TOWARDEX®
Your network success story.

Network Infrastructure Capital Programs Update

November 2017

Agenda

TOWARDEX®
Your network success story.

Network Strategy 2020 Overview

September 2017

Vision 2020

Split business into 3 different “business units” (BUs) by 2020 to increase focus in key growth areas:

- OPTICAL and FIELD OPERATIONS:

  TOWARDEX Optical Communications: Joint-venture management for outside plant – fibers, conduit leases, pole attachment & licensing, facilities etc. Everything from dark fiber to fixed network assets

- TRANSPORT:

  Packetsurge™ (AS 22147, AS 30071): Network-neutral interconnection ecosystem

  - Backhaul IP, cloud, voice & video throughout our network footprint
  - Network-dense ecosystem for cloud, edge computing and 5G wireless developments

- IP:

  TWDX IP™ (AS 27552): High performance, low latency IP transit

  - Develop a well engineered, high performance national backbone dedicated to providing low latency wholesale internet connectivity. Focus on high degree of peering & application-specific interconnections with economies of scale.

PACKETSURGE®

Middle mile programmable network for handling everything lit service and transport related. If you're using a lit transport service from TWDX, it's running on Packetsurge.

Major programs:

- EVOLVED EXTENSIBLE TRANSPORT PROGRAM (EETP)
  - Build into every major data center in the market and provide aggregation into TWDX IP hubs for colocation users purchasing IP bandwidth.
  - Build into network dense assets of tomorrow – wireless towers, edge data centers, high-tech offices, etc.

- INTERCONNECTION ECOSYSTEM

  Bring value to our data center partners – extend network ecosystem to every partnering data center.
  - Data centers are invaluable partners to TOWARDEX's business model. Extend the network-rich ecosystem by providing transport backhaul of carriers, cloud providers and exchange operators to partnering data center sites.
  - Partner agreement for carriers & data centers: providing volume discount rates for wholesale network access.

- MASS IX®

  Massachusetts Internet Exchange – Open-IX® certified carrier & data center neutral peering exchange in New England.
  - Distributed across 12+ data centers in the Boston metro – VPLS transport across Packetsurge backbone.

TWDX IP®

AS27552 – the #1 choice of blended IP bandwidth in Massachusetts. TWDX IP's mission is to move traffic from source to sink (content to eyeballs) as quickly and reliably as possible, leveraging economies of scale in regional interconnections.

Major programs:

- NATIONAL BACKBONE EXPERIMENTAL (NBX)
  - Build AS27552 hubs in New York, Washington DC and Chicago markets.
    Dallas under discussion for future.
  - Dense regional interconnections with last miles and large scale access networks:
    Comcast ibone, Comcast CRAN, Verizon today; Cox, TimeWarner, AT&T under discussion
  - Install 1 Tbps of new interconnection capacity
  - Software defined billing and usage metering: unique billing plans based on traffic flows, network egress
  - DDoS mitigation and alerting service for all customers

- NEXT GENERATION CORE (NGC)
  - Accelerate 100G adoption throughout TWDX IP, kill all 10G interfaces in the core by H2 CY 2018.
  - *10G is the new 1G* – move to 100G mentality in the core; enable ubiquitous 10G access for retail sales
  - PAM4 optical transmission – future proof for 400G backbone in the metro

- PEERED TRAFFIC SERVICE (PTS) - Partial Transit
  - #make_ip_great_again: IP transit industry challenges – "transit is dead", "peering is dead", "how low can you go", "who's buying who, @telecomramblings", etc.
  - Customers are increasingly looking for IP service that is tailored to their traffic profile. Some customers may want low latency access to last mile networks, some may want access to content at lowest possible price, etc. Why should all customers pay the same $ per Mb when they all have different expectations for IP transit? Settlement-free traffic ratios and Tier-1 are irrelevant in today's market.
  - PTS is a "commercial peering routes service" initially tailored for muni-broadband providers and K-12 & higher-edu institutions: it gets you access to all the content you want for a low simple price – pay for port charge, no usage billing, no 95th percentile, no non-sense & no frills.

TOWARDEX®
Your network success story.

Network Service Delivery Model

September 2017

PACKETSURGE Delivery Model

- Packetsurge® builds into data centers, central offices and wireless towers.
- TWDX IP (AS 27552) interconnects with Packetsurge at two redundant hub sites (head-ends) at each region. Packetsurge backhauls IP bandwidth subscriber via layer-2 to the nearest IP hub.
- Packetsurge also interconnects with cloud providers (Amazon DirectConnect, Megaport, Azure, Google Cloud Platform) at common interconnect facilities.

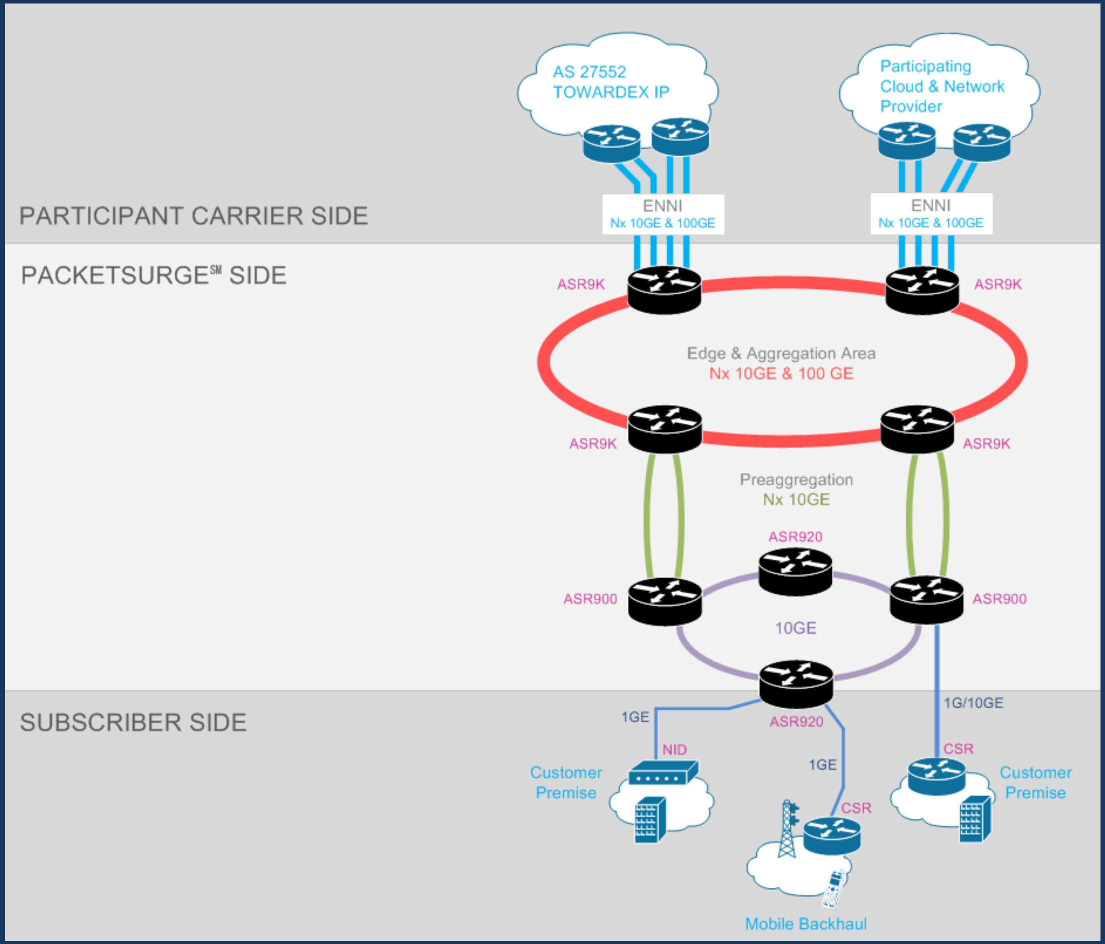

Packetsurge & TWDX IP Interconnection

- Packetsurge and TWDX IP meet in two common hub sites (called Joint Network Facility or JNF) in each region. 10G and 100G NNIs are setup between TWDX IP and Packetsurge at each hub location.
- When a customer orders TWDX IP service, Packetsurge backhauls to the nearest IP hub @ layer-2.
- Oversubscription is not permitted for EETP/Packetsurge interconnections into the TWDX IP network. No more than 10x 1G EVCs are permitted per 10G NNI.

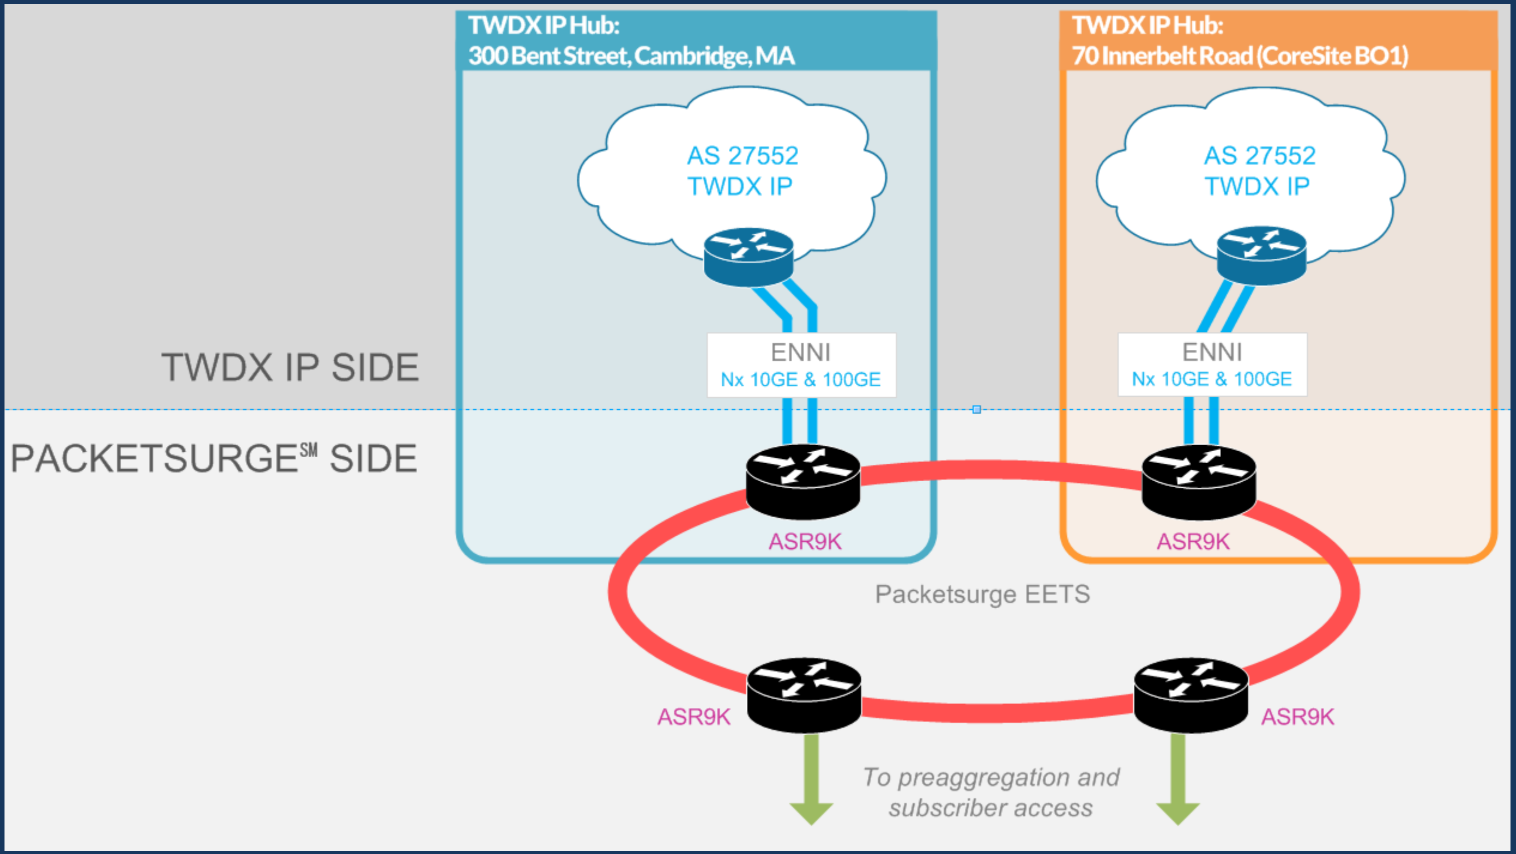

Traffic shaping and packet buffering

The topic of packet buffering is real for service providers:

- At natural points of network contention (i.e. speed step down – 100GE to 10GE or 10GE to 1GE as handing off to customer), any momentary traffic burst (microburst) could completely overwhelm customer's uplink port to the internet, even if average port utilization is very low.

- The problem can be exacerbated by ISPs selling links having same capacity as their own backbone to their customers:
  - For example, a WISP providing 20 Mbps service to a customer, and having 1GE port to their uplink ISP will have minimal to no problems – the majority of bandwidth contention will be at the wireless hand-off stepping down to 20 Mbps toward the subscriber. WISP can buffer as necessary at their wireless gear and TCP slowstart will take care of most issues.
  - By contrast, a small ISP selling Ethernet/GPON FTTH to customers at 1G while themselves only having 1GE uplink can exhibit significant contention issues. TCP will flood this ISP's own uplink at maximum possible speed, as downstream customer's link capacity is same as ISP's own backbone uplink.

- ISPs using low-buffered switches (many 1U switches in colo) cannot handle speed transition from 10G to 1G very well for internet traffic. These switches are meant to be used for low latency paths (i.e. server to server within the same data center) where TCP window sizes are to remain small. The result is users reporting poor performance for bursty data traffic over the internet, especially over longer distances.

Buffering is needed to drain packets at the point of interconnection, if you're selling IP transit.

However, buffering must be reasonable – the goal is to not eliminate packet drops in a packet switching network. Output packet drops are a healthy sign of a well functioning network – this is how TCP gets its feedback response during a network congestion.

Excessive buffering contributes to bufferbloat and makes the connection unusable. However, not providing sufficient buffering to the point of encountering high rate of packet drops would result in poor network throughput.

To address these issues, we use the following guidelines on NNIs which are interconnecting Packetsurge into TWDX IP. These are only performed on low speed (sub 10G) customer links:

- Each downstream subscriber EVC is traffic shaped to burstable capacity speed.
  Example: a 100 Mbps IP service burstable to 1G is traffic shaped to 1Gbps CIR at the NNI.
- Packetsurge side only admits traffic at specified traffic contract values – TWDX IP at upstream side has to traffic shape to the correct CIR. Excess entry is discarded by policer.
- Egress Traffic Manager (TM) runs at 2ms interval to smooth out bursts and pace packets. This helps low-buffered switches that may exist downstream to better handle bursts.
- Upper limit of 100ms worth of best-effort queueing to prevent excessive buffering. At 100ms queue fill, we start tail dropping – at this point, customer's link is clearly congested and further buffering is considered harmful.
- Customers may request Active Queue Management (AQM) as an optional configuration parameter by escalating to their SE. With this option, we begin WRED dropping at 50ms queue fill, then tail drop at 100ms fill.

TOWARDEX®
Your network success story.

Project Highlight:
Joint Network Facilities (JNF)

November 2017

Joint Network Facilities Program

We had too many single cabinet network colo sites in Boston metro.

- Eliminate unneeded colo and prepare our network for national expansion.
- Leave 1-rack and micro deployments to outlying data centers, central offices and wireless backhaul sites.
- Consolidate small cabinet deployments into two larger redundant hub locations where both TWDX IP and Packetsurge can jointly meet, and wholesale customers can bring in their own cross-connects to our hub facilities.
- TWDX IP (AS 27552) will only maintain POPs in JNF hubs. Packetsurge will backhaul all remaining on-net facilities in the region to the nearest JNF hub for customers ordering IP service.
- **Central hub locations must be redundant and diverse** – contract out space to two separate carrier hotel / data center providers in each market:
  - In the Boston market, **300 Bent Street, Cambridge, MA** (CenturyLink) – “BOS 01” and **70 Innerbelt Road, Somerville, MA** (CoreSite) – “BSN 05” have been selected as JNF hub sites.
  - Location study and site evaluation for NYC market are currently ongoing.
  - For Virginia/DC market, Equinix Ashburn, CoreSite Reston Campus and Level 3 Washington DC are being evaluated for JNF site selection. Two sites will be chosen.

JNF Site Specifications

- Each hub site is typically sized between 150 to 250 sq ft. with minimum of 4 racks emplaced:
  - 1x 4-post rack for heavy network equipment
  - 3x or more 2-post racks for remaining network deployment

- Both AC and DC power available
- Modest power density for network loads: 25 kVA per site

- Customer colocation of equipment is not permitted, as JNF sites are low density and only designed to accommodate TWDX network facilities. Customers looking for colocation are referred to data center partners.

- MDF for cross-connect meet-me use by customers:
  - Cross-connects to TWDX services do not incur a charge.
  - Pre-terminated trunk cables and inter-ducts to building MMRs.
  - Standalone cross-connects to third party carriers are provided. JNF standard cross-connect rates are $370 MRC, $350 NRC unless specified otherwise in your order form.

JNF Connectivity Requirements

- JNF hubs in the metro *must* provide fully diverse link protection for network loads. Verifiable diverse POE & fiber routes are a must.
- TOC (TWDX Optical Communications) is leasing wholesale fibers and contracting out vendors for outside plant construction & conduit leases.
  - 432-count cables from MPOE to outside points of interconnection, at commonly agreed manholes

- To facilitate bandwidth requirements for TWDX IP and Packetsurge networks, two 12-count fibers, each taking diverse route (24 strands total) are initially being provided between JNF hubs. More fibers could be spliced in later if needed.

  - TWDX IP requires 6-count unlit fiber to light 3x 100GBASE-LR4 directly into the dark fiber. (Cheaper to burn away fiber assets than buying into coherent 100G for short distances)

    - Where applicable, filters with wide-band 1310nm port is used to stack 100GE on top of 40 C-band channels.
    - Semiconductor Optical Amplifier (SOA) is used to indiscriminately amplify 1290-1330nm bands used by 100GE-LR4. Cheaper than 100GE-ER4 optics, but must watch out for chromatic dispersion - otherwise, BER will become intolerable!

  - Packetsurge requires minimum of 4-count fibers with DWDM for metro applications.
  - EDFA and Raman pump lasers are used on Packetsurge for medium & longer distance routes (30 to 120km) converging into JNF hubs.

Fiber routes between JNF hubs in Boston

Path A
Path B
Path C
Path D

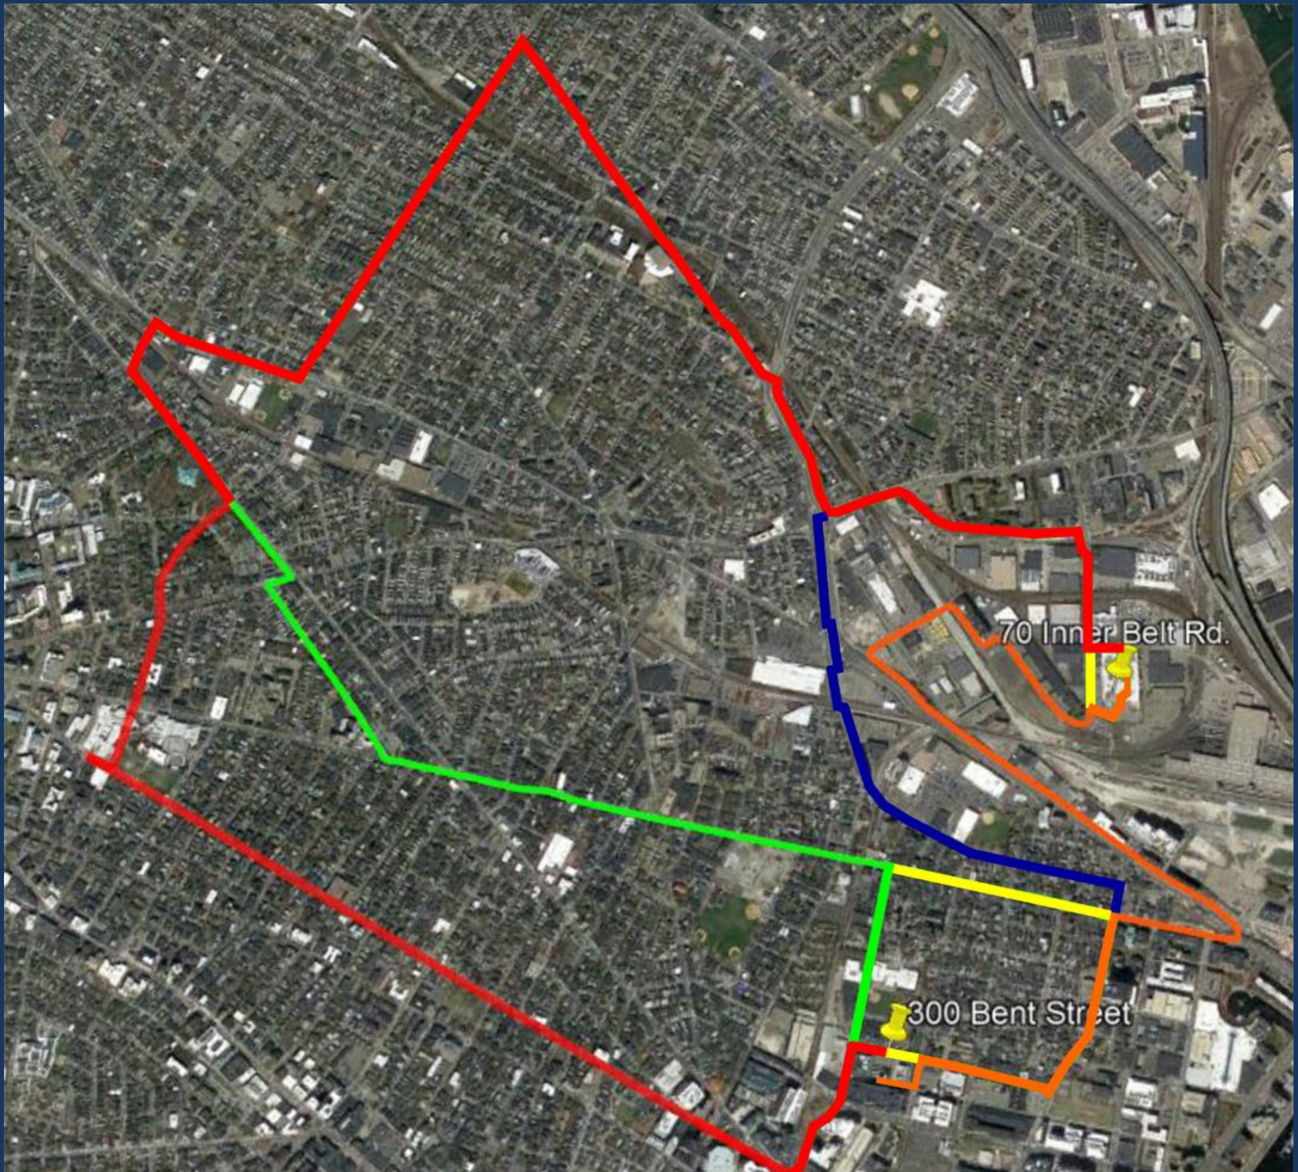

JNF Rollout – Boston, MA

- **May 2016:** Initial operating capability (IOC) achieved, reusing existing IT cage at 70 Innerbelt Road (BSN 05) in CoreSite.

- **March 2017:** CenturyLink (formerly Level 3) tapped as general contractor for design and development of JNF facilities (Level 3 Professional Services).
  - Sites are designed to meet NEBS guidelines and GR-63-CORE requirements.
  - Material suppliers selected: CommScope, Telect, Corning, Graybar

- **September 2017:** Buildout of 300 Bent Street (BOS 01) hub completed and suite turned over to TWDX. Site is now being kitted out for production launch.
  - Full operation of 300 Bent Street (BOS 01) site is expected to occur in Q1 2018.
  - Project coinciding with new TWDX IP core routers upgrade (Next Generation Core)
  - TWDX IP installing new head-end aggregation router (dcr02.bos01)

- **H2 CY 2018:** Packetsurge will migrate ar02.cambridge.ma.boston hub router to the new 300 Bent Street JNF.
  - Existing ASR 9006 located in a standalone cabinet will be decommissioned. ASR 99XX series and upcoming Lightspeed 32x100GE line card under evaluation for new replacement router.
  - MASS-IX peering LAN will also move together with Packetsurge into the new JNF.

300 Bent Street JNF - September 2017:
New suite constructed by Level 3 Field Services

300 Bent Street JNF - September 2017:
New suite constructed by Level 3 Field Services

300 Bent Street JNF – October 2017:
Terminating fibers on OSX frame

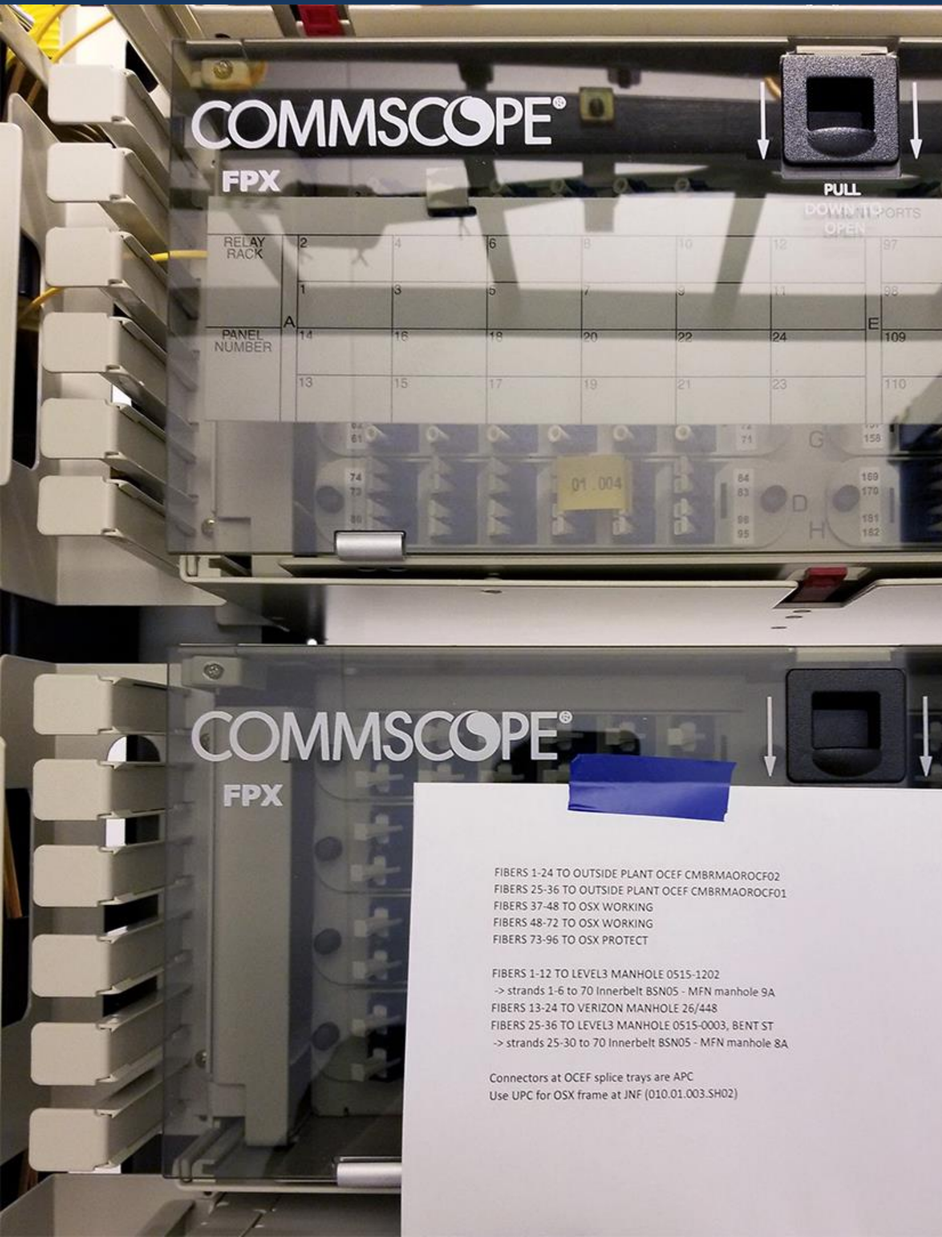

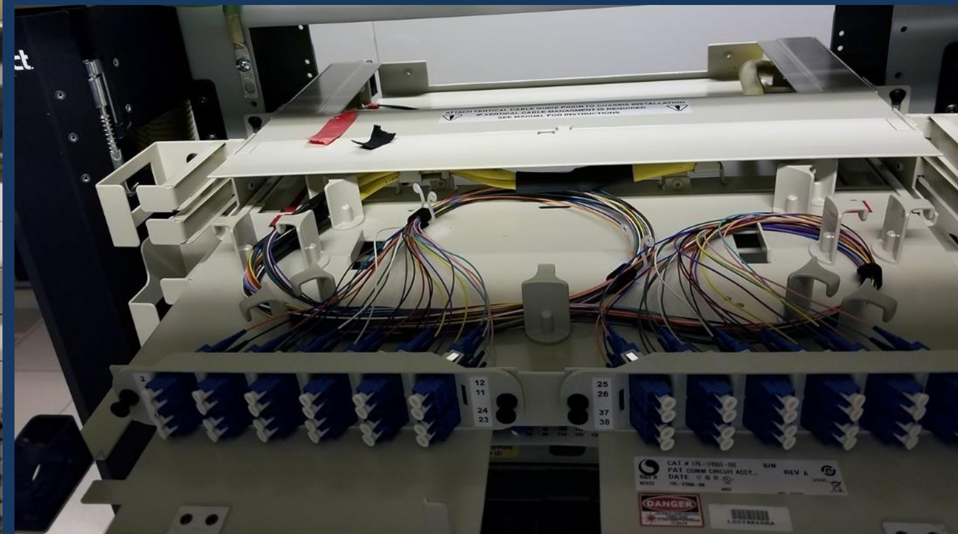

300 Bent Street JNF – October 2017:

OTDR testing fibers from JNF to a splice tray at outside manhole

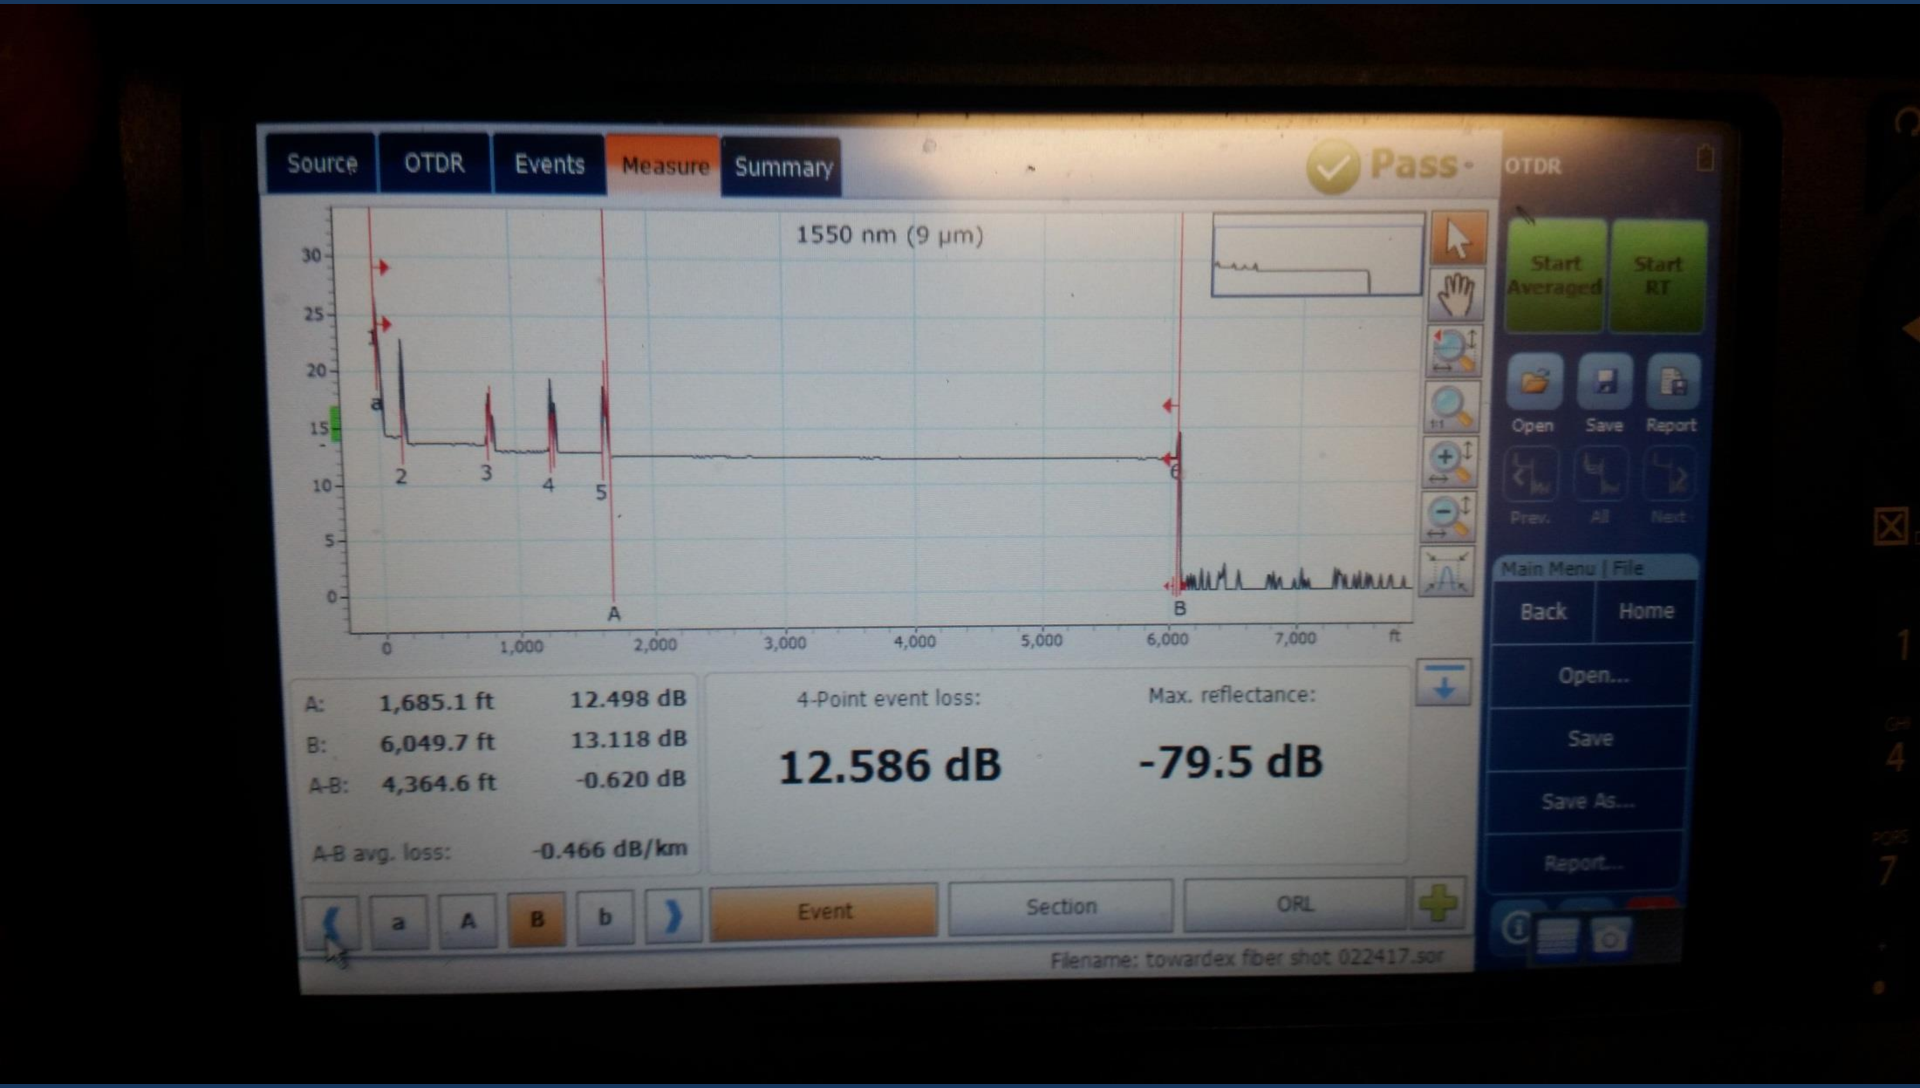

dcr02.bos01

- This is a new head-end aggregation router for TWDX IP going into 300 Bent St. to supplement the existing dcr01.bsn05 in 70 Innerbelt.
- New TWDX IP design requires two, independently diverse customer aggregation routers in each metro.
  - Each aggregation router must be sited in a different JNF facility:
    Customer requirements for diverse circuits (shadow ports / redundant dual-homed internet).
  - **dcr01.bsn05** is a Cisco ASR 9010 in 70 Innerbelt Road (CoreSite).
    Already in service and operational since May 2016.
  - **dcr02.bos01** is the new 2nd aggregation router going into 300 Bent St JNF.
- For **dcr02.bos01**, we've chosen the latest top-of-the-line ASR 9906 router from Cisco.
  - ASR 9906 (Torchwood) First Customer Ship (FCS): August 31, 2017
  - TWDX IP issued PO for pre-order on the following day.
  - We're one of the first 10 customers taking delivery of ASR 9906!

  - Initial deployment with MOD200 and Powerglide 24x10G-TR
  - Migrating to Lightspeed A99-32X100GE-TR in 2018.

  - 64-bit IOS XR on day 1
  - Production commissioning planned for Q1 2018.

300 Bent Street JNF – November 2017:
Installing dcr02.bos01 (ASR 9906):

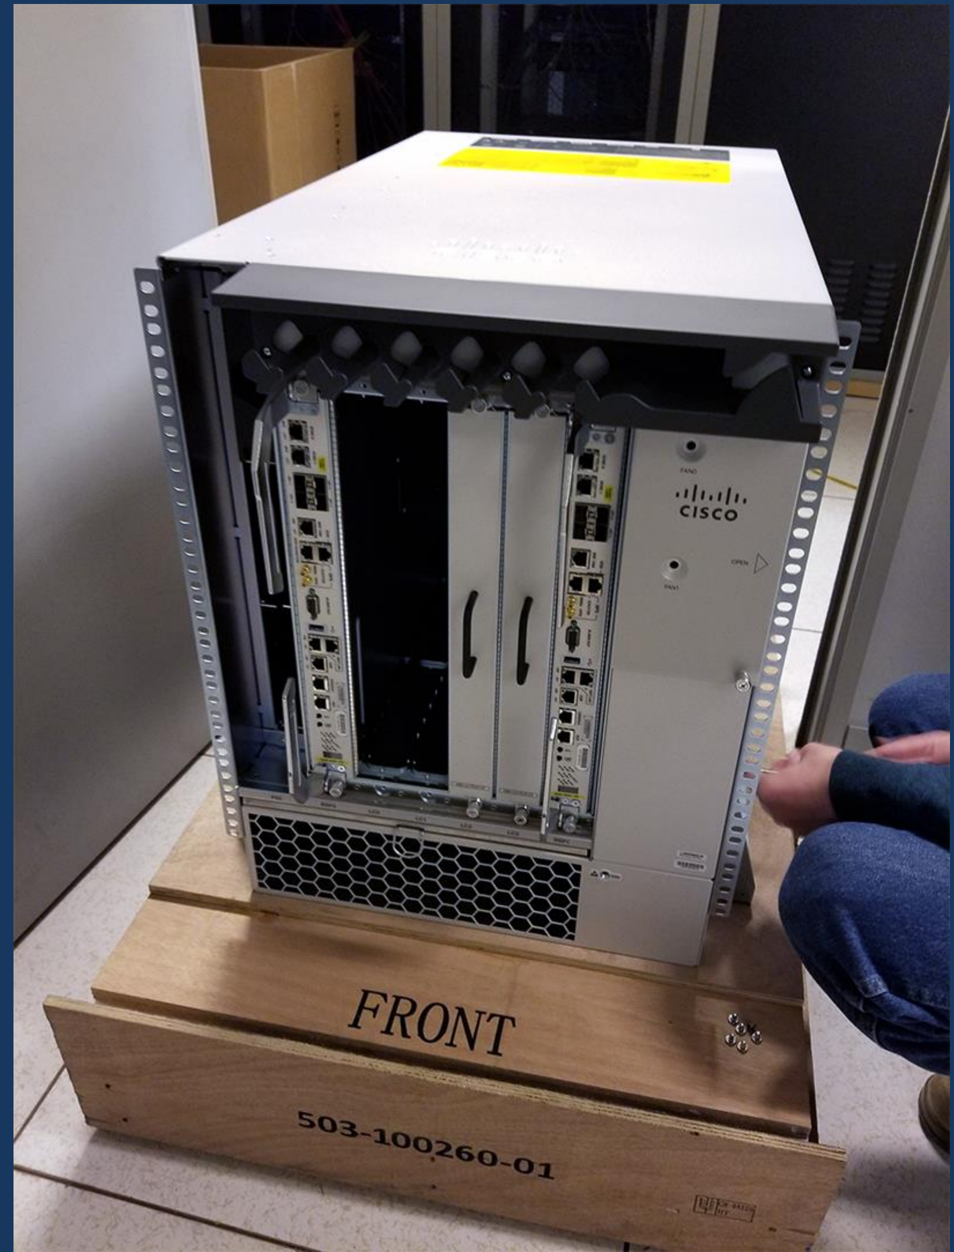

300 Bent Street JNF – November 2017:
Installing dcr02.bos01 (ASR 9906):

TOWARDEX®

Your network success story.

Project Highlight:
Next Generation Core (NGC)

November 2017

TWDX IP (AS 27552) Next Generation Core

Historically inspired by the former Exodus Communications backbone design (1998-2001), AS27552 operates a hierarchical architecture consisting of two-phase routing layers: distribution core (DCR) and backbone core (BBR). In busy markets, a dedicated peering router is also provided to separate out transit & peer interfaces - internet border router (IBR).

Updated to modern standards, the backbone uses MPLS/IP for packet forwarding.

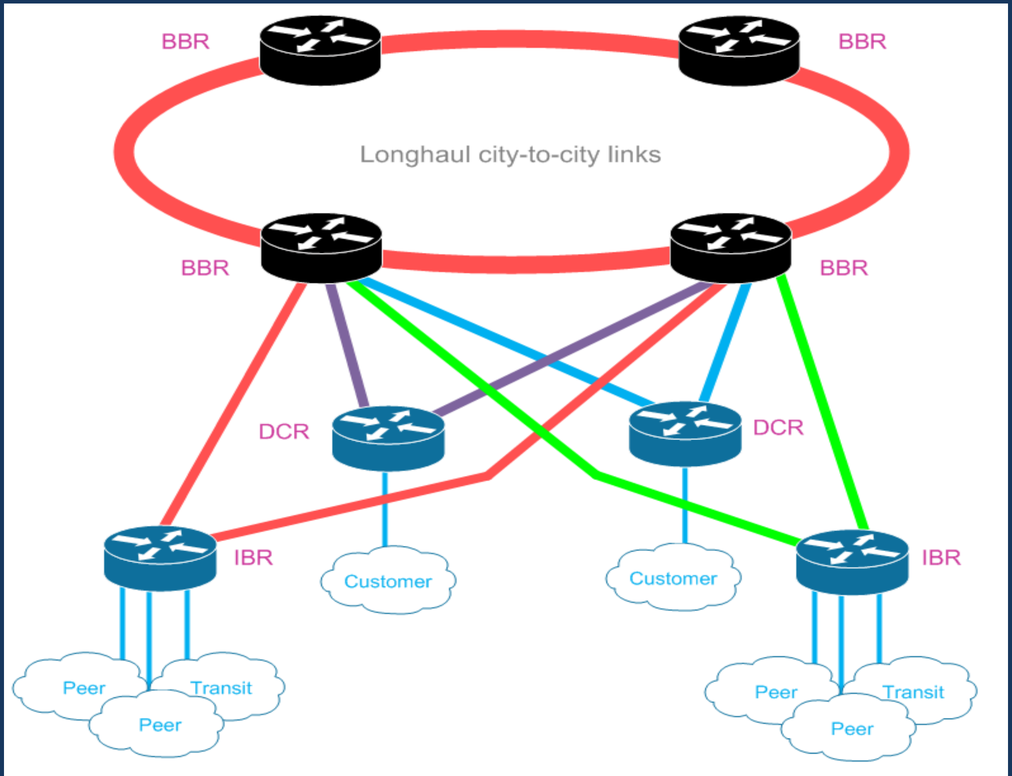

Next Generation Core

To lower cost and accelerate backbone upgrade initiatives, we operate a pair of high speed core routers called **BBRs** (backbone routers) in each region. The BBRs are tied together with longhaul 100G waves.

Using MPLS, the BBRs are pure label switching routers - AS27552 operates a “BGP-free core”. In other words, BBRs do not have any knowledge of the global BGP routing table, rather, they only maintain internal IGP-TE routes to operate the MPLS forwarding fabric.

Because BBRs have no knowledge of the internet routing table, DCRs and IBRs act as Provider Edge (PE) routers, imposing and disposing MPLS labels for IP traffic. This means that routing decision is made upon network ingress - the first PE (dcr or ibr) receiving the IP packet will make the routing decision on where to egress the network, then impose label that specifies which PE will be used to leave the network. BBRs will swap and forward labels blindly to the destination PE with zero understanding of the BGP/internet routing table.

- In each market, a pair of Cisco ASR 1001-X routers are provided as out-of-band Route Reflector (RR) appliances. Each ASR 1001-X functioning as OOB RR does **not** partake in any packet forwarding operation - these devices are simply dedicated servers announcing iBGP routes as BBRs don’t carry them. Yes, we could have used x86 servers for this.
- All OOB RRs are peered with each other, fully meshed iBGP-wise.

NGC Key Requirements

The use of MPLS allows us to significantly lower port cost in the core network fabric, accelerating deployment of higher speed interfaces (e.g 100G) and easing backbone upgrade initiatives. This creates relaxed hardware requirements:

- No need to carry full routing table.
- Higher density 100G and 10G ports at lower cost
- Only need to forward on MPLS labels imposed on packet

Hardware design wise, even a low cost layer-2 switch is capable of meeting such requirements, vastly lowering the cost of core network transport. However, with MPLS and service provider requirements, it gets more complex.

Our key requirements for NGC program are as follows:

- Economical port cost – cost per port must be significantly lower than run-to-completion based full featured routers.
- Must provide minimum of 50ms packet buffering to sustain speed transition.
- Must be able to perform ECMP and hash over LAG based on full L4 flow information.
- MPLS behavior testing – ICMP exceptions handling, pen-ultimate hop popping, MPLS over LAG, MPLS with ECMP, 2-3 labels, bottom of stack, etc – stable?
- Must support RSVP-TE, Fast Reroute, Facility Backup and minimum of 16,000 labels in LFIB.
- Minimum forwarding capacity of at least 1 billion packets per second.
- Roadmap support for Segment Routing (SR-TE)
- Must be able to respond on IPv6 traceroute when acting as P router under 6PE application.

NGC Vendor Evaluation

During the NGC evaluation process in 2016, we've shortlisted vendor offerings to the following, which once selected, will assume the role of new BBRs:

- Cisco:
  - NCS-5502
  - NCS-5501-SE

- Juniper:
  - QFX10002-36Q
  - PTX1000

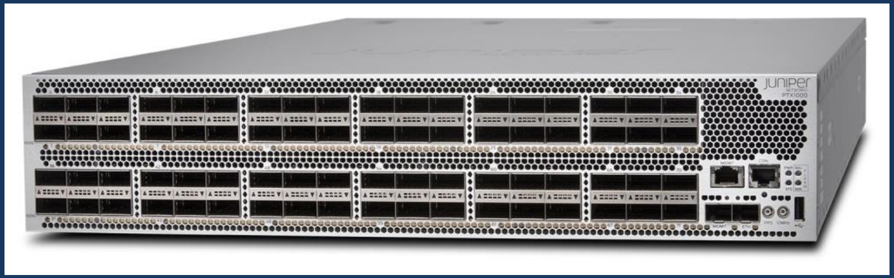

NGC Program Rollout

- **November 2016:** RFI responses received for evaluating NGC requirements. Technical evaluation started on product offerings from both Cisco and Juniper.

- **June 2017:** RFQ issued to shortlisted vendors for NGC proposals, including equipment, spare parts, licenses and support contracts.

- **July 2017:** Juniper Networks awarded program for NGC supply & delivery together with national backbone expansion (NBX). Deliveries and maintenance contracts to last through 2020.

  - Lot #1 order was two QFX10002-36Q with Advanced Feature License for Boston, delivered in August 2017. Production livening of new bbr02.bos01 will occur between December 2017 to March 2018.

  - Ongoing partnership & collaboration with Juniper on program improvement:
    Initial evaluation was completed by TWDX for QFX10002 in 2016. In October 2017, project modification review was started to evaluate MX10003 or PTX1000 for NGC, as 'lean internet core'.

- **November 2017:** Juniper MX10003 (Summit-3 with EA Trio) now under discussion for Lot #2 order. Order expected soon to cover BBRs in NYC metro.

Questions?

Contact us at ip-admin@twdx.net