Maximizing Agility at the Network Edge

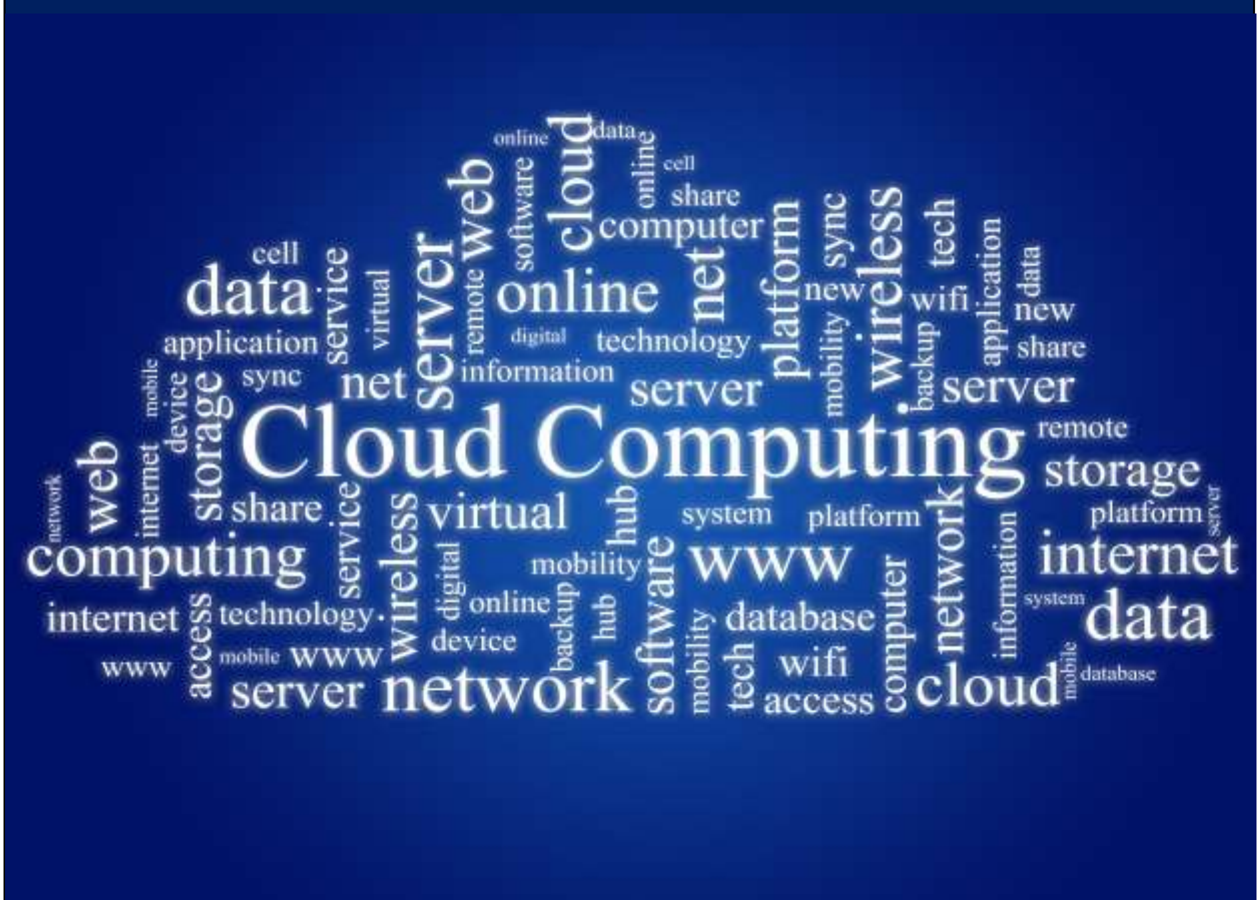

Table of Contents

This paper was commissioned by ENEA AB and incorporates original analysis and insight from ACG Research.

CREATING AGILE NETWORKS: BUSINESS NETWORKING AND THE ADVANCE OF THE CLOUDS

The industry's adoption of cloud-native designs is on. Cloud-native designs increase value by delivering applications rapidly and lowering the cost of delivery using open technologies at every layer of operation, from hardware to operating system and application software. In cloud-native designs the path to a more fully connected world is **open**.

Delivering applications in a cloud-native world demands an agility in networking that has not been available in the past. This is true in an expanding landscape of use cases, extending beyond conventional virtual private networks (VPNs) (which remain important), and incorporating new architectures across multiple clouds, and running applications in the Internet of Things in parallel with applications in the Internet for people.

These developments have stimulated efforts to bring new and forward-looking virtualization and software-driven functionality to networks to align them with the applications they support.

NFV 1.0

Early work on virtualizing network functions (VNFs) focused on using the framework that was most commercially accessible at the time. This involved running virtualized network functions as software modules in guest OS environments on virtual machines (VM) in x86 servers. During this first phase of network function virtualization (NFV) it was critical to gain practical experience deploying VNFs in such resource pools, as well as in how they would function in operators' networks and services.

As experience with design for NFV has grown, stakeholders have concluded they need to incorporate additional approaches in parallel with x86 and virtual machine-based deployments to tackle new application and economic requirements and still fit in with the framework of model-driven, open architecture solutions for multiple applications and services.

Toward this end, one observation has been that x86 processors are not optimal for every user's application. They provide a range of functionality and performance beyond the needs of certain types of applications and thus require a footprint that is larger than necessary in those cases. For example, industrial control and environmental sensing are often better served using smaller footprint advanced RISC machine (ARM) processors because of their smaller size and typically lower power consumption.

Implementers in the past two to three years have incorporated ARM processors into computing platforms and demonstrated that ARM-based configurations can fit nicely into an NFVI design using cloud management systems such as OpenStack, managing workloads in both x86 and ARM implementations. Being able to choose the right device for the right application, **either** x86 **or** ARM-based solutions (or **both),** has been identified as an important method for service providers and operators to reduce cost through lower capex and power consumption.

Similarly, with respect to virtual computing in NFV, experience has shown that virtualization is not always optimized when done using VMs. Virtual machines carry the burden of running a full copy of a guest OS in each VM to support a workload. Although VMs provide good isolation and, frequently, an efficient path for transitioning implementations from custom processor platforms to general-purpose servers, they also introduce overhead/cost and performance penalties people would like to avoid if a more efficient alternative is available. To this end, interest in using containers as a computing environment for VNFs in various use cases has emerged. Containers require less overhead per workload and are generally more performant than VMs and able to run multiple workloads on a given underlying server.

NFV Next

Taking these developments into consideration the number of communities engaged in delivering NFV solutions has expanded, and the number of opportunities available for stakeholders to pursue has expanded in turn. For example, the OpenStack community took on the task of supporting containers in parallel with VMs in a cloud infrastructure deployment. The OPNFV community embraced the integration of NFV solutions incorporating ARM servers in the family of open source implementations it is validating as usable in NFV use cases.

Enter Enea

Enea AB is an established, global supplier of software for mobile communications and other performance-critical applications (including automotive and aeronautical applications). The company has a heritage of contributing innovations to open source software communities such as Linux OS and of delivering efficient, highly scalable solutions in a variety of platform environments for production network services.

As the landscape of NFV 1.0 has evolved into the broader frameworks of NFV Next, and considering Enea's expertise delivering solutions for diverse network applications, the company has embarked on a program to expand the capabilities of NFV solutions and create opportunities to deploy network services in richer configurations and designs than have been available to date.

We describe the innovations Enea has made first at the level of overall significance in the market and then with a focus on the enhancements it is bringing to business networking services in its virtual CPE solutions.

Optimized Open Source Software Innovations

In its work developing solutions for telecommunications and other applications Enea has realized several enhancements for open source Linux deployments that also bring high value to NFV use cases. These include achieving fast boot times (of three seconds and under) for Linux configurations; running efficient, high throughput data planes in distributed network services capable of realizing physical NIC in to physical NIC out latencies of 10 to 15 microseconds for many VNF implementations; and realizing small footprint software for reliable operation in demanding, cost-constrained environments.

Complex System Integration and Validation

In delivering its OS and other software solutions for operators Enea has also developed extensive experience selecting, integrating and validating open source components appropriate for operator use cases alongside its own applications into functioning production networks. For NFV and especially in supplying virtual system infrastructures for a variety of use cases, this competence is especially useful and helps in all solution delivery phases: development, test/validation, deployment, operation and enhancement.

Understanding that the vision of NFV is to enable openly mixable combinations of functionality at each layer of the NFV architecture, Enea decided to bring its integration skills to the global NFV solution delivery community by offering to participate as one of the approved integration and test lab managers in OPNFV's Pharos network of interoperability labs.

In this process Enea has become an important participant in the OPNFV community. The purpose of OPNFV is to integrate open source distributions from the principal open source software communities producing software relevant to NFV in top-priority use cases. In its projects and labs, OPNFV produces tested combinations of software that can be further packaged by integrators for deployment in service providers' networks. Its work addresses all the critical NFV infrastructure components, as well as procedural and technical requirements involved in deploying NFV infrastructures that work, spanning virtual compute (Linux, KVM, Open Container, Cloud Native Computing), networking (Open vSwitch, FD.io, Open Daylight, ONOS), storage (Ceph) and orchestration for NFV (OpenStack, Open Source MANO, ONAP).

Leveraging its expertise of integrating software from multiple open source communities, Enea brought a distinct set of capabilities into the OPNFV process and opened the only lab in OPNFV's network that is available as a service to validate NFV software in ARM server configurations. In parallel Enea runs a heterogeneous lab environment for its own overall solution developments that is inclusive of both x86 and ARM server PODs so its solutions are being validated for a full mix of deployment scenarios. These commitments accelerate not only Enea's own solution developments, but those of its partners and of the NFV ecosystem.

Managing Distributed Networks Efficiently and at Scale

In addition to enhancing real-time software and extending the capabilities of the OPNFV reference designs and solution validation infrastructures, Enea has developed management software useful in deploying and operating highly distributed networks efficiently and at scale. Through its experience implementing distributed, real-time solutions with operators, Enea understands that managing distributed elements requires not only cost-effective elements themselves, but also an open, scalable and versatile approach to managing the elements remotely. Such management software is not an inherent part of the baseline NFV reference architecture, though it is essential for success in running large operator services.

Considering these requirements Enea has incorporated a set of open management functions in its NFV Access solution called Edge Link that enable it to be supported in a real-time, model-driven manner using a variety of mechanisms, including communications based on Netconf, an open, scalable and widely deployed protocol used in managing operators' networks. This allows operators to integrate VNFs and the platforms they are running on with management solutions flexibly, in an open, standards-based manner using a widely supported protocol and element information models developed in the closely correlated modeling language YANG. Of interest also in some cases is that remote NFV Access components can be deployed as OpenStack compute nodes from a centralized OpenStack implementation (Enea's or another's distribution that the operator may be using), enhancing both efficiency and scale. Figure 1 illustrates the components of Enea's Edge Link functionality within a whole NFV Access vCPE configuration. Note the versatility in communication mechanisms between the remote vCPE node and the management applications with which it needs to communicate, supporting both day-to-day and life-cycle management operations, as well as functionality suitable to the context (whether VM, container or underlying platform).

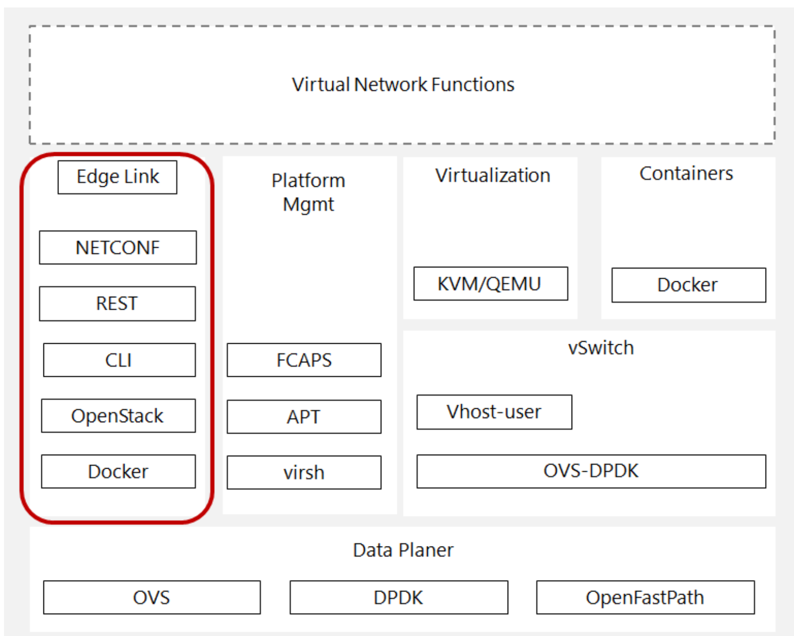

Figure 1. Edge Link Functions for Open, Standards-Based Management of NFV-Based Services at Scale

Sophisticated Traffic Management and Inspection for NFV

In addition to embracing open architectures and enabling their management at scale, Enea has invested significantly in expanding the traffic inspection and forwarding intelligence it places in its NFV solutions. This functionality is especially important in a core component of the NFV reference architecture and value proposition: creating service chain topologies among the variety of network functions employed in a given service. In service chaining the forwarding path between functions, such as virtual and physical routers, firewalls, malware inspectors and content caches, is uniquely specified based on user, service or application requirements. It is abstracted above physical network routing and switching domains and as such works as an overlay. The path for one user or application may be different, based on subscription parameters and policies, from another's. Having deep inspection to accomplish fine-grained control at high traffic rates helps the operator be more efficient and add value to services.

Recognizing this Enea has incorporated deep packet inspection and classification into its NFV solutions based on the ixEngine traffic analysis software it acquired when it bought the leading-edge traffic analysis firm Qosmos in 2016. By integrating ixEngine into its NFV solutions Enea optimizes the flow of traffic within the service chains that have been enabled and enhance the quality and differentiation of services its operator customers can offer by an order of magnitude beyond other available offerings. The operation of ixEngine in an NFV service deployment is illustrated in Figure 2. In it we can see customers' traffic belonging to an operator's deployment entering a node in the NFV deployment from the left, being evaluated by policy rules within the ixEngine module, and tagged for forwarding according the rules of the identified service chain to the sequence of VNFs that is required.

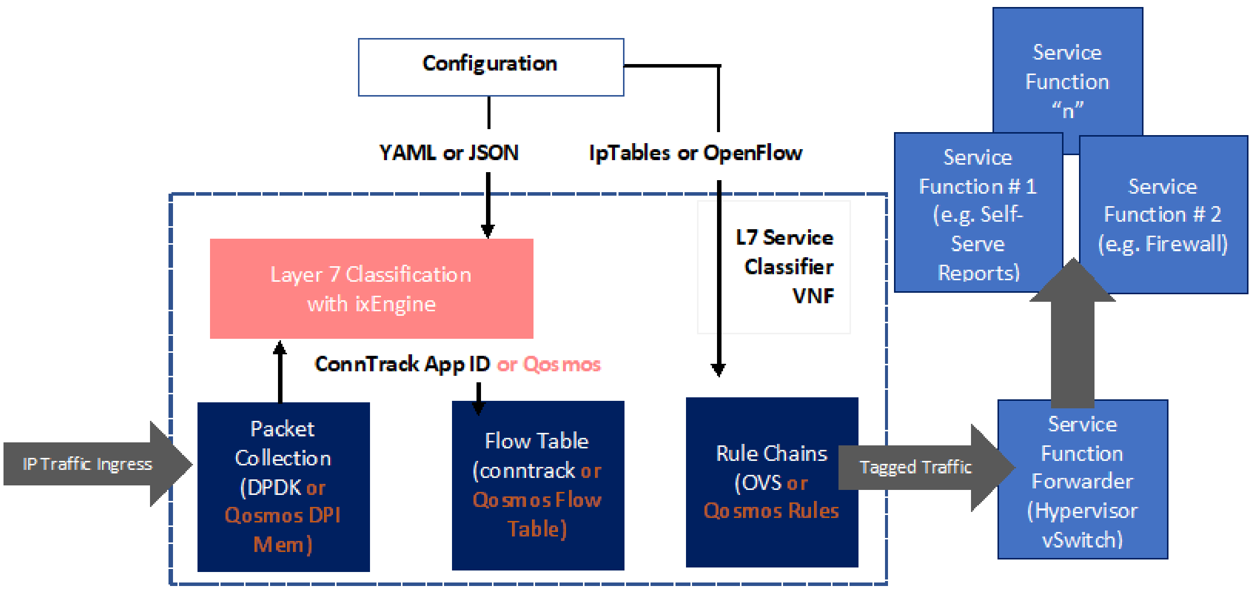

Figure 2. L7 Based Service Chaining Based on Enea Qosmos ixEngine

Having taken these visionary and forward-looking steps and anticipating their value in customer use cases, Enea began implementing deliverable solution offerings based on its open, cloud-native framework for customers' environments where the offerings would create distinct value.

Dynamic Virtual Networking for Businesses

One of the most important use cases in the early phases of NFV has been creating virtualized networking services for operators' business customers. Historically, business customers' networks have used individual, physically separate elements at a customer's premise for each network function being supplied. Elements have included purpose-built IP routers, security appliances, and application gateways. Deploying these individually and separately has proven costly and has inhibited the pace of innovation for both operators and solution suppliers. Overcoming these constraints has provided the stimulus for innovations in NFV to create new solutions using virtualized CPE (vCPE).

As with much of the initial phase of NFV, vCPE solutions were first developed using x86 platforms and servers for functions to run in VMs. Solution scales ranged from individual x86 units at a customer's premise running a handful of VNFs to significantly larger pools of servers running in operators' facilities and supporting a larger number of VNFs for the customers connecting to a given operator's site (PoP, CO, or data center).

Although this phase has proven valuable in gaining experience with virtualized infrastructure deployments, stakeholders have recognized the initial model supporting only x86 CPU and VM based designs is not sufficient for addressing the variety of applications customers want to run. New developments we discussed previously came into focus to enhance vCPE. Adding these capabilities expands the number of services operators can deploy and increases their differentiation. They also provide an ideal point of entry for Enea to the market for open, dynamic networking for businesses leveraging cloud-native designs in vCPE.

Architectural Innovations and Enea's vCPE

Enea embraced an open architecture perspective in developing its vCPE solutions. It has focused on the founding vision for NFV of supporting cloud-native, virtualized functions as software running in an open architecture framework. Functioning in such a framework increases the amount of **choice** operators have in creating services, as well as contributing to their efficiency and differentiation. Delivering a solution using these principles requires great diligence maintaining points of integration on many dimensions of an offering: **southbound** for choice in underlying hardware, **northbound** for integration with a variety of orchestration and management applications, **east-west** for integration with a variety of network environments, and **internally** to remain open for innovation in how components work together. This is a fundamental attribute of the NFV reference architecture, though it is difficult to realize consistently and at scale for both compact deployments at customers' premises, as well as in larger pools of virtual resources in operators' facilities enabling the same offered services.

As a result, Enea delivers its solutions for vCPE and business networking services as open software offerings **agnostic to the underlying hardware** and **agnostic to the management framework an operator chooses to employ**. Services can be enabled in an x86 or an ARM-based hardware infrastructure (or in both). They can run in virtual machines or in containers (or in both). And they can be managed by the heterogeneous mix of management applications operators employ, including FCAPS element managers, OpenStack and Kubernetes for VM and container support, and a range of orchestration and OSS/BSS applications using model-driven operations and REST-based APIs.

A landscape in the market emerges, then, in which solutions for delivering business network services in a virtualized architectural model differ as to the amount of flexibility and choice they offer across a range of application environments. In Figure 3, we see that the amount of flexibility and choice solutions can increase with the number of infrastructure technology environments they support, as well as by the degree to which they are deployable as pure software loads into that variety of environments versus being constrained by tightly integrated bundles of hardware and software that must be employed for the solution to be used.

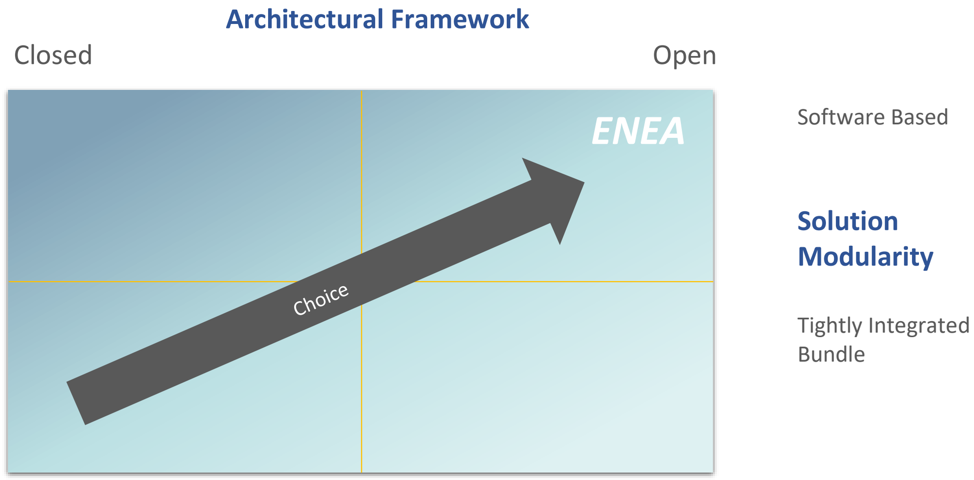

Figure 3. Architectural Openness and Degree of Choice Available in Deployment Models for Solutions

Enea has applied these principles and enabled agility and efficiency at scale via its NFV Access offering for customers' premises deployments and the companion NFV Core solution for deployment in operators' facilities. One result of applying these principles, as we see in the diagrams in Figure 4, is that an operator can deploy each offering by itself, if that is the requirement, or deploy in tandem. When deployed in tandem the operator can view them as an agile service delivery continuum. In true cloud native mode, that allows VNFs and related service workloads to be run in the location that best meets the requirements of the service.

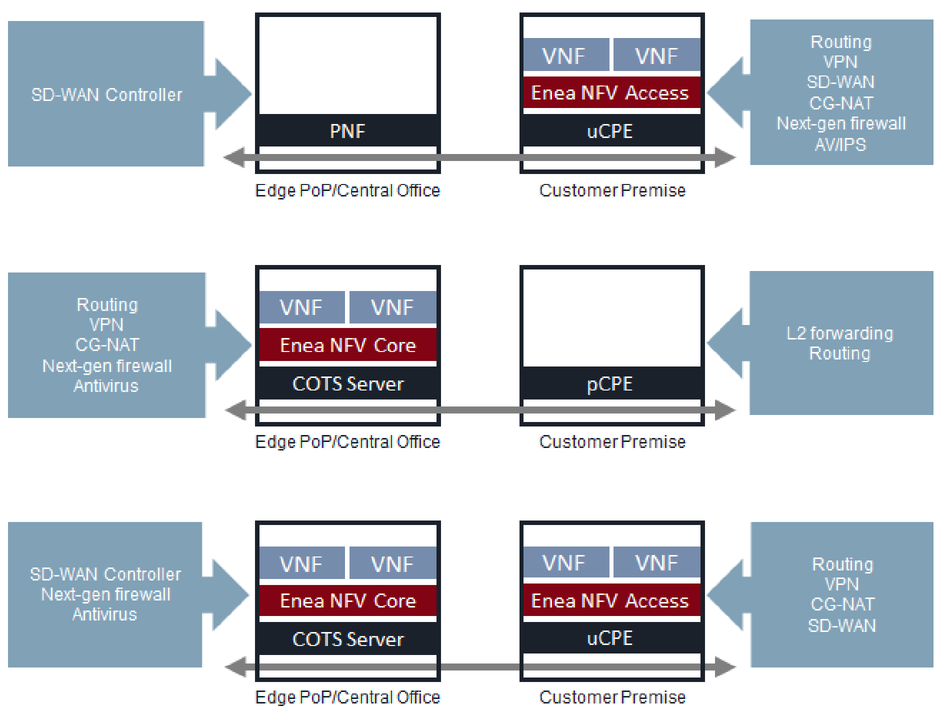

Figure 4. Flexibility in Deployment Options with NFV Access and Core

When deployed on their own, either NFV Access or NFV Core can work with other elements of the operator's deployment at any layer to enable a viable service. When deployed together, they combine to maximize both efficiency and choice for the widest range of environments.

Open Architecture, Efficiency and Scale in vCPE

Let's have a look at specific attributes of Enea's implementations that give it a position on the leading edge of offerings to evaluate in business networking services. Each draws on Enea's heritage in open, scalable, efficient and performant solutions in operators' networks and applies it to a portfolio of open, flexible and efficient solutions for vCPE. We look in turn at the complementary solutions, first for use on a business customer's premises with NFV Access and next in an operator's cloud deployments for business service offerings with NFV Core.

Efficiency, Openness and Flexibility at the Customer's Premise: NFV Access

NFV Access is a creative, insightful implementation of the NFV reference architecture built for widescale, cost-effective deployment in diverse services and applications. It is based on a **multi-architecture** design, allowing operators to decide whether x86, ARM or a mixed set of platforms is appropriate. The multi-architecture principle is also used in virtualization, allowing either virtual machines or containers (or both) to be used for workloads. When considering that business services are expanding into IoT in parallel with enabling XaaS in multi-cloud deployments, supporting a mix of processing options in an open, extensible design brings a new and larger value proposition into play for both operators and their customers.

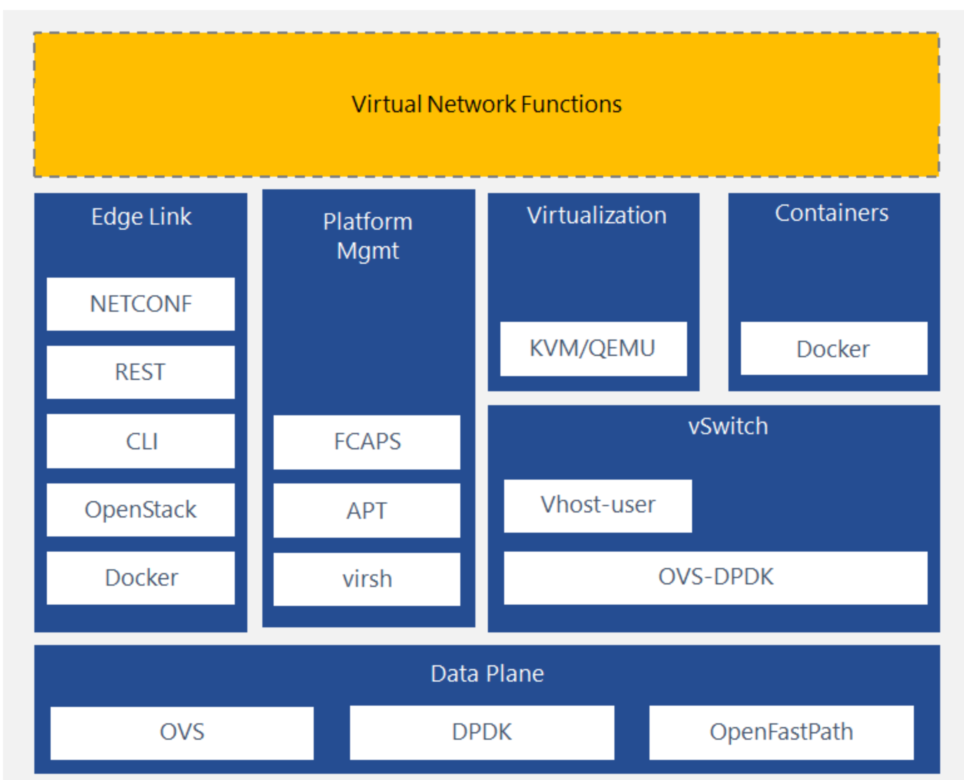

Figure 5. Open Architecture in the NFV Access Design

At the same time, NFV Access is built for management and operations at scale by incorporating the Edge Link remote management options we previously described. Via its Netconf functionality Edge Link helps operators integrate an NFV Access deployment into a range of FCAPS and element management applications. In parallel, if relevant, management of virtual resources in the NFV Access node can be performed by centralized OpenStack or container management platforms, reducing the cost and footprint of software needing to run at the customer's site.

NFV Access also employs the Enea Qosmos ixEngine traffic inspection and classification software to help deliver sophisticated value-adding services. In addition to enabling innovative services, ixEngine extends the operator's visibility into deployments, enhancing both debugging and analytics and enabling more insightful, responsive operations.

Combining this set of innovations with Enea's experience working with a variety of hardware and open source software ecosystems, its capabilities in system testing and integration, and its experience delivering production deployments in large operators' services, NFV Access for vCPE deployments is well-positioned to address the requirements of business networking services and applications moving forward.

Openness, Efficiency and Scale in the Operator's Cloud using NFV Core

In parallel with NFV Access, Enea has developed a complementary solution called NFV Core for use in service providers' own facilities. NFV Core enables cloud-native, flexible and scalable deployment of virtual network functions in a centralized fashion, which may suit the operator's goals for a given customer segment or set of network functions. NFV Core can be deployed on its own in support of other suppliers' vCPE or it can be employed in combination with NFV Access at the operator's choice; both options are feasible.

NFV Core incorporates all the multi-architecture and traffic inspection versatility present in NFV Access, affording operators the same benefits using either solution on its own or using them together. As a broader implementation of a cloud-native infrastructure, supporting a larger number of workloads and customers in its installations, NFV Core includes additional functionality to provide agility and scale in cloud-native environments. Toward this goal NFV Core benefits from Enea's participation in OPNFV and its ability to test full NFV solution environments in its Pharos labs.

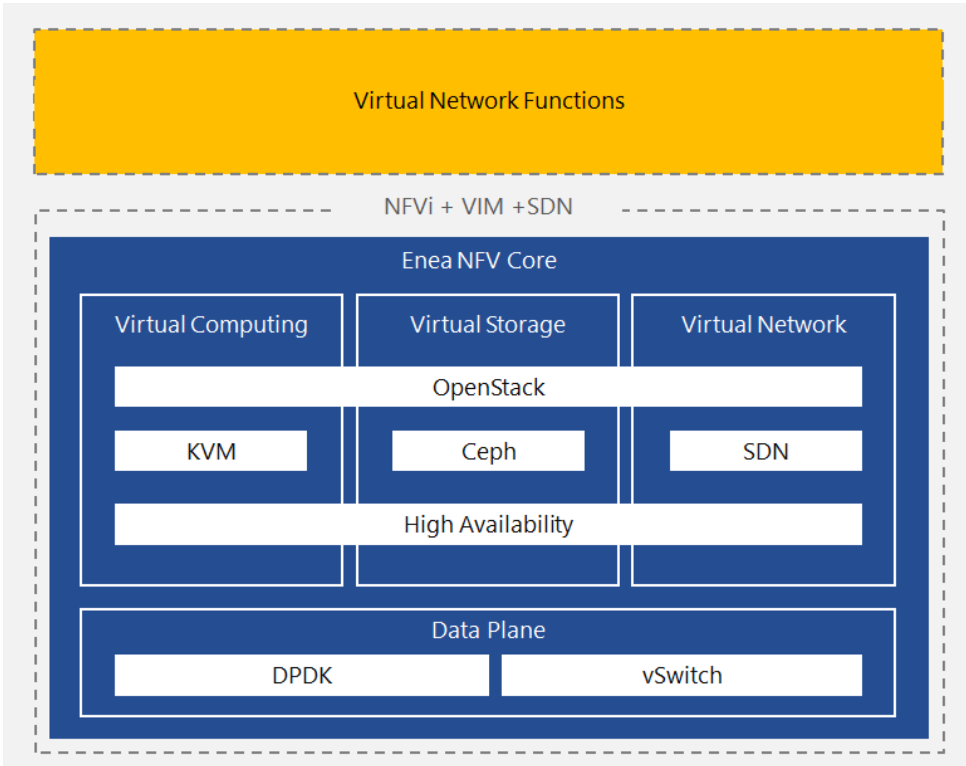

Figure 6. Agility and Scale in the NFV Core Design

As an implementation of complete OPNFV reference configurations, NFV Core is flexible in its support of a variety of approaches to enabling SDN in its NFVI to support service chaining as well as integration into operators' BGP VPNs. It also incorporates accelerated data plane functionality via open vswitch performance enhancements and supports DPDK and SR-IOV as appropriate to the VNFs being run. This is especially useful in larger system deployments, as in operators' data centers and PoPs. Further benefiting from OPNFV's emphasis on deployability and operations, NFV Core incorporates installer automation, high availability, fault diagnosis and management software to improve operations efficiency and responsiveness. The NFV Core high-availability solution is an open source implementation of the OPNFV Doctor framework that has been extended by Enea to meet the more extensive high- availability requirements of Tier 1 operators. To accelerate deploying VNFs and services, NFV Core incorporates OpenStack Tacker VNF deployment software as well as OpenStack HEAT templates for integration with orchestration systems such as Cloudify in a northbound direction.

By incorporating these additional full-cloud capabilities in its design, NFV Core serves as an efficient, versatile solution for delivering cloud-native business networking services at scale. Whether it is used in conjunction with NFV Access or on its own, it delivers the versatility and openness needed by operators in business networking services moving forward.

CONCLUSION

Businesses' ongoing pursuit of optimizations in their operations is continually introducing new applications and use cases for their ICT. New applications consistently stimulate wider use of cloud-based services from a range of SaaS and XaaS providers. The expanding number of enhancements in the intelligence of things generates the growing need to connect smarter business infrastructures to the Internet of Things. This combination of developments stimulates a parallel requirement for networks to be increasingly versatile, open, and responsive to their needs. This is especially in focus at the edge of providers' networks where the connection of businesses' users and infrastructures to diverse application services occurs.

Initial work bringing flexibility into customers' networks has occurred over the past several years in the first phase of NFV development. NFV 1.0 began to transition separate, purpose-built appliances into software modules (virtual network functions) to run in general-purpose x86 servers, functioning on the principles of modularity and scaling pioneered in cloud provider deployments. Services previously deployed as individual appliances could now be consolidated on a common platform and configured dynamically as software based on customers' needs.

As valuable as this has been, developers have continued to innovate, and the range of applications and use cases businesses need to support has expanded. Because of this, network solutions need to embrace a wider range of hardware and software options, and operators need the freedom of choice to decide which to use according to their requirements. This versatility is part of a cloud-native design model for deploying virtual networking which we think of as NFV Next. It is the next design paradigm crucial for enabling the services that end customers require.

Operators need to consider how ready their solution suppliers are to meet these requirements as they prepare themselves for deploying their services in this cloud-native mode.

Enea has embraced the cloud-native design paradigm fully in realizing its solutions for operators' business services and delivering solutions for deployment at the network edge. It has drawn on its experience contributing to open source technologies and integrating them in large-scale networks to guide its work. Its NFV Access and NFV Core solution offerings are based on thoughtfully selected combinations of open source technologies validated in NFV community labs and in Enea's own integration facilities. Enea has taken a visionary step in supplying its solutions as software-only offerings, allowing customers freedom of choice in determining the underlying hardware to employ in both customers' premises and in the operators' cloud. It has pursued this open paradigm at each level of implementation: VM and container, open integration for VNFs, open choice of element management applications, and open interfaces for integration with a wide range of orchestration systems. Coupling this with its heritage of supporting operators with comprehensive integration and support services in large network deployments allows it to function as a supplier of leading-edge solutions for business networking services that are ready to deploy and enable on a global scale.

By making these choices Enea has positioned its NFV portfolio at the leading edge of solutions that deliver the agility required in the fully connected, cloud-native environment businesses are embracing as they move forward. NFV Access and NFV Core deserve inspection by operators and systems integrators for their potential in addressing their agility and innovation goals in these deployments.

About ACG Research

About Enea

Enea develops the software foundation for the connected society with a special emphasis on reducing cost and complexity at the network edge. We supply open-source based NFVI software platforms, embedded DPI software, Linux and real-time operating systems, and professional services. Solution vendors, Systems Integrators, and Service Providers use Enea to create new networking products and services faster, better and at a lower cost. More than 3 billion people around the globe already rely on Enea technologies in their daily lives. Enea is listed on Nasdaq Stockholm. For more information: www.enea.com