BUSINESS CASE FOR MOVING DNS TO THE CLOUD

STRATEGIC WHITE PAPER | NFV INSIGHTS SERIES

The cloud and NFV are being touted as the means by which service providers can improve their current operations costs and contend with the rapid growth in demand. With the help of a Tier 1 service provider, Alcatel-Lucent conducted a business case analysis of this theory. The study used the service provider's actual figures and plans for expanding its DNS capacity to determine how a cloud-based future mode of operation would compare with the traditional present mode of operation. The study revealed how even a simple application like DNS can benefit enormously from running on an NFV platform. Processes such as scaling, software upgrading and healing are greatly simplified and infrastructure expenditures are reduced. These OPEX and CAPEX savings significantly lower the service provider's TCO and increase its agility – critical improvements for today's challenging telecom environment.

About the NFV Insights Series

NFV represents a major shift in the telecommunications and networking industry. NFV applies virtualization and cloud principles to the telecommunications domain, something that appeared to be impossible until recently due to the stringent performance, availability, reliability, and security requirements in communication networks. Many service providers are now keen to implement NFV to help them become more agile in delivering services, and to reduce equipment and operational cost. This series of whitepapers addresses some of the key technical and business challenges on the road to NFV.

TABLE OF CONTENTS

1. A WORD ON NFV AND INNOVATING IN OPERATIONS

In the current competitive landscape where windows of opportunity are rapidly shortening, gaining agility has become a key success factor. Companies like Google®, Amazon® and Facebook® are leading the way with a strong focus in operations innovation fueling their competitive advantage. Telecom legacy architectures operating in "siloes" have not helped service providers to adapt to the new market rules. However, the good news is that the new paradigm brought by Network Functions Virtualization (NFV) provides a unique opportunity to both catch up and prosper over the long term.

Much has been said about how virtualization and the cloud may be applied to the telecom world in order to improve the total cost of ownership (TCO) of current infrastructures and operations. Initial discussions analyzed the advantages that virtualization would bring in terms of hardware optimization. Lately, the focus has shifted to the operations field. While it is undeniable that virtualization of some network functions will bring savings in terms of CAPEX, NFV's greatest contribution is going to be that it enables a new way of approaching telecommunications. This means much more than just optimizing inefficiencies inherent in the current processes. Service providers can — and should — leverage NFV technology to redefine their current operations. Such an endeavor will require three major steps:

1. Map out every current process in detail
2. Analyze what can be automated (that is, handled by an NFV platform) to reduce complexity
3. Redesign operations to be much simpler and more agile

This white paper looks at the results of a study performed by Alcatel-Lucent's Cloud Consulting team in cooperation with a Tier-1 service provider. It compares the TCO of running Domain Name Server (DNS) operations in the service provider's present mode of operations (PMO) versus migrating to an NFV-enabled future mode of operation (FMO) model. Specifically, the service provider was planning to replace its existing infrastructure, which was reaching end-of-life. The company wanted to understand how migrating to an NFV model would improve its TCO compared to replacing older servers with new x86s.

The results of this study reveal how even a simple application like DNS can benefit enormously from running on an NFV platform. Processes such as scaling, software upgrading and healing are greatly simplified, which increases agility and significantly lowers TCO. Hence, migrating to NFV provides a clear competitive advantage.

2. MAIN DRIVERS FOR A DNS NFV BUSINESS CASE

Many parameters may be considered when developing a business case to analyze the impact of migrating an application like DNS to run on an NFV platform such as Alcatel-Lucent CloudBand™. Alcatel-Lucent's consulting team looked at the main drivers in order to shed some light on the magnitude of savings and agility gains that can be expected.

Figure 1 shows the main aspects that have been addressed in the DNS NFV Business Case study.

Cost drivers were clustered into three categories:

- **CAPEX**: one-time investments in fixed assets with a useful life extending beyond the taxable year
- **OPEX infrastructure**: ongoing costs directly related to the infrastructure (for example, maintenance)
- **OPEX processes**: ongoing staffing costs directly related to the daily management of activities or processes required to provide DNS services

Figure 1. Main NFV cost drivers addressed in the business case study

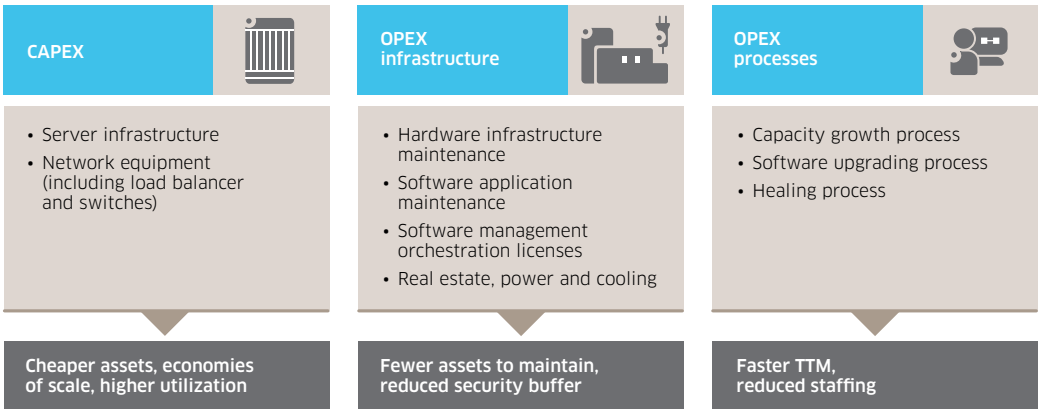

3. DNS MIGRATION SCENARIOS

The study team's analysis focused on comparing two scenarios. The first scenario consisted of migrating to new x86 servers and continuing to follow the PMO. The alternative scenario involved migrating to a cloud infrastructure and introducing an NFV-based FMO.

Present mode of operations

DNS operations included a total of 104 servers spread across 11 sites, with six different DNS applications running on them: consumer Cache Resolver (CR), consumer Authoritative Name Server (ANS), business CR, business ANS, other CR and other ANS. As shown in Figure 2, the service provider's initial plan involved grouping the servers in two blocks.

Figure 2. Five-year infrastructure replacement and growth plans

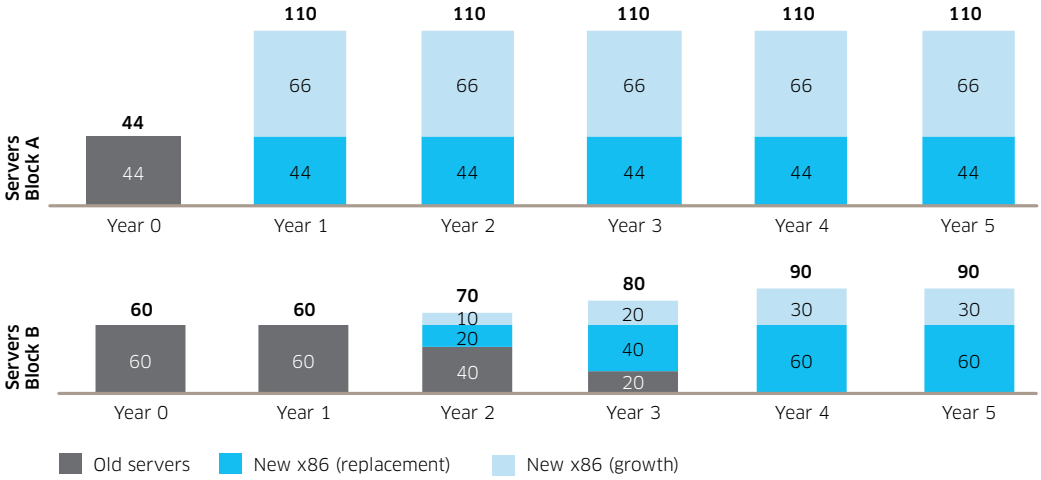

Block A servers, which includes 44 consumer CR servers, would all be replaced in year 1 and at the same time scaled up to 110 servers in total. The service provider expected this number of servers to consume all the DNS staffing resources usually available to manage this process on a yearly basis.

Block B servers, which includes the remaining five DNS applications, presented a greater replacement and migration process complexity. In fact, the service provider expected that replacing 20 Block B servers and adding another 10 in one single year would require a similar effort to that required for all Block A servers in year 1. Consequently, the plan was to replace and scale capacity of Block B servers in three batches over years 2 to 4.

Future mode of operations

The alternative FMO consisted of migrating existing DNS servers to a cloud infrastructure based on 11 small CloudBand Nodes (132 servers), one per existing DNS site. This scenario involved different deployment and scaling plans. First, all physical servers would be deployed in year 1 instead of the gradual deployment used in the PMO. Second, all existing DNS applications would be migrated in year 1 and would be virtually scaled in the subsequent year according to the growth in traffic.

The FMO looked at two different scenarios: dedicated and shared. The dedicated scenario assumes that the underlying cloud infrastructure will only be used for DNS. Thus, all CAPEX and OPEX infrastructure costs are allocated to the DNS deployment regardless of the capacity utilization. The shared scenario takes into account that the underlying infrastructure may be shared among several applications in a dynamic way. Consequently, a cross-application organization will likely be in charge of managing and providing capacity on demand, and will charge application owners on a pay-per-use basis.

4. CAPACITY GROWTH PROCESS

The study analyzed the lead time and staffing effort required for both the initial deployment and the increased capacity that would be required for DNS traffic over a 5-year period.

According to the PMO, the service provider would follow a four-step process (Figure 3) to deploy a new server infrastructure for DNS. First, a planning phase is kicked off with a site survey, followed by a high-level design and test plan. In parallel to the planning phase, the service provider can trigger the hardware and DNS application licenses PO process in order to reduce wait time, especially for the hardware delivery.

Once the PO is issued, the service provider needs to wait for a number of weeks until the hardware arrives. The third phase starts with the arrival of the hardware. In this specific case, the service provider receives hardware at a central location where its COTS experts perform assembly and cabling activities to create a number of "racks". Each of these racks will require some networking activities before the servers can be provisioned and the DNS applications installed. Once all of these actions are completed, the service provider has a number of "shippable racks", ready to send to each of the sites. Once the racks are at their destinations, the final cabling, networking and DNS configuration is performed.

Figure 3. PMO deployment process

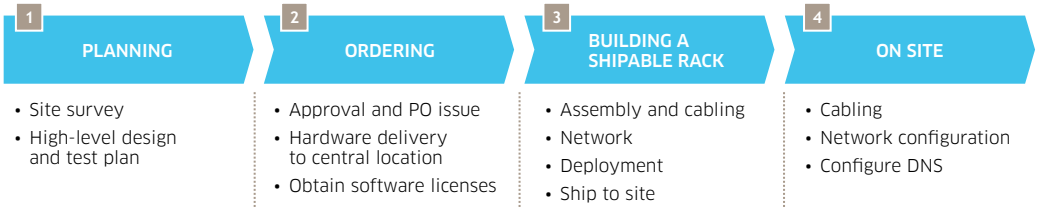

The FMO deployment process (Figure 4) differs from the PMO process in a number of ways. The planning phase remains the same. However, ordering becomes a bit longer since hardware arrives pre-assembled and pre-cabled by the vendor. This approach streamlines the assembly and cabling processes, and reduces both total lead times and staffing costs. (In most cases, the vendor's staff will take less time to assemble and pre-cable the services than the service provider's staff.) Also, application deployment can be pushed to the next phase and performed directly at the sites. Consequently, the service provider can skip the third phase and have the vendor ship the nodes directly to their final destinations. This provides two additional advantages. First, all networking is done in the same location, thus streamlining the process. Second, the considerable time dedicated to unpacking components, repackaging the racks and sending them to each of the sites is not required.

Figure 4. FMO deployment process

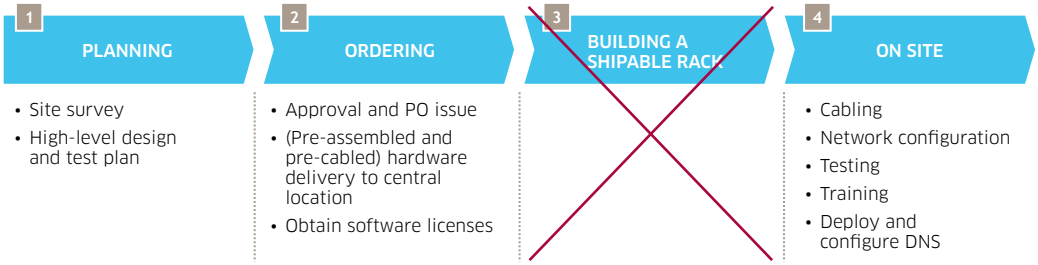

Once the nodes are on site and connected to the network, they are ready to be commissioned. Cloudband's automated node provisioning software suite does this within a few clicks. The service provider's staff needs to spend only a few minutes inputting key parameters in a web interface and clicking "provision". The node then automatically provisions itself while the service provider's staff focuses on something else.

After about 3-5 hours, depending on the size, the node is ready for application deployment. In the new NFV context, vendors will likely provide generic recipes for their applications that will be able to run on any platform. However, the operator will still need to perform some minor tuning to adapt the recipe to the service provider's local configuration.

Lead-time expectations for deployment phases

Figure 5 explains in a Gantt chart the lead-time expectations for both the PMO and the FMO in year 1. According to the numbers provided by the service provider for this specific deployment, a total of 199 days are needed every year to complete the new infrastructure implementation. CloudBand improves new physical deployment in year 1 by 2 percent, which is a minor difference. However, in the PMO, the service provider migrates DNS applications to the new infrastructure as the old hardware is replaced between year 1 and 4. In the FMO, it replaces all the infrastructure in year 1 and thus all applications are migrated as well. Hence, much more is achieved in about the same time.

As shown in Figure 6, the service provider can also realize substantial lead-time savings in subsequent years. Between year 2 and year 4, the lead time to scale the DNS service is drastically reduced since the physical deployment process is replaced by virtual scaling. In the PMO the service provider follows the same physical process each year, while in the FMO most tasks are automated. The planning phase becomes much shorter as site surveys only require minimal human intervention to verify capacity within the CloudBand Management System. Service providers also see a significant reduction in lead times when reordering as only software licenses have to be ordered. Even this cost can be further streamlined if a flat fee is negotiated. In addition, network configuration only requires verification of cloud networking availability (such as IP addresses) while application deployment becomes much faster with automation.

Figure 5. Lead time to replace and grow DNS services by adding physical servers

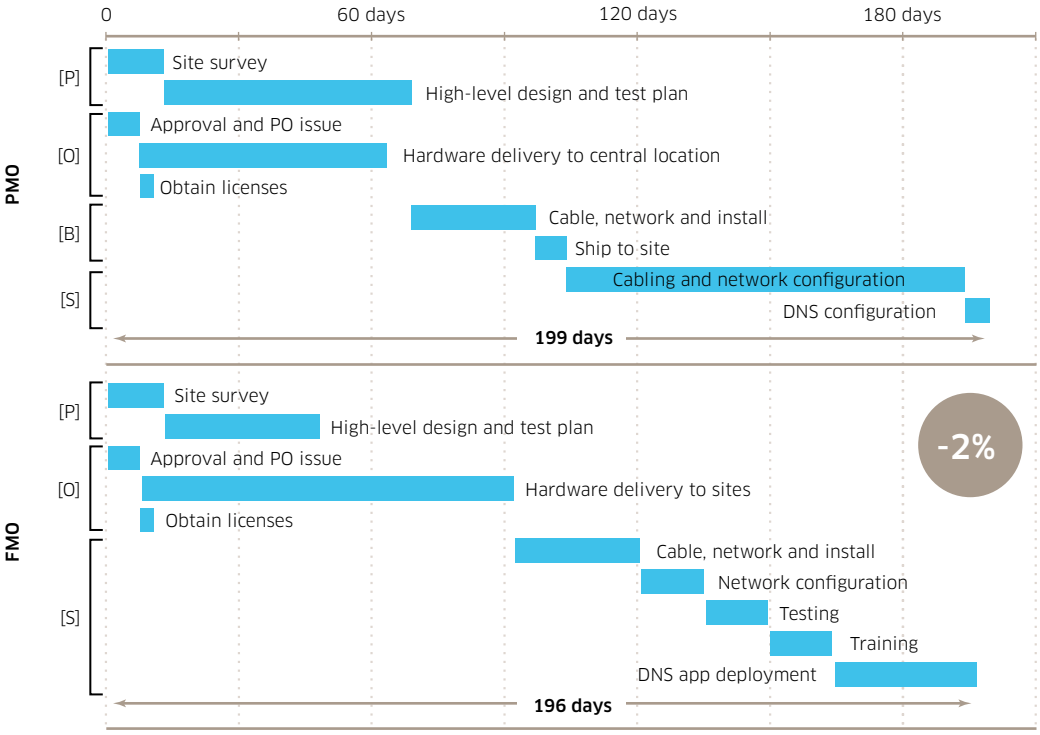

Figure 6. Lead time to scale DNS services by adding VMs

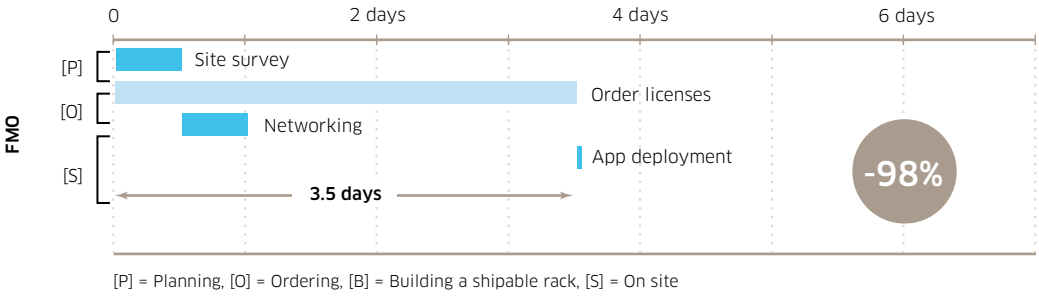

Deployment cost savings

The final numbers in Figure 7 show that the FMO can deliver considerable savings. Costs in year 1 are slightly higher in the FMO due to the additional cost for professional services incurred at initial deployment. However, this is a one-time cost attached to the introduction of the new infrastructure (cloud node provisioning, training). As the service provider becomes more familiar with the infrastructure, it will likely use its own operators and perform these tasks in house for future deployments. Note that in an NFV world, applications share the infrastructure, so the service provider's operations team will only need to be familiar with a very limited number of selected infrastructure elements.

1 All monetary units in United States dollars

In years 2 to 4 of the FMO, total server replacement and growth process costs are significantly reduced. On average this means 44 percent (from $0.66 million to $0.37 million) over a 5 year period. Clearly, CloudBand's virtual scaling and automated application deployment capabilities reduce capacity growth process costs significantly.

Figure 7. Cost to scale DNS services

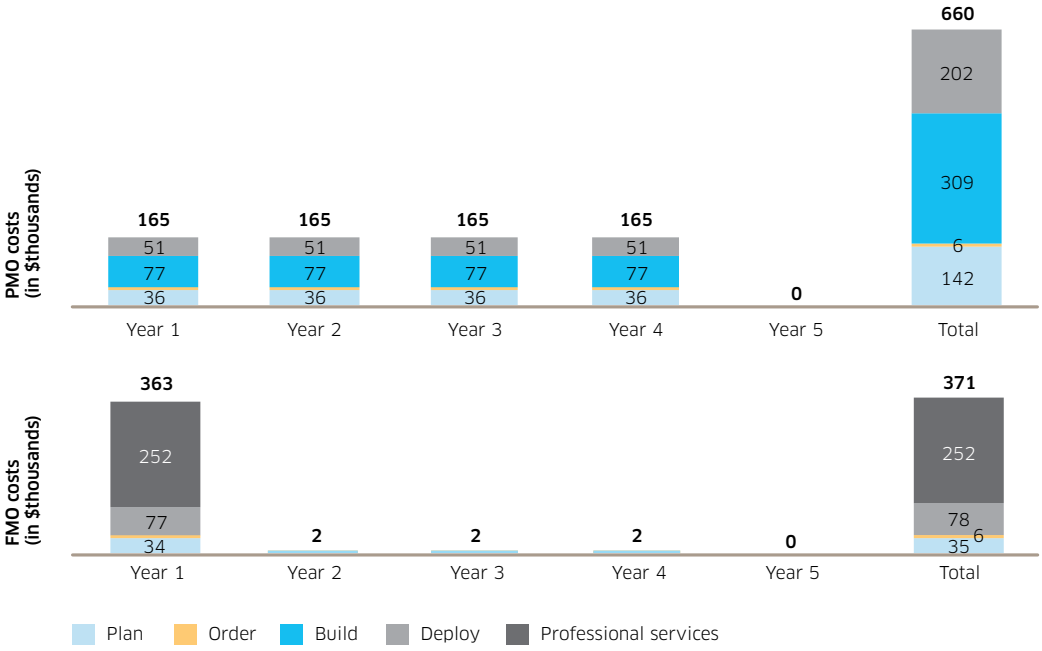

5. SOFTWARE UPGRADING PROCESS

Today, software upgrading of both new programmed releases and ad hoc patches can be a long and tedious process. Typically, the cycle follows four phases:

- Plan
- Obtain the new software
- Test the new software
- Install and configure

The last phase generally consumes the most time and resources.

The starting point to an analysis of any software upgrading process should be to look at the number of upgrades required on a yearly basis. For the study's service provider, six DNS license upgrades are generally subject to analysis per year by the operations team. The team usually approves four of the six for testing and deployment. Furthermore, the team also commonly carries out one operating system (OS) upgrade per year.

As shown in Figure 8, the DNS operations team currently dedicates, per year, approximately:

- 22 days to planning activities
- 10 days to obtaining new software
- 25 days to testing
- 133 days for installation and configuration

This business case assumes that both planning and obtaining the software will not change with the introduction of an NFV platform. However, testing becomes much simpler. NFV offers a reduced timeline and lower costs for test staging and environment creation as service providers can take advantage of "sandbox" testing environments without the need for dedicated equipment. The assumption is that this will enable simplified creation and parallel execution of test cases, and reduce by about one third the total time required for testing.

Comparing installation and configuration

NFV makes installation and configuration much simpler. Traditionally, service providers open maintenance windows at night to install and configure a predefined number of servers one by one. With the FMO, the service provider is able to upgrade four servers per night in a 5 hour maintenance window. The number of windows required depends on the total number of servers that can be upgraded per night by the operations team and the total number of servers required for the DNS deployment. As Figure 8 shows, the lead time required for maintenance windows grows over time as the service provider increases the number of physical servers to keep up with growing traffic needs.

Figure 8. Lead times to upgrade DNS/OS software

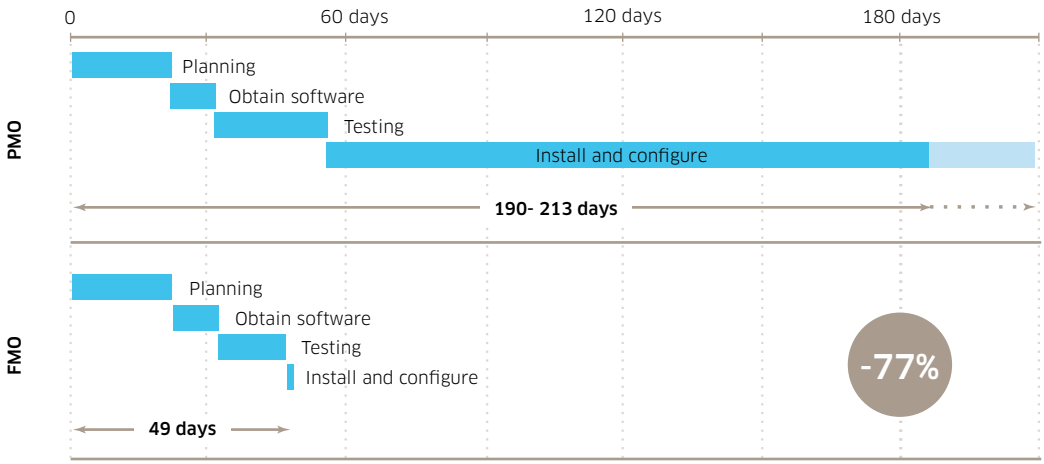

With NFV the whole process changes. The total number of servers is not relevant any more for installation and configuration. CloudBand leverages application recipes to push upgrades automatically, in a matter of minutes, to all servers in parallel. With NFV, application vendors normally provide the application's recipe and the service provider only needs to customize the recipe based on its specific configuration needs. Thus, the advantage is that the manual customization only happens once. The upgrade for the rest of the servers is automated.

This automation provides dramatic gains in agility. While the PMO requires 190 days in year 1 and up to 213 days by year 5, the FMO requires only 49 days to accomplish the full process. Hence, by migrating its DNS deployment to CloudBand, the service provider experiences a 77 percent reduction in lead time.

Cost comparison

In terms of costs, Figure 9 shows that in total, CloudBand reduces software upgrading costs by 83 percent from $0.69 million to $0.12 million, most of which comes from the automation achieved in the installation and configuration phase.

Figure 9. Software DNS/OS upgrading process costs

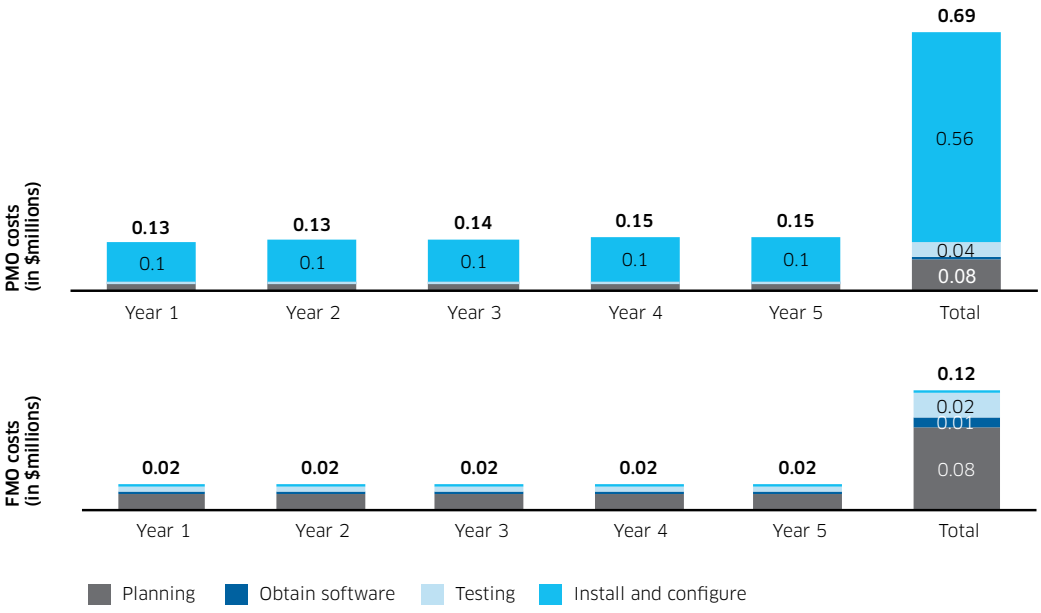

6. HEALING PROCESS

Device failures are a major issue for service providers, since they can result in loss of service for many users and can increase churn. To reduce this risk, service providers traditionally deploy fully redundant architectures. This security buffer comes at a high price, given that the investment requires double the amount of physical infrastructure, with much of it standing idle.

A device failure is not the only issue that can require a healing process. Service providers also need to be able to address OS failures, application failures and Distributed Denial of Service (DDOS) attacks.

PMO healing process
With the PMO, the healing process consists of three stages:

- Issue identification
- Trigger and execute solution process
- 'Post-mortem' root cause analysis (RCA)

As shown in Figure 10, lead times to identify and solve a problem vary depending on the issue at hand. Identifying the issue tends to be simpler and faster at the lower layers (the hardware and operating system layers), whereas actually solving an issue tends to be faster at the application layer. DDOS attacks are the fastest both to identify and solve. However, DDOS problems tend to consume more of the operations team's time, given that a DDOS attack is much more likely than the other three potential issues.

RCA is the final task performed by operators as soon as service continuity has been assured. By identifying the root cause of a problem, the service provider can make whatever changes are necessary to avoid reoccurrences.

Figure 10. Lead times for issues subject to healing process (PMO)

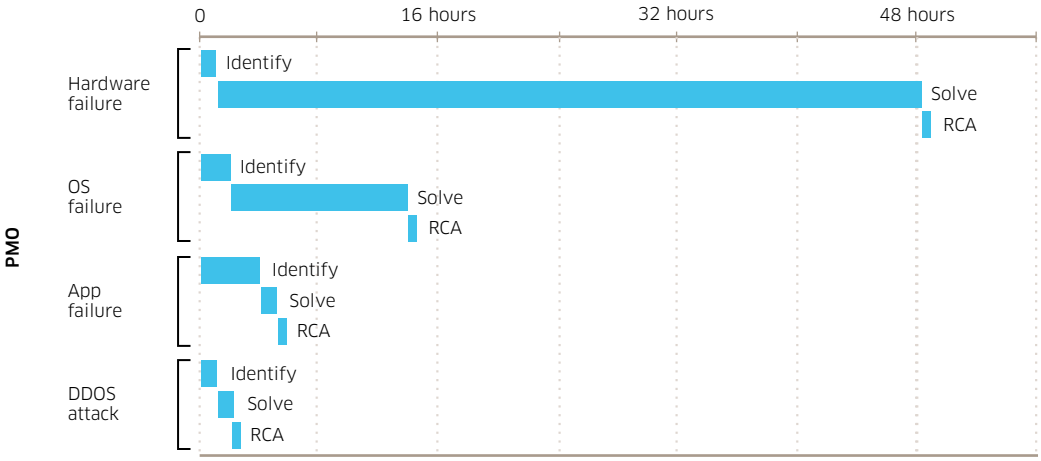

FMO healing process

With NFV, devices run as virtualized functions and are protected by the self-healing properties of the hypervisor and orchestration layer. The healing process is fully redefined as the business continuity process is decoupled from the problem itself. To provide end-to-end application resiliency and reliability, NFV platforms incorporate mechanisms for automated healing based on the monitored infrastructure and application-level KPIs. Should a failure occur, the system automatically creates a new instance with the same specifications to ensure application availability at all times.

Figure 11. Lead times for issues subject to healing process (FMO)

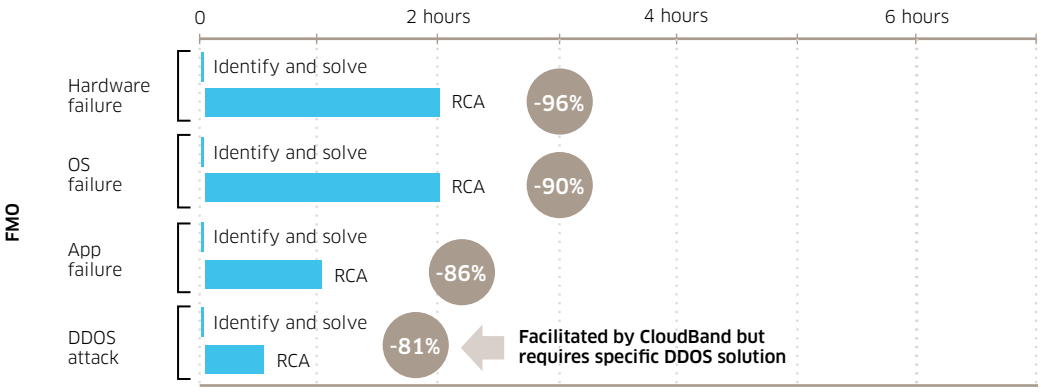

CloudBand lifecycle management automates the first two stages for hardware, OS and application failures. In the event of any of these three issues taking place, the system automatically identifies the problem and spins up a new instance of the failed virtual machine(s), making sure that business continuity is maintained.

The DNS operations team still needs to conduct a post-mortem RCA to ensure a problem does not reoccur. An RCA takes slightly longer in the FMO than in the PMO as identifying the cause of failure does not leverage knowledge gained in prior field analysis. However, overall the study identified lead-time reductions between 86 and 96 percent in all three cases. Furthermore, tools for RCA are under constant development and the lead time required to complete the process can be significantly reduced once the platform is able to provide a detailed description of the failure.

An NFV platform may bring additional benefits when it comes to handling DDOS attacks. Defense against these attacks requires a rapid response with packet filtering and deployment of additional capacities, so it is advisable to automate this process. Virtual scaling capabilities facilitate the development of a specific DDOS solution that would automate the process with very limited effort. According to our estimates, such a solution would allow service providers to reduce DDOS healing lead time by 81 percent.

Cost comparison

In terms of costs, Figure 12 shows that hardware, OS and application failures account for 53 percent of total PMO healing costs. With 50 percent of the ANS servers being attacked once every month, DDOS attacks account for the remaining 47 percent. CloudBand initially provides a saving of $0.38 million due to simplification of the healing process of the first three issues. An additional saving of $0.34 million could be achieved by developing a simple solution that leverages automated virtual scaling capabilities. The total 5-year healing process costs may decrease from $0.83 million to $0.12 million, an 86 percent reduction!

Figure 12. Healing process cost reductions

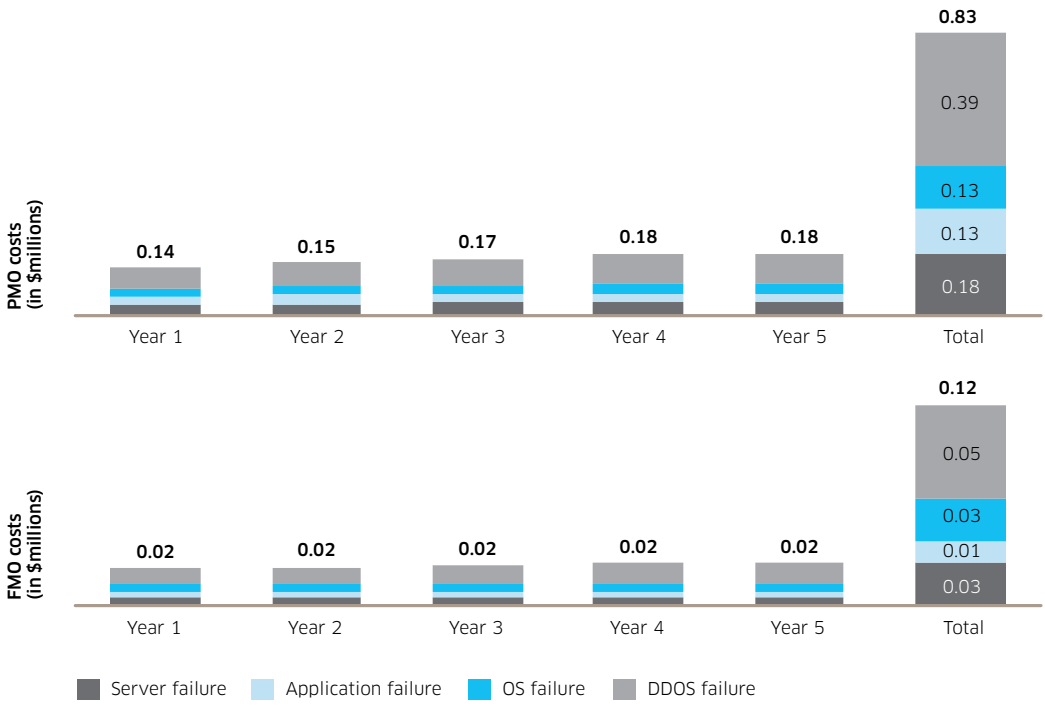

7. FLOOR SPACE, POWER AND COOLING

Real estate, power and cooling are OPEX infrastructure costs, that is, they are directly related to the number and characteristics of physical infrastructure items to be managed by the operations teams for a specific deployment. Consequently, provided that all constants remain equal, a lower number of physical hardware infrastructure items will result in a reduction in the total costs of real estate, power and cooling in the same proportion.

The main drivers considered when analyzing these costs are:

- **Real estate**: number and size of infrastructure items and square foot cost
- **Power**: rate of energy consumption and cost per kilowatt hour
- **Cooling**: a factor of 1:1 of power consumption

Figure 13 shows the real estate, power and cooling costs for the DNS deployment described in Section 3 (Figure 2). Clearly, CloudBand enables a significant cost reduction in all three categories.

Figure 13. Cost comparison for floor space, power and cooling

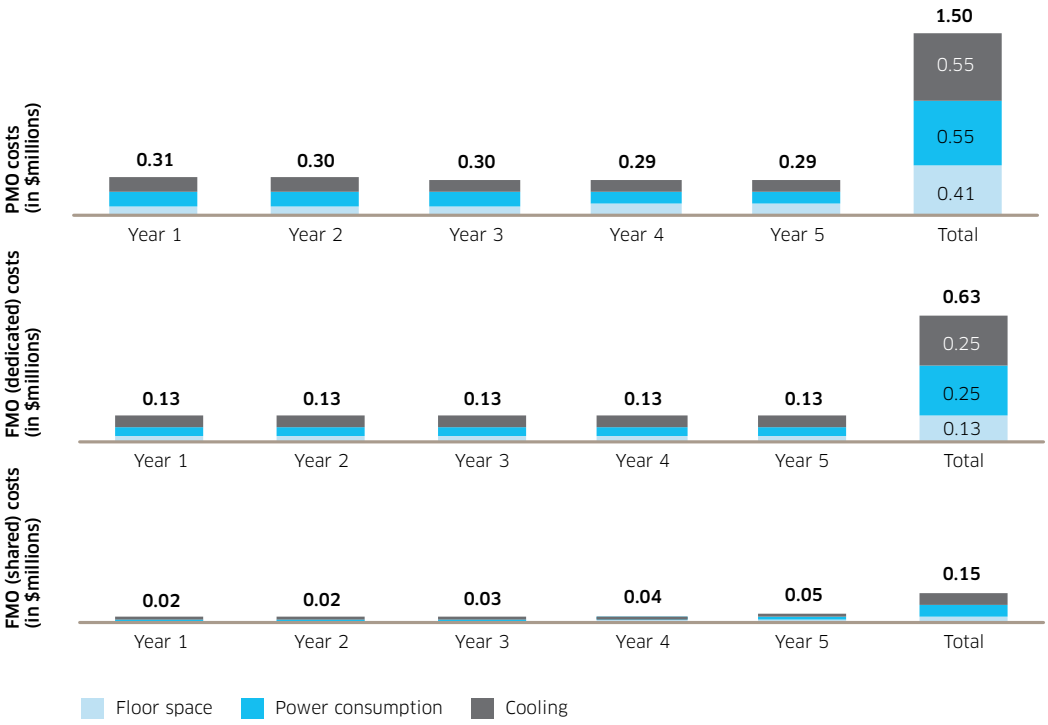

Real estate costs are reduced because the FMO requires fewer physical infrastructure items. In the PMO, load balancers and other networking equipment such as switches are placed separately from servers. In fact, the service provider stated it needed one rack for this additional equipment for every rack of servers. In the FMO, switches are integrated within the node and load balancers are virtualized, running on the same infrastructure as the DNS applications. Hence, though a CloudBand Node is slightly bigger than the DNS server racks (8.40 ft^2 vs. 7.74 ft^2), there will still be much less floor space required for the infrastructure as considerably fewer nodes will be required compared to the number of racks.

Power costs are reduced in the FMO because CloudBand makes it possible for the service provider to replace older servers sooner. A faster and simpler physical scaling process allows operators to migrate all existing infrastructure in year 1. While these savings could potentially also be achieved in the PMO by deciding to migrate all old servers in year 1, the service provider stated that it would not be able to manage a process of such complexity in one year with the current resources dedicated to DNS.

Older servers consume about twice as much energy as the new ones in the PMO (5 kW vs. 10 kW per rack of 10 servers). The CloudBand Nodes also consume slightly less power than the alternative x86 server racks (4.4 kW vs. 5 kW). Consequently, the bulk of the cost reduction comes from replacing old servers (all in year 1 in the FMO vs. gradual replacement until year 4 in PMO). The lower consumption of CloudBand Nodes compared to x86 racks adds a minor additional saving.

Lastly, cooling costs are generally calculated as a 1:1 ratio to power costs, hence, cooling costs decrease in the same proportions as encountered with power.

Looking again at the costs in Figure 13, the FMO differentiates cost scenarios: dedicated ($0.63 million) and shared ($0.17 million). The former, which constitutes a 58 percent reduction, is directly reflected in the bottom line of the company and is the result of comparing the total costs attached to such an investment. The latter, a 90 percent cost reduction, shows the new cost of running DNS. As explained in Section 3, the dedicated scenario allocates all costs to DNS, while the shared one allocates costs to DNS based on capacity utilization.

8. MAINTENANCE AND SOFTWARE LICENSES

Maintenance is also an OPEX infrastructure cost, that is, it is directly related to the number and characteristics of physical infrastructure items to be managed by the operations teams. In the case of a DNS deployment, the main infrastructure elements that require a yearly maintenance fee are the servers and the network equipment (including load balancers, switches and routing ports).

The software licenses category includes the costs of the CloudBand Management System licenses and maintenance fees. This is a new cost for the service provider, one that enables all the savings enumerated for the FMO. In fact, as Figure 14 shows, the extra cost of the CloudBand Management System and its maintenance ($2.64 million) is already offset by the savings in physical server and network equipment maintenance ($3.16 million).

Figure 14 differentiates the total costs of the dedicated FMO scenario with those of the shared FMO scenario, where only the used capacity is allocated to the DNS application. As mentioned in the previous section, the first cost ($4.64 million) reflects the impact on the bottom line of the company as a whole, while the second cost ($3.34 million) is that part of the expenditure that can be allocated to DNS alone. The difference is the cost of idle capacity ($1.3 million), which the service provider can allocate to other applications running in the system.

Figure 14. Software licenses and maintenance cost comparison

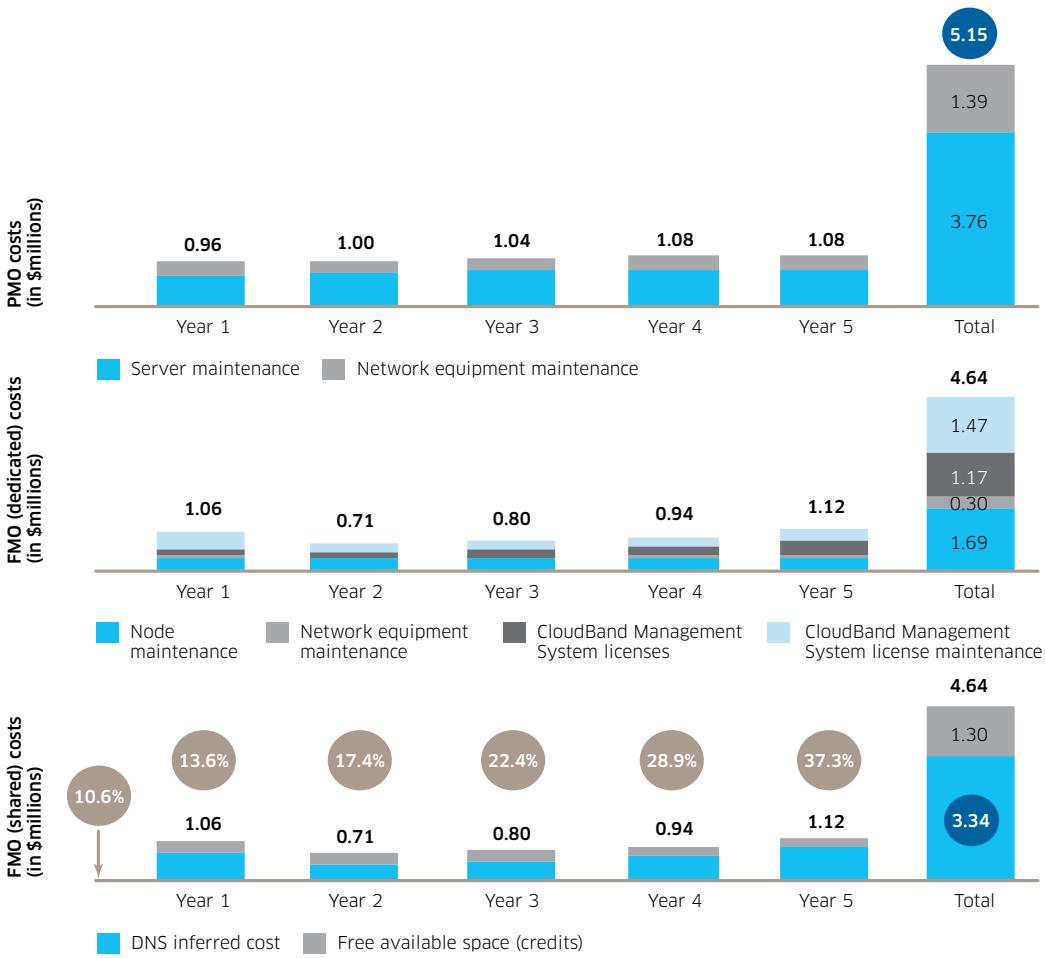

9. HARDWARE INFRASTRUCTURE

Virtualization of physical assets improves resource utilization by creating virtual machines, each with its own operating system on a single physical hardware asset. An NFV platform goes a step further. It enables dynamic placement of the virtual machines, which further improves hardware optimization.

The PMO DNS deployment operates in a “siloed” architecture as servers are dedicated to DNS. In fact, they are dedicated to only one kind of DNS application (the service provider was running six different applications for DNS). Given the need to split the DNS servers across 11 different sites, the result is a high number of servers running rather inefficiently in terms of capacity usage.

CloudBand enables a new model, where all underlying hardware forms a pool of resources that is shared by all the applications running on the same platform. Hardware optimization allows the service provider to reduce the number of servers from 200 in the PMO to 132 (in effect, 11 small CloudBand Nodes) in the FMO. The initial excess capacity provided by 11 small CloudBand Nodes to cover the expected traffic for a 5-year timeframe suggests that the number of nodes could be further reduced. However, topology requirements of the service provider demand that the infrastructure be divided into

11 different sites. Furthermore, the ability to share the infrastructure enables a new cost model in which the cost of idle capacity should not be allocated to DNS but to system capacity inefficiencies. The service provider can remedy these inefficiencies because the FMO makes it easy to identify available capacity, which other applications can use. As a result, the service provider may be able to avoid unnecessary investments in more physical infrastructure.

Looking at the numbers in Figure 15, the service provider can expect a significant reduction in server costs in the FMO since the CloudBand Nodes contain considerably fewer servers than proposed in the PMO. Nonetheless, the server cost only goes down by $0.03 million. CloudBand Nodes have integrated network switches so the switch cost is included in the Node cost — that is, the reduction in server costs is "hidden" by the reclassification of network switch costs.

A significant reduction in network equipment costs can be partially explained by the above-mentioned reclassification, but more importantly by the virtualization of load balancers, which eliminates the need for the physical equipment.

Again, the analysis provides two cost figures in the FMO: $1.81 million and $0.43 million. The former, which constitutes a 56 percent reduction, is directly reflected in the bottom line of the company and represents the total costs. The latter, an 89 percent cost reduction, shows only the new infrastructure cost allocated specifically to running DNS, that is, only the capacity that is used by DNS. The difference ($1.38 million) is the cost of idle capacity, which is available for running other applications.

Figure 15. Hardware Infrastructure cost comparison

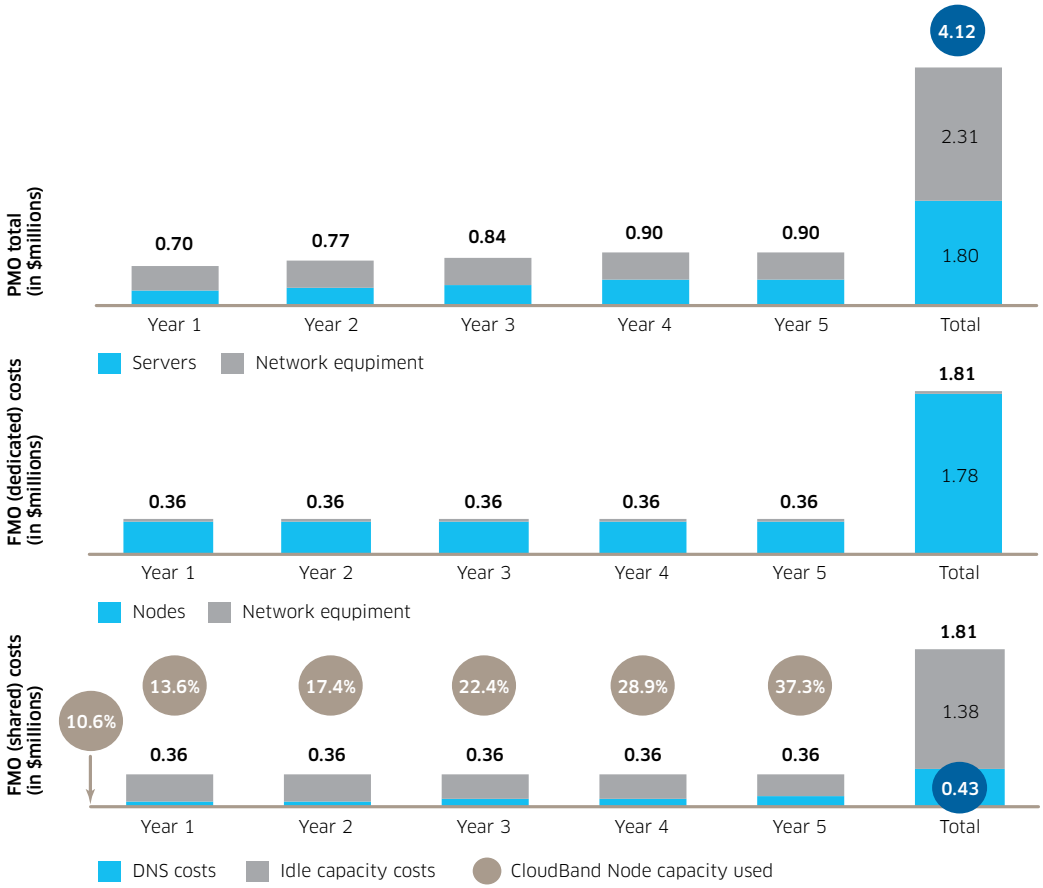

The total expected CloudBand Node capacity that the current DNS applications would require was calculated based on extrapolation of the data obtained in two load tests performed in the service provider's lab, with one CloudBand Node and a sample of DNS licenses. The results of this load test showed that the service provider could expect the current number of DNS applications in year 0 to use 10.6 percent of the total available capacity. With an expected 30 percent compound annual growth in traffic for the service provider DNS applications, the total capacity required in year 5 would be 37.3 percent.

The research team conducted a sensitivity analysis to understand the impact of a potential bias in the team's extrapolations to measure the capacity estimations. Figure 16 shows the results in the event of a bias of 25 percent in estimating the initial capacity required by the current DNS applications after being migrated to the CloudBand Node infrastructure (that is, 8 percent or 13.3 percent instead of 10.6 percent for year 0).

Figure 16. Sensitivity analysis

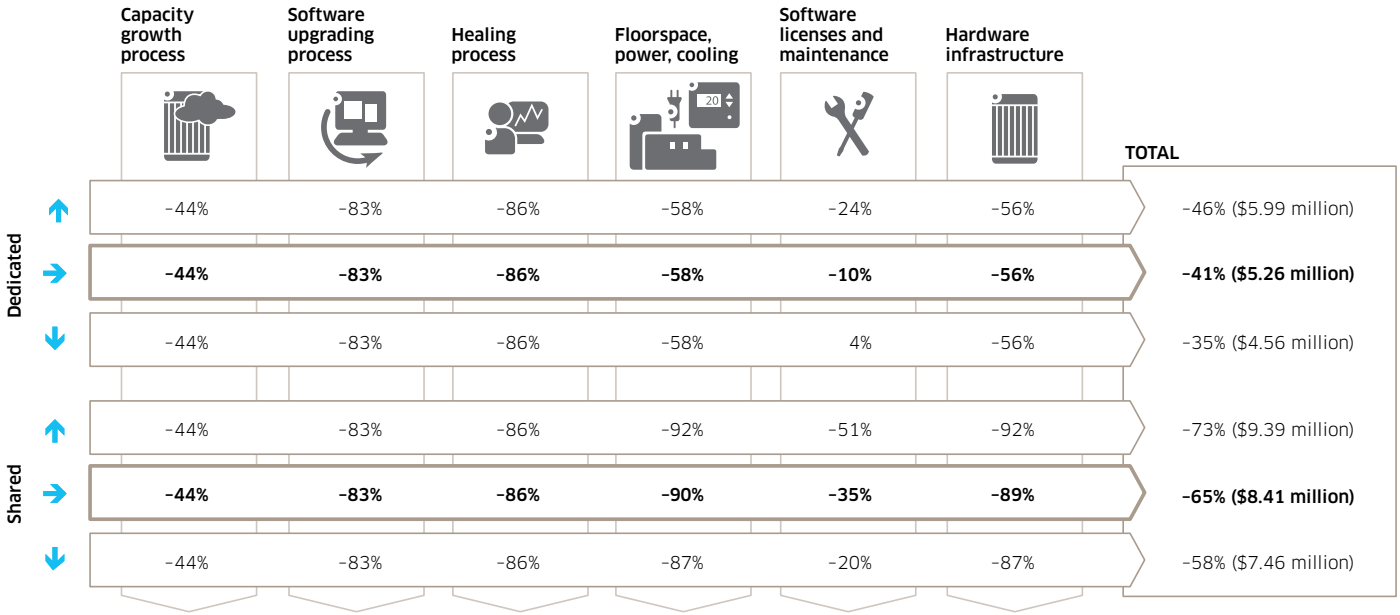

Figure 16 shows that even in the worst case scenario, considerable savings may be achieved with a direct impact of $4.56 million (35 percent savings) in the bottom line. In a shared scenario, DNS would experience a minimum of $7.46 million savings (58 percent savings).

10. CONCLUSIONS

The Business Case Study for DNS carried out jointly by Alcatel-Lucent and a Tier 1 service provider determined the cost advantages of migrating the DNS solution of a specific Tier-1 service provider to CloudBand. The cost savings analysis is related to three main cost drivers: CAPEX, OPEX infrastructure and OPEX processes. It shows the enormous benefits brought about by NFV even for a simple application such as DNS. In addition, apart from substantial monetary gains, the study determined that running the DNS application on CloudBand simplifies complex processes such as healing, scaling and software upgrading, which gives service providers greater agility and flexibility.

Figure 17 presents a summary of the 5-year total CAPEX and OPEX costs split in two scenarios. The first scenario (dedicated) shows the total costs that the service provider will need to incur to migrate from the PMO to the FMO. The second scenario (shared) shows the new cost of operating DNS in a more efficient way. The difference between the two is the cost of idle capacity, which other applications could use in a shared scenario.

Figure 17. Summary of the 5-year total CAPEX and OPEX costs

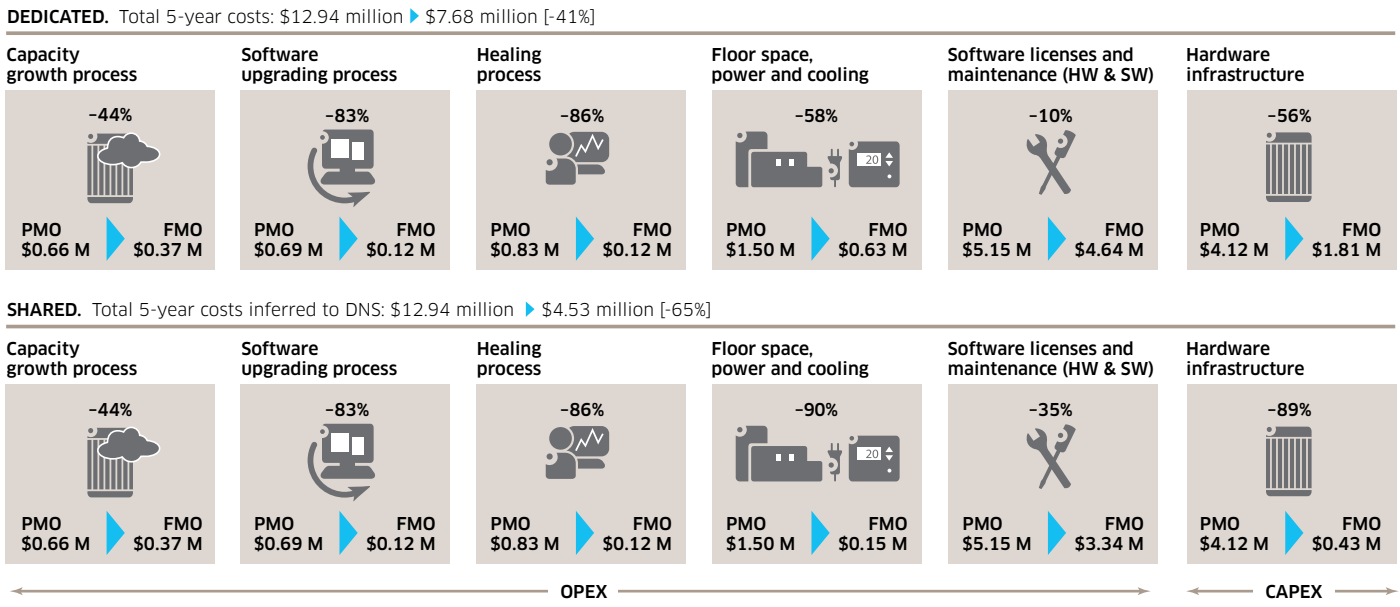

The total value created by CloudBand is the difference between the costs for the PMO ($12.94 million) and FMO ($4.53 million) in the shared scenario — $8.41 million. By migrating DNS alone, the service provider will be able to initially capture $5.26 million ($12.94 million minus $7.68 million). The balance of the savings will be obtained as the service provider leverages the free capacity for other applications thanks to the NFV model.

The analysis shows an impressive reduction in most of the process-related costs. In relative numbers, software upgrading and healing show an 83 percent and 86 percent reduction, respectively. This is a reflection of the operational process simplification enabled by CloudBand, which also shows in the lead-time improvements presented in previous sections.

Based on the absolute numbers, CAPEX initially contributes most to the total savings (44 percent) with $2.31 million in the dedicated scenario. Next are:

- Floor space, power and cooling (17 percent)
- Healing process (13 percent)
- Software upgrade process (11 percent)
- Software licenses and maintenance (10 percent)
- Capacity growth process (5 percent).

In the shared scenario, the CAPEX and the OPEX infrastructure-related costs increase their contribution to the savings as the cost of idle capacity is not allocated to DNS.

Although the scope of this business case did not include the revenue benefits that could be derived from agility and flexibility improvements (such as improved customer satisfaction and reduced churn), these savings should not be overlooked. Such benefits may prove to be the tipping point when evaluating whether to move certain applications to a cloud platform and the TCO business case is not convincing enough on its own.

11. ACRONYMS

ANS	Authoritative Name Server
CAPEX	Capital expenses
CR	Cache Resolver
DDOS	Distributed Denial of Service
DNS	Domain Name Server
FMO	Future mode of operations
KPI	Key performance indicator
NFV	Network Functions Virtualization
OPEX	Operating expenses
OS	Operating system
PMO	Present mode of operations
RCA	Root cause analysis
TCO	Total cost of ownership
TTM	Time to market
VM	Virtual machine

Alcatel·Lucent