

# What are we really measuring?

## The Hidden Complexity of Latency in FX Trading

By Phil Morris, Chief Executive Officer

*In a world where microseconds move markets, **how** we measure latency matters just as much as the latency itself.*

### 1. Understanding Latency Beyond the Number

---

In high-performance FX trading, latency is often presented as a single headline number. But that number only tells part of the story. Behind every latency figure lies a complex journey - across hardware, operating systems, queues and network layers. *How* we measure that journey can drastically change the result.

### 2. The Fast-Food Analogy: Measuring at the Counter

---

Imagine walking into a busy fast-food restaurant. You queue for five minutes before reaching the counter. Only once you get to the front does the timer start; it records that your burger is *prepared* in less than a minute.

From the system's perspective, you were served in just 45 seconds. But from your perspective, you waited nearly six minutes.

This is exactly what happens when latency is measured using software timestamps alone. The clock only starts when the software begins processing a message - ignoring the time already spent waiting in buffers, switches, NICs or the operating system.

### 3. Where Latency Can Hide

---

In modern trading systems, latency doesn't just live in the application. It can accumulate across multiple layers of the stack, especially under load:

- **Network infrastructure:** Switch buffering, cable latency, optical conversions
- **NIC level:** Receive queues, interrupt coalescing
- **Kernel space:** TCP buffers, context switching, scheduling delays
- **Application layer:** Internal queues, deserialization, thread contention
- **Outbound path:** Message formatting, NIC transmission

If we only measure latency *inside the application*, we miss all of this. Just like measuring your burger time only from when you start ordering, we're missing the full customer experience.

### 4. The Case for Measuring at the Network Edge

---

To truly understand system performance, we must measure **end-to-end latency**: from the moment a quote hits the network, to when the response is delivered to the client.

At Reactive Markets, we do this using hardware timestamping at the network boundary with tools like Corvil, which observe packets at the switch interface. This captures:

- Latency *before* software processing begins
- The software's execution time
- Any outbound queuing or transmission delays

It provides a true picture of real-world latency, under real-world conditions.

## 5. From Real-World Latency to Hypothetical Benchmarks

---

One common challenge in testing low-latency systems is a flaw known as **coordinated omission**. This happens when benchmarking tools only send the next message once the previous one has been fully processed. If a delay occurs, only that message experiences the latency, while in a real-world environment, a queue of messages would be backing up behind it.

The result? Benchmarks that look great on paper, but fail to reflect how the system behaves during bursts of activity or market volatility. This can dramatically distort high-percentile latency figures, which are crucial for understanding performance under stress.

To get an accurate view, benchmarks must simulate continuous message flow and be measured using the same tools you rely on in production. Software-based tools help, but hardware timestamping at the network edge - like Corvil - provides the full, production-grade picture.

## 6. Better Questions Lead to Better Systems

---

When considering a new technology provider, instead of asking:

**“What’s your latency?”**

The better questions are:

- **“How do you measure it?”**
- **“Do you include the time before the software sees the message?”**
- **“Do your numbers reflect performance during market stress?”**

## 7. Closing Thought

---

**At Reactive Markets, we believe that transparency in measurement is just as important as speed itself.** By making latency observable across the entire stack - and not just within the software - we help our clients gain a realistic, actionable understanding of performance.