ACG RESEARCH

An Economic Comparison of a DriveNets Network with Traditional Router Networks

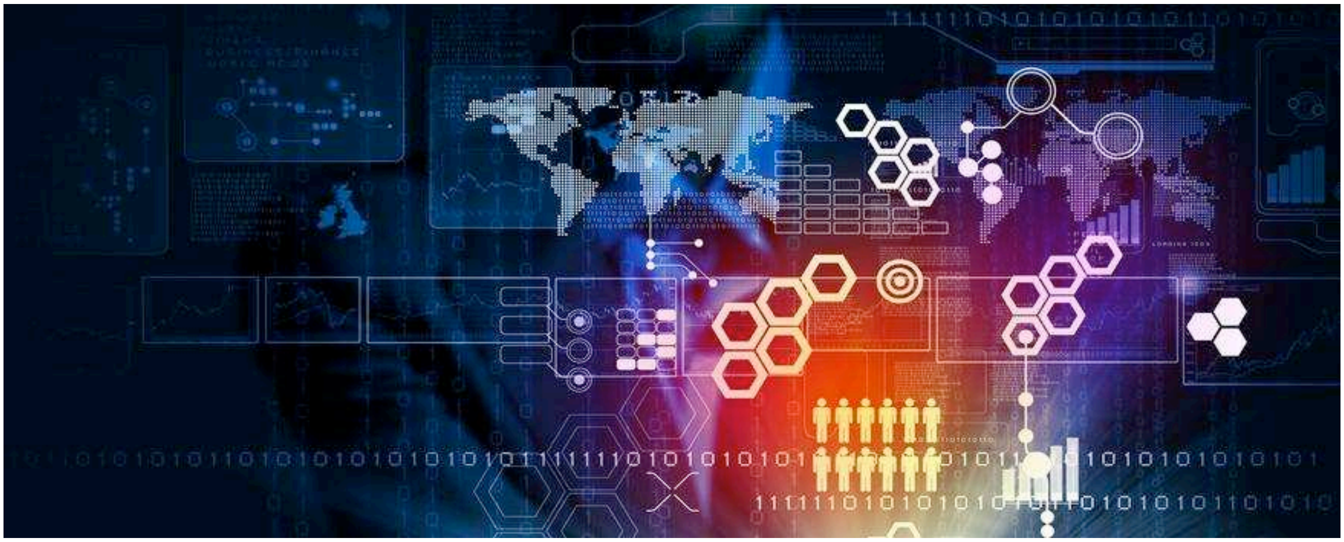

EXECUTIVE SUMMARY

IP networks are experiencing tremendous traffic growth that is expected to continue for the foreseeable future. Growth will primarily be driven by 4K/8K video, gaming services, and 5G mobile networks. In order to maintain or improve profit margins, service providers need a new approach to IP networking that will cost-effectively scale to meet future traffic growth requirements. DriveNets has pioneered a new software-based architecture for IP networks that is running on white-box building blocks, separates the control-plane and data-plane, and offers an innovative licensing model. The DriveNets architecture is based on the approach taken by hyperscalers, disaggregating the traditional router model. It allows network routing nodes to scale from a single white-box of 4 Tb/s up to an additional 192 white-boxes interconnected by a fabric made of a second type of white-boxes.

ACG Research developed a total cost of ownership (TCO) model to compare the distributed DriveNets architecture with the incumbent IP routing architecture. We assume that the incumbent router network uses the newest service provider routers available on the market today with the highest port densities and the lowest price per port. This ensures a fair comparison of the DriveNets architecture with current networking solutions in the market. We found the DriveNets architecture to have a 51% TCO advantage over the current architecture with a 61% savings in router hardware and software expenses and a 33% savings in operations expense (OPEX). The key drivers of these savings are that 1) the DriveNets white-box architecture dramatically minimizes hardware expenses, 2) the DriveNets attractive software pricing lowers the cost per bit, and 3) DriveNets has lower OPEX due to a common set of routing node building blocks.

KEY FINDINGS

- The DriveNets routing architecture for mid-size networks reduces five-year TCO by 51%
- For medium to large CSPs, DriveNets reduces five-year routing expenses by 61%
- DriveNets reduces five-year OPEX by 33%
- TCO savings are driven by:
  - Attractive software pricing
  - Lower OPEX due to a common set of routing node building blocks
  - White-box scalable hardware

NETWORK ARCHITECTURE

Over the last decade global IP traffic has grown at a high rate, and it is projected to continue to increase at a compound annual growth rate of 26% until 2022. Traffic growth will primarily be driven by 4K/8K video, gaming services, and 5G mobile networks. It is essential that service providers build networks that will scale to support these new services. Service providers' revenue per bit is declining; therefore, future IP networks must operate at a much lower cost per bit than existing incumbent networks. As networks grow there is also a demand for simplicity in network design and operations. Current networks are extremely complex and difficult to manage, and the architecture has not changed in decades. Future IP networks should take a new look and learn from the evolution of the data center into cloud. They must become more scalable, flexible, and cost effective to support the future growth of IP traffic.

Incumbent Network Architecture

Current IP networks have evolved using different families of routers in different layers of the network. Networks are usually designed with access, aggregation, edge, and core routers. The functionality and scalability requirements at each layer of the network are different; therefore, router hardware and software have been designed to meet the unique requirements of each network layer. Over time this has led to increasingly complex networks with many different router families, vendors, software releases, inventory parts, and complex and varied operations procedures. For example, different routers require separate test and certification procedures, distinct security patches, different network engineering approaches, and separate software upgrades. The incumbent architecture and technical challenges are illustrated in Figure 1.

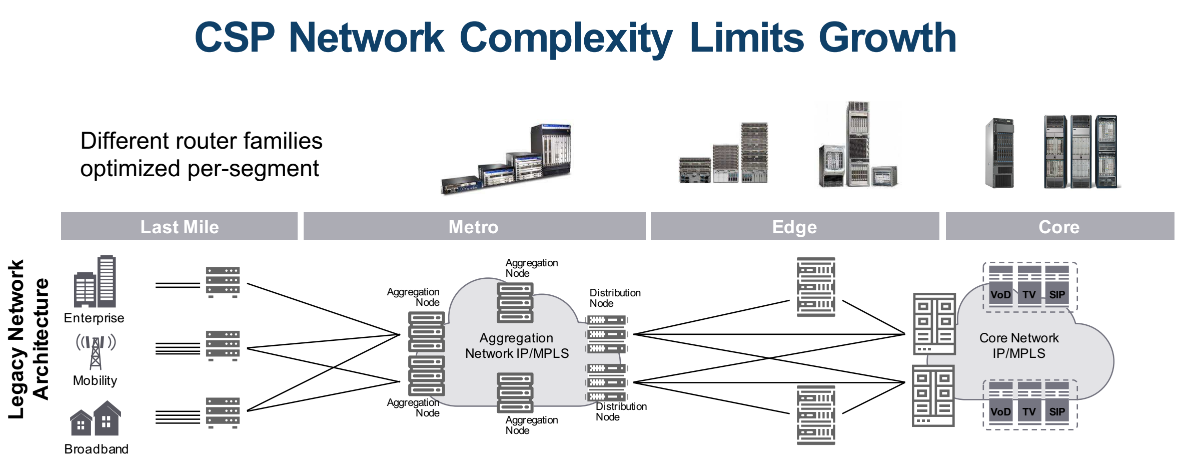

Figure 1. Incumbent IP Network Architecture

As a result of these factors the total cost of ownership (TCO) of IP networks has been increasing as traffic and networks continue to grow. Also, in core networks, it is possible to exhaust the ports in even the largest routers, which requires service providers to configure multiple core routers in a CLOS architecture to increase network scalability. The economic impacts of the CLOS architecture will be examined in another section in this paper.

DRIVENETS NETWORK ARCHITECTURE

DriveNets has developed a new approach to networking, Network Cloud, that is not based on traditional networking routers and addresses many of the problems with incumbent IP networks. The DriveNets architecture is depicted in Figure 2. The data-plane is based on two white-box building blocks: 1) the Network Cloud Packet-processor (NCP) forwarder and 2) the Network Cloud Fabric (NCF) fabric engine. The white-box architecture minimizes hardware expenses and allows service providers to ride the merchant silicon curve. The control-plane runs on a server that can be either local or remote and run on bare metal, a virtual machine or a container. Both the control-plane and data-plane can scale from small sizes to arbitrarily large sizes. Also, the control-plane is disaggregated from the data-plane, which allows the control-plane and the data-plane to scale independently of one another. A small router can be configured in a single-box configuration, where the control-plane runs on the NCP and forms a single box router at 4 Tbps capacity. As the network grows, additional NCPs can be added and interconnected with one or more NCF fabrics. The size of the router can grow up to 192 NCPs. The TCO analysis in this paper shows that the DriveNets architecture is very cost effective as networks scale to large sizes.

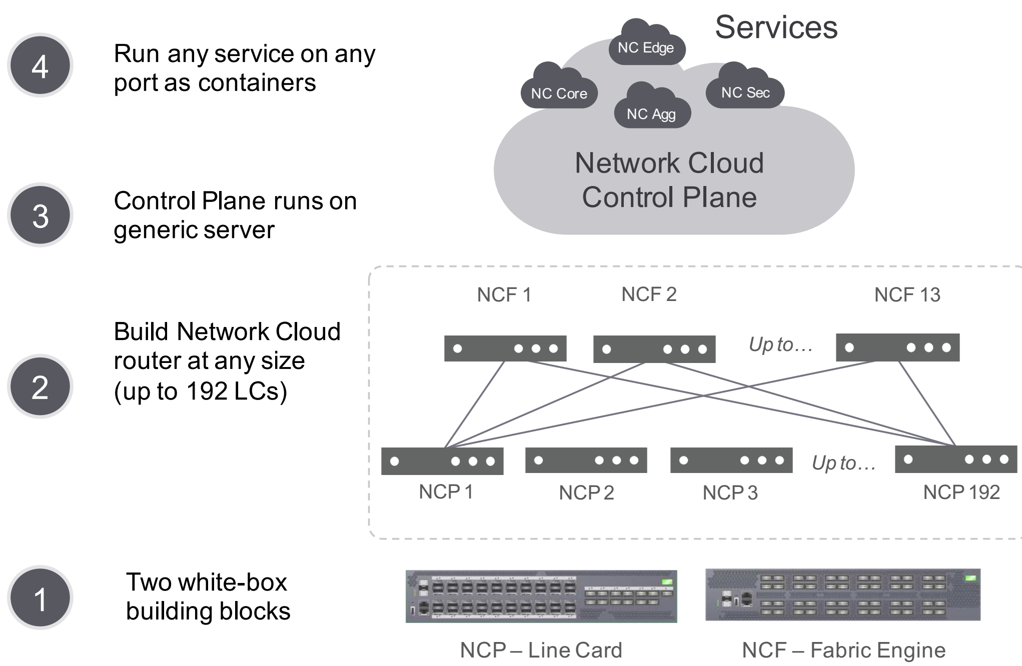

Figure 2. DriveNets Network Architecture

All router software is cloud native running in containers on both the control-plane and data-plane. Different control software is used for different functions, for example, edge routing and core routing. The container-based architecture allows a true converged network solution where a single unified data-plane is shared by different networks and services with separate control-planes. For example, mobile, broadband, and enterprise traffic networks that are traditionally separated over different infrastructure can share the same data-plane physical infrastructure but have separate control-planes. This way, the networks remain completely separate, allowing them to have different service level agreements, distinct security policy, and even a different maintenance team with limited access to specific service.

The ability to support this highly efficient, fully shared infrastructure while maintaining completely separate networks is enabled by the DriveNets software-based approach of Network Cloud, where networks are separated by different control-plane containers while sharing one data-plane. This software-based model also allows for any service on any port.

The key benefits of the DriveNets architecture are:

- Uniform white-box hardware components used in aggregation, edge, and core networks.
- Network services and protocols are cloud native running in containers.
- Cloud routers can scale from single NCP boxes to large interconnected routers using up to 192 NCPs connected with NCF fabrics, allowing for any size router from only two types of building blocks.
- Complete network convergence solution that covers both data-plane and control-plane requirements.
- Multiple separate networks can run over a single converged physical infrastructure in different containers.
- Ability to add new services and support service chaining by adding services in separate containers to the control-plane.

TCO MODEL FRAMEWORK

The TCO model compares the TCO of a traditional IP network with the DriveNets Network Cloud. The model uses the most recent versions of current routers with the highest available port density and lowest cost per port to make a fair architectural comparison. The comparison assumes that a converged IP network is used for mobile, residential, and business traffic because Network Cloud architecture enables full convergence where the data-plane is shared, and the control-plane is separated. A large nationwide network is broken up into regions where each region has aggregation, edge, and core nodes (Figure 3). Small, medium, and large regions are modeled for a typical Tier 1 service provider network. Aggregation nodes aggregate mobile, residential, and business traffic. Edge and core nodes are collocated in the same central office.

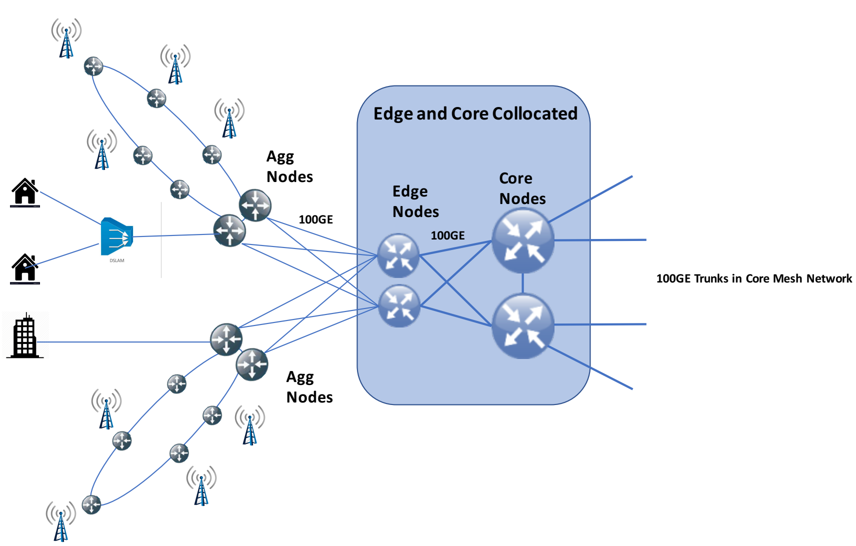

Figure 3. Converged IP Network

The DriveNets architecture allows edge and core IP services to be combined in a single router (Figure 4). Because the Network Cloud can converge different types of services to a single shared infrastructure separated by software containers, it can converge the edge and core services. The control-plane is separated from the data-plane. The physical routing elements (NCPs and NCFs) provide port density and traffic forwarding while the edge and core IP services are implemented in the software-based control-plane. Such converges of networks over a single physical infrastructure are different from the incumbent architecture that has traditionally separated the core and edge with separate routers providing different sets of services.

Note that in COs or data-centers where there are aggregation routers that are co-located with the edge/core routers, there is no need to build a separate Network Cloud cluster for the aggregation function and another one for the edge/core. One cluster can be used as a shared data-plane for both networks, supporting all networks and network services. To simplify the model, we assumed that the aggregation router is separated in all the locations. Additional TCO benefits could be achieved if the full benefits of the Network Cloud model were applied to aggregation, edge, and core routers.

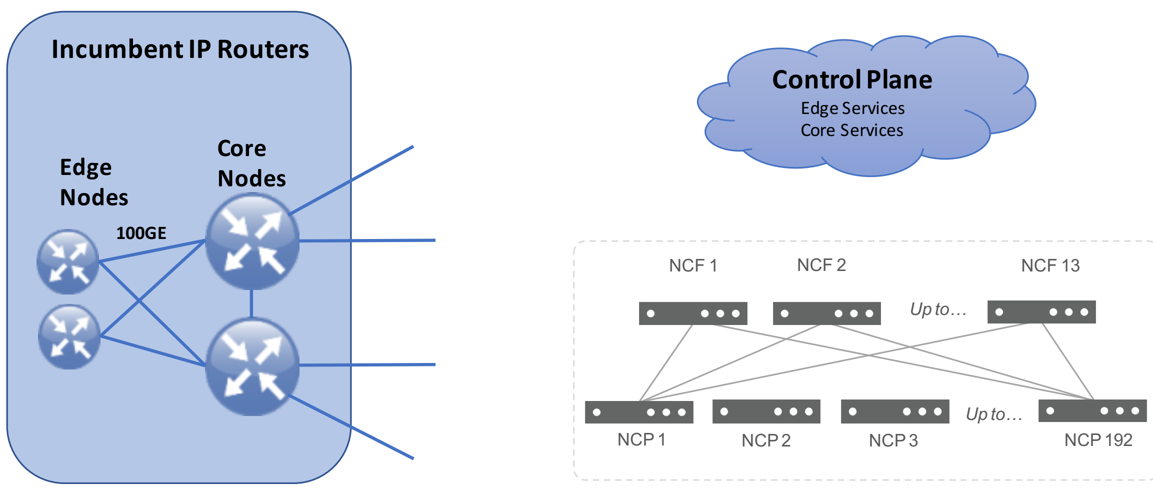

Separate Edge & Core Routers **Converged Core & Edge Routers**

Figure 4. DriveNets Combines Edge and Core Services

The core nodes provide a national mesh network interconnecting regional networks. Core nodes aggregate traffic from the region and carry transit traffic traversing the mesh. In typical mesh networks the transit traffic can be five times as large as aggregation traffic, which means that many more ports are required. If the number of ports required on the core router is exceeded then it is necessary to create a larger routing node using multiple core routers in a CLOS switching architecture as illustrated in Figure 5. The CLOS architecture consists of leaf nodes and spine nodes. Each leaf and spine uses a full featured core router. For example, if 1000 100 GE ports need to be interconnected at a core node, then an incumbent core CLOS architecture would require four leaf nodes and four spine nodes. This architecture is highly inefficient, expensive, and consumes large amounts of power. This is because the core routers are not designed for a CLOS architecture, and many service ports are required for interconnection between the leaf and spine routers.

The DriveNets architecture is designed to support very large numbers of ports using the NCF fabrics. It is much more cost effective and power efficient because it does not use service ports to connect additional units in the CLOS.

The TCO model considers three types of regions: large, medium, and small. In the large regions, the high levels of traffic drive the CLOS architecture for the incumbent core routers.

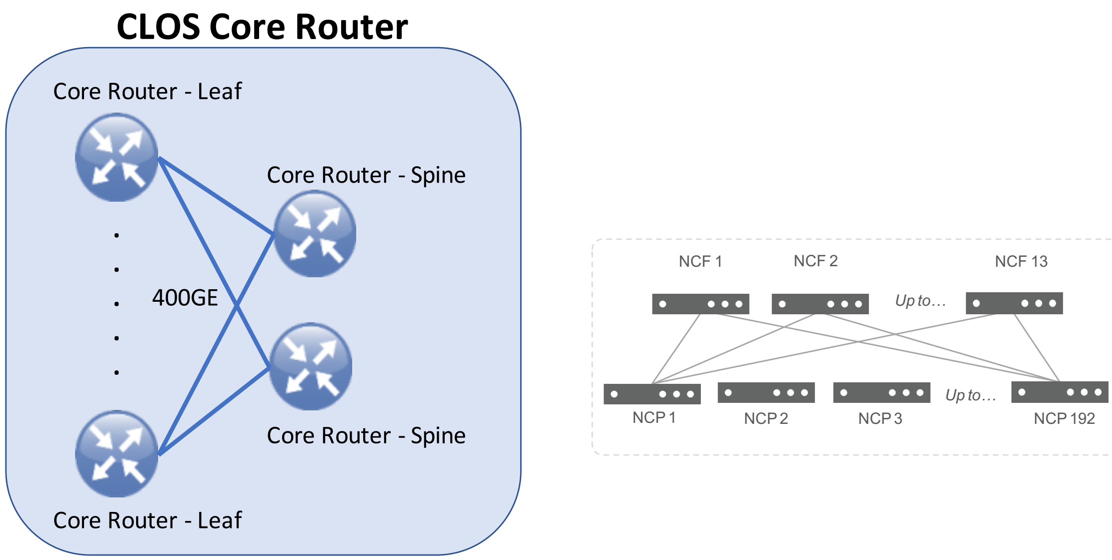

Scaling Incumbent Core using CLOS Architecture **Scaling DriveNets Core**

Figure 5. Scaling Incumbent Core Nodes with CLOS Switching

TCO MODEL ASSUMPTIONS

The TCO model is driven by detailed assumptions on:

- Regions served by the network
- Network traffic and growth
- Hardware, software, and equipment maintenance expenses
- OPEX assumptions

The model calculates the network TCO for small, medium, and large regions. The numbers of regions and number of aggregation, edge, and core nodes in each region is presented in Table 1.

Regions	Small	Medium	Large
Number of regions	12	5	2
Aggregation nodes	10	25	40
Edge nodes	2	2	2
Core nodes	2	2	2

Table 1. Number of Regions and Nodes

The size and growth of the network is driven by traffic requirements in each region. Traffic is generated by 4G and 5G mobile traffic, residential broadband traffic, and business data traffic (Tables 2–5).

Regions	Small	Medium	Large
4G traffic per cell site	100 Mbps	150 Mbps	150 Mbps
4G traffic growth rate	20%	46%	46%
Number of cell sites	1000	2000	5000
Cell site growth rate	3%	15%	15%

Table 2. 4G Traffic Assumptions

Regions	Small	Medium	Large
5G traffic per cell site	500 Mbps	800 Mbps	800 Mbps
5G traffic growth rate	20%	46%	46%
Number of cell sites	10	40	100
Cell site growth rate	80%	80%	80%

Table 3. 5G Traffic Assumptions

Regions	Small	Medium	Large
Residential subscribers	200,000	500,000	1,000,000
Traffic per subscriber	1.2 Mbps	1.2 Mbps	1.2 Mbps
Traffic growth rate	20%	20%	20%

Table 4. Residential Traffic Assumptions

Regions	Small	Medium	Large
Business CPE	500	1500	3000
Traffic per CPE	2 Mbps	2 Mbps	2 Mbps
Traffic growth rate	42%	42%	42%

Table 5: Business Traffic Assumptions

TCO RESULTS

The TCO model assumes that the incumbent router network uses the newest service provider routers available on the market today with the highest port densities and the lowest price per port. This ensures a fair comparison of the DriveNets architecture with the incumbent architecture.

The TCO model runs over five years, and the network capacity grows based on the mobile, residential, and business traffic. The TCO results are presented in Figure 6 and Table 6. The TCO is divided into two components:

1. Router Expenses: All expenses associated with routers, including hardware acquisition, hardware annual maintenance, and software licenses.
2. OPEX: All other operational expenses required to run the network.

The router expenses include both capital expense (router acquisition expenses) and OPEX (annual maintenance and annual software licenses). These are grouped into a single category to distinguish the long-term cost of the routers versus the cost of operating the network.

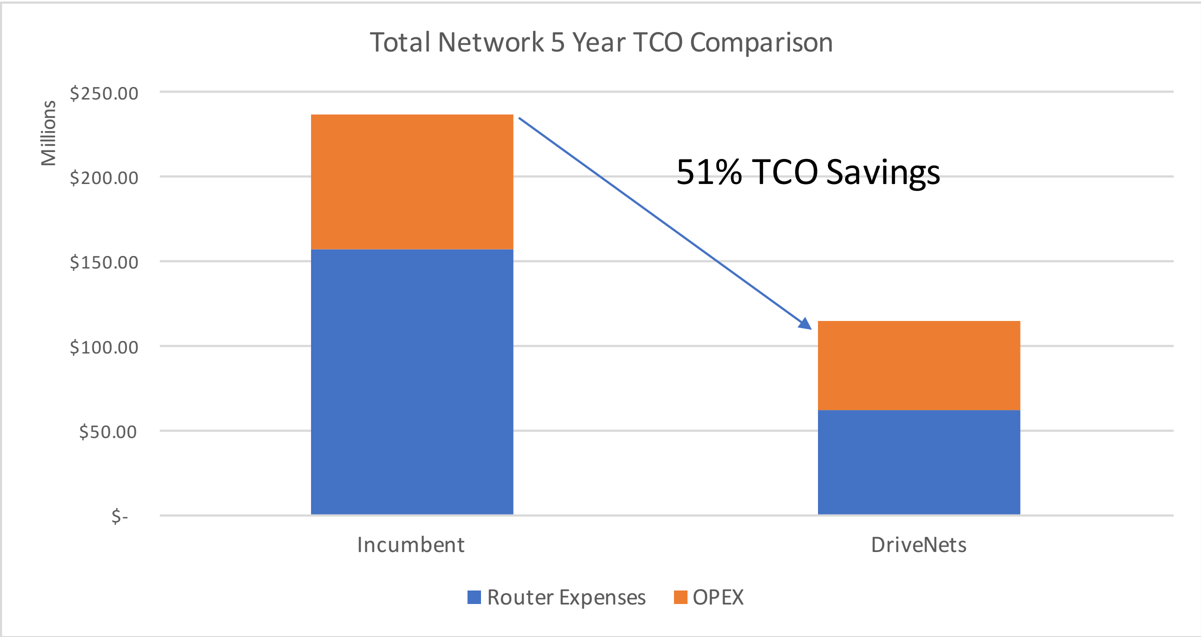

Figure 6. Five-Year Cumulative TCO Comparison

DriveNets TCO Savings	Savings
Router expenses	61%
OPEX	33%
TCO	51%

Table 6. Five-Year Cumulative TCO Savings

The DriveNets architecture has a significant five-year router expense savings of 61%. This is due to several key factors:

- DriveNets white-box architecture minimizes hardware expenses (due both to the cost efficiencies of white-boxes and the use of fewer boxes due to infrastructure convergence).
- DriveNets attractive software pricing provides lower cost per Gbps.
- Large scalability of the NCP/NCF architecture allows for:
    - Consolidation of edge and core routing services in a single routing node.
    - Elimination of CLOS switching architectures in the core node when core router ports are exceeded.

The five-year cumulative TCO is depicted in Figure 7. Note that the traditional network architecture has a steeper slope than the DriveNets architecture. This is important because it indicates that DriveNets costs will rise more slowly with increasing demand, which helps service providers maintain or improve profit margins on next-generation network services.

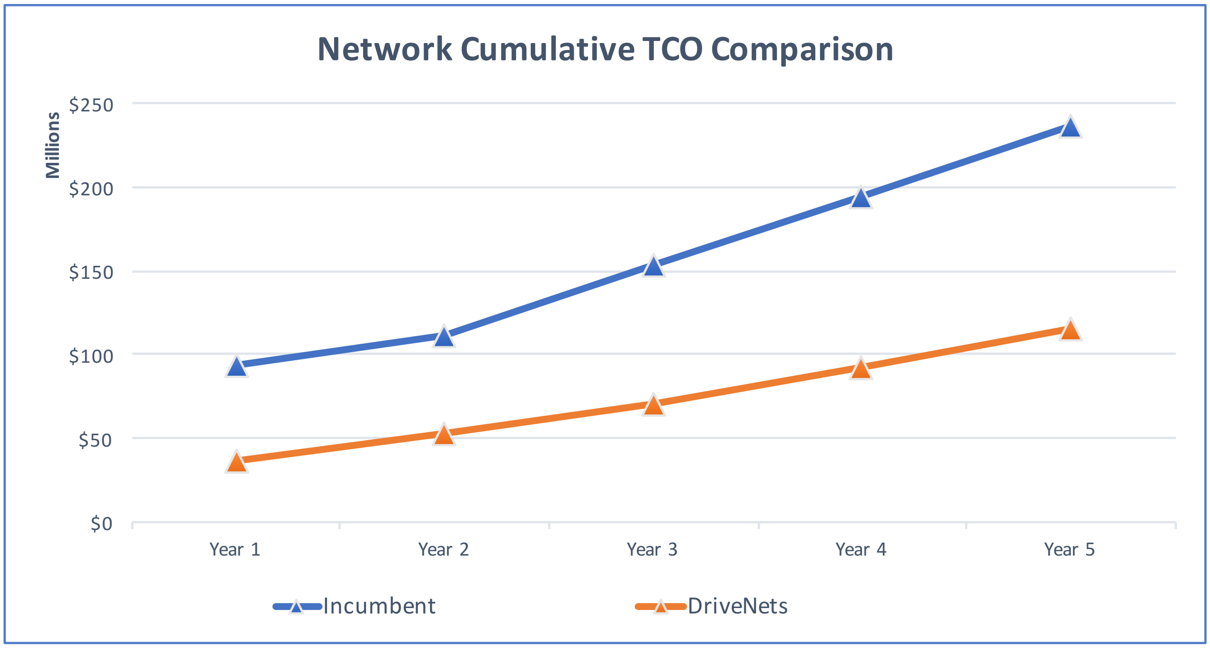

Figure 7. Five-Year Cumulative TCO

Figure 8 provides a breakdown of router expense savings by region. The greatest savings are in large regions. This is due to the number of 100 GE ports that exceeds the capacity of the core router, and a CLOS architecture with multiple core routers is employed to satisfy demand. This is an expensive and inefficient approach to scaling a core node as compared to the DriveNets fabric architecture, which allows for cost-effective scaling of core nodes as demand increases. However, the DriveNets architecture is also more cost effective in medium (51% savings) and small regions (48% savings) of the network where the CLOS architecture is not needed in the core.

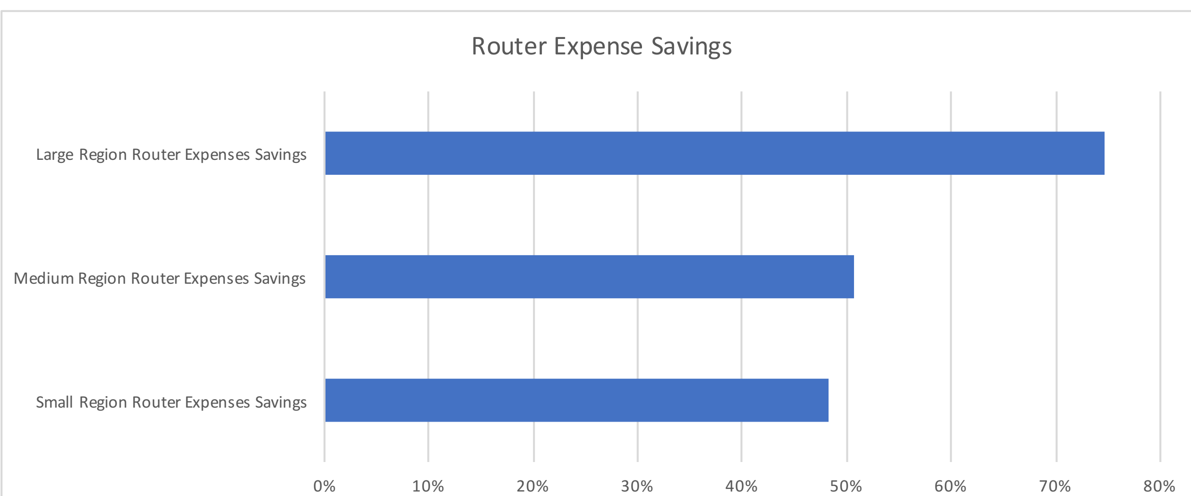

Figure 8. Breakdown of Router Expense Savings by Region

Figure 9 presents a breakdown of the five-year cumulative OPEX. In our OPEX we assume that there are one-time expenses for deploying the DriveNets network. Specifically:

- DriveNets OSS/BSS integration expenses of $1,500,000.
- DriveNets training expenses of $100,000.

The drivers of DriveNets OPEX savings are:

- Power and cooling of current routing CLOS core nodes is much higher than the DriveNets core nodes.
- Overall, more traditional routing nodes are needed to support demand, which also increases power and cooling expenses in the incumbent network.
- The functions of security, test and certification, network engineering, and software upgrades are duplicated for each type of traditional router; these functions are consolidated over all routing nodes in the DriveNets architecture.
- Expenses for service assurance, installation, moves, adds, and changes are a function of the number of routers deployed. DriveNets requires fewer routing nodes than the traditional architecture; therefore, these expenses are reduced.

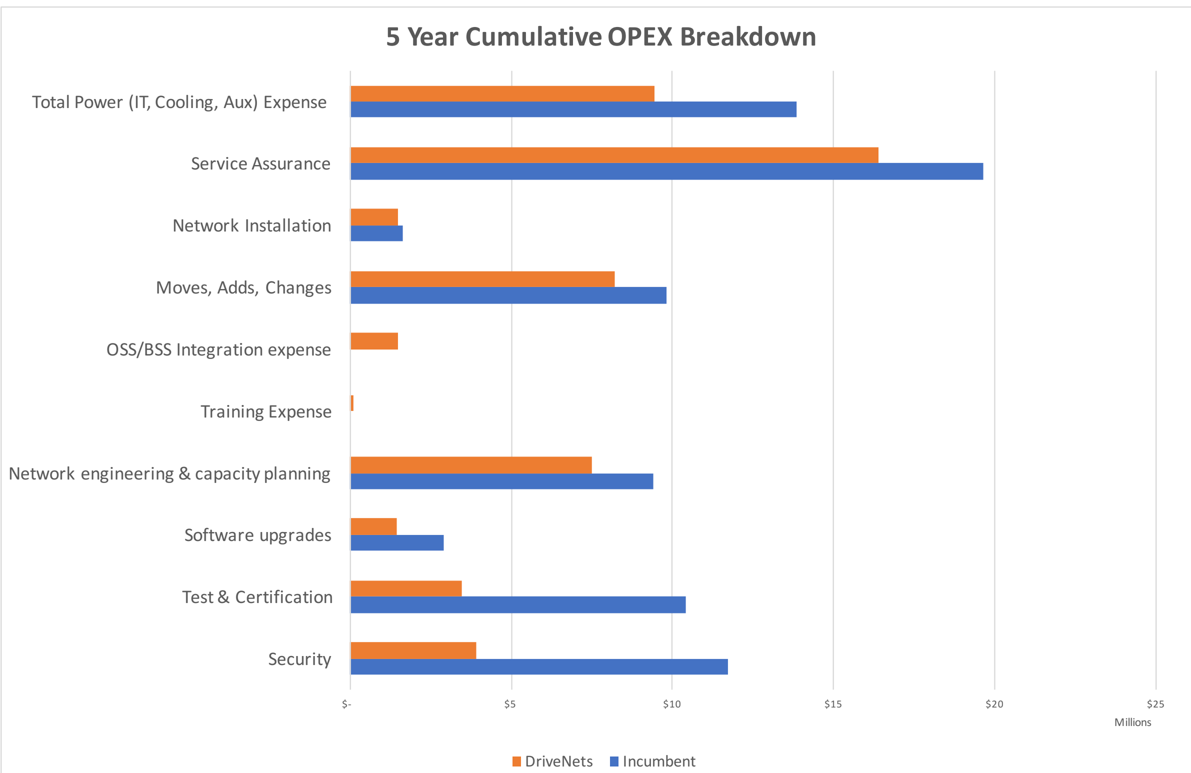

Figure 9. Five-Year Cumulative OPEX Breakdown

CONCLUSION

DriveNets has developed a new and innovative approach to high-scale networking. The DriveNets architecture supports multiple software-based networks over a single converged infrastructure that uses the same white-box building blocks to scale routing nodes from small access or aggregation nodes to extremely large core routing nodes. The white-box hardware allows service providers to take advantage of on-going cost reductions in merchant silicon. DriveNets has decoupled the control-plane and the data-plane and created an attractive software licensing model to further reduce the cost per bit.

Additionally, the white-box building blocks simplify OPEX by using a common routing platform throughout the entire network. The software-based approach also supports any service on any port, allowing multiple different networks to use a single converged infrastructure. This reduces the expenses of operating and managing different types of networks. The DriveNets Network Cloud architecture is also more power efficient, which reduces environmental expenses. Overall, DriveNets has 51% TCO advantage over traditional IP networks. This advantage increases further in larger networks.

Peter Fetterolf, Ph. D (pfetterolf@acgcc.com) is CTO with ACG Research. His primary focus is developing business models for next generation networks which includes IP transport, SDN, NFV, vEPC, vRAN, other mobile network functions, and optical transport networks. He is also responsible for software development of the Business Analytics Engine (BAE) software network economic simulation tool.