Why One-Way Latency Measures Are Critical For Distributed Databases, Microservices and AI Workloads

October 16, 2024 · clockwork.io/blog/why-one-way-latency-measures-are-critical-for-distributed-databases-microservices-and-ai-workloads/

In today's fast-paced world, a seamless digital experience can make or break your business. When customers interact with your applications, they expect instantaneous responses consistently. Any lag or delay can lead to frustration, diminished satisfaction, and ultimately, a loss of business. Accurate, real-time data on network performance isn't a "nice-to-have"—it's essential. Let's dive into why this matters and can now be realized cost-effectively for any real-world scenario utilizing distributed workloads.

The Real-World Impact of Accurate One-Way Latency

Imagine you're a crypto trader executing high-frequency trades—where even a microsecond delay could result in losses amounting to thousands, if not millions. Similarly, in streaming media, minor delays can lead to user churn, and in ecommerce, they can cause cart abandonment. For any

© 2026 CLOCKWORK. All rights reserved.

payment or transaction processing system, risk analysis must stay ahead of fraudulent activities. In telemedicine, fast, accurate data transfer is critical for real-time surgeries and continuous patient monitoring. Whether managing critical infrastructure—like hospitals, water supplies, power grids, or emergency systems—split-second decisions depend on ultra-low latency and precise one-way latency measurements.

The massive volume, velocity and variety of data used in these applications is powered by complex data pipelines that rely on distributed databases, data streaming platforms and real-time analytics that require fast and highly reliable data transfers to maintain service consistency and timely decision-making.

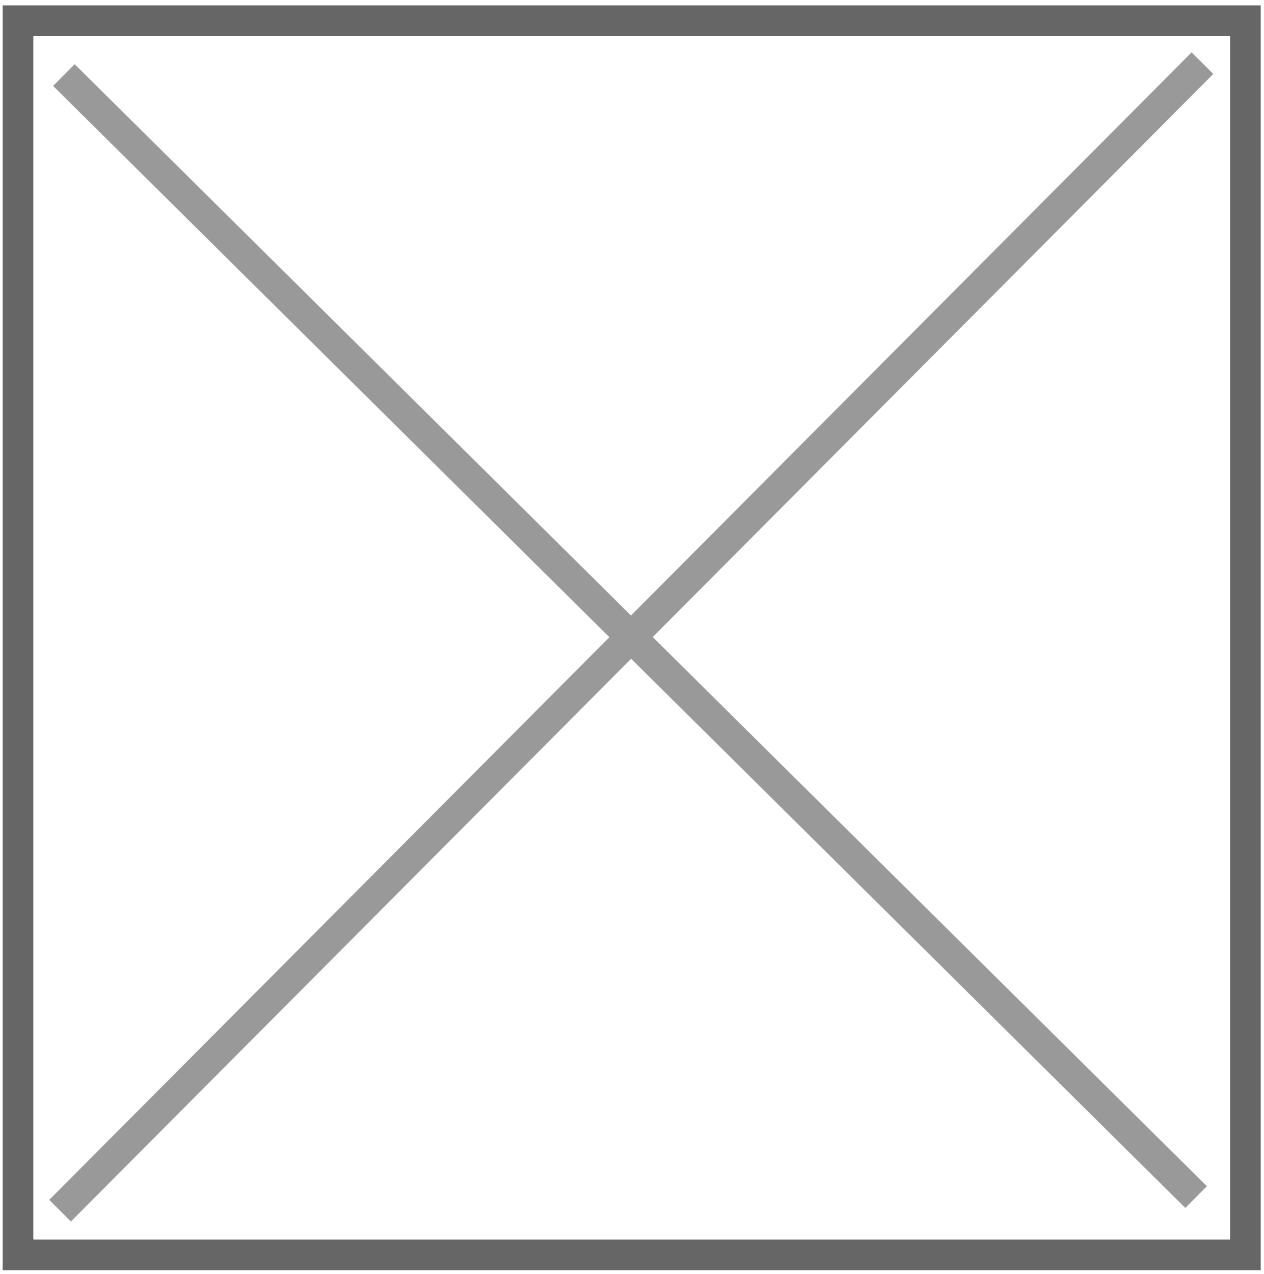
Yet, despite the critical need, accurately measuring latency between network hosts has historically been very cumbersome, relying either on rough estimates or expensive hardware. Conventional approaches simply don't cut it in today's environment.

Why Traditional Methods Fall Short

Historically, businesses tried to gauge one-way latency by estimating half of a round-trip delay (RTT) between sending and receiving nodes in the network or using Network Time Protocol (NTP)

© 2026 CLOCKWORK. All rights reserved.

synchronized clocks. But here's the problem: both methods are inherently flawed. Round-trip estimations are built on false assumptions of network symmetry in both directions and unchanging network conditions, while NTP clocks only give precision within tens of milliseconds—insufficient for latency-sensitive operations. Even Precision Time Protocol (PTP), while more accurate, requires costly hardware, making it impractical for many companies.



© 2026 CLOCKWORK. All rights reserved.

Moreover, cloud environments add another complexity, making it difficult to control network hardware and accurately measure latency. As a result, users often face unpredictable delays and time synchronization issues. So, what can you do to overcome these limitations?

Introducing Clockwork: Accurate, Affordable Latency Measurement

With **Clockwork's Clocksync technology**, measuring true one-way latency between hosts has never been easier or more precise. Imagine getting **sub-microsecond accuracy** without needing any special hardware—all you need is software. No matter how distributed your workloads are—whether on-prem, in the cloud, or across hybrid environments—Clockwork makes real-time latency tracking simple.

Here's why it stands out:

- **Sub-microsecond accuracy**: Precision without the hefty hardware costs.
- **Universal compatibility**: From VMs to Kubernetes pods, it works everywhere.
- **Seamless integration**: Connect effortlessly with monitoring tools like Datadog, Splunk, or New Relic.

© 2026 CLOCKWORK. All rights reserved.

It’s the easy, flexible solution for today’s complex and increasingly distributed digital workloads.

Case Study: How a Crypto Trading Firm Gained a Competitive Edge

Let’s look at a real example. A crypto trading firm specializing in algorithmic arbitrage faced challenges with network latency, as even microsecond delays could lead to missed opportunities. By implementing Clockwork, the firm could measure real-time latency with tens of nanoseconds accuracy, optimizing trade execution and outpacing competitors. No need for specialized hardware or complex deployments, just software integration and precise insights.

As a result, they increased profitability and maintained a technological edge, simply by understanding and controlling their latency.

How to Maintain Great User Experiences Across Distributed Workloads

© 2026 CLOCKWORK. All rights reserved.

In an era where cloud and multi-cloud deployments are the norm, maintaining low-latency performance isn't as simple as controlling infrastructure in a single location. But by accurately measuring one-way latency, you can gain deep insights into your network's performance and optimize accordingly. The benefits include:

- **Optimized Application Response**: Real-time measurements give you the ability to adjust application performance based on actual latency data.
- **Seamless User Experience**: With better insights into latency, you can proactively prevent lag, ensuring smooth, uninterrupted services for your customers.
- **Cost-Effective Scaling**: By identifying where latency is most problematic, you can optimize resource allocation across clouds, minimizing unnecessary expenses

Take Control of Your Latency

Whether you're managing a **distributed database**, running a **data streaming platform**, or ensuring smooth **application performance** in real time, measuring latency accurately will help you optimize network paths, improve application response times, and maintain great user experiences no matter how distributed your workload and where your data travels from on-prem to hybrid cloud to multi-cloud.

© 2026 CLOCKWORK. All rights reserved.

Clockwork helps you:

- **Save costs**: Measure accurately without heavy investments.
- **Gain visibility**: See what's happening in your network, in real-time, down to nanoseconds.
- **Scale easily**: From small setups to global infrastructures, Clockwork scales with your network's needs.

Ready to ensure consistent, low-latency performance? We'd love to hear from you to explore how accurate latency measurements can optimize your distributed workloads!

© 2026 CLOCKWORK. All rights reserved.