

Layer3 Network Architecture

R&E IP, TR-CPS, and Net+ Services

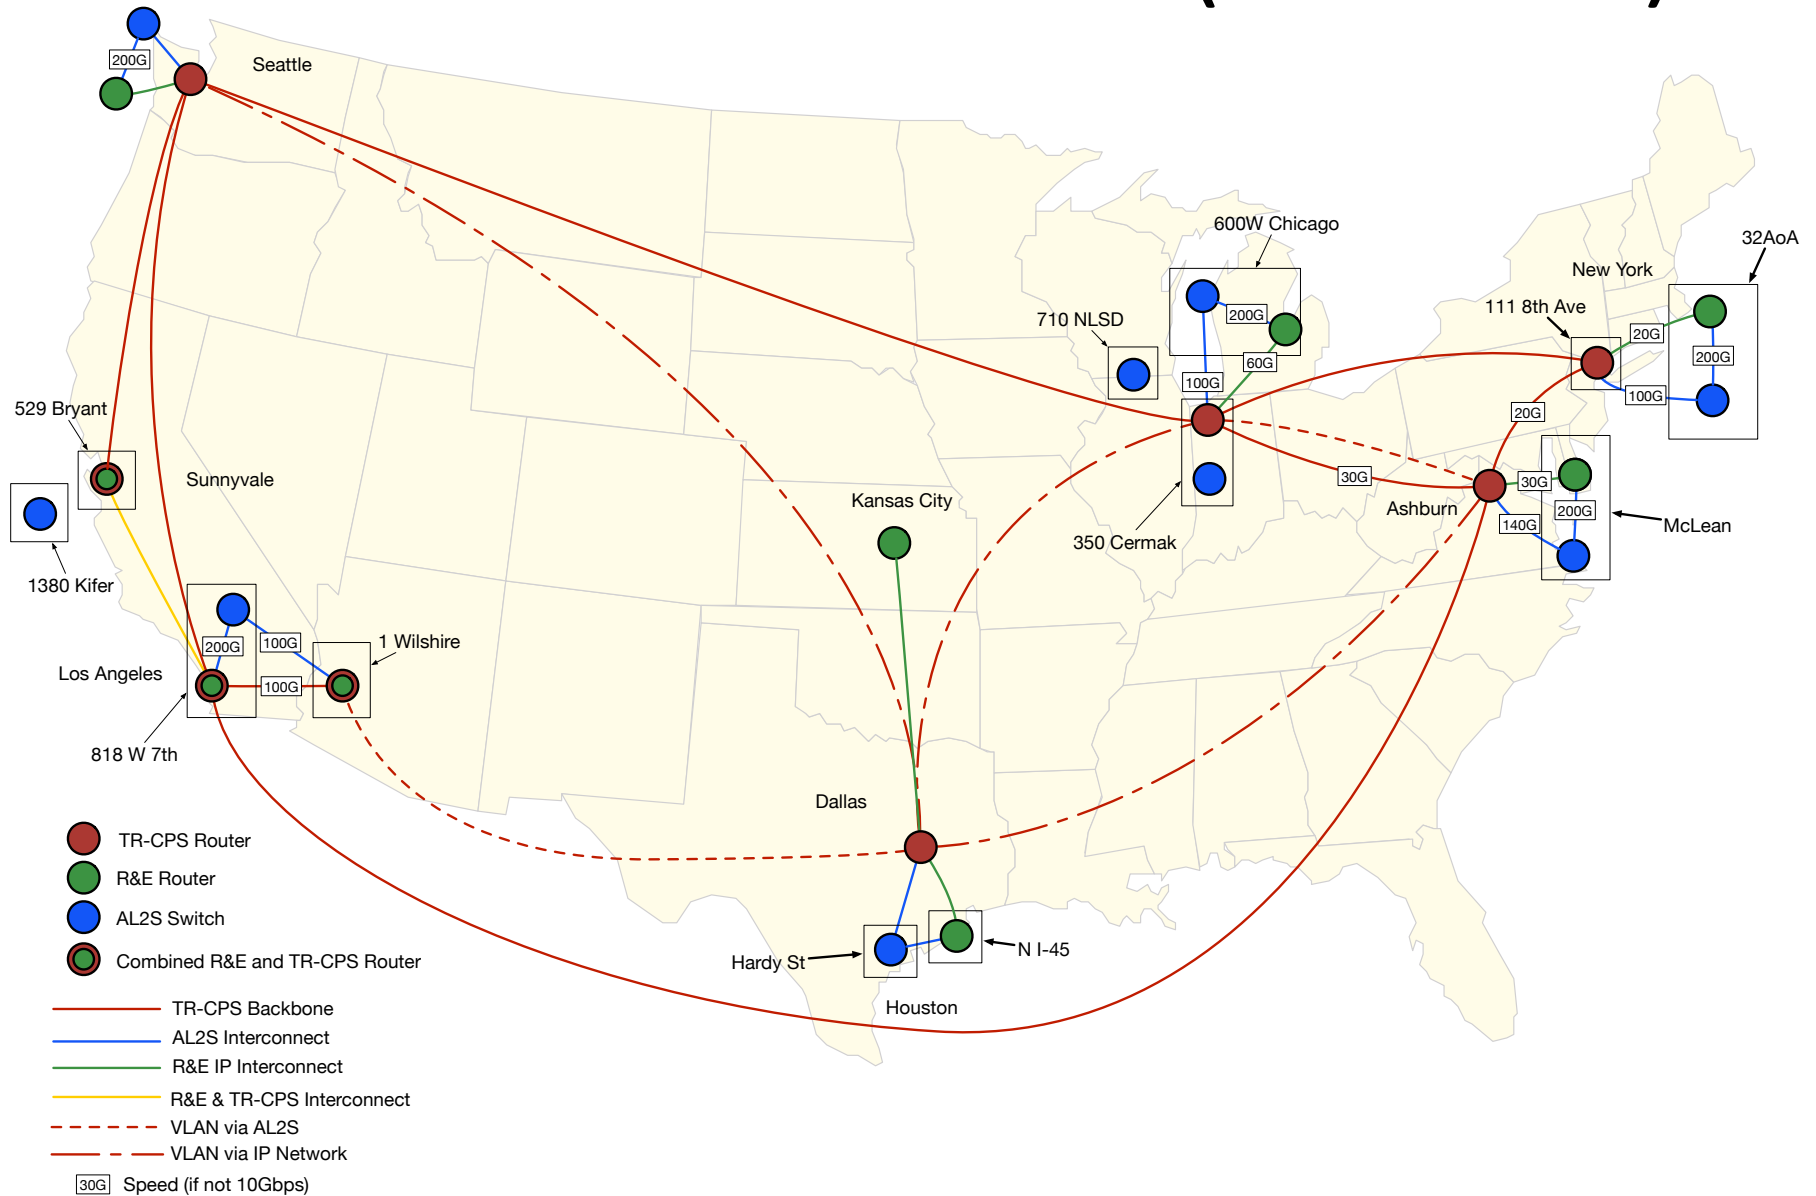
Architecture through 2014: Physical

- R&E and TR-CPS carried on separate hardware
 - Both networks Juniper MX-960 based
 - 10 R&E routers ([AS11537](#))
 - approx 14,000 IPv4 routes
 - 100G backbone via AL2S
 - 7 TR-CPS routers ([AS 11164](#))
 - approx 250K IPv4 routes
 - 10G (and Nx10G LAG) backbone via Internet2 optical network

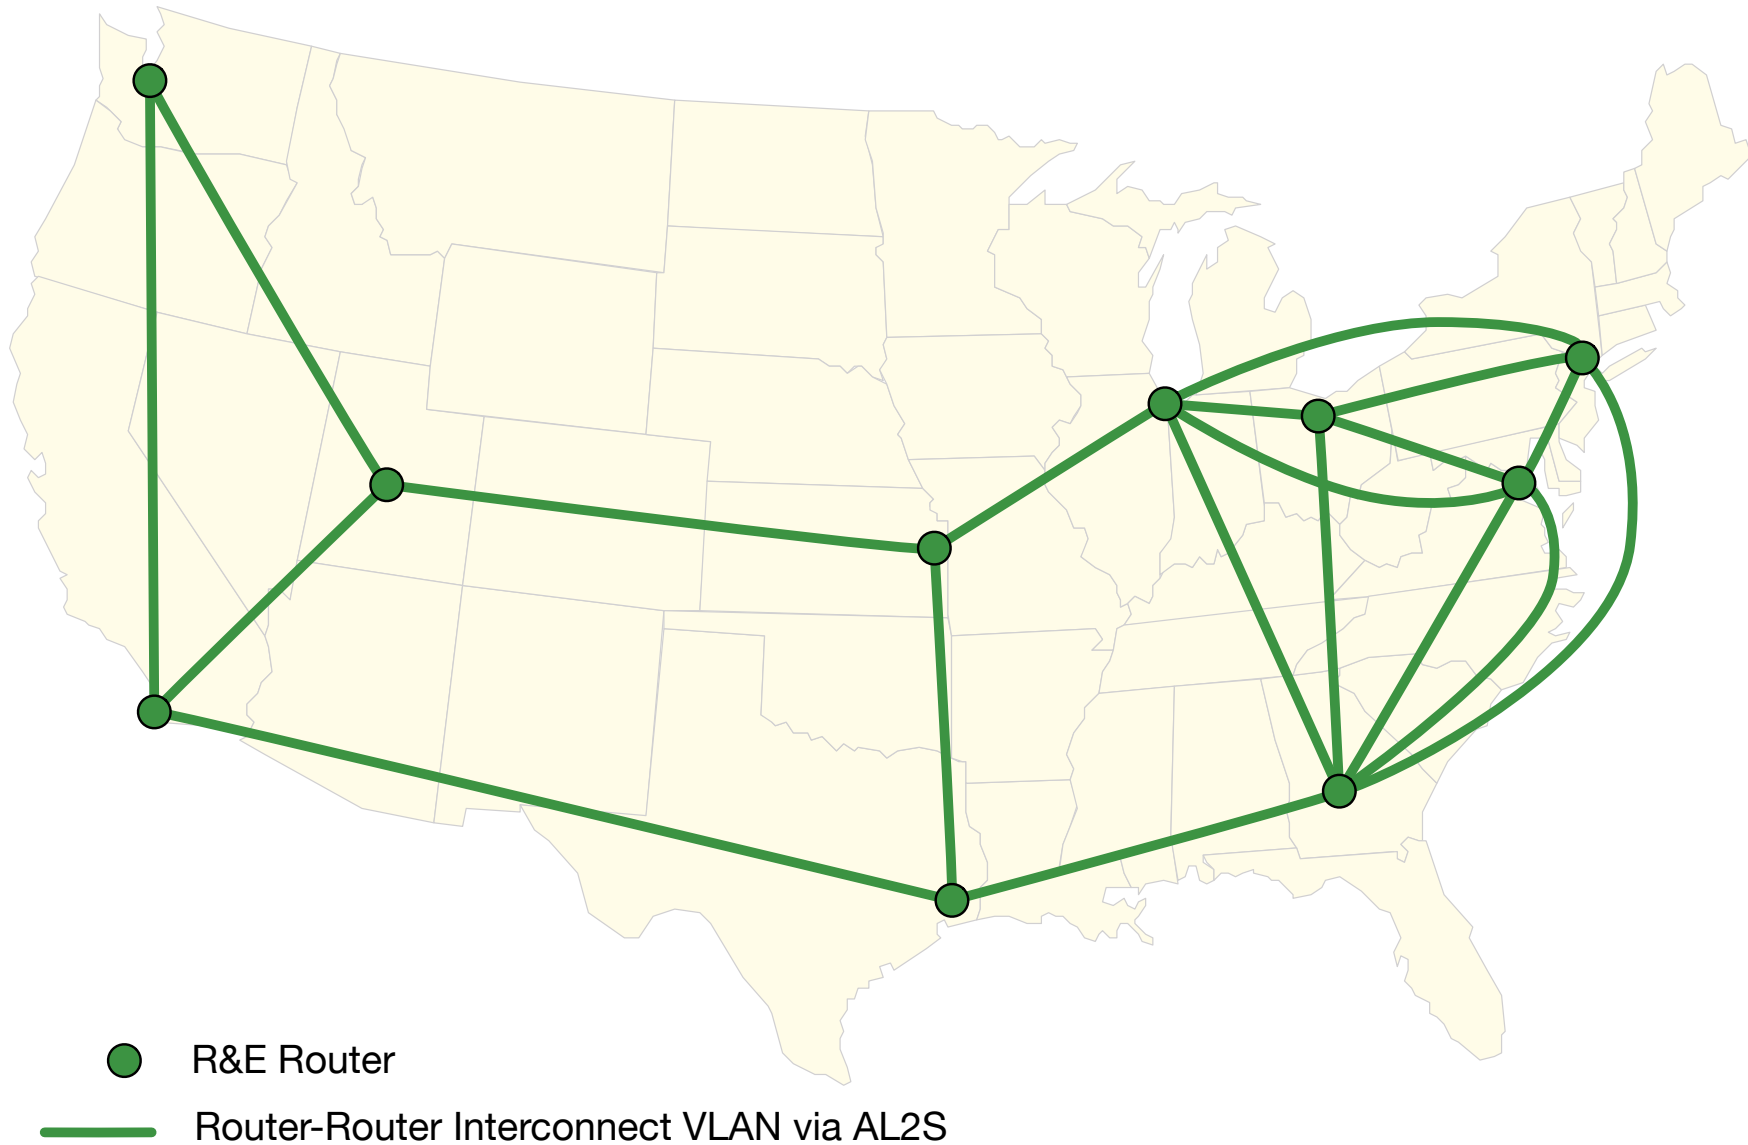
Architecture through 2014: logical

- Access to each service via separate BGP session
 - R&E accessed via AL2S VLAN or direct interconnect to **AS11537** routers
 - TR-CPS accessed via AL2S VLAN, IP Psuedowire over **AS11537**, or direct interconnect to **AS11164** routers
- Net+ carried only in **AS11537**
 - Interconnects to Net+ providers via direct peering with **AS11537** (physically connected via AL2S or direct router interconnect)

TR-CPS Infrastructure (Jan 2015)



Late 2014 Layer3 Infrastructure



Future Architecture

- Provide both **AS11537** and **AS11164** instances on most routers (“split brain”)
 - **AS11537** in the default routing-instance and configure **AS11164** as a VRF (routing-instance).
 - Since today **AS11537** is the default on R&E routers, and **AS11164** is expressed as a logical-system on TR-CPS routers, migration is reasonably easy.
 - Provides the future ability to interconnect and share TR-CPS, R&E, and Net+ routes in a flexible manner at more locations
 - Can reduce or eliminate the use of TRCPS pseudowires to the Connector

Future Architecture

- Consolidating onto one platform
 - This has been done already at 818 W 7th Los Angeles.
 - Makes sense today also in Seattle (reduce to one MX 960).
 - Probably does not make sense in Chicago, Washington DC, or New York; routers are in diverse locations, providing additional reach
 - Retain all other TR-CPS routers and offer **AS11537** services.
 - Add **AS11164** on R&E routers.
- Support for direct physical interconnect to “Split Brain” routers for the purposes of exchanging TR-CPS traffic
 - *NOTE: Need to determine costing model*

Future Architecture

- Backbone interconnect design
 - In the short term, continue to primarily interconnect [AS11164](#) traffic with Nx10GE LAGs.
 - Growth can be handled using available 10G capacity on the optical platform at low/no cost.
 - Allows us to examine advantages or disadvantages of combining over AL2S on a case-by-case basis.
 - Immediate judicious use of AL2S for rapid growth provisioning and restoration capabilities.
 - Load balance heavy Net+ traffic onto AL2S paths where needed.

Management

- Peering Steering Group
 - Advisory group that receives bi-weekly detailed updates on peering requests, architectural changes, and operational issues
 - Assist in monthly update to NTAC peering wg
 - Provide feedback to staff on direction
- Network consolidation will collapse public TR-CPS and R&E operational data into a single instance, though data can be pulled out if desired.
- Provide consistent metrics on monthly NTAC peering wg calls
 - *NOTE: need to determine metrics with peering steering group*

Management (cont.)

- Headroom policy updated in June 2014
- https://internet2.app.box.com/files/0/f/2913004843/1/f_24944921413
- Depending on the distribution of TR-CPS routes, we likely want to update with specific language on backbone links that exclusively carry commercial traffic
 - Less bursty and more deterministic. 40% is probably too aggressive
 - Need to accommodate restoration capacity