insight.tech
POWERING IoT AND EMBEDDED INNOVATION

Using Virtualization to Empower IoT Network Infrastructure

Alix Paultre

The continued growth of the IoT is putting a strain on traditional networks. And it's not just a matter of more bandwidth—IoT data is fundamentally different from the voice and video packets that comprise much of the traffic on modern networks.

For one thing, IoT traffic can be hard to predict, necessitating additional network appliances such as load balancers. But deploying equipment specifically for IoT traffic is challenging.

IoT data is asynchronous and event dependent, so network infrastructure must be flexible enough to adapt to sudden and/or dramatic changes in data traffic brought on by a variety of unpredictable factors. It must also be able to scale to meet future capacity and processing demands without requiring massive, repeated investment. In addition, some data types may be mission critical, with little tolerance for latency or unreliable connections.

Traditional network hardware struggles to meet these demands. These proprietary devices are designed for a fixed set of functions with relatively predictable traffic patterns, limiting the ability of operators to efficiently allocate and reallocate network resources as data demands ebb and flow.

Network functions virtualization (NFV) offers a superior alternative. This open, software-based approach to building and deploying elastic networks can address IoT communications infrastructure needs, now and in the future.

Challenges of Legacy Network Architectures

Network capabilities have increased dramatically over the past decade, with functions such as network acceleration, deep packet inspection (DPI), and remote access servers (RAS) providing performance enhancements and additional value-added services for users.

In traditional architectures, these functions require dedicated hardware appliances that must integrate seamlessly with existing network infrastructure. To ensure consistent quality of service (QoS), more of these appliances are required as networks expand.

In contrast, the IoT depends largely on an event-driven architecture. IoT network resource demands fluctuate based on the readings of hundreds, if not thousands, of endpoints that periodically require cloud-based intelligence for command-and-control operations. Adequate bandwidth must be available instantaneously when such events occur; unfortunately, these events are largely non-deterministic, which complicates the process of allocating network resources.

In addition, IoT data is distinct from the voice and video packets that comprise much of the traffic on modern networks, and may require specialized network appliances that are tailored to the demands of machine-to-machine (M2M) communications. Such demands include <50-ms latencies typical of voice networks, with some IoT applications requiring latencies as low as <1 ms or less. These varying parameters make it important to be able to also accommodate flexible, custom QoS agreements.

Conventional approaches to network expansion are therefore a costly proposition for network operators, given the unique requirements of heterogeneous IoT/non-IoT networks. Not only do dedicated network appliances carry

initial investment and management overhead with each additional unit deployed, but the static nature of these hardware platforms limits their effectiveness as demand rises and falls across what are now increasingly dynamic networks.

This dynamic environment also makes readily available, open, flexible, low-cost software support from many vendors increasingly important.

Advantages of NFV for Evolving Network Architectures

NFV provides an alternative to inflexible, hardware-centric network infrastructure development, management, and expansion. NFV (along with software-defined networking (SDN)) allows control-plane functions to run as virtual network functions (VNFs) on standard, off-the-shelf servers. Virtual switches are then used to orchestrate and route network traffic across a software-based data plane.

With functions running as virtualized applications, a single off-the-shelf server can perform multiple roles simultaneously and be reconfigured as needed on the fly. This is a major departure from the costly legacy procedure of purchasing, managing, and maintaining several different devices to address different applications **(Figure 1)**.

VNFs can replace multiple, dedicated appliances for more scalable IoT networks. Application acceleration, load balancing, policy management, network optimization, DPI, intrusion detection and prevention (IDS/IPS), distributed denial of service (DDoS) and other web-based protections, and firewalls are all applications that can be virtualized on such platforms.

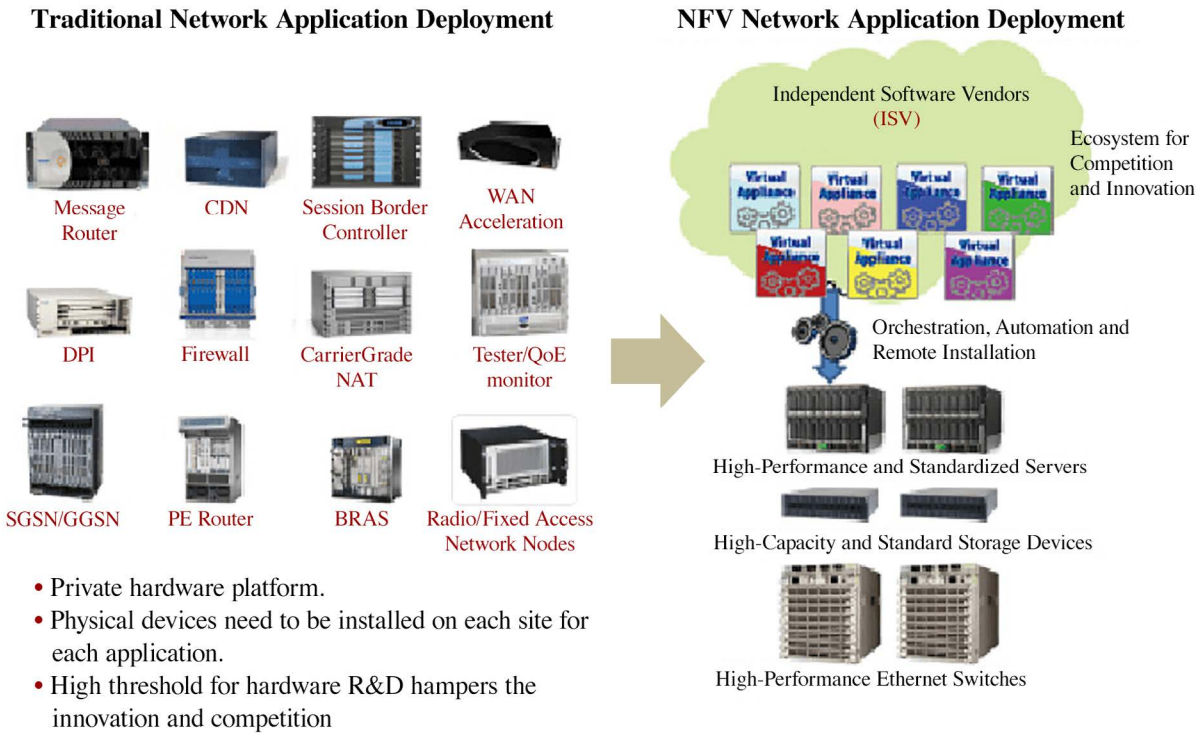

Figure 1. NFV technology allows many proprietary hardware appliances to be consolidated into virtual applications that run on standard hardware. (Source: ZTE)

NFV Standards and Systems

One reason NFV is cheaper, faster, and more flexible than traditional network implementations is the availability of open-source network virtualization tools that define an information model, set of APIs, and control protocols:

- **OpenStack** provides the structure for virtual machines (VMs), which are the operating systems that form the basis of VNFs. OpenStack Neutron, for example, is the networking component that abstracts Linux network configurations using a common API wrapper for VNFs such as Open vSwitch, virtual local area networks (VLANs), and iptables/netfilter.
- **OpenDaylight (ODL)** provides the code and architecture for virtualizing a network controller. Open vSwitch is a production-quality, multilayer, virtual network switch that can connect to an ODL controller. The Linux Foundation's Open Platform for NFV (OPNFV) project helps refine ODL as an SDN controller for NFV architectures.
- **The Intel® Data Plane Development Kit (DPDK)** is a set of data plane libraries and NIC drivers that provide a programming infrastructure for accelerated packet processing on off-the-shelf systems.

With the Intel® Architecture now ubiquitous in networking and data center environments, VM and VNF functions based on the NFV-enabling technologies mentioned can run on solutions based on Intel® processors. A good example of such an off-the-shelf network platform is the PL-8120A, a high-performance 2U rackmount solution from WIN Enterprises. The PL-8120A is based on the high-efficiency Intel® Xeon® processor E5-2600 V3/V4 and Intel® C612 series chipset, with variants that support up to 22 processor cores; 512 GB of memory using 16 DDR4 register DIMMs; 80 PCIe lanes; 64 gigabit Ethernet (GbE) ports; and ample additional I/O for demanding NFV applications **(Figure 2)**.

Figure 2. The PL-8120A is a 2U rackmount, high-performance solution for demanding NFV applications. (Source: WIN Enterprises)

In addition to reduced power consumption, cabling, rack/floor space, and inherent economies from off-the-shelf networking solutions like the PL-8120A, vendors like WIN Enterprises offer support for network engineers looking to overcome the challenges of NFV migration. This includes BIOS optimization and PCIe layout services that allow operators to increase reliability and fully leverage the features of NFV.

Dynamic NFV Architectures for Unpredictable IoT Networks

Taking network functions out of a hardware-oriented paradigm offers cascading benefits for network operators, from upfront investment to rapid resource allocation to future scalability. Thanks to NFV, the price of keeping pace in the IoT network evolution just went down.