We'll dive into why monitoring even a single CDN is hard (high data volumes and dizzying dimensionality), and why doing it across several CDNs *was considered* impractical without the right approach. Then, we expand outward: imagine full visibility across not just CDNs but every third-party service—your web analytics, advertising networks, security tools, and more—all stitched together into one live view.

The chapters that follow will show how forward-thinking teams are breaking data silos by using new technology to dramatically expand data storage volumes and retention periods at a fraction of the cost. We'll use real-world examples like FOX's streaming Super Bowl and TF1's unified observability journey to illustrate the art of the possible

By the end, you'll see why multi-CDN monitoring is just the tip of the iceberg. With the right strategy, you can gain visibility across all the services that power your digital business. The result isn't only fewer outages—it's a chance to unlock insights and optimizations that were previously out of reach.

What could you discover if all your data sources worked in concert, in real time?

Let's find out.

© 2025 Hydrolix, Inc. All rights reserved

17.4GB/s

PEAK DATA INGESTION

FOX Super Bowl streaming: Real-time log processing without performance degradation

200TB

TOTAL EVENT DATA

Processed during Super Bowl with sub-second query response times

75%

COST REDUCTION

Achieved vs traditional monitoring infrastructure

55,000

REAL-TIME QUERIES

Executed with average response time under 0.5 seconds

5x

DATA GROWTH

Average log volume increase over last 3 years

11x

MARKET GROWTH

Akamai's observability program customer adoption

© 2025 Hydrolix, Inc. All rights reserved

CHAPTER 1

The Multi-Service Ecosystem: A Hidden Monitoring Challenge

The modern CTO lives in a world of multi-service ecosystems. Your applications likely run on a microservices architecture, rely on cloud infrastructure, and integrate dozens of third-party APIs. On the front end, you might use multiple CDNs for content delivery, a web application firewall (WAF) for security, payment gateways for transactions, and external analytics and advertising platforms. These services are the lifeblood of modern digital experiences—they enable rapid development and global scale. But they also introduce a hidden challenge: **how do you monitor an environment with so many independent pieces?**

In the past, monitoring was relatively straightforward—you kept an eye on your own servers and applications. Now, a significant portion of your uptime and performance is in the hands of others. It's common to assume that each vendor will manage their part (and indeed, cloud and SaaS providers do offer status dashboards and some logs). However, trusting vendors' status pages is not enough for mission-critical systems. If an issue arises, you need to know *immediately* and in the context of your entire stack. Relying on an external provider to notify you (or worse, waiting for users to complain on Twitter) is a recipe for disaster.

© 2025 Hydrolix, Inc. All rights reserved

Consider a simple user transaction on an e-commerce site: A customer clicks a Google ad, which brings them to your site via a CDN. They browse products, your site calls third-party APIs (inventory, recommendation engine, etc.), they add items to their cart, and finally use a payment service to check out. This single flow touches maybe a dozen services. If the user experiences a delay or error, how quickly can you pinpoint the cause? Was it a slow CDN node in their region? A buggy response from the recommendations API? Or perhaps a hiccup in the payment gateway? Traditional monitoring might show your application's response time was slow, but it won't tell you *why*. The relevant clues could be buried in CDN logs, A/B testing analytics, or the payment provider's metrics—data sources typically outside your monitoring tools.

Modern digital businesses rely on countless third-party services, creating a hidden challenge: how to monitor an environment with so many independent, siloed pieces.

The conventional approach has been to monitor each component in isolation: use CDN provider dashboards for CDN metrics, use your APM tool for your app, use yet another tool for API performance, and so on. This creates silos of observability. Each team might have part of the picture, but no one sees it all. When everything is green, that's fine. When something breaks, teams scramble to compare notes across systems. It's a slow, painful process that can lead to extended downtime or degraded user experience while the root cause remains unclear.

Why haven't we solved this yet? One reason is that monitoring data from third parties often comes in torrents and in disparate formats. Every service generates logs and metrics in its own way. Trying to pull all that into one place is technically challenging. Many organizations simply haven't attempted it, clinging to a belief that "our tools can't handle that much data" or "it's not our responsibility to monitor X's internals." This mindset is exactly what we need to challenge. In an era where *your* business reputation is on the line regardless of which component fails, you *must* take a more unified approach. As we'll see, the barriers that made holistic monitoring daunting are starting to crumble.

In this chapter, we've surfaced the core problem: enterprises depend on a wide range of services that traditionally aren't monitored holistically. In the chapters to come, we'll drill into a specific example—multi-CDN—to understand why it's hard and how some have cracked the code. If we can learn how to monitor multiple CDNs as one, we can apply those lessons to the rest of our multi-service puzzle. First, though, let's appreciate why even monitoring a *single* CDN can be a formidable task.

© 2025 Hydrolix, Inc. All rights reserved

Traditional Monitoring

Siloed monitoring tools and dashboards

Data sampling required to manage costs

Limited data retention periods

Hours to identify and resolve incidents

Blind spots between service boundaries

Manual correlation across systems

Unified Observability

Single source of truth across services

Full-fidelity data ingestion capabilities

Extended retention at reduced costs

Minutes to incident identification

Complete end-to-end visibility

Automated cross-service correlation

© 2025 Hydrolix, Inc. All rights reserved

CHAPTER 2

One CDN, Many Headaches: The High Cardinality Problem

To grasp the complexity of multi-service observability, it helps to start with a concrete case. Let's take a single CDN—say you use CloudFront or Akamai to deliver your website's static content. Monitoring one CDN might sound simple: you want to track metrics like response times, cache hit rates, error rates, etc. The CDN provider likely gives you a dashboard with some aggregated stats. So what's the big deal?

The challenge lies in the granularity and sheer volume of data. A busy website or streaming service can generate billions of CDN log entries per day, each with dozens of fields (timestamp, URL, status code, user IP, location, device type, etc.). This is high-cardinality, high-dimensionality data—meaning there are countless unique values and combinations. For example, consider just one dimension: URL paths. Even a moderately sized site might have hundreds of thousands of unique URLs. Now multiply that by dimensions like user IDs, IP addresses, geolocations, and so on. The number of possible combinations (e.g. "downloads page, from Germany, using Chrome, on CDN node X, got status 502") is astronomical. Traditional monitoring systems (like time-series databases for metrics) struggle with this because each unique combination can become a separate series or entry, overwhelming the system.

© 2025 Hydrolix, Inc. All rights reserved

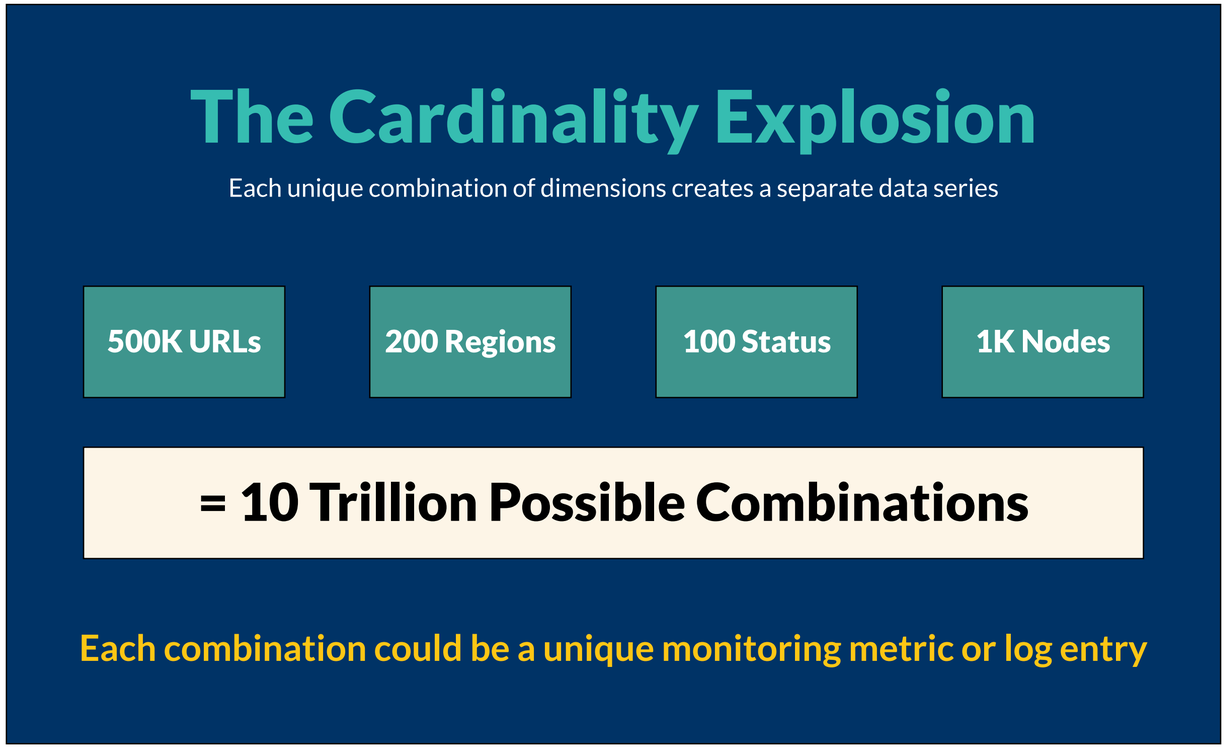

In practical terms, this means monitoring a CDN often forces painful trade-offs. Many teams resort to sampling or aggregation. They might only store 1% of the raw logs, or they roll up data into coarse metrics (e.g. overall errors per minute per region) and discard the details. While this avoids drowning in data, it comes at a cost: you lose visibility into anything outside the top-level aggregates. If a problem is widespread (say, *all* users are getting 500 errors), you'll see it. But if it's narrow—perhaps one ISP in Brazil has an issue connecting to your CDN node—it might get lost in the averages. High cardinality data means the needle can hide in a haystack of dimensions, and if you've pre-aggregated the haystack, you might not even know which needle to look for.

Real-world CDN incidents underscore this challenge. Imagine a scenario where a specific edge server in the CDN network starts misbehaving, causing slightly longer load times for users who happen to connect to that server. The overall CDN latency metric might barely budge (because it's one server among thousands), so your monitoring shows "all clear." Users in that region, however, are having a poor experience. Only by drilling into logs at the granularity of server ID or IP address would you catch it—and only if you retained those logs

© 2025 Hydrolix, Inc. All rights reserved

and can query them quickly. This is why one CDN can give you “many headaches”: the data is rich but unwieldy.

Another headache is real-time detection. It’s not enough to analyze CDN logs after the fact; you want to catch issues as they happen (or ideally, *before* users notice). But processing millions of events per second in real time is non-trivial. Many monitoring systems introduce delays (minutes or even hours) for heavy data, meaning you’re always a step behind live issues. High-volume streaming data can overwhelm databases not built for that ingest rate.

Finally, consider cost and retention. Storing months of full-resolution CDN logs is expensive in traditional systems. Companies often limit log retention (30 days is common) to control costs. This creates a “dark data” problem: if an issue’s clues reside in data older than a month, you’re out of luck. It also hinders trend analysis—you can’t easily compare this year’s Black Friday traffic to last year’s if last year’s detailed logs are gone. According to industry reports, the volume of log data has increased 5x over the last three years on average, pushing companies to jettison data earlier than they’d like. In short, monitoring one CDN is a big data problem that taxes both technology and budgets.

As we move to the next chapter, keep these points in mind. Providing observability for a single high-volume service like a CDN is hard because of high cardinality, massive throughput, real-time needs, and cost constraints. Now, what happens when we double or triple the problem by using *multiple* CDNs? It sounds like a recipe for “near-impossible” observability. Yet, as we’ll see, some organizations have tackled multi-CDN monitoring head-on—and succeeded.

© 2025 Hydrolix, Inc. All rights reserved

CHAPTER 3

Multi-CDN Monitoring: Once A Near-Impossibility, Now a Reality

Adopting a multi-CDN strategy (using more than one CDN provider at the same time) is a smart move for performance and reliability. Many global businesses do it for both economic and operational reasons. Perhaps a CDN has more points of presence in a key market, or offers variable pricing depending on load and time of day. Or perhaps alternate delivery is a high priority, such that if one CDN has an outage or slows down in a region, traffic can be routed through another CDN. Multi-CDN can reduce dependency on any single provider and even optimize costs by load-balancing between them. The flip side, of course, is complexity. If monitoring one firehose of CDN data is tough, monitoring two or three might seem ridiculously unmanageable. This is why, for years, multi-CDN monitoring in real time was regarded as *nearly impossible*.

Let’s break down the problem. With multiple CDNs, you now have multiple streams of telemetry, possibly in different formats. Each CDN has its own logging system, different APIs, and different ways of reporting metrics. Simply normalizing that data into a unified format is non-trivial. Next, the decision logic—many multi-CDN setups involve a real-time decision engine that chooses which CDN to serve each user. Monitoring needs to not only observe each CDN’s performance but also feed back into that decision engine. Ideally, you want to

© 2025 Hydrolix, Inc. All rights reserved

automatically switch traffic if one CDN's metrics degrade. In practice, many companies rely on basic round-robin or geographic splits and perhaps some synthetic monitoring to adjust the percentages. Truly dynamic routing based on real user telemetry requires up-to-the-second insights across providers.

Multi-CDN monitoring, once deemed nearly impossible, is now achievable with platforms capable of real-time ingestion, unified views, and automation, as shown by FOX's Super Bowl LIX success.

So how did anyone manage to solve this? The key is using a new breed of observability platforms capable of ingesting and analyzing huge volumes of log data from multiple sources in real time. One breakthrough example comes from FOX during Super Bowl LIX. The Super Bowl is notorious for being one of the highest-traffic streaming events every year, with millions of concurrent viewers. FOX employed a multi-CDN architecture to ensure a resilient stream, using four separate CDN vendors to ensure smooth deliverability during the event. To monitor it, they didn't rely on guesswork—they leveraged a unified log analytics approach that could handle the deluge from all CDNs simultaneously.

The results were remarkable: during the 2025 Super Bowl, FOX ingested an estimated 17.4 GB *per second* of log data from the various CDNs and other components, in real time, without losing a beat.

In total, nearly 200 terabytes of log data were streamed during the game, yet engineers could query that data interactively with sub-second latency—over 55,000 queries ran with an average response time under 0.5 seconds. This is an almost unfathomable scale by traditional standards, and it demonstrates that multi-CDN observability is not only possible but can be done *live* and under record-breaking traffic loads.

The payoff for FOX was huge. They achieved a near-flawless streaming experience for a record audience, with virtually no buffering or downtime throughout the broadcast. When you consider the stakes—the Super Bowl commands prime ad dollars and fan attention—the ability to see problems and reroute around them in real time is critical.

FOX is to be commended as well for their "origin-to-edge observability" approach — monitoring several key services beyond CDNs to assess the entire technical ecosystem supporting the live stream — and their investment in automated decisioning. In the lead-up to the event, they improved their multi-CDN routing engine logic by incorporating more real-time data; in addition to analyzing CDN logs, FOX's team combined data from video player QoE (Quality of Experience) and synthetic monitors from different geographies to paint a full picture of the stream's health. By funneling this data into

© 2025 Hydrolix, Inc. All rights reserved

their decision engine, they could make smarter choices on which CDN should serve each user session. If a viewer had a buffering issue, the monitoring team could dig into whether it was the origin, a CDN node, or the player software at fault. Every part of the delivery chain was observable.

In addition, FOX used executive-level and operational-level dashboard views to keep a keen eye on throughput metrics such as manifest and segment delivery latency, error trends, and traffic steering capabilities. This is a great example of monitoring not just as a passive observer but as an active participant in control systems. The monitoring system's output directly drove adjustments to the infrastructure in real time—truly closing the loop.

The FOX Super Bowl story is a dramatic example, but it's essentially proof by extreme case. If you can monitor that event across *multiple* CDNs, you can monitor just about anything. The key ingredients were: a platform capable of ultra-high ingest rates and query, a unified view across sources, and automation to act on the insights. This turns the "near-impossible" into a solved problem.

For those of us not running Super Bowl-sized events, the lesson is still relevant. Multi-CDN observability can ensure that your users always get the best experience by dynamically adapting to issues. But beyond that, it opens minds: if we can unify data across CDNs, why not unify data across all our third-party services? In the next chapter, we will zoom out and explore how the same principles can apply to a much broader scope—think *ecommerce + CDN + WAF*, or *ads + user behavior + backend*, etc. Multi-CDN is just one piece of a bigger vision of holistic, cross-service monitoring.

© 2025 Hydrolix, Inc. All rights reserved

CHAPTER 4

Beyond CDN: Stitching Together a Holistic View

The success of multi-CDN monitoring inspires a bold question: what else could we achieve with visibility into logs correlating multiple services?

Imagine having not just a “single pane of glass” but a single source of correlated data that shows your infrastructure metrics as well as your CDN performance, web analytics, security events, and even business KPIs—all in real time, correlated and contextualized. This is the essence of multi-service observability. It’s about stitching together insights across the many services and data streams that comprise your digital business.

Let’s walk through a few scenarios to illustrate the possibilities:

- **E-commerce, CDN, and WAF:** Suppose your online store uses a CDN for content and a cloud WAF for security. One day, you notice a dip in conversion rates (fewer customers checking out) in a certain region. There’s nothing obviously wrong with your application servers. With siloed monitoring, you might be perplexed. But a unified observability system could reveal a chain reaction: a WAF rule was misconfigured and started blocking some legitimate traffic from that region, which caused the CDN to serve more errors (or perhaps the WAF added latency), leading users to abandon their carts. By correlating security logs (WAF alerts) with CDN logs (errors/

© 2025 Hydrolix, Inc. All rights reserved

timeouts) and business metrics (cart abandonment), you pinpoint the issue quickly. In a traditional setup, the security team might not even be aware of the drop in conversions, and the business team might not connect it to the WAF. With holistic monitoring, all these dots connect in one place.

- **Ad Streams and Viewer Behavior:** This one is crucial for media and streaming companies. If you serve video content with ads (think of a streaming TV app or an AVOD platform), you care about both the ad delivery and the viewer experience. What if a particular ad campaign is causing viewers to drop off (maybe the ad fails to load or is long and viewers get frustrated)? By correlating ad server logs/metrics with video player analytics and CDN performance, you might discover that an ad tag from a third-party ad network is occasionally hanging, which in turn makes users give up on the content.

  One real example from the previous chapter: the FOX team correlated their ad tech data with QoE data during Super Bowl LIX and found that certain ad slots were leading to rebuffers, informing them which ad partner's tags to optimize.

  As described in "Modernizing the AdTech Stack with Next-Gen Data Solutions," another whitepaper from Hydrolix, the AdTech industry must harness and analyze billions of data points each day, often in real time, to optimize user experience and revenue. Only a unified approach can crunch that many signals together. The AdTech market is huge (over $845 billion in 2023, projected to exceed $3 trillion by 2030) and hugely competitive. With that growth comes a need to build differentiated products using the deluge of data. One forward-thinking ad platform adopted real-time log analytics to collect, enrich, and transform data so they can connect the dots between bids, impressions, and user outcomes in order to train AI models for bidding.

- **Business Metrics + Infrastructure Telemetry (Ads and Shopping Carts):** In many companies, marketing data and ops data live under the control of different business units. But consider how powerful it is to merge them. For instance, you run a flash sale campaign via Google Ads and suddenly drive a spike of traffic to your site. A holistic monitoring system could correlate the Google Ads click data with your site's telemetry and shopping cart behavior in real time. Maybe it shows that while clicks are surging (good news from marketing), your CDN cache hit ratio dropped and origin server load spiked (bad news from ops), which slowed page loads and caused a percentage of users to drop off before checkout. By seeing this unfold live, your teams can react immediately: ops can scale up resources or adjust CDN configurations, and marketing can modulate the ad spend if needed. After the fact, this combined dataset is a goldmine: it tells you exactly how ad dollars converted to revenue and where technical bottlenecks affected the funnel.

© 2025 Hydrolix, Inc. All rights reserved

- **Failed Logins and Security Data:** On the security side, consider an app where users occasionally have trouble logging in. Your auth service might log failed login attempts, while your security monitoring (SIEM) logs might separately flag unusual IP addresses or geolocation anomalies. If a spike in failed logins coincides with a surge in traffic from an unfamiliar country or a known malicious IP range, you could be witnessing the early signs of a credential stuffing attack. A unified view that correlates application logs (login failures) with security events (firewall alerts, identity provider logs) could alert you to a security incident much faster than either system on its own. Often, such cross-correlation is done manually (if at all)—an analyst might notice something in the auth logs and then pivot to check the security logs. But a platform that can ingest both streams and automatically highlight the correlation in real time turns hours of analysis into seconds. In the context of "dark data," security teams often collect far more logs than they actively analyze; integrating them with operational data helps ensure important clues don't stay hidden.

These scenarios all share a common thread: the need to break down data silos and look at multiple sources together. It's not just about dumping all data into a data lake and hoping for insights. It's about real-time, contextual stitching—so that when something happens in one corner of your system, you immediately see its ripples across the other parts.

In summary, *stitching together a holistic view* means you no longer have blind spots where one service ends and another begins. You can follow the user's journey or a transaction's path across every service in play. The payoff is not just faster troubleshooting, but new insights that drive improvement.

A traditional obstacle to this has been organizational as much as technical. Different teams own different tools: the ops team watches the APM dashboard, the marketing team watches Google Analytics, the security team watches the SIEM. Each tool is specialized, and it's challenging to give everyone a unified interface without stepping on toes or causing tool overload. However, the trend is moving toward unified observability platforms that can serve multiple stakeholders with customized views on top of a shared data foundation. The data can live in one place (or be federated logically) while each team sees the slices they care about—and crucially, they can also see the whole picture when needed.

Another traditional obstacle to unified observability is getting service vendors to provide log data. Again, this is not a technical obstacle

© 2025 Hydrolix, Inc. All rights reserved

(as you'll see in the next chapter); instead, it's a challenge caused by outdated business models. In the past, many service vendors have been reluctant to provide comprehensive log data, citing concerns about revealing trade secrets or Personally Identifiable Information (PII). However, today, customers are demanding transparency and including log sharing in their contracts and service level agreements (SLAs). The outdated log-hogging stance of vendors is giving way to a more progressive approach as service providers discover that log services are a valuable customer retention strategy and a revenue opportunity. Akamai, for example, launched TrafficPeak, its log observability program, in 2023, and achieved an 11x customer growth in the offering's first year of availability.

These mindset shifts we're seeing both internally and externally are game-changers: They mean that you can now ask bigger questions: "How do technical performance and business outcomes intersect? Where can we optimize both together?" That's a new perspective for many organizations, and it's exactly where multi-service observability leads you. But to realize this vision, we must tackle the elephant in the room: the **data deluge** and the traditional cost of keeping all this data. We address that next.

Stitching together a holistic view means correlating all data sources—infrastructure, CDN, web analytics, security, and business KPIs—to eliminate blind spots and enable new insights.

© 2025 Hydrolix, Inc. All rights reserved

CHAPTER 5

The Right Approach: Key Technical Ingredients for Multi-Service Observability

Throughout this eBook, certain themes have surfaced repeatedly as enablers for holistic monitoring. Let's distill those into the key ingredients of the "right approach" to multi-service observability. If you're a technical decision-maker envisioning this capability for your organization, these are the must-haves to look for or build toward:

- **Unbounded Ingest Capacity:** The platform must accept streaming data at high volume from many sources simultaneously. We're talking millions of events per second (think back to FOX's 17.4 GB/s ingestion during Super Bowl) without choking. Whether logs, metrics, or traces—the firehose should flow in real time. If adding a new data source (another CDN, a new microservice, etc.) requires special tuning or cannot be handled in real time, that's a red flag. Hydrolix's streaming data lake, for example, is noted for ingesting massive quantities of data live while keeping costs low.

- **High-Cardinality Query Engine:** It's not enough to store the data; you need to retrieve insights from it on demand. You should be able to slice and dice on any dimension (user ID, video ID, product SKU, IP address, etc.) without pre-indexing each combination manually. The system should use techniques like micro-indexing and distributed querying to give sub-second, interactive query responses even on billions of records. In other words, asking a question of your data (no matter how granular) should feel instantaneous. Traditional brute-force scanning or rigid index limitations won't cut it.

© 2025 Hydrolix, Inc. All rights reserved

Hydrolix, for instance, built on an optimized version of ClickHouse with a custom engine, can deliver this kind of performance, eliminating the need to restrict queries to pre-aggregated metrics.

- **Long-Term, Low-Cost Storage:** As we discussed, keeping data for the long haul is crucial. The right approach leverages cheap storage to retain raw or enriched data for months or years at a sustainable cost. It should automatically compress and organize data so that yesterday's logs and last year's logs are both readily accessible. The goal is to avoid the traditional hot/cold split where anything older than X days is painful to retrieve. A next-gen platform will let you query last year's event with almost the same ease as last minute's event. The economics of log management have to make this feasible. Look for solutions claiming significant cost reduction (e.g., 75% less cost at scale)—and those that specifically advertise eliminating data retention trade-offs.

- **Schema Flexibility and Enrichment:** Data from different services comes in different shapes. Your monitoring platform should be able to ingest various formats (JSON, CSV logs, binary telemetry, etc.) and normalize or join them on common keys (like timestamp, user ID, session ID) on the fly. It should allow enriching one stream with data from another—for example, tagging CDN logs with the name of the CDN provider, or merging in geo-IP lookups for IP addresses, etc. Flexibility is key; you don't want to be locked into a rigid schema that was designed for server logs when you start ingesting marketing data. Some platforms achieve this by storing data in a semi-structured form and letting you run SQL or custom code to combine sources at query time. The ease of integrating a new data source into your observability schema is an important consideration.

- **Unified Portal with Custom Views:** While the backend is unified, the frontend can still offer role-specific dashboards. The right approach provides a single portal where a CTO or SRE can see everything at a high level (the "big board" of system health), while specialists can drill into their area. For example, network engineers might have a dashboard focusing on CDN and DNS performance, security analysts have one for security events and anomalies, product managers have one showing user engagement alongside system metrics. But all of these draw from the same data pool, so when an issue arises, everyone is looking at a single source of truth. No more "my dashboard says it's fine, so it must be your system" finger-pointing—if something's wrong, it will surface across views consistently.

- **Real-Time Alerting and Automation:** Observability isn't just passive; it should trigger actions. The platform must support sophisticated **alerting rules** that can draw on multiple data streams. For instance, you could set an alert that fires if *login failures increase by 50% AND CloudFront cache hit rate drops below 80%*, correlating two conditions. Furthermore, those alerts

© 2025 Hydrolix, Inc. All rights reserved

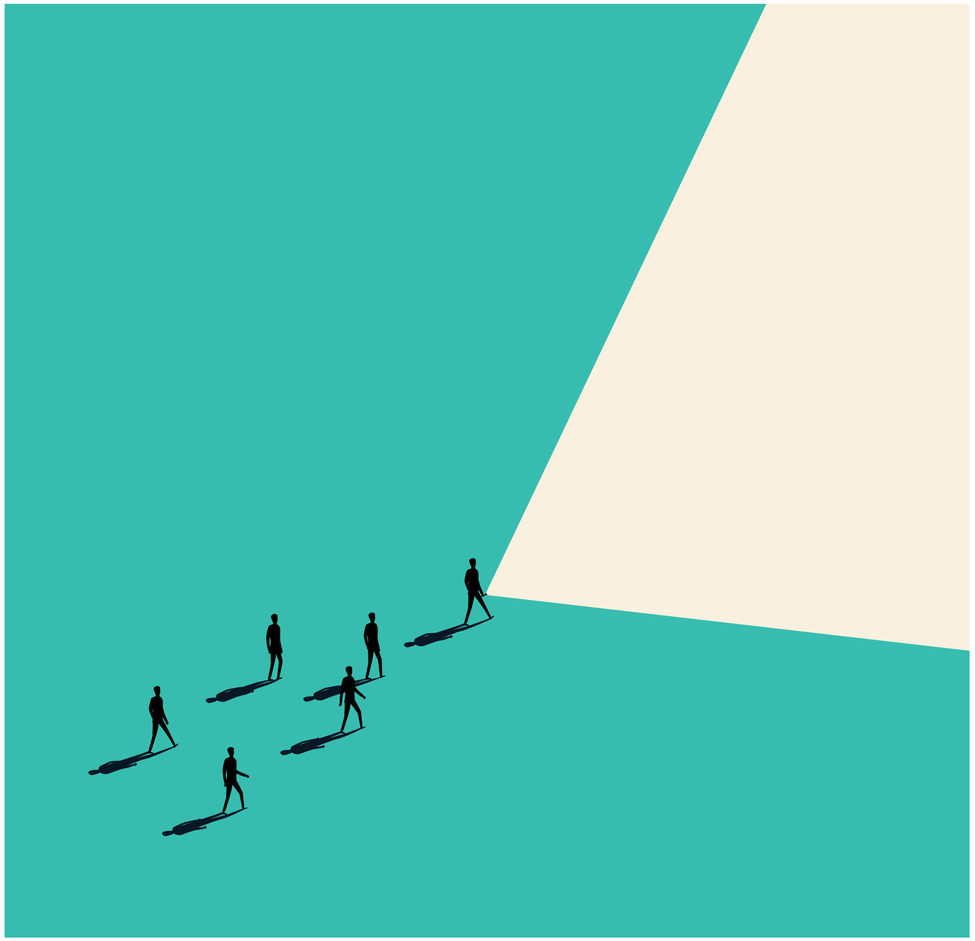

Key technical ingredients for multi-service observability include unbounded ingest capacity, a high-cardinality query engine, long-term low-cost storage, schema flexibility, and real-time alerting.

should ideally plug into automation (AIOps). Perhaps that alert triggers a script to automatically purge a CDN cache or disable a WAF rule. The system could even interface with traffic management—e.g., automatically shifting load from one CDN to another if latency crosses a threshold. Essentially, the observability solution becomes part of the control plane of your architecture.

- **Scalability and Reliability:** It goes without saying, but the observability platform itself must be rock-solid and scalable. If it can ingest 10 TB/day, it should scale to 100 TB/day with more resources but without a fundamental redesign. Cloud-native services (Kubernetes-based or serverless architectures) help here, allowing the monitoring system to expand as your business expands. Also, consider its own resilience—it should be distributed so that no single point of failure or network glitch will cause you to lose visibility when you need it most.

© 2025 Hydrolix, Inc. All rights reserved

One might ask: do such ideal solutions actually exist, or is this a wish list? The good news is they do exist and are already being used by cutting-edge organizations.

In practice, companies like TF1 have benefitted from this approach using the Hydrolix platform. The French broadcaster TF1 unified their observability data (across CDN, streaming, etc.) and as a result were able to reduce the number of monitoring tools it needed, lowering complexity and cost, while actually improving end-user experience (better QoS). They also found they could afford to store more of their data for longer periods without budget issues, eliminating dark data and giving them historical insight. This kind of success story underscores that the right ingredients, when put together, yield real-world benefits: TF1 achieved better customer experience and cost savings simultaneously—a near holy-grail in IT ops.

This example shows that the right approach isn't just theoretical. If you're evaluating solutions, look for evidence like *customer stories, benchmarks, or case studies* that demonstrate these capabilities under real workloads. For instance, ask vendors: *Can you show me a case of ingesting 100 TB of data and querying it in <1s? Can you integrate data from my CloudFront CDN, my Kubernetes logs, and my Salesforce API logs in one query?* The ones that can will be proud to share such numbers.

In conclusion, multi-service observability requires a rethink of tooling. The traditional mix of monitoring tools—each good in its niche—is giving way to unified platforms built on modern data architectures. When evaluating or designing your approach, ensure it checks the boxes above. You want a solution that scales, that doesn't flinch at complexity, and that ultimately *makes your team's life easier.* After all, the goal is not just to collect data for its own sake, but to empower your team to solve problems faster and make smarter decisions. With the right approach in place, you'll be well on your way to achieving that.

© 2025 Hydrolix, Inc. All rights reserved

CHAPTER 6

Outcomes and Business Impact: Why It Matters for CTOs

We've journeyed through the technical landscape of multi-service monitoring—now let's talk about the prize on the horizon: what does this mean for your business and strategy? As a CTO or technical leader, implementing holistic observability is not just a tech project; it's a strategic initiative that can deliver tangible business outcomes. Here are some of the major benefits and shifts in perspective that come with full-spectrum visibility:

1. **Dramatically Improved Customer Experience (CX):** At the end of the day, all this monitoring magic aims to keep customers happy. When you have unified real-time insights, you catch and fix issues *before* they cascade into customer-facing problems. Think of TF1's experience—by unifying observability and reducing tool sprawl, they were able to spot issues quickly and ensure a smoother viewer experience for their streaming services. Similarly, FOX's Super Bowl success meant millions of viewers enjoyed a flawless stream, which boosts brand reputation and trust. For an e-commerce site, this could mean higher uptime and faster page loads during critical sales (no more surprise slowdowns during your big promotion because you have eyes on every component). In an era where users have little patience, this proactive stance directly correlates to revenue retention. Happy users complete transactions, consume more content, and stick around.

© 2025 Hydrolix, Inc. All rights reserved

2. **Faster Troubleshooting and Incident Resolution:** When incidents do happen, unified observability can cut resolution times from hours to minutes. No more war room calls where each team is looking at different charts and arguing over which system is at fault. With a single source of truth, your teams collaboratively zoom in on the chain of events. It's not uncommon for companies to report major reductions in MTTR (Mean Time to Resolution) after adopting a more unified monitoring approach. One reason is the elimination of the "swivel-chair" debugging—switching between tools and context—which wastes precious time. Also, having long-term data means if an incident seems similar to something that happened 6 months ago, you can query that historical event and recall how it was solved (or detect a recurring pattern). Essentially, you build a collective memory in your data.

3. **Cross-Team Collaboration and Alignment:** When everyone can see the impacts across the stack, silos between teams start to break down. The marketing team can observe how a campaign is affecting site performance in real time (and appreciate the ops team's challenges), while the ops team can see how technical issues impact business metrics (and appreciate marketing's stakes). This fosters a culture of shared responsibility for end-to-end outcomes. A unified observability system becomes a communication hub—often the dashboards are up on big screens in the office or shared in

© 2025 Hydrolix, Inc. All rights reserved

chat during incidents, serving as a common reference. Over time, this leads to more integrated thinking: developers consider how their code affects not just the app but also CDN usage and business KPIs; network engineers see how a latency drop might improve user engagement. Everyone starts to speak a bit of a common language grounded in data.

Holistic observability dramatically improves customer experience, accelerates troubleshooting, fosters cross-team collaboration, streamlines costs, enables innovation, and supports strategic decision-making.

4. **Cost Optimization and Efficiency:** It might seem counterintuitive that storing *more* data and monitoring *more* stuff could save money—after all, tools can be expensive—but the key is the new economic model we discussed. By leveraging a platform that lowers storage costs and consolidates capabilities, companies often find cost savings in multiple areas. First, you might be able to retire several point solutions (APM, log management, analytics) and replace them with one unified platform—cutting licensing and operational overhead. TF1, for example, reduced the number of observability tools, which presumably lowered their tooling costs and maintenance effort. Second, with better visibility, you can optimize your use of third-party services. For instance, by analyzing CDN performance and costs across providers, you might adjust traffic allocation to get the best bang for buck (maybe pushing more traffic to the CDN that is performing equally well at lower cost in a region). Or you might find wasteful cloud resource usage by correlating application load with infrastructure spend, leading to rightsizing. One specific angle: log data economics—many companies pay through the nose for log ingestion (in legacy tools) and thus limit what they send; a cheaper solution means you can ingest everything for the same budget, or ingest the same and save money. Either way, the CFO will be happier.

5. **Enabling Innovation and New Services:** When you have comprehensive observability, you build confidence to innovate. You're more willing to adopt new technologies or services because you know you can monitor them properly. For example, you might try out a new CDN provider, or a new microservice architecture, without fear of losing insight, because your observability layer can absorb the new data source easily. Additionally, the rich data you collect can spur new product ideas—it's not just for ops. Perhaps you realize users are using your product in a way you didn't expect, or an underutilized feature becomes evident, by analyzing logs. Some companies even monetize insights from their operational data (with privacy safeguards, of course). For instance, a streaming service could provide content producers with analytics derived from correlating network performance and viewer engagement, creating a value-add service powered by the observability data.

© 2025 Hydrolix, Inc. All rights reserved

6. **Strategic Decision Making:** This is more forward-looking—as observability data starts to merge with business intelligence, CTOs and other execs can use it to guide strategy. You begin to answer questions like: which investments in infrastructure yield the most improvement in user experience? Should we spend more on optimizations or on marketing, based on which drives revenue more? You have data to back these decisions in a holistic way. Moreover, being able to "see everything" means you can anticipate problems. It shifts IT from reactive firefighting to proactive improvement. You might discover latent issues (like that one ISP that has intermittent packet loss to your CDN, affecting a small segment of users) and work with partners to fix them before they grow. This predictive and preventive posture is where everyone wants to be, and unified observability is a step in that direction.

As a CTO, implementing multi-service monitoring is akin to turning on the lights in a room that was dimly lit. Suddenly, you see the corners clearly. Yes, it requires investment—in technology, in time to set up, and in shifting some cultural mindsets—but the return is far more than just technical metrics. It's a stronger business. It's the ability to ensure that when your company makes promises to customers (like "our streaming service won't buffer that big game" or "our website will handle the holiday rush seamlessly"), you can keep those promises because you have the telemetry to manage reality to meet expectations.

We should acknowledge that this journey can start small. You don't have to boil the ocean at once. Many organizations begin with a high-impact slice (say, unify CDN and application logs first) and then expand. Each success builds momentum and buy-in for the next integration. Over time, you find that not only has your tech stack become more observable, but your team has developed a habit of data-driven operation.

In the next and final section, we'll wrap up our discussion and leave you with a parting thought. If you've come this far, you likely see the value proposition. The question remains: how will *you* leverage these ideas to transform observability at your organization? Let's conclude with a look at the road ahead.

© 2025 Hydrolix, Inc. All rights reserved

CONCLUSION

Reimagining Observability: Seeing the Whole Chessboard

Observability is undergoing a renaissance. What used to be about checking server health has evolved into a quest to understand complex, distributed systems holistically and in real time. In this eBook, we used multi-CDN monitoring as a microcosm of that evolution—a challenge that, once solved, opens our eyes to a much broader horizon of possibilities. We saw that by harnessing the right technology and approach, even the most daunting data problems (like ingesting hundreds of terabytes of logs in minutes) can be overcome. More importantly, we learned that doing so yields far more than bragging rights; it delivers concrete business value in customer satisfaction, operational agility, and strategic insight.

If there's one takeaway, it's this: full visibility is achievable, and it's a game-changer. No longer do we have to accept blind spots in our digital operations. We don't have to accept that some data is just too costly to keep or too unwieldy to analyze. The tools and methods now exist to shine a light on every corner of our systems. The question moves from "*can* we monitor everything?" to "*what will we do* now that we can see everything?"

This brings a new mindset for CTOs and technical leaders. It's an opportunity to rethink what "monitoring" means. It's not a mundane IT task—it's the foundation of digital trust. When you

© 2025 Hydrolix, Inc. All rights reserved

Full visibility is a game-changer, moving beyond “can we monitor everything?” to “what will we do now that we can see everything?” This shifts IT from reactive to proactive.

have observability baked in across all services, you operate with confidence. It’s like seeing the whole chessboard versus just a few pieces; you can anticipate moves, weigh trade-offs, and execute strategy far better. Issues that used to be mysterious become clear cause-effect stories told by data. And once the mysteries are gone, you can focus energy on higher-level innovation rather than firefighting in the dark.

We encourage you to take inventory of your current state: How many separate monitoring or logging tools do you have? How much of your architecture is effectively a “black box” to your team right now? How quickly can you detect and diagnose a multi-faceted issue that spans systems? If the answer to any of these makes you uncomfortable, that’s exactly the conventional situation this eBook has challenged. The good news: you don’t have to rip everything out overnight. Start by identifying critical integration points—maybe it’s those CDN logs, or maybe linking your application metrics with business metrics. Experiment with a modern observability platform in a small scope. See the difference it makes. Even a single unified view can start paying dividends in cost and insight.

One thought-provoking idea to leave you with: **What might you discover if you retained *all* your data and correlated *anything* with**

© 2025 Hydrolix, Inc. All rights reserved

everything? This isn't a theoretical question—it's one that leading companies are beginning to answer daily. They're finding patterns in user behavior tied to infrastructure events, revealing security holes before they're exploited, and optimizing complex supply chains of services on the fly. In short, they are turning observability into a competitive advantage. When you fully instrument and observe your business's digital nervous system, you gain the ability to respond faster and smarter than those who don't.

In the coming years, we expect the gap to widen between organizations that embrace this holistic monitoring vision and those that stick to siloed thinking. As systems get more complex and interconnected, partial visibility will be akin to flying with one eye closed. We hope this eBook has sparked ideas and given you a roadmap to pursue true multi-service observability. The technology is here, the case studies prove it works, and the rest is up to bold leadership.

In summary, multi-CDN monitoring showed us that no problem is "too hard" if approached with the right tools. Now, extrapolate that ethos: no aspect of your digital business should remain in the shadows. By unifying observability across all services, you ensure that you're never caught off guard. You'll deliver better experiences, make data-driven decisions, and ultimately build greater trust with your users and stakeholders. That's the new world of observability—one where seeing the whole picture becomes the norm, and with it comes the power to shape outcomes, not just react to them.

© 2025 Hydrolix, Inc. All rights reserved

FIELD GUIDE | Multi-Service Monitoring

› Visit us at hydrolix.io to learn how Hydrolix can turn log data into your competitive edge

› Email us at info@hydrolix.io to arrange a demo

› Watch Hydrolix in action on YouTube