On the Economical Benefit of Service Orchestration and Routing for Distributed Cloud Infrastructures:

Quantifying the Value of Automation.

ALU Bell Labs Business Modeling White Paper.

The cloud, interconnected and distributed software cloud computing platforms, virtualized real-time-aware streaming services are incumbent technologies that present service providers and telecom operators with some challenging business opportunities and business case questions.

This paper, in a back of the envelope style, reveals and quantifies possible benefits from virtualization, automated provisioning, deployment, management and scaling combined with optimal multi-metric placement of a simple personalized streaming application, in combination with service routing.

With input coming from qualitative and quantitative studies, data collection including interviews with a telecom operator and ALU SMEs, analysis of existing NPPE records, academic and business consultancy papers, as well as Bell Labs Business Modeling sources and expertise, a business case study was conducted, examining revenue, TCO, CAPEX and OPEX when deploying a multi-input/single-output video streaming application in different technical architectures.

This study provides an indication for possible economical and financial benefits generated through virtualization and automation as it applies for a multi-media streaming service.

Table Of Content

Overall technological observations

From application hosting point of view, several technological architectures can be chosen from, e.g.

- deploying an application onto a hardware appliance at customer premises,
- application virtualization in a centralized datacenter or
- opting for a cloud-enabled application closest to an end-user's location.

Depending on the application, the hosting platform needs to exhibit specific functionality with respect to performance, isolation, and several other factors and this at the level of compute, network and storage infrastructure.

Different cloud architectures can be considered, some examples being distributed cloud, over-the-top cloud architectures, heterogeneous clouds or any combinations. For each of these architectures, from a business point of view, optimal placement of a service towards its end-users is a necessity to ensure financial benefits for an operator and can cover parameters like minimizing networking resources, or maximizing server utilization or efficient allocation of hardware acceleration resources for services over time.

Depending on where an application is hosted in the cloud, the network connecting the service to the end-user is impacted at the levels of resource consumption, QoS guarantees etc. Given the type of application, these impacts are more or less stringent. e.g. web services versus real-time aware services such as gaming, differing in requirements for compute (performance, GPU adapter presence, compute isolation, etc.), network (latency, jitter, BW etc.) and storage (image size, storage resource consumption etc.).
The application itself is characterized with input and output networking resource requirements such as bandwidth, QoS, latency, jitter etc. that need to be accounted for in the network.

From a business perspective, the question raises whether "given a set of interconnected DataCenters and user demand for a service, what is the optimal location to host a service and at what TCO, CAPEX, OPEX, peering costs etc?".

From a technical perspective, the cloud's orchestration logic deals with optimal placement of services over a set of clouds and cloud nodes assuming a fully automated roll-out of services towards DCs. In a scheme of dynamic and mobile service deployment, service routing needs to be added to the mix to ensure that end-users connect to their requested services.

Clearly, the service placement algorithm that accounts for all of the above mentioned parameters will be a prime factor in determining the outcome of the overall economical equation and possible benefits.

Today's cloud computing architectures are centralized and network-agnostic and given the current state of the art, making them unfit for geo-graphically distributed services with tight QoS-constraints and/or high bandwidth and computation demands. Some of the reasons enumerated are

- Platforms or services or hosting platforms are focussed on web services. These applications are mostly non-real-time, nor latency bound, nor high-bandwidth oriented.
- Some emerging cloud platforms support streaming such as Amazon's CloudFront which is the commercial offering for a Content Delivery Network. Amazon's AppStream is a web service that deploys a multimedia application on Amazon Web Services (AWS) infrastructure and streams input and output between the application and end-user devices. The platform is realized only on windows server technology and architecture, provides automatic scaling but omits technologies like android, linux and other platforms, nor does it provide an answer for distributed clouds or service routing.
- Frameworks like ContentCentricNetworking (CCNx) only cover content and service routing.
- In the cloud, virtualisation focuses on generic HW and is not dealing with heterogeneous HW and networks on a general availability level. Early signs of the use of heterogeneous technology are mostly applied on a proprietary basis.
  One example is bing and baidu search engines using FPGA accelerators in search technology to speed up web-searching with a factor of 2 with only 30% power increase [1].
  Another example is the use of GraphicsProcessingUnits for hardware accelerated rendering in networked video games.

There is currently no framework that offers a combined approach on the topics mentioned above.
A framework that provides all of the above functionality is envisioned in this study and will be referred to as Future Service Oriented Networks (FUSION) (figure on the right).

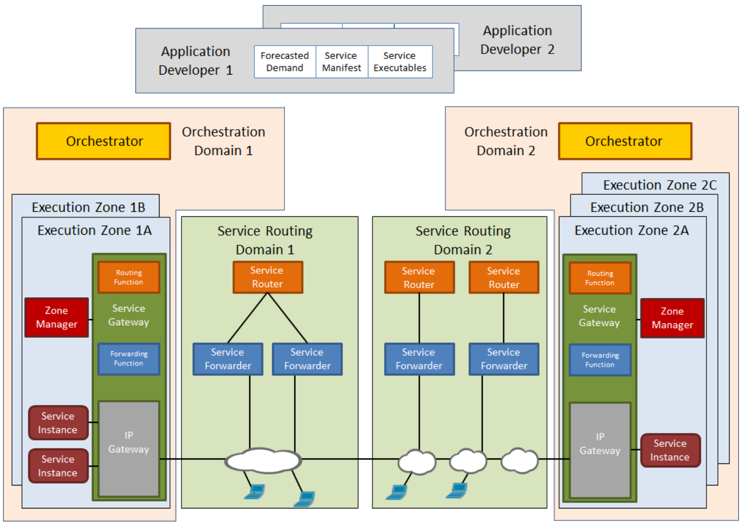

Positioning of FUSION

The goal of FUSION it to provide a framework that intends to facilitate

- the deployment of
- the placement of
- the scaling/elasticity in a configurable manner of
- the optimal service routing to
- taking into account HW accelerators in the cloud for placement and routing of

any kind of service in an **scalable** and **elastic** way, tackling the problem by integrating service provisioning and service-centric networking.

The scope of FUSION is on Interactive Services offered to a large number of geographically distributed users with services that require either real-time processing with low-latency or non-real-time processing with low-latency (e.g. labelling pictures submitted by a smart-phone for scene identification). Fusion will instantiate and route to the most "optimal **"** service for clients, where optimal is determined by:

- the service itself, possible factors are, but not limited to, cost, locality, legal, specific HW adapter, dependency, bandwidth, latency.
- the client request parameters, such as lower cost, high quality, fast response, ...
- the operating conditions like network occupation, data centre utilization, presence of hardware accelerators on compute platforms.
- the availability of automated intelligent service routing based on the response time and not necessarily to the nearest service instance, e.g. accounting for the overhead to launch a new instance taking too much time and therefore routing to an instance farther down the network, but overall will react faster to user request.

From business point of view, FUSION offers the roll out of new applications in an automated manner with optimal service provisioning, placement and routing and therefore enables AppStore like concepts that could be extended with cloud computing service offering.

FUSION allows for the business case to assume that from service hosting point of view, the most optimal conditions can be achieved given the availability of FUSION's functionality.

Some business case questions

An immediate question that arises from business perspective is which compelling and appealing application justifies the investment into a platform as envisioned by FUSION? Which service hosting and networking architecture should be retained? Which business relations accompany a given service or application, at which conditions and how are these effectuated? How does one determine and select a popular service? Given an AppStore-like concept, complemented with service hosting, where developers post their applications in a store, what is their benefit/opportunity of using such a platform and what burden are they relieving by outsourcing operational aspects? Should network operators host datacenters themselves and invest in datacenter technology or team up with cloud service providers to provide service offering? A multitude of questions arise and they are elaborate.

This study takes a pragmatic approach by focusing on some specific questions that are raised by service providers and/or operators when selecting hosting infrastructure for deployment of real-time aware services:

1. Hosting of a service onto a hardware appliance (e.g. SetTopBox) at customer premises?
2. Hosting a virtualized service in a DataCenter (aka virtual STB service), then is a centralized DC enough and what about network effects?
3. Hosting a virtualized service in a Distributed DataCenter using (vSTB) , but then how much distribution does one need to obtain maximized economical benefits?

The definition of "any kind of application" ranges from simple text-based request-reply services to multi-user cloud gaming with video conferencing capabilities (a service that likely does not exist yet). This leads to a plurality of possible use cases that determine the outcome of their respective business cases.

In order to obtain an initial, back of the envelope outlook of financial and business figures when opting for FUSION, the study selected a midrange application containing multi-tier video streaming and processing with a rather “short-lived” usage (about ½ hour per day). This service is described in the following chapter.

This paper further focuses on the business case of this specific service and used input coming from qualitative and quantitative studies, analysis of existing NPPE records, academic and business consultancy papers, Bell Labs Business Modeling source and interviews with a telecom operator and ALU SMEs.
The business case study solely examines revenue, TCO, CAPEX and OPEX when deploying a multi-input/single-output video streaming application in different technical architectures.

Use case

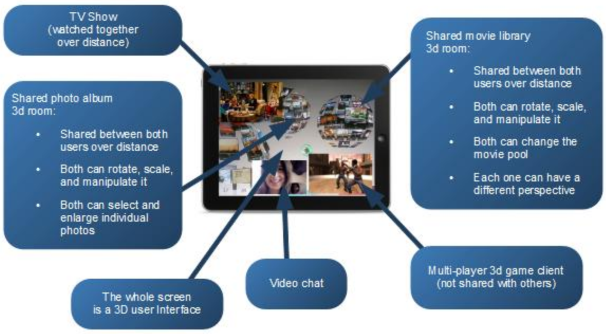

A Personalizable Dashboard service, consuming several input video streams and producing a single output stream including a user feedback channel, is selected in this business case study. The BC considers 3 possible technical realizations:

1. Service running on HW appliance (e.g STB)
2. Virtualized service running on dedicated datacenter technology
3. Virtualized service running distributed cloud infrastructure with FUSION including orchestration and service routing.

Some practical scenarios for operators/service-provides that will be considered:

1. Use HW appliance (possibly a powerful STB) (referred to as **P**resent **M**ode of **O**peration, **PMO**)
2. Move to virtual STB in centralized heterogeneous datacenters (**F**uture **M**ode of **O**peration **1, FMO1**) that uses hw acceleration in servers for improved application density.
3. Move to vSTB in distributed DC with FUSION technology (**FMO2**) allowing roll-out of all kinds of services alongside vSTB service and whereby server technology uses heterogeneity albeit more generic (e.g. reprogrammable FPGA’s).
4. If operation is already in mode FMO1, move to FMO2. Furthermore, the level of distribution needs to be considered. In this use case, low, medium and high distribution indicating DCs at core, aggregation and access network.

This paper will look primarily at choices [2-4] and at parameter sensitivity of the different options.

Some functional observations

Cloud and virtualization (FMO1) as well as FUSION (FMO2) both enable the launching and trialing of new services and open up new revenue sources in more flexible ways. The time to market for new service introduction reduces significantly compared to a hardware appliance approach (PMO). That translates into the capability of faster deployment of new services as well as time-, geographically- or population-based variations of services – yield management made possible in application's world. FUSION technology specifically accelerates new service deployment with a reduction in time, of roll-out of a service onto cloud technology, taking 1 month down to 15 minutes of new service/application deployment time. (cfr. RedHat figures for it services).
Due to the automated service deployment, hosting and routing, FUSION supports AppStore kind of service models. For an operator, this allows to circumvent the need to select an appealing service. Moreover, it enables an operator to enter business models where services can be gathered and offered in a service-store and where end-users determine the popularity of an application. This allows for additional revenue for an operator.

From business modeling point of view, revenue and revenue increase are modeled based on Netflix revenue CAGR observations (2008-2013) and have been normalized for the take-up rate used in this study.

The different technological architectures have an impact on a service providers or operators operational processes which need to be adapted to align with fusion-alike fast pacing technology. This will require operators to change their internal organization and processes as to take maximal advantage of the enhanced flexibility and innovation.
While considering the results of this business case, it became clear that a fusion architecture would help improve an operators' business allowing for new business development towards OTTs and enterprises.

Virtualization (and so FUSION) shift CAPEX from STB to network infrastructure (to handle the increased traffic due to processing performed in central or distributed data centers). That allows operators to amortized network investment (otherwise delayed), allocating some costs to vSTB service.

Business case Assumptions

The business modeling for hosting a multimedia dashboard service is projected on a 5 year basis assuming the availability of FUSION-alike technology by mid 2015.

A hypothetical network is assumed consisting out of a backbone with routers, an aggregation network with Ethernet switches and an access network with access multiplexer. The referenced operator is assumed to have 2.5M subscribers. The different networks are directly interconnected having a single hop between them.

Network assumptions	[units]
Routers	5
Switches/router	50
Access multiplexer/switch	100
Subscribers/access multiplexer	4000

The Personalizable Dashboard service described above consumes up to 8 video input streams and produces 1 video output stream with an upstream feedback channel. Each user has its personalized service instance and uses the service for about 30 minutes per day spread over 8 sessions.

On a high level, the Personalizable Dashboard service is comparable to an ElectronicProgramGuide service and the business case used EPG related metrics as known by ALU business units.

An endpoint is a (v)STB or end-user connection that will be served by a server.
Server disk resources associated with storing and retrieving user context are accounted for in the server cost itself.

Service assumptions	
Input streams/ app	8
Video Stream BandWidth	4 Mbps
Feedback channel	10 kbps
Input streams per server	200
End-points per server	4000
Services per server (highly optimized)	300
Service per server (optimized)	80

In the case of virtualization with distributed data centers, the service needs to be downloaded to the appropriate DC. On the downloading or moving of images, the following observation can be made for this business case: Assume that the service has a 50MB image and caching is assumed to reduce the number of downloads and is estimated at 1 download per month (either for provisioning or application update purposes). Considering the service usage for 1 month then, for service output only, the following calculation can be maintained:

(4Mbps/8bitsperbyte)*60secpermin*30min*30dayspermonths=27000MB.

Image download is only ~0.023% of the service own download quota and can be omitted since this is in the margin of error / noise of this business modeling exercise.

In FMO2 case, FUSION offers service routing functionality that is discounted with a session license cost of .01€.

Some key cost parameters related to data centers and cloud exploitation are enumerated in following table. The STB needs to have multiple tuners and is medium to high end.

It is assumed that in case of FMO1, server technology can be architected more efficiently using heterogeneous hw acceleration specifically for vSTB application and due to centralization and higher volumes compared to FMO2 with more fragmented server technology that likely needs to deal with a multitude of services.

datacenter assumptions		
STB	150	€
Floorspace DC	200	€/m2
Energy cost DC	0,112	€/kWh
man day (installation etc.)	600	€/day
IT Admin per 500 servers	150.000	€/year
Server CAPEX	7350	€/server
Power consumption per	405	Watt
Sessions per server (FMO2)	80	
Sessions per server (FMO1)	300	

Concerning deployment cost, FMO1 is considered to have a dedicated deployment of servers for a specific application where in the FMO2 case, the deployment of FUSION servers is dedicated. FUSION enable the automated roll-out of any kind of application requiring no dedicated application provisioning/roll-out.

Application field-trialing costs/expenses were not accounted for in FMO1 nor FMO2.

Maintenance cost of servers has been modeled as a percentage of installed base.

Network transit costs were modelled based on [2].

Business Case method.

The framework described in “Positioning of FUSION” allows for the business case to assume that, from service hosting point of view, the most optimal conditions can be achieved given the availability of FUSION’s functionality.
Depending on a service’s characteristics (e.g number of input channels and required bandwidth, the amount of compute resources needed etc.) , FUSION orchestration is able to determine the most optimal placement of a service and routing to a service for any kind of underlying architecture (from centralized to highly distributed interconnected datacenters).
For the selected “Use case”, the business calculations can assume that, given an underlying service architecture, the minimal required resources can be allocated and used to deploy, host and route the use case service.
The business calculations are presented in the next chapter.

Results

The evaluation of the results focuses on the *relative* improvements that can be gained in the different scenarios since this helps in drawing initial conclusions. As a note, the absolute values are of less importance and should be correlated against the business case assumptions

Cumulative discounted cash flow

On a 5 year basis, the cumulative discounted cash flow shows negative returns for the hardware appliance approach (PMO). Both virtualization (FMO1) and FUSION technology (FMO2) show positive returns within 1 year whereby FUSION high distribution shows some additional benefits in longer term. Moving from PMO to virtualization offers highest gain with additional gains when moving from virtualization to automated orchestration, provisioning, deployment. In case of FMO2, with a higher TCO due to automation investment costs, the CDCF is positive due to even higher Revenue increase.

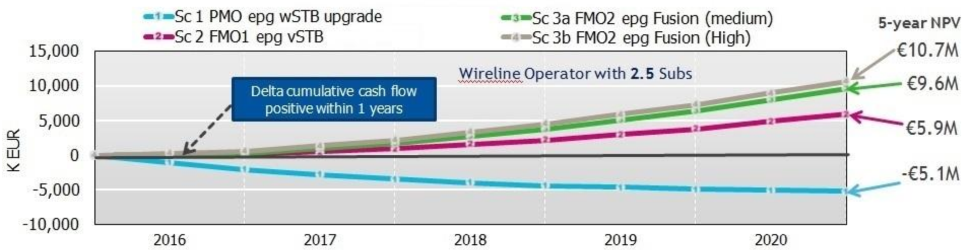

NPV

A Comparative 5 year NPV evaluation shows a ~200% improvement when moving from PMO to FMO and even a ~300% when comparing FMO2 against PMO. When comparing FUSION technology against virtualization, at least a 60% increase can be obtained with as high as 80% in case of FMO2 applied to high distribution.

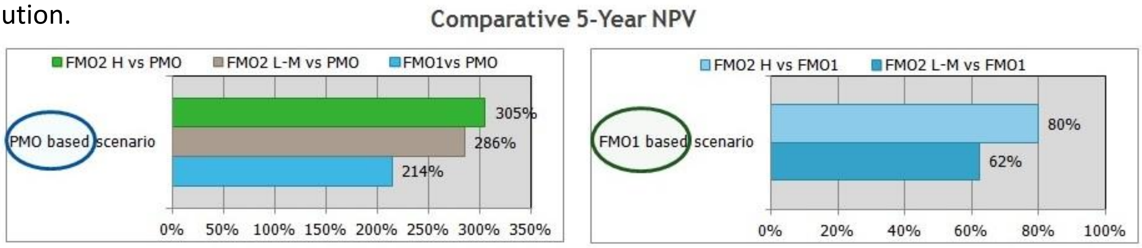

Revenue

The diversity of virtualized services not explored previously triggers additional revenue since virtualization allows activating untapped revenues.
A revenue growth estimation was modeled based on observed revenue growth from Netflix between '08 and '13 as indicated in fig to the right.

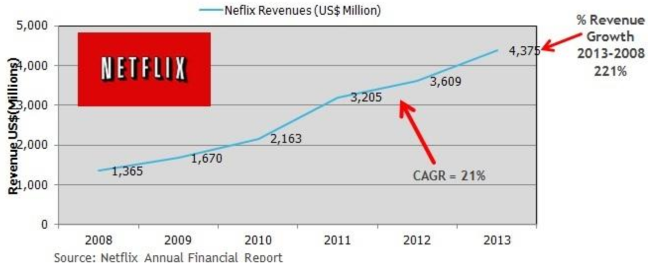

DataCenter distribution and automated provisioning with optimized service routing as offered by FUSION allows for improvement of overall service quality and enables additional real-time/QoS aware services that could further increase this effect up to about ~120% revenue increase compared to an hardware appliance roll-out of services.

revenue vs	PMO
PMO	
FMO1	+55%
FMO2(L-M)	+119%
FMO2(H)	+119%

A 5-YEAR revenue analysis as a percentage of the total revenue is shown in the figure on the right. The improved service offering allowed by virtualization would obviously generate higher usage of the (v)STB application (corresponding to the contribution of the application to the subscription fees), and would probably contribute to attract additional subscriptions to the operator's service, and would also allow to offer a higher class of (v)STB applications for additional subscription fees. This is materialized in the table by the 27 and 44 % increase against PMO.

Incremental Revenue analysis (%)	FMO1 vs. PMO	FMO2 low & medium vs. PMO	FMO2 High vs. PMO	FMO2 low & medium vs. FMO1	FMO2 High vs. FMO1
Generated subscription part	27%	44%	44%	10%	10%
EPG Generated additional sales	23%	62%	62%	25%	25%
EPG-generated advertisement	4%	14%	14%	6%	6%
Incremental Revenue (%)	55%	119%	119%	42%	42%

In addition, the increased flexibility of the (v)STB service would allow the creation of new applications (including personalized versions of existing applications that could be offered as additional paid options), the creation of targeted advertisement (4 and 14 % increase figures versus PMO in the table) and the creation of promotional links with other services (e.g. VoD catalogs, not-user-subscribed or user-subscribed but not used channels and packages, ..., translated in the table, with 23 and 62% increase versus PMO).

The possibility of offering a variety of services clearly drives incremental revenue and as indicated by the results, the value FUSION could add is much beyond that enabled by virtualisation.

TCO

Looking at the 5 year TCO comparative analysis, maximal TCO savings in the range of ~50% can be made when upgrading from HW solution towards virtualization.

5Y-TCO *versus*	PMO	FMO1
PMO		
FMO1	-48%	
FMO2(L-M)	-35%	+24%
FMO2(H)	-44%	+7%

In the case only 1 service is hosted, then, due to FUSION's additional investments for orchestration servers, service routers etc., a TCO increase can be observed when going from virtualization towards FUSION of up to ~20% for medium distribution. The additional TCO when moving from virtualization to FUSION is rather limited to a mere 7% in the high distribution scenario.

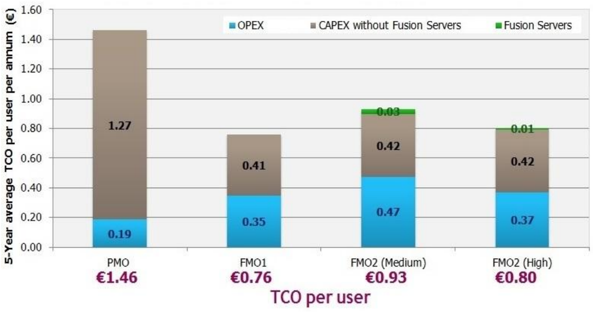

CAPEX

Introduction of virtualization over hardware appliances will cut CAPEX with about ~60%. Due to the additional cost of automation in the case of FUSION, CAPEX gains drop with a mere ~2% with almost being invariant against the level of distribution.

5Y-decr. CAPEX *versus*	PMO	FMO1
PMO		
FMO1	-59%	
FMO2(L-M)	-56%	+5%
FMO2(H)	-58%	+2%

Comparing FUSION against centralized datacenter virtualization incurs additional CAPEX with as low as ~5% with only a ~2% increase in case of high distribution due to network savings.

OPEX

An increase in OPEX is to be expected since services are being hosted in datacenters and all operation expenditures are at the expense of service provider or operator. The increase is about ~10% for virtualization compared to hard appliance approach and ranges between 15 and 20% in case of FUSION against PMO. However, again considering only 1 application/service, the increase in OPEX is less in the FMO2 vs FMO1 scenario then the FMO1 vs PMO scenario.

5Y- OPEX *versus*	PMO	FMO1
PMO		
FMO1	+11%	
FMO2(L-M)	+21%	+19%
FMO2(H)	+14%	+5%

On overall, FUSION Servers operations increase OPEX between 10%-20%.
However, FUSION squeezes OPEX due to bandwidth savings when considering distribution. This is one of the main drivers to implement FUSION.

Business case sensitivity

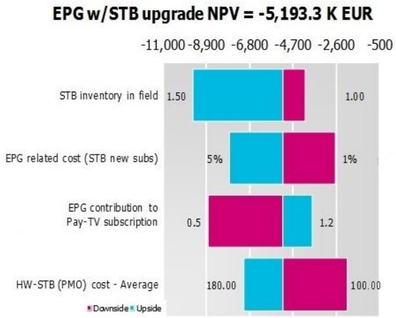

The Net Present Value of the Present Mode of Operation is related to CAPEX investment due to STB upgrade and is therefore STB inventory sensitive. Concerning the STB inventory in the field, the assumption for the base calculation, is that there is 10% more STBs than subscriptions, due to the fact that few subscribers have multiple TV at home and need service replication. Therefore, the figure of 1.10 in the model and the pivot of the item in the graph nearby. In order to test the sensitivity of this parameter, a variation of 1.00 (= 1 STB only per subscription) and up to 1.50 (= 50 % of subscribers have a second STB at home for an additional TV set) was modeled. Not surprisingly, reducing the number of STBs distributed reduces the cost (please note that at least one STB per subscription must be maintained, so the possible reduction is rather limited) and increasing the number of STBs (which is the market trend, due to the increasing equipment of households), increases the cost..

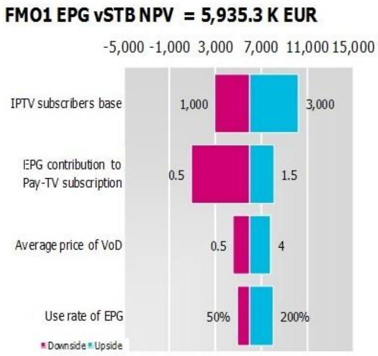

The Future Mode of Operation 1, being virtual STB hosted in centralized datacenter, is sensitive to the revenue factor including take-up rate and pricing.

The second main parameter is the calculated contribution part of the application, to the subscription fee (that is the calculated value of the application within the subscription). As stated above, the higher quality of the application raises its use rate and therefore its value.

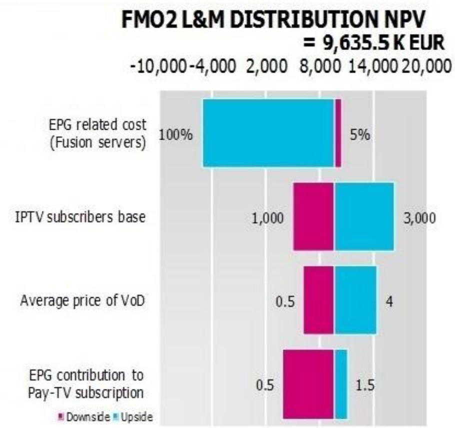

Given that FMO2 represents automated provisioning and roll-out of services in an optimized way, FUSION makes sense for several applications and not just one.

From the results obtained in the study, the 5-Year NPV for Fusion is highly sensitive to the server investment and therefore also to achievable application density per server, performance etc.
The impact can be further neutralized by serving a number of apps/services.

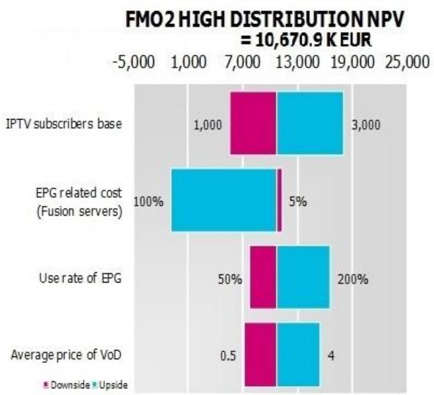

Please note that it is not possible to quantify a correlation on these figures as different apps have different needs (processing, storage, latency, ...) and would be handled differently by FUSION and therefore would not give any meaningful results beyond stating that costs are diluted and revenues are improved.

Conclusions

Virtualization in a centralized, heterogeneous datacenter approach contributes to a flexibility in application management that allows exploration of new revenue territories. However, its introduction requires changes in an operator's structure and processes whereby:

- Processes for SW upgrades have to be reviewed
- Processes to quickly test lots of new features, price packages... have to be invented
- Building eco-systems of partners becomes of particular importance and quality and personalization of applications and services rules over content quantity offer.

Virtualization reallocates CAPEX from STB to network equipment (increased traffic) and to OPEX with a CAPEX reduction of about ~60% and an OPEX increase of about 10% driving revenue up with as much as ~50% result in a positive Cumulative Discounted Cash Flow, over a 5-year period, within 1 year. Virtualization is also about structural changes needed (in network and end-user solutions management) to leverage flexibility.

FUSION offering automated service deployment (provisioning, orchestration, scaling, placement and service routing) and targeting distributed interconnected heterogeneous datacenters, benefits from the above virtualization advantages and pushes virtualization and heterogeneity a step forward by:

- Lowering the OPEX costs added by virtualization through minimization of overall consumed network bandwidth;
- Improving users' QoE (especially latency) by allowing to position services closer to the user from a network point of view;
- Easing apps' roll out and processing management

Specifically compared to virtualization in centralized, heterogeneous datacenters, FUSION increases CAPEX by 2 to ~5% and OPEX increase within a range of 10% to 20%, resulting in a overall TCO increase ranging between 7% and 24%, according to the degree of distribution. The benefits FUSION brings is a ~40% revenue increase resulting in positive Cumulative Discounted Cash Flow, over a 5-year period.
Considering FUSION and a higher degree of distribution with respect to distribution, only rather limited benefits were obtained.

... FUSION allows operators to explore new B2B businesses: creation of apps distribution and processing management services for OTTs (entertainment...), CDNs (reach extension), enterprises (own use and customers' service)

... FUSION contributes to improve ROI of network equipments' upgrade, as FUSION would use processing and storage capability needed for other purposes, the more FUSION distribution granularity, the closer to access.

Contacts

luc.vermoesen@alcatel-lucent.com,
jean-paul.belud@alcatel-lucent.com, emilio.lastra@alcatel-lucent.com, luc.ongena@alcatel-lucent.com,
Frederick.vandeputte@alcatel-lucent.com

Acknowledgements

The research leading to this paper has been partially supported by the European Commission under the contract FP7-318205, "FUSION: Future Service Oriented Networking". We would also like to thank the anonymous reviewers for their valuable feedback to improve the overall quality of this paper.

References

[1] http://www.pcworld.com/article/2464260/microsoft-baidu-find-speedier-search-results-through-specialized-chips.html

[2] http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php, http://www.michaelgeist.ca/content/view/5727/125/

Acronyms

BC: Business Case
BW: BandWidth
CAPEX: CAPital EXpenditure
CCNx: ContentCentricNetworking
CDN: Content Distribution Network
DC: Data Center
EPG: Electronic Program Guide
FPGA: Field-Programmable Gate Array
FUSION: Future Service Oriented Networks.
GPU: GraphicsProcessingUnit
HW: HardWare
NPV: Net Present Value
OPEX: OPerational EXpenses
QoS: Quality of Service
ROI: Return On Investment
STB: Set Top Box
TCO: Total Cost of Ownership
vSTB: virtual Set Top Box

About the Authors

LUC VERMOESEN is a research engineer in the IP Platforms Research Program in Bell Labs in Antwerp, Belgium. He graduated in engineering in 1989 and studied computer science in 1995. He joined Alcatel-Lucent back in 2000 where he worked on projects involving 3G Mobile, VDSL, Asynchronous Access Multiplexer and IP Service routing and switching. In 2007, he joined the Bell Labs Fixed Access team where he was involved in Home Networking research and also contributed to the Broadband Forum standardization activities. In 2009, he focused on multimedia-related research topics like novel graphical user interfaces for IPTV and network-based rendering techniques using dedicated HW acceleration. From 2011 onwards, he is involved in cloud computing research with specific interest in virtualization and performance and the applicability of heterogeneous hardware in the cloud. He holds over a dozen patents and is author and co-author of at least 2 publications.

JEAN-PAUL BELUD is a Network Strategy Professional in the CTO organization, Bell Labs in Nozay, France. He obtained his MBA from ISG, Paris, France in 1984 and completed the Entrepreneurship program at ESCP-Europe, Paris, France in 2012. He has 33+ years of management, innovation and entrepreneurship experience in all aspects and stages of B2B, B2C for various industries and products. Within Alcatel-Lucent, he brings a 5 year sales, 12 years of marketing and profit center management, completed with 12 years of international purchasing and project management. The last 5 year, he specialized in business modeling, leveraging his experience in multimedia, internet and research and innovation areas. He is author and co-author of at least 2 publications

EMILIO LASTRA is an EMEA business modeling manager, Bell Labs, UK. He is a PhD candidate in Management - Information Systems and Innovation and holds an MSc in Economic History/Research from The London School of Economics and Political Science in London, UK. He has a MBA degree from the University of Bath, UK, a MSc in Management and a BSc in Finance and Accounting from Monterrey Institute of Technology and Higher Education in Monterrey, Mexico and English History from Oxford University, UK. He has 14+ years of experience in Telecommunications in various positions including Business Modeling, Bids&Proposals, Business Strategy and Processes, Quality and Project Management. He specializes in modeling business cases for operators and Alcatel-Lucent internal projects with emphasis on broadband access networking, wireless and services. More recently he is supporting Bell Labs Research projects.

LUC ONGENA is a Senior Consultant in Bell Labs Advisory Service. He has a Master degree in Telecommunications at the University of Antwerp and has a Marketing and Communication Post-graduate from Insead business school, Fontainebleau/Paris, France. He was elected twice as an Alcatel-Lucent Technical Academy member. His experience include 30+ years in software engineering, network architecture, product lifecycle management, solution integration and operations analysis. He brings extensive industry knowledge and technical breadth and depth in developing solutions for business and operations transformations, network migrations and IT architectures. He has led several innovation assessments, market opportunities and revenue enhancements studies. His recent interest is in consulting with European service providers, analyzing their IMS and access network architectures, troubleshooting and outlining future strategy options.

FREDERIK VANDEPUTTE is a research engineer in the IP Platforms Research Program in Bell Labs in Antwerp, Belgium. He obtained his master's degree in Computer Science in 2003 at Ghent University, Belgium. In 2008, he obtained his PhD in Computer Science at Ghent University, doing research on characterizing and exploiting the time-varying behavior of applications on multi-configuration processor architectures. He joined Bell Labs in 2009, further specializing on software parallelization for many-core as well as heterogeneous architectures for media processing applications. From 2012 onwards, he started working on efficiently deploying media processing applications in the network and studying the feasibility and impact of using hardware accelerators in such environments. His current research topics include heterogeneous clouds, network functions virtualization and performance optimization. He is a work package leader of the FUSION FP7 project involving service-oriented networking for demanding real-time media applications, and is author or co-author of over a dozen papers in international conferences and journals.