

RETHINKING CORE ECONOMICS

QUANTIFYING THE POWER AND SPACE ADVANTAGES OF THE ALCATEL-LUCENT 7950 XRS PLATFORM

FINANCIAL WHITE PAPER

Broadband traffic projections continue to show exponential growth, forcing service providers to balance network capacity and operational costs on a continual basis to profitably stay ahead of escalating demand. The reality of mounting power and space requirements places significant pressure on service providers, and makes the status quo untenable.

The move to 100G technology within the network infrastructure represents a major inflection point. With 100G links as the new currency in core networks, service providers are seeking to maximize the impact of their network infrastructure in delivering a full spectrum of services in the most cost-effective manner.

The Alcatel-Lucent 7950 Extensible Routing System (XRS) is an innovative core routing platform designed to ensure efficient network scaling for years to come. The industry's first network processor-based core routing platform powered by groundbreaking 400G silicon (FP3), it delivers industry-leading density for 10GE, 40GE and 100GE interfaces as well as a clear path to 400GE and terabit Ethernet interfaces.

This financial network modeling study from Alcatel-Lucent Bell Labs compares the costs associated with addressing core network growth using the 7950 XRS to existing solutions, with particular focus on key operational cost parameters such as rack space, power and cooling requirements.

TABLE OF CONTENTS

Executive summary / 1

Business case assumptions / 2

Study objectives / 4

Business case results / 5

Rack space / 5

Power / 6

Weight / 7

Capacity headroom / 8

Network consolidation savings / 8

CAPEX efficiencies / 9

Space / 10

Power / 12

Capacity headroom / 12

Conclusion / 13

Acronyms / 13

References / 13

EXECUTIVE SUMMARY

This paper quantifies the operational cost advantages of the Alcatel-Lucent 7950 XRS relative to traditional core routing alternatives that are currently deployed in service provider networks.

The business case model shows that the superior scalability, port density and power efficiency of the 7950 XRS result in significant savings of recurring operational costs associated with rack space, power and cooling. Traditional platforms are shown to incur nearly twice the cost when compared to the 7950 XRS.

In a realistic network model spanning a 16-node topology:


- The 7950 XRS required less than a third of the space used by legacy solutions, which in this business case resulted in approximately 60 percent cost savings in rack space over the five-year study period (at \$3,000 per rack annually).
- The 7950 XRS consumed up to three times less power, resulting in 50 percent savings (at 12.5c per kWh) over the entire network during the five-year study period.
- The 7950 XRS accommodated traffic growth over the five-year period with virtually no overhead associated with multi-chassis switching shelves, and with more than 50 percent headroom for growth available on most sites. Legacy solutions necessitated multi-chassis implementations to meet the capacity need for all but the smallest locations, and required frequent additions of new chassis to cater to growth.
- The tremendous capacity and port density of the 7950 XRS enable service providers to rethink their core network, so that they consolidate and concentrate core routing capacity in fewer, albeit larger, core routing nodes. A study of a consolidated eight-node topology shows that the 7950 XRS yields double the cost savings compared to competing solutions, while still maintaining ample headroom for growth.

Although this financial study examines a single business case, every care has been taken to ensure a balanced and fair comparison between competing solutions. Alcatel-Lucent believes these results are representative of the superior performance and economics that the 7950 XRS platform offers. We encourage operators to assess these results for their own business case and offer our full assistance in doing so.

BUSINESS CASE ASSUMPTIONS

The study is based on a 16-node topology for a national backbone that spans the United States (US). Core routing nodes are placed in major metros, totaling about 100 million subscribers across the network footprint. Aggregate traffic to and from each core routing node is assigned using a gravitational model based on subscriber population densities.

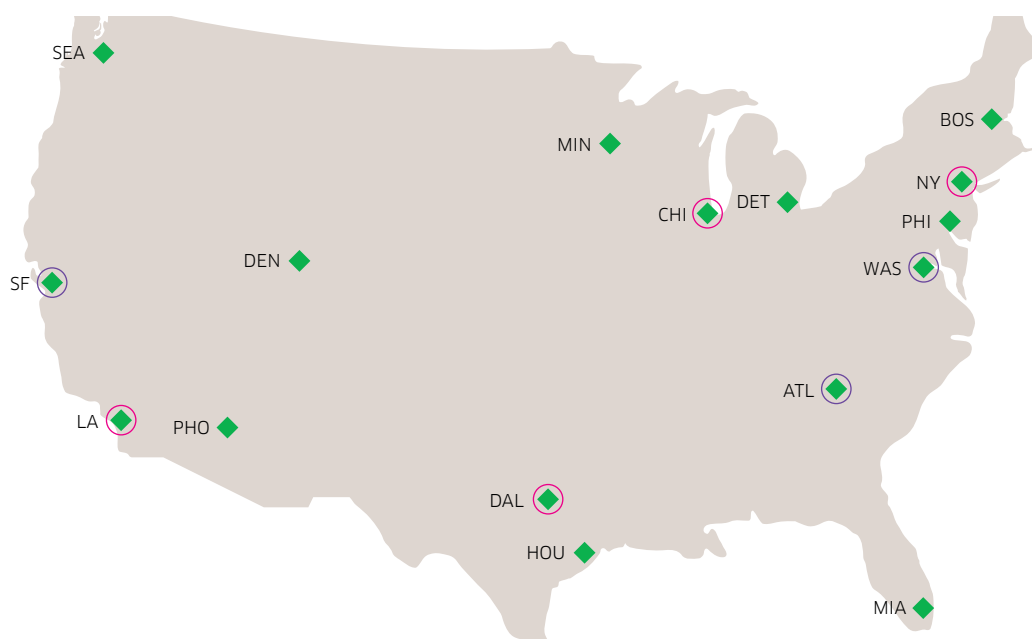
Figure 1. Node topology



Source: Bell Labs Analysis

Figure 1 provides the list of core routing nodes and subscriber populations served, ranging from small (2.5 million) to very large (18.9 million) metropolitan areas, with an average of 6.4 million subscribers. Figure 2 shows the node topology.

Figure 2. Network topology



Source: Bell Labs Analysis

The four largest metro regions are marked with red circles, the next three largest are marked with purple circles. We assumed that local traffic is routed within the Provider Edge (PE) routers, thereby not loading Provider (P) routers in the core. We also assumed that the amount and distribution of outgoing traffic is proportional to the population sizes served by the target nodes. The study is conducted over a five-year period and assumes that each subscriber entity generates a symmetric traffic flow of 100 kb/s on average during busy hours, with an annual growth rate of 38 percent. This is a conservative but reasonable estimate for user behavior in the recent past. The service provider is assumed to address one out of three (33.3 percent penetration) of the 102 million subscribers in the geography served. For redundancy reasons the carrier provisions twice the required capacity to serve busy-hour subscriber traffic.

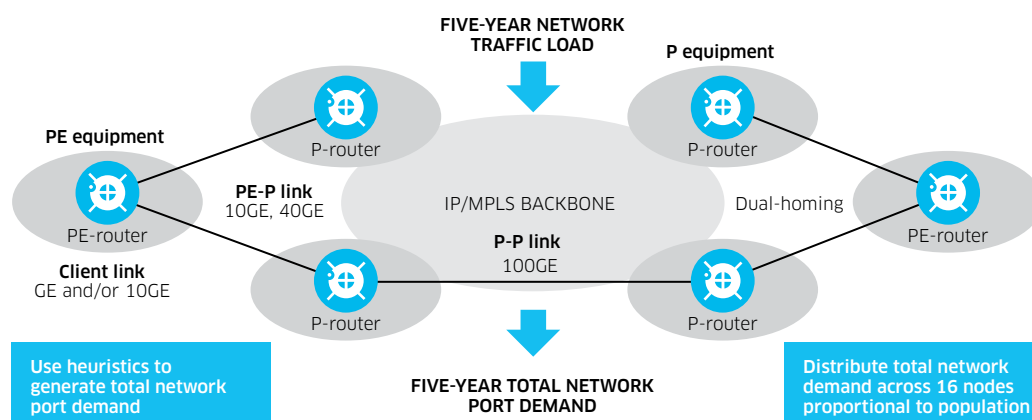
Connectivity of edge (PE) routers to the core (P) routers is assumed to be using 10GE interfaces, though migration to 100G connectivity from PE routers is a reality that would further enhance the Alcatel-Lucent gains. Connectivity between P-routers is assumed using 100GE links over an ROADM/DWDM optical transport network. The underlying transport network is identical for each of the modeled scenarios to ensure a level playing field. No use of router bypass for transit traffic has been made. Although router bypass can be applicable to further optimize IP core network cost, the decision regarding use of router bypass is in principle independent of the chosen core routing equipment.

To determine the optimal link topology, a service provider will likely deploy network design aids such as network planning and analysis tools (NPAT). For the purpose of the Total Cost of Ownership (TCO) comparison we were less concerned with finding an optimal link topology, because the same node and link topology is applied for each of the competing solutions. Bell Labs deployed an internally developed, heuristics-based tool that does not involve manual steps, as is the case with NPAT, to conduct a complete sensitivity analysis of the variables.

The Bell Labs tool estimates the core network cost based on:

- *Level of meshing*, reflecting the number of remote Server Nodes connected by each Server Node
- *Core capacity indicator* (τ), defined as total capacity available in the core, divided by the minimal capacity required, calculated statistically based on given level of meshing and ingress traffic at Server Nodes. Set at 1.5.

Figure 3. HLN tool input and output



Source: Bell Labs Analysis

The outputs of the dimensioning tool for an “Average P-Node” and for the total network are shown in Figure 4. From this “Average Node” configuration we derived the sixteen individual core routing node specifications in proportion to their actual subscriber size divided by average subscriber size. For example New York has 18,897,109 divided by 6,410,260 or 2.95 times the “average” node dimensions.

Figure 4. Average node and network dimensions

PER P NODE (“AVERAGE” NODE)

	YEAR 1	YEAR 2	YEAR 3	YEAR 4	YEAR 5
PE-P ports (10GE)	42	54	72	102	138
P-P ports (100GE)	12	16	22	31	42

TOTAL NETWORK PORTS (16 NODES)

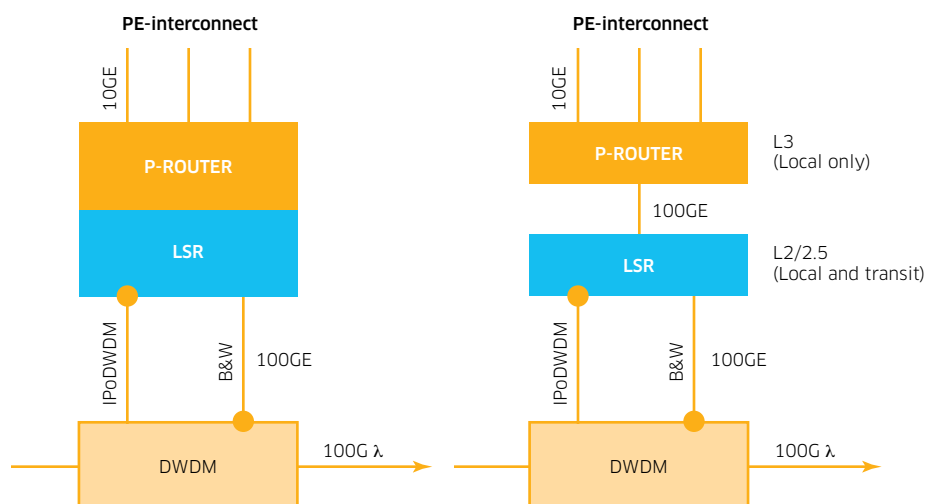
	YEAR 1	YEAR 2	YEAR 3	YEAR 4	YEAR 5
PE-P ports (10GE)	672	864	1152	1632	2208
P-P ports (100GE)	192	256	352	496	672

Source: Bell Labs Analysis

STUDY OBJECTIVES

The study compares the cost of deploying the 7950 XRS with alternative solutions from the two other leading incumbent core router suppliers (Suppliers B and C). The 7950 XRS – as well as incumbent routing platforms from Suppliers B and C – combines IP routing and Label Switch Routing in a single core routing platform, as shown at the left of Figure 5. To achieve the required capacity and efficiency, some have recently proposed single-function LSR-only platforms that would have to be used in combination with their existing P-router platform in a hybrid solution as shown at the right of Figure 5. Supplier C claims additional cost efficiencies in a super core from switching transit traffic in a dedicated LSR platform, and offloading the P-router of this traffic. This study will compare the Alcatel-Lucent 7950 XRS against these alternatives as well.

Figure 5. Converged P-router and LSR (left) versus separate platforms (right)



Source: Bell Labs Analysis

Table 1 lists the system dimensioning data inputs that were used to determine the required system configurations and related rack space and power usage. Only line cards with non-oversubscribed port capacity were used in all cases. Supplier data was used where publicly available. Conservative best estimates were made where this was not the case.

Table 1. System dimensioning inputs

	7950 XRS-20	P-ROUTER B	P-ROUTER C	DEDICATED LSR C
Shelves per rack	1	1	1	1
Slots per shelf	20	16	16	16
10G ports/slot (line rate)	40	14	12	24
40G ports/slot (line rate)	10	3	3	6
100G ports/slot (line rate)	4	1	1	2
Multi-chassis support	Yes	Yes	Yes	No
Power consumption	2 W/G	2.75 W/G	2.7 W/G	~1 W/G
Weight (lb), chassis, MC-shelf	1000/900	1600/1630	606/900	1000/NA

Source: Bell Labs Analysis

The objective of the Bell Labs modeling study was to quantify and compare operational costs associated with space, power/cooling and weight between the 7950 XRS and competing products.

The study includes a sensitivity analysis related to node topology and link capacity choices.

BUSINESS CASE RESULTS

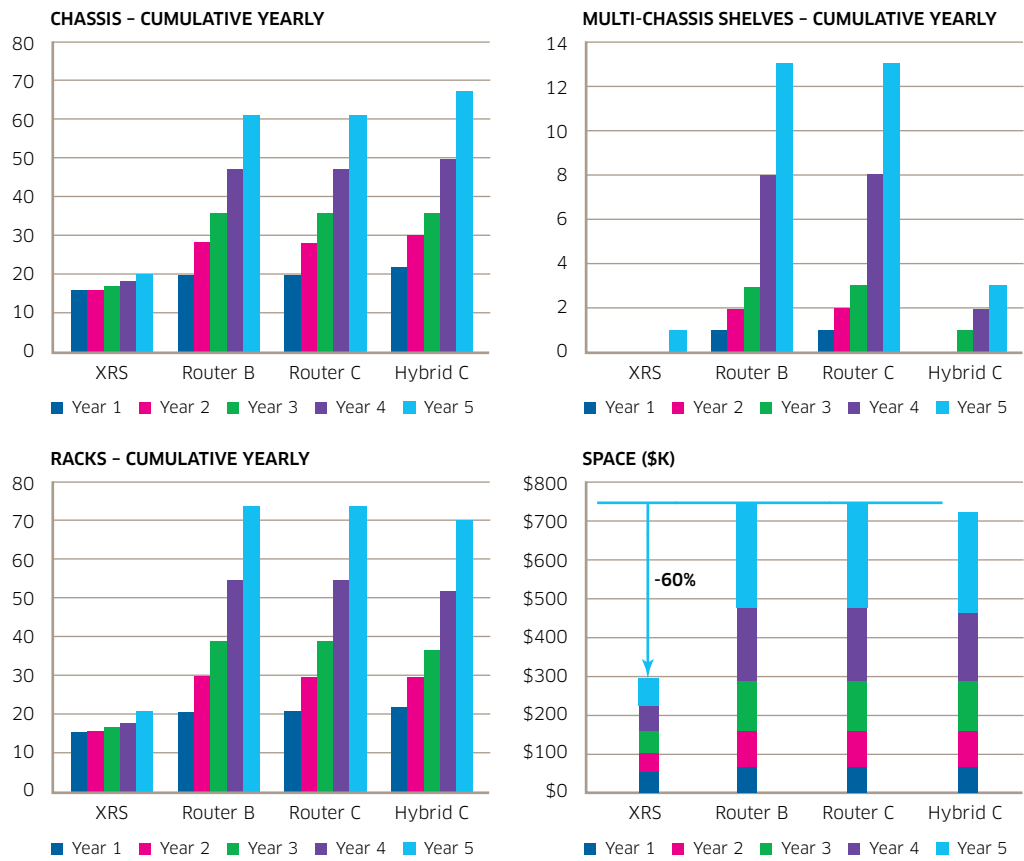
Rack space

The first dimension compares the number of chassis and hence the resulting rack space needed in order for the service provider to meet core traffic demands over a five-year period of network growth. Real estate to house network equipment comes at a premium, and savings in rack space translate into long-term savings on facilities costs. It is not easy to allocate more space for additional chassis because most central offices are already crowded. It is also not possible to put a core router wherever there happens to be surplus space. Operation, maintenance and installation (cabling) impact equipment placement, and therefore require careful consideration. This also requires planning for future space requirements and their associated costs.

Figure 6 shows the number of chassis required on the top left, while the need for multi-chassis switching shelves is shown at the top right. Due to its much higher port density, the 7950 XRS can scale within a single chassis in most locations by simply adding line cards. Only in one location (i.e. New York) would the 7950 XRS require a multi-chassis configuration to accommodate traffic in Year 5, assuming that no higher density line cards become available before that time. For site planners this is an ideal scenario as the 7950 XRS occupies by far the smallest amount of real estate. However, flagship core routers from Suppliers B and C both require significant chassis growth over time, and eventually resort to multi-chassis configurations for all but the smallest nodes.

From initial deployment to the complete Year 5 roll-out, traditional routers must more than triple the number of chassis to keep up with traffic demand. In this context the hybrid solution of separate P-router and LSR offered by Supplier C did not demonstrate any scaling advantages in rack space, though the LSR-only product defers the need for multi-chassis P-routers in most node locations.

Figure 6. Rack space usage and cost




Source: Bell Labs Analysis

The bottom right shows the translation of rack space into dollar savings. Competing solutions are 150 percent more expensive in terms of real estate costs as compared to the 7950 XRS. At an estimated \$3000/rack/year and five percent yearly inflation, the cumulative savings of the 7950 XRS amount to \$400,000 over five years. This does not include extra costs associated with freeing space to accommodate chassis growth in the case of Suppliers B and C.

Power

The next point of comparison is power consumption. Lowering power consumption contributes environmental benefits, and yields substantial cost savings too. There are practical limits to the amount of power that can be delivered to a physical location due to building safety regulations. Power efficient equipment saves both on the power it consumes as well as on the power required to cool the equipment. Referring back to Table 1, it is evident that line cards are responsible for roughly half the total system power consumption. To calculate power consumption the model took into account how many cards were actually required and used in each chassis to carry the required traffic load.

Figure 7. Power consumption and cost



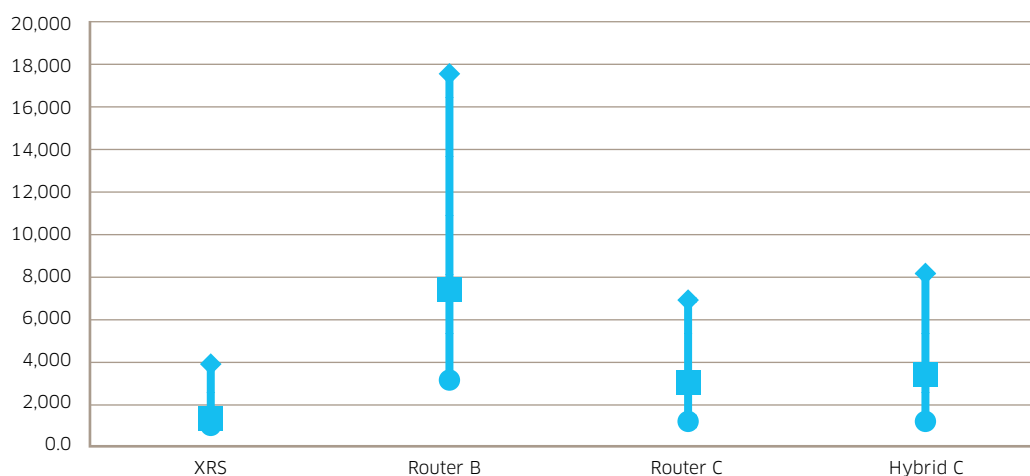
Source: Bell Labs Analysis

Figure 7 compares power consumption and savings. As with rack space usage, competing core routers from Suppliers B and C are virtually identical in power use, while the 7950 XRS has a sizeable 50 percent efficiency advantage. At \$0.125/kWh and 5 percent yearly inflation, this 7950 XRS advantage amounts to a savings of more than \$1.5 million over the study period compared to implementations of Suppliers B and C. Although this does not include additional savings in cooling, it can be assumed that savings in cooling are proportional to the savings in power consumption.

Weight

Weight is a limiting construction factor because a floor can only bear so much physical weight. Excessive weight may therefore limit the full use of available floor space. Extremely heavy equipment may even be limited to installation at ground floor level only, even though a central office may have multiple floors available. Figure 8 gives the comparable weight of a small, average and largest configuration used in the network under study.

Figure 8. Weight comparison




Source: Bell Labs Analysis

Capacity headroom

Headroom (capacity for growth) is the relative amount of reserve capacity available before additional chassis are required. Headroom in deployed systems is attractive because it allows the operator to quickly add more capacity when needed, at the lowest cost (just add more line cards). Figure 9 gives an overview of the headroom per configuration in the model. Green areas depict greater than 50 percent spare capacity remaining, yellow indicates 25-50 percent remaining, and orange shows less than 25 percent chassis capacity left. The 7950 XRS is the only system that provides ample headroom throughout the five-year study period.

Figure 9. Headroom



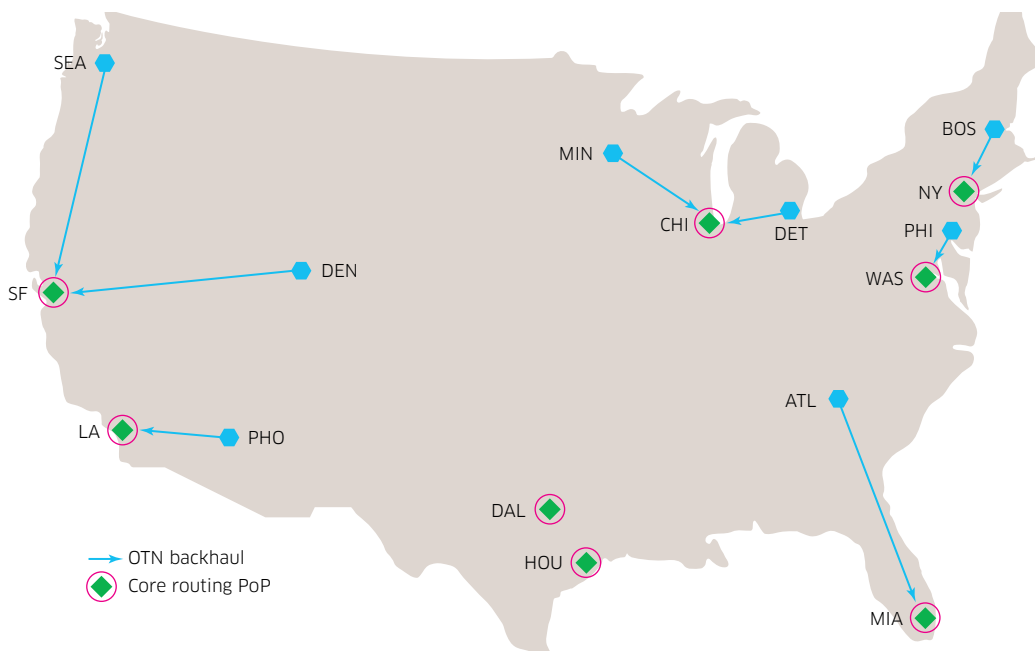
Source: Bell Labs Analysis

NETWORK CONSOLIDATION SAVINGS

The tremendous scalability of the 7950 XRS offers service providers sufficient headroom to stay ahead of traffic growth for many years to come, while keeping to a minimum the recurring operational expenditures related to space, power and cooling. In addition it opens new possibilities to consolidate the core network into fewer larger locations and leverage more cost-effective OTN equipment to aggregate and backhaul traffic from smaller “outer core” locations to the larger “inner core” nodes.

Consolidating a core topology in fewer but larger core routing nodes is an interesting contrast to competing strategies that essentially keep the same network footprint and attempt to save costs by removing Layer 3 capabilities from routers for the inner core (note that the 7950 XRS can also be configured as LSR-only). The potential cost of not having packet grooming in edge locations without a (P-) core router should be limited, as at this stage in the network at least three grooming stages have occurred over an aggregate of three to four million sessions.

Figure 10. Consolidated core topology



Source: Bell Labs Analysis

To determine the potential cost savings of network consolidation we reduced the number of core routing nodes from sixteen to eight. OTN equipment is assumed to backhaul Provider Edge traffic from smaller core locations to the larger core routing nodes (blue arrows in Figure 10). The cost of the OTN layer is not factored in, but this cost is assumed to be neutral across alternative solutions for this study. All the other assumptions, such as traffic volumes, match the 16-node model. The analysis that follows again focuses on direct cost savings related to lower space and power usage. Potential savings due to the reduced operational complexity of managing a smaller core network topology are not factored in. In addition there may be slight performance improvements since traffic will on average see fewer core routing hops, while link transmission speeds will be higher on average.

CAPEX efficiencies

Due to the meshed nature of core routing networks, the number of links and relative transit traffic grow with the number of nodes. When fewer core routing nodes can be deployed, there are also fewer links required to interconnect these core routers, while the average capacity on these links goes up; and cost per bit on a high speed link (e.g., 100 Gb/s) is lower than that on a 10 Gb/s link.

Figure 11 compares the port requirements for the 16-node topology (top) and 8-node topology (bottom). While the total number of access ports remains the same (row in beige), the number of network ports (P-P trunks) shown is a relative drop of about 90 percent (row in green). In other words, by reducing the number of nodes in a network topology there are concrete savings on network ports.

Figure 11. Comparing 16- and 8-node network port requirements

TOTAL (16 NODES)

	YEAR 1	YEAR 2	YEAR 3	YEAR 4	YEAR 5
PE-P ports (10GE)	672	864	1152	1632	2208
P-P ports (100GE)	192	256	352	496	672

TOTAL (8 NODES)


	YEAR 1	YEAR 2	YEAR 3	YEAR 4	YEAR 5
PE-P ports (10GE)	672	864	1152	1632	2208
WRT 16 node	100%	100%	100%	100%	100%
P-P ports (100GE)	168	232	320	440	608
WRT 16 node	88%	91%	91%	89%	90%

Source: Bell Labs Analysis

Space

Figure 12 compares space requirements of the 16- and 8-node topology and shows that having fewer but larger nodes results in 50 percent fewer chassis for the 7950 XRS-based deployment, without any change in multi-chassis configuration needs. Competing core routers only see marginal savings on the number of chassis required.


Figure 12. Chassis requirements of 16- versus 8-node topology



Source: Bell Labs Analysis

Figure 13 translates these differences in dollar values. It is surprising that incumbent routing platforms do not show any reductions in chassis required in the initial deployment. While the 7950 XRS cut by half the number of chassis required, the incumbent routing platforms still require the same number, but they are now concentrated in eight locations rather than sixteen. While all platforms eventually gain from a reduced topology, the results demonstrate far better economics for scaling 7950 XRS-based deployments, with a 40 percent drop in space costs.


Figure 13. Comparing rack space usage and cost of 16- and 8-node topology



Power

Similar cost savings are obtained with respect to power consumption. Again we see the 7950 XRS-based deployments benefiting twice as much (30 percent less instead of 10-15 percent less) from a consolidated node topology (Figure 14).

Figure 14. Power consumption



Capacity headroom

The last remaining question is how a 50 percent reduction in the network topology impacts utilization results. Insufficient headroom for growth is a planning concern because each time more capacity is required additional chassis must be deployed, resulting in higher costs for space and power. Most operators will consider placing a new chassis once a utilization of 80 percent is reached.

Figure 15. Capacity headroom - 8-node topology




Figure 15 shows capacity headroom for the 8-node topology. The pattern for the 7950 XRS has changed slightly; utilization is a bit higher (some green fields changed to yellow) but there is still ample headroom for graceful capacity growth within the confines of the initially deployed chassis footprint. Competing solutions that provide integrated routing now see a consistently higher utilization across the board (orange depicts less than 25 percent spare capacity).

CONCLUSION

The Alcatel-Lucent 7950 XRS is a purpose-built core routing platform designed to efficiently meet the scaling needs of the 100G era and beyond.

This business case demonstrates how the industry-leading system capacity and port density of the 7950 XRS, in combination with a perfect system geometry that enables 100 percent efficiency for 10, 40 and 100G line rates, translates into concrete and considerable cost savings on space, power and cooling. The 7950 XRS scales cost effectively in a small footprint with minimal overhead costs, as opposed to flagship incumbent platforms that resort to premature expansion to more costly, cumbersome and power-hungry multi-chassis configurations.

Further, in combination with the Alcatel-Lucent 1830 Photonic Service Switch, the XRS allows service providers to consolidate and simplify their core routing networks to reap even more cost efficiencies through network designs that optimize the IP and optical domains of their transport infrastructure. The 7950 XRS provides ample headroom for growth throughout the deployment lifecycle, with full flexibility to support high-scale IP routing or MPLS switching in a single platform on common hardware.

Relative to existing alternatives, the 7950 XRS offers superior scale, efficiency and versatility.

ACRONYMS

CAPEX	Capital Expenditures
DWDM	Dense wave division multiplexing
HLN	High Leverage Network
LSR	Label Switch Router
MPLS	Multi-Protocol Label Switching
NPAT	Network planning and analysis tools
OPEX	Operational Expenditures
OTN	Optical transport network
PE	Provider edge
PSS	Photonic Service Switch
TCO	Total Cost of Ownership
XRS	Extensible Routing System

REFERENCES

Please refer to www.alcatel-lucent.com/solutions/ip-core-routing and www.alcatel-lucent.com/7950-xrs for more information.

