TOWARDEX
Your network success story.

2021 Customer Updates

State of TWDX Network Improvements

February 10, 2021

Table of Contents

TOWARDEX
Your network success story.

70 Inner Belt Road Gateway Update

Joint Network Facilities Program

January 18, 2021

70 Inner Belt Road Status Update


Pictured: JNF BSN 04

We're pleased to announce that our new gateway in 70 Inner Belt Road, known as "Joint Network Facilities (JNF) BSN 04", is now in production and passing customer traffic! Remaining outstanding items on the project are:

- **Continued tenant fit outs**

  TWDX IP has taken occupancy of the new JNF BSN 04 space. We have now completed migrating all customers to the new Cisco ASR 9922 (dcr03.bsn04) aggregation router. Occupancy for Packetsurge metro-E network is expected to commence in late 2021.

  Work on optical transport infrastructure at JNF BSN 04 continues – Installation of DC power plant will soon commence to support deployment of new ROADM shelves.

- **Recommissioning of the old DC5 colocation space (known as "BSN 05")**

  Instead of vacating the old DC5 computer room we were in previously, we will be recommissioning this space as redundant / diverse POP for data center tenants in CoreSite BO1.

  Contractor is completing construction of the new Inside Plant infrastructure (internal conduits and pull boxes) to support this initiative. Redundancy features between JNF BSN 04 and BSN 05 sites are described in the next few slides.

- **JNF South Fiber Entrance Construction**

  We've now completed construction of our new fiber entrance at the South Gate Driveway area of 70 Inner Belt Road. Installation of innerducts and fiber optic cable plant will soon commence.

  The new entrance provides a 4 – 4" duct bank connecting the CoreSite South Fiber Room onto TWDX Hub Express (HEX) Conduit System out in the street.

TABLE OF CONTENTS

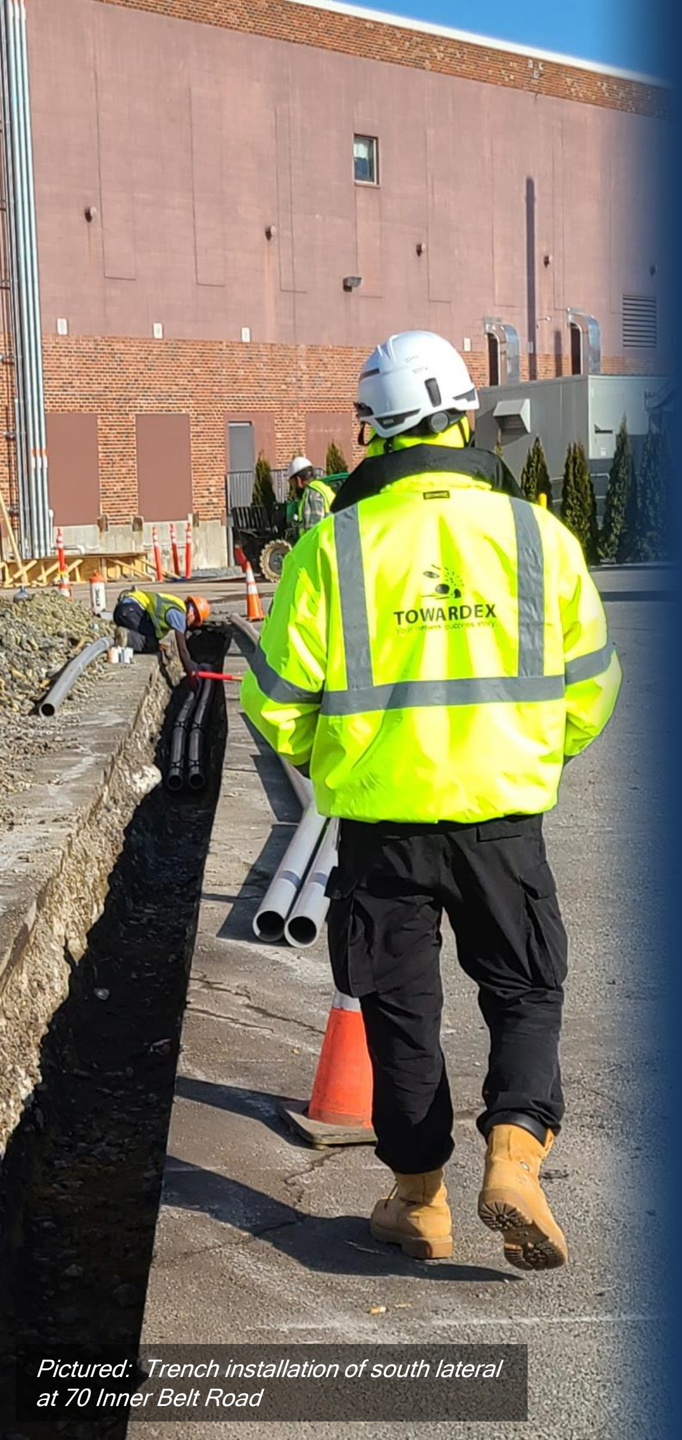

Pictured: Trench installation of south lateral at 70 Inner Belt Road

Integration with Hub Express Program

New fiber construction at 70 Inner Belt Road is being conducted in conjunction with the Hub Express (HEX) Program. Hub Express is a TWDX-led capital initiative to develop targeted fiber infrastructure in the metro, with focus on improving interconnection between participating carriers and data centers.

The key aims of the Hub Express program are:

- **Increased interconnection density through ubiquitous access to dark fiber**

  Boston has been an underdeveloped metro for density in data center connectivity and interconnections, making it expensive for new ISPs and networks to operate. The Hub Express program provides a blueprint for massively increasing dark fiber capacity between data centers and telecommunications infrastructure, helping to lower the cost of backhaul in the middle mile.

- **Improved choice and diversity for interconnection facilities**

  Through trench partner engagements and conduit licensing programs, TWDX fosters interconnections between member utilities of the Hub Express System. TWDX is also constructing laterals from HEX conduit system to neutral network elements, such as vertical assets for 5G and small cell attachments.

  70 Inner Belt Road is at the center of the HEX system route, being connected with 1,728 strands of new fiber and provisioned with additional conduits to accommodate other telecom carriers.

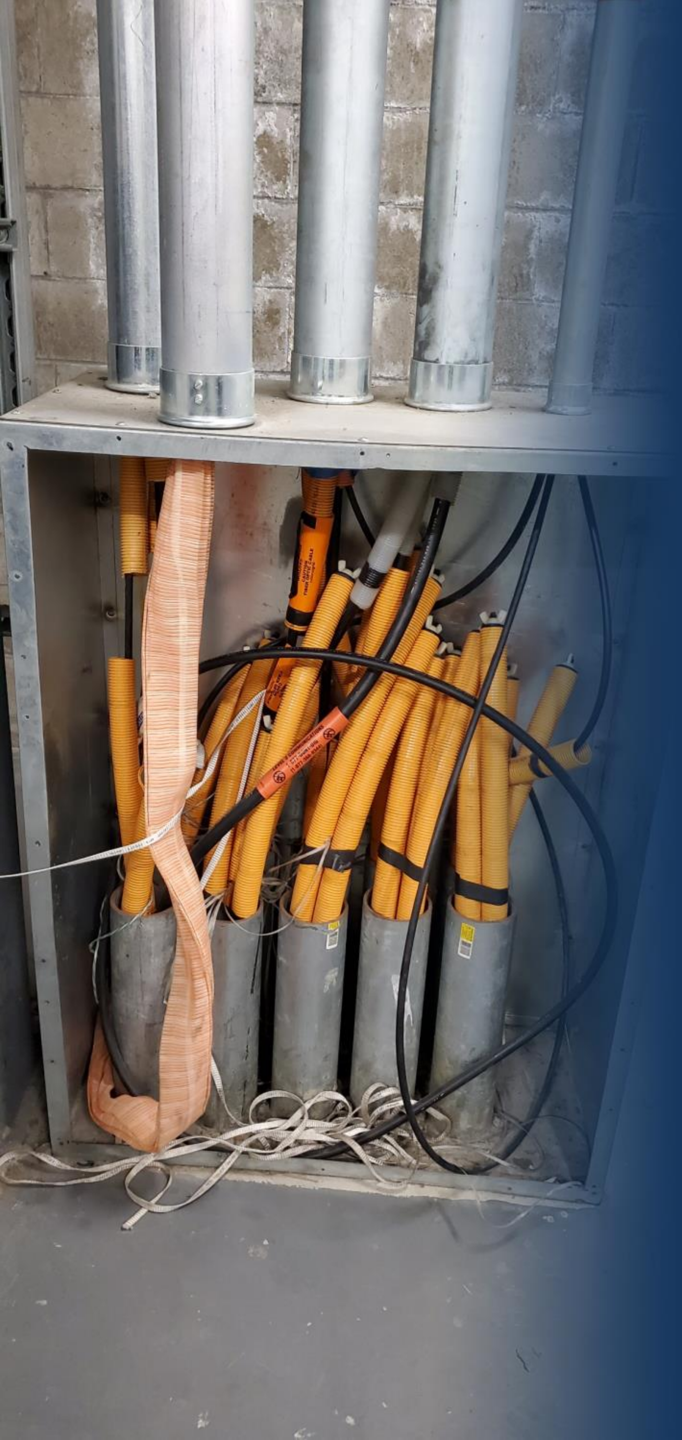

New TWDX Fibers at 70 Inner Belt Road

In close partnership with CoreSite, we're making the most significant investment into fiber infrastructure at 70 Inner Belt Road, perhaps arguably more so than any other carrier since 2003. This work is being conducted under following parts:

- **Cross Connect Diversity via two Meet-Me Rooms (MMRs)**

  CoreSite's primary MMR is located in the second floor (original CoSpace MDF). In response to increasing demand from enterprise customers, CoreSite has recently constructed a brand new MMR in the first floor of the building. We're integrating our two POPs in the building to connect into both MMRs.

- **Inside Plant Diversity**

  To improve resiliency, our two redundant POPs in the building are being connected back to back using completely separated conduits taking different interior paths in the building, maintaining minimum common area separation of 50 feet. Conduits are also required to transverse through different rooms/suites to prevent both A/B sides from running through the same room.

- **Outside Plant Diversity**

  In conjunction with the TWDX Hub Express Program, we've constructed new diverse laterals and entrance facilities through North & South Fiber Rooms in the building, with new conduits and duct penetrations owned by TWDX.

  The North Fiber Room lateral to outside street has its own dedicated manhole on Third Avenue (JNF MH C19E1-104), purposed for feeding into the CoreSite data center.

TABLE OF CONTENTS

Inside Plant Diversity

This is absolutely critical for data center customers receiving redundant IP connections. Customer facing gateways should maintain maximum separation to ensure resiliency.

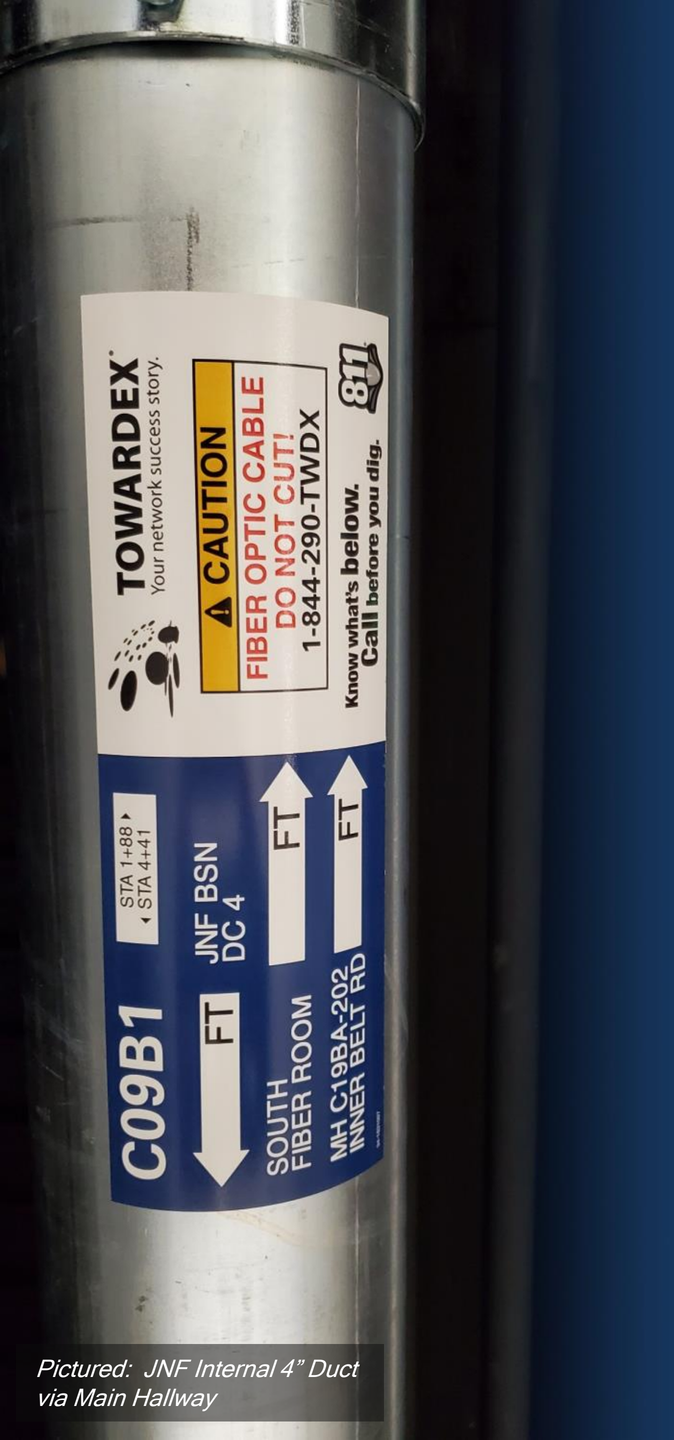

Pictured: JNF Internal 4" Duct via Main Hallway

Two POPs in the data center:

- Each POP connects to different CoreSite MMR in each floor.
- **Primary and Protect side POPs are tied together using:**

A. 192-strands internal fiber cable for most connections

B. By going out to the street then coming back, using completely diverse conduit in different hallway & pathways throughout. 12 strands pre-spliced at outside manhole for backbone circuits that require assured diversity.

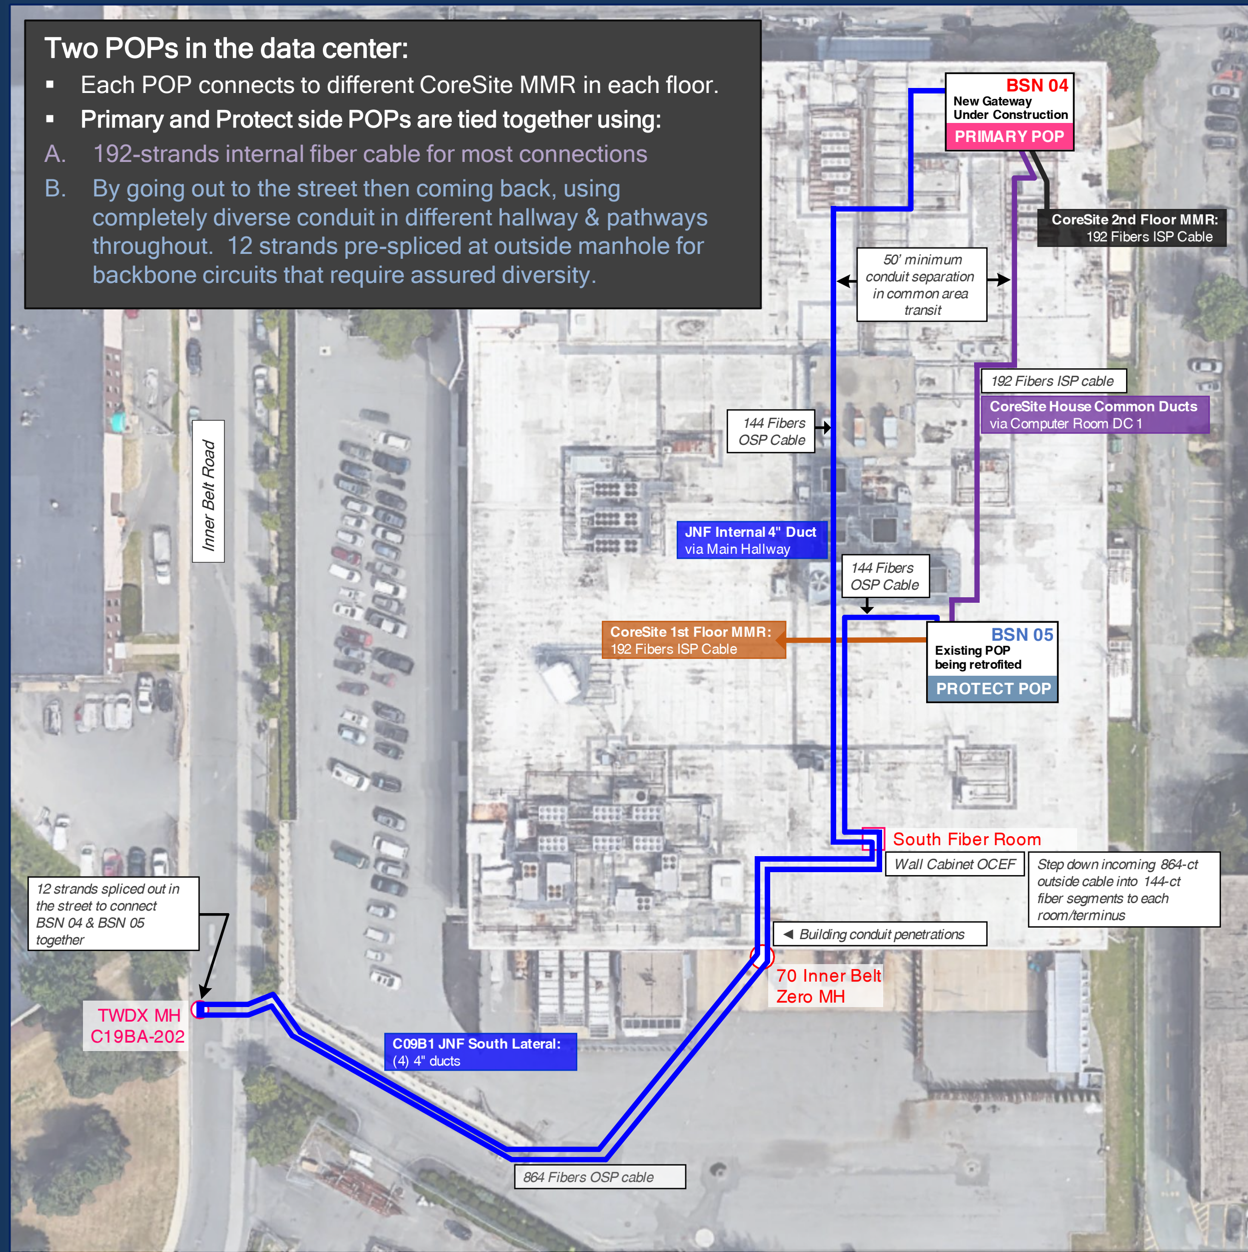

TABLE OF CONTENTS

Outside Plant Diversity

Diverse entrance from both north and south sides, with TWDX constructed utility laterals.

Pictured: A Hub Express System (HEX) Manhole

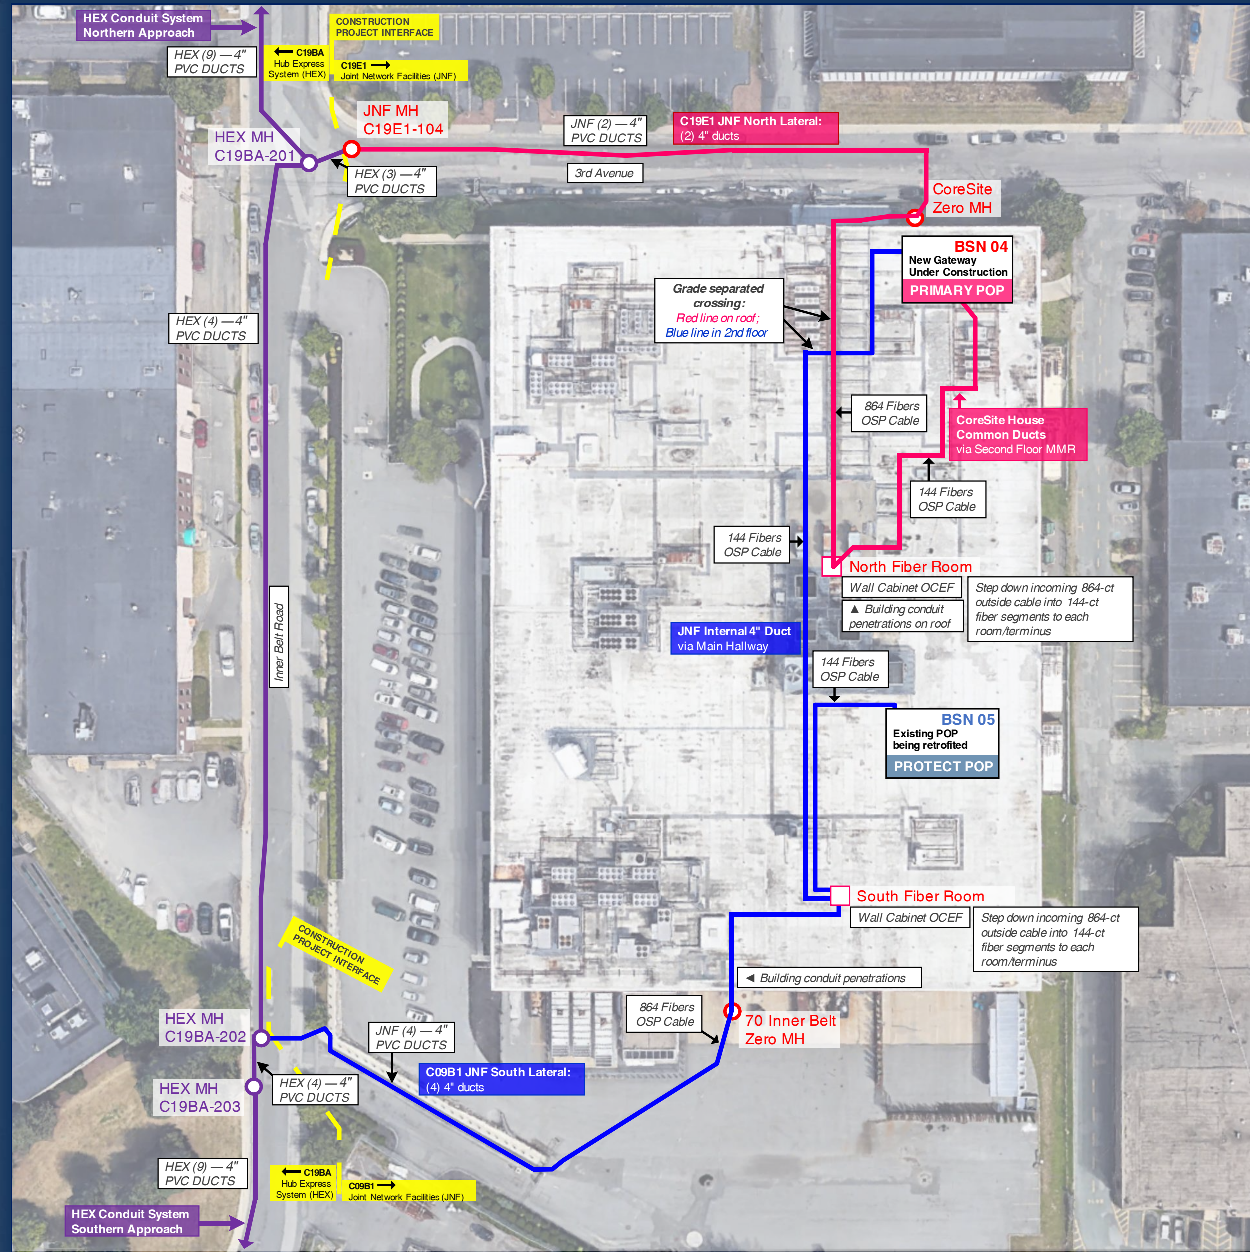

70 Inner Belt South Entrance Construction

The interface between 70 Inner Belt South Entrance and Hub Express System (HEX) out in the street turned out to be more complex than expected. The project area is congested with rights-of-way and easements dating back to 1970, requiring rigorous efforts to locate existing utilities.

A number of change orders had to be made to accommodate existing utilities, including changes to manhole locations and re-design of entry into CoreSite property.

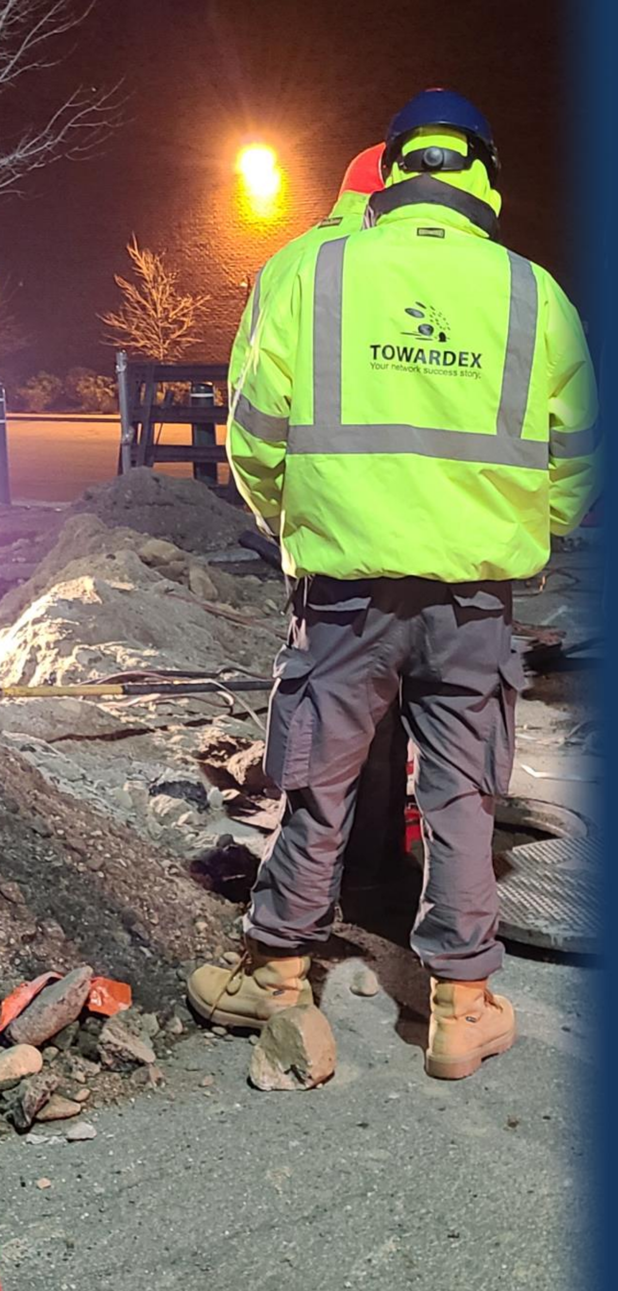

Pictured: Crews are double-checking extents of a utility vault to ensure separation from existing underground facilities.

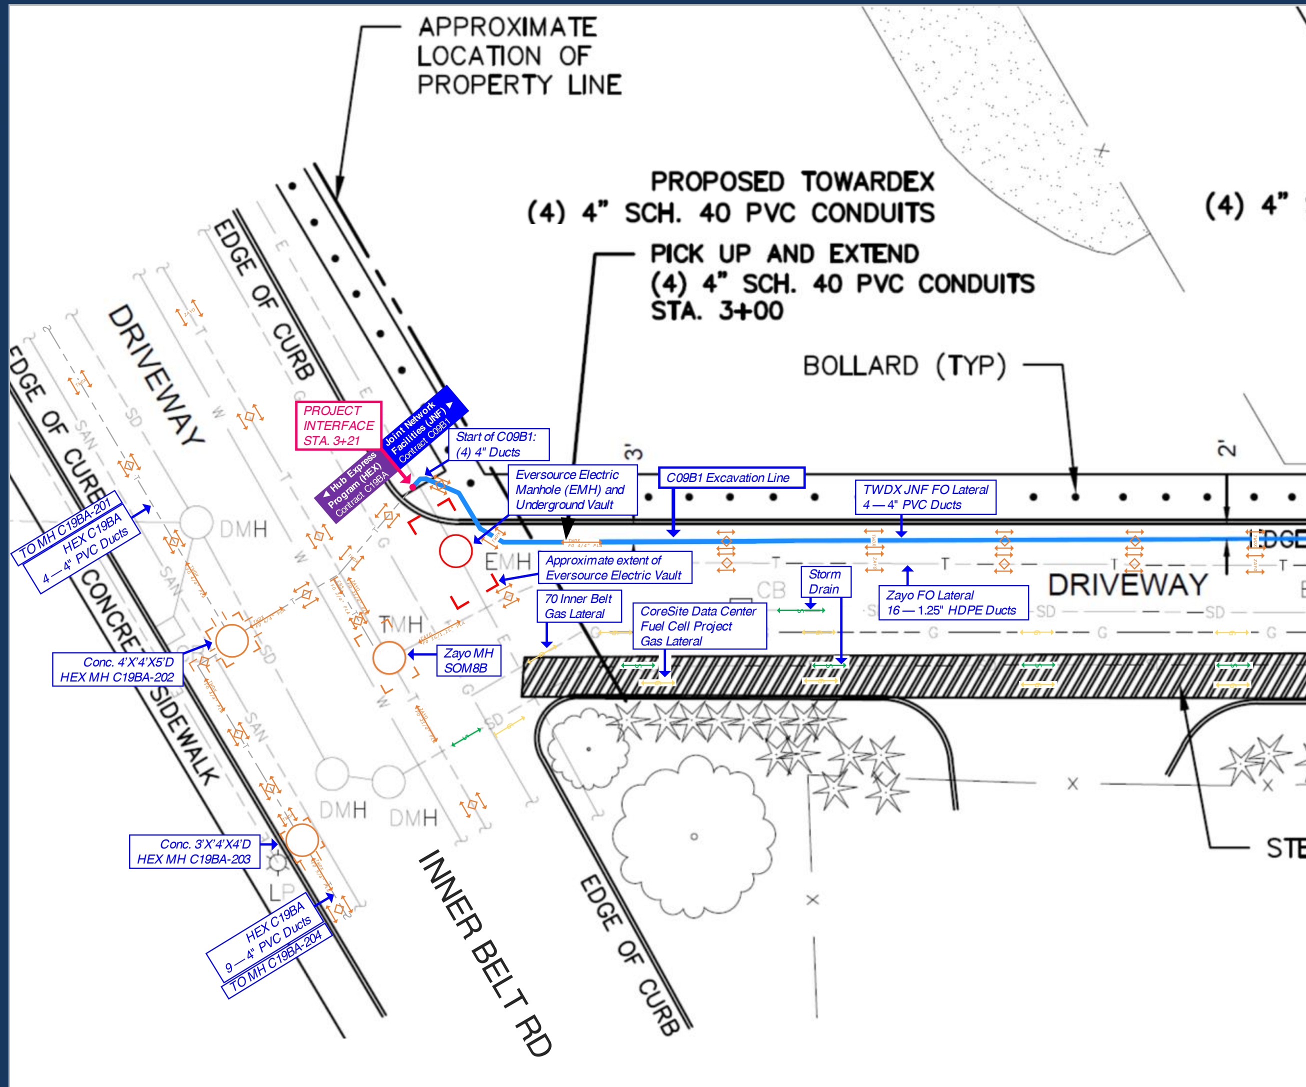

TOWARDEX®
Your network success story.

TWDX IP Updates

AS27552

#PushDaStuff

February 10, 2021

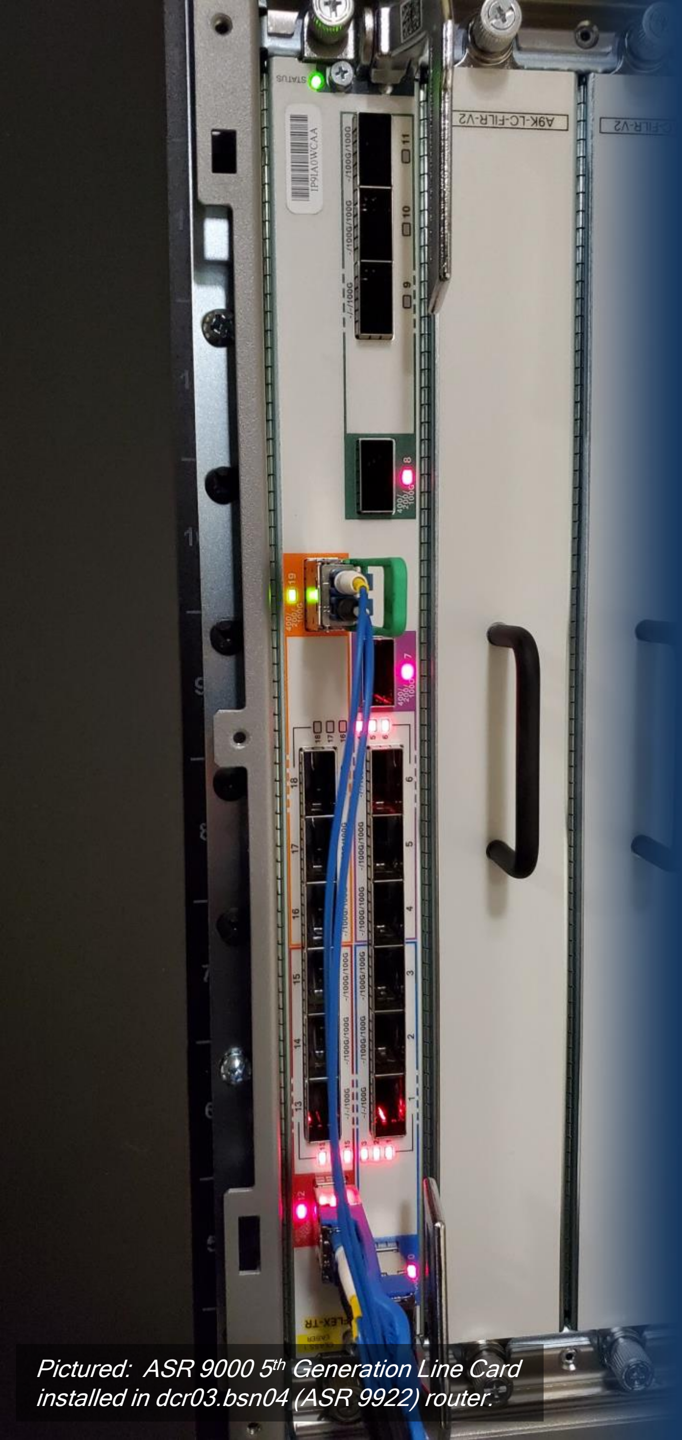

Pictured: ASR 9000 5th Generation Line Card installed in dcr03.bsn04 (ASR 9922) router.

TABLE OF CONTENTS

400G Begins..!

We've begun deploying our first native 400GbE backbone interfaces in Boston. We selected Cisco ASR 9000 Series 5th Generation Line Cards (known as "Lightspeed Plus") and 8200 Series Routers for our 400G deployment.

- **ASR 9000 Series 5th Generation Line Cards (Lightspeed Plus)**

  The following line cards and systems are used to deploy 400GbE at TWDX IP:

  https://www.cisco.com/c/en/us/products/collateral/routers/asr-9000-series-aggregation-services-routers/datasheet-c78-743580.html

  For use in modular chassis-based platforms (ASR 9906-9922):
  - **A9K-20HG-FLEX-TR:** 2 Tbps multi-rate card providing 5x400G or 20x100G
  - **A9K-8HG-FLEX-TR:** 800 Gbps multi-rate card providing 2x400G or 8x100G

  For use in rack space-constrained sites (e.g. New York):
  - **ASR-9903:** High density non-chassis (pizza box) peering router

  Lightspeed+ NPU is a 7nm evolution of Cisco's silicon (nPower run-to-completion chip) previously made for CRS & NCS 6000 carrier core routers. This means that outsourced EZChip silicons are now being phased out for the ASR9K family.

- **Cisco 8201 Router**

  Targeted to function as high density Label Switching Router (LSR) at the core of our IP network (BBRs), 8200 Series routers utilize Cisco's new Silicon One/Q100 chip (known as "Spitfire"). As the new platform matures over time, 8200 Series routers will likely be gradually expanded to cover peering and edge roles as well in the future.

- **Spotting native 400G ports in traceroute**

  `FourHundredGigE` is abbreviated as "`FH`" in IOS XR. Hops showing "`fh-`" in their interface numbering is a native 400G interface:

```
3  ibr01-fh-0-3-0-7.bsn04.twdx.net (198.160.62.5)  0.968 ms  1.017 ms  0.912 ms
```

```
▲ Denotes FourHundredGigE0/3/0/7
```

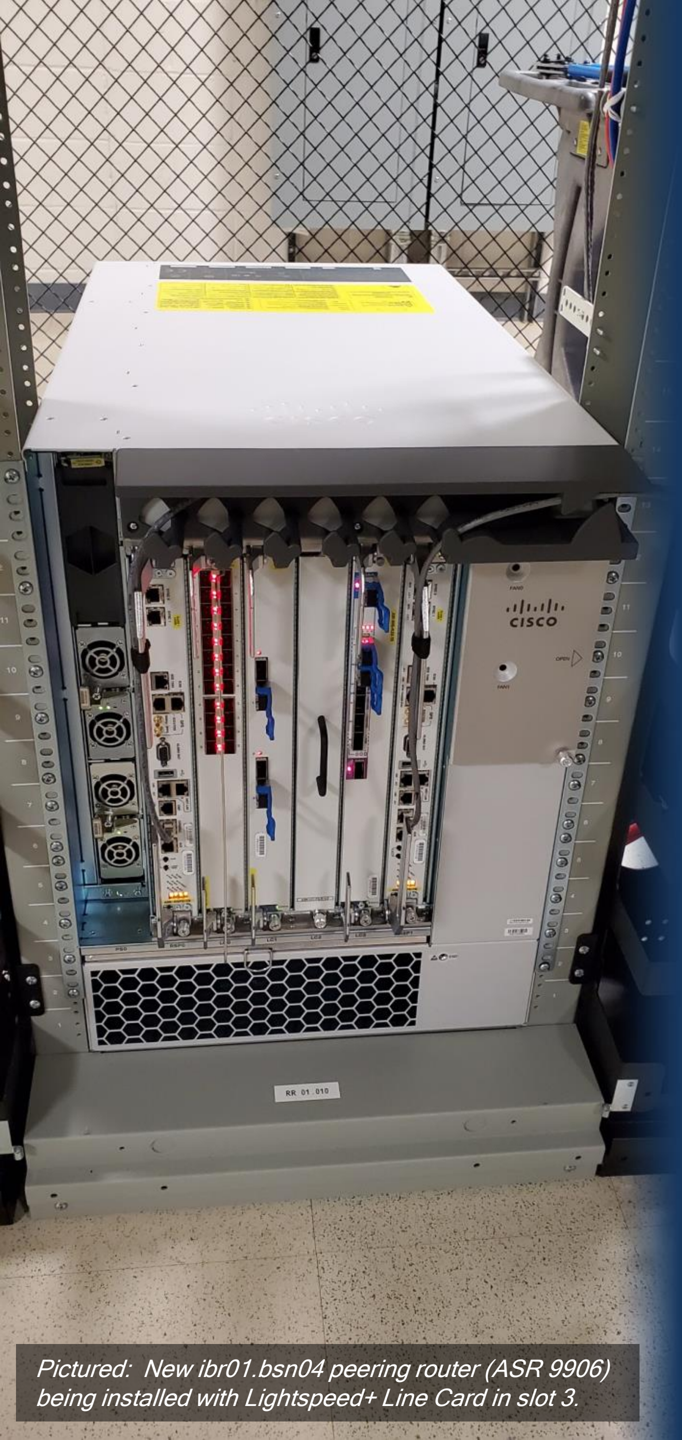

Pictured: New ibr01.bsn04 peering router (ASR 9906) being installed with Lightspeed+ Line Card in slot 3.

Support for BGP Large Communities

RFC 8092 (BGP Large Communities Attribute) is now supported throughout AS27552. Existing RFC 1997 classic communities will continue to remain supported for backward compatibility.

```
For complete list of supported classic communities:
  whois -h whois.ripe.net as27552

'CCC' means Region Code (e.g. 701 = Boston, 702 = New York)
 For global scope (all locations), use '000' for CCC

'Target' is AS number of the targeted adjacent peer
 For TWDX transit customers, use '27552' for Target
 For global scope (all targets, including TWDX customers), use '0' for Target

-----------------------------------------------------------------------------------------------------
Large Community Attribute Format:

27552:4CCC0:Target                 No-export to target
27552:4CCC5:Target                 Force-export to target (countermand wider scoped no-export)
27552:4CCC1:Target                 Prepend 1x to target
27552:4CCC2:Target                 Prepend 2x to target
27552:4CCC3:Target                 Prepend 3x to target
27552:4CCC8:Target                 If supported by the target, request to no-export to its peers
27552:4CCC9:Target                 Reset MED to 0 upon export (enforce hot-potato)
27552:911:CCC                      Blackhole prefix at region 'CCC'


Use Case Examples:

To only advertise to TWDX IP customers, but not to our peers and transits, send these two communities:
  (27552:40005:27552, 27552:40000:0)

Global no-export to everything (including TWDX customers; provides NO_EXPORT 0xFFFFFF01 function):
  27552:40000:0

To Prepend 1x to Comcast (AS7922) at everywhere:
  27552:40001:7922

To Prepend 2x to Lumen (AS3356) at everywhere, but NOT in Boston (701), send these two communities:
  (27552:40002:3356, 27552:47015:3356)

To blackhole a prefix (IPv4 /25-/32, IPv6 /64-/128):
  Blackhole only in Boston region (701):   27552:911:701
  Blackhole at *everywhere*:               27552:911:0

The following classic communities can also be used to request blackhole at *everywhere*:
  27552:911 (classic TWDX blackhole community); or
  65535:666 (RFC 7999)
```

Warning: Blackhole (null route) drops connectivity to the announced prefix. Only use for DDoS mitigations– service outages caused by customer (mis)use of BGP communities are not covered under the SLA.

Automation enhancements

Backend provisioning tools have been updated for adding peers and customer adjacencies into AS27552. Written in Python, the new toolbox provides built-in integrations with IRR and NAPALM for network automation of IOS XR platforms.

```
mc@blackmesa> /usr/local/libexec/ex-bengi/ioxr_adj_customer ~bengi/config_files/customer_docdemo.in -df
  TWDX Routing Policy Generator (IOS-XR) 3.31
  $twdx: blackmesa/bengi/ioxr_adj_customer.pyc,v 3.31 2021-01-09 23:57:12Z summers $
  [>>>>>>>>>>>>>>>>>>>>>] 100.00% - DOCUMENTATI

  Configuration Receipt:
!! Building configuration..[OK]
!! ** Interface Configuration **
!
interface HundredGigE0/10/0/15
 description CUST/INET/DOCUMENTATI/28KJLS092001/NE/100G.bps
 mtu 9216
 load-interval 30
 ipv4 mtu 1500
 ipv4 address 100.74.113.149/30
 ipv4 access-group inet_border-ingress ingress
 ipv6 nd suppress-ra
 ipv6 mtu 1500
 ipv6 address fc00:1eab:701::100.74.113.149/126
 ipv6 access-group inet6_border-ingress ingress
!
!! ** BGP Routing Configuration **
!!
!
prefix-set CUSTOMER:64496
 # inet filter list auto-generated from object TESTDB::AS-DOCDEMO
 192.0.2.0/24 le 32
end-set
!
prefix-set CUSTOMER6:64496
 # inet6 filter list auto-generated from object TESTDB::AS-DOCDEMO
 2001:db8::/32 le 128
end-set
!! Building configuration..[OK]
!
!
router bgp 27552
 address-family ipv4 unicast
  network 100.74.113.148/30 route-policy adv_mpd_customer(ipv4_no-export_smalls)
 !
 address-family ipv6 unicast
  network fc00:1eab:701::100.74.113.148/126 route-policy adv_mpd_customer(ipv6_no-export_smalls)
 !
 neighbor 100.74.113.150
  remote-as 64496
  use neighbor-group inet-customers-full-w-default
  description CUST/INET/DOCUMENTATI/28KJLS092001//NE
  address-family ipv4 unicast
   route-policy mpd_customer_in(CUSTOMER:64496, ipv4_no-export_smalls, 1) in
   !
  !
 neighbor fc00:1eab:701::100.74.113.150
  remote-as 64496
  use neighbor-group inet6-customers-full-w-default
  description CUST/INET/DOCUMENTATI/28KJLS092001//NE
  address-family ipv6 unicast
   route-policy mpd_customer_in(CUSTOMER6:64496, ipv6_no-export_smalls, 1) in
   !
  !
!! %EXMaster: received valid candidate config (XR_SYSDB_SYNTAX_PASS=True), starting push
!! %EX_Python_Interpreter> ex_open_napalm_driver(): Opening network driver to dcr01.lab01
!! %EX_Python_Interpreter> m.open()
!! %EX_Python_Interpreter> m.load_merge_candidate(config=M_DELTA)
!! %EX_Python_Interpreter> m.commit_config()
!! %EX_Python_Interpreter> m.close()
!! %EXMaster: commit operation on dcr01.lab01, 1610993293 is complete
!! %EXMaster: DDNS update to con-auth102.ns.twdx.net for 149.113.74.100.in-addr.arpa
!! %dns.update: failed for 149.113.74.100.in-addr.arpa: SOA query unsuccessful
!! %EXMaster: DDNS update to con-auth102.ns.twdx.net for 150.113.74.100.in-addr.arpa
!! %dns.update: failed for 150.113.74.100.in-addr.arpa: SOA query unsuccessful
!! Execution completed for DOCUMENTATI
```

Improved automation of config push and activation onto customer & peer attachment routers

Pictured: dcr03.bsn04 (ASR 9922) customer aggregation router

TABLE OF CONTENTS

A Terabit of connectivity to the internet.

For a small blended IP provider in New England, we're pretty well connected – biased as we may sound, we're probably better connected than most other providers in Boston ;-).

This was realized by constructing a modern peering fabric which vertically integrates all layers: from automation to router interfaces, to optical transport, and all the way down to dark fiber and physical conduits & cable vaults. We will discuss this more in next slides.

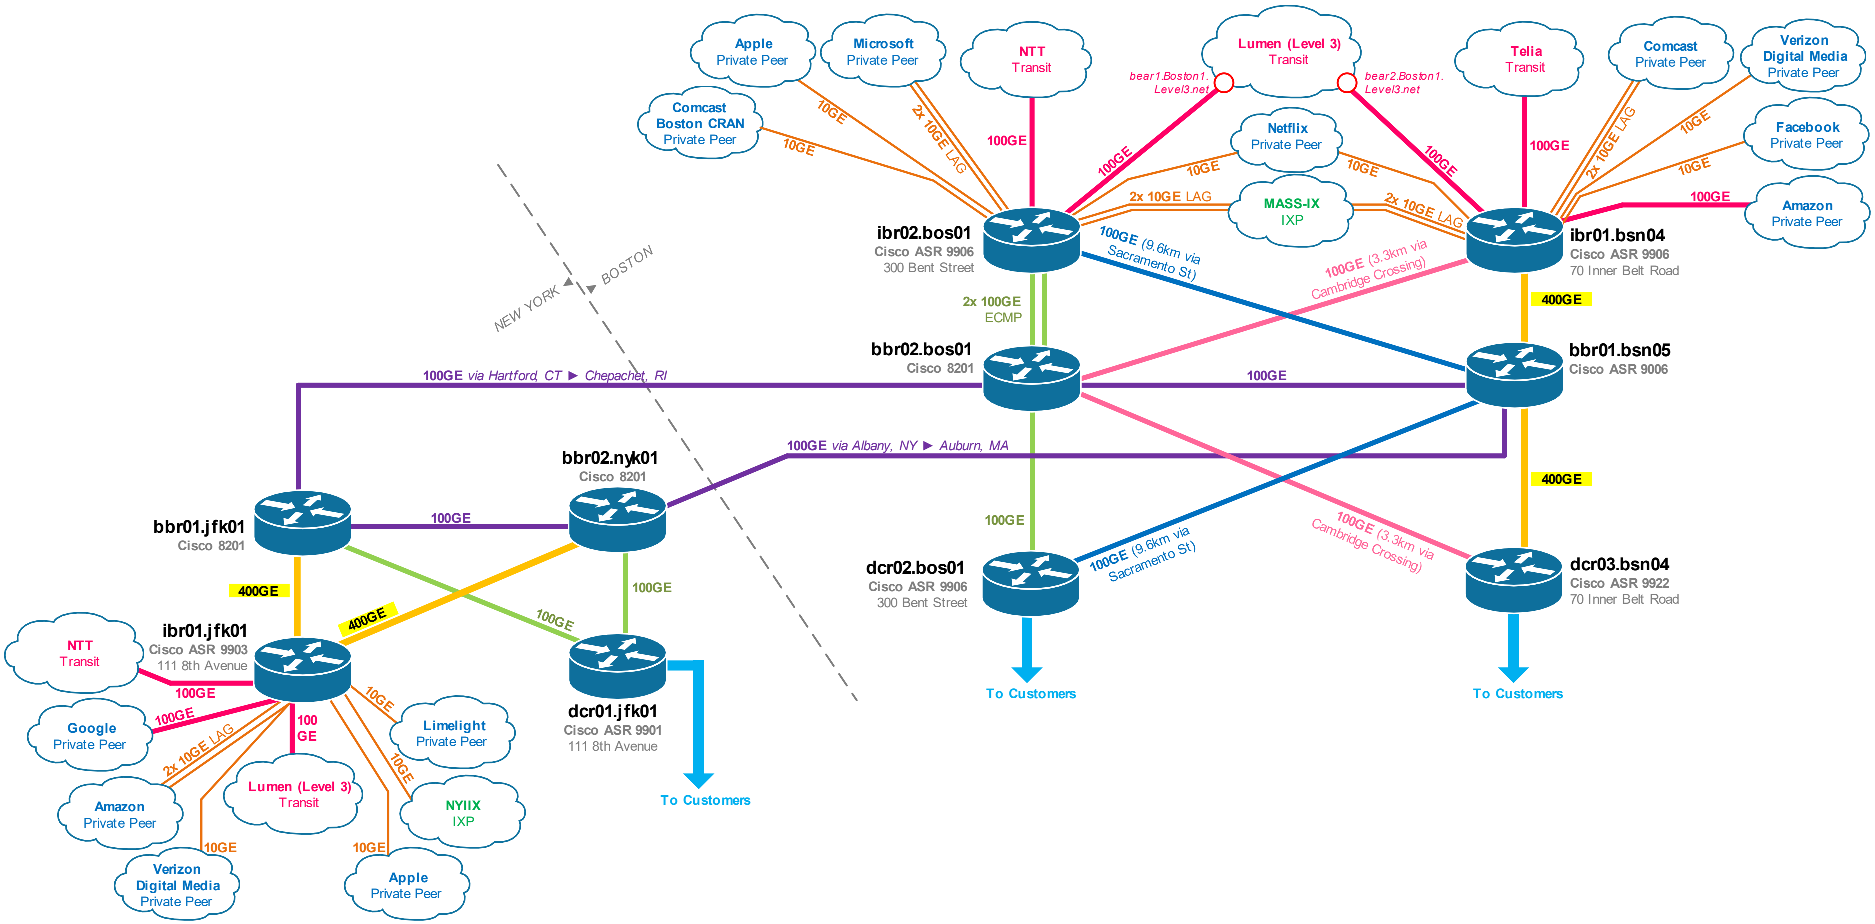

TABLE OF CONTENTS

TOWARDEX®
Your network success story.

Building our new Metro Peering Fabric

From physical plant to optical layer

January 7, 2021

TABLE OF CONTENTS

Vertical Integration for Peering

Pictured: Crews overseeing manhole penetration of 2 - 4" conduits for interconnection from another network provider into TWDX underground utilities.

Disrupting IP transit in a less developed market like Boston isn't always simple as just plopping your POP and lowering price. We'd need to design the architecture around getting the customer traffic off of our network as quickly and locally as possible.

The internet is a series of tubes after all..
Peering, transit, automation and Big Expensive Routers are all cool toys, but they're all overleveraged by access to something so seemingly basic and simple as dark fiber.

To vertically integrate processes for a metro peering architecture, we need to first analyze different methods we use to make physical connections with others. Let us review:

- **Recruit Level: Cross Connects via Meet-Me Rooms (MMRs)**

  Cross connects are highly convenient and quick way of making physical connections between networks. When available, this is by far the easiest mode of connection.

  However, Cross Connect is typically economical when both networks are situated in the same common facility. Cross Connect is also vulnerable to third party interference, where service can be hindered when competing colocation providers are fighting.

- **Veteran Level: Outside Plant (OSP) Interconnections**

  Interconnections can also be made out in the street by running a cable between each party's enclosure and splicing mutually agreed fiber strands. This works best when both parties have unrestricted access to dark fiber, or when at least one party has significant control over network assets. OSP Interconnection is often time consuming, has higher upfront costs and is inflexible. However, it is generally immune to third party interference and typically does not incur recurring costs of Cross Connects.

- **#YOLO Level: Asset swaps and counter-leveraging**

  An advanced concept from OSP Interconnections above, where utility asset owners mutually agree to provide fiber loops or license conduits to each other at pre-agreed rates, or on settlement-free basis. This is pure power play involving leveraged negotiations and long-term strategy and is often the hardest to achieve. Access to capital and advanced know-how of right-of-way operations, heavy construction, permitting and regulatory processes are a must.

Pictured: Construction of a small demarcation manhole (DMH) which functions as the border interface between TWDX underground conduits and another utility company.

All of them are good options.

In the previous slide, we went over the three common methods to physically interconnect networks in the metro. The reality is that there is no right or wrong or superior way out of all three options. We're not going to discuss virtual interconnections or internet exchanges here, as these are out of the scope of our objectives.

Each option has its own share of benefits and complications. Thus, we needed to build a solution which integrates all 3 options to maximize our interconnection flexibility in the metro.

To build upon this goal, we first start from the physical layer:

It's no secret that acquisition of rights-of-way and new constructions for TWDX are often most influenced by IP network engineers, with targeted strategy in expanding options for interconnections with other networks.

These targeted constructions of utility infrastructure and trench partner engagements don't just improve our connectivity. They also naturally increase competition for interconnection facilities and broadband infrastructures for everyone, helping to level the playing field for all. That is the beauty of interconnection – it takes two to tango.

When manholes become Meet-Me Holes..

Telecommunications manholes are an important nexus for modern communications. They provide convenient attachment between middle and last-mile infrastructures for broadband, and keep data center ecosystems well supplied with connectivity options.

When constructing new TWDX underground utilities, multi-tenancy is an important design specification (i.e. building additional ducts and vault space to accommodate other carriers). This not only helps enable other providers to more cheaply install fibers, but also greatly enhances our ability to interconnect with them, by welcoming them into the same physical facility as where our own backbone cables are located.

For facilitating manhole transit, 864-ct fiber optic cable is our preferred workhorse– they offer a good balance between cost vs. capacity.

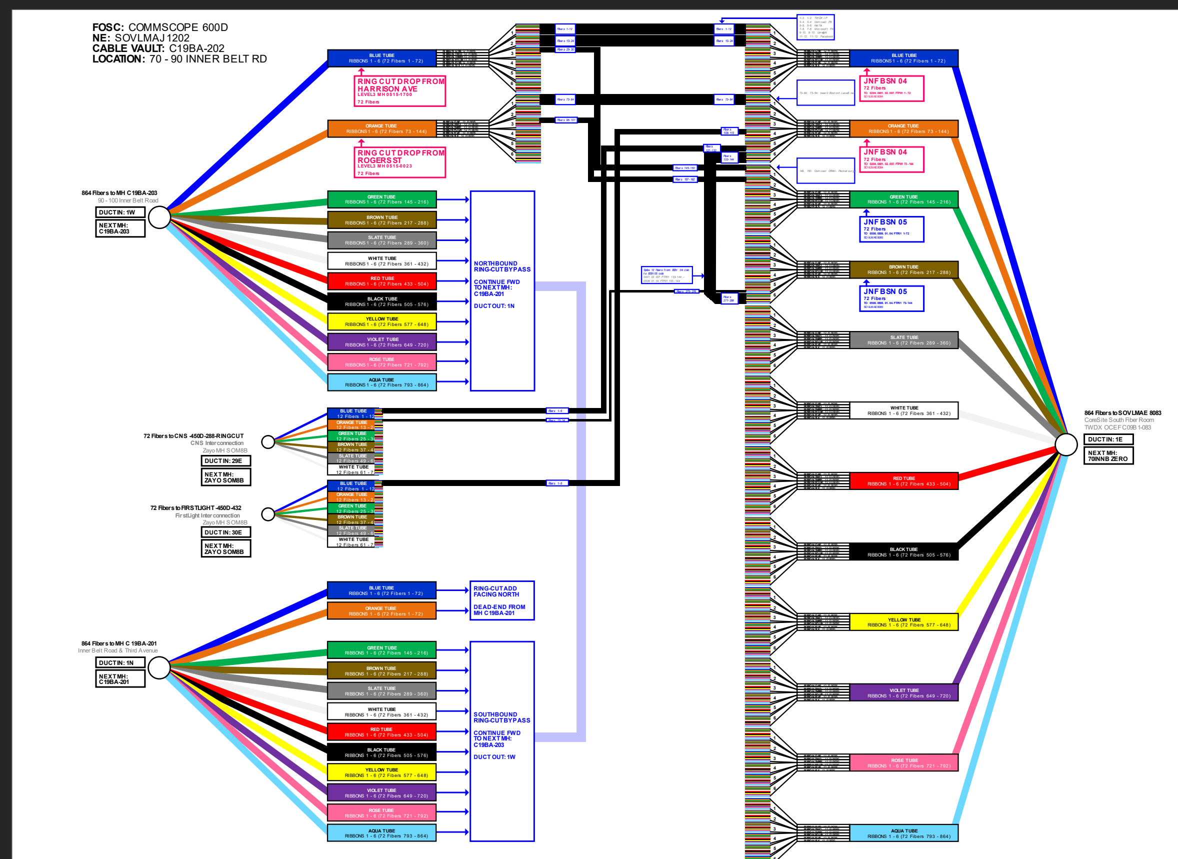

TABLE OF CONTENTS

Optical Peering Fabric for TWDX IP in Boston Metro

To tie all of this together, optical transport architecture (ROADMs and DWDMs) are integrated together with outside plant infrastructure, maximizing our interconnection options in the metro.

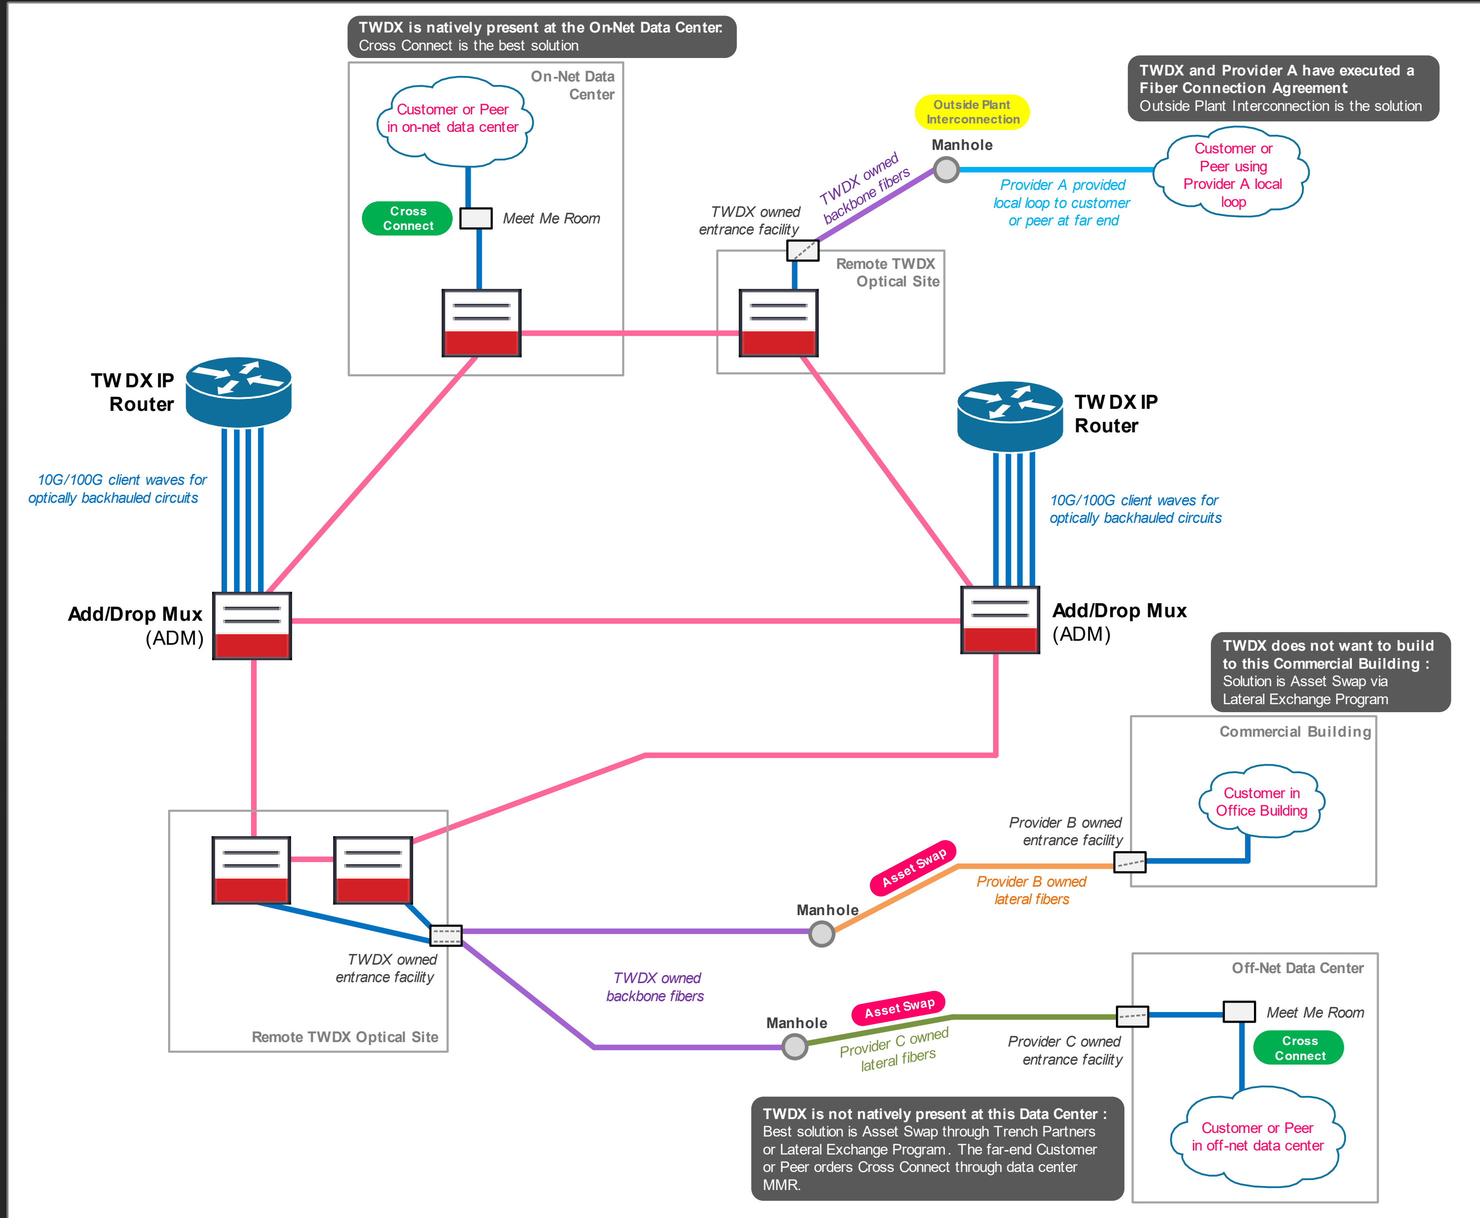

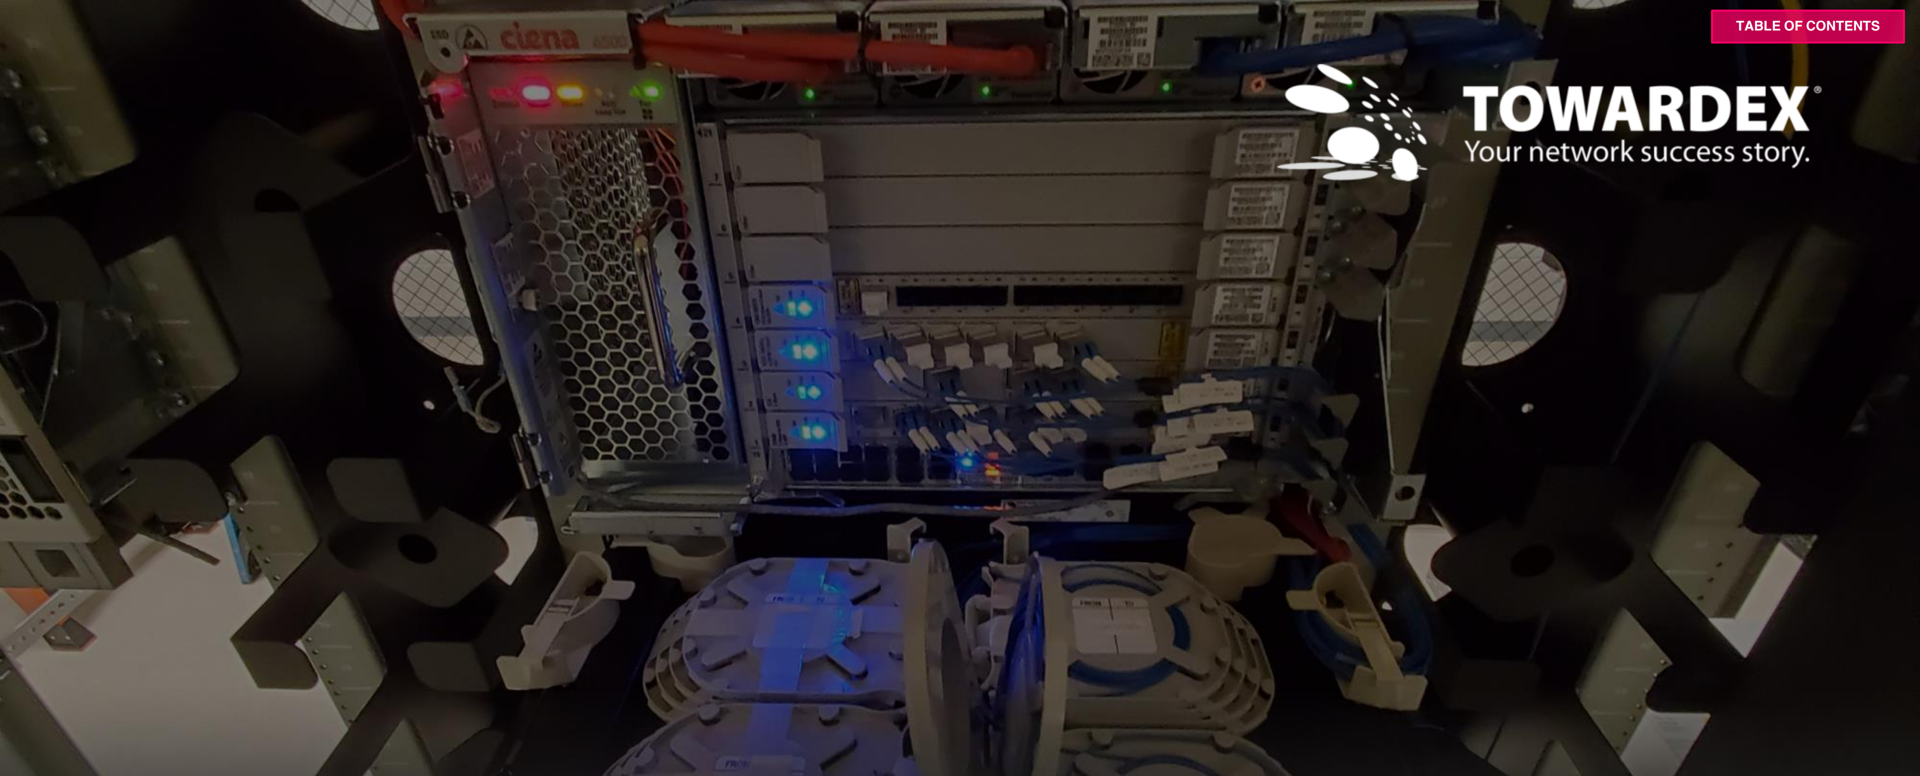

TOWARDEX
Your network success story.

Examples of some interesting interconnections into TWDX IP

#KeepTrafficLocal

December 17, 2020

TABLE OF CONTENTS

Pictured: New 400G muxponder installed to terminate bear2.Boston1.Level3.net

Lumen - TWDX IP Interconnection

Both Lumen (AS 3356) and TWDX IP (AS27552) operate redundant IP routers in Boston. The parties have engineered redundant private exchange (PX) connections to each side, taking advantage of all appropriate facilities which are available.

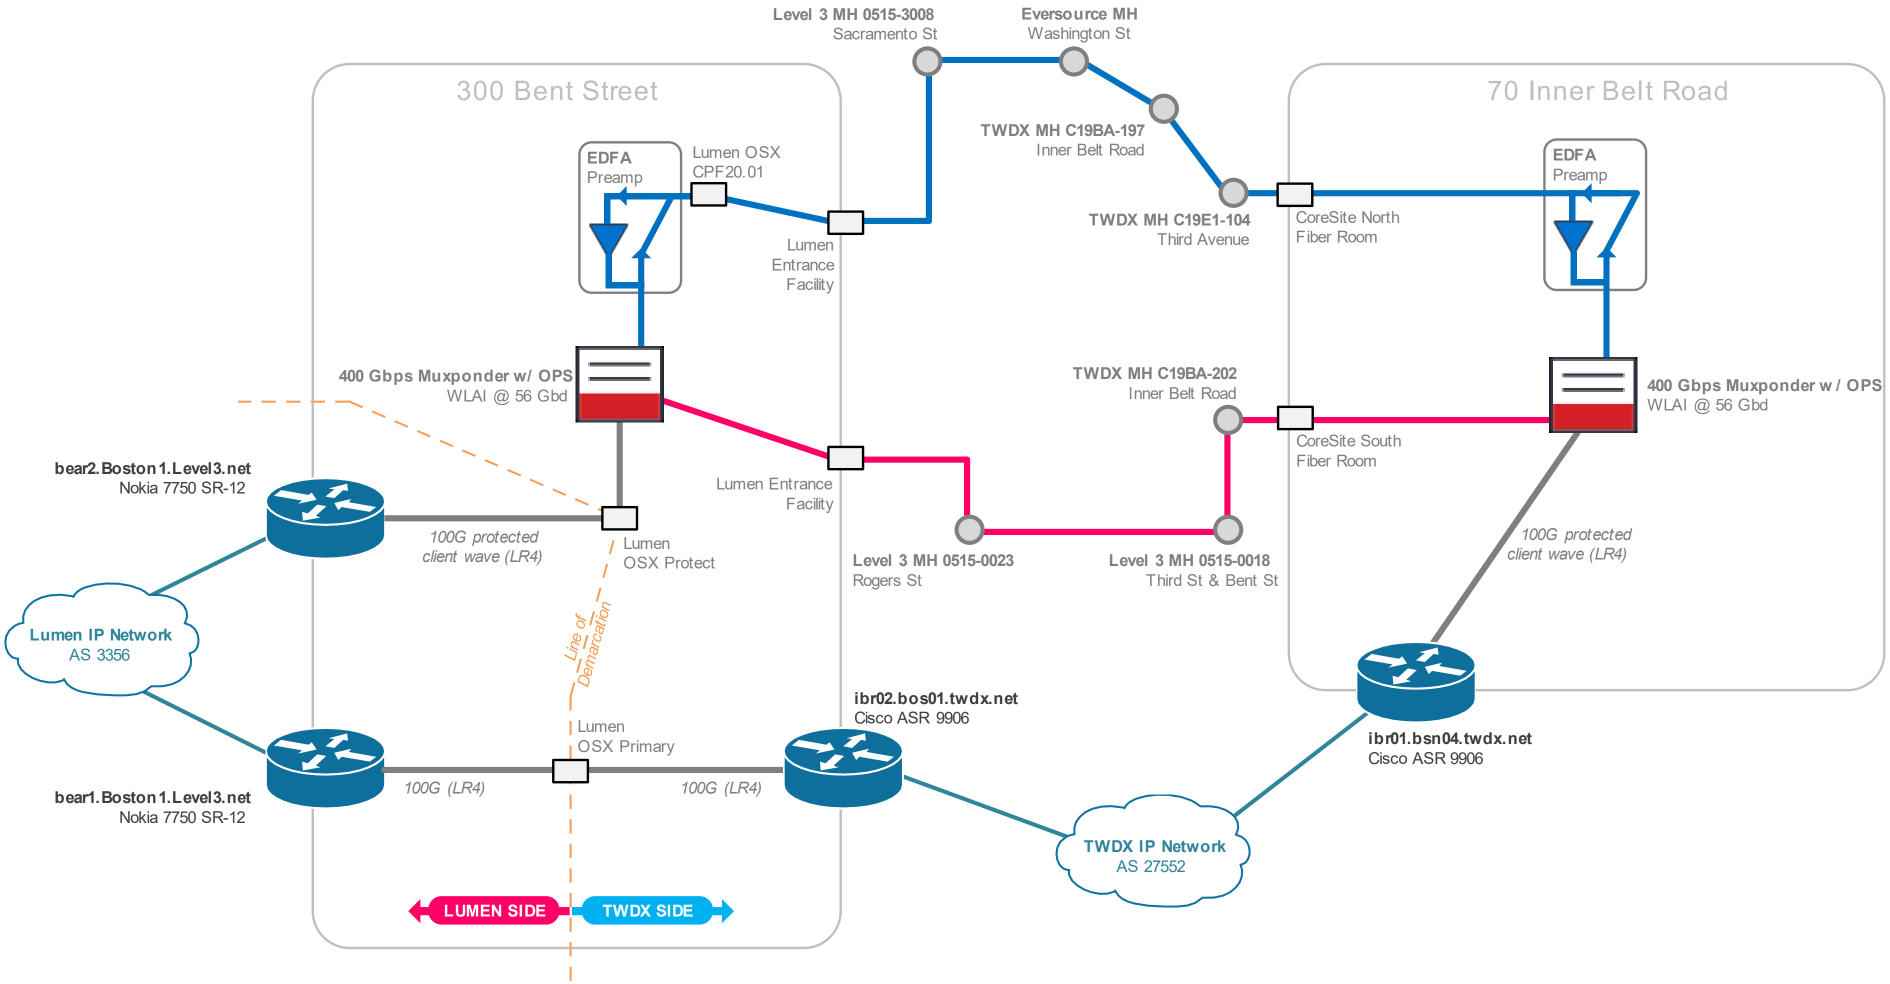

TABLE OF CONTENTS

Comcast - TWDX IP Interconnection

Pictured: 400G muxponder being installed for PX links from 22 Linnell

Comcast operates two networks in Boston - national IP backbone (ibone) and the Boston Converged Regional Area Network (CRAN). The Boston CRAN is where all of the regional eyeballs are connected to, and it's centered in Needham & Woburn out in the suburbs. This topology is taken into account when developing the PX implementation plan between Comcast and TWDX.

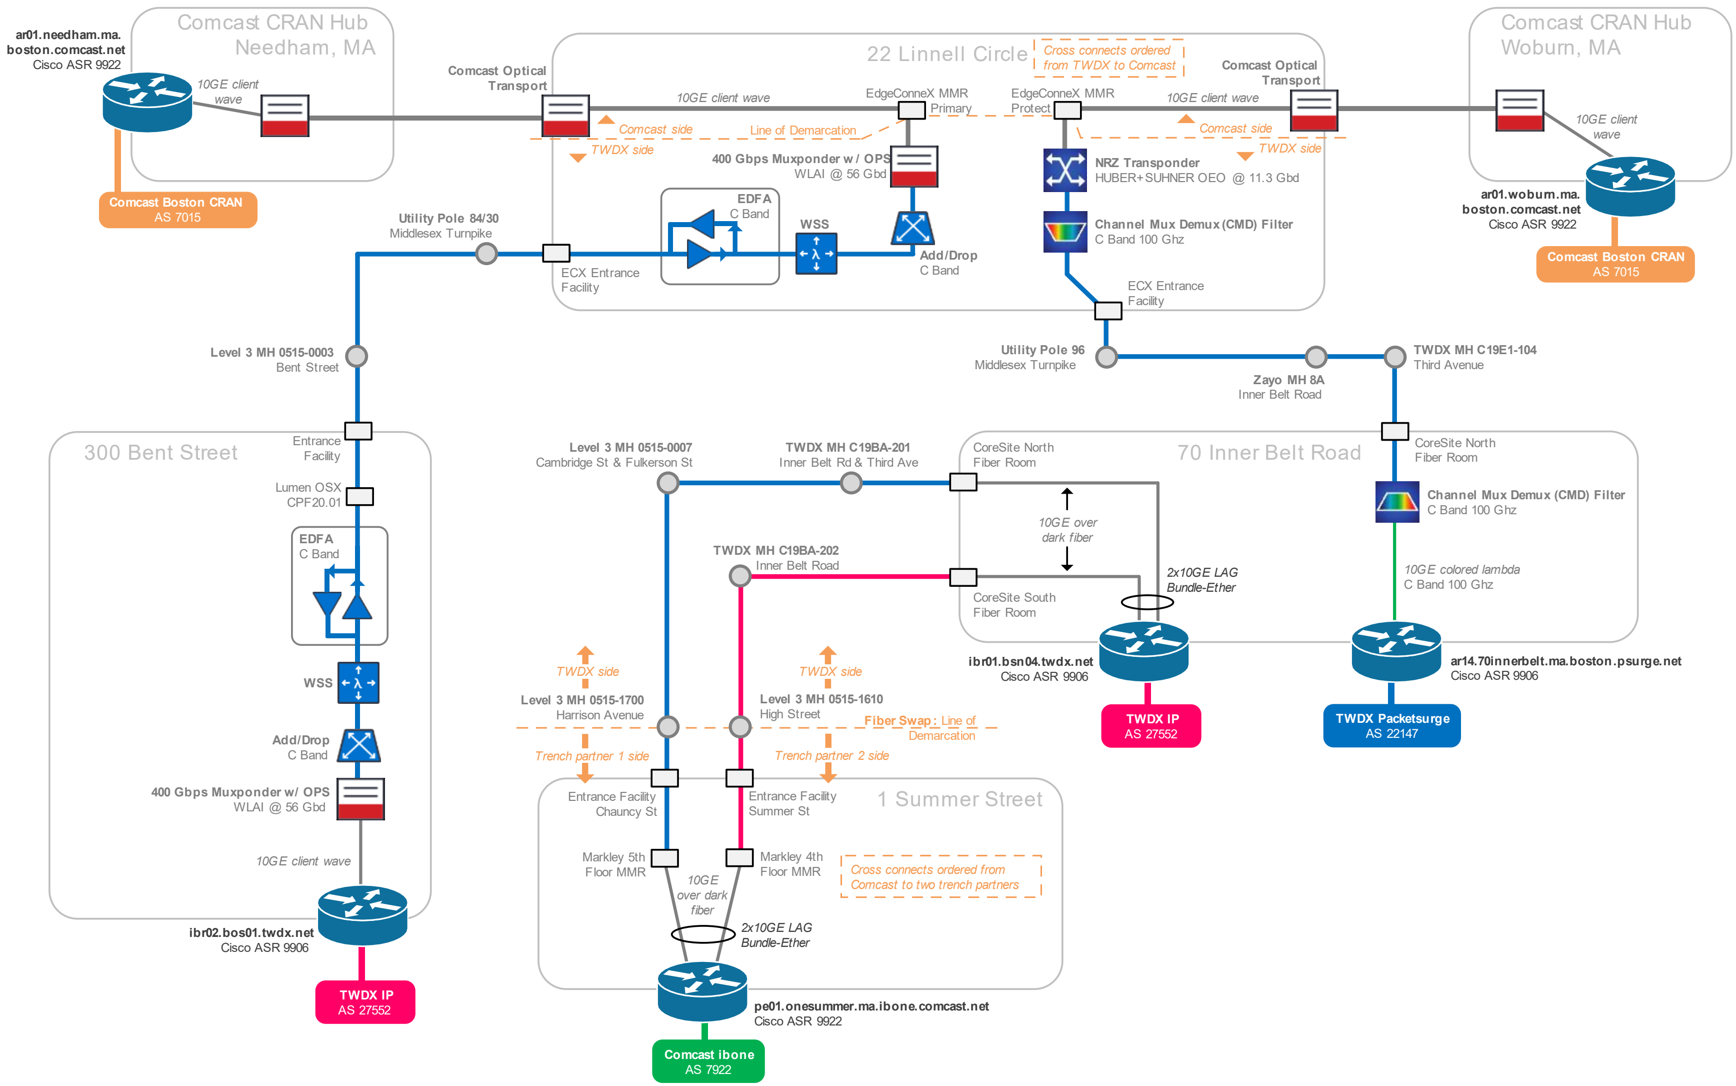

Any Questions?

Corporate Headquarters:
TOWARDEX Technologies, Inc.
One Marina Park Drive 14th Floor
Boston, MA 02210
617-849-7278

For general inquiries and sales:
inquiry@towardex.com

For Network Operations:
ip-admin@twdx.net
1-844-290-TWDX

For MASS IX questions:
ops@mass-ix.net