TOWARDEX
Your network success story.

2020

State of IP Network Interconnections in Boston

James Jun
TWDX IP
April 11th, 2020

In Boston, #interconnection is now improving!

TOWARDEX
Your network success story.

- **Boston used to be a pretty decent location for peering back in the dot-com era** (BBNplanet, Exodus, UUnet, Sprint, Verio, Qwest, etc were mostly interconnecting in one form or the other locally in Boston). We also had the BostonMXP for internet exchange.

- **After the dot-com bust, Boston became arguably one of the worst places for IP transit and internet routing, especially for a major region (New England).** NYC was practically our default gateway for everything, even for lot of street to street traffic.

- **However, in recent years, we are now starting to see a revival and serious development of IP transit connectivity in Boston**, where private interconnection (PNI) density is suddenly climbing like we've never seen before. The region is also now served by three internet exchanges that replaced BostonMXP (MASS IX, Boston IX and NNENIX).

- **There is still lot of work that needs to be done in Boston, and you can still find MANY examples of hairpin routing through NYC in Boston.** But, we are starting to see significant improvements, to the point where Boston is beginning to look like a Tier-1 city in some respects.

- **What do I mean by "beginning to look like a Tier-1 city?"**
  When enough number of large networks (Tier-1 ASNs and national sized ASNs) are peering with each other in a region, they will start improving traffic to/from nearby smaller networks that are not peering, thereby creating a "herd immunity."

  In my opinion, the day when you can simply buy IP transit from a major ASN, and get to most other important networks locally in the same city, is when that city is starting to exhibit a "Tier-1 market" effect. Because, in NYC, Chicago and many other major "Tier-1" cities, you can buy transit from mostly anyone and (for the most part) not worry about latency.

Tell me more!

- **An interesting side effect of COVID-19 is that perhaps IP transit is great again.** With much of the population working remotely from home, connectivity between work and residences and underlying paths used by VPNs are becoming increasingly important.

  Two important observations to draw from COVID-19 shifts in traffic patterns are:

  1. **No, the internet is not broken.** COVID-19 is only proving that Internet is scaling very well, and it's actually performing much better than we thought. Yes, there are latency issues from bad routing (and that has always been the case in remote markets), but **rarely capacity issues to deliver traffic**. This is an important distinction.
  2. **For New England, both incumbent and new providers now have significant opportunity to enhance their IP networks and create competitive advantage.** Time will tell – infrastructure investments are easiest to justify when there is scale in traffic (*"To engineer traffic, you must first have traffic" – Vijay Gill*). But I say, *the scale* is now coming to New England!

- **Previously, before COVID-19, only paths that mattered to an eyeball network were paths to content.** Content providers (Google, Netflix, Amazon, Facebook, Akamai, etc) were the only ASNs that you needed to care about and worry frequently about capacity and routing issues. Yes, we can argue about exceptions to the rule, but I'm talking *in general*.

- **Now, with COVID-19, not only traffic usage hours have shifted noticeably, it's making us think about "golden packets" scenarios (lower traffic apps that demand good latency and increasingly generate end-user complaints).** Suddenly, enterprise customers are finding about why it takes 27-54ms for their home connections to get to their VPN gateway, when it should be more like <15ms.

- **On the next slide, we've gathered round-trip delay (in milliseconds) between all of the well-known and important network providers serving Boston and tabulated them into a matrix.**

  The round-trip delay (RTD) figures were gathered by running series of traceroutes between each network provider, either using their provided looking glass, or from a bounce box that sits at a downstream customer premise. The RTD figures are based on approximate latency to the other provider's backbone router in Boston (backbone to backbone latency), and does not include end-user's last mile latency. We've also made slight adjustments (within 1-2ms range) to account for latency added by target router's control plane when pinging and tracerouting to it ("exceptions traffic" latency).

The Raw Data

Test data captured 2020/04/11	**Level 3 (CenturyLink) 3356**	**Verizon 701**	**Comcast 7922, 7015**	**Cogent 174**	**Telia 1299**	**Zayo 6461**	**NTT 2914**	**GTT 3257**	**TWDX IP 27552**	**Internap (BSN) 14742**	**FirstLight 13536**	**Crown Castle (Lightower) 46887**	**Sprint 1239**	**Hurricane Electric 6939**	**AT&T 7018**	**NOX (TR-CPS/I2PX) 10578 (cps VRF)**
Level 3 (CenturyLink) 3356	N/A Self	0 ms	0 ms	11 ms	13 ms	11 ms	0 ms	11 ms	0 ms	11 ms	0 ms	0 ms	13 ms	11 ms	12 ms	N/A No CPS route
Verizon 701	0 ms	N/A Self	0 ms	12 ms	14 ms	13 ms	13 ms	12 ms	0 ms	12 ms	15 ms	0 ms	12 ms	12 ms	13 ms	N/A No CPS route
Comcast 7922, 7015	0 ms	0 ms	N/A Self	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	15 ms	0 ms	12 ms	0 ms
Cogent 174	11 ms	12 ms	0 ms	N/A Self	0 ms	12 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	12 ms	0 ms	13 ms	N/A No CPS route
Telia 1299	13 ms	14 ms	0 ms	0 ms	N/A Self	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	9 ms	0 ms	0 ms	15 ms	N/A No CPS route
Zayo 6461	11 ms	13 ms	0 ms	12 ms	0 ms	N/A Self	0 ms	0 ms	0 ms	6 ms	7 ms	0 ms	11 ms	10 ms	11 ms	11 ms (via I2PX transit)
NTT 2914	0 ms	13 ms	0 ms	0 ms	0 ms	0 ms	N/A Self	0 ms	0 ms	0 ms	0 ms	0 ms	15 ms	9 ms	12 ms	N/A No CPS route
GTT 3257	11 ms	12 ms	0 ms	0 ms	0 ms	0 ms	0 ms	N/A Self	0 ms	0 ms	7 ms	0 ms	10 ms	12 ms	12 ms	N/A No CPS route
TWDX IP 27552	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	0 ms	N/A Self	0 ms	0 ms	0 ms	0 ms	0 ms	12 ms	0 ms
Internap (BSN) 14742	11 ms	12 ms	0 ms	0 ms	0 ms	6 ms	0 ms	0 ms	0 ms	N/A Self	N/A No Test Data	N/A No Test Data	0 ms	0 ms	12 ms	13 ms (via I2PX transit)
FirstLight 13536	0 ms	15 ms	0 ms	0 ms	0 ms	7 ms	0 ms	7 ms	0 ms	N/A No Test Data	N/A Self	N/A No Test Data	6 ms	0 ms	12 ms	0 ms
Crown Castle (Lightower) 46887	0 ms	0 ms	0 ms	0 ms	9 ms	0 ms	0 ms	0 ms	0 ms	N/A No Test Data	N/A No Test Data	N/A Self	11 ms	5 ms	13 ms	N/A No CPS route
Sprint 1239	13 ms	12 ms	15 ms	12 ms	0 ms	11 ms	15 ms	10 ms	0 ms	0 ms	6 ms	11 ms	N/A Self	10 ms	12 ms	N/A No CPS route
Hurricane Electric 6939	11 ms	12 ms	0 ms	0 ms	0 ms	10 ms	9 ms	12 ms	0 ms	0 ms	0 ms	5 ms	10 ms	N/A Self	12 ms	0 ms
AT&T 7018	12 ms	13 ms	12 ms	13 ms	15 ms	11 ms	12 ms	12 ms	12 ms	12 ms	12 ms	13 ms	12 ms	12 ms	N/A Self	N/A No CPS route
NOX (TR-CPS) 10578 (cps VRF)	N/A No CPS route	N/A No CPS route	0 ms	N/A No CPS route	N/A No CPS route	11 ms (via I2PX transit)	N/A No CPS route	N/A No CPS route	0 ms	13 ms (via I2PX transit)	0 ms	N/A No CPS route	N/A No CPS route	0 ms	N/A No CPS route	N/A Self

- **Green cell represents #KeepTrafficLocal is happening!** It means routing is symmetrically local in the Boston metro. This is good! More green cells, the better.
- **Yellow cell represents possible detection of asymmetric routing:** one side is routing in Boston, the other is returning via NYC.
- **Orange and red cells represent routing is symmetrically backhauling through NYC.** This is bad. :-(

The Raw Data - continued

IP networks ordered from best to worst:
(scored by having the most number of networks locally routed in Boston, ordered by the most # of ***green*** cells)

TWDX IP (27552)	14
Comcast (7922, 7015)	13
NTT (2914)	10
Telia (1299)	10
Cogent (174)	9
Crown Castle (46887)	8
FirstLight (13536)	8
GTT (3257)	8
Internap (14742)	8
Hurricane Electric (6939)	7
CenturyLink (3356)	6
Zayo (6461)	6
Verizon (701)	4
Sprint (1239)	3
AT&T (7018)	0

Note that we did not score NOX (10578). Except for some of the I2PX provided routes, CPS service on NOX only provides best peered (local peering) routes only to the benefit of its participants – it does not provide catch-all IP transit. Therefore, scoring NOX against commercial transit providers would erroneously generate a low score that will be both factually incorrect and out of context.

Some conclusions?

- **Look at the number of green cells!** **This did not happen before in Boston!**

  256 cells - 40 no tests = 216 scored cells. There are 118 green cells vs. 98 yellow-red cells.
  55% of detected local interconnections in Boston amongst scored networks. Could be better, but this is a #WIN!

- **If you were to select two transit providers with the best latency to every other network in Boston:**
  Based on this data, I would choose TWDX IP and Comcast. These two are the best scoring networks for keeping traffic local.

- **Credit is due where it is due:** Out of all of the major (national+ size) IP networks, Comcast (7922) has literally beat every other large IP provider in keeping traffic local. They have gone from being one of the worst peered networks in Boston to literally one of the best.

  👋 A BIG round of applause and congratulations to Comcast for their hard work into improving connectivity in recent years!

- **NTT and Telia appear to be best performing Tier-1 networks in terms of interconnection density in Boston.**
  Big THANK YOU to these two networks for their demonstrated commitment to our region!

- **Observation of "Herd Immunity" and "city is starting to look like a Tier-1 market" effects:**
  Crown Castle (Lightower) scores pretty high (coming into 6th place) even though they do not have any peering (they only buy transit).

  This is because Lightower's upstream transit providers (Cogent, GTT, CTL/Level 3 and Zayo) are now beginning to sufficiently peer with each other and other important networks in Boston. Thus, we have a situation where even a Tier-3 ISP (a network that does not peer with anyone and solely purchases transit, by Wikipedia definition) is having generally good connectivity to many other networks in Boston. **This is what I mean by "herd immunity" and "city is starting to look like a Tier-1 market" effects** – enough interconnection density by large ASNs will start covering for Tier-2 and Tier-3 networks that do not peer with each other.

Other key indicators of #interconnection improvements in Boston

- **Cost metrics in Boston are now nearly competitive as every other major NFL/"Tier-1" city**. Sure, data center power is very expensive in New England, but power cost is arguably not the most important factor that'll determine interconnection density.

- **You now have choices for interconnect locations in Boston:** In downtown alone, you have 1 Summer (Markley Group), 70 Inner Belt (CoreSite) and 300 Bent (CenturyLink), all of which are popular interconnect locations to setup your POPs in. EdgeConneX has now also joined the fray – although they are out in suburb, that particular area has access to cheap fibers.

- **Long-haul waves out of Boston are as cheap as "Tier-1" cities.** It's cheaper to buy a 100G lambda from Boston to Los Angeles than 100G from Boston to Boston. 100G long-haul lambdas out of Boston from major carriers are consistent with national market rates.

- **IP Transit is wicked cheap in Boston.** If you are finding IP transit in NYC is cheaper than Boston, you're doing it all wrong. There are **multiple** major and regional providers in Boston that will *never* lose on price, period.

- **Metro dark fiber is now cheap in Boston:** Access to cheap fiber is a **significant factor** to improving interconnection density, because cheap fiber allows more choices in colocation for your POP (thereby driving competition amongst colocation providers), and decreases cost of interconnection when the PNI needs to span between physically separate facilities in the metro. It's all about them tubez! Fiber is arguably the most important platform that forges density in interconnection.

  Access to dark fiber used to be very difficult (if not, outright hostile) in Boston, but that's no longer the case now. More on this on next slide.

- **In conclusion, from infrastructure and cost perspectives, I argue that Boston now has all of the main ingredients to become a major hub for the Internet.** All it remains are individual motives and willingness to invest their time by remaining networks that chose to neglect New England as a market. "High cost and low demand" are no longer acceptable excuses in my opinion.

So about that metro dark fiber..

“In my country, there is a problem, and that problem is transport”

This absolutely used to be the case in Boston, with few incumbent providers raising price and limiting access to dark fiber. If you are a network provider that needed to purchase dark fiber to run circuits between existing on-net carrier hotels in Boston metro, this is usually how the conversation would go:

To be fair, this is representative conversation from few incumbent providers at the time – I am not singling out any one provider.

No. That is not the case. Depends on the situation & existing revenue at risk & who the buyer is.

I am not saying we won't sell you Dark. Just saying it wouldn't even be anywhere close to $1,500/month.

No one else has & will sell Dark between these POP's that I am aware of.

When we do sell it we tend to get a premium if it is someone that sells competing services.

As far as I know we are one of the few Carriers that still sell Dark Fiber in MA.

We just take a strategic approach to selling Dark.

Happy to sell you a 10 Gig or 100 Gig or Dark but sounds like if $1,500 is your target we are not a good fit at this time.

[redacted]

That makes sense. [redacted] must have reached out.

When I'm trying to not do business with you I just can't sell you this at that rate.

Figured I needed to send you an email to prevent any confusion for my other sales channels.

I understand your position. Please try to understand ours as well.

Thanks

[redacted]

Direct: [redacted]

Cell: [redacted]

[redacted]

So about that metro dark fiber.. - continued

- **Not anymore!**

  Now, there are **several** providers constructing new facilities and selling dark fiber at competitive rates between popular data centers in the Boston metro.

- Even large carriers are now offering dark fiber at competitive rates where facilities and capacity are available (and no new construction would be required, which is typically the case anyway for on-net facilities).

- Per-pair price of dark fiber in Boston is now becoming as competitive as NYC, to the point that you can afford to burn away fiber pairs to cheaply light 100G using LR4 optics between metro facilities.

- **There is no longer "I am the only guy in town so you pay 3-5x the market rate" problem for dark fiber between carrier hotels in Boston.** We now have a healthy competition in this small, but very important sector needed in our industry to support dense interconnections.

What is up with Maine?

- My friend and a trusted advisor in our industry, James Troutman had correctly noted that improvements in Boston is absolutely not reflected at all in rest of New England, especially in Maine.

  In Maine, simply speaking, commercial Internet still sucks (but I argue, not broken), and this is very true:

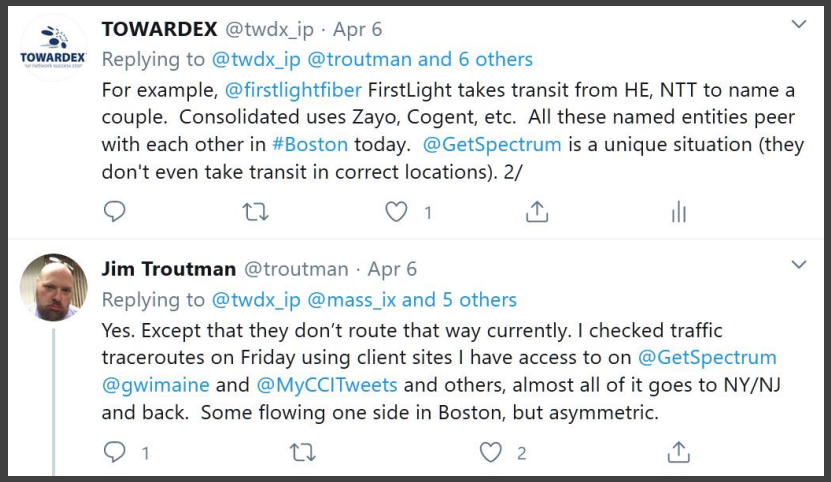

- **I'm not picking on GWI. I only use them as an example referenced from the above tweet, as I couldn't readily remember end-user IPs in Maine from Spectrum and Consolidated. But anyway, why is this the case!? Is transport between Boston – Portland that much more expensive than to NYC? May be so, but I doubt it.**

  Or perhaps it's possible that GWI was looking to diversify where they pick up IP transit from, which is a very reasonable position. But in my opinion, if you want diversity away from Boston in Portland, a better choice is to go up north, to Canada.

```
4  ae-2.r00.bstnma07.us.bb.gin.ntt.net (129.250.4.103) 1 msec  1 msec  1 msec
5  ae-6.r20.nwrknj03.us.bb.gin.ntt.net (129.250.4.114) 7 msec  7 msec  7 msec
6  ae-0.a01.nycmny17.us.bb.gin.ntt.net (129.250.3.153) 7 msec  6 msec  7 msec
7  xe-0-0-60-1-104.a01.nycmny17.us.ce.gin.ntt.net (129.250.204.14) 14 msec  19 msec  14 msec
   ^^^ GWI border router, opposite side (core-facing interface) is cr01-ptld-po25--cr01-340c-po25.gwi.net (207.5.143.33), so
       NTT transit is purchased in NYC and terminated in Portland, ME using a long-haul transport circuit.
8  cr01-340c-po25--cr01-ptld-po25.gwi.net (207.5.143.34) 14 msec  19 msec  14 msec
9  ns1.gwi.net (207.5.128.9) 14 msec  14 msec  14 msec
```

Final slide:
What can we do to improve connectivity further?

- If you are an enterprise customer, demand your network providers to provide you better service, because you deserve it. Include performance requirements for regionalized routing and to specific destinations in your RFP.

- If you are a service provider in MA and/or Northern New England, stop thinking in context of NYC. NYC is no longer the center of universe for New England, Boston is. Remember that as far as long-haul topologies are concerned, Boston attaches onto NYC and Chicago, and most providers already have excellent connectivity facing NYC. Installing a “diverse” transit provider off of NYC makes little operational sense. Also, there is no reason for an ISP in Maine to haul traffic to NY/NJ.

- **Networks are only strong as the community that supports them!**

  If you care about development of the Internet infrastructure in Boston, you need to come join us at the New England Peering Forum (NEPF) (http://www.nepeeringforum.org/). We meet every year (though likely not this year, due to COVID-19) to discuss how to improve our networks and #KeepTrafficLocal here in our region.