Network Design Evolution -- (T-Mobile USA Case Study)

Issa AbuEid

NANOG68

October 18 2016

Abstract

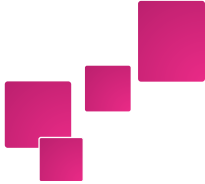

- In this session, we will describe the main design model and components used to build the T-Mobile USA network, present major evolutions in the last ten years such as architecture, technology, and routing protocol changes. Finally, some lessons and recommendations learned from our experience will be shared.

Agenda

- Hierarchical Network Design Model
- T-Mobile Network at a glance
- T-Mobile Design Model Evolution
  - Architecture Evolution
  - Technology Evolution
  - Routing Protocol Evolution
  - BGP RR Evolution
- Lessons Learned and Recommendation

Hierarchal Network Design Model

- T-Mobile USA decided to use *The three-tiered hierarchy* model as the preferred approach to network design. In the three-layer network design model, *network devices and links* are grouped according to *three layers*:
  - Core
  - Distribution
  - Access
- A *hierarchical network design* model breaks the complex problem of network design into smaller, more manageable problems.
  - *Optimize* network hardware and software to perform specific roles in each *tier*
  - E.g., devices at the access tier accept traffic into a network and pass to the higher layers
- The three-layer model is a *conceptual* framework – similar to the concept of the (OSI) reference model

The Core Layer

- The *core layer*, aka *backbone,* provides an *optimized* and *reliable transport structure* by forwarding traffic at *very high speeds*. In other words, the *core layer switches packets as fast as possible*.

- The major features of Core layer :
  - Resilient
  - Few changes
  - Few features
  - High bandwidth

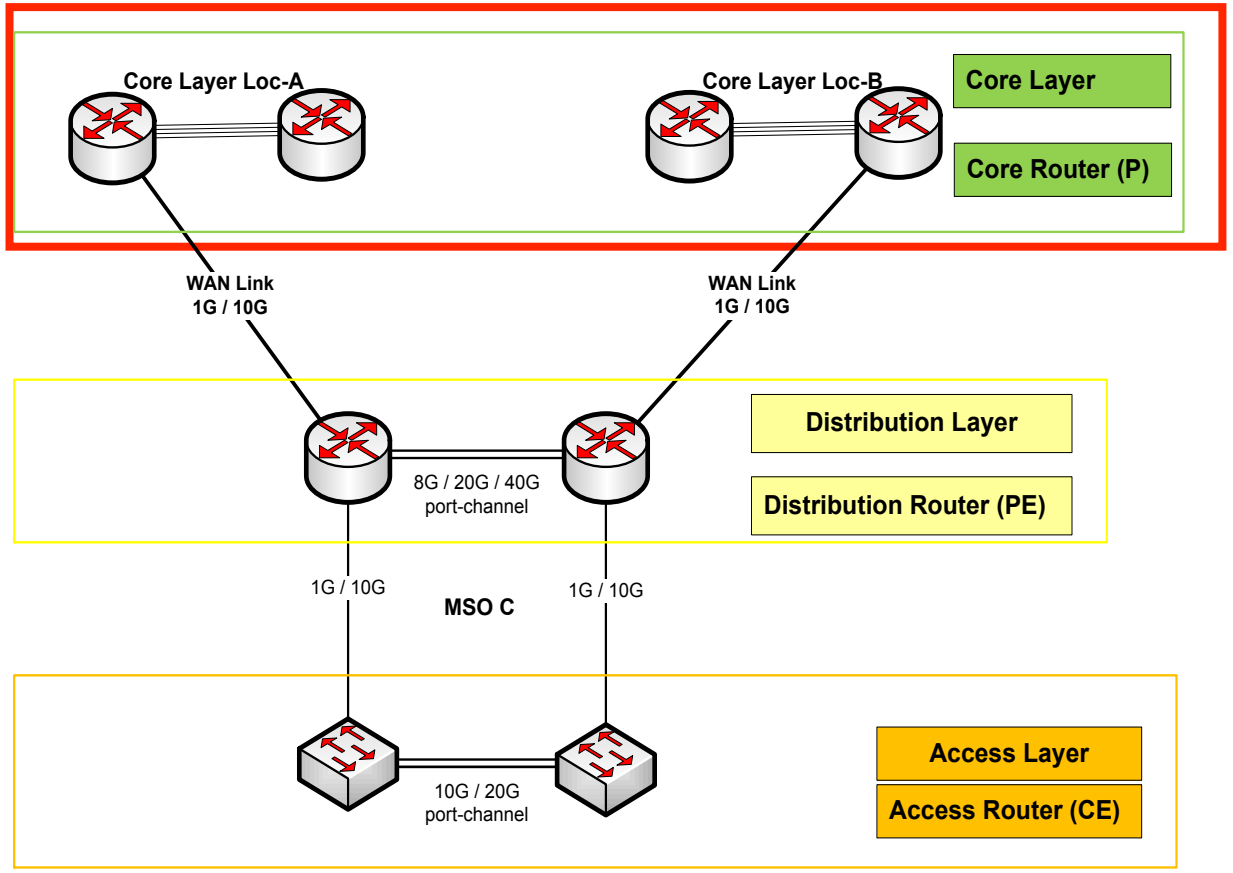

The Distribution Layer

- The *distribution layer* is located *between the Access and Core* layers and helps differentiate the core from the rest of the network.
- The purpose of this layer is to ensure packets are *routed* properly between subnets and VLANs in the enterprise.
- Moreover, this layer provides both *aggregation* and *redundancy*.

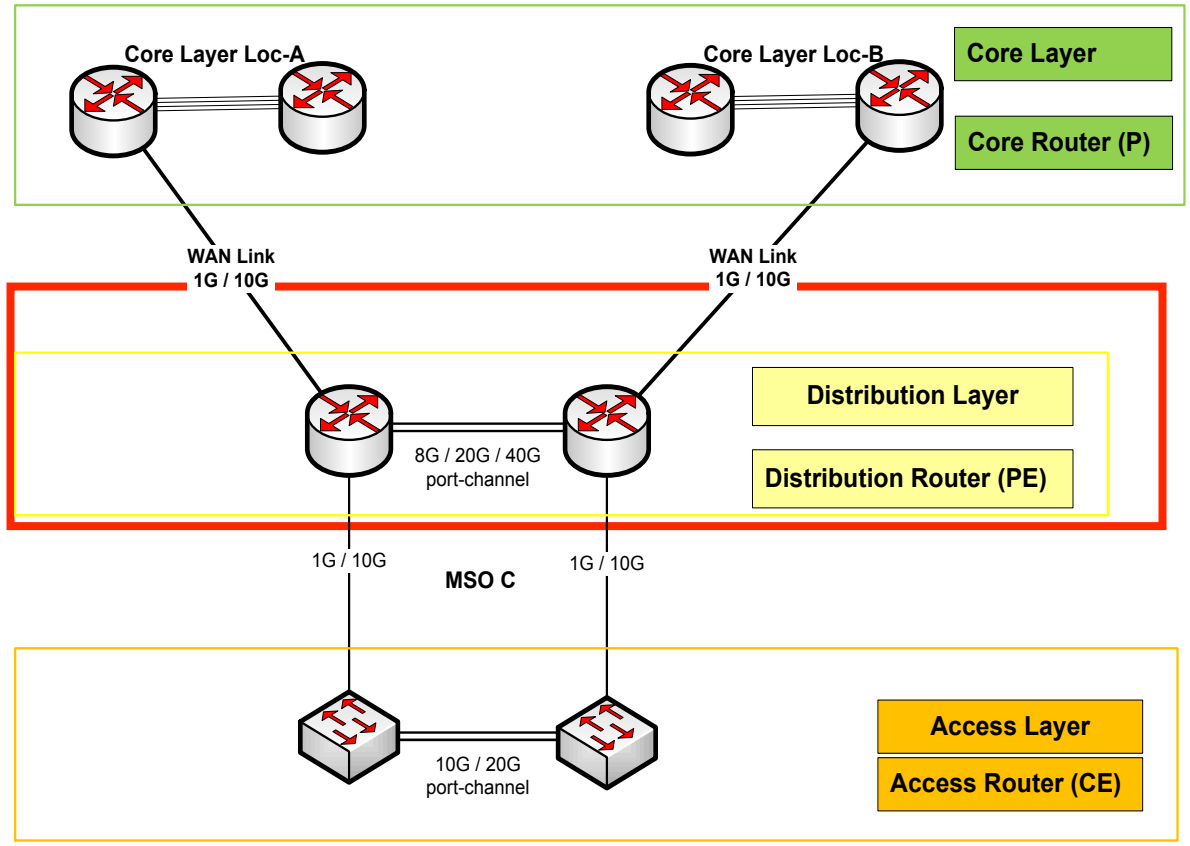

The Access Layer

- The *access layer* supplies traffic to the network and performs *network entry control*.
  - End users/services access network resources by way of the access layer.

- Acting as the *front door* to a network, the access layer employs access lists designed to prevent unauthorized users from gaining entry.

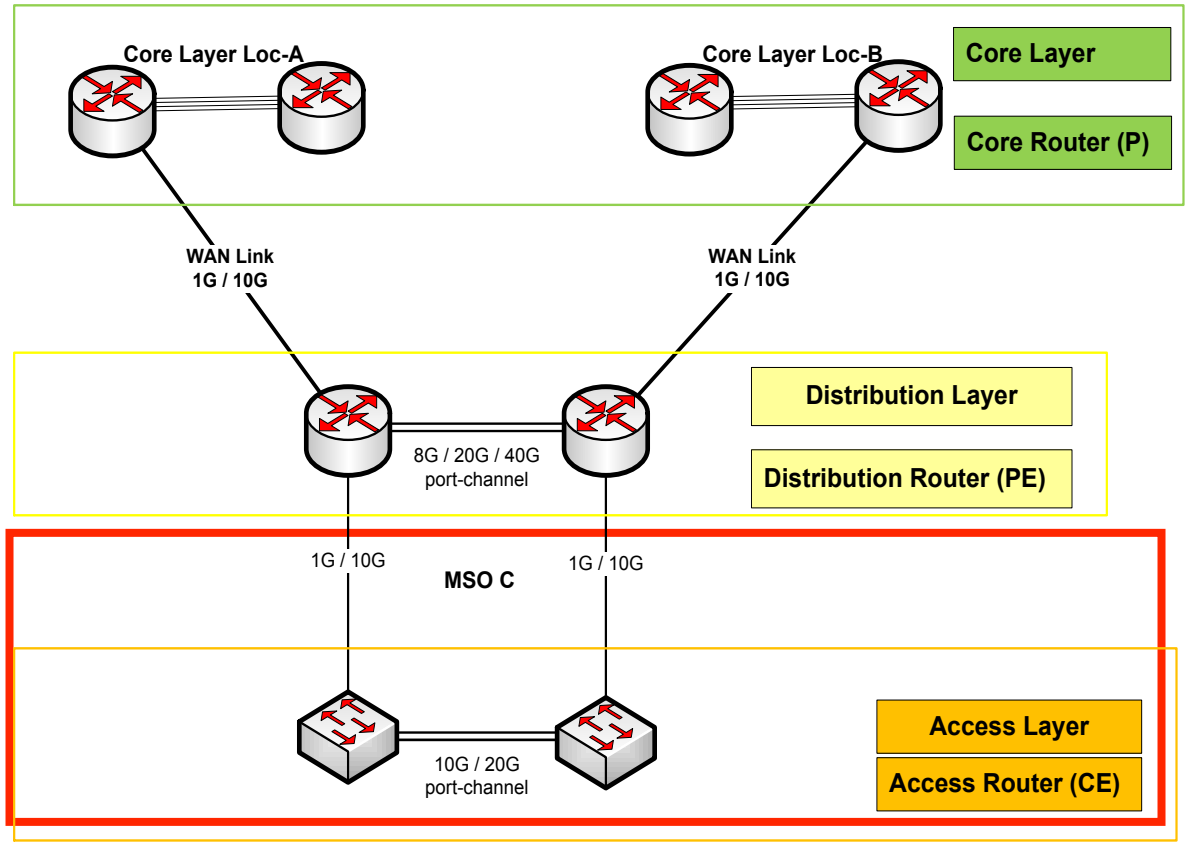

Two-Tier Collapsed Design

- The *three-tier hierarchical design* maximizes performance, network availability, and the ability to scale the network design.

- However, many networks do not grow significantly larger over time. Therefore, a *two-tier hierarchical design* where two layers are collapsed into one layer is often more practical:
  - Option 1: The *core and distribution layers are collapsed into one layer*
  - *Option 2: The Distribution and Access layers are collapsed into one layer*

- The primary motivation for the collapsed core design is *reducing network cost,* while *maintaining* most of the benefits of the three-tier hierarchical model.

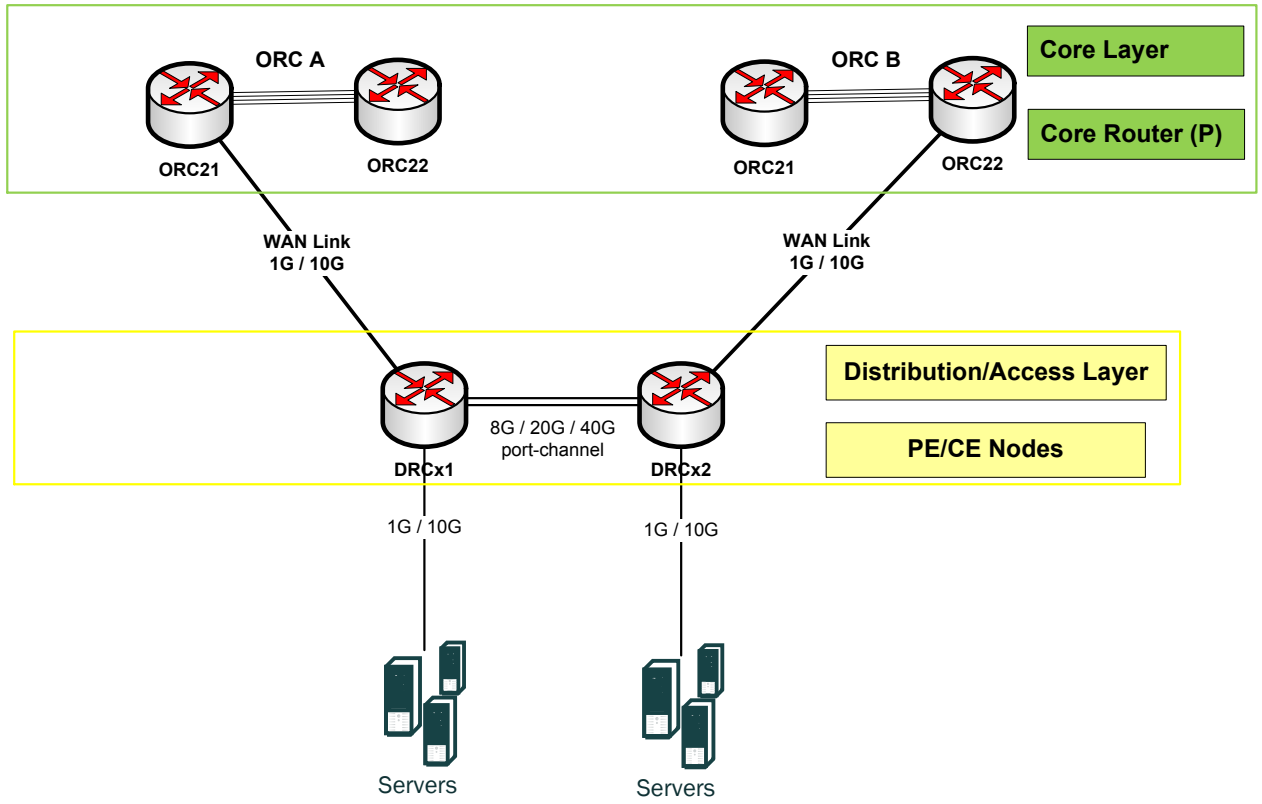

T-Mobile's Network At a Glance

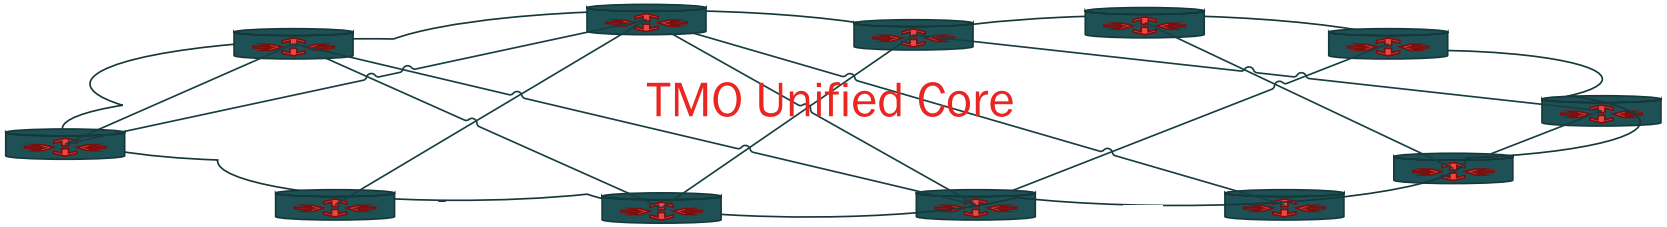

T-Mobile LTE Service

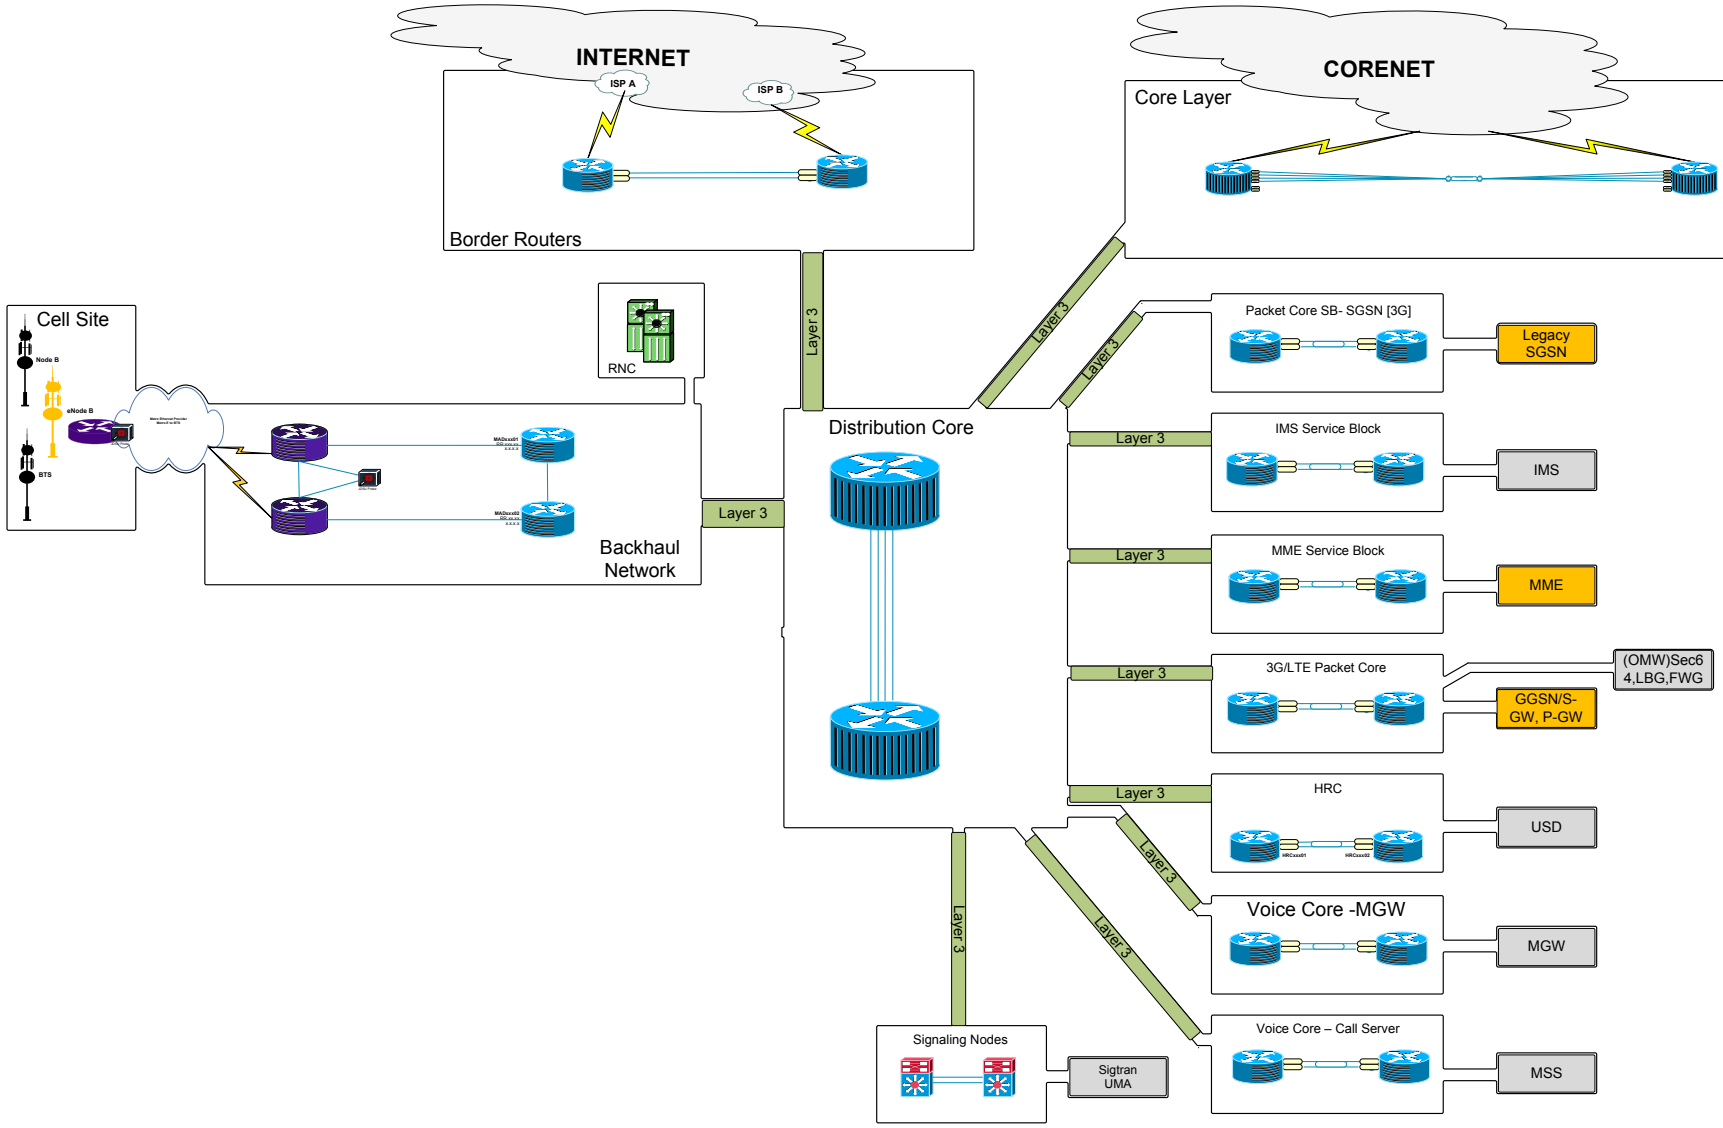

T-Mobile Design Model Evolution –*Architecture* Evolution

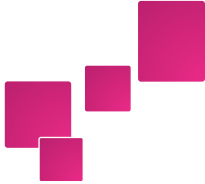

Phase 1: Three different Core Networks –

logically and physically separated

- Customer Core Network, designed for subscribers' traffic
- Management Core Network, designed for Management access to TMO infrastructure nodes
- Enterprise Core Network, designed for Enterprise applications and services along with Extranet/Partner Networks

- Customer Core Network and Enterprise Core Network communicates using a Firewall complex known as EIT-CoreNet Stitch (to comply with IT regulations, e.g., PCI)

T-Mobile Design Model Evolution –*Architecture* Evolution

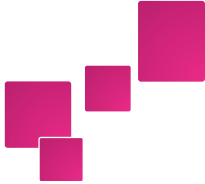

Phase 2: Dual Core Networks

- Once MPLS Layer 3 VPN technology was introduced to T-Mobile network, management traffic was integrated into Customer Core Network, and *Management Core* Network was decommissioned.

- Still, Enterprise Core Network was running due to the need to support ExtraNet/Partner networks.

- Using MPLS Layer 3 VPN to support Management traffic was not trivial at the beginning due to a limitation in existing technology where they were not able to support out of band traffic sourced from a VRF. To overcome this issue “hairpin” workaround was used.

T-Mobile Design Model Evolution –*Architecture* Evolution

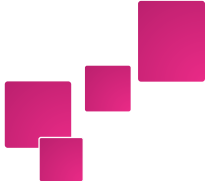

Phase 3: Unified Core Network

- Run Single Core Network that can support all T-Mobile Application, e.g., Customers, Management, and Enterprise.

- This Core was equipped with several technologies (Traffic engineering, Fast Re-Route tunnel, and auto-bw) to guarantee resiliency and the ability to accommodate failures without impacting signaling and real-time applications.

T-Mobile Design Model Evolution – *Technology* Evolution

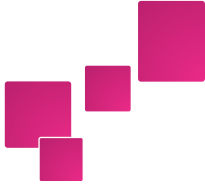

Phase 1: IP Core Based

- Forwarding packets is done using traditional Layer 3 routing protocol (mainly OSPF as IGP protocol).
- In cases, where traffic engineering was needed, Policy based routing was used to forward packets based on their Source/Destination IP addresses, and/or Layer 4 port information.
- Only IPv4 packets were supported during this phase.

T-Mobile Design Model Evolution – *Technology* Evolution

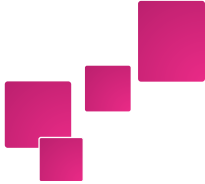

Phase 2: MPLS Core Based with Layer 3 VPN TE/FRR Tunnel

- MPLS technology was implemented in T-Mobile Network with the following applications/services:
  - OSPF routing protocol in Core and Distribution Layer
  - LDP as signaling protocol
  - Layer 3 VPN through using VRF and MP-BGP protocol
  - Dynamic Traffic Engineering in the Core along with auto-backup tunnels to support NHOP and NNHOP Fast Re-route during failures
  - BFD protocol was utilized when possible (routing protocols, RSVP, links, etc.)
  - IPv6 was introduced in T-Mobile network using VPNv6 along with IPv6 VRF. This technology impacted only Distribution and Access layer, however Core Layer was intact

T-Mobile Design Model Evolution – *Technology* Evolution

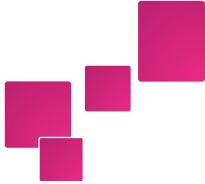

Phase 3: MPLS Core Based with TE/FRR Tunnel and Auto-BW

- Introduce the concept of Auto-BW in T-Mobile Core Network:
  - Auto-Bandwidth: It runs directly on the router; it can respond to changing traffic conditions more rapidly, with less overhead
  - Every statistics interval, bandwidth over an LSP is measured
  - Every adjusted interval, the largest sample from the process above is used to calculate the new LSP bandwidth (e.g., 15 samples and adjust every 3600 seconds)
  - If the change is larger than a user-configured minimum threshold, the new bandwidth value is re-assigned across RSVP.
  - Ideally using a make-before-break configuration

T-Mobile Design Model Evolution –*Routing Protocol* Evolution

Phase 1: OSPF as IGP and BGP as EGP:

- OSPF was the main IGP protocol in T-Mobile Network
  - Multi-areas were first used, however due to some code issues, we moved a lot of our sites to area 0
  - If needed, static routes were also used to accommodate some requirements and they were injected in OSPF process
  - BGP was mainly used in our Border routers to exchange routing updates with ISP

T-Mobile Design Model Evolution –*Routing Protocol* Evolution

Phase 2: OSPF as Core protocol and BGP everywhere else

- A New design was implemented in our Core with the following enhancements:
  - OSPF Backbone area will be used between T-Mobile Core and Distribution layer only
  - All Routing protocols between Distribution Layer and Access layer is based on BGP protocol (eBGP)
  - In Typical cases, Distribution Layer nodes will inject Default route only to Access Layer nodes
  - If Access layer nodes need to receive routes from different T-Mobile Routing contexts (e.g., VRF), Lite VPN will be implemented
  - In some services, access layers can run internal OSPF process to exchange local updates

T-Mobile Design Model Evolution –*BGP RR* Evolution

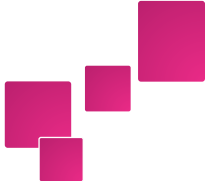

Phase 1: Using BGP Route Reflectors in all Core Locations

- Since BGP RR Servers until recently, will advertise best updates from their perspective and not based on RR clients, T-Mobile decides to implement BGP RR servers in all Core Locations

- Each node in T-Mobile network will have two BGP peers with RR servers in two different markets

T-Mobile Design Model Evolution –*BGP RR* Evolution

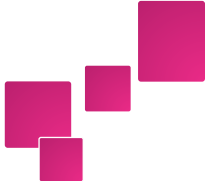

Phase 2: Using BGP Route Reflectors in Four Core Locations only

- Introduced the concept of *Add Path* feature to BGP process where RR Servers can advertise not only the best update but up to all updates to BGP RR Clients

- As a result, T-Mobile decided to reduce number of BGP RR servers to four Core locations only

Lessons Learned and Recommendation

- Design standards must be enforced everywhere; deviations from standards (due to some corner cases and application needs) must be documented and approved before deployment.
- Use Bundles (aka Port channels) from day 1: it is much easier to add more links to an existing bundle than to transition a production link to a bundle.
- Validate your design by running failover test in services and infrastructure nodes. This test should be part of any service/node acceptance in your network.
- Evaluate WAN circuits you are using to connect Core Nodes and/or Core/ Distribution nodes and make sure they meet diversity requirements.

Lessons Learned and Recommendation

- When Bundles are used (aka port channels) use minimum links feature properly to decide when Bundle will go down (e.g., min-links = # of links, or min-links = # of links –x)
  - For LAN links with N+1 redundancy, Set minimum links to # of links
  - For WAN links, set min links based on capacity and diversity level (avoid site isolation)

- Architecture/Design reference documents must include proper data flows and operation support sections.

- Your design should avoid using propriety features (e.g., vPC) to "…keep your options open…"

Lessons Learned and Recommendation

- When choosing what vendor/platform to use and where, give this process a lot of thoughts, examples are:
  - Understands your need: Technology should advance the business
  - Engage right people: e.g., Engineers, finance, legal, security
  - Establish your requirements: Architecture, maturity, cost, serviceability, reliability of the product, support, and company stability
  - Focus on the total cost of ownership, not just initial cost
  - Keep your options open: Negotiate, deal with multiple vendors simultaneously
  - Be a good customer. The win/win approach is at the heart of all good vendor agreements. "Both sides must benefit from the relationship”
  - Anticipate the future
  - Don't let technology turn your head

Q & A