INTRODUCTION TO CLOUD SYSTEMS

Lecture 11 – SE-O SD-WAN, Overlay provisioning, SDN essentials, SE-O Private Wireless

INTEL SE-O – SD-WAN

Software-Defined Wide Area Network (SD-WAN)

An SD-WAN is a set of network functions that enable application-aware, intelligent, and secure routing of traffic across the WAN. An SD-WAN typically uses the public internet to interconnect its branch offices, securing the traffic via encrypted tunnels, basically treating the tunnels as “dumb pipes”. Traffic at the endpoints can be highly optimized, because the network functions at a branch are virtualized and centrally managed. The SD-WAN manager can also make use of information about the applications running at a branch to optimize traffic.

Smart Edge Open provides an edge computing-based reference architecture for SD-WAN, consisting of building blocks for SD-WAN network functions and reference implementations of branch office functions and services, all running on an Smart Edge Open edge node and managed by an Smart Edge Open Controller.

INTEL SE-O – SD-WAN

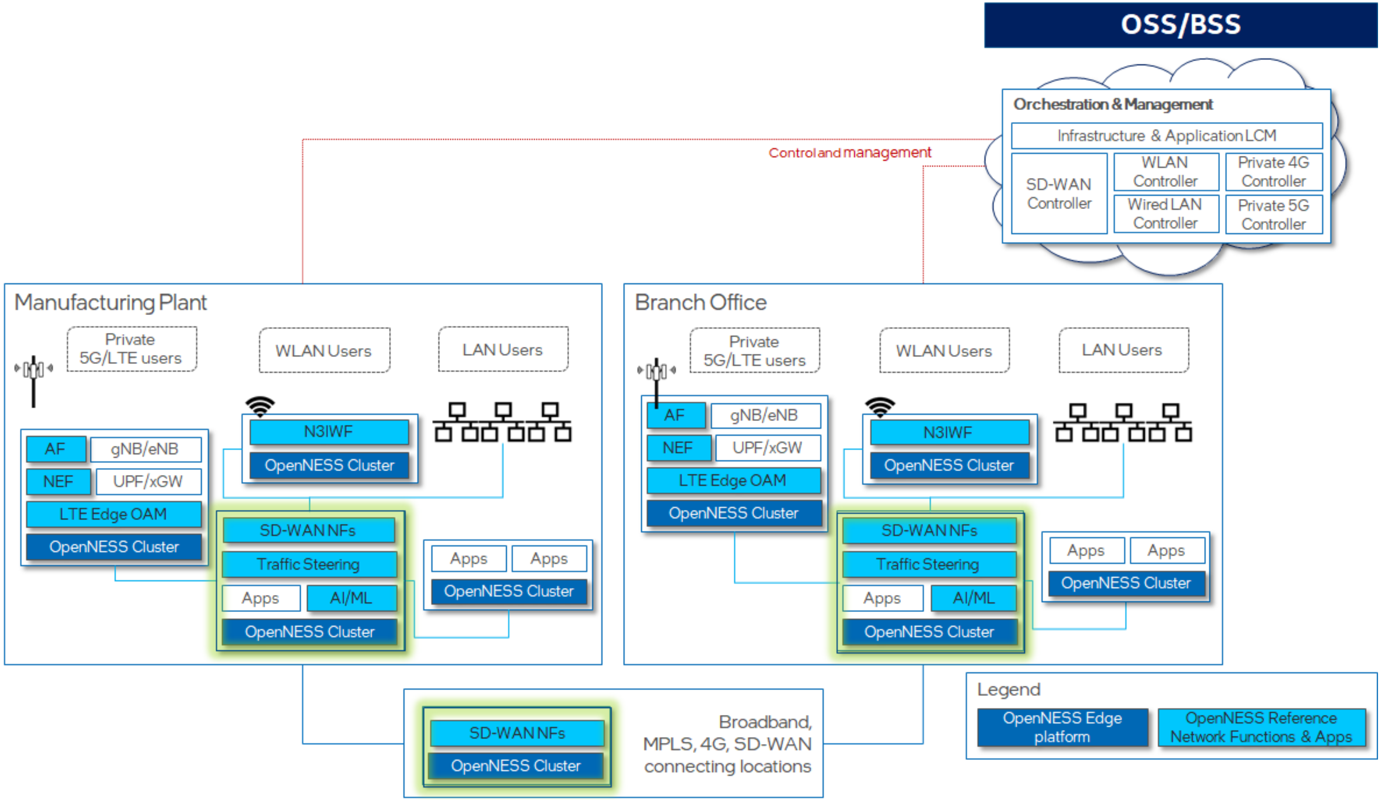

https://smart-edge-open.github.io/ido-specs/doc/reference-architectures/smartedge-open-experience-kit_sdwan/#software-defined-wide-area-network-sd-wan

INTEL SE-O – SDN essentials

https://smart-edge-open.github.io/ido-specs/doc/reference-architectures/smartedge-open-experience-kit_sdwan/#software-defined-wide-area-network-sd-wan

INTEL SE-O – SDN essentials

Software-defined networking (SDN) capabilities, enable us to virtualize the enterprise data center network similar to the way we virtualize servers. SDN separates the control plane from the data plane to provide a centralized, programmable interface for the SDN enabled portion of the network. By making network components programmable, SDN provides for newer multitenant models and distributed security access control, enhances scalability, and enables rapid service innovation through network applications.

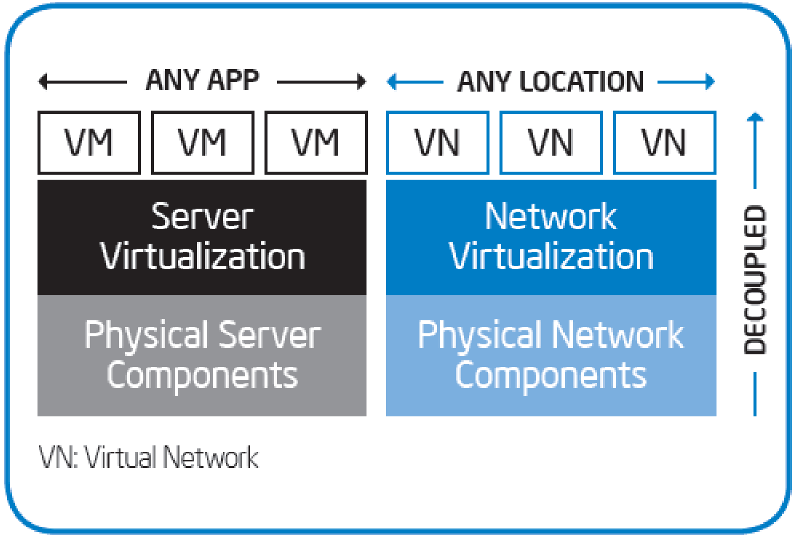

INTEL SE-O – SDN essentials

The SDN overlay network configuration provides the most value in terms of optimizing data flow in data centers. In the second half of 2013, SDN architecture was deployed an overlay network in production environment. The three primary components of this architecture are the SDN controller(s), gateway nodes, and hypervisor nodes, which connect to the VMs using virtual switches.

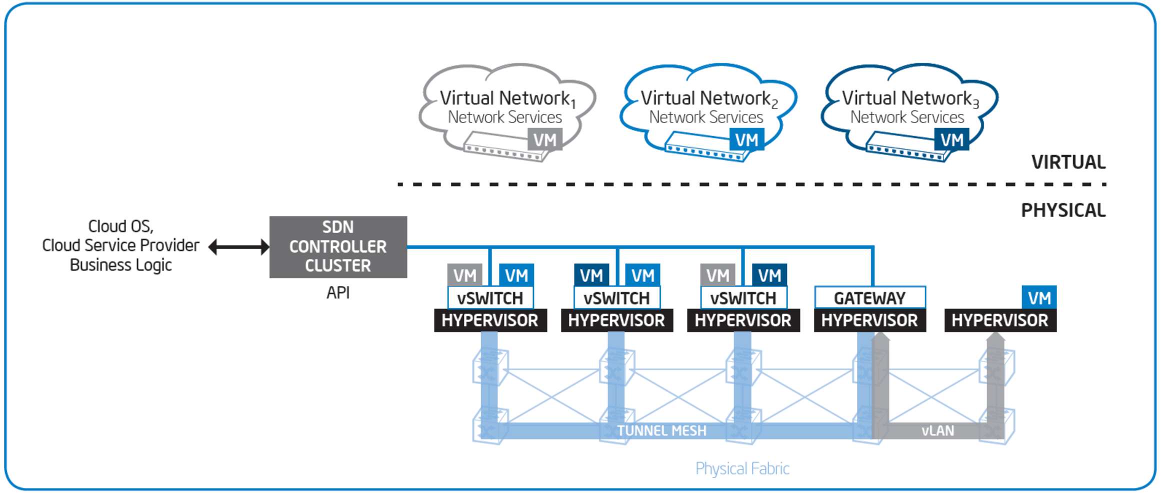

INTEL SE-O – SDN essentials

SDN Controller

The SDN controller makes an SDN network possible by abstracting the network from the hardware. The controller - in reality a software application, not a piece of hardware - manages the communication between applications and network devices. In a large implementation, there can be more than one SDN controller, which creates an SDN controller cluster.

The centralized-console approach results in intelligent networking characterized by optimized communication flows and automated configuration, along with an overall view of the network from a single console:

- Optimized data flows. A traditional network has a single path from the source of communication flow to its destination. The SDN controller can identify multiple paths for a flow and can split the flow's traffic across multiple nodes. The SDN controller optimizes the network path for a particular data flow based on the source and destination nodes. These capabilities increase network performance and scalability.
- Automated configuration. In contrast to a traditional, manual, device-by-device network configuration, the SDN controller saves time and increases configuration accuracy and consistency by automating the configuration of network devices. This approach helps to easily adjust configurations when network conditions change. Essentially, the SDN controller helps to manage the entire network architecture as if it were a single device.
- Overall view of the network. The SDN controller's console provides network administrators with a global view of the entire network—improving decision making and management efficiency.

INTEL SE-O – SDN essentials

The SDN controller’s programmable interface gives network administrators greater control over network traffic than is possible in a traditional network. For example, our security policies may require inbound traffic for a particular server to pass through a firewall. But outbound traffic does not pose a security threat and would not necessarily have to pass through the firewall. In a traditional network, this level of granular control is difficult to achieve. With SDN, an employee who is not a network engineer can easily program the controller to redirect outbound traffic around the firewall. These policies can be applied through the SDN controller at the session, user, device, and application levels.

INTEL SE-O – SDN essentials

Gateway/Service nodes

Gateway nodes provide a way to integrate VLANs’ physical network fabric with SDNenabled cloud fabric, which is configured and managed by the SDN controller. Gateway capability allows the integration of the two environments, allowing the SDN-enabled hypervisor to communicate with the non-SDNenabled network. Services nodes process the broadcast, multicast, and flow establishment between VMs.

INTEL SE-O – SDN essentials

Hypervisor nodes

The hypervisor nodes host the VMs. As illustrated in Figure 3, each hypervisor node runs the virtual switching module, which is where the SDN controller builds the virtual networks for tenants. A single hypervisor might host VMs from several virtual networks. Policy control is managed by a virtual switch between the hypervisor and the VMs.

INTEL SE-O – SRIOV

The Single Root I/O Virtualization (SR-IOV) specification is a standard for a type of PCI device assignment that can share a single device with multiple pods.

SR-IOV can segment a compliant network device, recognized on the host node as a physical function (PF), into multiple virtual functions (VFs). The VF is used like any other network device. The SR-IOV network device driver for the device determines how the VF is exposed in the container:

- netdevice driver: A regular kernel network device in the netns of the container
- vfio-pci driver: A character device mounted in the container

You can use SR-IOV network devices with additional networks on your OpenShift Container Platform cluster installed on bare metal or Red Hat OpenStack Platform (RHOSP) infrastructure for applications that require high bandwidth or low latency.

E810 features: https://cdrdv2-public.intel.com/630155/630155_E810%20Feature%20Summary_rev3_4.pdf

INTEL SE-O – Private Wireless

Intel® Smart Edge Open Private Wireless Experience Kit provides a reference blueprint for on-premise 5G deployment with edge services. Combining Intel cloud-native technologies, wireless networking, and high-performance compute, the Private Wireless Experience Kit delivers AI, video, and 5G network functions with optimized performance on Intel edge platforms.

The Private Wireless Experience Kit accelerates private 5G deployment with Intel Edge platforms.

INTEL SE-O – Private Wireless

Intel Smart Edge commercial software can connect all on-premises devices and applications to local compute and storage over private LTE and 5G wireless networks. This turnkey software also enables and governs far cloud resources and public-carrier services.

INTEL SE-O – Private Wireless

Key high-level benefits of the Intel Smart Edge commercial offering include:

- Secured operational environment
- Higher-performing compute and storage
- Ease of network deployment and use
- Faster application responses and lower latencies

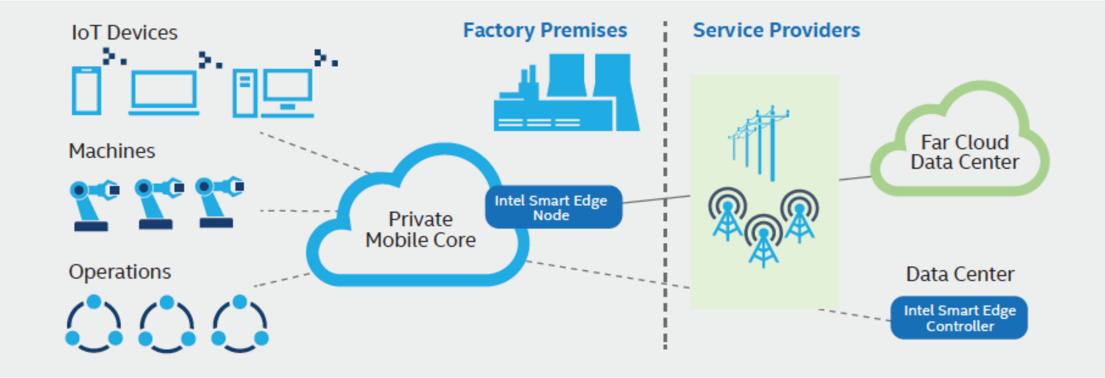

INTEL SE-O – Private Wireless

A comprehensive edge compute solution comprises both software and world-class silicon. Intel Smart Edge provides complete application lifecycle services for the network edge. This enables ready-to-deploy applications as well as the zero-trust security in the unified Intel® architecture required to onboard devices and help protect resources from unauthorized access.

Through an extensive set of deployment and management tools, this software enables any enterprise to set up a private, on-site LTE network. Only identified devices associated with the factory have access to these facilities, thereby better ensuring performance and security. Local hosting of applications provides submillisecond access to compute and storage by authorized devices located within the premises.

INTEL SE-O – Private Wireless

To further enhance your edge capabilities, servers based on the Intel® Xeon® processor provide embedded technologies for operationalizing and securing the network and compute edge. As an example, Intel® Secure Device Onboard (Intel® SDO) ensures automated deployment of endpoint devices, enabling a chain of trust between these devices and other resources.

INTEL SE-O – Provisioning

The Private Wireless Experience Kit hosts the 5G access network and core network functions on a single cluster. There are two ways to deploy the cluster.

- Autonomous Deployment Through ESP The Edge Software Provisioner (ESP) enables ODMs, System Integrators and Developers to automate the installation of a complete operating system and software stack (defined by a Profile) on bare-metal or virtual machines using a “Just-in-Time” provisiong process. The software stack can include software components, middleware, firmware, and applications. Automating this process increases velocity by focusing resources on rapid development, validation of use cases and scalable deployment. ESP simplifies customer adoption through confidence gained validating Profiles. Profiles are cloned and distributed through GitHub containing the human readable prescriptive literature to deploy the complete operating system and software stack. In summary, this a scalable, simple bare metal provisioning process including virtual machine provisioning.

INTEL SE-O – Provisioning

- Staged Deployment Through ESP Staged Deployment is carried out in two phases. Phase 1 only deploys infrastructure through ESP without deploying any vendor’s 5G CNFs this is similar to Autonomous Deployment through ESP. In Phase 2 of the deployment, the customer needs to install the vendor 5G CNFs manually by following the vendor guide after the infrastructure deployment through Phase 1 is finished.