White Paper

Building Cost Effective and Scalable CORE Networks Using an Elastic Architecture

Executive Summary

Cisco® CRS Elastic Core is based on the Cisco Carrier Routing System (CRS). The Elastic Core is a scalable and flexible solution providing the lowest total cost of ownership when compared to core solutions provided by other vendors. The Elastic Core provides:

- **Scalability** with up to 400 Gbps per line card, and single-chassis, back-to-back, and multichassis configurations
- **Flexibility** with multiple options for packet forwarding cards including full IP core, lean core, and IP services
- **Support for a wide variety of physical interfaces** including Ethernet, SONET/SDH, and T3
- **Clean integration of IP and optical** with Cisco nLight™, allowing integration with the optical network using 100 Gigabit Ethernet IP over dense wavelength-division multiplexing (IPoDWDM) and an integrated IP and optical control plane

As Cisco CRS scales from single-chassis to multichassis systems, no rewiring, network topology, or routing changes are required as additional chassis are added. This allows service providers to easily scale their networks while minimizing operating expenses (OpEx). Alternative solutions that do not support multichassis equipment can require extensive rerouting, topology, and routing changes if additional core routers need to be added to a node to support demand.

A TCO analysis shows that the CRS core is more cost effective than alternative solutions, with up to 46 percent capital expenditures (CapEx) savings, 55 percent OpEx savings, and a total of 49 percent TCO savings.

Background Information

Core IP networks aggregate large amounts of traffic from residential, business, mobile, and wholesale networks. With the growth of “over the top” (OTT) video, mobile broadband, and the prospect of an Internet of Everything, core networks need to continue to scale beyond current boundaries. The Cisco Visual Networking Index™ (Cisco VNI™) forecast predicts a massive growth in IP traffic, which will directly impact the core. Some highlights from the VNI forecast are:

- Annual global IP traffic will surpass the zettabyte threshold (1.3 zettabytes) by the end of 2016. In 2016, global IP traffic will reach 1.3 zettabytes per year or 110.3 exabytes per month.
- Global IP traffic has increased eightfold over the past 5 years, and will increase threefold over the next 5 years. Overall, IP traffic will grow at a compound annual growth rate (CAGR) of 29 percent from 2011 to 2016.
- In 2016, the gigabyte equivalent of all movies ever made will cross global IP networks every 3 minutes. Global IP networks will deliver 12.5 petabytes every 5 minutes in 2016.

As a direct result of the growth in global IP traffic, Infonetics Research is projecting 1 Gigabit and 10 Gigabit Ethernet ports to grow at double-digit rates and 40 Gigabit and 100 Gigabit Ethernet ports to grow at triple-digit rates.

Cisco CRS Elastic Core solution provides a framework for scaling the core while minimizing capital and operational expenses. The next sections describe the key value propositions and present a TCO analysis of the CRS Elastic Core.

Cisco Elastic Core Value Proposition

Cisco CRS Elastic Core, powered by CRS routers, provides a scalable and cost-effective framework for next-generation core networks. The elastic core is:

- Scalable
- Flexible
- Supportive of a wide variety of physical interfaces
- Cleanly integrated with optical networks through Cisco nLight

Scalability: Up to 400 Gbps per Line Card

The Cisco CRS family of routers is highly scalable while supporting a wide variety of line cards that target various applications and price points. Three classes of line cards are available: 40 Gbps, 140 Gbps, and 400 Gbps. The 400-Gbps line cards and switch fabrics are new additions to the CRS product family, and include these configurations:

- Four 100 Gigabit Ethernet Ports per Card
- Ten 40 Gigabit Ethernet Ports per Card
- Forty 10 Gigabit Ethernet Ports per Card

Scalability: Multichassis Architecture

The CRS routing system is a highly scalable multichassis system. The key benefit of a multichassis router is that additional port capacity can be incrementally added to the router while maintaining a nonblocking switching fabric.

This is not the case for a traditional single-chassis router, where additional routers need to be added and interconnected to add service ports to a routing node. This problem is illustrated in Figure 1.

Figure 1. Capacity Additions with Single-Chassis Routers

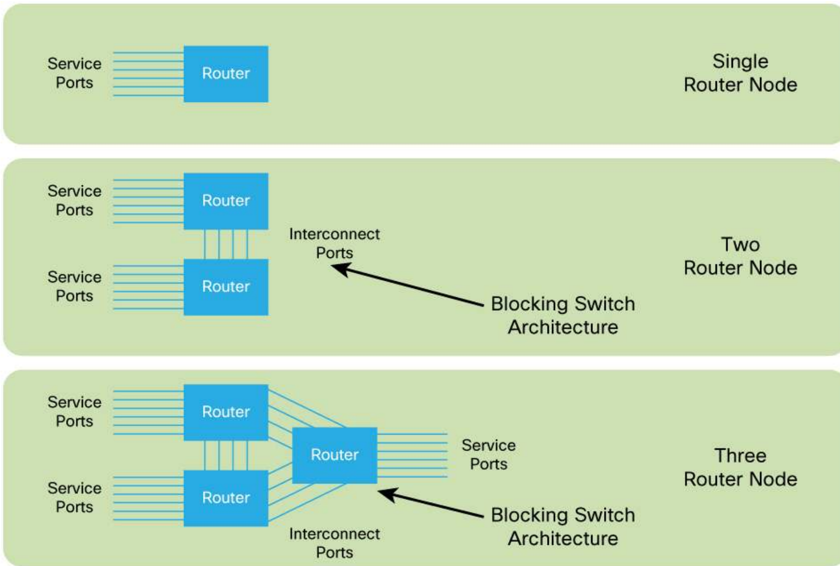

The top node in the figure depicts a single node router. In this case the number of service ports at the node fits within a single router. As the number of ports grows, it might be necessary to add another router as depicted in the two-router node. If the ports on the two-router node are exceeded, a third router must be added. Each time a router is added to the node, interconnect ports must be used to interconnect the routers. There are several problems with this approach:

- Interconnect ports consume service ports that could be used to generate revenue.
- To meet traffic requirements, between one third and one half of the ports on the router must be used for router interconnect.
- The interconnect architecture is a blocking-switching architecture; this means that in some traffic conditions, there may not be enough bandwidth in the router interconnect to serve demand.
- As routers are added, service ports need to be moved from one router to another, causing service disruption and increasing operational complexity and cost.

This traditional router architecture leads to serious scalability and operational problems. It also consumes a high level of CapEx due the expenses for ports interconnecting multiple routers. As the number of routers grows, these problems become more severe.

Cisco CRS solves these problems by using a scalable multichassis architecture. CRS routers come in many sizes: 4-, 8-, and 16-slot chassis. Additionally, CRS routers can be scaled by combining two chassis using a back-to-back architecture or combining up to nine chassis in a multichassis system.

A Cisco CRS multichassis system consists of two major elements: a line card chassis (LCC) and a fabric card chassis (FCC). The LCC hosts Performance Route Processor (PRP) cards, the first and third stage of switch fabric cards, and line cards that provide the physical interface and process data packets. The FCC hosts the second stage of the switch fabric cards. The LCC and FCC are connected by a set of optical cables. Expanding your core network capacity is a smooth process for Cisco CRS, supported by in-service hardware and software upgrades. Upgrading from a single chassis to a back-to-back or multichassis system does not require any rewiring or disruption to service ports, simplifying the operational process and reducing service downtime.

This architecture is illustrated in Figure 2. As the number of service ports grow, a second LCC is added to increase port capacity. The two LCCs are connected back-to-back using a fabric interconnect over optical cables. No service ports are required for interconnect and the switch fabric is nonblocking. For a two-chassis system no FCC is required. If three or more chassis are needed, LCCs are connected using an FCC. Again no service ports are required for interconnect and the switch fabric is nonblocking.

Figure 2. Capacity Additions with Multiple Chassis

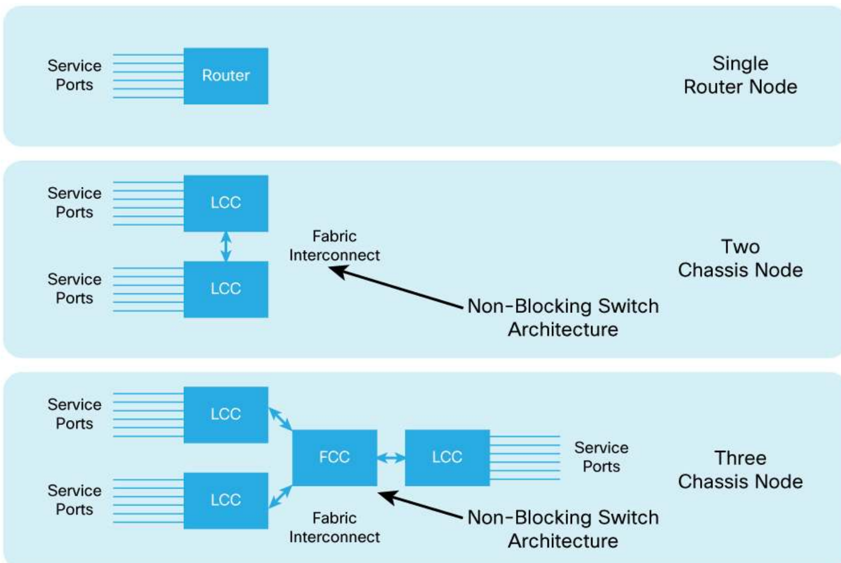

The key benefits of the CRS multichassis architecture are:

- No service ports are required for LCC interconnect.
- The switch fabric is nonblocking.
- Adding back-to-back or multichassis capacity with an FCC is a “hitless” upgrade, with no service disruption or system rewiring.

This architecture allows service providers to scale network capacity from 12.8 to 922 Tbps in accordance with demand. The major CRS system components are depicted in Figure 3. The system can scale to up to 72 chassis. However, in the field the largest deployment is 9 LCC + FCC.

Figure 3. Cisco CRS System Components

Chassis	CRS 16 Slot	CRS 16 Slot 2+0 Back-to-Back	CRS 16 Slot X+1 Multichassis
			
Aggregate Switching Capacity	12.8 Tbps	25.6 Tbps	Up to 921.6 Tbps
Number of Forwarding Slots	16	32	Up to 1152

Scalability: 100 Gigabit Ethernet IPoDWDM

The Cisco CRS router provides additional levels of scale with integrated, coherent 100 Gigabit Ethernet IPoDWDM optics. Integration of DWDM into the CRS line cards eliminates the need for a set of 100-Gb grey optics interconnecting the router card to a transponder in a reconfigurable optical add-drop multiplexer (ROADM). Figure 4 depicts a router connected to a ROADM using 100-Gb grey optics. This scenario requires a set of grey optics on both the router and the transponder. The integrated 100-Gb solution is illustrated in Figure 5. In this scenario, the grey optics are eliminated, reducing both CapEx and OpEx, with a TCO savings of up to 29 percent. The details of this analysis can be found in a white paper published by ACG Research: Why Service Providers Should Consider IPoDWDM for 100G and Beyond.

Figure 4. Router Connected to a ROADM Using Grey Optics

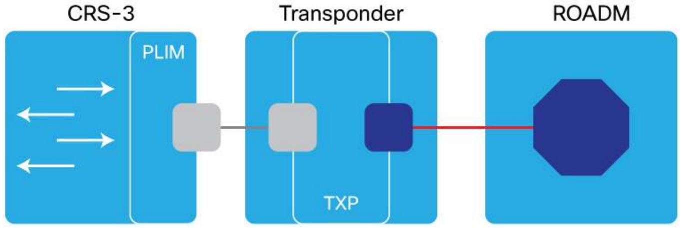

Figure 5. Router Directly Connected to a ROADM

Flexibility

Flexibility is a fundamental part of the CRS architecture; a variety of packet forwarding and service cards can be selected to perform different functions, including edge services, Carrier Grade Network Address Translation (NAT), IPv6 conversion, distributed denial of service (DDoS) protection, Internet peering, and lean core switching. The various cards allow service providers to configure large routers with a variety of characteristics and functions, optimizing services while reducing TCO. Cisco CRS uses a midplane architecture in which a Physical Layer Interface Module (PLIM) plugs into the midplane and connects to a packet-forwarding card. There are many types of PLIMs with multiple interface types, including SONET/SDH and Ethernet. The forwarding cards provide the appropriate services at the right price points. The service provider can mix and match the following packet forwarding cards:

- Modular Services Card (MSC): A scalable IP edge service card providing massive VLAN scale and wire-rate hierarchical quality of service (HQoS)
- Forwarding processor (FP): A scalable IP core forwarding card targeted at core routing and Internet peering

- Label-switched path (LSP): A scalable lean IP core forwarding card targeted at lean-core Multiprotocol Label Switching (MPLS) deployments
- Carrier-Grade Services Engine (CGSE): A scalable services engine that can run IPv6 services, anti-DDoS services, and other network applications and services

The Cisco CRS cards can be configured as edge routers, services engines, Border Gateway Protocol (BGP) core peering routers, or lean-core MPLS switches. This flexibility allows service providers to mix and match functions in a single chassis or multichassis CRS, thus optimizing their network and reducing overall TCO. This approach contrasts with that of some competitors that use different products for lean core routing, full core routing, and edge routing.

Wide Variety of Physical Interfaces

The CRS router allows service providers to mix and match a wide variety of physical interfaces using PLIMs. The PLIMs can be matched with the various packet forwarding cards to optimize the network architecture and services. The interfaces include:

- T3/E3
- SONET/SDN: OC3/STM through OC768/STM256
- Ethernet: 1-, 10-, 40-, and 100-Gigabit Ethernet
- Integrated 100-Gb IPoDWDM or grey optics interfaces

This variety of physical interfaces gives service providers the flexibility to serve existing as well as future interfaces with the same CRS router

IP and Optical Integration with nLight

Another key value proposition of the Cisco CRS system is that it is closely integrated with the optical network. This provides a wide range of financial benefits that are articulated and quantified in an ACG white paper entitled: The Economics of Cisco's nLight Multilayer Control Plane Architecture.

The fundamental problem is that because the IP and optical control planes are not integrated, service providers' networks are always overengineered and underutilized, with 50 to75 percent of network capacity deployed for network resiliency, not for passing traffic. Cisco has addressed this problem with nLight: a multilayer routing and optimization architecture that focuses on IP and optical integration, increasing network agility and flexibility while improving network utilization. By enabling information flows between the routing and optical layers, nLight provides an end-to-end protection and restoration approach that meets performance constraints, such as "five 9s" (99.999 percent) availability, at much lower TCO than present methods, such as the widely used 1+1 optical protection scheme. Cisco's broad portfolio of DWDM enablers, such as colorless, contentionless, and omnidirectional add/drop, and Flex Spectrum transmission, are enablers of nLight.

Cisco nLight is the optical and IP layer of Cisco Open Network Environment (ONE), a portfolio of Cisco technologies and open standards that provides programmatic control of the network. Cisco nLight provides programmatic interfaces between the routing layer (Layer 3) and the optical transport layer (Layer 1). Figure 6 shows where nLight resides within the overall Cisco ONE framework.

Figure 6. Cisco nLight Within Cisco ONE

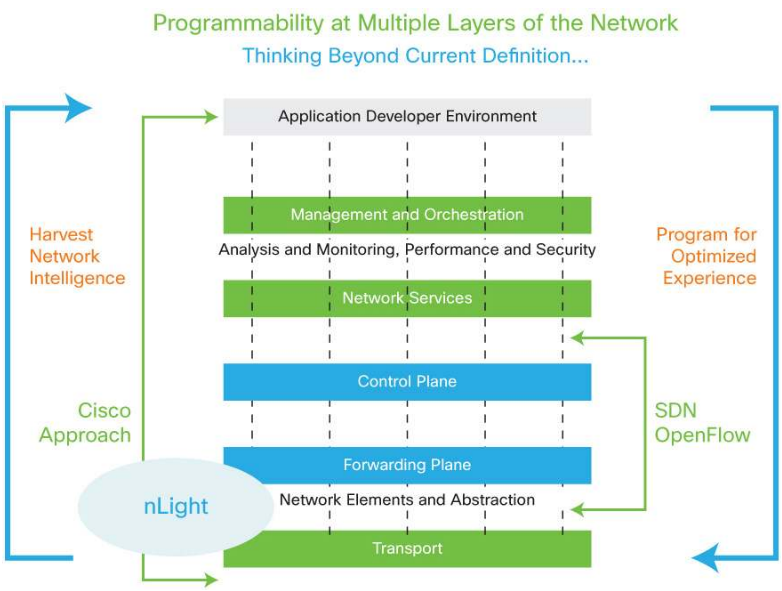

As shown in Figure 6, nLight provides an integrated optical and IP control plane. The integrated control plane produces a virtual and agile IP and optical topology. Using the integrated control plane, it is possible to simultaneously optimize optical and IP total cost of ownership subject to constraints such as:

- Latency
- Shared Risk Link Groups (SRLG)
- Restoration
- Protection

Optimization of the IP and optical topology is further enhanced by the Cisco ONE controller, which analyzes network traffic, topology, and failure conditions and uses this information to run optimization algorithms to provide the lowest-cost network that satisfies network resiliency and performance constraints. The controller architecture, which enhances the capabilities provided by nLight, is an optional component of the Cisco ONE architecture.

Business Case for Cisco CRS Elastic Core

This section uses a TCO model to quantify the financial benefits of a CRS Elastic Core.

Model for TCO Comparisons

This model uses a hypothetical greenfield network to compare a CRS network with core networks provided by from two competitors called Vendor A and Vendor B. The key characteristics of the CRS Elastic Core as compared to solutions provided by Vendor A and B are presented in Figure 7. The CRS system scales from single chassis to multichassis, supports a large variety of packet forwarding and physical interfaces, and integrates tightly with the optical network.

Figure 7. Cisco CRS Elastic Core Compared to Competitor Solutions

Cisco	Vendor A		Vendor B	
CRS-3 	Core Router	Edge Router	Core Router	MPLS Switch
• 400 Gbps per Slot • IP Edge Services • Full IP Core • Lean MPLS Core • Ethernet • SONET/SDH • Single Chassis, Back-to-Back, and Multichassis • IPoDWDM • nLight	• 400 Gbps per Slot • Full IP Core • Lean MPLS Core • Ethernet	• 120 Gbps per Slot • IP Edge Services • SONET/SDH • Ethernet	• 240 Gbps per Slot • Full IP Core • Ethernet • SONET/SDH • Single Chassis and Multichassis	• 400 Gbps per Slot • Lean MPLS Core • Ethernet

The core router provided by vendor A is an Ethernet-only core routing solution. Therefore, if SONET/SDH interfaces are present in the core network, a second services router is required. Also, Vendor A does not support multichassis systems, so if additional chassis are needed in a core node, service interfaces must be used to combine multiple routers into a larger routing node. Vendor A does not support IPoDWDM, so 100-Gb grey optics and 100-Gb transponders are required for DWDM transport.

Vendor B has separate products for core routing and core MPLS switching (lean core). The MPLS switch is targeted at core packet transport functions and does not support IP peering or SONET/SDN interfaces. Therefore two core routers are needed if the network design demands IP peering, SONET/SDH interfaces, and lean MPLS links. Vendor B does not support IPoDWDM, so 100-Gb grey optics and 100-Gb transponders are required for DWDM transport.

For each of the three solutions, network configurations are generated based on a set of assumptions:

- Core network containing 10 nodes
- Demands evenly divided among all the nodes and the total demand specified in Table 1
- 400-Gb Cisco CRS line cards and switching fabric
- FP 400 cards used for access ports on the CRS router (demand)
- LSP 400 cards used for trunk-side ports on the CRS router

The demand includes Ethernet, SONET/SDH, and full and lean IP ports. Full IP ports are defined as ports that require IP routing with full support for routing protocols and Internet peering. Lean ports are layer 2.5 transport ports that only require MPLS and minimal IP routing.

Table 1. Port Requirements for Total Demand

Port Requirements	Number of Ports	CAGR
10GE Full IP	500	10%
100GE Full IP	500	10%
OC192 Full IP	500	2%
10GE Lean IP	100	10%
100GE Lean IP	100	10%

Each of the architectures is presented in the figures that follow. Figure 8 shows the Cisco CRS core architecture. The demand for full IP ports is served by the FP 400 forwarding engines. Demand for lean core ports and the core links interconnecting core nodes is served by LSP 400 forwarding engines.

Figure 8. Cisco CRS Core Architecture

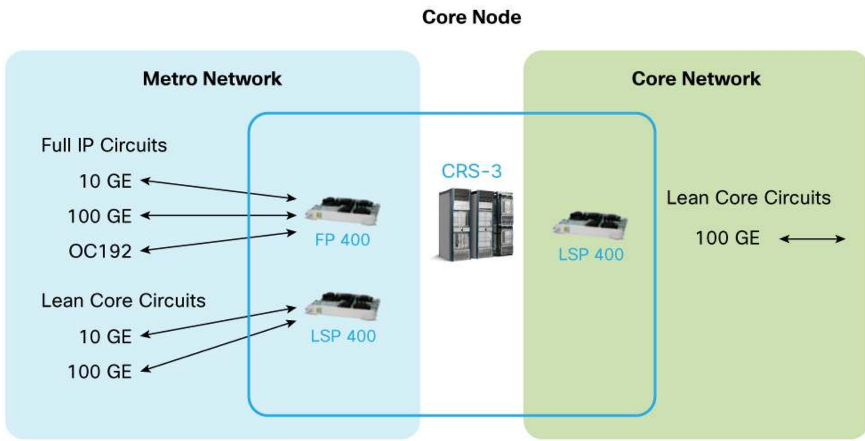

The architecture for Vendor A's solution is presented in Figure 9. All 10 Gigabit Ethernet and 100 Gigabit Ethernet full and lean IP ports are serviced by the large core router. However, SONET/SDH ports must be served by an edge router that is connected to the core router. Also, Vendor A uses 100 Gb grey optics with 100 Gb transponders for core links interconnecting routers over a DWDM network.

Figure 9. Vendor A Core Architecture

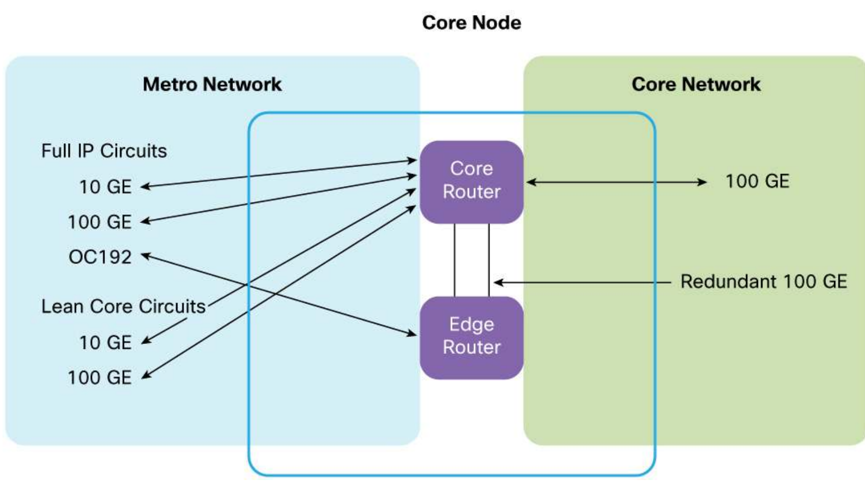

Vendor B's solution is presented in Figure 10. Vendor B uses a large core router for traditional full IP routing and has both Ethernet and SONET/SDH interfaces. Vendor B has a different product for lean core switching that uses a scaled-down MPLS switch. Lean core circuits are terminated on the MPLS switch and all core links are interconnected through the MPLS switch.

Figure 10. Vendor B Core Architecture

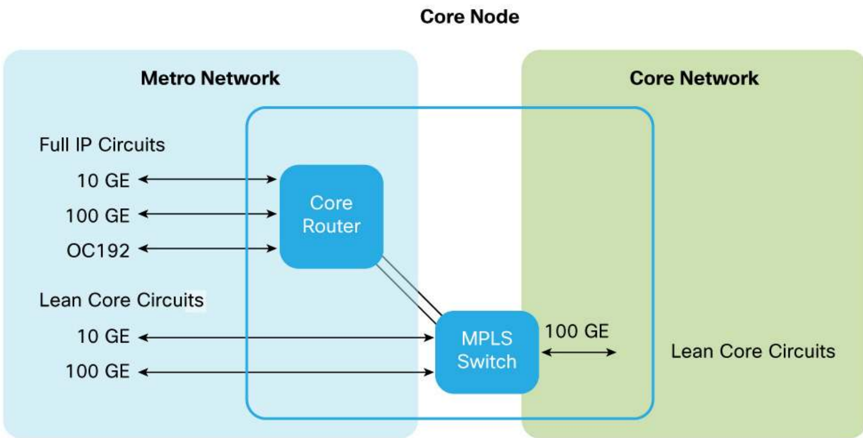

TCO Results

A TCO model was used to analyze and compare CapEx and OpEx of the three solutions.

Figure 11 and Table 2 present the 6-year cumulative CapEx and OpEx of each solution. CapEx consists of all equipment and systems integration expenditures. OpEx consists of environmental expenses, service provider operations expenses, and vendor technical service expenses. The analysis shows that for the assumptions presented earlier, the overall TCO of the Cisco solution is 49 percent lower than that of Vendor B and 38 percent lower than that of Vendor A. The TCO benefits are due to a variety of factors including the flexibility of Cisco CRS to support both lean and full IP interfaces, Ethernet and SONET interfaces, back-to-back and multichassis scaling, and integrated IP and optical transport with Cisco nLight.

Figure 11. Six-Year Cumulative TCO Comparison

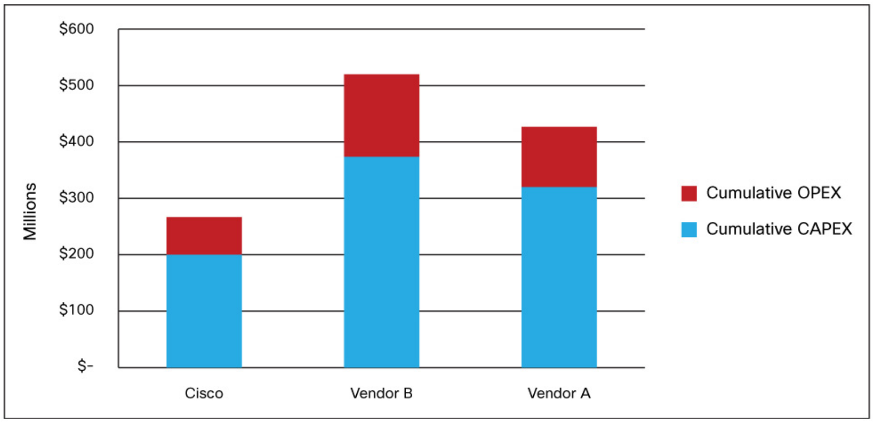

Table 2. Six-Year Cumulative TCO Comparison

	Cisco	Vendor B	Vendor A
Cumulative CAPEX	$202,713,399	$378,032,049	$323,119,414
Cumulative OPEX	$60,572,050	$135,196,826	$98,369,411
Cumulative TCO	$263,285,449	$513,228,875	$421,488,825
TCO Savings	N/A	49%	38%
CapEx Savings	N/A	46%	37%
OpEx Savings	N/A	55%	38%

Total annual power consumption for each of the solutions is shown in Figure 12. The efficiency and flexibility of Cisco CRS to support demand in a single multichassis system rather than with separate routers and switches is one of the key factors leading to power efficiency.

Figure 12. Annual Power Consumption Comparison

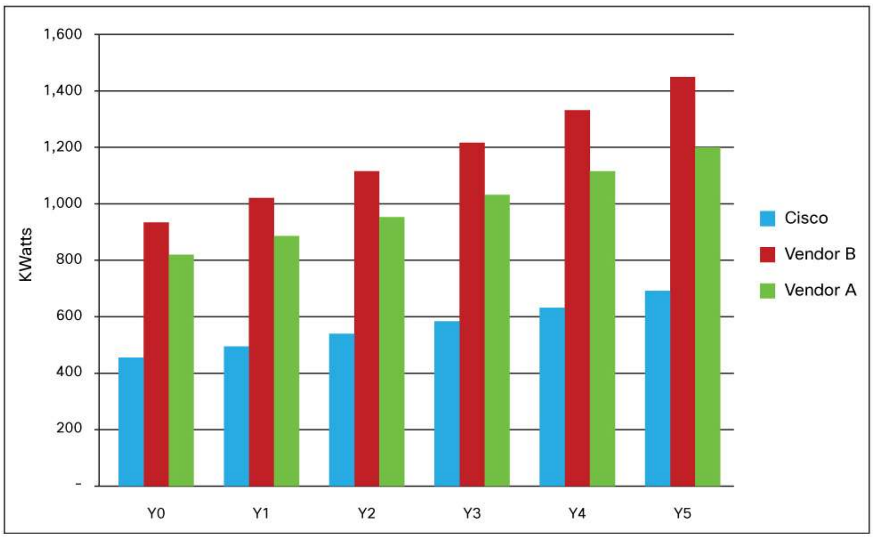

A detailed comparison of the OpEx of each solution is presented in Figure 13. OpEx categories include:

- **Certification and upgrades**: processes required for new software releases and new system deployments
- **Service**: cost of technical support and service provided by the vendor
- **Network care**: labor expenses incurred by the service provider for managing and operating the core network
- **Space**: Cost of floor space in the point of presence (POP)
- **Cooling**: Cost of cooling core router equipment
- **Power**: Cost of power for core router equipment

A single integrated CRS elastic core reduces expenses in each of these categories as compared to the solutions provided by Vendor A and Vendor B.

Figure 13. Operating Expense Components

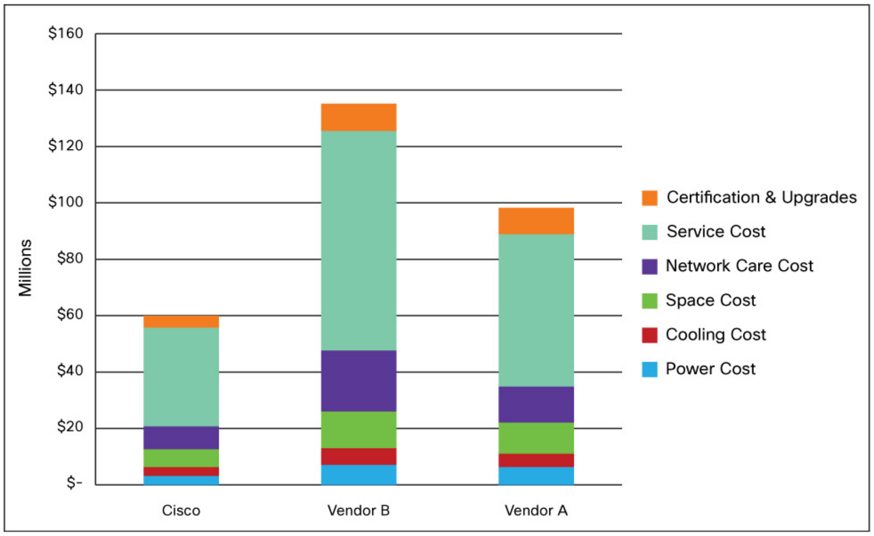

Conclusion

The key benefits of Cisco CRS Elastic Core are scalability and flexibility. The CRS can scale from single chassis to multichassis with no additional wiring, topology, or routing changes, minimizing operations disruptions and reducing OpEx. The wide range of packet forwarding engines and physical interfaces provides the flexibility to cost-effectively scale the core with a CRS solution - adding chassis, forwarding engines, service engines, and physical interfaces as demand for various services grows. Solutions from other vendors require different routers or MPLS switches to provide different core services. Cisco CRS Elastic Core provides all of these interfaces and services on a single, scalable, multichassis system.

A TCO analysis shows that the CRS Elastic Core solution is more cost effective than alternative solutions, with up to 46 percent CapEx savings, 55 percent OpEx savings, and a 49 percent TCO savings.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-727983-00 06/13