The Cisco Digital Network Architecture Vision – An Overview

Table of Contents

As enterprise business processes become increasingly digitized, new demands on the enterprise network architecture arise.

This overview presents the vision for Cisco® Digital Network Architecture (DNA). The architecture is built to facilitate fast and flexible network services that support digitalized business processes. DNA centers around a network infrastructure that is not only fully programmable and open to third-party innovation, but can also fully and seamlessly integrates the cloud as an infrastructure component. The DNA controller facilitates simple, automated, and programmatic deployment of network services. It brings the notion of user-and application-aware policies into the foreground of network operations. With DNA, the network can provide continuous feedback to simplify and optimize network operations and to support digitalized applications to become inherently network-aware.

Index Terms—Digitalize Services, Cloud, Virtualization, Controllers, Fabrics, Policy, GBP, Application-aware networking, Simplicity, Orchestration, Analytics, Telemetry, Security, Threat-defense, Automation, Openness.

I. Introduction

ENTERPRISE network architectures are being challenged by the evolution of digital businesses that are embracing mobile, cloud, video, and Internet-of-Things (IoT) technologies. Often referred to as the digitization of businesses, the adoption of these technologies can impact our lives much as the introduction of the Internet and World Wide Web did 20 years ago (see [1] [2] for details). The advent of big data and analytics helps to enable better real-time decision making, automation, and efficiencies necessary to deliver such digitalized applications competitively. Additionally, the wide availability of cloud services is supporting these trends. The cloud is rapidly equalizing the opportunities for small and medium companies to cost-effectively and rapidly access advanced IT services that were previously only within reach of established organizations.

However, the evolution toward a dynamic, ubiquitous digitalized business does not come without risks. Universal connectivity opens up new, and often, devastating security risks. Complexity of network operations is expected to increase as such operations become more dynamic, generating more events and data for an increasing number of users and applications. Unless addressed, the network will become harder to manage and operate.

These trends mandate a need for a significant shift in the enterprise network architecture. Although the traditional network virtues of availability and performance are not disappearing, the digital business network needs to quickly evolve to support new more dynamic forms of worker and customer interaction, as well as simplification and automation of network operations.

Cisco Digital Network Architecture (DNA) offers a new blueprint for the digital organization. The network is extended to embrace data center, cloud and IoT infrastructures while maintaining the traditional high availability, scalability and performance characteristics. The DNA infrastructure is designed to deliver services: network services to enable ubiquitous connectivity; security services to protect data and user integrity; and digital services to optimize business applications. As a result, Cisco DNA provides the platform for digital solution delivery to enable a superior workforce and customer experience while simplifying business operations. Figure 1 shows an overview of the Cisco Digital Network Architecture.

Figure 1. Cisco Digital Network Architecture Vision

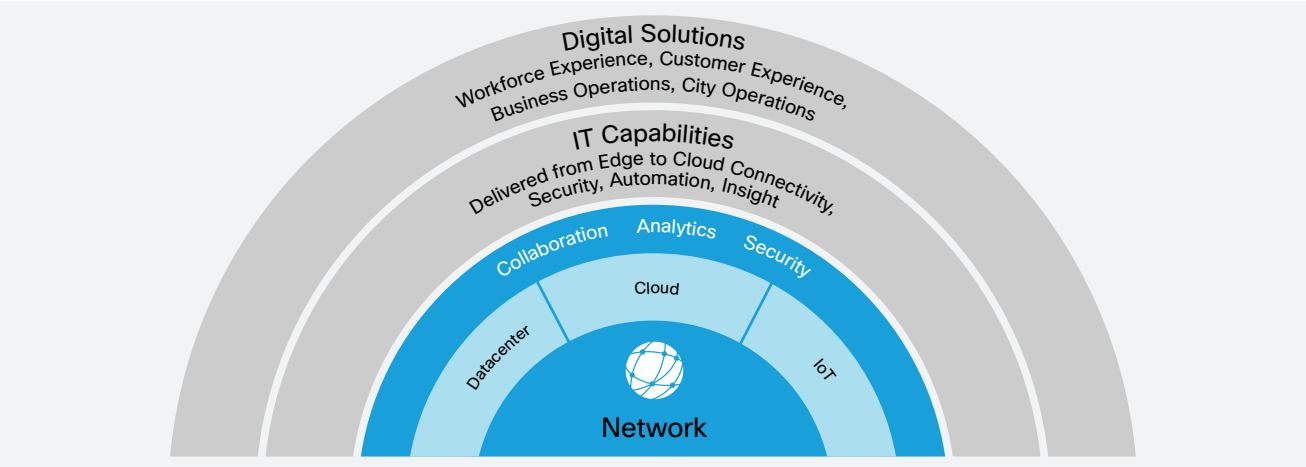

This paper is written for CXOs, network architects, and network engineers who wish to understand the technical details underpinning DNA. Engineering-minded decision makers who want to understand how the architecture supports the business benefits of DNA may also benefit from this paper. Readers will learn the fundamental concepts of DNA. Section II offers a summary of the requirements that an evolved enterprise network has to meet to support digitalized business processes. Section III then outlines the main philosophy of DNA – a network where services can easily be automated, which provides actionable insights for operators and application developers, and which offers full security and regulatory compliance. The network architecture details are then introduced in section IV. The glossary in the Appendix A provides definitions of the main concepts introduced in this paper.

II. Requirements for Enterprise Networks in the Digital Age

Over the last few years, the demands on network functionality have evolved significantly. Wireless is the dominant access method in enterprise networks, enabling the proliferation of endpoint types. According to recent reports [3], [4], over 40% of all network traffic generated originated from non-PC devices in 2014. Today, many employees can bring their own devices into the workspace, which requires the necessary functionality to authenticate and control these devices, as well as to shield corporate data from unwanted access. Multimedia communication has also increased dramatically, with video being the standard communication, whether for direct communication or for meetings.

What is more, many businesses are taking advantage of digitalization to transform their processes. Digitalized business processes require rapid service deployments to meet the ever-increasing demands and expectations of customers, partners, and employees. Enterprises can identify market opportunities and offer business services to customers to capitalize on these opportunities by using rapid, agile application development. Cloud computing allows the deployment of such new services without lengthy and costly investments in server or networking infrastructure. Application delivery is also accelerated by using existing software stacks and drawing additional data from a multitude of information sources. Cloud computing is further used to collect rapid feedback from customers at a large scale through digitalization (i.e., electronic surveys, business analytics), and to engage in shortened improvement cycles for the digitized services to adapt business services according to consumers' feedback.

These trends have not only led to a significant increase in overall IP traffic volumes; they have also increased network complexity. For example, network operators need to manage an increasing number of end devices, guarantee security of the network, reduce the time to enable new digitalized applications on the network, ensure application service-level agreements (SLAs) are met, and provide increasingly granular reports on application performance. As a result, the requirements for the enterprise network architecture have shifted significantly to the following four categories:

1) Faster Innovation Requirements

- **Flexibility:** The number and types of endpoints seeking transport services from the enterprise network are expected to continue to grow. A fundamental requirement for an evolved enterprise network architecture, therefore, is to accommodate a wide variety of different endpoints, regardless of how they wish to connect. The same network and security services must be available to wired, wireless, or mobile hosts. Flexibility is also required from the network to support a variety of network and business applications that are increasingly dynamic in nature and no longer tied to a geographic location. As a result, the network must support network services and applications that can run anytime, anywhere.
- **Intelligent Feedback mechanism:** Significant improvements in network state feedback mechanisms has also become a main requirement. Traditional mechanisms based on Simple Network Management Protocol (SNMP), screen scraping are no longer viable since they are highly manual and prone to errors[1]. The network is expected to provide the correct metrics to support business-critical and transport-policy decisions at the application level. Such data needs to support active troubleshooting, facilitate trending and health- monitoring of network elements and applications, and be easily extendable. Analytics capabilities for the enterprise network are required to be at the user and application granularity, and of course, support a myriad of device types that are expected in a network that supports the IoT.
- **Application and user awareness:** As an enterprise network becomes more and more an integral part of every aspect of the business processes, operators are looking for application awareness to properly treat and prioritize traffic flows or apply the desired transport policies. The network needs to be able to identify applications, even if they are encrypted. Furthermore, the network needs to keep pace with rapid application development cycles and offer the flexibility to introduce transport services for new applications in minimal time. This typically implies a requirement to also support third-party application development, or to have the capability of grouping applications according to a joint transport policy.

2) Requirements to Reduce Complexity and Costs

- **Simplicity:** The simplification of network operations and deployment of new network transport services have become major requirements to offset the increased complexity caused by the trends around device, application proliferation, and dynamics. Simplicity extends throughout the network services lifecycle, from day 0 design and installation of the infrastructure components to day 1 service enablement and day 2 management and operations. Requirements for an evolved enterprise network include fast and error-free service deployment, consistent operations across different element types (routers, switches, access points, etc.), automation of bootstrapping and configurations, automation of licensing and authentication, and thorough reporting in case of failures.
- **Automation:** As the complexity of network operations increases with a growing number of end devices and end users, so do operational expenses. Automation is a primary requirement to limit or even reduce the network operating expenditure (OpEx). The requirement for automation is to support the trends toward programmability, open APIs, standards-based protocols and simplified network operations.

3) Compliance and Technology Requirements

- **Security:** Security has to keep pace with the dynamic application environment and the endpoint proliferation. The ability to segment traffic and users from each other is a challenge for any enterprise network architecture. The network should offer a single segmentation strategy that can be applied to different user groups and applications and be part of the operators' policy framework. The segmentation capabilities of the network should be complementary to the above mentioned requirements of simplicity, application-awareness and endpoint flexibility. Segmentation, however, has to be augmented by tools that can react to security breaches in real time, detecting anomalous network states while still complying to the security requirements of regulators.

[1] E.g., minimal changes by vendors in the output screens may trigger significant operational changes in the network. These techniques are also vendor-specific.

- **High availability:** Finally, ensuring the availability of network services continues to be a baseline requirement. Traditional high-availability mechanisms are expected to improve to offer better predictability and determinism. Any failure of a transport service must be graceful, with configurable orders of operations to prioritize applications according to their business priority. Again, the demands for simplicity, application awareness, and endpoint flexibility need to be supported by high availability.

4) Cloud-Enablement Requirement

The enterprise network needs to be fully integrated with the cloud to support rapid application development, prototyping, and hosting of digitized applications. Cloud infrastructures must integrate seamlessly with the rest of the network from both an operations and a transport perspective. Enterprise network operators no longer wish to distinguish between applications hosted in their own data centers and applications hosted in the cloud. Network operations, analytics, and policy should apply with equal simplicity to either enterprise-hosted or cloud-hosted applications. Digitalized Enterprise Networks are able to take full advantage of the power of the cloud to speed the delivery of innovation and incremental services augmentation. Further, they are able to use cloud-based analytical analysis of crowd-sourced information (telemetry, utilization statistics, etc.) to facilitate rapid enhancement development and service augmentation. ("We know what is working and what is not, and can innovate with that information and deliver it quickly, on a software scale).

III. DNA Tenets

Cisco Digital Network Architecture is designed to meet the requirements already discussed. The design philosophy behind DNA centers around the concept of service delivery, openness, and software centricity. Services in the Cisco DNA fall into the following three categories:

1. Network services deliver the transport to allow end users and applications to communicate over the network. Network services in DNA are designed inherently to be simple and fast to instantiate without loss of functionality.

Figure 2. Tenets of Cisco Digital Network Architecture

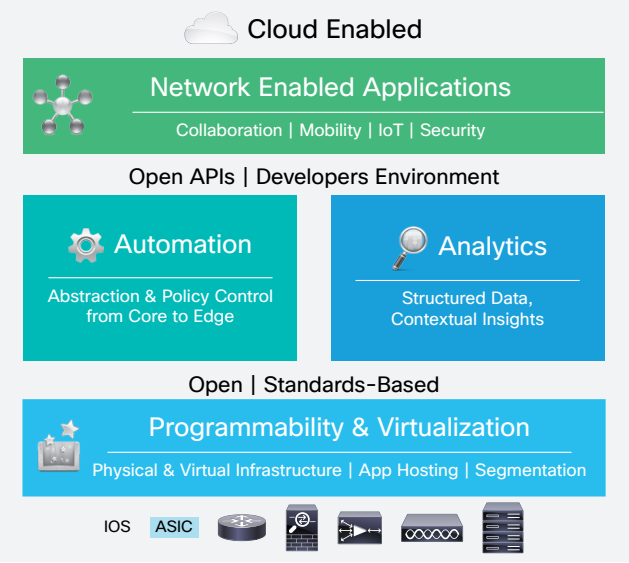

2. Digital services are embraced in DNA to support the digitalized business applications. The network can assist in providing intelligence to the applications, for instance, by providing location information in addition to network instrumentation and telemetry (counters, NetFlow, Application Visibility and Control, Perfmon counters, etc.) or other forms of telemetrics.
3. Security services can ensure that the integrity of the communication relationships between end users and applications, or the information carried by the network, are protected at all times and not compromised.

A Cisco DNA-based infrastructure is open by design. Openness in this context is defined as the ability for customers and partners to guide and influence the operation of the network, by:

- Allowing enterprise operators to integrate with existing or third-party OSS systems (open management). An example of open management is the DNA controller, which provides customers or partners with the ability to develop custom applications.
- Allowing third-party virtualized network functions (VNFs) to be integrated into the DNA architecture (open functionality).

- Relying on standards-based protocols for the operation of the fabrics and the underlay network, thus allowing third-party network elements or hosting servers to co-exist alongside Cisco products (open interoperability)
- Providing open and standards-based APIs so that the network elements can be controlled in a programmatic manner (open programmability).

The majority of Cisco DNA functions are driven by software. Functions to forward and manipulate the IP traffic flows are increasingly being provided in a virtual form factor, thus decoupling the packet-processing software from the underlying hardware (routers, switches). Control of the network elements has traditionally been a software-driven process, but is now being elevated in importance by enhancing this control plane with a centralized controller function, being de-coupled from its underlying hardware. Since such a centralized control point of the network elements is responsible for the entire network infrastructure (possibly domain-based), additional intelligent algorithms can be developed to optimize the operation of the network. Such algorithms may even replace complex protocol interactions between the network elements and therefore contribute to the simplification of network operations[2]. Software also plays an increasing role in collecting telemetry data from the network, processing it, and presenting it to network operators or business applications in the right format at the right time.

The DNA architecture is further guided by the tenets around cloud enablement, network-enabled applications, automation, analytics and telemetry, and virtualization. The remainder of this section elaborates on these tenets.

A. Cloud Enablement

The DNA platform fully integrates different cloud models, including private clouds, virtual private clouds, hybrid clouds, or public cloud environments. This integration is designed, both at the transport and the control layer of the network. At the transport layer, the various cloud environments can be fused into the enterprise network seamlessly by use of tunneling techniques. For example, a virtual IPsec gateway can be instantiated in a virtual private cloud (VPC) and connected to the enterprise-operated fabric by means of an IPsec tunnel. Configuring the virtual router with similar policies and functions as the enterprise-operated branch environments then makes the VPC behave like other branches that are part of the network infrastructure. At the control layer, management and orchestration tools can be used to operate the network. Analytics and telemetry engines accessing and storing information in the cloud are also part of the cloud enablement aspect of DNA. By making the cloud an integral part of the enterprise network, the DNA is addressing the requirement for cloud enablement to drive rapid application prototyping, fast application deployment, and continuous feedback loops to speed digitized business processes.

B. Network-Enabled Applications

Empowering business applications with support from the network infrastructure is a key tenet in DNA. Because the digitalized business applications are increasingly dynamic – being developed in an agile manner – the network has to provide the capability to flexibly enable communication paths between end users and applications while supporting key transport characteristics such as security, high availability, and quality of service. It provides the granularity and scale to offer and instantiate the network services to digitized endpoints or endpoint groups at the application and user identity level. Anchoring a service policy based on the application type or user identity allows business services to be appropriately prioritized throughout the network. In DNA, applications are also supported by a network feedback mechanism, allowing the applications to be optimized based on the usage or communication patterns. For example, the network can provide continuous feedback for voice applications to the controller, which in turn may trigger the instantiation of an updated QoS policy to ensure voice quality. A network running DNA continuously collects data that can be mined to improve network operations or to support digitalized applications. The concept of a DNA controller outlined later in this paper plays a critical role in this function.

[2] For example, the optimal path for a particular endpoint group can be algorithmically determined by software based on full view of the network state. This can replace protocol interactions to exchange state between network elements and the execution of a distributed path optimization algorithm.

Enabling applications to be inherently network-aware is particularly relevant to:

- **Collaboration:** Networked collaboration applications allow employees and partners to communicate effectively regardless of geographic location. The DNA supports collaboration with voice and video (e.g. Cisco TelePresence®, instant messaging) in real time with high-quality, helping to ensure that stringent SLAs are met. What is more, these collaboration applications are less frequently tied to particular end devices and are enabled on a multitude of end hosts. Non-real-time applications such as document sharing are also increasingly network-based, often taking advantage of the cloud to foster collaboration (for example, SmartSheets, Box file sharing)
- **Mobility:** Access to digitalized applications by employees, partners and customers is increasingly mobile. Cisco's DNA not only supports such mobile access: it also can act as a sensor to provide feedback to support and enhance such applications. Mapping out movement patterns of mobile users or collecting data on device types and location are examples where the network can assist to enhance mobile applications.
- **IoT:** An increasing number of business applications are running on devices other than traditional compute hosts. A Cisco DNA infrastructure extends the network to seamlessly include **any** connected device types used in various vertical industries (i.e., healthcare, transportation, retail, etc.). Applications controlling such devices or telemetry data collected from them can be networked to run in the enterprise data centers or the cloud, or they can even be distributed onto network elements with compute resources.
- **Security:** Digitalized business applications can be supported by network security services. Application traffic flows can be encrypted when passing over untrusted domains in the end-to-end communication path. Furthermore, the network can detect communication anomalies or allow policies to be associated with the application to restrict individual users or groups from access.

C. Automation

Network automation is not only essential in the network to reduce OpEx as outlined above. It is also required to deliver the network, digital, and security services within minutes rather than weeks or months, as is traditionally the case. Central to automation is the concept of the DNA controller. Triggered by the orchestration applications, new digital services are deployed by the controller onto the relevant network elements at the right time. The controller offers a network abstraction layer to arbitrate the specifics of various network elements in this automation and toward the orchestration and analytics engines. Consistent, machine-based interfaces are being developed as an integral part of DNA to enhance automation techniques. These interfaces will also be used by DNA controllers to provide scale beyond what is available today. Speed and agility are critical requirements for service rollout and network operations, and a controller-based system helps to facilitate this. Network functions can also be instantiated by the controller.

The DNA controller is vital to drive the **policies** associated with digitalized services consistently throughout the network infrastructure. It translates the business intent for the digital services into actionable and verifiable network policies. Such business intent of a digital service may result in multiple (possibly domain-specific) network policies that need to be instantiated and monitored to instantiate the service. The DNA controller, therefore, implements policy in any part of the network that a service instance reaches, such as the cloud or the campus, WAN, or data center domains, for all variants of policies, such as those governing access, transport or path optimization (see below).

D. Analytics and Telemetry

The ability of the network to assist in a feedback loop to the applications is a novel enhancement to the network infrastructure as compared to traditional enterprise networks.

Telemetry provides for an automated process to communicate various measurements about the network and service state to an analytics engine. Support for telemetry and analytics comes in the following forms with the DNA infrastructure:

- **Network sensors:** Enhance the network elements in the DNA fabric with the mechanisms to monitor and collect data for infrastructure events and service delivery state so as to allow for their respective characterization. The network elements for a telemetry-enabled network are expected to generate a vast amount of data, depending on the monitoring policies of the operator.

- **Communications infrastructure:** Support the efficient transport of the telemetry data for analysis by the analytics engine in the network. Depending on the granularity of the analytics required, this may require a substantial amount of bandwidth and associated SLAs.
- **Analytics engine:** The analytics engine is also part of the Cisco DNA, typically hosted in the enterprise data center, or even in the cloud. DNA telemetry and analytics are, therefore, an important tenet to address the requirements to support pro-active feedback loops.

E. Virtualization

To support the requirements for flexibility, analytics and telemetry or geographic independence of networking functions, DNA fully embraces virtualization. Virtualization allows the creation of logical networks at the transport layer by means of VLANs, virtual routing and forwarding (VRF), or logical tunneling techniques, helping to enable the separation of different service flows at Layer 2 and Layer 33. Transport virtualization is enhanced in DNA by the use of network function virtualization (NFV), which allows network functions like Network Address Translation (NAT), firewall, deep packet inspection (DPI), etc. to run in a virtual machine. The key benefit of virtualization is to de-couple the service transport architecture for both IP packet forwarding and virtualized network functions from the underlying hardware (physical links, physical hosts). This, in turn, gives operators the flexibility to rapidly deploy functions anywhere in the network based on factors other than physical location, and significantly increases the speed of deployments. For example, a firewall function required by a service policy may be instantiated in a branch, rather than the enterprise's data center within minutes. VLANs, virtual routing, or overlay networks can then be deployed to steer the service flows through the firewall. This contrasts with the deployment of physical firewall functions, which today take weeks or even months to install in enterprise branch environments.

Figure 3 shows further details on the main tenets of the architecture. It illustrates that applications run by users or things are connected to digital services in a manner that honors the service intent (and implementing a policy). It shows that the cloud can enable multiple aspects of the architecture: applications driving the digital business processes can run in the cloud in a transparent manner. Network functions such as web security, WAN optimization or intrusion prevention or detection systems (IPS/IDS), or cloud management applications can also run on cloud infrastructure. Analytics applications are cloud-based to process network telemetry and to offer the results as feedback into the application development cycle.

Figure 3 also depicts the responsibility of the controller to automate both the network infrastructure and digital service instances. The figure highlights that the controller offers a network abstraction layer to arbitrate the specifics of various network elements in this automation and toward the orchestration and analytics engines.

Finally, the network infrastructure layer is also represented in Figure 3, by showing both physical network elements and virtualized network functions. For those functions that require high throughput levels and supply long-term static resources, physical elements may be more cost-effective. Other functions may be provided in a virtual form factor, such as for rapid and flexible deployments, increased service granularity, or rapid prototyping, or in cases where the resources are required for a limited period of time. Both physical and virtual infrastructure elements are part of the transport fabric to connect the applications running in user-devices, data center hosts or in the cloud.

IV. Digital Network Architecture Overview

A network architecture-centric view of Cisco DNA is depicted in Figure 4. In the architecture, endpoints[4], including applications, connect to the network to seek transport services. All endpoints are outside of the network domain, and are crossing a user network interface (UNI). A transport service in DNA is defined as the transmission of IP flows from an ingress to an egress UNI, that is, as the transport of IP packets making up the flow between applications running on end-user devices or hosts[5].

[3] Transport virtualization is often synonymous with segmentation.

[4] See the glossary in section A for definitions.

[5] As a consequence, a network service often touches multiple network elements of functions in the architecture that need to be configured in support of the service. Also, a network service is typically bi-directional.

Figure 3. DNA Concepts in Detail

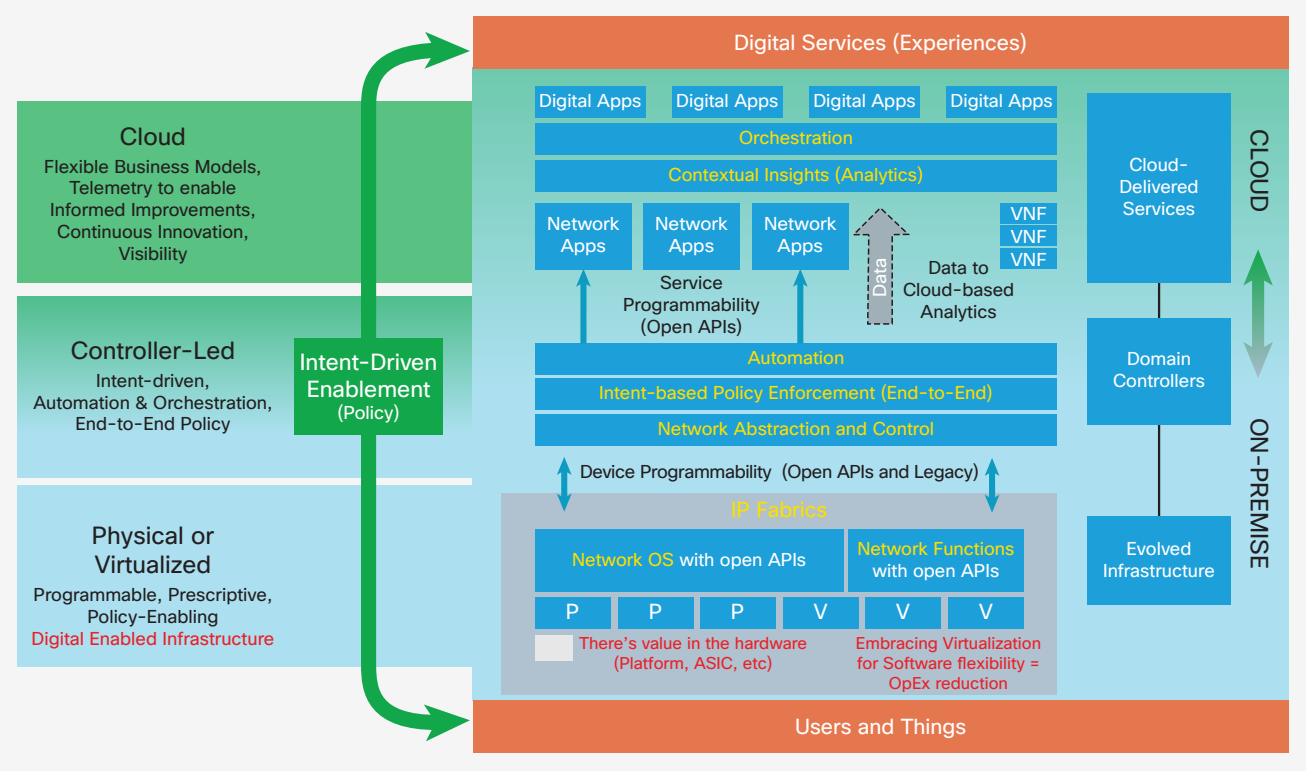

These IP flows may be modified by the network by applying network-based Layer 4 to Layer 7 services. Any host or application seeking services from the network are governed by a policy enforcement point (PEP), specifying which services a host, application, or a related communication relationship can receive from the network. The DNA then provides the connectivity and transport services, allowing endpoints or applications to communicate, regardless of whether these are hosted within the enterprise infrastructure or in the cloud.

The main building blocks of a Cisco DNA infrastructure include:

A. Network fabrics

B. Virtualization

C. Cloud enablement

D. Network controller

E. Service definition and orchestration

F. Analytics and telemetry applications

The concepts are introduced in the remainder of this section and elaborated upon in sections IV-C, IV-D, IV-B.

A. Network Fabrics

The transport infrastructure is provided by the concept of an enterprise IP fabric, defined as a network exhibiting the following characteristics:

- Any-to-any network element connectivity provided by an IP-based underlay[6] infrastructure
- Flexible service delivery to endpoints (users, devices, applications) based on programmable overlays
- Policy enforcement at the user or application level at the edge of the enterprise network and optionally between domains

[6] See appendix A for terminology definitions of "underlay" and "overlay".

- Localized encapsulation to allow hosts and applications to connect to the network using a variety of Layer 2 protocols (i.e., different Ethernet encapsulation types)
- Location–independent forwarding to help enable hosts and applications to move without requiring changes in the underlay infrastructure
- Controller–based management to simplify network operations, in particular to make them networkwide instead of individual network–element–based
- High availability to assure resilience of the network against failure of its network elements or software functions

The fabric can be organized into different domains to provide an administrative boundary around a set of network elements. In this way, the end–to–end transport communication path associated with a service in DNA can be super–imposed on various domains, each providing a separate leg and offering domain–specific policies. The pertinent domains in DNA are the campus, the data center, the cloud, and the WAN (which includes WAN aggregation and branches), as illustrated in Figure 5. In the campus, multiple endpoints connect to the network using either wired or wireless access. The campus may also provide connectivity to the data center or WAN domains, helping endpoints to reach applications residing the data center, or to reach hosts located in branches, respectively. The WAN domain typically consists of many branches that are aggregated across an enterprise– or service–provider –managed WAN into multiple WAN aggregation sites. The WAN aggregation sites connect into the campus domain. The data center domain provides the infrastructure to connect servers hosting enterprise applications to the network. A detailed description of the DNA fabric architecture is described in [5].

The network elements in the fabric are predominantly hardware–based to offer high–speed transport at scale. Assuring low latency packet transport or quality of service (QoS) at high speeds typically requires sophisticated ASICs to be deployed in the network elements, such as Cisco Unified Access™ Data Plane [6] or Cisco QuantumFlow Processor™ [7]. Providing security in the fabric through encryption or imposition of security–group tags is another reason for hardware support in DNA network elements. Furthermore, to meet the requirements for fast and dynamic service instantiation, these ASICs are fully programmable in the DNA fabric. Finally, the requirement for a high volume of telemetry data to be collected and delivered to the analytics engine is also a reason for the use of high–speed ASICs in the fabric. For a more comprehensive overview of ASIC technologies, see [8].

Figure 4. Overview of DNA

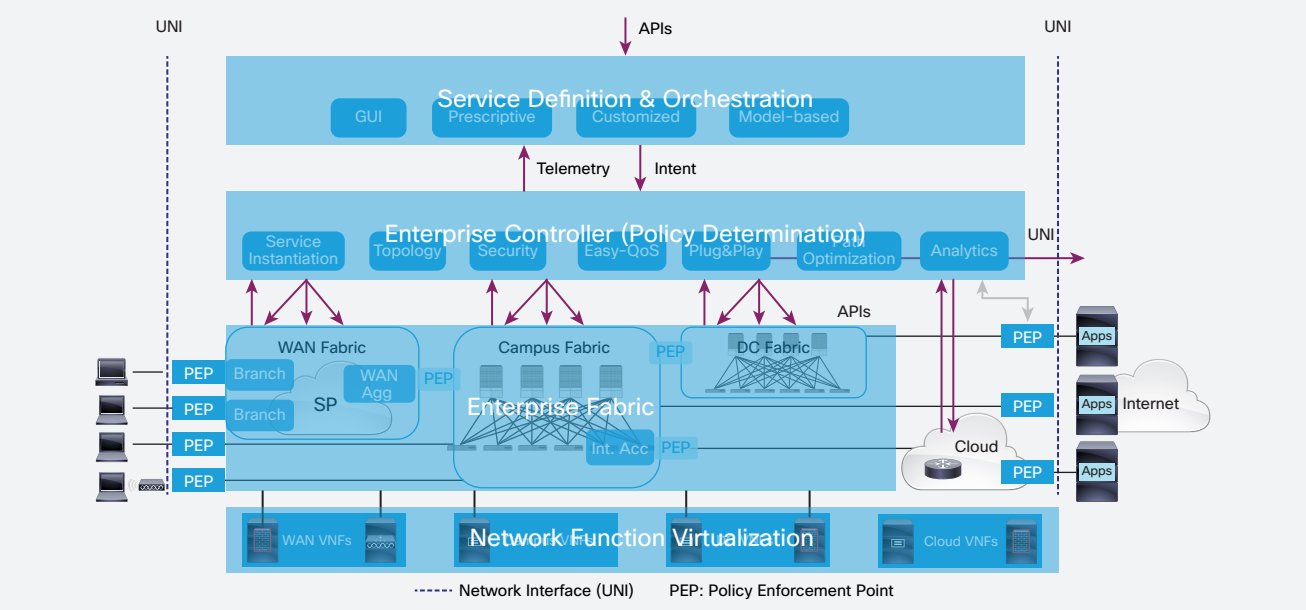

Figure 5. Network Fabrics in DNA

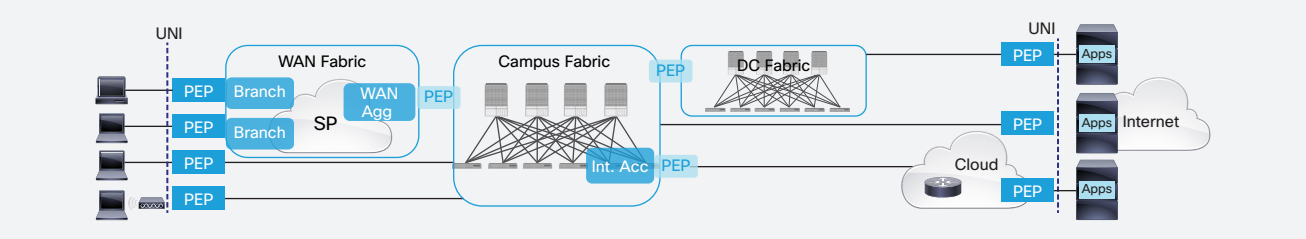

B. Virtualization Architecture

Virtualization plays another key role in DNA, since it is a crucial mechanism to deploy services fast and with minimal dependencies on the underlying hardware infrastructure. The virtualization architecture that is part of DNA can be broken down into two main components:

- **Transport virtualization:** The logical separation of traffic by means of VLANs or VRF is already a standard tool in enterprise network architectures. The network elements comprising the enterprise fabric are fully capable of logical separation of traffic at both the Layer 2 and Layer 3 [9]. These concepts carry forward to DNA, but gain even more importance to ensure that services are associated with the right segmentation policies within the network. Recall that a service connects endpoints or endpoint groups; the service may, therefore, not only be application-aware, but also user-(identity) aware. The logical segmentation of groups of users and applications, then, is an essential component of the enterprise fabrics overlay architecture. [10] elaborates on the technical details of transport virtualization.
- **Network function virtualization:** Network function virtualization (NFV) is part of the architecture that can enable network functions to run anywhere in the network infrastructure based on the availability of x86 compute resources. The virtualization of network functions that manipulate the IP traffic flows according to the policies is essential in DNA to provide a completely virtualized environment. Transport virtualization is extended in DNA by allowing these network functions to run in virtual machines or containers (for instance, LXC, Docker), on any available x86-resource based on the operator's policies. Network elements increasingly offer x86-based compute resources for this purposes. In cases where these resources are insufficient, dedicated x86 servers may be co-located throughout the DNA infrastructure, as illustrated in Figure 6. The operating systems of the network elements in DNA are enhanced to support this virtualization, providing the ability to spinup or tear down network functions within minutes, to monitor their status, redeploy, restart, or even support their move from one server to the next.

In this way, NFV allows for novel architectural approaches where the functions applicable to a service flow can be distributed in the infrastructure based on optimal placement policies. They no longer are tied to physical network elements, nor are these functions required to be executed in a single process on a network element. The concept of chaining a service flow through multiple virtualized network functions based on policy is thus also supported by the DNA architecture. Note that virtualization in DNA is thus seamless between the different fabric domains. A virtualized network function such as DPI or a firewall may be deployed in a virtual machine (VM) on a host connected to the campus, or it may be deployed in a VPC. The seamless extension of the enterprise fabric architecture into the cloud by means of encrypted tunnels carries forward to VNFs and gives the operator the flexibility to decide where to best instantiate a network function[7].

Service function chaining [11], as specified by the IETF, provides an open mechanism to implement such chains of virtualized network functions. As in the case of standards-based north-bound and south-bound interfaces in the controller, this allows third-party vendors can contribute in DNA. Figure 6 shows the two aspects of virtualization, transport virtualization and network function virtualization, in the context of DNA. Technical details on the DNA NFV Architecture can be found in [10].

[7] E.g. based on cost or SLA criteria

Figure 6. Virtualization in DNA

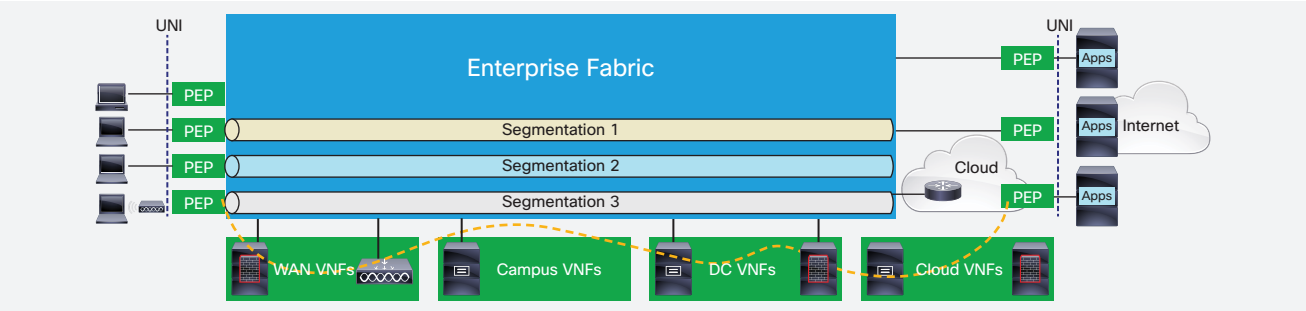

C. Cloud Enablement Architecture

Fusing the cloud with the enterprise-operated infrastructure is an integral part of the DNA. This cloud integration is crucial in order to deliver several benefits. First, the cloud provides resources to host digitized applications for business process delivery as if they were hosted in the enterprise-owned infrastructure. In the case of virtual private cloud environments, making use of a virtual IPsec gateway[8], for example, allows the establishment of a secure tunnel into the VPC, effectively making it another branch. Different VPC providers can simultaneously be integrated into the corporate network. And, the same network functions that can be applied in a virtual form factor in the branch can also run in the VPC, providing additional consistency from an operations perspective between the enterprise-hosted infrastructure and the cloud. This consistency is depicted in Figure 7a by example of a firewall, Cisco Wide Area Application Services (WAAS), and IPS. In case of public cloud environments, applications can be accessed by the devices on the enterprise network by means of a public cloud gateway, which provide a demarcation between the public and the enterprise domains. The public cloud gateway implements the policies pertaining to those applications hosted in the public cloud. Figure 7c depicts an evolution to the architecture shown in Figure 7a. Here, the networking functions can even be packaged into one entity and deployed either on a standard x86 host in the branch or in any VPC environment. This packaging further simplifies the operational consistency between the enterprise-hosted environment and the cloud, because the interactions between the networking functions applied to a service can be packaged once but deployed in various environments. Figures 7a, 7b, and 7c show different types of cloud application hosting architectures.

These applications can also be orchestration and management applications to run the enterprise network, rather than applications to support the digital business processes. The advantages of the cloud, which include dynamic provisioning of resources, resilience, and global reachability, are extended to the DNA orchestration and management tools:

- Compute and storage resources can be dynamically added or removed, depending on the requirements of orchestration and management tools. Enterprises no longer need to consider the investments in procuring and operating the servers for their management and operations tools.
- The operation and management tools benefit from the availability and resilience of cloud architectures. Enterprise operators no longer need to provide highly available (and geographically distributed) infrastructures to operate the network.
- By running a public cloud gateway in a VPC (see Figure 8), orchestration and management functions could be accessed from anywhere. This allows the operators to perform their duties from anywhere.

The third application for a full integration of the cloud into the DNA platform is to supply telemetry and analytics to provide application developers and network operators with real-time feedback loops. This concept is described in detail in section IV-F.

Details on the fusion of the cloud into Cisco Digital Network Architecture can be found in [12].

[8] E.g. by deploying a Cisco CSR 1000v

Figure 7. Cloud Application Hosting Architectures

(a) DNA Extensions into VPC

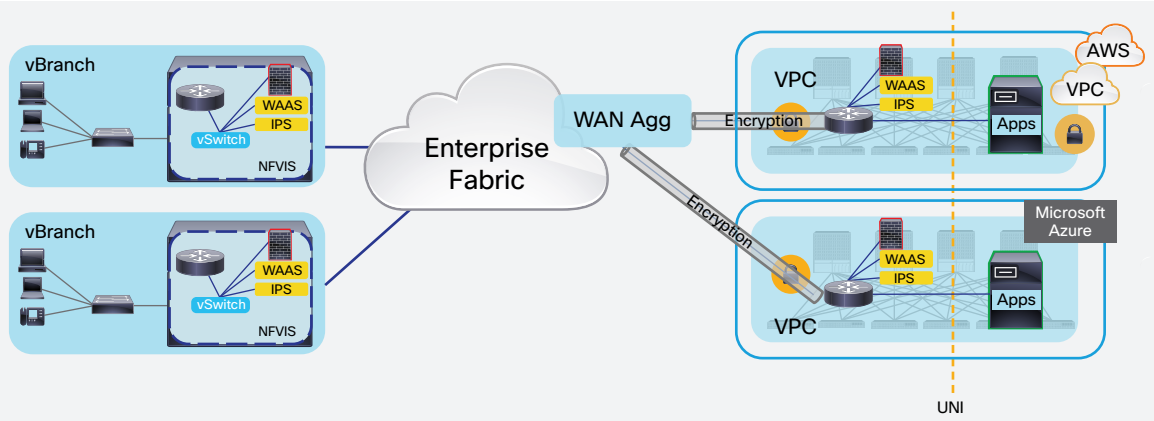

(b) DNA Extensions into Public Clouds

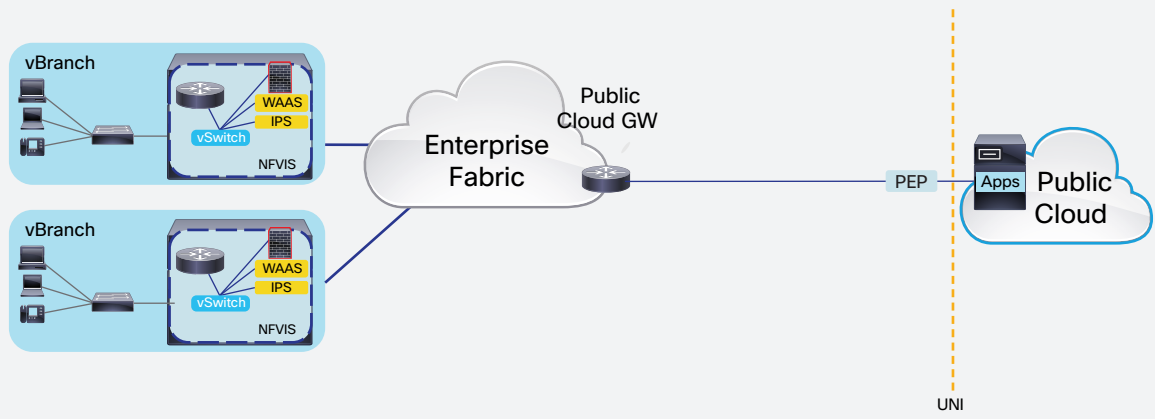

(c) Enterprise–NFV packaged functions for cloud and branch consistency

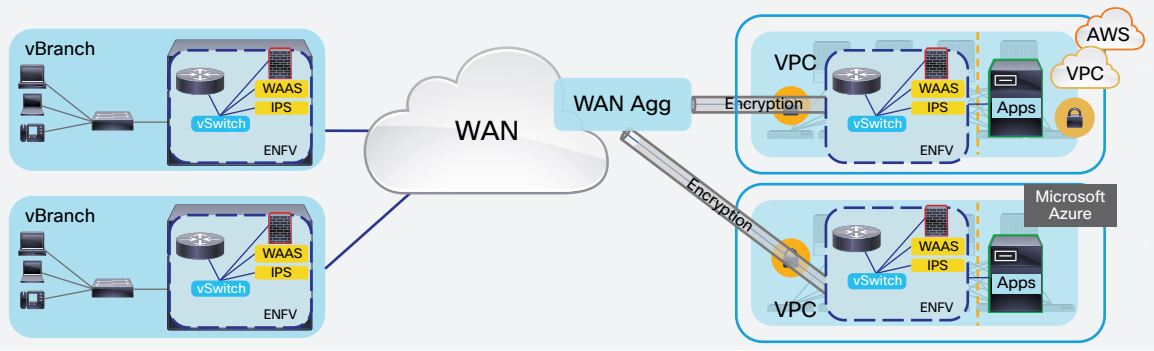

Figure 8. Cloud Management and Orchestration

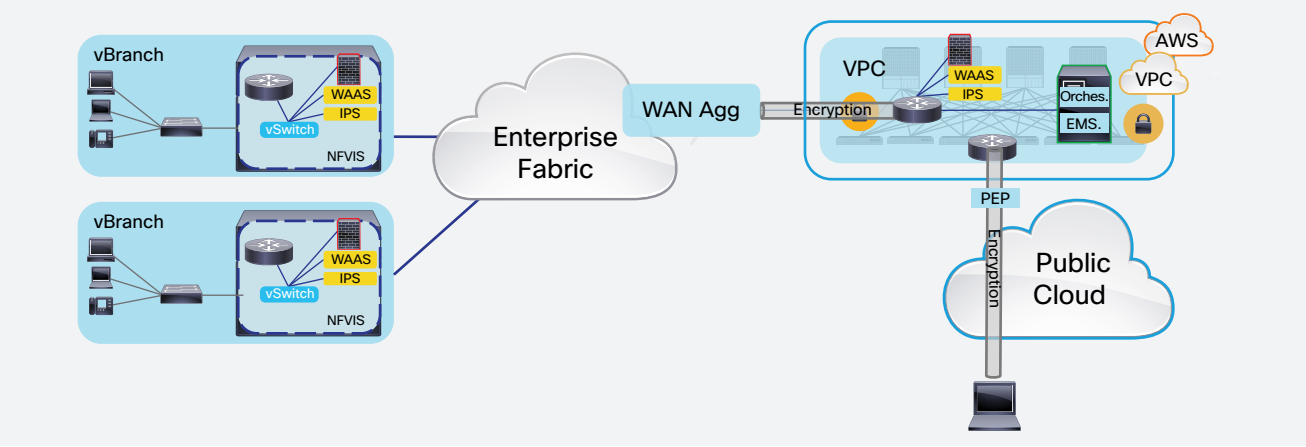

D. Controller-Based Networking Automation Architecture

The enterprise IP fabric in the Cisco DNA is governed by a controller that oversees the configurations and operations of its network elements. As such, the controller has a domain-wide view of the state and the operations of the associated network elements (Figure 9). The DNA controller is responsible for the configuration of the network fabrics – the underlay communication and overlay services architectures. It configures the user- or application-visible network services (for example applying a service policy to a particular IP flow, or applying one or more Layer 4-7 service functions to a user-application communication pattern). Most importantly, the controller is the element in the architecture to implement policy.

Policy and intent is one of the key architectural shifts supported by DNA. Recall that each service offered by DNA implements a business intent and is instantiated by a transport path with the associated policies. The policies are correlated with the services, and implement the intent of the business service. Policies for network services can fall into the following categories:

- Network transport policies regulate the traffic flow relationships between users, applications, or hosts. A transport policy determines who can communicate with who, as well as the details of the transport (i.e., with which VPN the service flow should be associated with). The network policy is typically applied at the UNI into the network.
- Security and encryption policies help to ensure that the service is treated with the right security in different parts of the network. For example, when a service traverses an insecure network domain, application traffic may need to be encrypted. Security policies also include applying additional functions such as firewalling, anomaly detection, and blacklisting or whitelisting to a service.
- Traffic engineering policies determine the optimal path that a service flow should take trough the network with application granularity.
- Network function policies. In addition to transporting service flows between endpoints and applications, the network may offer additional network services that manipulate the IP packets. For example, a flow may be optimized using a WAN optimization function, may be forced through an IPS or IDS service, or may be tagged to be inspected at the payload level (DPI). The network services policies determine which additional network services should be applied to an IP flow, and they may rely upon NSH-based service function chaining (SFC) to force a flow through the required services.

The policy layer allows the specification on how a particular network service defined above is to be treated by the network. This specification of policy is done with the assistance of a policy engine and an associated application. Standards-based policy descriptions, e.g. based on group-based policies [13] or [14].

The benefits of the DNA controller's policy functionality can be illustrated by example of application-aware services. The network operator may specify in the policy application that a particular service applies only on a per-application basis. The controller must then translate this service policy into access policies to be applied at the policy enforcement points (PEP), ensuring that the right applications are filtered out at the UNI. Such a filter may be based on standard DPI techniques, such as Cisco Application Visibility and Control [15]. Alternatively, mechanisms such as DNS as authoritative source (DNS-AS) [16]) can be deployed to create such filters. DNS-AS helps to classify applications with the help of the DNS server, where the application type can be specified as part of DNS record and returned to the requestor in the DNS response[9]. The controller may also need to deploy policies between domains or within a particular fabric domain to affect the right treatment for the service, such as mapping the service to a QoS transport class. The domain-wide view of the controller, therefore, helps to instantiate the right policies in the right places of the network, and ultimately delivers on policy-based service delivery. Figure 9 depicts the main functionality of the controller: taking intent as an input from the orchestrator, computing the right policies and service constructs and pushing those into the DNA fabric and NFV components using APIs.

Programmability is a critical supporting aspect of the DNA controller. Configuration of the underlay network, the overlay architecture, or even specific services is handled by a south-bound interface between the controller and the network elements and functions. This south-bound interface may rely on command-line interface (CLI), or other standards-based mechanisms such as YANG [17], representational state transfer (REST) [18] or REST Configuration Protocol (RESTCONF) [19]. Supporting standards-based southbound interfaces contributes to the openness of the DNA and also allows also third-party vendors to participate in the network infrastructure. Towards the orchestration layer, the controller provides a level of network abstraction by exposing northbound APIs, which again may be standards-based. The controller can enable a simplified operation of the domain fabric, and supports rapid and elastic provisioning of network services. The programmability aspect of the DNA controller is vital to delivering on the automation, as well as fast and flexible configuration, of DNA.

From an architecture view the entire enterprise network is governed by the DNA controller, which regulates the operation of all network elements using southbound APIs. The DNA controller thus spans network elements in the enterprise branches, WAN, campus, data center, or the cloud.

Figure 9. Controller in DNA

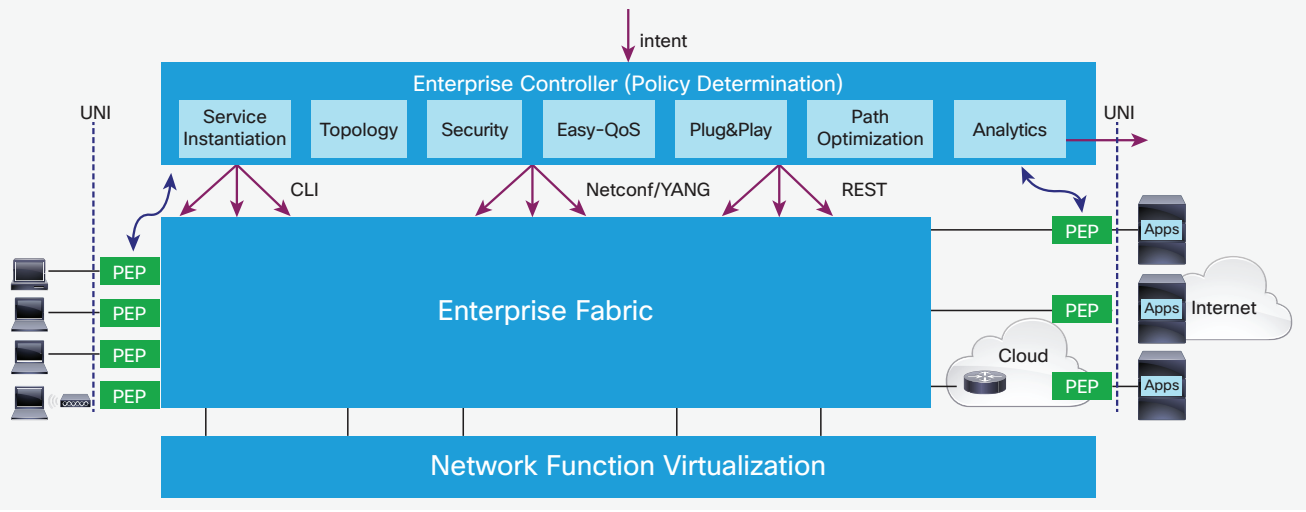

[9] The DNS requests from clients can be snooped to trigger a network element's own DNS request, thus acquiring an application's type for classification.

Note again that multiple domains may be governed by a single controller instance. For example, the Cisco Application Policy Infrastructure Controller Enterprise Module (APIC-EM) [20] may serve as a controller for both the WAN and campus domains. Similarly, a controller may be further divided into several subcontrollers, each fulfilling a distinct role for a particular domain.

Details on the DNA controller architecture can be found in [21].

E. Service Definition and Orchestration Architecture

Orchestration in DNA plays the critical role of allowing the network operator to specify the services offered to applications and end devices and to associate the desired characteristics against those services (i.e., the transport or security policies). The orchestrator thus enhances the concept of network abstraction offered by the controller, abstracting from the details of network elements or virtualized network functions and providing the focus on the services. This is important for the network to deliver business intent. The operator focuses on the specification of the service details from the service consumer (i.e., the applications or end devices) point of view and express the service intent. The orchestrator then interfaces to the controllers using standard APIs to instantiate the specified services at the right time in the right network element in the right sequence of operations.

Stipulating network services and their associated characteristics can be done in DNA in multiple ways. The operator can create standard templates for services using a graphical user interface. Alternatively, a service can be characterized using a graphical tool. For example, the Cisco Enterprise Services Automation (ESA) tool allows the customized specification of enterprise branch profiles. A network extension service (adding a new branch to extend the network reach) can be defined using ESA, empowering the operator to not only list the functions required in a particular branch, but also influencing their order and associated policies for the service. Model-based approaches to service declarations are also part of the orchestration architecture in DNA for those operators who seek further abstraction. Figure 10 shows the relationships between policy and orchestration.

Note that Figure 10 also illustrates the feedback loop from the controller into the orchestrator. Data collected by the network elements and pushed back into the controller can thus be analyzed and presented to the controller in a concise and condensed form, focusing on the relevant data sets required to optimize or refine the service operations. Figure 10 underlines the importance of APIs in DNA, helping to enable a programmatic and extensible platform to continuously innovate around services.

Details on the DNA orchestration architecture can be found in [21].

Figure 10. Policy and Orchestration Relationships

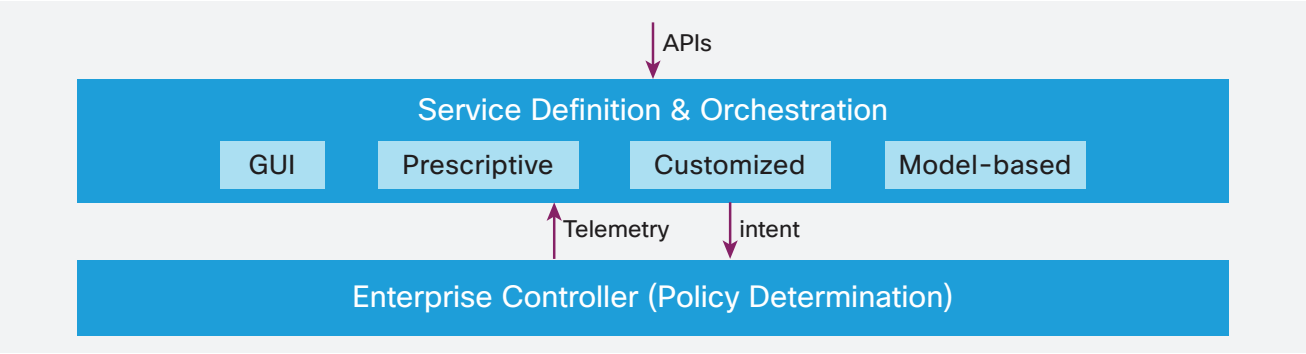

F. Analytics and Telemetry Applications

An infrastructure of analytics and telemetry is another building block in DNA. Feedback mechanisms are built into the architecture to offer continuous and relevant information about the operational state of the network. Such states can be harvested to optimize the network and security services (delivered to end users and applications). Feedback mechanisms in the DNA are also accessible to digitalized business applications, supporting the dynamic environment that these applications require, and fostering the cycle of continuous application development and improvement.

Analytics and telemetry support is offered in the following three ways:

- Data collection
- Data analysis
- Feedback and control

First, the DNA network elements are enhanced to collect data about all aspects of the network, including the state of the network elements and the traffic flows pertaining to the services offered by the network. This data collection is not just restricted to the transport network elements (routers, switches, access points, etc.). It also extends to supporting network functions such as AAA servers and virtualized or physical Layer 4–7 functions. Network elements and virtual network functions in DNA are inherently designed to collect a vast amount of data. As an example, well-established mechanisms such as SNMP and Flexible NetFlow are enhanced in a DNA infrastructure to provide statistics on user groups[10], in addition to the standard fields on source and destination IP addresses and ports; protocol types, bytes, or packets sent; or TCP flags. Furthermore, firewalls such as the Cisco Adaptive Security Appliance (ASA), authentication servers, such as the Cisco Identity Services Engine (ISE) [22] or URL filtering functions collect data at similar granularity. The continuous collection of such telemetry data always includes timestamps, allowing time-series analysis of all events.

Second, the DNA analytics engines can analyze and correlate the data collected by all network elements and functions. This is illustrated in Figure 12 by example of Lancope's Stealthwatch [23]. In this solution the network acts as a security sensor and provides network anomaly detection services to the operator. Data from various sources is correlated to make inferences about the security aspects of the network. The analytics engine within DNA offer the capability to not only analyze the events over time – they are also capable of filtering the relevant data provided under step 1 to the correct granularity or by correlating traffic flows (i.e., directionally). By looking at the events and data over time, both user- and application-specific, intelligent insights into the state of the network and the applications can be made.

Such analytics capabilities can also be deployed in a virtualized form factor within DNA, and as such, be distributed in the infrastructure where the required compute resources are available. See [10] for details on this concept.

Figure 11. Analytics & Telemetry in DNA

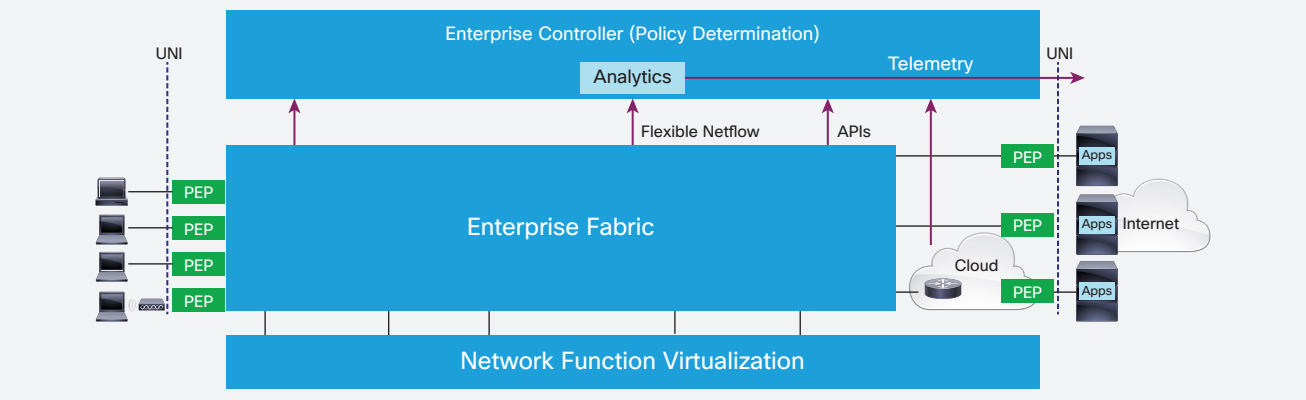

[10] For example, TrustSec® tags

Figure 12. Analytics Support in DNA

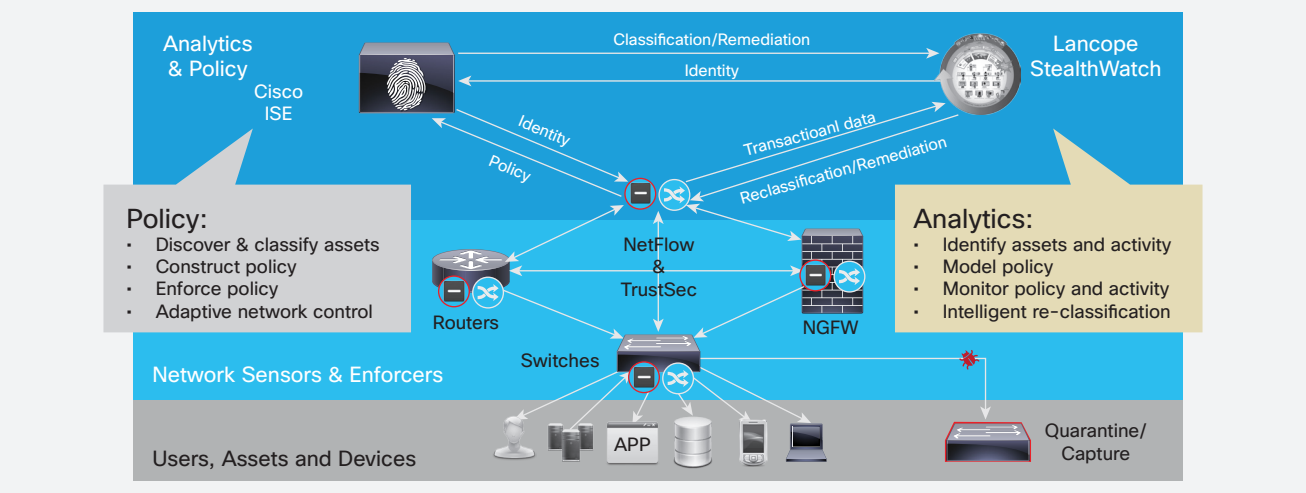

Third, the insights gained by an analytics engine can be used to drive events and actions. Network or security services can be optimized based on the analytics results, for example, by instantiating new policies for particular applications or user, or by modifying existing ones. In the example illustrated in Figure 12, the network functions work in unison with the identity engine, firewalls, and the Lancope Stealthwatch Analytics engine to deliver continuous adjustments to the security policy.

The DNA telemetry and analytics infrastructure does not just optimize available network services and operations. Analytics can also be exposed to the digitalized applications to enhance their operations or behavior. Information about user–behavior on the network, such as their communication or location patterns can enhance the applications riding on the network themselves, thus powering ever –improving and valuable client experiences.

V. Resilience and Security Considerations

As introduced in section II, any network architecture supporting business–critical digitalized applications must be resilient against hardware and software failures, as well as offer security and service integrity. The technical aspects in DNA around high availability and security are thus further elaborated upon in this section.

A. Security

1) Endpoint Security Services

Network services offered to endpoints or endpoint groups can be enhanced by security and encryption functions. First, the inherent capability in DNA to instantiate policy with user and application granularity assures fundamental security in the network. By applying a mandatory policy enforcement point at every edge into the DNA domain the network regulates which user and applications get access to the network, as illustrated in Figure 13. The communication relationships between endpoints can be governed using the segmentation offered by transport virtualization.

Second, security functions such as encryption or IPS and IDS can further be specified for a service flow as part of its service definition. The DNA orchestrator and controller components can ensure that the right security functions are applied in the right domains to instantiate security. For example, if a service IP flow traverses an untrusted domain (for instance, a service provider WAN, Internet, or a public cloud), an encryption function can be applied to the service at the (untrusted) domain edges. Similarly, anomaly detection or intrusion prevention services can be instantiated in a VNF, and the service flows requesting such functions can be chained to traverse these functions.

Figure 13. Policy Enforcement in DNA as Security Mechanism

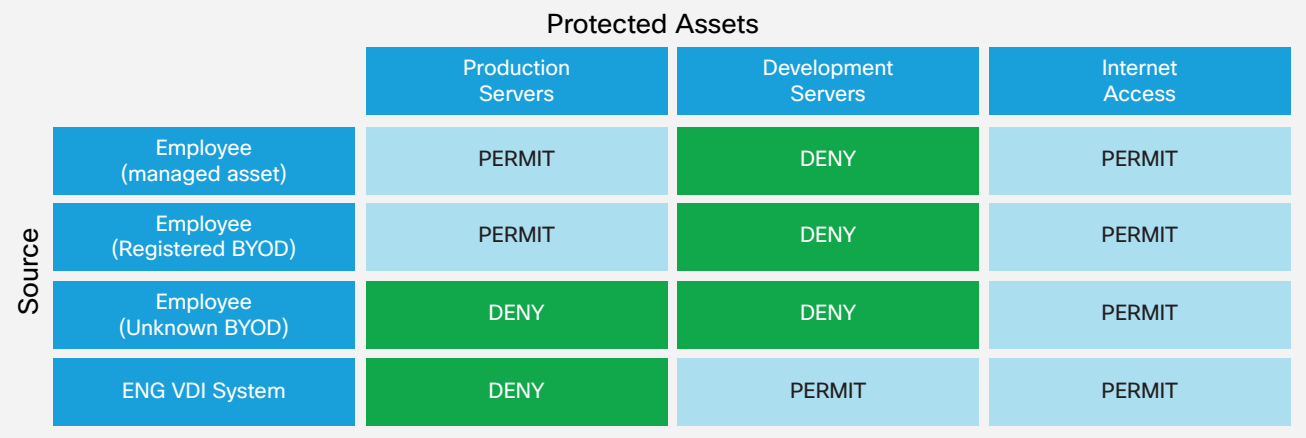

2) Security of the DNA Control Plane

Securing the control plane in DNA also relies on several mechanisms. First, access to the control-plane components is strictly authenticated. Network operators need to submit username and password credentials to any control-plane function. Groups of network operators are supported using role-based access mechanisms, allowing different classes of operators to execute only well-defined actions associated for that group[11].

Second, all communication between the control components in DNA and the network elements and VNFs are passing through a secured channel (i.e., HTTPS). API calls between components are subject to an authentication mechanism. All REST interactions between the controller and the network elements, for example, are associated with a security token generated upon successful authentication at the start of the REST session.

B. High-Availability

1) High-Availability of the Controller and Orchestration Components

The controller and orchestration components in DNA fulfill a critical function for the operation. Ensuring the continuous availability of these particular components becomes essential. Both the controller and orchestration software are supported in a cluster configuration to guard against outages of the underlying servers. Any state generated by these components is kept in persistent and distributed databases.

2) High-Availability of Network Services and Functions

Well-established resilience mechanisms for IP networks are applied to provide resilience for the service IP flows in DNA. Any transport element in the DNA fabrics can be deployed with redundant data and control planes. Link redundancy in the fabric offers protection against link failures, and can be supplemented by IP high-availability techniques such as Hot-Standby Router protocol (HSRP), IP fast convergence optimizations, or fast re-route supported by Bi-Directional Forwarding Detection (BFD).

Stateful Layer 4–7 functions offered as part of the network services are typically deployed using VNFs in DNA. High-availability mechanisms for virtualization can be used on top of the traditional stateful redundancy mechanisms for such services. As an example for the former, virtualization functions such as fault tolerance and VM high availability can be exploited. An example for the latter is the stateful inter-chassis redundancy feature for IOS® -XE and the Cisco virtual Adaptive Security Appliance (ASAv), which can be configured to continuously synchronize the firewall state between an active and a standby firewall VNF.

[11] For example, superusers may have full access to all operations, whereas other groups of operators are able to execute only a subset of the control functions.

Note that the infrastructure in the cloud is assumed to be highly available. One of the key benefits of the integration of cloud computing and cloud services is the outsourcing of network and compute infrastructure to a cloud provider. Assuring the availability of the compute and transport resources within the cloud, therefore falls into the responsibility of such providers.

Multicast is increasingly important to many business applications today. It is used for such diverse applications as multipoint video delivery, software updates and imaging, device and service discovery, and real-time financial applications. The high-availability characteristics of DNA apply to multicast flows, just as they do for unicast traffic. A DNA network provides robust protection for multicast traffic flows, with rapid recovery of end-to-end data transmission in the face of network node or link failures, As a result, it protects the critical business processes that these multicast traffic streams support.

VI. Conclusion

Enterprise networks are evolving rapidly to support the requirements arising from digitalizing business processes. The Cisco Digital Network Architecture provides the network infrastructure to support this evolution. The architecture consists of several main building blocks. **Transport fabrics** connect users, applications, and things seamlessly to greatly simplify network operations. **Virtualization** in DNA allows for the decoupling of network or transport functions from the underlying hardware elements, and offers the flexibility and speed required to instantiate services in the network. **The cloud** becomes an integral part of the Cisco DNA infrastructure. Network operators are empowered to run applications where it makes sense from a business perspective. They can take full advantage of cloud infrastructure to operate the DNA, or to offer advanced analytics services. The **DNA controller** centralizes the network control plane of the infrastructure and plays a critical role in automating its operations. It configures policies that govern network access and transport to instantiate the intent of the network services delivered to business applications. **Analytics and telemetry** offer the feedback mechanisms to support network-enabled applications, providing real-time data to application developers for continuous improvement cycles.

Based on these building blocks, Cisco Digital Network Architecture delivers a flexible and innovative environment to deliver transport, security, and digital network services. Being fundamentally a software-driven and fully open environment with APIs, programmability, and virtualization, the Digital Network Architecture platform allows network operators, partners, or network function vendors to drive required innovation in the network to keep pace with the developments in digitalizing business operations for years to come.

Appendix A: Glossary

- **API:** Application Programming Interfaces facilitate network elements or functions to be controlled by outside applications. Typically, they are a set of functions that can be called into a software program with specified parameters and formats to provide input data or receive outputs from the function. APIs can enable the openness and flexibility in DNA by allowing third-party vendors to contribute to the network operations and to increase the speed of deployment for new services.
- **Cloud:** The cloud is the compute, storage, and networking infrastructure offered by the aggregate set of cloud providers as services to run enterprise applications. The cloud is made up of all the data centers by cloud providers to host applications. The advantage of cloud computing is that these services can be acquired instantaneously upon demand, so that enterprises do not need to invest in data center infrastructure (both from a CapEx and OpEx perspective). The cloud provider services the resources to ensure that the capacity meets demand, that cloud services are offered with redundancy and security. Cloud computing offers a different consumption model for compute (subscription pricing).
- **Controller:** A controller is a component in the network that manipulates the network elements in the DNA according to the policies (to instantiate the services). The controller has the full holistic current state of the network, that is, it keeps an abstracted network state. The controller interfaces with the orchestration and policy layers via a northbound interface. Controllers instantiate configuration entries into the network to create transport paths or to implement the services that are offered to the endpoints using a southbound interface.In DNA, a single controller may span multiple domains (cloud, WAN, campus, data center). Alternatively, multiple controllers may collaborate, each fulfilling a domain-specific role. The functions executed by a controller may even be broken out

into multiple sub-controllers, each specializing in a particular control task, such as monitoring and manipulating the QoS state of the network.

- **Devices:** Devices are physical systems capable of running applications and sending or receiving network traffic.
- **Digitalization:** Digitalization is the process of bringing digital technologies into all aspects of an enterprise's business, and offering digitalized services to the business's consumers. Digitalization also refers to the act of converting business processes into digital formats and take advantage of automated algorithms and processes to optimize and simplify either internal operations or interactions with consumers. See also [24].
- **DPI:** Deep Packet Inspection is a technique to determine the type of the application carried in an IP flow by inspecting the payload, often over multiple packets, and making inferences based on the payload. This contrasts with application characterization based on TCP port numbers or payload type fields, which are often insufficient to adequately characterize an application.
- **Endpoint:** Endpoints are consumers of network services defined at the **application** level. These can run on traditional hosts like PCs, notebooks, and IP phones, but also increasingly on IP-enabled devices that drive digitalized business processes like robots, point-of -sale (POS), scanners, inventory tracking devices, vehicles, and more. Applications running on servers in a data center are considered consumers of network services and are therefore endpoints in the architecture.
- **Endpoint Groups:** These are the categorization of multiple endpoints into one logical entity receiving a service from the network. An endpoint group could bundle all endpoints from a particular end device, for example. Alternatively, a group could bundle all endpoints from a particular end user where the endpoints are on multiple physical devices. The term 'endpoint' is used loosely in this paper to mean both individual endpoints as well as endpoint groups.
- **Fabric:** A fabric is a collection of network elements that offers communication paths between its outbound-facing ports, providing any-to-any connectivity by use of overlay tunnels. A fabric is governed by a controller to simplify the operation. A service can be instantiated by applying the right policy enforcements characterizing the service instances, and making use of the any-to-any connectivity offered by the fabric to reach the services remote ends.
- **Intent:** An intent is the association of a business process with a service. The purpose or intent of a service delivered by the network is to fulfill the business processes relevant to the enterprise. Services, therefore, implement the business intent of the enterprise, and a policy is a specification on how the service is implemented and enforced by the network.
- **Internet of Things (IoT):** The IoT is the extension of the Internet to reach not only traditional compute resources such as PCs, notebooks, and servers, but also **any** device. This implies that IoT devices become networked, i.e., they are extended to have compute resources on which a networking software stack can run. It typically also implies that IoT-enabled devices run digitalized business applications to benefit from the Internet connectivity. Examples of IoT devices would be refrigerators, cameras, vehicles, parking meters, production robots, elevators, RFID-equipped hardware, sensors and so on. The advantage of networked objects equipped with software stacks is that these can be controlled remotely, so they help to digitalize business processes and services.
- **Intrusion prevention systems and intrusion detection systems (IPS and IDS):** IPS and IDS are a set of network functions that monitor the network on a continuous basis. The systems look for behavioral anomalies or malicious activities and to take appropriate actions (such as blocking, or reporting) when detected. IPS and IDS often rely on baseline behavior to be characterized (signatures), against which the anomalies are then defined. Statistical methods may also be employed.
- **Open networking:** Open networking allows the operation of the network architecture to be influenced by third-party software or hardware vendors. It relies on standards-based and published APIs to facilitate such integration. Open networking empowers the community of developers outside Cisco to contribute to network operations and functionality. Open networking increases the innovation speed and introduces more flexibility into the network.

- **Orchestrator:** This component in the network allows the specification of **services**, and the initialization of the instantiation or ongoing modification of those services in a controlled manner. The orchestrator in DNA is focuses on service definitions from a network consumer point of view, thus abstracting any low-level details of how, and sometimes where, those services would be configured. The orchestrator determines the service intent and communicates this to the controller. The controller in turn manipulates the network elements to provide the transport between the relevant PEPs, to instantiate the policies (transport, security, etc.) and to ensure services are monitored on an ongoing basis.
- **Overlay network:** This network is based on network tunneling that sits on top of an underlay network. Tunneling techniques are used to decouple a network service from the underlying transport infrastructure. The state of the service is only kept at the edge of the network. For all network elements that make up the underlay network, the tunneled service traffic appears simply as the encapsulated tunnel traffic. Overlay networks are characterized in [25] by:
  - Segregating traffic between users
  - Supporting different address spaces
  - Supporting dynamic device or VM placement (independent of the underlay topology and addressing)
  - Supporting large scale

A network is fully virtualized if virtual network functions, such as routers or firewalls, are connected with each other using VLANs or VRFs.

- **Policy:** Policies are applied to **each** service to govern how it is treated by the network. A policy is an attribute of each service, and the policies are instantiated in one or more places in the network. Policies fall into the following sub-categories:
  - Access policies determine which endpoint is allowed access to consume the service and to regulate the communication relationships between endpoints. An access policy regulates the communication relationships for an endpoint or between endpoints.
  - Transport policies are a specification on how an IP flow consuming the service should be treated by the network. Examples include encryption policies to traverse untrusted parts of the network, path optimization policies (such as traffic engineering), or segmentation policies.
- **Policy Enforcement Point (PEP):** PEP refers to the functions in the network that instantiate and implement the policy associated with a service. Each service must have an access PEP, even if this specifies "transport all". The PEP may also specify additional transport policies if they have been defined for a service. PEPs may be instantiated between domains in the network (for example, to apply encryption or decryption as the service flows traverse an untrusted part of the network). A PEP may therefore, be instantiated by the network operator as the service traverses the network-to-network interface.
- **Software-Defined-WAN (SD-WAN):** SD-WAN is the use of software-defined networking to control and instantiate an enterprise WAN. This is particularly attractive for a geographically distributed branch network, where SD-WAN can help to determine the most efficient transport between locations. Virtual overlay networks, by using tunnels, centralized intelligence using controllers, and APIs to automate the network configuration, are aspects of SD-WAN. See [26] for a more elaborate description.
- **Service:** Service refers to the transport of IP traffic between two or more endpoints or endpoint groups provided by the DNA network. The service is offered by the network to the endpoints or endpoint groups. That is, a service is delivered across a UNI. The endpoints are the consumers of the service. The network realizes the service by instantiating the right transport connectivity to connect the endpoints associated with the service to each other. The service is **always** associated with a PEP, and the service may manipulate the IP flows (for example, apply NAT, DPI, or WAN optimization) anywhere along the path according to the service description. Services are offered to implement the enterprise's business intent. The term "services" in the DNA implies a customer-facing entity, as opposed to a resource-or network-facing construct. In DNA, resource-facing functions are not referred to as "service". Instead, they are called "functions" to remove confusion. Examples of functions, elsewhere referred to as "network services", are NAT, firewall, and domain name server (DNS).

- **Service-level agreement (SLA):** The network offers SLAs to the traffic flows it transports to characterize upper bounds on transport delay, jitter, or packet loss. Other measures may also be incorporated in SLAs. In DNA, a service is always associated with an SLA. Telemetry: In the context of DNA, telemetry is the collection of measurements or other data about the state of the network, i.e. the network elements and functions that define its operation. Data collection and measurements can also be applied at the user and application level. The data and measurements are transmitted to monitoring systems (analytics) for processing. These measurements are done remotely from the point of view of the analytics engine without processing, hence the use of the term "telemetry" in a networking context.
- **Traffic engineering:** Traffic engineering is the determination of an optimal path in the network and its instantiation using protocol-based mechanisms like DiffServ-TE or Performance Routing.
- **Underlay network:** An underlay network is a transport network that can enable IP connectivity between its constituent elements. Underlay networks are realized using traditional routing techniques.
- **User-network Interface (UNI):** A UNI is the demarcation point between the enterprise network domain and the applications and endpoints that consume network services. In DNA, all network elements and functions between two UNIs are under the responsibility of the network operator. The UNI defines the services that elements outside of the network may consume and how this consumption is realized. A PEP is associated with all endpoints crossing the UNI in DNA.
- **Virtualization:** In the DNA, virtualization refers to the concept of creating virtual networks that are de-coupled from the underlying hardware systems. Virtualization can be categorized into network function and transport virtualization:
  - Network function: deployment of network functions in software on top of standard x86-based hardware hosts, either using virtual machines or network function containers.
  - Transport: deployment of VLANs or virtual routing concepts to create a logical Layer 2 or 3 forwarding constructs.
- **VRF:** Virtual routing and forwarding is the logical separation of routing tables to segment the traffic forwarding at Layer 3. VRF is an important construct in DNA to realize virtual private networks at Layer 3.
- **WAAS:** Wide area application services. The Cisco implementation for WAN optimization.
- **WAN optimization:** WAN optimization is the deployment of techniques to improve the efficiently of traffic flows across the wide-area network. Examples of techniques include:
  - Compression of IP payloads
  - TCP/IP flow control optimizations
  - Data redundancy elimination and caching

References

1. IMS Research. The World Market for Internet Connected Devices, 2012 Edition. [Online]. Available: http://www.iot- now.com/wp-content/uploads/2012/07/Internet-Connected-Devices-World-2012-Proposal.pdf
2. S. Brennen and D. Kreiss. Digitalization and Digitization. [Online]. Available: http://culturedigitally.org/2014/09/digitalization- and- digitization/
3. Cisco. Visual Networking Index. [Online]. Available: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html
4. The Zettabyte Era – Trends and Analysis. [Online]. Available: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/VNI HyperconnectivityWP.html
5. Zacks, D. and Montañez, M. The Cisco Digital Network Architecture Vision – Details on Fabrics. Available upon request. Cisco.
6. Cisco Unified Access Technology Overview: Converged Access. [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-3850-series-switches/white paper c11-726107.pdf
7. The Cisco Flow Processor: Cisco's Next Generation Network Processor Solution Overview. [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/routers/asr-1000-series-aggregation-services-routers/solution overview c22- 448936.html
8. D. Zacks. (2015) Cisco Enterprise Silicon – Delivering Innovation for Advanced Routing and Switching. [Online]. Available: https://www.ciscolive.com/online/connect/sessionDetail.ww? SESSION ID=81811&backBtn=true
9. (2015) Case Study in Network Infrastructure Virtualization. [Online]. Available: https://www.ciscolive.com/online/connect/sessionDetail.ww?SESSION ID=81846&backBtn=true
10. Falkner, M. and Arena, S., "The Cisco Digital Network Architecture Vision – Details on Virtualization". Available upon request.
11. J. Halpern and C. Pignataro. Service Function Chaining (SFC) Architecture. [Online]. Available: https://datatracker.ietf.org/doc/rfc7665/
12. Falkner, M. and Montañez, M. The Cisco Digital Network Architecture Vision – Details on Cloud. Available upon request.
13. Openstack. Group Based Policy. [Online]. Available: https://wiki. openstack.org/wiki/GroupBasedPolicy
14. O. Daylight. Group Based Policy User Guide. [Online]. Available: https://wiki.opendaylight.org/images/9/90/Gbp-lithium-user-guide.pdf
15. Cisco. Cisco AVC Solution Guide for IOS XE. [On- line]. Available: http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/avc/configuration/xe-3s/asr1000/avc- xe-3s-asr1000-book.html
16. W. Riedel. BRKCRS-2321 DNS-AS: Done with the SDN and Tired of Dealing with Snowflake Network Complexity? Change the Game with a simple TXT string. [Online]. Available: https://www.ciscolivelatam.com/connect/sessionDetail.ww?SESSION ID=3157
17. M. Bjorklund. RFC 6020, YANG – A data Modeling Language for the Network Configuration Protocol (NETCONF). [Online]. Available: https://tools.ietf.org/html/rfc6020
18. R. Fielding. Architectural Styles and the Design of Network-based Software Architectures. [Online]. Available: http://www.ics.uci.edu/%7Efielding/pubs/dissertation/top.htm
19. A. Bierman, M. Bjorklund, and K. Watsen. RESTCONF Protocol. [On- line]. Available: https://tools.ietf.org/html/draft-ietf-netconf-restconf-05

20. Cisco. Cisco Application Policy Infrastructure Controller Enterprise Module. [Online]. Available: http://www.cisco.com/c/en/us/products/cloud-systems-management/application-policy-infrastructure-controller-enterprise-module/index.html
21. Riedel, W. and Zacks, D. and Falkner, M. The Cisco Digital Network Architecture Vision - Details on Controller and Orchestration. Available upon request. Cisco.
22. Cisco Identity Services Engine. [Online]. Available: http://www. cisco.com/c/en/us/products/security/identity- services- engine/index.html
23. D. Miller and M. Robertson. Cisco Live BRKSEC-2026: Network as a Sensor and Enforcer. [Online]. Available: https://www.ciscolive.com/online/connect/sessionDetail.ww?SESSION ID=83690&tclass=popup
24. BusinessDictionary.com. Digitalization Definition. [Online]. Available: http://www.businessdictionary.com/definition/digitalization.html
25. S. Salam. Overlays, Underlays and the New World Order. [Online]. Available: http://blogs.cisco.com/getyourbuildon/overlays-underlays-and-the-new-world-order
26. TechTarget. SD-WAN (software-defined WAN). [Online]. Available: http://searchsdn.techtarget.com/definition/SD-WAN-software-defined-WAN

C11-736842-00 03/16