

Understanding Enterprise SDN

Introduction

The simplest way to describe Software-Defined Networking (SDN) is to call it ‘an approach to networking in which control is decoupled from hardware, and given to a software application called a controller’. However, there is a lot more to it than that.

SDN, like most fields of technical advancement, has quickly developed its own terminology. It has generated a lot of interest and excitement, and therefore much has been written about it, including a huge amount of information about OpenFlow, network virtualization, control plane separation, Network Functions Virtualization (NFV) and more.

This White Paper gets back to the core of SDN, and addresses the simple questions:

- ▶ What are the drivers behind the emergence of SDN?
- ▶ What effect will SDN have on enterprise networking today, and in the future?

Contents

Introduction	1
How did SDN come about?	2
A ‘new era’ of networking.....	2
1. The Internet has distributed control among network nodes.....	2
2. Data centers have changed the game.....	2
Two co-existing modes of networking	3
What does this mean for the enterprise/campus network?	3
1. Will moving to SDN lower capital costs?	4
2. Will moving to SDN lead to faster introduction of new capabilities into the network?.....	4
3. Will moving to SDN improve network security?	4
4. Will moving to SDN improve network reliability?	4
5. Will moving to SDN provide a better experience for the end user of the network?	5
6. Will moving to SDN lower operating costs?.....	5
How to simplify network management and save OPEX, today	6

How did SDN come about?

After decades of data-networking as a decentralized and distributed process, recent years have seen the development of methods for centralizing data network control. The umbrella term for this centralizing of network control is “Software Defined Networking” (SDN). This term refers to the fact that software applications, like those that move virtual servers around in a data center, or orchestrate file transfers between data centers, define the operation of switching and routing nodes in the network.

A ‘new era’ of networking

A look back over the last few decades can identify a first era of data networking, which was dominated by short-haul communication between mainframes; and a second era, dominated by TCP/IP over the Internet and over LANs.

The current consensus is that a third era in data networking is beginning, and that this new era will see a split—between the distributed protocols used on the public Internet, and the centralized way that local and private networks operate.

Why is this split occurring? There are 2 main reasons:

1. The Internet has distributed control among network nodes
2. Data centers have changed the game

1. The Internet has distributed control among network nodes

The distributed nature of networking intelligence in the Internet has been incredibly successful. An Internet that had just one central point of control could not have grown in the way the current Internet has. The Internet operates in a manner that is astonishing, because it seems almost organic—new network nodes graft on with a minimum of trouble, broken links self-heal, and the entire network spreads seamlessly over international borders.

This lack of a central point of control is absolutely key to the Internet’s reliability. The idea of ‘shutting down the Internet’ or ‘breaking the Internet’ has become the source of plenty of jokes, due to its sheer impossibility. These same control protocols that are used in the Internet have also taken hold in local and private networks over the last 15 years or so. There is no fundamental reason why an office LAN should use the same network control methods as the global Internet. It has just evolved that way because of ‘economies of scale’ in technical knowledge—engineers skilled in developing and operating Internet networking equipment can apply those same skills to LANs and private WANs.

2. Data centers have changed the game

The rise of the ‘data center’ has changed the landscape. Data centers are facilities used to house computer systems and associated components, such as telecommunications and storage systems. Data centers are different to what came before them, because:

- ▶ They contain virtual servers that come in and out of existence, and move physical location, at a rapid rate.
- ▶ Their commercial success rests on achieving remarkable levels of responsiveness and reliability.
- ▶ Data can be frequently transferred in bulk between data centers.

Because a data center is a discrete, contained entity (in contrast to the distributed nature of the Internet), centralizing the control of data center networks can positively enhance the data center’s performance. Centralizing network control makes it possible for the system that moves virtual servers to different locations in the data center, to also be able to reconfigure the network to match the servers’ relocation. The resulting unified control of servers and network makes a very compelling model for a system that, by necessity, must operate with extreme efficiency. Waiting for the network to learn the new location of a virtual server is just too slow. It is far more efficient to ‘rewire’ the network to match the new layout when a virtual server moves.

Furthermore, centralized control of the networks that connect data centers to each other is also beneficial. Bulk data transfers need to traverse the inter-data center links as quickly as possible, so links must be utilized with maximum efficiency. Network-based load sharing mechanisms, like Equal Cost Multi-Path (ECMP), are fine for per-packet or per-session balancing. However, with ECMP, the switching nodes are not aware of the relative sizes and relative importance of the files being transferred in any given session. A centralized application that orchestrates data transfer can have this knowledge of file sizes and priorities. If the orchestrating application can tell the switching nodes which sessions to put onto which links, then files can be transferred with maximum efficiency.

Two co-existing modes of networking

In the new world of networking that is emerging, two distinct modes of data networking will co-exist:

1. Large public networks will continue to operate in a decentralized manner, as centralized control of such networks is simply not feasible. These networks, although generally reliable, operate on a 'best effort' basis. Responsiveness to faults, optimization of bandwidth use, and network latency will be variable in such networks. But, decentralization means that they are highly resilient, and able to tolerate unregulated growth and alteration.
2. Smaller purpose-built commercial networks will utilize SDN to achieve levels of performance, resource optimization, and user-responsiveness that are unthinkable in decentralized networks.

This sort of divergence is inevitable. As data networking becomes increasingly integrated into more spheres of human activity, the range of network requirements keeps expanding. One model of network control no longer fits all.

What does this mean for the enterprise/campus network?

Which model is more suitable for the enterprise/campus network? Do enterprises benefit more from staying with networks that operate on the decentralized, self-organizing model of the Internet? Or, are they better to move to the centralized model of SDN?

To answer these questions, consider the pros and cons of moving an enterprise network to using SDN.

Will moving to SDN:

1. Lower capital costs?
2. Lead to faster introduction of new capabilities into the network?
3. Improve network security?
4. Improve network reliability?
5. Provide a better experience for the end user of the network?
6. Lower operating costs?

These questions are explored in more depth in the next sections.

1. Will moving to SDN lower capital costs?

Proponents of SDN frequently refer to a future in which networking nodes are low-cost commodity units. There are plenty of reasons to be sceptical of this claim. Networking equipment must be reliable, secure, high performing, power-efficient, and more, so the cost of building the equipment cannot dramatically decrease. The chip vendors will still need to recoup the costs of developing new Application-Specific Integrated Circuits (ASICs) for SDN. So, there will not be a dramatic reduction in the cost of ASICs in the near future.

Even though a lot of network intelligence will move to the controller and application software, the sum total of intellectual property in the network will not decrease, so the cost of the network cannot really decrease dramatically. Perhaps some vendors who currently charge a premium for their products might find they need to reduce that premium as the playing field is levelled out, but it is unlikely that the capital cost of data networks will plummet.

2. Will moving to SDN lead to faster introduction of new capabilities into the network?

Moving the intelligence out of the switches and into application software means that the switching equipment vendor is no longer necessarily the provider of the network functionality. The writer of the software puts the functionality into the network.

This opens the field up for multiple third parties to provide network control applications that add specific pieces of functionality. So, rather than waiting for a switch vendor's next software release—possibly months away—functionality could be added to a network in minutes, just by downloading a piece of application software.

The idea of the data network becoming a system that runs applications is an exciting scenario, which is now starting to become a reality. New services are emerging that can:

- ▶ **Detect and block various types of security threat**
The flexibility provided by SDN means that attack blocking can be performed in a very precise and dynamic manner.
- ▶ **Facilitate Bring Your own Device (BYOD)**
Dynamic updates to access rules to accommodate users' changing devices are easily automated by SDN.
- ▶ **Perform scheduled user access permission changes**
The programmatic approach means that a workplace timetable can be automatically converted into a set of scheduled network configuration updates.
- ▶ **Balance WiFi load**
A variety of factors can be measured in the network, to make decisions about where connections should terminate.

3. Will moving to SDN improve network security?

Security depends on blocking malicious users and malicious traffic. The faster the malicious activity is detected, and the closer to the source it is blocked, the better. SDN opens the door to new methods, which ensure that the detection and blocking of malicious traffic within the network is rapid and automated.

Monitoring internal traffic enables suspicious activity to be detected. Then the switch or Access Point directly connected to the traffic source can be instructed to block data from that source. This is a new, and effective, approach to dealing with intrusions and infections. It enables a very rapid and granular blocking of the precise sources of threats.

Similarly, newly connected devices can be interrogated regarding their security status, and automatically quarantined if found to not be up to scratch.

SDN provides a standardized interface for applications to control network nodes, which opens up the opportunity for innovative security applications to enter the market quickly. The applications can focus on the accurate detection of security threats, and leave the SDN OpenFlow-based infrastructure to carry out the process of blocking threats.

4. Will moving to SDN improve network reliability?

Reliability is a combination of:

- ▶ Physical reliability of the network equipment
- ▶ Reliability of software running on the switches and controllers
- ▶ Effectiveness of the protocols that effect failover around broken links, like the EPSR protocol implemented in Allied Telesis switches

SDN does not provide an inherent increase in reliability. SDN has no effect on the physical reliability of the network equipment. Importantly, SDN relies on stable communication between network nodes and the controller. Ensuring the reliability of this critical communication path is very important. Allied Telesis have a unique solution to this requirement—namely the AMF management channel. This auto-created, self-healing communication path presents the ideal way to achieve reliable OpenFlow communication.

5. Will moving to SDN provide a better experience for the end users of the network?

Users' expectations for their data networking experience evolve rapidly. As new devices, connection options, services and communication methods become available, users want to adopt these facilities as soon as possible.

As such, the level of flexibility required in the network is steadily increasing. Providing a high level of flexibility, while still maintaining good security, requires a network infrastructure that is dynamic and responsive.

SDN, with the ability to detect user identities and their network activity, and automatically adapt to dynamic changes, has great potential for enhancing the user experience.

6. Will moving to SDN lower operating costs?

The bulk of the operating expenditure (OPEX) in an enterprise network is in network management. SDN, among other things, provides new ways to centralize and automate network management.

Network management tasks are made easier by centralizing network control. These tasks include:

- ▶ detecting and replacing failed links and failed units
- ▶ adding new nodes to the network
- ▶ resolving users' connectivity issues
- ▶ investigating network slow-downs
- ▶ maintaining security
- ▶ changing switch configurations—even open-flow controlled switches need configuration
- ▶ upgrading switch software—even open-flow controlled switches need software upgrades

The key to truly reducing enterprise network OPEX, is through simplifying network management. Automating routine management tasks, and facilitating insight into what is happening in the network, are key to reducing the cost of network management.

Employing SDN in the management plane, as well as the control plane, delivers these capabilities.

How to simplify network management and save OPEX, today

In conclusion, it is clear that SDN is a maturing technology. It will steadily change the way that networks operate over the coming years, delivering ever more flexible and dynamic data networks. Allied Telesis have embraced this technology to enable organisations to simplify network management, improve security, enhance the user experience, and save costs—right now.

To this end, Allied Telesis have taken a unique, two-fold approach to bringing SDN into the Enterprise:

- ▶ Allied Telesis Management Framework (AMF)
- ▶ OpenFlow

The first part of the Allied Telesis SDN offering is Allied Telesis Management Framework (AMF). This technology focuses on providing the required network management advances.

By embedding management intelligence into the network itself, AMF automates management tasks, like:

- ▶ Backing up network device configurations and software images
- ▶ Adding new units to the network
- ▶ Replacing failed units with new units
- ▶ Making configuration changes to multiple units
- ▶ Rolling out a firmware upgrade

AMF provides truly zero-touch integration of new and replacement units into the network, and single-command automated software upgrades of a whole network. AMF greatly reduces the time and error risk involved in performing repetitive configuration tasks across multiple network nodes.

In addition, AMF operates seamlessly across large networks distributed over several sites. The zero-touch device replacement operates just as effectively at a small remote office as it does in the Head Office. The automated firmware upgrade radiates across the whole network, upgrading units on all connected sites.

The second part of the Allied Telesis SDN offering is OpenFlow. Embedding an OpenFlow control interface into the Allied Telesis network equipment opens the door to bringing standardized SDN applications into the Enterprise environment. The innovative combination of AMF and OpenFlow represents a highly effective foundation upon which to build comprehensive Enterprise SDN solutions.

Visit <http://www.alliedtelesis.com/solutions/networkmanagement> for more information about how Allied Telesis delivers innovative SDN offerings in Enterprise Networking, today.

About Allied Telesis, Inc.

For nearly 30 years, Allied Telesis has been delivering reliable, intelligent connectivity for everything from enterprise organizations to complex, critical infrastructure projects around the globe.

In a world moving toward Smart Cities and the Internet of Things, networks must evolve rapidly to meet new challenges. Allied Telesis smart technologies, such as Allied Telesis Management Framework™ (AMF) and Enterprise SDN, ensure that network evolution can keep pace, and deliver efficient and secure solutions for people, organizations, and “things”—both now and into the future.

Allied Telesis is recognized for innovating the way in which services and applications are delivered and managed, resulting in increased value and lower operating costs.

Visit us online at alliedtelesis.com



NETWORK SMARTER

North America Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895

Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830

EMEA & CSA Operations | Incheonweg 7 | 1437 EK Rozenburg | The Netherlands | T: +31 20 7950020 | F: +31 20 7950021

alliedtelesis.com

© 2015 Allied Telesis, Inc. All rights reserved. Information in this document is subject to change without notice. All company names, logos, and product designs that are trademarks or registered trademarks are the property of their respective owners.
C613-08033-00 RevA