

# The Markley Network Fabric, Boston Internet Exchange, and You: **A Practical Example**



by R. Leigh Hennig  
*Chief Network Architect*

Explaining the importance of the BOSIX on the Fabric  
and what it could mean to your organization

**The Markley Group**

One Summer Street  
Boston, MA 02110

P: (617) 451-6464

[www.markleygroup.com](http://www.markleygroup.com)



# INTRODUCTION

Suppose your commute from home to your office was a lousy one, with no direct route, and the roads heavily congested with traffic. Maybe you're 'lucky' in that like most Americans, your commute is just under 30 minutes. Or maybe you're not as lucky, and the circuitous path you're required to take gets you there in 45 minutes, or an hour. And if the main road you take is closed for any reason? Good luck.

When you access websites like [office.com](#) or use other public services from Microsoft (SharePoint, OneDrive, etc.), that's exactly what you're doing: taking a public path that may be subject to things like congestion, high latency, packet loss, and more. Your traffic is at the mercy of every interconnection between you and your destination—network engineers call this "Internet weather."

But what if you could build a portal from your home to your office, where all you had to do to get there was simply step out of your door and walk to the end of your driveway? Instead of being subjected to problems on the road (bad drivers, traffic jams, construction delays, icy roads in the winter, expensive tolls, etc.), imagine that your office was literally the next door over. That would make your commute pretty painless, wouldn't it? That's what it's like to use the Boston Internet Exchange (BOSIX).

I've written before on [the benefits of the BOSIX](#) and what you can do with it, but to summarize: the BOSIX provides a way for participants to peer directly with one another, enabling traffic to bypass the Internet and land directly on the destination network.

Sounds great, right? But what does that actually *mean*? Can we take a look at actual data to quantify this?

**Yes.**

## The Markley Group

One Summer Street  
Boston, MA 02110

 @MarkleyBoston

 [www.linkedin.com/company/markley-group/](http://www.linkedin.com/company/markley-group/)

## Before we move on, it's important to understand a few technical terms:

### Latency

If you've ever heard anyone refer to 'lag' when talking about video games, this is what they're talking about. Latency (sometimes seen as 'RTT,' or 'Round Trip Time'), usually expressed in milliseconds (ms), is the length of time it takes packets to reach a destination. The lower, the better. If latency is too high—if it's taking too long to get to your destination—your application will feel slow, glitchy, non-responsive, or it may not work at all. In our commuting example, it's how long it takes you to get from home to your office and back.

### Jitter

This is the difference in delay between packets. A lot of jitter means there's significant variance in latency from one packet to the next. Many applications need a consistent, smooth experience, with each bit of traffic taking as long to deliver as any other. If it takes you between 18 and 19 minutes to drive to the office each day, we would say you don't have a lot of jitter, and you can reliably plan your commute. However, if on one day it takes you 18 minutes, and the next it takes you 40 minutes, and after that it takes 30 minutes, then you'd be experiencing a lot of jitter.

### Packet loss

Quite simply, packet loss is traffic that never reaches its destination. There can be many causes for this, but ultimately what is being experienced is a total failure to deliver information. If you set out from your house one morning but fail to make it to work completely and have to go back home (maybe your car died, or the only bridge you can use is out), packet loss would be the networking equivalent.

### Hop count

Getting your traffic to a destination is kind of like a game of telephone: your computer tells its neighbor (your router) to pass along a message, then that neighbor tells its neighbor to pass along the message (your service provider), so on and so on until it reaches the final destination. The number of networks your data must pass through before its destination is the hop count. And, like a game of telephone, the more times your message has to be passed up the line, the greater the chances for problems. Packets can get lost, corrupted, or delayed at any hop. Thus, the more hops there are along the path, the more chances for bad things to happen to your traffic.



### To summarize:

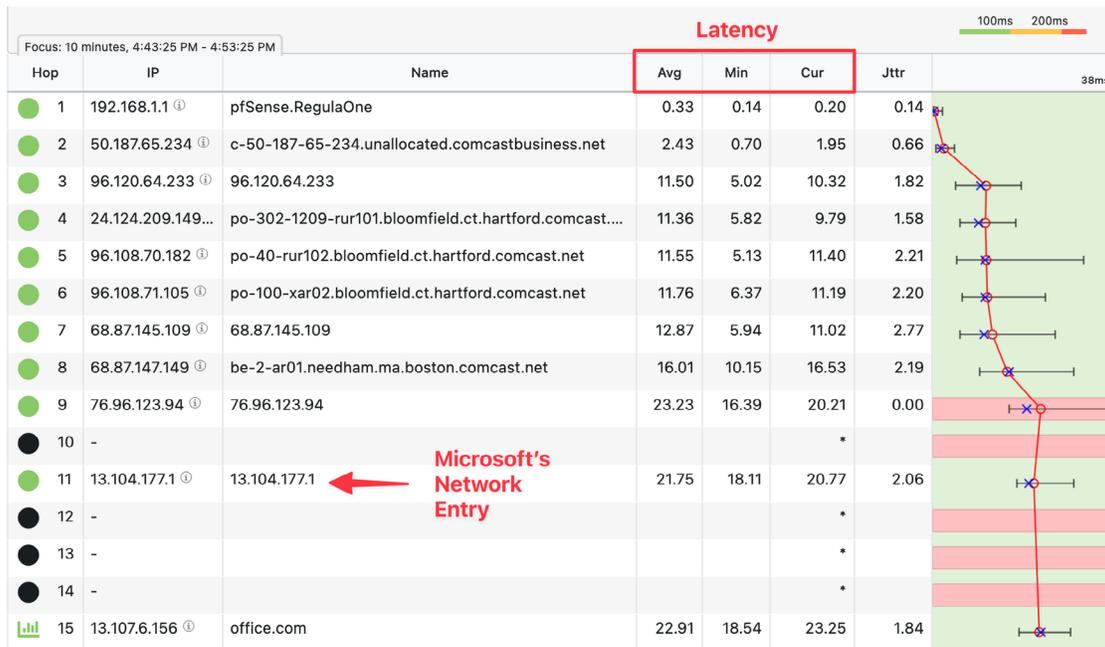
Latency is how long it takes your traffic to get to its destination. Jitter is the latency difference in subsequent trips. Packet loss is a failure to get there at all, and hop count is how many stops there are along the way. You want all of these values to be as low as possible.

# A PRACTICAL EXAMINATION

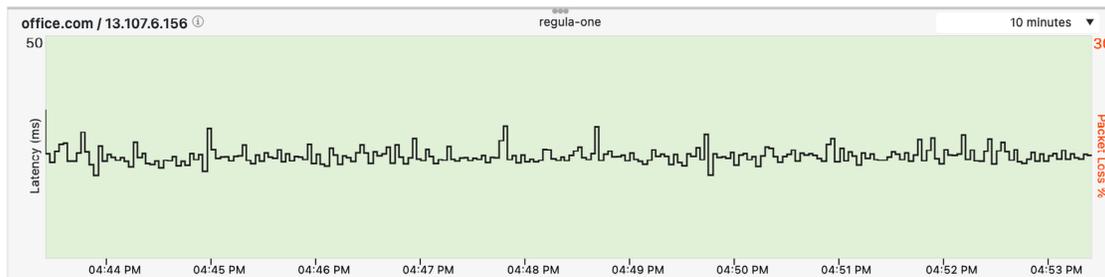
Now that we're speaking the same language, let's look at a real example of what this can look like in practice, which will help understand why Markley's ownership of the BOSIX and its presence on the Markley Network Fabric is such a big deal.

## Case #1

A business in Hartford, Connecticut uses Comcast Business as their service provider. This is a business-class service from a Tier 1 provider offering all the usual features you might expect: Service Level Agreements (SLAs), static IP assignment, an account manager, priority technical support, unlimited bandwidth with no rate limiting, etc. Take a look at what the path looks like:



Latency and jitter, graphed



Comcast Business → Microsoft O365 (office.com)

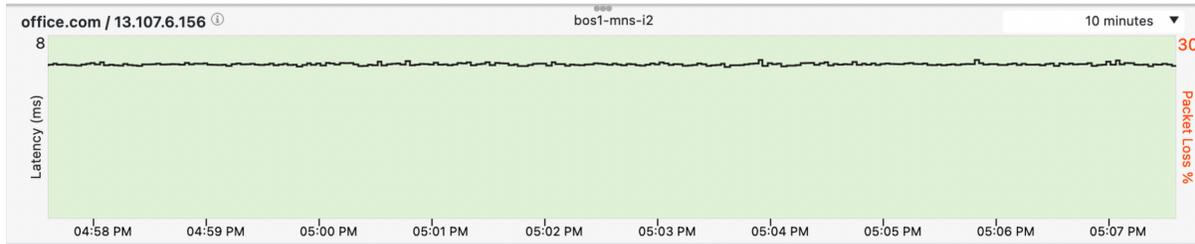
A grand total of 15 hops start to finish, with an average latency of 22.91ms, and a maximum jitter of 2.77ms. This is a good example of a clean network using 1 Gigabit wired connections from server all the way to the service provider. Not bad! But can we do better?

Case #2

A business connected to the Markley Network Fabric uses the BOSIX to reach the exact same target (office.com) at the exact same time of day. Here's what *their* telemetry looks like:

Focus: 10 minutes, 4:57:35 PM - 5:07:35 PM								Latency	
Hop	IP	Name	Avg	Min	Cur	Jttr			
1	185.172.188.29	185.172.188.29	0.96	0.28	0.45	1.05		26ms	
2	206.108.236.110	206.108.236.110	6.96	5.58	6.25	1.93			
3	-				*				
4	13.104.177.4	13.104.177.4	6.92	6.81	6.92	0.06			
5	-				*				
6	-				*				
7	-				*				
8	13.107.6.156	office.com	6.73	6.61	6.72	0.05			

Latency and jitter, graphed



Markley Network Fabric + BOSIX → (Login | Microsoft 365)

There are some really interesting things to notice.

First, **hop count**: rather than getting to Microsoft's network in 11 hops, with the Markley Network Fabric, we can get there in only 2. That's a lot of middlemen we've completely eliminated from the picture!

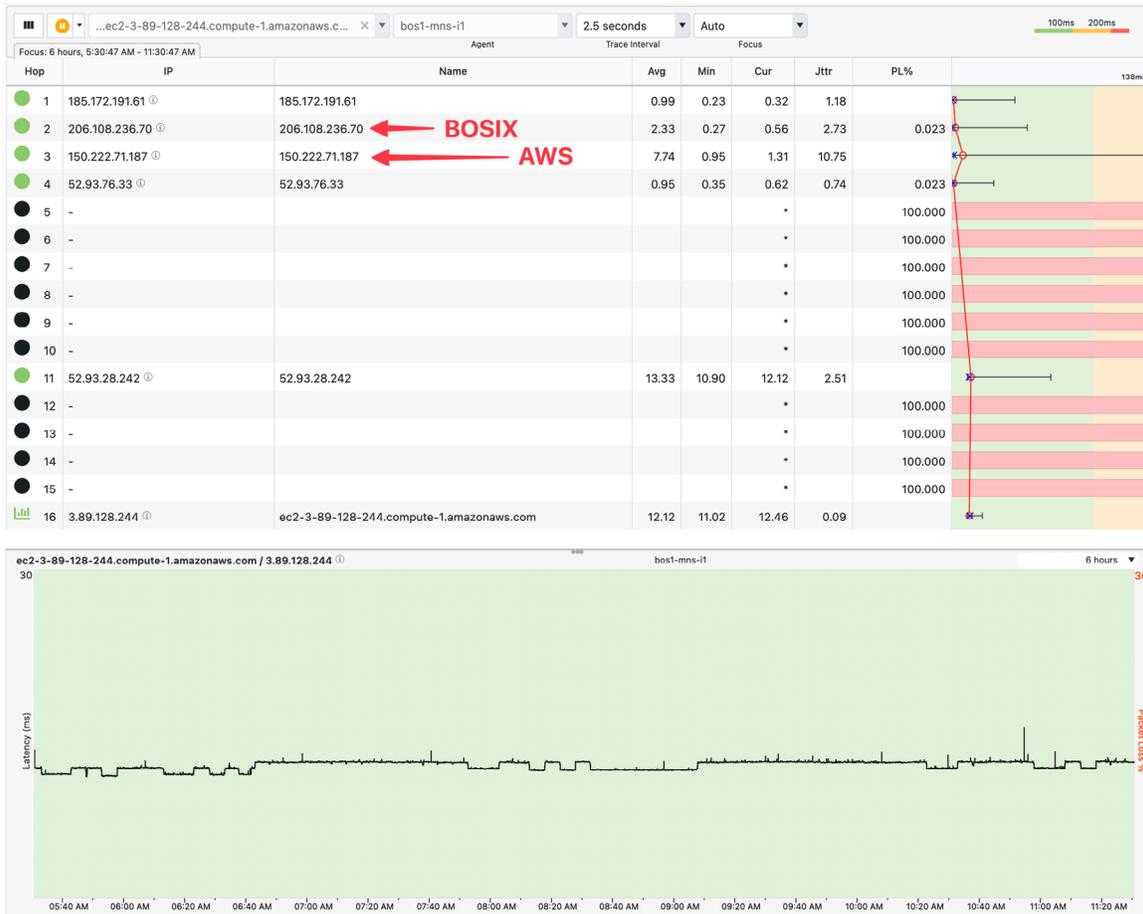
Second, **latency**: We're looking at an average time of 6.73ms. Compared to the first example of 22.91ms, a difference of 16.18ms—that makes traffic going to [office.com](https://www.office.com) through the Fabric nearly 70% faster!

There's one more thing about this second example that stands out: **jitter**. You may have noticed that the latency graph in this example is much flatter than the first example. While in our first example the jitter is quite low, here the jitter is almost nonexistent.



### Case #4

The same 6-hour window captures traffic destined for the same instance in EC2, but this time utilizes the Markley Network Fabric to get there:



As shown, the difference is immediate. The latency and path through the business-class Comcast circuit is pretty good...but it's still *more than twice* as slow as getting to our AWS target in EC2 through the BOSIX.

### Path Summary

In each of the previous examples, we did a deep dive on the specific path to [office.com](http://office.com) and an instance hosted in AWS.

Let's now take a look at a summary of the latency, jitter, and hop count to targets [amazon.com](http://amazon.com), [office.com](http://office.com), and the public IP of an instance in AWS. As with each case before, we are using instances with 1G wired connections end-to-end, with the network and source/destination instances unloaded over a 6-hour period covering morning business hours on a Thursday in October.

## AWS and Microsoft - via Tier 1 Business Class Carrier



## AWS and Microsoft - via BOSIX



The difference in latency (and jitter) is shown in the following table:

	<b>office.com</b>	<b>amazon.com</b>	<b>AWS EC2 instance</b>
via BOSIX	1.47ms (0.07)	11.97ms (0.09)	12.13ms (0.09)
via Tier 1 ISP	19.86ms (4.73)	31.04ms (4.61)	30.18ms (4.43)

While our Tier 1 connection is decent, when you can bypass the public Internet and get connectivity to the destination network in just two hops, there's simply no comparison.

### Want to drill into the metrics for yourself?

We've enabled live links to each summarized dashboard above, which you can find below:

[Path Summary - Tier 1 ISP \(non-BOSIX\)](#)

[Path Summary - BOSIX](#)

## Closing Thoughts

The difference between traffic taking advantage of the Boston Internet Exchange on the Markley Network Fabric vs. traffic taking a standard business-class Internet circuit is striking—but we're not done yet. There's one more advantage that may not be as immediately obvious: any traffic that is able to take advantage of the BOSIX is traffic that is not taking up space across your upstream WAN link to the Internet. If, for example, you have a 1Gb Internet connection, and 30% of that is traffic that can be diverted to the BOSIX, that's freeing up 300Mbps of throughput to your upstream provider. In our example we used Microsoft 365 and AWS, but Microsoft isn't the only participant available on the Boston Internet Exchange. As of this writing, there are over 80 organizations with a presence on the BOSIX—CDNs such as Amazon, Fastly, Cloudflare, and Akamai, carriers such as Verizon and T-Mobile, and content providers like Netflix, Facebook, and Apple. For a complete list, check out our [list of participants](#).