



OCP Enabled Switching

SDN Solutions Guide


NEC's ProgrammableFlow Architecture is designed to meet the unique needs of multi-tenant data center environments by delivering automation and virtualization on commodity switches.

What is driving the need for Software Defined Networking (SDN) in the Data Center?

As the pace of transformation from enterprise to digital enterprise increases, the need for IT alignment within the enterprise also increases. Market forces are driving the pace of change and to stay competitive, enterprises need to respond. In order to meet these new demands, many firms are adopting Rapid Application Development (RAD) tools, combined with server virtualization. While server virtualization provides more efficient allocation and flexibility of compute resources, RAD increases enterprises' agility by reducing the development time of new applications through prototype acceleration, allowing for on demand project changes. But, traditional networking technologies in the data center fall short of meeting the demands of RAD and server virtualization.

RAD and server virtualization requires seamless deployment of network services in support of new virtual machine deployments, and orchestration capabilities to replace labor intensive processes with deployment automation, network visibility, dynamic scalability and simplified administration. Additionally, new demands on the data center requires simplified access to network programmability to customize orchestration and automation tools, as well as leverage the growing ecosystem of application partners offering network services.

Existing Network Model



- Access layer connects to devices.
- Aggregation layer forwards traffic between racks and the core when required (internet/external).
- Core layer basically routes traffic to or from data center devices.
- Spanning tree is required on all L2 switches which disables one of the uplink ports.
- Every device needs to be configured with spanning tree protocol which otherwise will result in loop.
- Requires significant time to add a new switch into the network.
- More devices = more configuration required. Management becomes cumbersome especially when scaling the network up.
- Aggregation/Core devices are generally chassis-based devices which are expensive and consume high power.
- Replacing/maintenance window for chassis are very large and will impact the network with more downtime.
- Device-based policies are individually configured/managed on each device which results in a long down time required to troubleshoot when failure occurs.

Traditional network environments operate through distributed intelligence, with logic maintained individually by every switch or router in an environment, producing silo'd capabilities and potential dependency on a particular vendor's hardware. With SDN, intelligence and associated visibility is moved from individual switches and routers to a centralized, software driven controller server. Using open standards protocols, the SDN controller empowers administrators to



set and maintain network-wide policies from a centralized location, enabling the abstraction of network control from the devices in the network while routing and switching hardware receive and execute the instructions. Moving from a distributed, hardware driven intelligence environment to a centralized, software driven intelligence environment with SDN networking, the service dependency on hardware is removed, administration is consolidated to a single platform and GUI, and network wide visibility is achieved. A software driven controller provides administrators with new opportunities to leverage orchestration, automation and other specialty applications.

Software Defined Networking Meets Open Compute

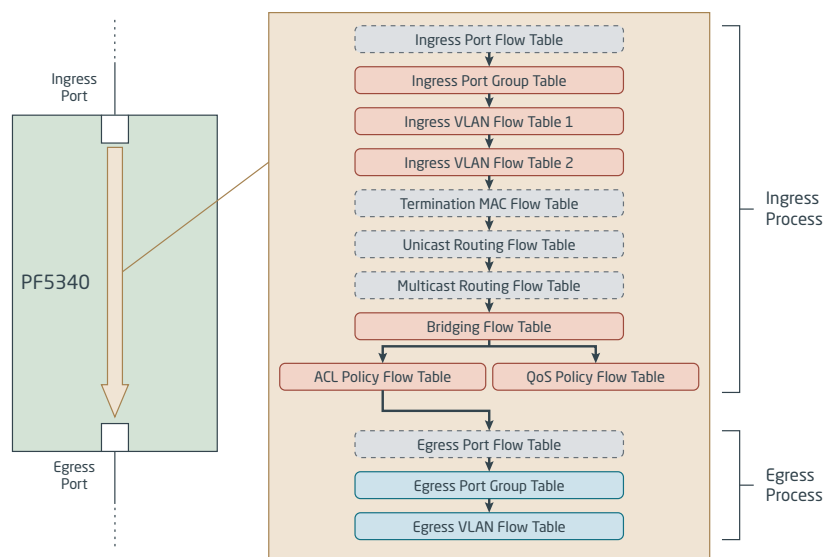
Datacenter operators have recognized that proprietary vendor hardware leads to higher costs and less service flexibility. For this reason, many large datacenter operators such as Facebook®, Google®, and Bank of America® are taking part in the Open Compute Project (OCP) Initiative. The Open Compute Project Initiative seeks to standardize hardware designs, providing economies of scale that extend to all participants.

Most organizations, however, do not have the software development resources to take advantage of Open Compute Hardware. These organizations need production ready software that works with Open Compute Hardware.

NEC's ProgrammableFlow® Network Suite version 7 (OpenFlow-based SDN solution) is specifically designed for Open Compute Switching. The ProgrammableFlow Controller (PFC) extends the service agility of Software Defined Networking to cost efficient 10G server connectivity, switch line-up enhancement for high port density and OCP hardware for switch line cost savings.

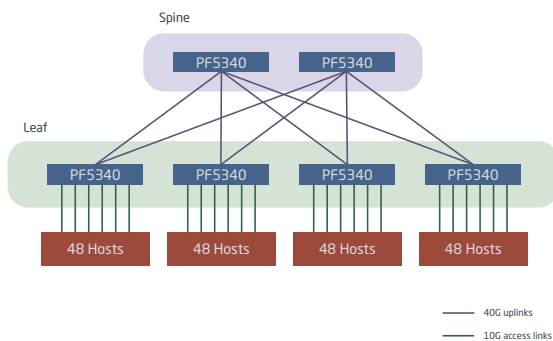
Traditional networking approaches have fixed packet pipelines, limiting the capabilities of the switch to the services and protocols that can be delivered by the device.

Unlike traditional approaches, the ProgrammableFlow Network Suite introduces a pipeline structure for packet handling. The ProgrammableFlow Open Ethernet Fabric (OEF) Architecture leverages enhancements to the ProgrammableFlow Controller software to deliver software defined services built on software packet pipeline structure. This pipeline structure is dynamically programmed by the OpenFlow protocol to leverage the hardware resources of Open Compute hardware.



As a centrally managed, software driven solution, the ProgrammableFlow OEF Architecture provides administrators with a simplified graphical user interface capable of sending execution instructions to switching and routing hardware by a drag and drop policy on the GUI. Also, for large networks, NEC's ProgrammableFlow OEF Architecture supports two controllers (in redundant configuration) while maintaining a single GUI for administration. In addition to ease of administration, NEC's ProgrammableFlow OEF Architecture provides dynamic scalability and lower cost of ownership when compared to data center networking alternatives available today.

New Network Model - Software Defined Network Fabric




- Leaf-Spine design allows all devices/hosts predictable distance and consistent delay for traffic.
- Leaf layer can connect to hosts and/or appliances (or virtual appliances) such as firewalls, load balancers and routers.
- Spine layer handles traffic between leaf switches.
- Complete network can be configured from a centralized management tool - PFC.
- PFC optimizes the path for devices communicating East-West where traffic contains data designed to be handled within data center and not outside network.
- This architecture eliminates Spanning tree while handling Broadcast, Multicast and Unknown-Unicast traffic.
- In this design, the oversubscription is 6:1 (48x10G per leaf switch = 480G connecting to 2x40G=80G). This value is reasonable for a small or medium sized data center.
- This design supports:
 - Mobile virtual machines
 - 2 way ECMP
 - Up to 40,000 networks

Scalable Multi-Tenant Architecture

While the physical constraints of the application deployment are reduced through virtualization, new scalability constraints were introduced. The IEEE 12-bit VLAN ID standard has limit of 4096 usable VLAN ID's. Prior to the explosive use of virtualization the 4096 limitation was acceptable. However, proportional to the server resource capacity increases, the amount of virtualized application a server can support also increases. When these applications require dedicated broadcast domains, the amount of VLANs created in a server's virtual switch can increase rapidly until a wall is hit and it becomes a network wide problem. NEC's ProgrammableFlow OEF Architecture addresses the scalability constraint associated with VLAN ID's.

Virtual Network configuration



- 4000 Virtual Tenant networks can be configured with 40,000 L2 segments
- Allows VLAN duplication
- Virtual Tenant network spans across the entire physical switch environment. This allows flexible configuration irrespective of device location.

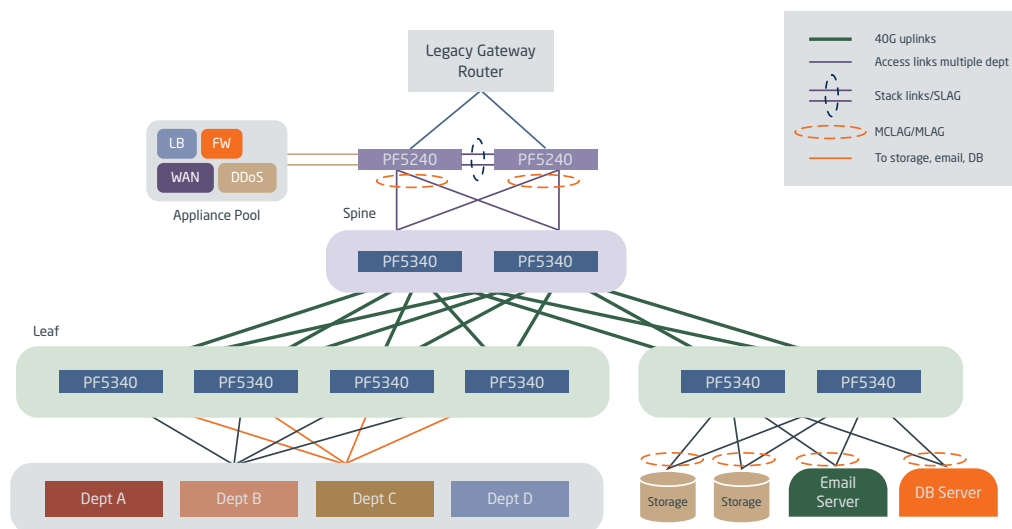
Server compute capacity has outpaced traditional data center networking technology. As result, physical servers hosting large numbers of virtualized applications compete for insufficient bandwidth at the network and storage level. To support application demand, high density servers require switch port capacity of 10Gbps and up. NEC's PF5340 switches were designed as wideband switches with multiple 10GbE and 40GbE ports capable of accommodating high density server virtualization needs now and in the future.

ProgrammableFlow OEF Architecture Features

NEC's ProgrammableFlow OEF Architecture is a 10G/40G data center networking solution with a L2 Fabric for private or public cloud, L2 cross connect and multi tenancy support. The ProgrammableFlow OEF Architecture is made up of multiple components, including:

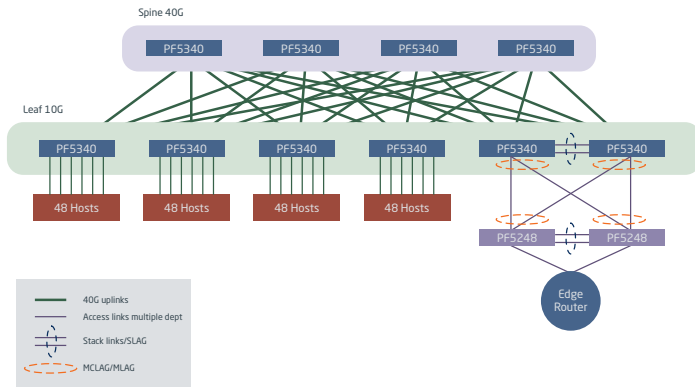
Performance

- Wideband switches with multiple 10GbE and 40GbE ports
- Up to 16 way multipath



Scalability

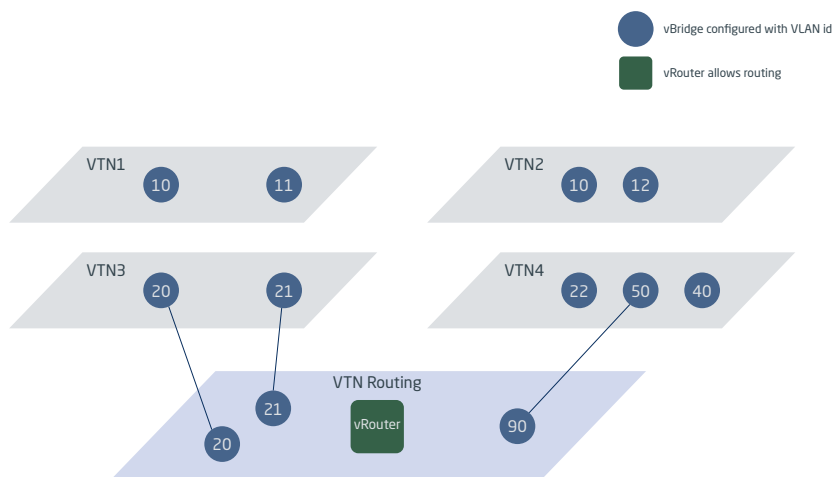
- Up to 250 OEF switches per ProgrammableFlow Controller
- Up to 4000 Virtual Tenant Networks
- Up to 4000 layer 2 networks per Virtual Tenant Network, 64,000 per system
- vBridge provides virtual L2 networks with scalability beyond 44K VLANs between/within OEF-Leaf domains that cross switches



- Additional path for fabric by adding two spine switches. All active paths providing better bandwidth and resource utilization.
- 2xPF5248 are connected to PF5340 (redundant) to provide Layer-3 between racks within data-center. Those PF5340 switches can be used to connect other devices in addition to PF5248.
- Oversubscription is well managed in this design. $48 \times 10G = 480G$ with $4 \times 40G$ uplinks = $160G$ providing 3:1 that is sufficient for medium sized datacenters.
- Virtual networks allow separating L2 networks and connect to L3 networks when necessary.

Flexibility

- VLAN duplication is supported between Virtual Tenant Networks on a OEF switch
- Adding network equipment or changes to wiring of physical networks does not require changes to settings on virtual networks when using Virtual Tenant Networks
- VLAN Transparency - All VLANs (#1 to #4094) can be configured in a lump for ports of an edge switch connected to a server, etc.



Reliability

- Link Down Relay - In case there is no output destination in an OEF-Leaf domain, link down is reported to connected equipment; as a result the external equipment that supports LAG, etc. performs re-routing.
- MC-LAG implemented by switches (PF5340 method) - redundant configuration by two switches is realized at the network edge (server connection point between an OEF-Leaf domain and an OEF Spine domain).
- Link Load balancing (+ Multiple flooding path) - Up to 16 flooding paths can be implemented.

ProgrammableFlow OEF Architecture Differentiators

NEC's ProgrammableFlow OEF Architecture delivers 10G to 40G Data Center switching for large datacenters with over 1,000 racks. With centralized network control and network virtualization features inherited from Open Flow based SDN, the ProgrammableFlow OEF Architecture utilizes a flow based forwarding method that converts the MAC address of each terminal to a flow. This new forwarding method is based on a hybrid between the flow based forwarding and the forwarding database method, which learns MAC addresses automatically, resulting in higher speed and control with lower cost in large datacenters compared to alternative offerings available today.

Compared to traditional networking environments, ProgrammableFlow OEF Architecture improves performance, cost and multi-tenancy. Additionally, operating costs are significantly reduced through automated policy enforcement while network agility and service delivery timing are improved substantially.

Physically dependent networking environments are challenging when attempting to right size capacity and performance. To meet this challenge, traditional networks are many times over provisioned as a means to ensure performance levels are maintained, which results in excessive cost of product, services and maintenance. With NEC's ProgrammableFlow Software Defined Networking, services are expanded seamlessly and configured automatically, through pre-configured policy, resulting in the elimination of over provisioning and associated costs.



Summary

NEC’s ProgrammableFlow OEF Architecture is designed to meet the unique needs of multi-tenant data center environments. The PF5340 products support 10G and 40G interfaces with attractive benefits, including high scalability at a competitive price.

Corporate Headquarters (Japan) NEC Corporation <i>nec.com</i>	North America (USA & Canada) NEC Corporation of America <i>necam.com</i>	NEC Enterprise Solutions NEC Europe Ltd <i>www.nec-enterprise.com</i>	APAC NEC Asia Pacific Pte Ltd <i>sg.nec.com</i>	Latin America NEC Latin America <i>www.lasc.necam.com</i>
--	---	--	--	--

About NEC Corporation of America: Headquartered in Irving, Texas, NEC Corporation of America is a leading technology integrator providing solutions that improve the way people work and communicate. NEC delivers integrated Solutions for Society that are aligned with our customers' priorities to create new value for people, businesses and society, with a special focus on safety, security and efficiency. We deliver one of the industry's strongest and most innovative portfolios of communications, analytics, security, biometrics and technology solutions that unleash customers' productivity potential. Through these solutions, NEC combines its best-in-class solutions and technology, and leverages a robust partner ecosystem to solve today's most complex business problems. NEC Corporation of America is a wholly-owned subsidiary of NEC Corporation, a global technology leader with a presence in 140 countries and \$29.5 billion in revenues. For more information, visit necam.com.