The Economic Benefits of Distributed Broadband Services

Peter Fetterolf, Ph.D.

EXECUTIVE SUMMARY

There are several key market trends that are driving the need for new distributed, disaggregated, and cloud-based broadband network architectures. Traffic is growing rapidly because of 4K/8K video, cloud gaming, work at home, and other new applications. High-bandwidth, low-latency applications and services are driving the requirement for service intelligence in the metro access network. New CUPS architectures are disaggregating the user plane and the control plane in both mobile and fixed broadband networks, and we expect to see increasing convergence of mobile and fixed IP networks.

Although legacy centralized BNG architectures have served many service providers well for decades, there are a few challenges with this legacy approach in existing networks:

- BNG ports provide rich features combining aggregation and service functions and are relatively more expensive
- BNG expenses rise with increasing port density requirements
- Complex feature sets combining BNG features and aggregation functions in legacy BNGs can require considerably more coordination and planning for upgrades

Traditionally, BNGs use 1:1 redundancy for high-availability implementations. A new approach to fixed broadband BNG service uses a spine-leaf design similar to today's cloud data center approaches. The spine-leaf design consists of access leafs and service leafs that disaggregate access node functional requirements from BNG service functions. This architecture supports distributing edge nodes or service leafs to the access network providing BNG services for high-bandwidth and/or low-latency intensive applications such as video caching, IoT gateways, AR/VR servers. Likewise, this spine-leaf design can be implemented in centralized or regional data centers with similar advantages.

This paper compares the legacy centralized BNG architecture with the spine-leaf design including:

- Identification of some challenges with legacy centralized BNG architectures
- Details of the new spine-leaf design
- Assumptions and results of the total cost of ownership (TCO) model

In our analysis we compare the legacy centralized BNG architecture with a distributed spine-leaf BNG architecture and show a five-year TCO savings of 35%. We also present scenarios where both the legacy BNG and the spine-leaf BNG designs are deployed in central office data centers to compare their TCO. These approaches produced the most significant five-year cumulative TCO savings of 64%.

Drivers for Distributed Broadband Services

Multiple trends in network services are driving the need for distributed broadband services. Over the past five years we have seen significant growth in broadband traffic, and we expect this growth, which is driven by current and emerging broadband applications, to continue for the next five years:

- Video streaming
- 4K TVs and future 8K TVs
- Cloud gaming using HD & 4K video
- Hybrid work using video conferencing (Zoom, Teams, WebEx, etc.)
- Social networks with video and images

ACG Research projects residential average broadband traffic to increase from an average busy hour data rate of 13Mbps in 2021 to 21Mbps in 2025 with an annual growth rate of 12% . This level of traffic is driving the requirement for higher network capacity and BNG scalability in the metro access network.

One of the key technology trends that we expect will accelerate is edge computing in the metro access network. There are several key reasons for Communication Service Providers (CSPs) to deploy edge computing:

- Reduce traffic in core and aggregation networks
- Provide services to latency-sensitive applications
- Allow for massive IoT scale providing IoT gateways
- Deliver security to the metro access network

Some of the applications driving edge computing are:

- Video caching is moving increasingly to the metro access network to reduce traffic in the aggregation and core networks as video bandwidth increases
- Cloud gaming has both high-bandwidth and low-latency requirements that necessitate edge computing
- IoT applications require edge gateways for scalability

- Many Industry 4.0 applications require edge computing, for example, AI/ML driven manufacturing or video monitoring in large warehouses
- Emerging AR/VR applications will drive edge computing
- New edge applications emerging over the next five years

In conjunction with traffic and applications, network technology transitions are also driving the requirement for distributed broadband services. Network service control planes are being separated from forwarding planes (CUPS) in both mobile and fixed broadband networks. Separation of the forwarding plane from the control plane allows operators to place user plane functions closer to the metro access network to support the emergence of new services and applications. CUPS also enables operational efficiencies where a single control plane manages many downstream forwarding planes. In addition to CUPS, forwarding planes for mobile services (UPF) and fixed broadband services (BNG) are starting to converge using controllers to separate mobile and fixed service forwarding functions. These technology trends are driving some CSPs to distribute broadband services to the metro access network.

Challenges with Legacy Centralized BNG Architecture

BNG nodes provide network services and control for fixed broadband networks. Legacy BNG network architectures implement BNG service nodes at the service Points of Presence (POPs) in centralized data centers. The BNG aggregates traffic from the access network and provides BNG services that allow fixed broadband connections access to the internet (Figure 1).

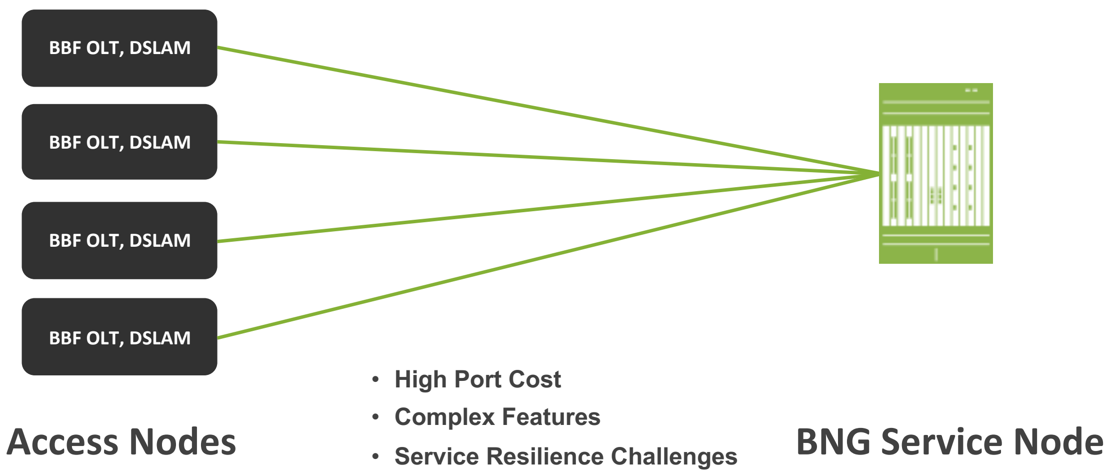

Figure 1. Current BNG Architecture

There can be several challenges with this traditional BNG architectural approach. Aggregation of broadband access traffic on dense BNG service nodes can be expensive. BNG nodes supporting both aggregation and service functions concurrently on the same port have high port density requirements with higher per port costs. Delivering both access aggregation and BNG services on the same node can be expensive. Complex feature sets impact network operating expenses. Changing both aggregation and service features can create additional coordination, testing, and planning to provision changes in a timely manner and with minimal service impact. Feature growth also has an impact on device life-cycle management, posing the risk of premature system obsolescence. Legacy BNG architectures pose challenges to service resilience because service attributes are tied to physical ports. Moving services requires similar ports to be available. This creates challenges managing IP VPNs and moving services between ports. One approach to solving these challenges is to disaggregate BNG service node functions from access port functions in a distributed broadband services architecture.

Distributed Broadband Services

Disaggregating access nodes and BNG service nodes by leveraging a spine leaf design addresses several challenges posed by the legacy centralized BNG architecture. The spine-leaf design is commonly used in data centers to efficiently scale and increase availability of cloud services. An example of this architecture is Juniper's Distributed Broadband Access Solution (Figure 2).

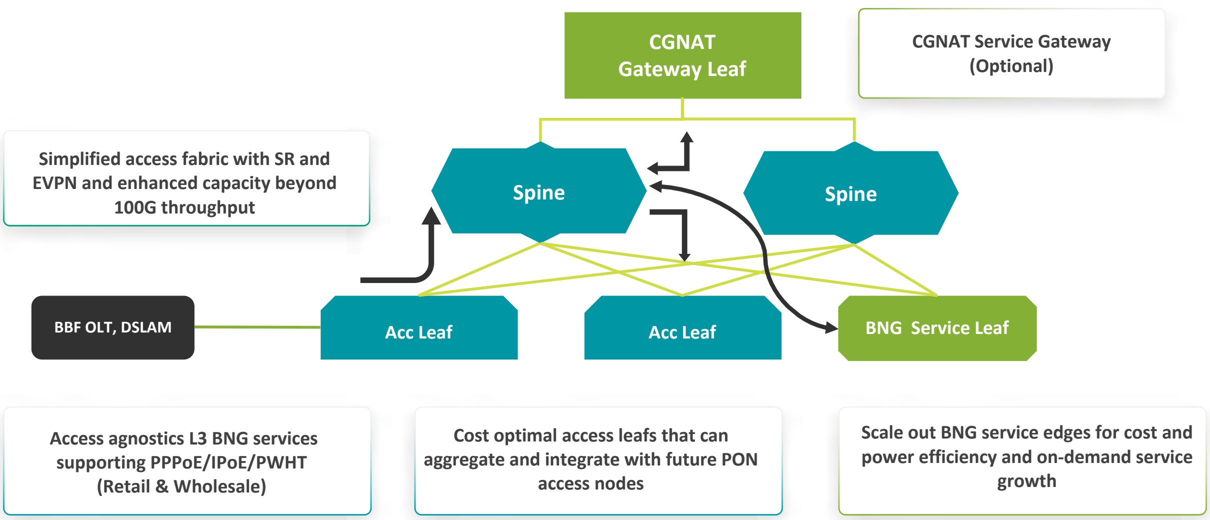

Figure 2. Distributed Broadband Access Solution

Some of the key characteristics of this architecture are:

- Distributed BNG service leafs can be distributed to the metro access network
- Data center-like spine-leaf design uses a highly scalable IP service fabric
- Access ports are separated from BNG service ports, enabling a horizontally disaggregated architecture

The spine-leaf design uses access leafs to connect to access nodes (OLT, DSLAM). The access leafs are interconnected with BNG service leafs using spine switches with MPLS fabrics. Access leafs provide cost-effective, scalable ports, which can often leverage merchant silicon. The key characteristics of access leafs are:

- High-access port density
- Access feature set simplified and are relatively less expensive

Spine switches provide interconnectivity between access and BNG service leafs. The key characteristics of the spine are:

- EVPN-VPWS with segment routing provides connectivity to service leafs
- EVPN resilience allows for N+1 service resilience
- The service interface is abstracted, enabling both resilience and placement flexibility

The BNG service leaf uses 100GE ports to aggregate multiple access ports in a single BNG service port. This provides cost-effective BNG service ports that scale based on network traffic, not the number of access ports required. The BNG service leafs also allow for large session scale up for up to 96,000 sessions per BNG service leaf. The architecture enables incremental and cost-effective scale-out. Access ports and leafs can be added as necessary and are driven by port density requirements. Spine-switch fabrics can be added as necessary to scale out the system, and service leafs can be added based on network traffic and session demand. Disaggregating BNG service ports from access ports enables for cost-effective scaling of BNG services.

TCO Model Framework and Assumptions

ACG Research has developed a total cost of ownership (TCO) model comparing the present mode of operation (PMO) with the future mode of operation (FMO). The PMO architecture uses legacy BNG service nodes as depicted in Figure 1. The FMO architecture uses a disaggregated spine-leaf architecture as represented in Figure 2. We analyze two types of scenarios:

1. A centralized PMO BNG versus a distributed FMO spine-leaf BNG
2. PMO BNG versus FMO spine-leaf BNG deployed in all central offices

The first comparison analyzes the benefits of distributed versus centralized BNG architectures. The PMO solution is depicted in Figure 3. In this architecture access nodes are located in access central offices, and the centralized BNG is located in an aggregation central office. Access nodes are interconnected to the aggregation node using the DWDM optical transport network. Because the BNG function is centralized, redundant BNGs are used for a high-availability architecture.

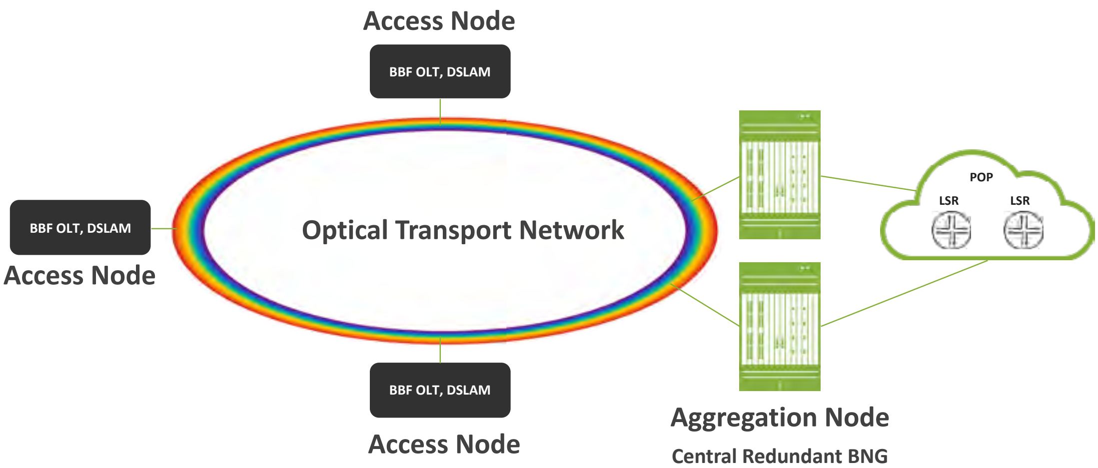

Figure 3. Centralized PMO BNG Architecture

The FMO solution in the first comparison is depicted in Figure 4. This architecture uses spine-leaf BNGs at distributed locations. The spine-leaf architecture avoids single points of failure so 1:1 redundancy is not required.

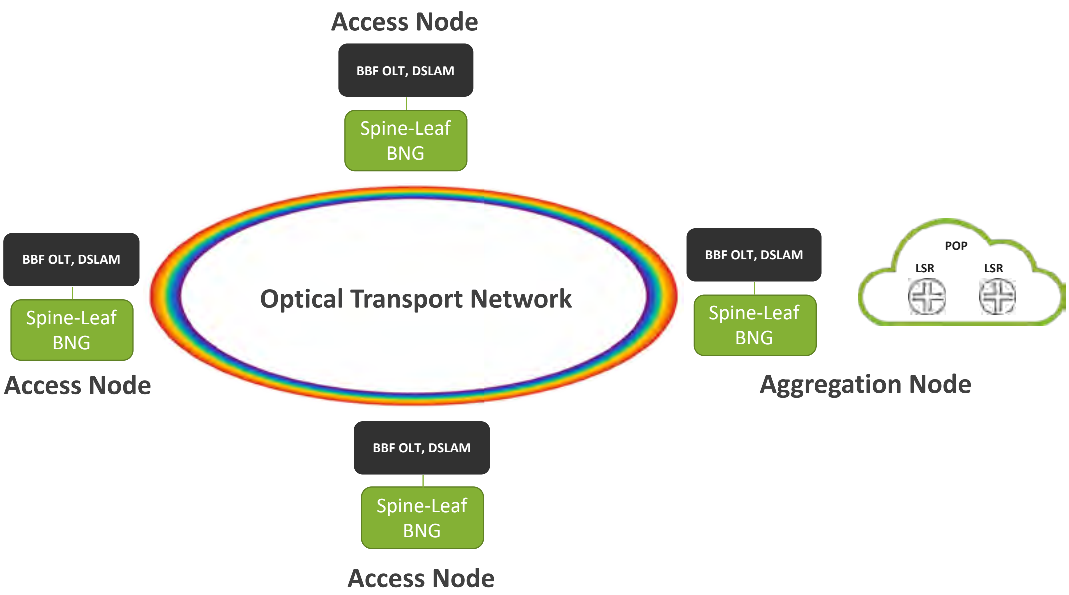

Figure 4. Distributed FMO BNG Architecture

The second comparison analyzes the cost difference between the PMO BNG and the FMO spine-leaf BNG in the case where BNGs are deployed in all central offices. This comparison is depicted in Figure 5 and provides a TCO comparison of the BNG PMO and FMO solutions in each central office.

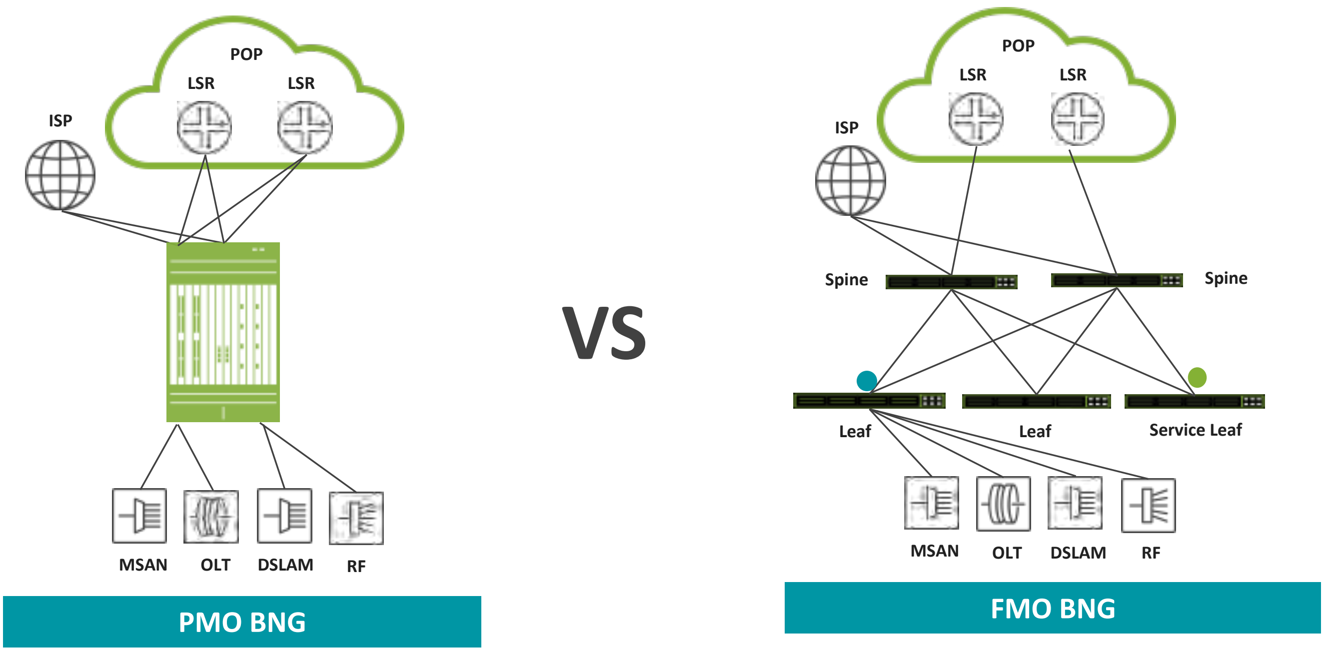

Figure 5. Comparison of PMO BNG and FMO BNG in All Central Offices

For each of the scenario comparisons we assume a medium-size network with:

- 10 large central offices
- 100 medium-size central offices
- 1000 small central offices

The number of residential and business subscribers and average traffic per subscriber is presented in Table 1. These assumptions are used to configure the PMO and FMO BNG solutions for both scenario comparisons.

Network Demand	Year 1	Year 2	Year 3	Year 4	Year 5
Residential subscribers per large CO	18000	23400	30420	39546	51410
Residential subscribers per medium CO	10000	13000	16900	21970	28561
Residential subscribers per small CO	1500	1950	2535	3296	4284
Residential subscriber traffic (Mbps)	13	15	17	19	21
Business subscribers per large CO	4000	5200	6760	8788	11424
Business subscribers per medium CO	2000	2600	3380	4394	5712
Business subscribers per small CO	300	390	507	659	857
Business service traffic (Mbps)	5	7	10	12	15

Table 1. Residential and Business Fixed Broadband Demand

The TCO model considers both capital expense (CapEx) and operation expense (OpEx) associated with the BNG solutions. The FMO CapEx includes:

- Leaf nodes
- Spine nodes
- BNG service nodes
- Software licenses
- Pluggable optics

The PMO CapEx includes:

- BNG chassis
- BNG line cards
- Software licenses
- Pluggable optics

The OpEx assumptions for both the PMO and FMO include:

- Equipment technical support services
- Power and cooling
- Floorspace

TCO Results

Our comparisons use a five-year TCO model. In the first scenario we compare a centralized PMO BNG solution with a distributed FMO BNG solution as represented in Figure 3 and Figure 4. The five-year cumulative TCO results are presented in Table 2. The total TCO savings over five years is 35%.

	CapEx	OpEx	TCO	TCO Savings
FMO	$ 129,874,500	$ 90,106,767	$ 219,981,267	35%
PMO	$ 235,542,000	$ 101,452,518	$ 336,994,518	N/A

Table 2. Five-Year Cumulative TCO Comparison of Centralized PMO BNG versus Distributed FMO BNG

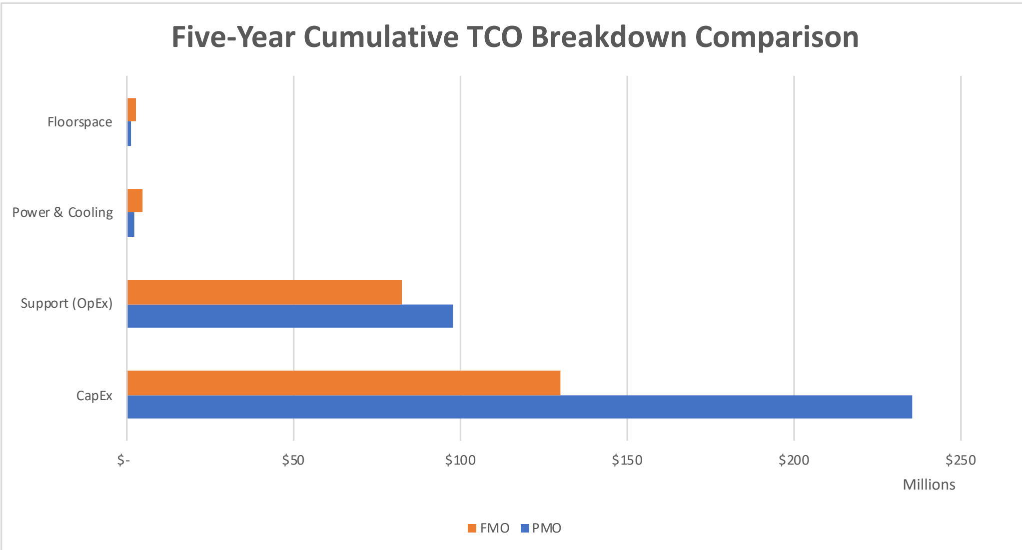

Figure 6. Five-Year Cumulative Expense Breakdown of Centralized PMO BNG and Distributed FMO BNG

The next TCO compares PMO BNG and FMO BNG configurations both in central offices (small, medium, and large), Figure 5. In this comparison we do not consider centralized versus distributed BNG architectures. Instead, it is a direct analysis of the PMO and FMO BNG architectures in a central office scenario. The five-year cumulative TCO is presented in Table 3, and the five-year cumulative TCO breakdown is presented in Figure 7. The FMO has a 64% TCO savings over five years. This is primarily a result of the spine-leaf architecture cost-efficiency and the benefit of disaggregating access ports from BNG service ports.

	CapEx	OpEx	TCO	TCO Savings
FMO	$ 129,874,500	$ 90,106,767	$ 219,981,267	64%
PMO	$ 409,049,000	$ 203,339,465	$ 612,388,465	N/A

Table 3. Five-Year Cumulative TCO Comparison of PMO and FMO Architectures in All Central Offices

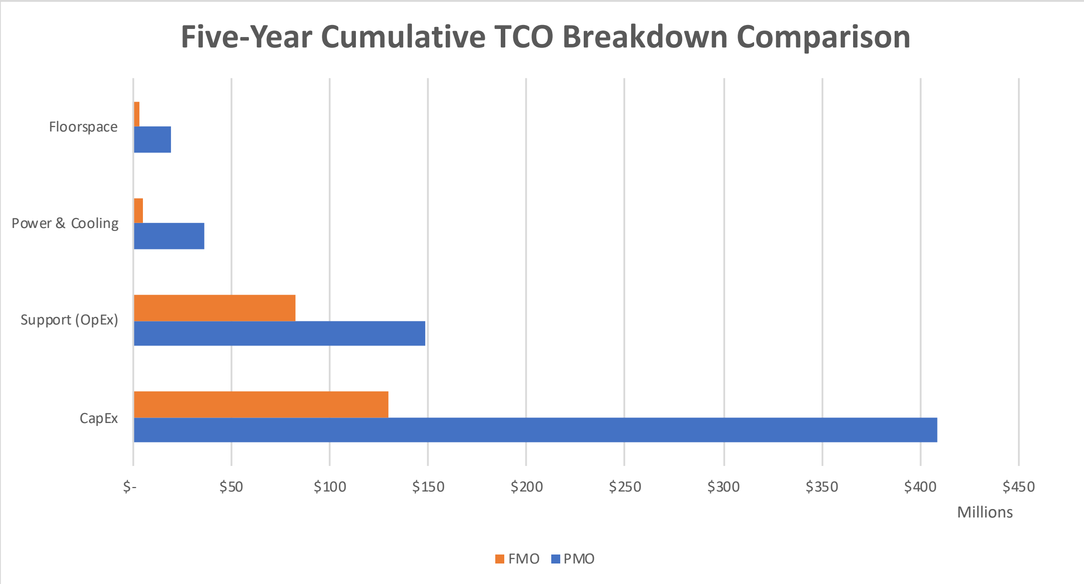

Figure 7. Five-Year TCO Breakdown of FMO and PMO BNG Architectures in All Central Offices

Conclusion

The analysis and results demonstrate the need for a new distributed, disaggregated, and cloud-based BNG architecture. As network traffic growth and edge services drive the requirement to deploy BNG service nodes closer to the subscriber, it is clear a new architecture is advantageous to allow for cost-effective scale-out of BNG service nodes. The new distributed BNG architecture uses a spine-leaf architecture that horizontally disaggregates access ports from service ports, improving cost efficiencies in scaling to meet increased port density and traffic growth requirements while incrementally scaling BNG service ports to optimize BNG service delivery. Our analysis shows that a distributed BNG architecture using the new approach has a 35% TCO savings over a legacy centralized BNG architecture. When directly comparing the PMO BNG with the FMO BNG in strictly central office scenarios, we show that this new spine-leaf design has a 64% TCO savings over the traditional PMO approach.