The 2019 Guide to

Modern OSS

Produced by

Table of

Contents

Welcome

Getting involved in your first OSS project?

Start here.

This introduction to OSS assumes you know very little about either telecommunications operational processes or the software and systems used to run the network.

Starting with the basics, it will explain why OSS plays an important role in ensuring users of communication services enjoy a good experience while the operators of the networks run an efficient, profitable business.

Any questions?

Ask the author

If you have any questions after reading this guide, go ahead and contact me via email (ossline@outlook.com).

Just ask.

I'll respond to as many questions as I can, and if I don't know the answer, I probably know someone who does.

Chapter One

What is OSS

and why is it needed?

Operational Support Systems (OSS) is IT for running communications networks.

OSS includes software, hardware, integration between systems, and business processes. As a collection of integrated applications, OSS supports the design, build, monitoring and assurance of both the communications network as a whole and the individual customer services that make use of that network.

OSS encompasses many highly technical network management processes but ultimately its purpose is to ensure the network is efficient, services are profitable, and customers are happy.

A lot of concepts have been introduced there already, so let’s start by looking at the acronym OSS: What does OSS stand for?

What does OSS stand for?

OSS stands for either “Operational Support Systems” or “Operations Support Systems”.

“Operational Support Systems” is perhaps more commonly used. But don't worry, you won't sound dumb if you use either *operations* or *operational*.

Let's break it down further...

Operational/Operations

Relating to the day-to-day tasks of supplying and supporting communication services. Getting technical and infrastructure jobs done. Running the network and services. As opposed to the business of selling, marketing or billing (which, as we will see later, are tasks that belong to Business Support Systems, BSS).

Support

Enabling and improving the service provider's operational activities: Automating operational tasks; executing them faster; making them consistent; and tracking progress/results.

Systems

One or more distinct software applications, that are responsible for doing specific OSS jobs, running on servers, or on devices installed in the network, or executed in the Cloud.

The role of OSS

OSS includes many applications that a service provider requires in order to perform 'back-office' activities. The service provider's 'OSS environment' will include many (from tens to hundreds) of separate OSS applications, each responsible for their own part of the businesses' operation. Usually OSS applications are function-specific, owning part of the entire operations process such as monitoring for faults, fulfilling service requests, and so on.

Being focused on the network and services, OSS is used by network planners, service designers, service operations, network operations and network engineering teams. Increasingly marketing, product managers and senior staff under the CTO or COO also rely on OSS to gather data on the network state to plan strategic changes.

***What is** Telco, Service Provider, Operator, DSP, CSP?*

These are all common names for what most people in the industry now call "communication service providers" or "digital service providers": A business that maintains a network and sells telecommunication services.

OSS & BSS

If there are 'back-office' tasks there must also be a 'front-office' and indeed there is. More accurately called Business Support Systems (BSS), these are a separate set of applications supporting commercial, revenue and customer-relationship activities.

Combined, OSS and BSS cover the entire footprint of specialized IT applications a service provider will use to run a network and sell services.

To provide a good customer service – and operate a profitable business - OSS and BSS must work together.

***What is** Back Office and Front Office?*

An analogy to businesses that have customer-facing staff on the front-desk and technical boffins making stuff in the workshop out the back.

The analogy with front-office is stretched to include other points of interaction with the customer, such as call-centers and web portals.

The basics of

Communications networks

Let's introduce an example now, and also explain why a modern communications network needs OSS.

A communications network is responsible for getting data from A to B. A and B might be mobile or land-line handsets, computers or servers. The data being carried from A to B includes our voices, a text message, video, pictures, emails and files. In modern networks everything being communicated between two points is data.

***What is** a device?*

We will use the term *device* to mean any piece of equipment that is responsible for receiving, routing, switching, processing or transmitting data. Other commonly used terms are *Network Element (NE), Managed Element (ME),* and *Managed Device (MD)*.

So, the things that make up a network are mainly concerned with receiving and transmitting data to carry it across the network.

A network includes hundreds or thousands of devices whose job is to move our data around.

Generally, the further the data needs to be carried and the more people or things that are communicating, then the more devices there are in the network.

Each device in the network is relatively simple. It either prepares or receives data and then determines where to send the data next, to get one ‘hop’ closer to the destination. It might also do other tasks such as checking the integrity of the data, making security checks, caching or compressing data and manipulating the data flow to put it on to a different type of network (for example, to take data off a radio interface and on to a fiber-optic network).

Devices come in different ‘sizes’ both in terms of how many other devices you can connect it to and how much data can be sent down each connection (commonly called bandwidth).

You will often find many smaller devices in the 'access' network where customers connect to the network, then fewer (but larger) devices in the 'core' that carries vast quantities of data between regions.

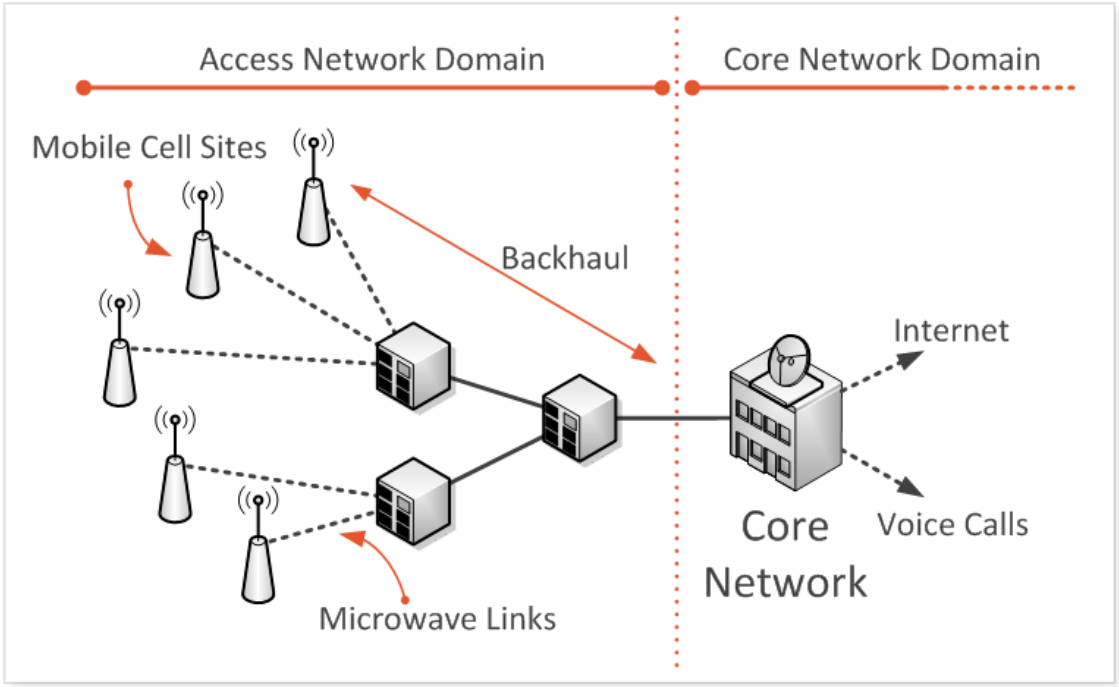

Domains

The concept of ‘access’ and ‘core’ networks has already been mentioned, so it’s worth explaining the terminology used to describe parts of a large communication network.

Access Network Domain

This network is concerned with connecting to customers. It includes the likes of radio networks for mobile phones, DSL and cable networks for landline services. Access networks can be characterized by many relatively low capacity connections to individual customer sites and radio access points. A higher capacity ‘backhaul’ aggregates access traffic and connects to the core network domain.

Core Network Domain

The core network handles routing of data between networks, between CSPs, and over geographical regions. It also provides centralized resources for delivering services and content to customers. The core network is made up of fewer devices, but devices able to support far greater bandwidth and numbers of services.

User Domain

Customers have their own network or, at the very least, devices such as smartphones and routers that connect to the CSP’s access network. The user domain has become increasingly relevant to the CSPs operations as more types of services are configured to the same customer or location, and increasingly management of the user domain becomes more frequent thanks to Network Functions Virtualization.

Edge Network

While all the above domains comprise further sub-domains, the edge network is perhaps the most commonly used refinement. It refers to the equipment and services exposed at the edge of the core network. This is often where most of the service delivery functions live (rather than long-haul or transit traffic carrying devices). It is a term particularly used by those CSPs that provide network transit, high bandwidth and ‘leased line’ services to business customers and other CSPs by connecting customers directly in to the core network, rather than via consumer access technologies. In such cases the edge describes the CSP’s service technologies and points-of-presence for connecting to them.

Network Virtualization

Until recently, the devices in a network were all individual, single purpose devices built to fulfil their specific role. If you wanted to change a device's role, or even just scale it up to support more services, this would mean replacing it or upgrading its hardware.

Similarly, the network was built using fixed technologies and protocols, with routes across the network being defined once and changed rarely. If you wanted a new network protocol to handle traffic differently, you would need to go through a major network upgrade, including adding new devices and management systems.

Today, networks can increasingly be *virtualized*.

Network Functions Virtualization (NFV) and Software Defined Networking (SDN) share a common set of goals, while offering quite different capabilities.

NFV can be thought of as the ability to create network devices in software, running on a generic server, with the ability to scale resources up or down as needed. The network function is separated from the physical device on which it runs. In the NFV world, the virtualized device is commonly called a Virtual Network Functions (VNFs).

SDN allows connectivity across the network to be defined and reconfigured by software. This allows the physical underlying network to be abstracted, with connectivity operated as a virtual resource to deliver services.

With virtualization, network assets can be used more flexibly, capacity allocated, reallocated, overbooked and so on, to ensure it is available where and when it is needed.

Being software driven, both the network and data centers can be smarter. Smarter means analyzing end to end traffic, predicting future trends and dynamically adjusting the network as needed, rather than simply employing static protocols and static resources.

We take a more detailed look at both NFV and SDN in the Chapter 4 section *Rebuilding the network with software*.

Managing

the network

To operate a network, you need to know what devices you have, where they are, be able to check their configuration and reconfigure them to 'route' data across the network (get it from A to B in the quickest or cheapest or most reliable way).

You need to know how the devices are being used by customers, so you can be sure you're billing them correctly.

And you need to keep an eye on the devices' state to check there are no faults while ensuring their secure operation.

The *International Organization for Standardization* has a name for these network management processes – FCAPS – and it's a good place to start.

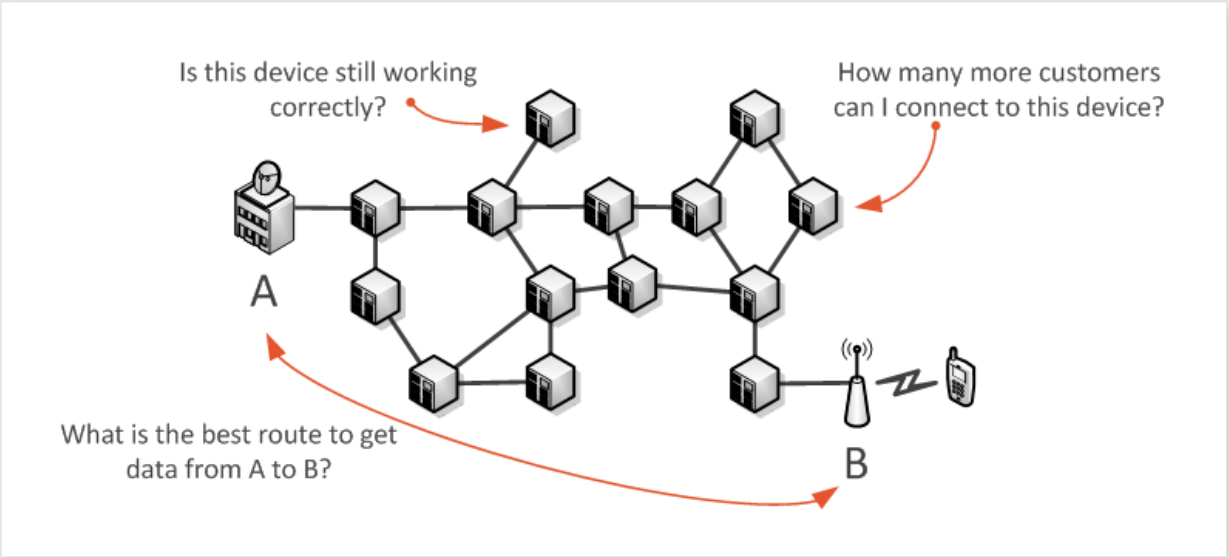

The FCAPS guide

to network management

FCAPS is an acronym for Fault, Configuration, Accounting, Performance, and Security.

These five simple categories cover a broad set of network management processes that must be addressed by the CSP to ensure the network is operated reliably.

Fault

Things go wrong. Devices fail. Connections between devices get cut. People accidentally reconfigure one thing which breaks another thing.

Fault management is about responding to ‘alarms’ from the network (as many devices can remotely report a fault, which is helpful) and diagnosing harder to find faults.

A big part of fault management is tracking problems, identified by the network devices or the customers, through to resolution.

Configuration

Devices need to be configured to function as required and to enable them to talk to other, networked devices. Much of this must be done remotely, as a device may be hundreds of miles away from the engineer responsible for carrying out the work.

Configuration can include minor changes to improve performance, routing of a new customer service, updates to connect new bits of network, and major changes to devices’ software or firmware.

Accounting

Account management is the act of collecting statistics on how a service is being used, primarily for billing purposes. BSS applications are most involved in this management process.

Increasingly there is a need for OSS to be involved with Accounting - More sophisticated services and particularly their ‘Service Level Agreements’ require greater understanding of how a service is being supported by multiple devices across the network. Therefore, Fault and Performance management data will feed in to the Account management.

Performance

Network devices and individual services must be monitored to ensure they are delivering the capabilities they were configured to deliver. Statistics can be collected from the network to aid understanding of how well data is flowing. In a modern network, performance management is as critical to operations as fault management. If service performance degrades for a customer, the impact for them can be as disruptive as a physical fault.

Security

Networks must be secured to permit only authorized staff to configure devices or services. Security also ensures customers only have access to resources they are entitled to use. Malicious attempts to use services or access customer data are blocked.

In modern networks there is a cross-over in the roles of Security and Configuration management: As services become less fixed to a single location or device, customers are accessing services from many locations on many devices. Services no longer benefit from the simple security of being 'hard wired' to a single premise or terminal. Security must be managed for every customer interaction with the network.

Sounds simple enough?

If you run a network with just a few devices from the same supplier, you probably don't need sophisticated OSS applications. You can get by with the tools that the supplier gives you and some simple spreadsheets or databases to keep track of things.

For example, from your workstation you could log-on to each device directly, through a 'command line', to see its state and reconfigure it. And you could manage many aspects of network security with simple desktop IT applications.

But in a network run by even a small telecommunications company, things are rather complex.

Scale - the number of things you need to manage and the number of decisions you need to make - quickly makes running the network extremely challenging.

Big networks

and how OSS can help

How big is a CSP's network? The answer obviously depends on the region they are serving and how many customers they have.

Even a small European telecommunications company probably has thousands of devices to manage, across several different technology types like IP, Ethernet, optical and radio. In addition, they need to manage each service across that network too, for hundreds of thousands, or millions of customers.

A typical national North American CSP would be at least ten times bigger.

A large Asian operator, another five to ten times bigger.

You can't keep track of a quarter of a million network devices and one hundred million customers using a spreadsheet.

Here's why...

Many locations

Network devices are distributed over a wide area, such as a city or country: Device can be installed in roadside cabinets, on antenna towers, in telephone exchanges, and many other types of site.

The CSP needs to know what equipment is installed, where, so that planning can be carried out without requiring a visit to site to determine what devices are available or how much physical a site must accommodate more devices.

OSS applications employing databases, floor plans and mapping services exist to manage these records.

Many vendors

Most CSPs are supplied devices by two or more device vendors. Each vendor can supply their own *Network Element Manager* (NEM) or *Element Management Systems* (EMS) which connect to the devices to support basic configuration and fault management.

However, these applications only offer their full functionality for the vendor's own devices. To configure a service across the entire network, the user must tap details in to each vendor's NEM or EMS. Each has its own user interface, its own way of working, and this naturally can lead to a complicated process.

Many OSS applications are designed to do their job on multiple vendors' devices, ensuring consistency and streamlined network management processes.

Many technologies

Most CSP networks are made up of several different technologies. Networks still rely on technology deployed many years ago to deliver services today.

A new, modern network would likely comprise radio devices for mobility, IP devices for flexible data routing, optical devices for its high-capacity bandwidth, and servers processing data and hosting content.

Since a service uses several technologies to carry data from A to B, how these technologies interact can be very complicated and subject to many rules and constraints.

OSS is used to manage the complexity of network planning and service design by using templates, rules and step-by-step processes to guide users.

Many users

If you've ever tried to get multiple people to work on the same spreadsheet file, you know how many problems arise. The file may be locked when you want to update it. One person might overwrite another person's updates.

In a CSP, dozens or even hundreds of people are involved in network operations at the same time. Planners, engineers, support staff, marketing, all needs to manage the network and report on the current state of devices and services.

OSS applications are designed to share data with many users, implement data update rules, and control processes that need input from different users at different stages.

Many tasks

FCAPS might only be five letters, but that hides a huge number of tasks to be carried out by the CSP. TM Forum's 'map' of CSP business processes – Enhanced Telecoms Operations Map (eTOM) – comprises over thirty process categories. Each category will itself involve several monthly, daily or constant management tasks.

Automation of these tasks is frequently the responsibility of OSS applications. In many cases OSS applications can cut the time it takes to complete a repetitive task, like routing a customer's service, from hours of manual work to just a few seconds.

OSS also makes tasks consistent, enforcing rules about how they are carried out, and providing updates to other systems such as the BSS (for example, to start billing) and management reporting.

Many options

Each individual device may be simple, just routing or switching data, but they are configurable and flexible. Which devices should you connect? How much capacity should you put in each connection, considering the cost and customer demand? What's the best 'route' through the network for a service considering cost, bandwidth, and reliability? What happens if there's a sudden peak in demand somewhere? Which customer services have priority over others?

These are difficult decision to make if the network is bigger than you can draw on a single piece of paper. A big part of OSS is analysis of networks and services, resulting in intelligent decision making, optimal network design and efficient service routing.

Many services

Let us not forget that the network of a CSP exists to make money by selling communications products to customers.

In the not-too-distant-past, most CSPs made money from just one communication product: A plain, old telephony service.

The modern trend is for more services, more customization of services, and more ways to access those services. CSPs want to quickly create new services, and test them in the marketplace, to keep up with rapidly changing consumer trends.

With its detailed network records, OSS can be used to identify services that can be sold in particular regions. By simplifying and automating operational tasks, OSS can enable new services to be constructed from 'building blocks' of network resources.

OSS to the rescue

In summary, Operational Support Systems are IT tools and applications used to monitor and manage large communications networks.

The challenge of running an efficient and profitable network, with even just a few hundred devices and customers, means OSS is a necessity for all communication service providers.

In this chapter we have briefly introduced how OSS can address the operational challenges of network management.

In the next chapter, we'll look at the different types of use cases for OSS and the business processes they support.

Chapter Two

OSS use cases

What jobs does OSS take care of?

Chapter One of The Guide introduced the challenges of managing a modern communications network. You may already be thinking that there are quite a lot of things that OSS needs to do to address those challenges.

You'd be right.

This chapter will give you more detail about the role of OSS and will introduce some typical use cases that an integrated OSS environment can support.

OSS is flexible, supporting processes and requirements unique to each CSP. The use cases in this chapter serve as an example of the role of OSS. They also set the scene for the next chapter's examination of individual OSS applications and functions by demonstrating the scope and complexity of these processes.

You will notice as we step through these use cases that most cut across the FCAPS categories described in Chapter One. The focus may shift depending on the specific OSS process, but doing OSS right means taking in to account each of the FCAPS categories.

Two basic processes,

many tasks

Let's look at two network management processes that highlight the jobs, data and integration that is typical of an OSS environment.

OSS was once described in very simple terms as Orders Down, Faults Up.

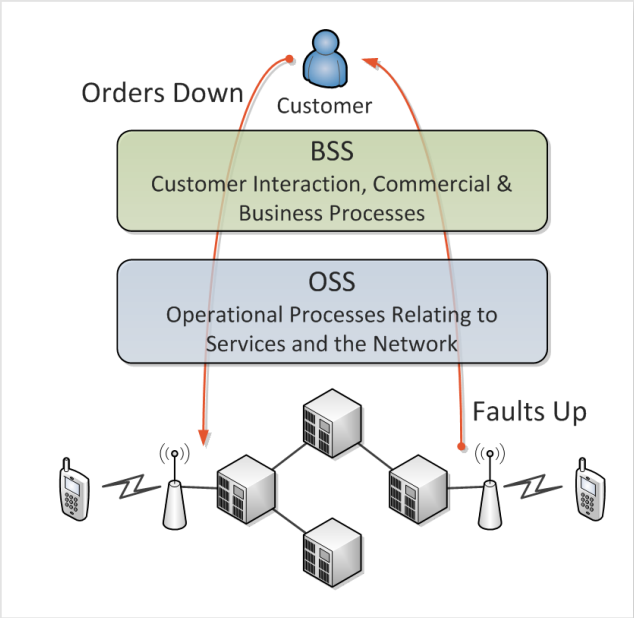

Orders down

Visualize the process flow of a new service order from a customer.

The order 'flows down' from the customer at the top (where, of course, they belong), through the BSS layer, through OSS, and finally down to the network where the necessary configuration and activation changes are made.

These sorts of processes, called *service sulfilment*, are a key part of FCAPS Configuration. But it's much more complex than simply connecting to a few devices to configure them.

The BSS applications are responsible for customer-facing jobs like correctly capturing the order details, notifying the customer of costs, timescales and engineer visits, and starting the billing process when the service is turned on.

As far as the OSS layer is concerned, the BSS layer is required to provide a detailed customer order comprising data such as service type, optional service parameters, the customer ID, and the origin and termination of the service.

Sunny day service fulfilment

Not every service fulfilment enjoys a ‘sunny day’ where everything goes right. In the event of an issue such as a shortage of network capacity (either physical connections or available resources) then additional jobs would be started to either redesign the service or upgrade the network capacity. And, of course, the delay and relationship with the customer would have to be managed.

With a well-designed network and a well-understood service fulfilment process, fulfilment can be automated to involve little or no human intervention, other than making physical connections. But equally, these jobs can be completed by staff at computers working directly with individual OSS applications.

Depending on how ‘service-ready’ the network is, and the automation of the OSS applications, a fulfilment process could take minutes or days to enable the customer’s service.

Faults up

Still with the customer at the top of our imaginary stack, and the network at the bottom, *faults up* refers to processes that are started by a problem being reported by a network device.

A more common term used today is the more positive sounding *service assurance.* It's also a more accurate description as the primary objective of managing a network fault is minimizing disruption to customers' services.

These processes naturally fit in to FCAPS Fault management, but also involves Accounting, Configuration and Performance tasks. So, a bit of everything.

A working device can report all sorts of problems such as a bit of its equipment failing or a loss of signal on a connected cable. Even if the entire device fails and cannot report its own state, other devices connected to it will start reporting faults.

A single network fault can result in many 'alarms' being reported. Getting to the 'root cause' of a fault and solving the problem is the first challenge faced by OSS.

A cable is accidentally cut by some workmen. Seconds later, the IP devices that are making use of the cable start redirecting their data packets, finding an alternative route around the cut.

The network has done what it can to maintain service, but it's up to the OSS to identify the problem, take smarter remedial action, and start the process of fixing it.

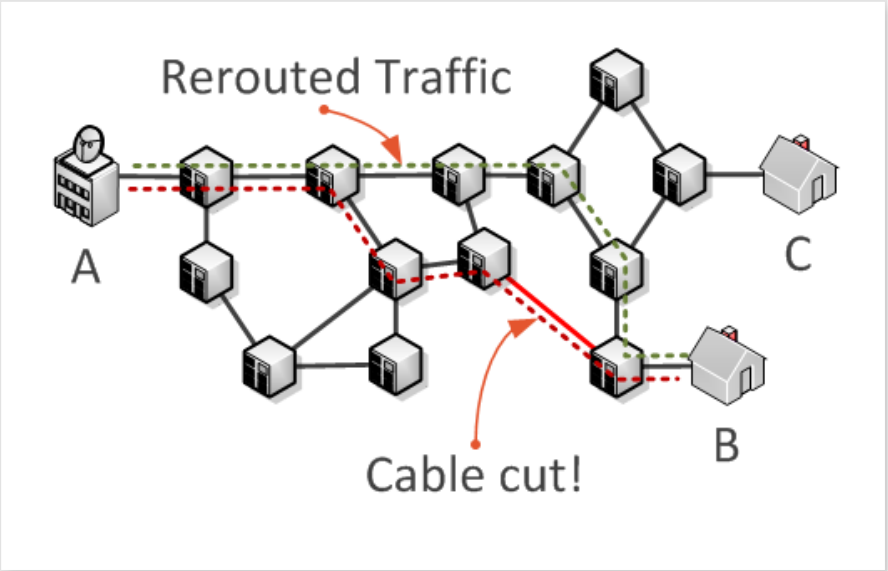

The wrong way to do it:

Faults down

Let's not forget that there's also a *faults down* scenario, where the customer reports the fault. That means the customer found the problem first, the network wasn't designed to handle whatever the fault was, and the CSP didn't react quickly enough to reduce its impact.

Not a good scenario.

With modern OSS applications and a well-planned network, that should be a rarity.

The appeal of CEX

When we look at specific OSS applications in the next chapter, we will see that there are solutions for discovering faults, fixing faults, restoring services, and so on.

Recently the telecommunications industry has been considering how these processes and applications can work together to ensure a good *Customer Experience* (CEX).

Also known as *Customer Experience Management* (CEM), these uses case cross OSS, BSS and network boundaries to support:

- Understanding of the individual needs of customers
- Fast, error-free delivery of services
- Identifying and resolving problems before they happen
- Measuring and reporting on customer experience metrics

While every interaction with a customer and their services is a chance to deliver a good experience, the principal focus for OSS is on a more customer-centric and pro-active approach to planning, fulfilment, service assurance and performance management.

Planning to succeed

To be sure that your service fulfilment and service assurance processes go according to plan, it is essential that the network is correctly designed and that the CSP has planned for the unexpected.

Consider the previous example of the cut cable. The traffic on the network was rerouted by the devices which noticed that the existing route was broken. This inevitably meant that the traffic running over the cut cable was moved to routes that were already carrying their own traffic.

If the network wasn't designed with this possibility in mind, the rerouted traffic may overload the devices on the new route. An even bigger issue would be if a combination of multiple faults meant that there was no valid route at all for the affected traffic, resulting in a total loss of services for many customers.

What if?

What if there is a major event like a concert or festival with a big increase in local mobile and WiFi traffic for a short period of time?

What if a sales campaign is unexpectedly so successful that demand grows faster than the strategic plan for capacity upgrades?

What if a big business wanted to connect all their offices with high-quality, reliable, broadband links?

These are good problems to have! But a CSPs needs to be ready for success too, as much as they need to be ready for a network fault.

Rise of the machines?

The promise of many modern network technologies, even the now very well-established IP technology, has been greater self-configuration, self-optimization, and self-healing.

IP, for example, will attempt to reroute data around major network faults. And mobile networks can self-adjust properties of its cell sites to compensate for interference, congestion and outages.

This built-in capability is valuable because it means the network can react quickly, without human intervention.

But CSP networks are big and their paying customers demand a high quality of service. Leaving critical decisions like traffic routing and fault management *only* to the devices themselves does not result in sufficiently predictable or reliable network management.

Intelligent design

Even the smartest devices have a narrow view of the network: Themselves; their neighbors; maybe the end points for a traffic route. They certainly have no understanding of the business context of an issue, and no predictive capability to get ready for a future problem.

Those issues of network scale and complexity discussed in Chapter One mean some greater intelligence is needed to ensure the network is running at its best.

Strategy and Tactics

Planning is simply taking time to think, to design, before doing something.

Some planning, let's call it *strategic*, considers requirements months or years down the line. Other planning tasks search for optimal solutions to an immediate challenge, which we'll call *tactical* planning.

The best networks are planned. They are efficient, reliable and *profitable*.

The foundations of OSS:

Data and engineering

For any of the OSS use cases described in this chapter, and for any FCAPS task, two fundamental elements are required to get the job done:

- A source of information about the network, services and customers
- A means of making a change to the network and customer services

However smart the analysis or automated the process, without accurate data to make decisions and an efficient way of configuring devices, services, and resources, the CSP's OSS cannot hope to do a good job.

In the next chapter we will look at specific OSS applications, each of which supports some part of the use cases described in this chapter. And, as with these use cases, we will see how having the right data management and engineering tools in place is an essential prerequisite to the deployment of OSS applications for service fulfilment, service assurance and planning.

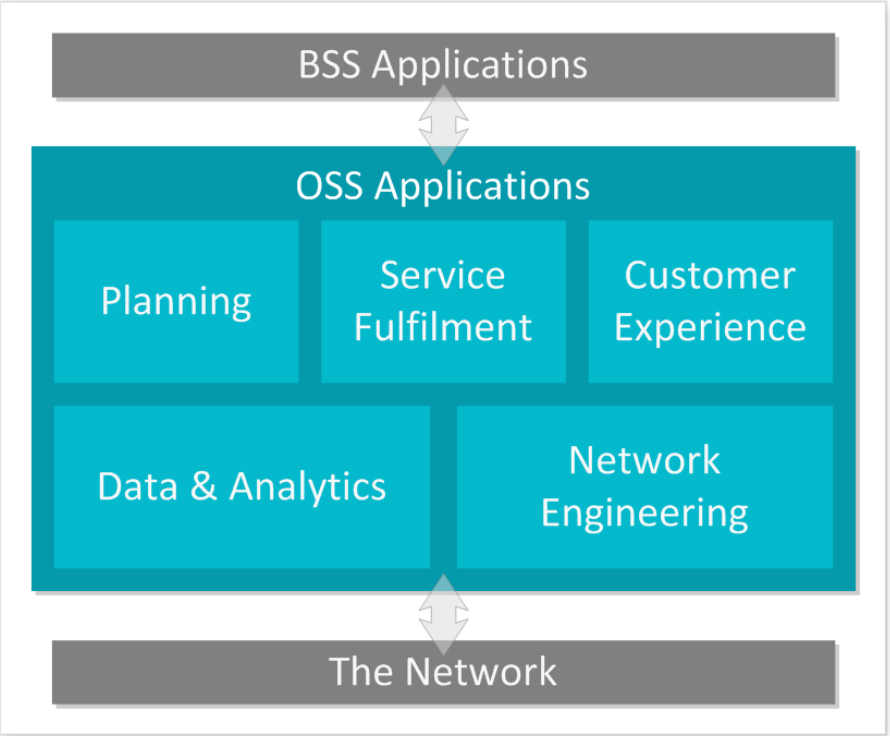

Chapter Three

OSS applications and functions

The building-blocks of OSS

In Chapter Two of the Guide we looked at several example use cases, which demonstrated that typical CSP operational processes carry out a large number of tasks, and most processes will comprise some element of each high-level FCAPS category.

We also saw in Chapter Two that many CSP processes go in one of two 'directions':

- From the customer, through BSS, then OSS, eventually affecting a change on the network – Service orders being the classic example
- From the network, to OSS, then to BSS where it affects customer-facing activities – Examples being fixing network faults, or the roll-out of new technology offering new services

There are distinct applications that support each step of a CSP's processes. Some are universal; relied on no matter what the service, network or objective. Others are very much specific to a particular job.

This chapter introduces the major categories of OSS application. It also adds some meat to the bone with a list of the most common applications in each category.

Standard definitions

of OSS applications

In this chapter we want to look at OSS applications and functions. What OSS applications exist, and what do they do, to support the CSP's use cases and processes?

Communications is a highly technical and complex science, so it should come as no surprise that people have attempted to formally define the role of OSS applications.

FCAPS

We looked at FCAPS in Chapter One. FCAPS broadly defines the responsibility of OSS and BSS systems to implement fault, configuration, accounting, performance, security management. But it does not define processes, functionality, or applications.

TMForum eTOM

TM Forum spends a lot of time defining and naming things.

Their eTOM Business Process Framework provides a very detailed 'map' of operations and network management processes.

Comprising over thirty process categories, a discussion of eTOM is out of the scope of the Guide. If you want to get a more comprehensive view of OSS/BSS processes than the examples in Chapter Two, eTOM is a good option.

eTOM, however, does not describe the applications responsible for carrying out these processes.

TMForum TAM

The Application Framework (TAM) is as detailed as eTOM and attempts to identify specific applications to manage each stage of each process. Should you find yourself defining the architecture of an integrated OSS environment, TAM can give you a means of describing each component.

However, for this Guide, we need a higher-level and more pragmatic way to organize and describe OSS applications.

A pragmatic definition of OSS applications

The level of detail that some definitions go in to, particularly eTOM and TAM, is quite breath-taking and, for many purposes, too detailed.

A big challenge you will face is the fact that commercially available OSS applications do not neatly map on to any of these definitions. A suite of applications from one vendor will most likely address a few functions completely, but maybe not a full end-to-end process. The features they do have may also partly address other functionality.

OSS products have evolved over many releases to complimentary features as well as their 'core' functionality.

Keep it simple

It is therefore best to take a high-level view of the OSS application landscape when you are in the initial stages of an OSS project. Work with enough details to identify each vendor's product, to understand the primary roles of the systems you already have, and to create your project strategy.

Later, when there is a need to specify the system integration of OSS applications, and mapping business functions to processes and data, then an OSS Solution Architect or Business Analyst might want to use a tool as detailed as the eTOM.

OSS in the middle

OSS is concerned with the operation of the network and the technical aspects of services. Broadly speaking OSS sits between a CSP's other IT systems: BSS applications, and the management systems and controllers of the network layer.

The OSS layer is heavily influenced by demands and trends in the customer-facing BSS layer and the technology of the network layer, so we will take a look at these two layers first.

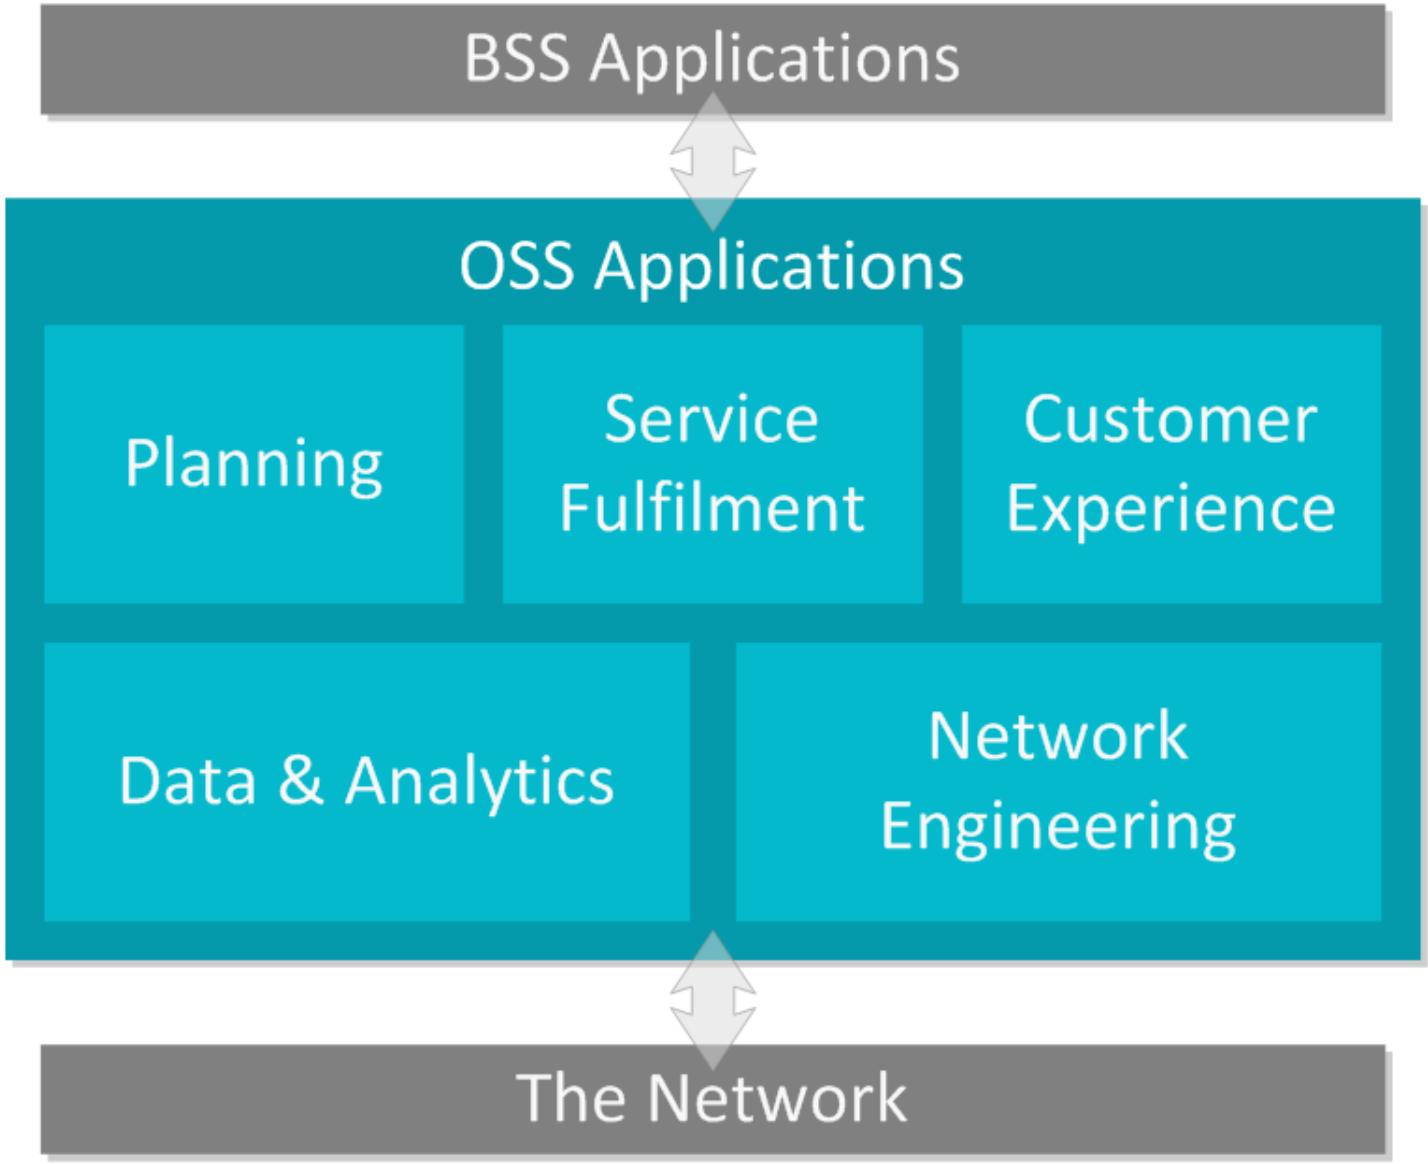
BSS Applications
OSS Applications
Planning
Service Fulfilment
Customer Experience
Data & Analytics
Network Engineering
The Network

The BSS layer

The BSS layer is concerned with the service's customer, commercial and contractual considerations.

The interface between OSS and BSS is usually pretty clear:

- Orders for new or modified services are collected by BSS and passed to OSS for activation on the network
- Network state and usage is passed from OSS to BSS for tracking quality and service usage
- And, increasingly, technical capabilities of the network are shared between OSS and BSS to facilitate service catalogs and design of products for customers.

It is convenient to assume that the BSS layer is completely isolated from the network by the OSS layer. Certainly, OSS applications are responsible for many tasks bridging BSS and the network, but bear in mind that the BSS does have its own interfaces to the network layer.

A classic example is mediation.

Mediation applications are responsible for taking data from the network relating to how individual services use resources. CDR (call detail records) and IPDR (Internet Protocol Detail Record) are two examples of 'raw' statistics available from network devices. Mediation applications are responsible for pulling this data from network devices, filtering and processing it in to a common format for use by BSS charging and revenue assurance applications.

There are only a few examples of BSS interfacing to the network, and those that do exist are almost all 'read-only', importing and analyzing data.

That said, analysis of network data and the trend of Big Data means this is a growth area: Understanding the current state or change in the network can be used to gain valuable insight in to both the customer's behavior and the quality of service they are enjoying.

Any non-trivial process that requires an understanding of how devices and services actually work, particularly if a change to the network is required, will reside in the OSS layer.

So, job #1 of the OSS layer is to mediate between the network and the business of keeping customers happy.

The network layer

While OSS is concerned with the operation of the network, the network itself is capable of significant 'self-optimization' while also supporting at least basic functions to enable manual configuration, fault finding and monitoring of individual pieces of equipment.

In a traditional optical, switched and IP network, the Network Layer is the devices, their operating systems, and their built-in interfaces.

These interfaces support integration between the network and OSS applications. Typical integration supports: Modification of the network for planning and service fulfilment purposes; Reporting of alarms from devices to the OSS layer for tracking and resolution; Real-time statistics on network utilization and performance, and service quality.

Emerging 'virtualization' technologies like SDN and NFV are starting to take some functionality and logic out of the devices, to be executed and controlled centrally. For example, today an IP device is responsible for working out the 'next hop' to pass data to get it from A to B. An SDN enabled device would defer this decision to a centralized controller.

This simplifies the line between the OSS layer and network layer, since OSS applications can now interface with a centralized controller for the domain. An SDN controller could make design and routing decisions that would previously have been carried out by an OSS application.

We will discuss the implications to OSS of new network technologies like SDN and NFV in a later chapter of the Guide.

The OSS layer

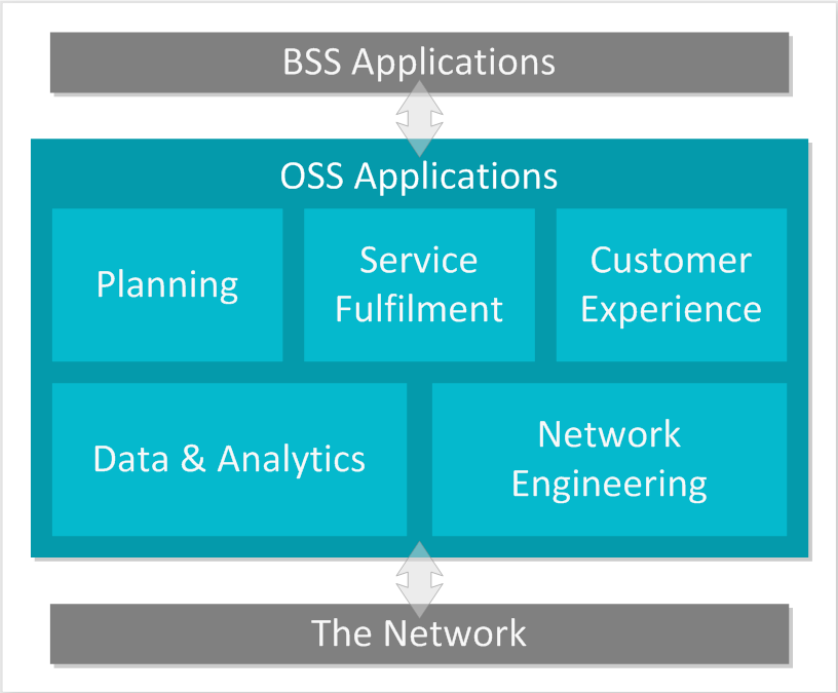

Let's get back to the OSS layer.

The trends in customer-focus, service quality and network virtualization must be reflected in the capabilities of the CSP's OSS environment.

OSS is as responsible for the results of customer satisfaction surveys as it is for a row of green lights in the network operations center.

The customer comes first, but we're going to start by looking at the network-facing applications. Why? Most CSPs don't start with a blank page: They have a network from which they need to deliver customer services, build new types of service and generate revenue. And many OSS applications depend upon or build on the underlying network-facing application.

We will look at the major categories of OSS applications, from network-facing to customer-facing. For each, we will introduce the main functions which approximate to the individual software modules you can buy or build.

So, as we present each category of OSS applications, we'll be putting together the nuts and bolts to build an OSS environment for an efficient, profitable, modern CSP.

Network engineering

A communications network contains devices. It is the configuration of these devices that enables the network to carry traffic and services efficiently from one place to another.

Network Engineering must address two challenges:

- These devices are highly complex and often use proprietary interfaces and configuration syntax
- These devices could be anywhere: In your office basement, in a roadside cabinet, on a customer's premises, or virtualized in a datacenter

It is therefore the role of Network Engineering applications to enable connections to devices to facilitate the tasks of monitoring and configuration while, as much as possible, simplifying these highly technical tasks.

Touching the network

Network Engineering comprises systems and processes that directly affect, or touch, network devices.

As these applications manage device configuration and state, many Network Engineering applications are also concerned with the services that use the network devices.

We use the term 'engineering' for this category of OSS applications because they are the most technical, 'hands-on', tools for operating a network. As such they are typically used by teams within the CSP that carry titles such as Network Engineers and Network Operations.

Network Element Manager

The NEM provides remote connectivity to network devices (or ‘elements’) offering essential configuration and monitoring tools.

While for many devices it may be possible to use a command line to connect and perform these tasks (for example using TELNET) the NEM makes the process quicker and more consistent by maintaining lists of available devices, automatically checking devices’ state, maintaining the user’s security credential, formatting the output in a graphical user interface, and so on.

An NEM is concerned with the individual device and its current state. Managing the coordinated configuration of multiple devices, for example to route a service across a network domain, is usually done by ‘higher-level’ OSS systems such as a Network Management System, Orchestration or Fulfilment application.

An NEM may also be referred to as an Element Management Systems (EMS).

Configuration Management

A Configuration Management (CM) application is designed to remove the risk associated with modifying a device’s configuration.

In the simplest case, the CM maintains a record of configuration changes and historic configuration files to enable roll-back in the event of a change causing an unexpected problem.

A CM may also monitor devices to log changes that occurred via any user accessing the device using an NEM or command-line interface. A sophisticated CM can identify configuration changes that breach a policy, such as a valid parameter range, and report this issue to the user.

Security Management

Security Management provides an application dedicated to monitoring and configuring the security credentials of network devices and services.

The Security Manager supports various intrusion detection techniques, either by identifying inappropriate access patterns, questionable configuration changes, or by examining data packets for suspicious data.

Security Management will monitor traffic and devices, and also interface to other common security appliances in the network, which provide hardware-based security features such as firewalls, VPN, and SSL.

Modern Security Managers are increasingly not just concerned with the network but also the specific applications being accessed. For example, a customer may be accessing an application in a data center via the network. The Security Manager is able to monitor the security end-to-end to give a complete picture of the security situation.

Network Management System

While an NEM is concerned with individual devices a Network Management Systems (NMS) is concerned with a whole network domain, usually comprising a single vendor's equipment.

An NMS will typically provide access to all the functionality of an NEM, either directly connecting to devices itself or by interfacing to the NEM.

Where the NMS improves on the NEM is in providing a more complete view of how devices work together to carry traffic and services across the network.

The NMS can display configured service routes and simplifies the process of configuring new routes across multiple devices in the domain.

Many NMS incorporate Configuration Management and Security Management capabilities in a single application.

Increasingly, traditional NMS are being superseded by SDN Controllers which have greater centralized software control of the devices.

Workforce Management System

At some point someone has to go out in to the world and install equipment and physically connect customers.

Workforce Management Systems (WFMS) are responsible for assignment of work to engineers to maximize efficiency and ensure timely delivery of services to customers.

Sending an engineer to a site, sometimes called a ‘truck roll’, is expensive and time consuming. The most significant benefit of WFMS is the ability to intelligently organize work so as to minimize this cost.

A WFMS can only optimize work when other applications optimize their output to the same goal. Often Service Fulfilment and Planning applications will comprise some high-level ‘technical’ WFMS function. They will produce ‘work-orders’ that detail the sequence of work, the resources that are required and the sites that may need a visit. This leaves the work of managing staff time and motion to generic scheduling functions that are often found in a large company’s Enterprise Resource Planning systems.

Fault Management System

A Fault Management System (FMS), also called Event Management, is responsible for collecting ‘alarms’ from network devices that are indicative of faults and outages.

An FMS collects alarms from many types of device, usually supporting multiple vendors, to provide a single view of the state of the network.

The FMS may provide some initial analysis of faults, with basic prioritization and correlation, but will pass the raw details to a Service Assurance application or some other Customer Experience application which can cross reference alarms with other data sources to perform more intelligent analysis of the cause and effect of a network fault.

Data & analytics

While you can learn a lot by examining devices and services directly via Network Engineering applications, there are many management processes that require a richer set of data.

Gathering, storing and making available data about the network, service and customers is essential infrastructure at the heart of many OSS processes.

This data is stored in a specially designed database known as an inventory.

In many cases, the primary benefit of storing data is to remove the need to go elsewhere to find the answer to a question. If a CSP wants to know if ten new customers can be connected to an access point in the network, it would be expensive and slow to send someone to the site to check for available ports.

Even with an NMS in hand, it can be slow to click through a very technical interface, perhaps one NMS for each vendor or technology, to determine if capacity exists for new services.

So, data or 'network models' are pulled in to dedicated OSS Inventory applications to make Planning, Service Fulfilment, Fault Management and Analytics simpler and faster.

Once that data is there, supporting core processes, there's a huge opportunity to mine it for insight in to customer behavior, network trends, and performance characteristics.

While analysis of data has existed in OSS applications for some time it is usually limited to addressing application-specific questions. The emergence of Big Data, Machine Learning and open source tools means more questions can be asked of the data and the option to ask ad-hoc or one-off questions becomes viable.

Inventory

An Inventory systems' core component is a database of network equipment, their configuration and records of services supported by the network. The Inventory provides an abstracted view of the real world, containing an as-built or as-designed 'model' of the network augmented with customer-facing data.

Unlike an NEM or NMS, the Inventory will minimize vendor-specific syntax to provide a consistent view of the network.

The Inventory also provides an off-line view of the network, historic and planned network configuration, unlike most network engineering tools which only provide the current network state.

Faulty or temporarily out-of-service equipment remains visible in the Inventory even if its fault renders it temporarily 'invisible' to the NEM/NMS, allowing design and planning to continue unaffected.

Changes committed to Inventory but not yet activated or deployed in the network are recorded, allowing user to base future designs on these imminent changes.

Resources such as equipment or circuits can be set to a reserved state, before any live services make use of them, reducing resource conflicts in Service Fulfilment and Activation processes.

Network resources can also be associated with customer services in the Inventory's records, to aide tasks such as fault analysis and service reconfiguration.

Inventory's 'as-designed' view of the network is essential to Service Fulfilment, Fault Management and Planning processes. Inventory is therefore tightly integrated with, or even embedded in, these other OSS applications.

GIS

GIS, or Geographic Information System, is similar in principal to an Inventory. Where Inventory focuses on a model of the network and services, GIS is more concerned with the world the network occupies.

GIS is the data source at the core of Spatial Planning and is increasingly used by Analytics.

In the simplest case, a GIS provides a mapping capability to record the precise location of devices and customer premises. In addition to mapping, the GIS will hold network-specific physical properties such as the route and capacity of a duct holding cables, or the position of devices within each of the CSP's sites.

A GIS network model will include a record of other physical dependencies like available power and air-conditioning capacity at a site, of the presence of buildings and foliage in the line-of-site between two microwave receivers.

The adoption of GIS-based OSS applications was primarily driven by network engineers responsible for installing and maintaining equipment. But again, as is the case with Inventory data, CSPs now recognize the value of the data in GIS for more commercial purposes. Augmented with demographic data, and able to model the 'reachability' of customers, GIS is increasingly the basis of marketing analytics to determine which customers can be profitably connected to the network.

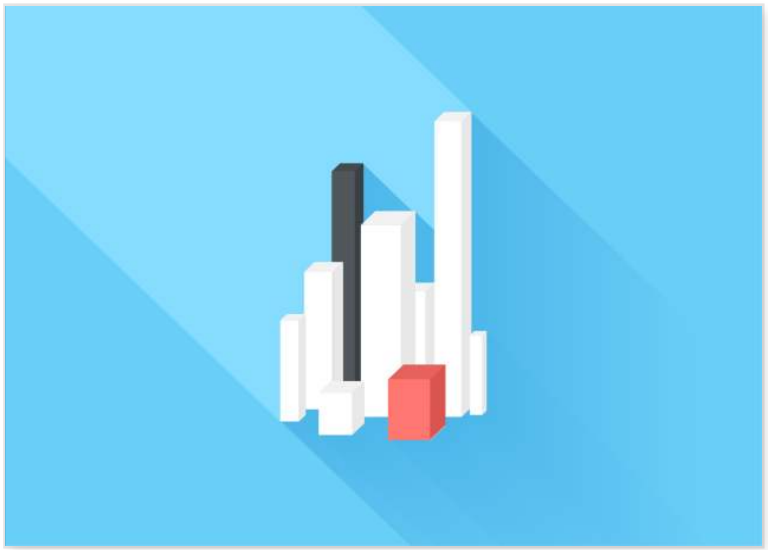

Discovery

Discovery applications are responsible for pulling in data from the network for either immediate access or, more commonly, storage in Inventory, GIS or other databases for subsequent analysis. The discovered network view is modeled as a topology, and it continuously updated as information is gathered from the network.

While some applications have discovery capabilities built-in, particularly if they address just a single technology or an industry-standard network interface is available, in most CSPs there are several different network technologies from two or three different vendors.

The challenge Discovery meets is to interface to multiple vendor's devices, NEM or NMS systems, via different interfaces, extracting data in varying formats. More recently, this discovery function is often integrated with orchestration which oversees the network topology across multiple domains and multiple vendors.

Data goes through an extract, transform and load (ETL) process, eventually being made available to the target OSS application via a single interface and single data format.

Federation

OSS processes are only as good as the data they have available to them. Data represents an OSS view of the world, describing networks, services, customers, future trends and past events. Data is available directly from the network (via a Discovery application) or in databases belonging to Inventories and Monitoring applications. There is certainly no shortage of data to be extracted from communications networks, OSS and BSS systems.

But because data comes from many sources – different OSS systems, network domains and network technologies – it's not immediately easy to work with. The data content and format can vary considerably. The data may be physically in different places. Access to the data will be through a range of different APIs and database types.

The role of Federation is to bring this data together either by duplicating dispersed data to a single convenient database or by providing a consistent API to reach through to multiple data sources. By using a combination of these techniques, Federation can make large amounts of data useable by OSS, often without the need to go through the process of replacing these established data sources or introducing additional Inventory solutions.

Analytics & Business Intelligence

Data has been at the heart of OSS applications for years, but the analytics capabilities have tended to be rigid – fixed to support the primary propose of each application.

The ever-reducing cost of IT hardware has made traditional data analytics platforms like data warehouses much more affordable. Furthermore, the emergence of Big Data platforms like Hadoop has further reduced the cost in terms of money and time.

These generic (at least, not OSS-specific) analytics and business intelligence platforms are now being deployed within OSS for ad-hoc reporting purposes and data mining to gain better understanding of the CSPs business.

Inventory, GIS, Discovery, Network Engineering data sources, merged with customer data from BSS applications offer tremendous opportunity to learn about how the network is working and what keeps customers happily paying their bills.

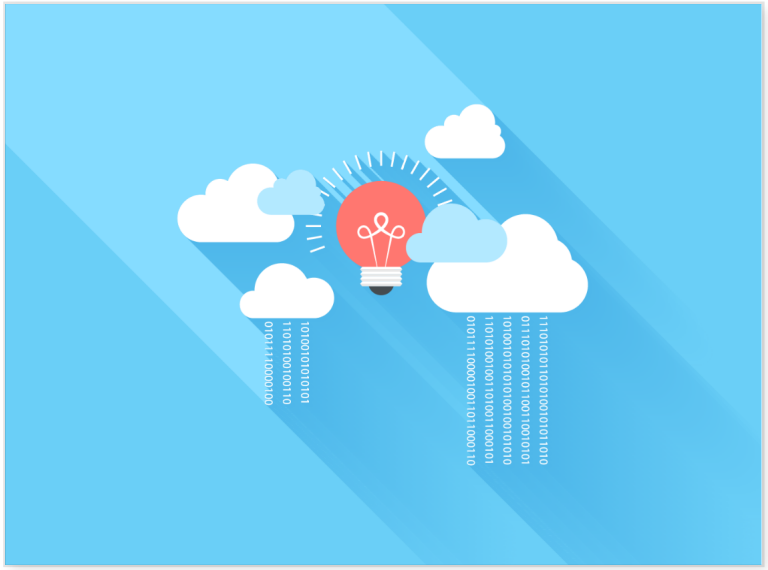

Planning

A CSP must be able to engage in strategic planning for growth in customer traffic, and tactical planning to design bespoke services or resolve performance issues.

When new technology becomes available, such as LTE, the CSP must also be able to plan its roll-out and the associated upgrade of the existing network.

Planning, more than any other operational task, considers not just the network capacity and capabilities, but also the cost, revenue potential, and reliability of the network.

Planning is ultimately responsible for ensuring that all other operational tasks are following rules and policies that ensure the CSP is building and running an efficient, profitable network that meets the business' strategic objectives.

In the past, networks were relatively 'static'. Device functions were fixed; traffic routes were inflexible; resources were allocated to specific services for the long term. As such, planning a network was a relatively slow, infrequent task carried out maybe every few months to prepare budgets and schedule major works.

Increasingly, as network become more dynamic and technologies that enable software-driven networks like MPLS-TE, SDN and NFV are adopted, planning is being carried out daily, hourly or in real-time: The results being used to reconfigure virtual resources to prepare for shifts in traffic throughout the day or to quickly react to network performance issues.

Capacity Management

Capacity Management (CapM) applications are used to analyze how traffic across the entire network results in high or low utilization of specific resources, and what action can be taken to more efficiently manage this traffic.

With the move to all-IP networks there is no longer a simple way to model how much bandwidth a certain number of customers will use, or even what route that traffic would take in the event of one or more network faults.

Without a simple statistical model, CapM applications have emerged that simulate how traffic flows across the network, analyzing this demand for bandwidth while also considering the effect on cost and service quality of a network upgrade or reconfiguration.

Currently, as with most planning application, CapM is carried out periodically to support strategic investment in the network.

Increasingly, used in real-time, CapM outputs configuration details directly to Activation, Network Element Managers or Network Management Systems.

Spatial Planning

Often called inside-plant and outside-plant applications, Spatial Planning provides a means of finding space to install new devices and run new cables.

Sending an engineer to a site, digging holes in the ground, and installing devices are all time-consuming, expensive tasks. Spatial Planning allows much of the initial planning to take place without a site visit and can identify the most suitable location to deploy new resources.

Inside-plant applications provide a schematic of a building, tower, or cabinet, identifying the exact location of devices and the routing of cables through the site. An additional important part of inside-plant planning is identifying the availability of enough power and air-conditioning; both being limited and expensive resources.

Outside-plant Spatial Planning is, as the name suggests, about physical resources located in the street, fields and on other peoples' buildings. It is supported by map-based schematics diagrams or fully-fledged Geographical Information Systems (GIS) capabilities. Again, it is concerned with identifying exactly where a resource is and finding space for new devices and cabling. Outside-

plant Spatial Planning is also used for identifying potential new sites and some initial civil engineering planning tasks.

Spatial Planning applications are not generic mapping or drawing tools; They include built-in rules and models for how communications resources are affected by the environment and physical topology of the network.

GIS-based Spatial Planning application used by mobile network operators will often include some RF optimization capabilities to enable the design of point-to-point microwave links (a wireless alternative to using cables to connect cellular sites to the network). The performance and configuration of microwave links is heavily influence by physical factors such as line-of-sight, distance, angle, and so on meaning that a GIS is the natural place to carryout related planning and design activities.

Similarly, fixed line physical planning, such as placing new fiber optics cable, will include a model for signal degradation caused by distance or the number of connected customers.

Forecasting

By analyzing traffic volumes, sales figures, and usage patterns, Forecasting applications will predict the future demands on the network, providing a firm basis for planning activities.

Like other Planning functions, Forecasting may be used strategically or tactically.

Identifying network-wide trends for new customers and growing data usage patterns can be used to ensure there is always enough capacity in the right part of the network to meet that demand.

In the event of a network fault or performance issue, Forecasting can be used to identify when, or if, there will be a knock-on effect on other parts of the network due to the rerouted traffic

Forecasting can generate useful insight on its own, but by providing input, such as demand forecasts, to a Capacity Management application the solution can produce traffic trends, simulate the impact on the network, and identify the optimal plan for resolving any predicted issues.

Service Optimization

Service Optimization is concerned with the proper allocation of resources and routes across the network, to make best use of available resources while also meeting criteria to ensure a good quality of service.

It is therefore a similar function to Capacity Management, however while that function is more concerned with the design of network resources as a whole, Service Optimization is concerned with how those resources are used by each individual service.

Service Optimization may be part of a Service Fulfilment process. In this case optimization is just one, important, step in the larger fulfilment process.

Optimization may also take place at any time during the service's lifetime, perhaps to re-configure a service as the environment around it changes resulting in the current design being sub-optimal.

Service Optimization is often a tool of the CSP's network planners as a stand-alone tool, but as network resource become virtualized and software-driven, this function can be automated and executed in near real-time as part of an SDN control stack.

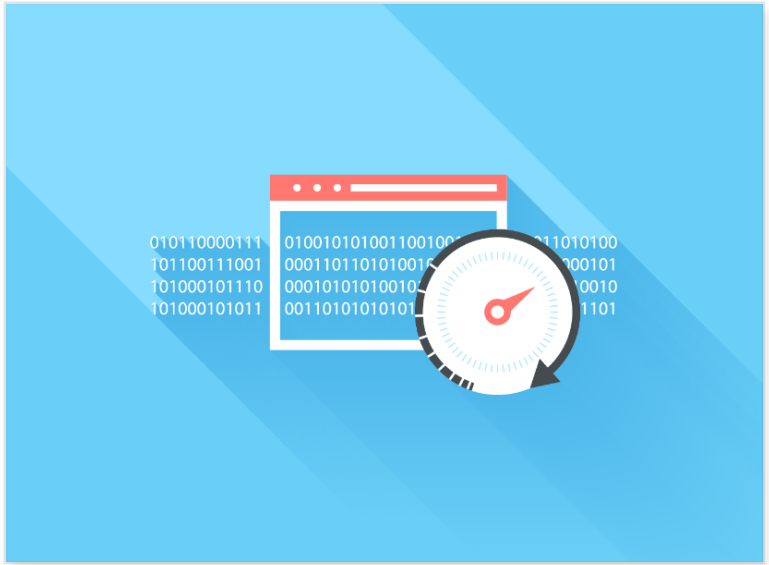

Service fulfilment

A CSP does not make any money from a network until they can deliver services to customers and start billing them.

Service Fulfilment is therefore perhaps the most fundamental of all operational processes, taking orders from the BSS layer and ultimately ensuring they are successfully deployed and activated in the network.

Timely delivery of an active, billable service, deployed adhering to technical and business rules, while minimizing cost and engineering effort, is the goal of Service Fulfilment.

Achieving this goal can be challenging.

It requires multiple OSS applications to be coordinated, the reservation and allocation of resources across multiple network technologies, and extensive process automation but with the ability to involve multiple teams of people in choices or exception resolution.

Or, to put it another way, Service Fulfilment must embody the knowledge and expertise of dozens of network planners, designers and product managers, completing a task in seconds that would take days or weeks to complete manually.

Though still a critical part of any OSS environment, it is worth noting that Service Fulfilment is no longer the ultimate solution that it positioned itself as in the 1990s/2000s. As networks, services and customer have changed, so has the need to consider proper design of network resources before fulfilment, and management of customer experience throughout the service's lifetime.

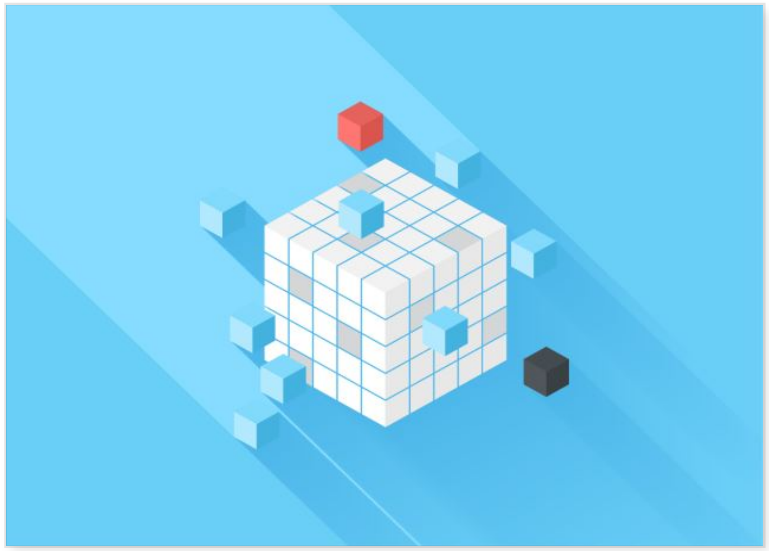

Service Management

In the recent past, CSPs offered only three or four types of service to residential customers, and a similar number to businesses, supported by a small number of network technologies, accessed with devices supplied by the CSP.

Today, CSPs introduce more service types every year, with an array of options for customization, which are offered to customers across different networks and accessed by many and varied devices.

To fulfil or modify a service requires an understanding of the customer's current services and infrastructure. It must be possibly to quickly determine what compatible services and upgrades can be delivered.

Service Management applications provide a source of this high-level customer and service data to determine an overall service fulfilment process for a customer and to ensure BSS tasks such as order entry and order management are offering suitable types of service.

Service Management applications do not typically hold a lot of detailed information about customers, services or the network - Just enough to support initial decision making and to determine a fulfilment process.

Service Catalog

Service Catalogs provide a model of how network resources are combined to deliver a type of service offering to customers. This serves two purposes:

Firstly, the Service Catalog acts as a design tool, allowing the CSP's product managers to combine network capabilities to create new types of service.

Secondly, the model can be used as a high-level guide for Service Fulfilment or Orchestration, choosing the right network resources for an individual service order, considering the customers location, existing services and service requirements.

An important role for the Service Catalog is to abstract the various network technologies in to 'capabilities' that can be combined to create service offerings. A consistent range of services can then be offered to the market, based on the capabilities of the underlying OSS to fulfil the service and the network to support the service.

Service Catalogs may also be referred to as Product Catalogs.

Service Orchestration

A modern service comprises components from many technologies and network domains. It's increasingly common for multiple OSS systems, network virtualization platforms and data centers to be coordinated to create an end-to-end service.

Service Orchestration is responsible for the overall delivery of a service offering. This involves integrating with the Service Provisioning functions of OSS systems and control systems responsible for the underlying network resources that will be required. Unlike the Service Provisioning, Service Orchestration does not know *how* to configure network resources. Instead it ensures each system is delivering the right component, to the right specification, at the right time.

Service Orchestration benefits the CSP by allowing 'recipes' for new service offerings to be easily created by piecing together available network services and resources, as offered by Service Provisioning and SDN Controllers. Service Orchestration also provides a platform for automation, reducing or removing the need for a network engineer to manually manage coordinate the overall fulfilment process.

Service Provisioning

Where a network domain does not have a controller function (such as an SDN Controller), Service Provisioning is responsible for the design of a service in terms of identifying and allocating network resources, to create a complete specification that can be activated on the network.

Service Provisioning delivers process automation to ensure business and technical rules are applied to a design, minimizing the need for a planner or engineer to be involved.

Service Provisioning is tightly integrated with OSS Inventory to enable access to the necessary information to select available resources and reserve capacity.

The result of Service Provisioning is a detailed, technical design for a service which can be supplied to engineers, NMS, or Activation applications.

Activation

Activation applications are responsible for modification of the network to turn on or modify a service.

Activation of a single service may involve orchestration of configuration changes via multiple NMS, NEM, and Configuration Management systems. Activation must ensure the correct sequencing of this work to ‘stitch’ together the service. In the event of a problem, Activation must be able to gracefully resolve the issue or rollback all the changes.

Activation applications differ from NMS as they are multi-vendor and often multi-technology. Also, Activation applications tend to be automated and driven by APIs, while an NMS is primarily driven by a user.

In complex OSS and network environments, Activation isolates higher level Service Fulfilment applications from the network complexity while offering a consistent API for them to interface to for service configuration purposes.

Activation applications are usually applied to legacy network technologies that support a low-level API or command-line for configuration but lack a controller layer to abstract this for access by Service Orchestration.

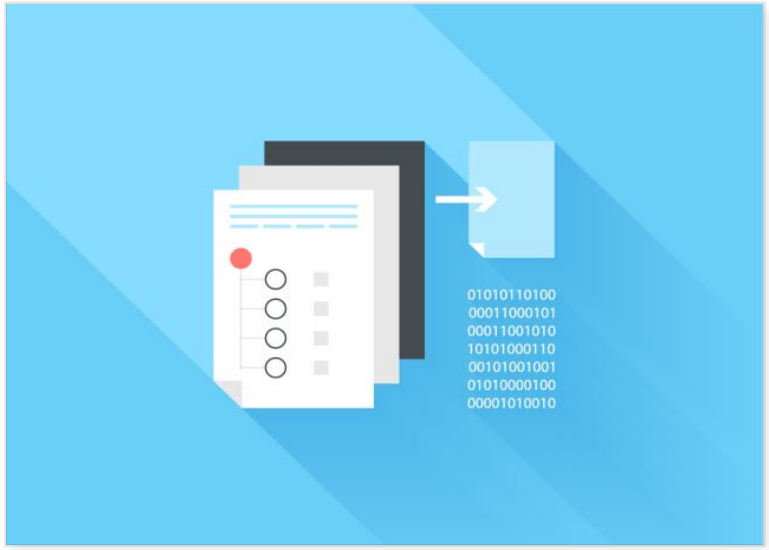

Customer experience

Everything a CSP does should avoid a negative customer experience.

Doing OSS/BSS right and in a timely manner is a good start.

In modern CSPs a far greater focus is put on proactively improving the customer experience, with new applications emerging that improve the quality of the networks and services for all customers or focus on specific regions or even individuals to address their issues and needs.

In some markets, customer turn-over can be as high as 30-50% annually. This is a big problem for CSPs when the cost to acquire a new customer is very high and maximizing lifetime value from a customer is critical to achieving profitability.

Customer experience management (CEX or CEM) emerged out of traditional OSS fields like Service Assurance and Performance Management. These were largely engineering, network-focused applications originally. The modern twist is the refocusing on the customer rather than on the network.

While an optimized, functioning network is a good thing, one must be mindful that a technically 'good' design is not necessarily the same as a 'good' solution for customers.

If nothing else, CEX focuses effort and budget on network planning and engineering tasks that have the greatest benefit to customers.

At its best, CEX drives company strategy, defines operational policies, and identifies issues and opportunities to make customers loyal, well before a network alarm is raised.

CEX is a broad discipline that by its nature must traverse the CSP's many business functions as well as spanning BSS and OSS processes.

In this Guide, we will discuss the network operations aspects of CEX. Be aware that this is just the tip of the iceberg.

Service Assurance

Service Assurance (SA) applications provide an intelligent response to network and service issues.

Taking input from Fault Management, Performance Management and direct feedback from customers via BSS, it is the responsibility of Service Assurance to determine an issue's impact on customer services, the root cause of the problem, and provide data to support resolution.

Service Assurance applications leverage Planning applications for their ability to redesign the network to resolve problems.

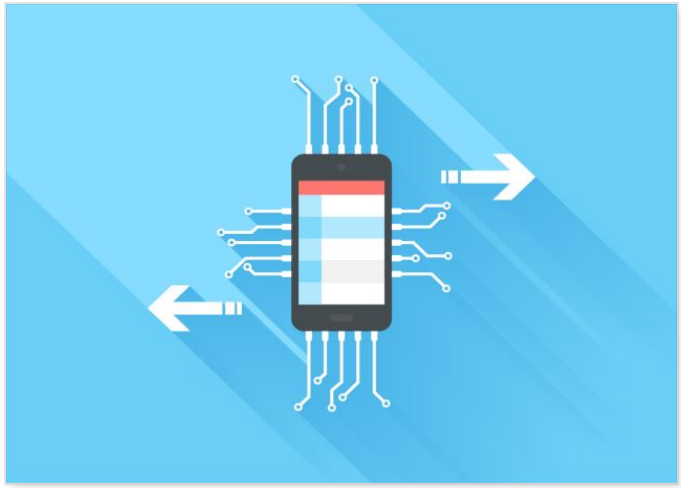

Performance Management

Performance Management (PM) applications monitor the flow of traffic across a network, and the ability of devices to support the traffic while meeting service quality metrics. They gather key performance indicators (KPIs), such as the length of time it takes data to get from one point in the network to another point.

Modern IP, MPLS, Ethernet, optical and SDN networks will route traffic according to a variety of policies and the current network state. Performance Management monitors the situation over time, allowing the CSP to make informed decisions to optimize those routing policies and support strategic network planning activities.

Performance Management applications will acquire data about the network either directly from devices themselves, from specialized devices – known as probes - that collect KPIs.

Application Performance Management

While Performance Management applications traditionally focus on network devices, Application Performance Management (APM) is concerned by service applications, and their infrastructure's ability to meet an expected service quality.

APM has emerged as a separate application as it monitors IT resources such as server resources, server processes, and transaction performance, rather than 'traditional' telecommunications devices.

For example, a database supporting a massaging service may be monitored to check system up-time, storage capacity, CPU utilization, and the performance of queries on the data.

Modern APM can combine monitoring of a CSP's IT and telecoms devices to provide a full end-to-end view of service performance.

Where IT resources are being used as a platform for Network Functions Virtualization (NFV), the role of APM is consumed by NFV's MANO (Management and Orchestration) function.

MANO monitors the Virtual Network Function's (VNF) performance needs, to ensure that server resources (compute, memory, ports) are sufficiently allocated for the VNF to do its job.

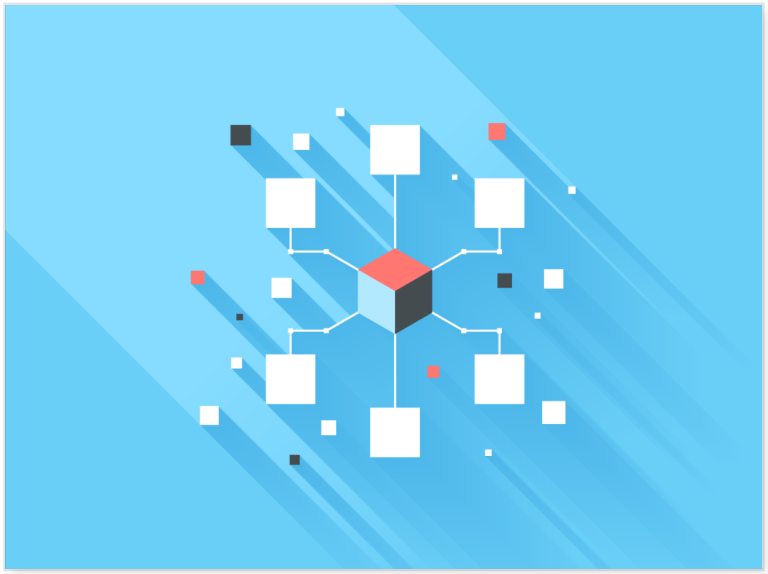

Experience is everything

Looking at individual OSS applications it may look like Customer Experience is just a layer of gloss on well-established functions.

But CEX is about the big picture.

Not only does each OSS application need to be aware of customers and their expectations, they also need to work together as part of a CSP's operational processes, crossing all OSS functions.

'Right first time' is a good target for a CSP when fulfilling services. In practice, between service order, service design and fulfilment, problems can occur. The design and delivery of services to customers should not be a 'black hole': Order goes in, a service comes out…

Eventually.

It must be predictable, transparent, and responsive. If there's a problem, most customers can tolerate it, if the communication is good and the resolution is prompt.

Changing engineering

Engineering must support risk-free changes to device configuration and services with less customer-affecting downtime. This goes beyond the traditional approach of giving customers notice of engineering works and trying to minimize downtime.

Instead, an understanding of the impact of changes, prioritization, and intelligent scheduling of work around the least disruptive traffic patterns is required.

Experience is in the data

A CSP's data repositories must be open to allow their data to be leveraged for CEX analysis by BI and Big Data applications.

This requires more than sharing APIs, XML data or permitting SQL queries. Databases at the heart of Inventory and Service Fulfilment applications must be able to support high-volume, real-time analysis. This means adopting data structures and interfaces that are often quite different from those designed for their primary role of processing transactions.

What next?

OSS vendors and CSPs have made great progress in becoming more customer-centric, but clearly there is more to do to advance both the technology and the industry's culture.

Being customer-first, rather than network-first, is *nearly* a reality.

And while CEX is re-focusing the industry, there are plenty more trends emerging today that will further change the way CSPs deliver great experiences.

In the next chapter we will look at these trends and their impact on networks, on CSP business models, on operations and on each of us as consumers of communications services.

Chapter Four

Network trends and the future of OSS

Traffic and topology – The trend driving all trends

A focus on good customer experience is great, but that experience must be delivered by the network in an efficient way to ensure the service is profitable. At the same time, more than ever, customers' expectations are changing, the network is changing, and services are changing.

From a network planning and operations perspective, there are two major, related trends driving change at a network-wide level: Customers (and devices) are using the network in a fundamentally different way, and the network architecture is evolving to support those customers' changing habits and needs.

In this chapter we consider these trends, and how they converge, resulting in a step-change in the flexibility of network operations and the economics of service delivery.

More everything

You are probably familiar with the statistics and forecasts that show massive increases in demand for data and bandwidth, from both mobile and fixed-line customers.

For instance.

Cisco's Global Fixed and Mobile Data Forecast in 2017 suggested that Internet data carried globally will grow from 96 petabytes per month today to over 350 petabytes per month by 2022.

That demand from customers drives change: CSPs are having to deliver an increase in network capacity and must handle that traffic in more intelligent ways.

Growth in capacity is a familiar challenge: In the previous ten years the spread of DSL and cable broadband, and the rollout of mobile 3G services, all resulted in massive investment in networks. But today the economics are different. CSPs do not have the almost unlimited funding they had 10-15 years ago, but there is still an expectation that the latest generations of fixed-line and mobile technology - *particularly* mobile technology - will result in a major improvement in bandwidth and service quality.

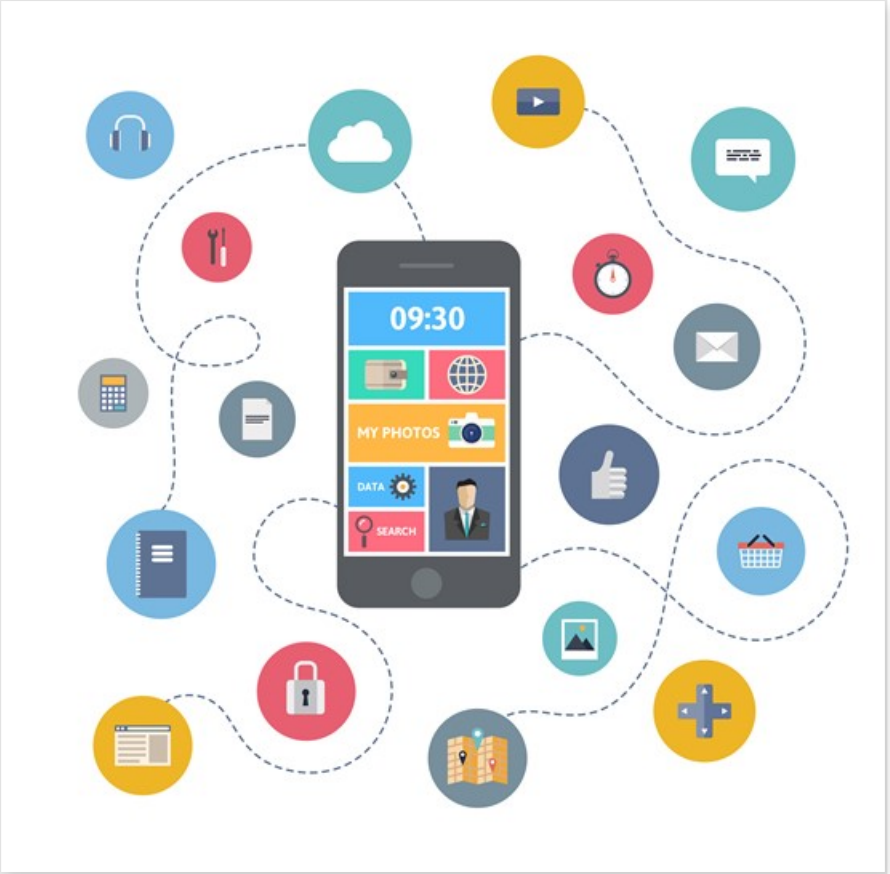

Mobile trends

Whether 3G, 4G or (coming soon!) 5G – rollout of mobile networks is one of the biggest areas of expense for CSPs today. Long after most countries issued the rights to the radio spectrum used for 4G/LTE, delivering that service to the majority of customers continues to be a challenge.

The primary technical challenge is delivering good performance while achieving both *coverage* and *capacity*.

Coverage

Coverage is about providing a signal, or a connection to the network. With 4G, some radio frequencies are great at covering a wide area but are not good at penetrating inside buildings. Other frequencies are best for urban areas because they allow good signal indoors, but the range of the signal is much smaller.

But coverage alone is useless, unless the CSP also delivers enough capacity.

Capacity

Capacity to hold voice calls and broadband data capacity, for all the people (and machines) served by the coverage. Capacity itself is dependent on two things – the capacity of the radio signal and the capacity of the backhaul connection from the antennae to the CSPs core network, then on to the public Internet.

Let's consider the fundamentals of mobile network design, then look at the options available to CSPs to address these challenges.

The mobile network

A mobile cellular network is made up of ‘cell sites’ containing towers, masts or other infrastructure that contain a transceiver. In a cellular network each cell site’s transceiver uses a radio frequency that differs from other nearby cells, to avoid interference. The result is like a patchwork quilt of cells that provide mobile coverage.

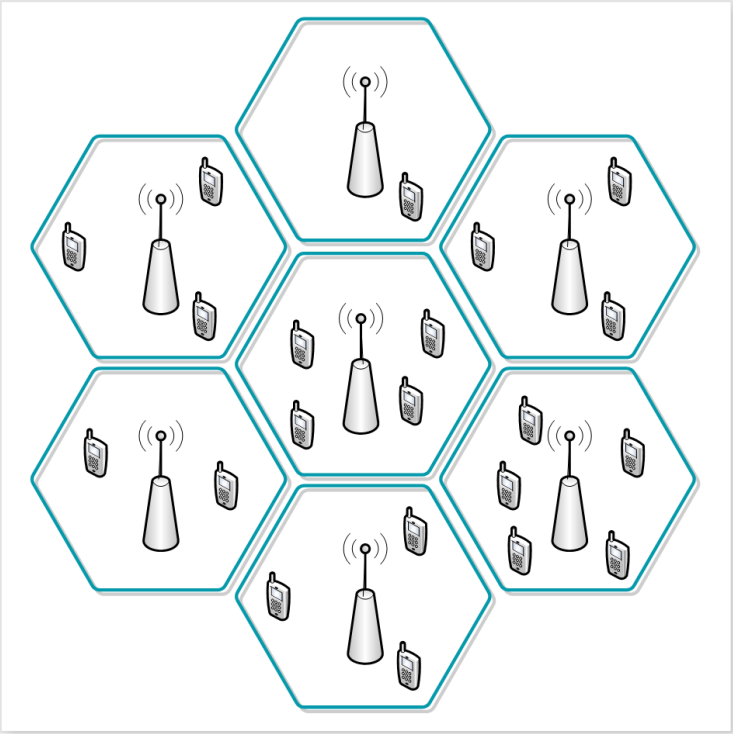

Until recently most mobile networks were made up of cells which covered a range of about half a mile to five miles from the tower. These big cells that are served by big towers, and big antenna, are called *macro cells*.

The problem in mobile networks is that you can’t put up enough of those big towers. And even if you could, wide ranging, overlapping signals are inefficient at solving the challenge of capacity and coverage – both in terms of radio frequency efficiency and economically.

Macro cells are what made up the vast majority of coverage for 2G and 3G network.

But because 4G need much more capacity to deliver the step-up in bandwidth, and bigger coverage to make it an attractive, usable service, the answer is to go…. Small.

Small cells

As the name suggests small cells have a much smaller range than macro cells, delivered using smaller transceivers that can be easily located on walls, up lamp posts or inside equipment like a broadband router.

Small cell technology exists to cover various ranges and are deployed either within the CSP's network or on the customer's premises. You may therefore see terms like micro cell, pico cell and femto cell used to cover a variety of use cases – But collectively they can be described as small cells.

More coverage

Small cells improve coverage by being able to fill in gaps in the network. If, as a CSP, your 4G signal isn't penetrating certain buildings, you can supply a small cell to be installed inside. Or, if macro cell roll-out isn't economically feasible in a region but you have a dense population of users, such as a train station or business park, small cells can be deployed without the logistical hassle of finding a site for a big macro cell tower.

More capacity

Small cells also deliver greater capacity. A single 4G macro cell, using a single frequency, covering a square mile would offer a share of 40Mbps of data capacity between all customers in that coverage. The same area served by four small cells, using four different frequencies, would offer 160Mbps in total radio capacity.

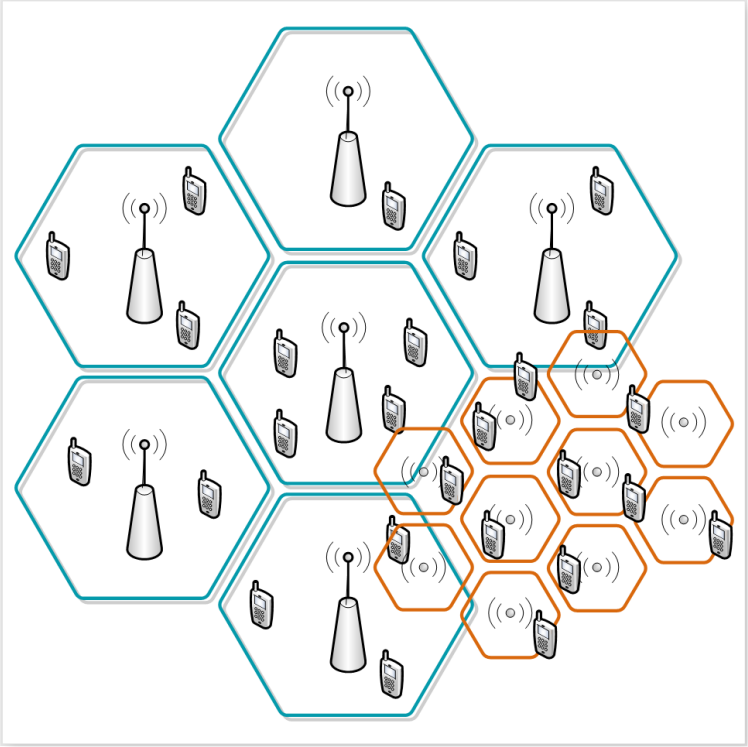

Changing OSS for new mobile networks

Small cells are a powerful tool in the CSP's armory for delivering coverage and capacity. But like any network technology, determining when, where and how to use it most effectively and economically is not a simple task.

Small cells are therefore putting pressure on existing OSS systems that were designed around the macro cell model: Occasional, predictable rollout of large cell sites with a predictable traffic load.

Now, the flexibility of small cells and their fast rollout means OSS must keep up.

The core OSS capabilities are well established. Radio frequency planning and simulation, traffic analytics, physical GIS planning for cell site connectivity, and network planning/fulfilment for backhaul.

But if, as estimated, a CSP is rolling out thirty small cells for every macro cell site, it is more essential that these OSS applications are joined-up and able to repeat and iterate to an optimal network design. One cannot simply go through a single long planning cycle and assume 'it's done' once a small cell is up and running.

Now, OSS must be able to perform on-going analysis and planning reviews of not just radio coverage but also the ability to deliver sufficient backhaul and core capacity to those cells.

This change in the way mobile networks must be planned and operated has been reflected in OSS vendors' product offerings and company acquisitions. Big OSS vendors have augmented their capabilities to cover the mobile network end-to-end including radio, cell sites, backhaul and core transmission network operations.

But while demand for network capacity is increasing, and increasing faster than ever, it is not what is causing the biggest changes to the way CSPs build their network.

Machine Learning

The next data revolution

Artificial Intelligence (AI) is the broad field of science and engineering attempting to perform tasks with (apparent) intelligence rather than relying on the pre-programmed rules. Within the field of AI, Machine Learning (ML) has emerged in recent years as the most practical, useable application of the science by a wide range of industries.

Machine Learning is a type of AI that uses techniques like neural networks to analyze data with the intention of categorizing or rating it, usually against term that are hard to clearly encode in traditional methods, but usually easy for a human to complete. For example, software using traditional coding can determine how much red is in a picture by statistical analysis of the data encoding the picture. However, answering the question 'does the person in this picture look happy' is much harder. There are no bits and bytes in the picture indicating the quantity of happiness. Nor whether the photo is even a picture of a person or not.

Machine Learning answers these sorts of queries first by being taught. It is dependent on a large dataset and a feedback loop the tunes the neural nets based on whether it is accurately analyzing the training dataset. Once trained on a sufficiently varied dataset, it can start answering questions based on data it's never seen before.

And while answering a question like 'is this person happy' is trivial for a human, Machine Learning has the advantage of being able to it millions of times per second.

Deep Learning

Like most AI today, Machine Learning usually solves a narrow, well defined problem: Answering one or two questions for a consistent data set.

Deep Learning is to Machine Learning what Big Data is to, well, normal data. Deep Learning is trained on massive data sets with broader learning criteria. With the intention of being able to answer more types of question or even questions you haven't thought of yet.

Now you can ask whether a picture shows not just 'a happy person' but 'a happy child at a birthday party'.

Deep Learning can also infer relationships and outcomes. For example, the same application that analyzes photos could be asked 'what makes people happy'. And the answer might come back as 'birthday parties and petting puppies', based on what the Deep Learning application learnt by analyzing the photos.

While the subject matter for inference necessarily remains constrained - images of things; sensors on moving vehicles; the human voice - the application for Deep Learning is much broader, to the degree that the Deep Learning system can identify the characteristics of the data with accuracy like that of a human, but with the ability analyze millions of data points to form a conclusion.

So, while Machine Learning could be described like having a million people in a room doing a repetitive task, Deep Learning is compared to having one person analyzing a million data points to bring new insight in to the dataset.

Machine Learning today

Machine Learning has already had significant success in BSS, where it can be found targeting products and services at customers it determines are most likely to make a purchase. Machine Learning is also increasingly used by a customer support and marketing teams to analyze service use, service-affecting issues, and even the customers sentiment when they phone the help desk or post on social media.

While BSS has been where the source of low hanging fruit for AI in communications, Machine Learning was slower to achieve success in OSS. Why?

Data found in BSS systems, like call records and purchase history, is well-structured and easily interpreted by Machine Learning. OSS data is often highly complex, interdependent, and of varying formats.

Understanding a customer's sentiment and their behavior is a complex problem for Machine Learning, but one that is relatively common across many industries. OSS data is unique to OSS, and the modelling of communication networks is unique to this industry – There are far fewer similar successful systems to build on, compared to BSS.

In other specialized industries like aviation, manufacturing, and vehicle automation, the approach has not been to create a broad, autonomous solution in one step. Instead, Machine Learning has been applied to improve small but critical functions. In this way, benefits are realized quickly, the implications of autonomy are easily managed, and a small success can provide insight and data the next Machine Learning solution will benefit from.

A similar approach has been successful in OSS, leading to several Machine Learning applications:

Data traffic forecasting based on historical data volumes and real-world events, such as holidays, sporting events, TV events, and major security incidents.

Responding to **network capacity** demands (either current or forecast demand). Machine Learning can assist in rerouting traffic to maintain service quality or it can identify the most cost-effective way to meet changes in capacity demand.

Spotting patterns in service use, attempted network intrusions and transactions to **identify security or fraud** risks.

Imagine…

Across BSS and OSS functions, a combination of traditional solutions, AI-enabled functions and orchestration can deliver a level of customer experience seldom seen in this industry.

Imagine… A fault occurs in the CSP's network. Two business processes are immediately started.

One performs a traditional root-cause analysis to determine the reason for the fault. This is fed in to a Machine Learning function that can both identify how to predict a similar failure in future and proactively advise the CSP about how to mitigate the impact on customers.

In the second process, a Machine Learning function predicts the effect on each individual customer, based on their service use and positive/negative views of the service to-date. A tailored customer-satisfaction action plan is created, meaning the next time the customer uses any contact points with the CSP (online portal, phone, social media) an appropriate bonus service or discount is offered, carefully chosen to compensate the customer and minimize any negative feedback on social media.

Machine Learning can push OSS further than ever to meet its commitments to increase network efficiency and reduce operational costs.

However, as we'll see later in this chapter, the trend towards software defined networks and virtualization might make Machine Learning *mandatory*, not just a nice-to-have. The way network resources are deployed, expanded, and even moved, has gone from taking weeks to just seconds. How can network engineers make the most of that agility? The answer is, they can't…

But the machines can.

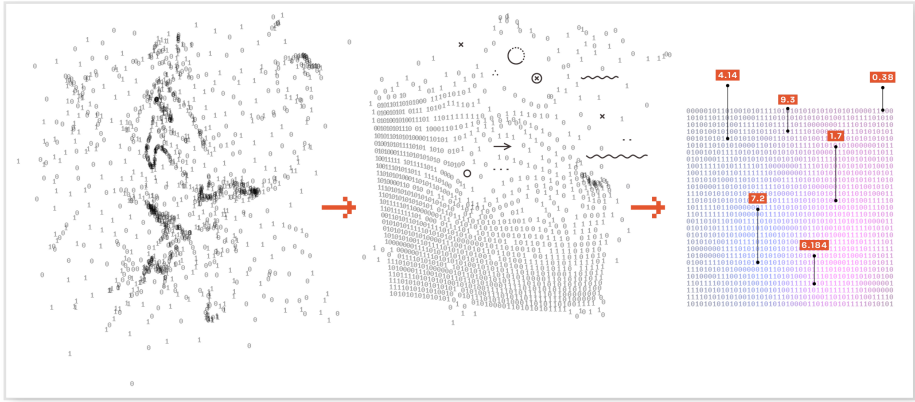

Changing network architectures

The increase in customer data demand is largely driven by new types of service being available to them. Some services are supplied by the CSP; others are supplied by third-parties, often called Over the Top (OTT) services (as they run over the top of the network rather than being run by the network operator).

North and south

In the past, most communication services were point-to-point, connecting two customers or two locations. This was achieved by connecting each ‘end’ of the service to a central point in the network, a telephone exchange for example, and if necessary, connecting two exchanges with a big fat pipe to carry traffic between network regions.

The network looked like this...

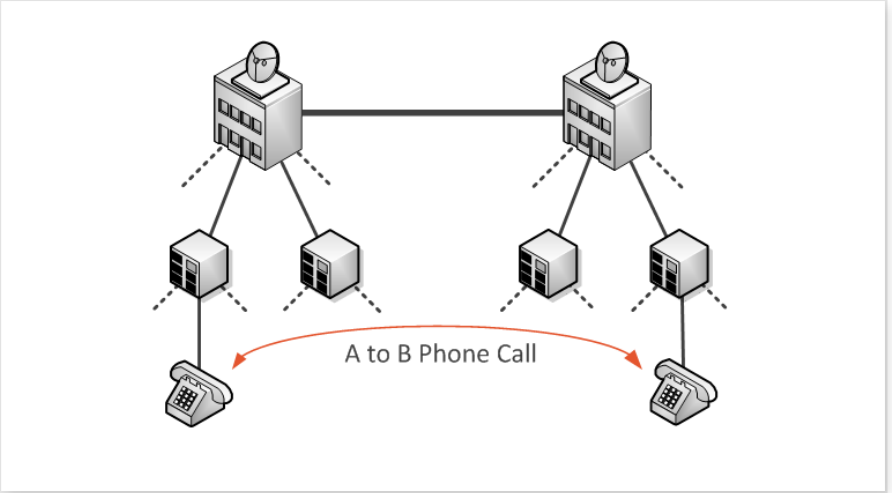

The connection was the service, more or less.

If one phone could connect to another phone reliably, the customers on each end were happy. The network was therefore designed for this ‘north-south’ traffic between customers and centralized network devices in fixed locations.

East and west

Today, in the case of traditional telephony and data services this network architecture and connection-oriented service is still relevant. However, modern services are more than a simple point-to-point connection. New services depend on content or applications that may be delivered from one or more data centers, belonging either to the CSP or OTT service providers.

New services

For example, a comprehensive business communications service could include Cloud connectivity, VoIP telephony, video calling, remote meeting services, email and file sharing all supported by several servers, perhaps in a number of different data centers. It may also include services dedicated to ‘things’ rather than people: Vehicles, meters, industrial sensors – The *Internet of Things* has its own communications service requirements.

Now, consider that the office workers may be mobile between sites, and may want to access these services from multiple devices: Laptops, mobile phones, desk phones and so on… Users may be connecting on dedicated broadband lines, via the mobile network, via the public internet…. And they may be connecting to one, two or more servers in different locations and also have a direct connection to another user for the duration of a phone call.

The service is no longer just a connection between two similar devices.

New networks

On the other side of the coin, the service provider wants to offer these services in a way that maximizes their network efficiency and gives the customers a positive experience. The applications and data that these services rely on could reside in any of their servers, in any of their data centers. The CSP can move applications and data in order to balance traffic on across servers or increase responsiveness and quality.

The building blocks of the service may be moved around on a daily, hourly or near real-time basis, in response to demand and network conditions.

Some or all the service's building blocks may actually be hosted by a partner organization, with the CSP responsible only for connectivity.

The impact on the network is that there is a greater emphasis on 'east-west' traffic: Connectivity between CSP data centers, network gateways, and partner networks.

Also consider that a customer is as likely today to be using a service hosted in a data center in a different continent, as they are to be calling a relative in the next town over. The implication is that traffic is being carried long distance across multiple networks, rather than between two nearby end-points.

CSP networks now need to look more like this…

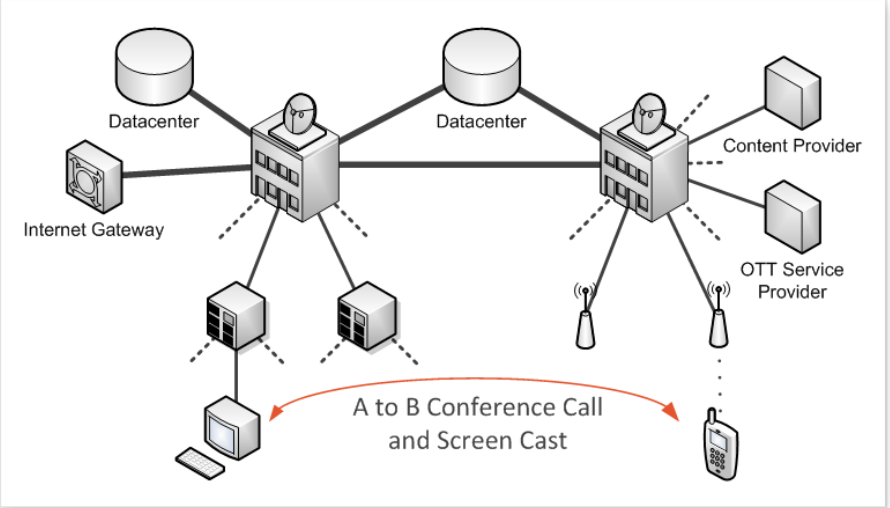

Rebuilding the network…
…with software

Users depend on the well-established north-south network, but clearly there is a need to invest in east-west connectivity, improving both capacity and flexibility.

Two techs are emerging to assist with this network change:

NFV and SDN

While raw capacity will always need to increase, CSPs can be smart about how they deploy capacity and use capacity, to increase network efficiency.

While SDN and NFV enable the network to be smarter, it's still up to OSS to ensure that the technology is used in the right way, in the right place.

SDN

Software Defined Networking changes the way that traditional IP networks make decisions about carrying data, moving the logic from the physical IP devices to a software-based controller.

In traditional IP networks, each device has a relatively static configuration defining how it connects to other IP devices in the network. When the IP devices needs to forward data across the network it decides the best route, combining the constraints of its configuration and common routing protocols like OSPF and IS-IS.

OSPF and IS-IS both allow the IP device to build an internal map of its part of the network, so the device knows the 'least cost' route to any other IP device. This route only changes only occasionally: When devices are added or removed; when there is a fault; or when reconfigured by a Network Element Manager.

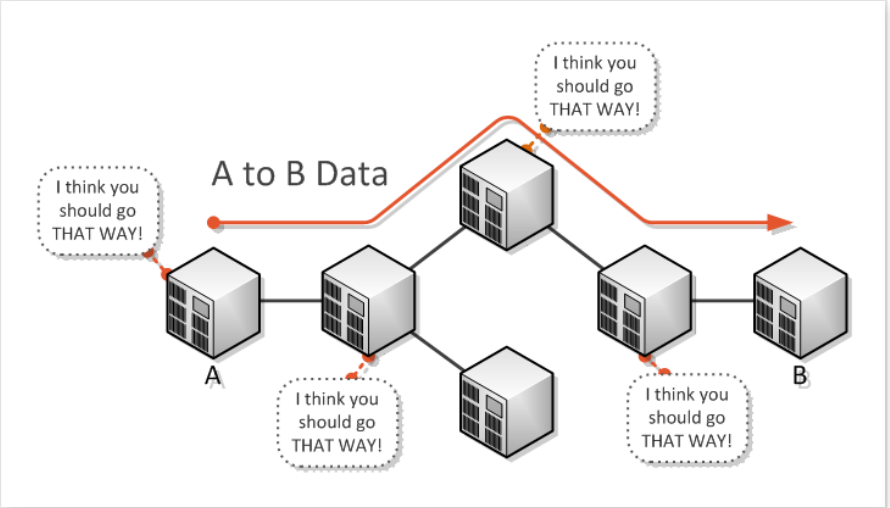

This autonomy was one of the great benefits of IP networks as it meant that the network itself could react to reconfiguration and faults extremely quickly, without needing an engineer to intervene.

The downside is that each IP device is smart locally, but relatively dumb network-wide.

The protocols it uses will mean that a route for data *will* be found but it may not be the optimal route, when you consider things that the individual IP device is not aware of like network congestion, network latency, or an imminent planned outage.

The demands on traffic routing has grown ever more sophisticated. Just getting data from A to B is no longer enough. On that journey from A to B, some traffic might be required to take a different route. For instance, some traffic may need to take a route via a specific network device.

And traffic may need to arrive at alternative destinations to enable load-balancing.

And traffic may need to be quickly redirected when the resources of a service or application are migrated between datacenters.

And traffic may need to be handled differently depending on priority, quality or reliability constraints associated with some services or customers.

SDN takes the relatively simple decision-making about how traffic is routed out of each physical device and puts it in a centralized SDN Controller.

The Controller determines how each device should handle traffic and configures them accordingly. By centralizing the routing intelligence, SDN can make smarter decisions with its broader view of the network, and its ability to be aware of non-technical rules and policies.

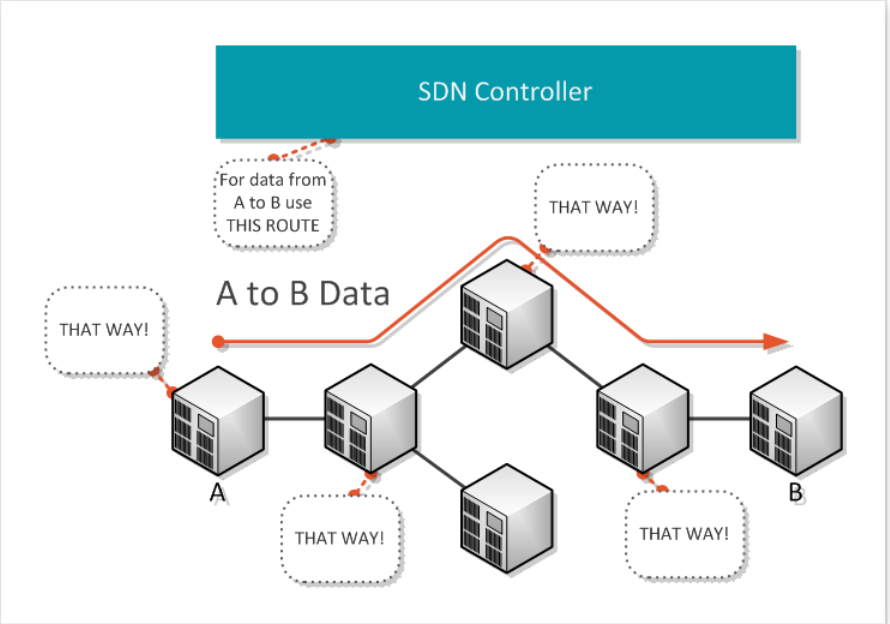

Traffic routing is not all that SDN improves. SDN is also a movement away from proprietary management tools to open API standards. A lot of the IP hardware vendors' work went in to implementing and maintaining protocols in their routers and switches, and that drives up the cost of traditional devices.

With SDN, CSP's get the flexibility they need to deliver today's services, while driving down the cost of operations and integration.

NFV

Network Functions Virtualization (NFV) moves the logic and function of proprietary physical network devices in to a virtual machines running on general purpose servers.

That sounds very clever and complicated. It is, but to put it another way: Rather than installing big boxes to do communications jobs, you install software on generic servers instead.

Consider one example: A router on a business customer's premises. Previously, this would be a dedicated box that the customer's server plugs in to, and then connects to the CSP's network. The business customer might be using an enterprise-class network from the CSP, employing MPLS for example, requiring a compatible router from a big vendor like Cisco or Juniper.

NFV would replace the physical router with a Virtual Network Function (VNF) running on the customer's server. The server would then just need a small interface card fitted to terminate the network connection.

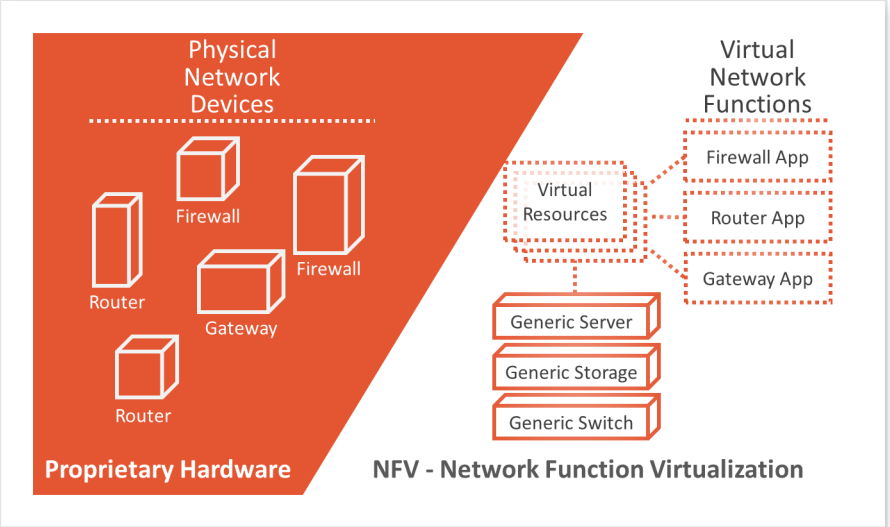

Now, replacing just one box with a virtual router might not sound like a big deal. But if that CSPs is supplying IT services to a customer that has 500 retail outlets, that's 500 less boxes that need to be bought, shipped and maintained.

NFV can also replace network functions at the edge or core of the CSP's network, as well as on customer premises.

NFV means CSPs, and their customers, benefit from virtualization in the network in much the same way that data centers benefit from it: The flexibility to deploy a workload anywhere and scale as needed (quickly if necessary).

So, instead of ordering a physical device, getting it shipped to site, and having an engineer install it before it comes online, NFV enables a software image of the device to be downloaded and initiated on the servers already at that location. NFV can therefore provision network functions in minutes, rather than days or weeks.

Furthermore, NFV allows CSPs to take that workload off expensive proprietary hardware and allows it to run on lower cost hardware, such as Intel x86 based servers, sporting cheap off-the-shelf chips for switching and routing tasks.

MANO – OSS for the virtual bits

A Virtual Network Function can expose the same management and monitoring interfaces as a physical device. In this way, the OSS required to manage a VNF, once it is provisioned and configured, is not much different from the OSS that manages a physical device.

But what about all that generic hardware and virtualized resources? Traditional OSS systems don't know that a VNF is running in a virtual machine, so how are virtualized resources assigned and monitored?

An NFV management and orchestration application (MANO) is required.

MANO looks after the hardware that hosts VNFs, manages compute and storage resource allocation to VNFs, and provides workflow processes for managing the lifecycle of a VNF, including provisioning, resizing, and upgrading.

SDN and NFV

The implications for OSS

The way CSPs architect their network is changing, forced by the change in customer expectations, and the change in the way services and content are delivered to those customers.

With the introduction of SDN and NFV to enable this new architecture, CSPs are going through the biggest technology change since the introduction of carrier grade IP networks.

With this change comes a great challenge for OSS.

Such a fundamental change in data traffic and network architecture demands rigorous analysis, capacity planning, optimized network design and management of change.

Large CSPs will need to invest millions, probably billions, in the long-term to change their network to meet traffic demands for the next decade. Establishing the right strategy to meet business objectives is the responsibility of network architects and planners supported by the right OSS applications.

Traditional planning tools like GIS and inventory will play an important role, but more crucially there is a need for network-specific analytics that can model and predict traffic as it flows around the more dynamic and flexible network topology that NFV and SDN enable.

Forecasting and Capacity Management will play a big part in strategic network planning, with Service Optimization being necessary to determine how best to use the network efficiently on a day to day, or even minute by minute, basis.

Network Activation and Service Provisioning applications will need to integrate with new interfaces, and be able to make far more frequent, software-driven changes to the network.

We've talked about all these OSS applications before, but don't let that mislead you in to thinking that a CSPs OSS environment, with its 10-year-old systems, is ready for this change.

While the role of OSS remains the same, with some shift in emphasis between functions perhaps, the underlying processes require a radical reengineering.

Scope, control and agility

Modern networks and modern services are as much built on servers, databases and applications as they are on routers, packets and paths. So, in addition to being more analytical, OSS must also branch out in to modelling data centers and IT resources with as much sophistication as they have modelled routers and switches in the past.

Network virtualization, and the use cases it enables, means that protocols that once were fixed and predictable are now flexible, and subject to change as demands on the network change. This introduces a new layer of control to the network. This control layer is usually found in the software stack the CSP chooses when deploying technologies like SDN and NFV. The control layer provides a whole new point of integration for OSS, which is a challenge in itself, because these software stacks and their associated standards are yet to fully mature.

More fundamentally, the control layer provides a means of influencing how the network operates in near real time, where in the past this operation was changed either occasionally or entirely fixed by the device type deployed. This breaks a lot of assumptions and processes currently in use by CSP's OSS. Assumptions that traffic-engineering a data route is a hands-on process to be done once during service provisioning; assumptions that when a device is installed in the network, that it can't almost instantly double its capacity to handle traffic. Assumptions that devices and service termination points stay in the same place.

The old rules that OSS software and the CSP's processes were built on – hundreds or thousands of development years of effort – have changed.

The future of OSS is already here

There will be a significant change to the very DNA of many OSS applications and a need to introduce, and automate, new operational processes to meet a CSP's specific business objectives for their new network.

In the short term, as network virtualization is introduced, there will be a temptation to use enterprise IT tools and self-built tools, using network management APIs to fill in the gaps, creating silos of just-good-enough OSS.

Single use, point solutions for managing new technology inevitably are developed, in advance of being integrated into a carrier grade OSS. Such Enterprise software and DIY tools are fine in the early days, possibly the only option, until the technology and services they support reach carrier-scale and become common-place.

Such an approach will lead to disjoint operational decision making and inefficient processes.

OSS took a big leap forward at the start of the 21st century when inventory and service fulfilment applications were able manage the complete stack of physical, logical and service resources.

Learning from history, it's important that CSPs have a strategy for OSS change which is in-step with network change.

The really good news is that more suppliers than ever are focused on OSS. With the shift away from proprietary hardware and protocols to software and virtualization, the telco industry now sees its value coming from the software stack, not the boxes. This has resulted in traditional hardware vendors investing in OSS for their next generation networks, and traditional OSS vendors raising their game to maintain their leadership.

Truly integrated, carrier grade software takes time to develop by the vendors and time to deploy by CSPs.

To paraphrase sci-fi author William Gibson – The future's already here, *it's just not been installed by the CSPs yet*.