HEAVY READING

**White Paper**

# Optimization & Automation: A Business Objective-Driven Approach

Prepared by

Caroline Chappell
Practice Leader, Cloud & NFV, Heavy Reading
www.heavyreading.com

on behalf of

AriaNetworks

www.aria-networks.com

**June 2016**

# Introduction

Telcos are facing an unprecedented amount of change in their business environment, in customer expectations and as a result of introducing virtualization to their networks. They want to able to adapt in a lean and agile manner, understanding what needs to change in their organization optimally to meet those goals and then driving those changes in a disciplined and automated way. Business-driven optimization is the adaptation mechanism that will help operators survive an era of rapid and unpredictable change. It is the practice of determining the business goals behind change, finding the best solution to meet them and ensuring that the objectives are achieved. In a telecom context, a business goal might be the most profitable way to deliver a service to a particular customer or the most profitable way of deploying a network.

Business-driven optimization is not a one-time activity, however. Due to the rapid pace of environmental change, initial solutions can suffer from entropy, a challenge that is compounded by the dynamic nature of a virtualized network. So optimization must be carried out continually, adjusting for and adapting to any changes that come along. Nor is it sufficient for operators to optimize brilliantly in silos – for a single layer of the network or one specific type of service, for example. In a complex world with competing pressures and priorities, business-driven optimization must be able to take into account multiple business, technical and regulatory variables across the entire network and service landscape. This is a massively demanding task but one that reflects the reality of operating a sophisticated business today.

Given the pace and scale of the optimizations that operators will need to apply, business-driven optimization will need to be highly automated. Operators are already exploring the automation needed if the virtualized network is to deliver the agility and cost benefits promised for it. In the automation scenario for the virtualized network, business-driven optimization is the "brain," analyzing input from the virtualized environment, modeling the input against business objectives, generating a new optimization solution and triggering other automation tools to make the required changes.

The fact that the optimization process needs to be fast, holistic and continuous not only implies automation but also that it should be driven by artificial intelligence and machine-learning techniques. Business-driven optimization powered by artificial intelligence (AI) can cope with multi-variable complexity at speeds impossible to achieve with conventional optimization tools. Machine-learning techniques enable business-driven optimization to acquire and process knowledge without the need for coding, making it highly agile and responsive to change.

This paper describes a business-driven optimization approach and why it's different from the optimizations that operators carry out today. It assesses the features and benefits of this approach and explains the key requirements for a centralized optimization platform that supports it.

# Adapting to Change: Why Optimization Matters to the Business

## Adapt or Die

Every business today is engaged in a Darwinian struggle to adapt rapidly to changing conditions. Agility is a top requirement as the world evolves at an ever-accelerating pace. Telcos are no exception: Their businesses are changing as customers become more demanding, value-chain relationships grow more complex and regulation is rewritten for the digital age. The types and volume of customer-facing services they are expected to deliver is exploding. They are deluged with new technologies in every area of their organization, from network virtualization, next-generation mobile and fixed access and core networking technologies to those that support the Internet of Things (IoT), information management and new payment methods.

As part of becoming agile enough to embrace and exploit change, telcos want to operate in a lean manner. At a business level, they want to execute just in time to eliminate waste, build in quality and use knowledge holistically to optimize their processes, systems and key assets. Lean implies a continual, incremental optimization of all three to accommodate change smoothly rather than disruptively.

Two areas of business change in particular are driving operators toward a lean and agile mode of operation: customer experience management and the need to virtualize their key asset – the network – to achieve agile business goals. As **Figure 1** shows, customer expectations of product performance are rising as the services they want become increasingly latency-sensitive.

**Figure 1: Changing Customer Experience Expectations**

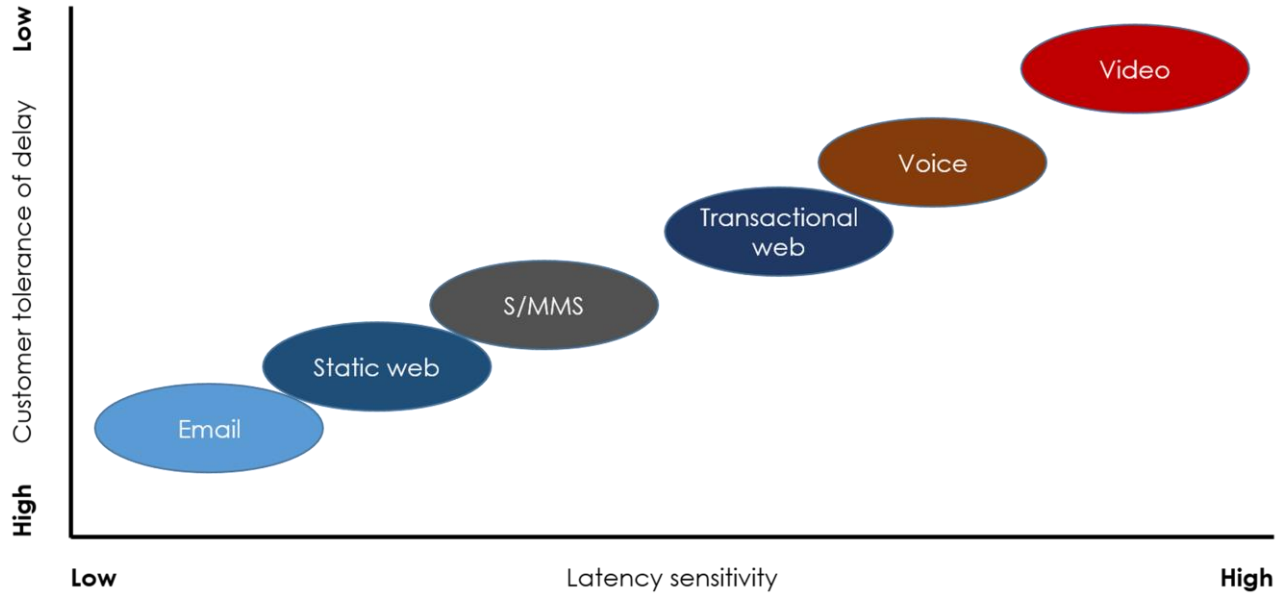


*Source: Heavy Reading*

According to Cisco's Visual Networking Index, mobile video traffic now accounts for more than half of all mobile data traffic (55 percent in 2015) – a figure that is set to soar as the trend for live streaming in the fixed network moves to smartphones. How well an operator manages its assets to deliver the right experience to all subscribers has an immediate impact on its bottom line. The right customer experience enables an operator to achieve business goals such as higher profitability, churn reduction and lower operational costs.

Improving customer experience is itself one of the drivers for network virtualization. Virtualization supports the on-demand, self-service and personalized delivery of customer-facing services that a Web-savvy generation expects. Operators that can't provide this experience will quickly find themselves disadvantaged in the market. Telcos expect network virtualization to give them greater flexibility over capacity deployment, more control over the capex and opex costs associated with service delivery and better alignment between revenue and network investment as traffic continues to grow. As AT&T's Domain 2.0 initiative illustrates, network virtualization is a business-driven program designed to transform a telco's operational capability and profitability in a changing world.

## Optimization & Automation Are Key to Adaptation

In an era of rapid and unpredictable change, optimization is emerging as a critical step in the adaptation process (see **Figure 2**).

achieve the right set of business goals. As we have said, an operator's business goals are increasingly oriented toward meeting new customer experience management requirements and achieving new levels of profitability through network virtualization.

Optimization is that guide, and in a customer-facing service delivery context, it means making the delivery of each service as fully perfect, functional and effective as possible according to a specific set of objectives. It is the step where business goals for change are determined, and that ensures these business objectives are achieved. For this reason, it should be closely aligned with operators' automation programs.

The dynamic nature of the network means it must be optimized continuously to fulfil such goals. This is especially the case with virtualization, which introduces the potential to change network design in real-time as resources – virtual network functions (VNFs) – are spun up and down to fulfil new service requests and/or handle the ebbs and flows of customer demand. The virtualized network, already a complex system, becomes one that must adapt to customer and service requirements in a lean and agile manner. Optimization is the intelligent management function that determines how the virtualized network should change in response to service delivery and customer experience objectives.

It is well understood that the virtualized network must be managed in an automated way. Network function virtualization (NFV) and software-defined networking (SDN) depend on automation to achieve the cost and agility benefits operators expect from them. As a function that critically drives service lifecycle management from a business perspective, optimization will also need to be automated in a virtualized network context. Optimization cannot be applied fast or frequently enough to changing virtual network topologies without automation.

The network will be a hybrid of physical and virtual elements for many years to come. Although the spotlight is on virtualization as a driver of optimization and automation, the network as a whole will require an adaptive and automated approach that is very different from approaches operators have undertaken in the past.

## New Requirements for Optimization

Operators must increasingly weigh up multiple factors and take complex business decisions that affect service delivery in shorter and shorter timescales. For example:

- A **provider of enterprise services** wanting to introduce a new, secure on-demand video conferencing service for businesses, using a mix of owned and leased resources, must decide on the price point that will make delivery profitable, how to avoid the conditions that could damage profitability and identify the growth scenarios that will justify additional build.
- A **global content provider** planning to roll out 4K streaming of live music events must decide on the best way to use its global capacity, to identify hotspots and the failure conditions that could compromise the subscriber experience and come up with ways of managing service degrade.
- A **mobile network operator** hoping to monetize a dedicated slice of its 5G network to carry live feeds of a city race will need a means of dimensioning the slice and plans to change its network configuration to protect revenue and service levels in the event of a partial failure.

All of these are optimization challenges that require sophisticated, business-driven plans to guide an optimal, automated reconfiguration of the network.

# Breaking With the Past: A New, Business-Driven Approach to Optimization

## Shortcomings in Optimizing Service Delivery Today

Telcos' current optimization practices come from the days when they could afford leisurely planning and implementation cycles for long-lived systems and customer-facing services. They typically focus on the network with the business aim of deferring capex and opex by producing an optimal design for asset planning purposes. They are fragmented, manual and ad hoc. As a result, they are not applicable to – or agile/lean enough for – the dynamic, service-driven virtualized network and the completely different customer experience requirements operators have to contend with today.

- Different layers of the network are optimized separately by different domain specialists. Services, however, run across multiple layers and need a holistic optimization for their needs. Joining the dots between different service and network domain optimizations is extremely challenging. Operators typically overbuild by throwing "just in case" bandwidth at the problem, but this is the antithesis of a lean and agile, "just in time" solution.
- Different organizational roles optimize for different goals. The CTO, for example, wants the network optimized for reliability; the COO is concerned about its operational aspects and product managers' optimization variables reflect business issues, such as regulatory compliance and margin. This leads to conflict and the lack of a coordinated, business-driven optimization strategy. Each role covers the risk for their individual department but doesn't determine holistically what's good for the business.
- Optimization is typically a one-time activity carried out as part of a new customer-facing service design and deployment. A substantial amount of planning and effort goes into a new service, as operators want to drive the optimal amount of value from their investment, but optimization decisions are not subsequently revisited despite the changes that inevitably take place after service deployment, so it doesn't take long before a well-optimized network starts running in a non-optimal way. For example, if customers are connected in a non-standard way to "get the job done" and equipment is added and moved, the result can be stranded assets, inefficient connections and under-utilization of network resources. This situation is compounded In a virtualized network, where change is more frequent than in the physical network, and costs may quickly spiral out of control.
- Excel is the main tool used to support optimization. Optimization models built by different people can be varied and idiosyncratic and have limited reuse. Tooling shortcomings means that optimization is rarely automated and tends to be unsophisticated, relying on one or two variables, such as cost and/or latency. It is simply too complicated, time-consuming and expensive to build multi-variable optimization models using "brute force" manual means.

## New Business-Driven Optimization Requirements

The need to adapt rapidly to change at a business level, to improve customer experience as a business goal and to operate the virtualized network in real time and with extreme automation for competitive business advantage are driving the need for a new approach to optimization.

Adaptive, business objective-driven optimization is significantly different from the localized, siloed optimization that operators carry out today. It uses a richer set of information, a business focus and new computing techniques to produce solutions aligned with business objectives. Key features include:

***Automation***

With automation, operators can carry out continuous optimization rather than manual, one-off optimization at the asset planning stage. Continuous optimization is a must in a virtualized network, for example. Even with the best predictions and engineers working on an initial optimization in the virtualized network, the dynamic nature of business, service and traffic trends will destabilize it almost immediately. Optimization also needs to be a continuous process in a customer experience context where operators want to provide personalized experience on a per-user, per-service basis, given the variable nature of customer demand.

***Flexibility***

The optimization approach should enable operators to change their business goals and the variables they are optimizing for quickly and easily. Services are not born equal from a revenue-per-bit perspective, and operators may have good business reasons to re-optimize the network to reflect different service priorities several times a day. The optimization required for a video service, for example, is not the same as that needed for email or server backup.

Sudden spikes in demand for capacity, caused by a sporting event, the weather or other predictable and unpredictable phenomena may cause the virtualized network to spin up new machines and functions. This action changes the design of the network, increases power consumption and needs to be intelligently guided by optimization to ensure the most efficient placement of new workloads according to their latency, security, cost and other requirements.

***Ability to Handle Multiple Variables***

As a corollary of flexibility and in contrast to Excel spreadsheets or linear tools that can only optimize across one or two variables, operators need an optimization approach that can handle multiple variables. There are three main types of variable – business, technical and regulatory – but the number of potential variables within each category is growing, and operators must be able to choose any combination of them that help them achieve specific business goals. For example, when operators decide what their objectives are for the delivery of a specific customer-facing service, they should be able to create sophisticated models that can optimize for multiple variables, such as latency, capital cost of equipment, revenue, margin and operational cost, at the same time.

In fact, the weighting of different variables may change over time: Capex will be important at service deployment stage when deciding on the minimum hardware required to roll out the service while meeting performance objectives, but opex-related variables will come into play when the service is running in production. Continuous optimization requires the constant adjustment of multiple variables and the adaptive optimization approach should support this. Network virtualization is introducing new variables that must be taken into account, such as power consumption, geographic footprint and last-mile placement. Multi-variable support gives operators much richer, more fine-grained control over their business objectives.

***Centralization***

This is the point where all stakeholders, whether they are in a business, service or network domain, can come together and decide holistically on their optimization objectives and how each of their domains contribute to them. These objectives should be business-driven (see **Figure 3**), and this approach will require constructive collaboration between all parts of the organization to understand which optimizations drive value and the trade-offs between different sets of variables. This interaction should create a knowledge feedback loop that continuously improves the optimization process and ultimately, business performance.

**Figure 3: Business Objective-Driven Optimization**

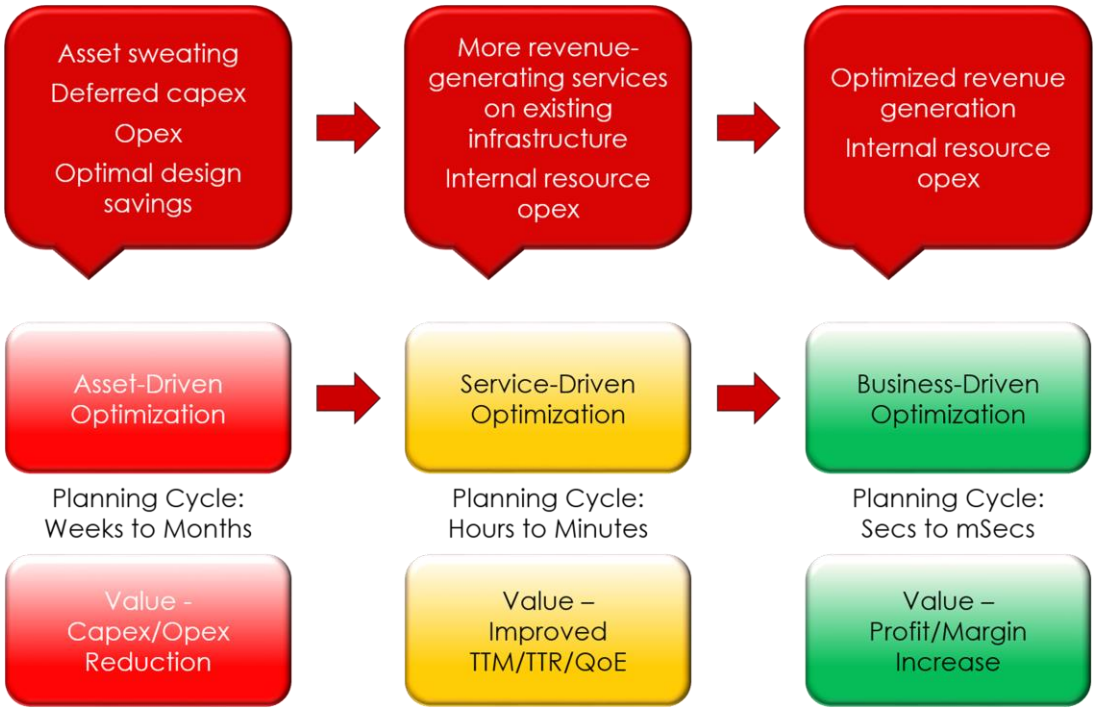


*Source: Aria Networks*

# A Platform for Business Objective-Driven Optimization

## Key Features of a Centralized Optimization Platform

In order to adopt an adaptive optimization approach, operators need a centralized optimization platform to support it. The platform should serve those stakeholders engaged in defining what "good" looks like for the business – its profit objectives for customer-facing services, for example – and those who turn business objectives into automated optimization models that can be applied in a continuous, institutionalized manner to the virtualized network.

A successful platform for continuous optimization has several key attributes:

- It can be fed automatically by analytics collected from the various data sources required by the different optimization goals it supports. It should be able to link to **any** data source in **any** domain – business, technical and regulatory – to support the flexible combinations of multiple variables that operators need to work with. Because the optimization goals are set *a priori* and driven by clear business goals, the optimization platform can selectively collect and use analytics: it should not need to sift through the big data noise.
- It enables the creation of sophisticated **optimization models** on a per-customer, per-service basis that support optimization across a range of variables from business, technical and regulatory domains. Such optimization models provide a common way of handling multiple variables and setting the policies that will drive the optimization of individual customer-facing service instances depending on customer experience, revenue, margin, opex, supply chain cost and other considerations.
- It supports **automation**. Through its optimization models, the optimization platform automatically drives the actions needed at different layers of the network and across the virtualized network to fulfil each model's objectives. The models trigger other operational systems in the lifecycle management cycle, for example, the service orchestration system. The models decide why, what and where changes should be made while the service orchestrator executes those network changes. As the function is increasingly responsible for monitoring the network, the service orchestration system is also a source of analytics that feed the optimization platform and alert its models to deltas as their optimizations decay over time. The link between optimization platform and orchestrator should be automated, enabling a feedback loop between the two systems that drives continuous optimization.

In future, orchestration is likely to be highly distributed and a certain amount of autonomous optimization will be delegated out to individual network elements. But a central "brain" function that can balance optimization needs across the entire environment, make holistic decisions and capture them in automated models will still be needed at a higher level of abstraction.

## The Role of AI

Critically, an optimization platform will be most effective if it uses **artificial intelligence** (AI) and machine-learning techniques to arrive at the "perfect" optimization.

Machine learning is the field of study that enables computer systems, such as an optimization platform, to learn without the need to explicitly program them.

When humans program systems, they inevitably introduce their own assumptions about the environment and data that, in an optimization context, could bias the results and lead to sub-optimal choices. An AI-driven system can be given a set of goals – for example, the combination of variables it needs to optimize for – and then left to make its own way toward a solution without the need for coding. Any change of inputs immediately changes system processes, again without any programming.

And an AI-based optimization platform can reach its conclusion at speed. AI techniques can support massively parallel operations compared to conventional, linear programs, making an optimization platform that uses them significantly faster at finding optimized solutions. For example, a service provider optimizing IP topology across a global optical network to reduce capex estimated that it would have taken 35 years to derive a solution using linear techniques. Using an AI-based optimization platform, a solution was reached in just under three hours.

Machine-learning techniques do what they say: they learn. Frequently, human experts use insights and experience to carry out an action without being able to articulate exactly why or how they have done it. An AI-based system can capture and learn that behavior, helping to codify complex and subtle tasks. Again, it can do so at speed: it can emulate complex time-consuming processes, such as building algorithms orders of magnitude faster than humans.

An AI-based optimization platform should have a machine-learning algorithm that can take in a definition of what "good" looks like in terms of meeting business objectives for the delivery of a particular customer-facing service and then generate likely candidates, learning from and refining them until it comes up with the best solution. This is best achieved using an algorithm and learning technologies based on genetic and evolutionary techniques, which also play well with the adaptive response to change that operators want – through optimization – to achieve.

It is critical that the optimization platform is based on AI and machine learning if it is to meet successfully all the requirements of an adaptive approach.

## The Need for Culture & Process Change

Having an adaptive optimization platform is not in itself sufficient to enable operators to handle rapid change in a non-disruptive, continuous manner. They need to commit to continuous optimization as a key business goal and guiding principle: this will require top-level support from the business. Senior management must ensure that optimization models become a common language between roles, a key driver of business performance and a motivator for cross-domain collaboration.

# The Benefits of a Business Objective-Driven Approach

The benefits of an adaptive optimization approach are already clear from its application in the physical network:

- Two European operators that have investigated the value of incrementally and **continuously** re-optimizing the network to prevent stranded resources calculate that they could save 12 percent of network capex annually.
- One large service provider using a **centralized** approach to cross-layer optimization estimates that it is making capex savings of between 15-20 percent compared to a time when it optimized different layers of the network separately. This service provider has also used optimization to significantly reduce the overbuild of expensive optical links.
- A North American operator and a global content provider also overbuilt their networks because they couldn't accurately predict network behavior under failure conditions. Applying **AI-driven** optimization enabled them to evaluate billions of possible configurations in hours to come up with solutions that resulted in a substantial increase in network return on investment.

Operators are also applying optimization in the virtualized network with encouraging results:

- A leading European operator carrying out a proof of concept for the automated delivery of virtualized enterprise services used machine learning to identify traffic trends and anticipate capacity problems in real time as new services were being delivered. This helped reduce a 14-week delivery cycle to minutes, providing timely instantiation of additional NFVs to avoid degradation of existing services.
- A Tier 1 Asia/Pacific operator used business objective – driven optimization in a 5G proof of concept to find a solution that maximized profitability when designing network slices to support a city road race. The multiple variables taken into account included commercial commitments to sponsors, regulatory obligations to support emergency services, low-latency service demands and dynamic traffic loads. The proof of concept demonstrated optimization guiding the automated flexing of the network in real time to create and remove 5G network slices to protect revenue, customer experience and license conditions.

# Conclusion

In a time of massive and relentless change on all business and technology fronts, operators need a means of evolving their services and networks quickly and efficiently. Optimization is a critical mechanism that helps operators adapt to and intelligently drive change. But it needs to be driven by business objectives, which means elevating optimization as a practice out of its current narrow remit within network planning and into the organization as a whole, where multiple stakeholders can contribute to it. Holistic optimization across domains and organizational units is vital for effective service delivery and customer experience management. Continuous, automated optimization is a necessity for virtualized networks where network design is dynamic and constantly changing.

Operators that adopt an adaptive, business objective-driven optimization capability will have greater visibility of and control over their operations, resulting in better decision-making and an ability to exploit new market opportunities. It is already clear that this approach drives significant capex and opex savings compared with traditional, siloed optimization practices, but this is only the beginning.

As optimization becomes the focal point for business intelligence in the operator organization, driving a range of performance, cost and quality improvements into the network, it will create competitive advantage and differentiation opportunities. These are increasingly difficult to achieve by other means and are immensely valuable in the fast-commoditizing, ultra-competitive telecom market.

# About Aria Networks

Aria Networks enables service providers in the Digital Economy to maximize their profit potential by automating and optimizing how services are delivered, according to customer demand, operational constraints and business goals. Aria helps service providers rapidly deliver services profitably and at scale, eliminate unnecessary capex and opex, and ensure a high-quality customer experience. Powered by pioneering AI technology, Aria enables service providers to identify business-optimal solutions for service and network designs from billions of possible scenarios, in near real time.

For service providers deploying SDN and NFV, Aria provides the "brain" required to drive orchestration platforms with "best for business" resource and design choices. The company is headquartered in the U.K. with offices in Silicon Valley. Customers include BT, Level 3, and some of the world's most valuable Internet brands. Aria partners with leading ISVs, integrators and vendors, including Amdocs, Ciena, Ericsson, Packet Design and TierOne OSS.