



Network Partitioning (NPAR) for HPE ProLiant Servers

Enhancing Network Scalability and Flexibility Solution Sheet

HPE and Marvell combine to deliver adapter-based virtualization with NPAR capability for HPE Flexible Network Adapters in HPE ProLiant servers.

OVERVIEW

The explosion of applications, rich multimedia content, increasing use of virtualization, and the growth of cloud computing are all placing heavy demands on networking efficiency. To handle increased workloads, many organizations have transitioned from 1Gigabit Ethernet (GbE) networks to 10Gigabit Ethernet (10GbE) networks. This bandwidth increase can provide significant performance and cost benefits while supporting efforts to consolidate data centers as well as consolidating multiple 1GbE adapters into a single 10GbE adapter. However, as organizations increasingly transition from 1GbE to 10GbE networks, traffic flow for application workloads may not require or be able to utilize the entire bandwidth available in a 10GbE pipe.

Historically, many organizations have built their IT infrastructures by adding separate servers, networks, and storage capacity for each application workload. These silo application deployments helped simplify management in some ways, but also left excessive pockets of unused I/O capacity and contributed to network sprawl. While network and storage convergence and server virtualization have helped address these problems, the challenge for 10GbE networks has been provisioning the appropriate bandwidth and data center resources to meet varying application demands in both physical and virtual environments.

One powerful approach to help eliminate these problems, which has been implemented by HPE® and Marvell®, is Network Partitioning (or NPAR), in both physical and virtual environments. This approach enables administrators to split up the network pipe to divide and reallocate bandwidth and resources as needed at the adapter level. NPAR further helps by reducing the number of separate adapters needed to stream different traffic types, including Fibre Channel and iSCSI (in HPE FlexFabric adapters).

HPE and Marvell have worked together for many years providing network virtualization using HPE Virtual Connect technology within HPE BladeSystem servers. Now, working with Marvell, HPE provides 10/20GbE NPAR capabilities in the HPE FlexFabric® 530/630 Series adapters, HPE Synergy® 2820C/3820C, and HPE StoreFabric CN1100R/ CN1100R-T Converged Network Adapters for the ProLiant® rack and tower servers, giving IT organizations enhanced flexibility and capacity as they migrate to 10/20GbE networks. Customers deploying HPE ProLiant rack and tower servers with GbE adapters can greatly benefit from consolidating multiple network adapters and freeing up PCI slots for other add-in card upgrades.

NPAR is an operating system and switch agnostic technology that allows customers to reduce the number of I/O adapters required to support different application workloads. Traditional best practices require separate LAN or SAN connections for different aspects of application workloads. As HPE FlexFabric 530/630 Series, HPE Synergy 2820C/3820C 10/20GbE adapters, and the HPE StoreFabric CN1100R/ CN1100R-T Converged Network Adapters already support widely used SAN protocols like Fibre Channel over Ethernet (FCoE) and iSCSI, administrators can already reduce the number of adapters needed for separate protocols, including separate Fibre Channel and Ethernet adapters. With NPAR, these adapters can now partition their network bandwidth further into multiple virtual connections, making one dual-port adapter appear as eight adapters to the operating system for use by the applications. This greatly simplifies the physical connectivity to the server, reduces implementation time, and lowers the acquisition cost of the solution.

NPAR VS. VIRTUAL CONNECT

As stated earlier, HPE pioneered network virtualization with HPE Virtual Connect within HPE BladeSystem servers. HPE’s Virtual Connect is a robust solution for HPE BladeSystem and HPE Synergy servers that include adapter virtualization, network orchestration and management as well as Quality of Service (QoS) bandwidth management in a single unified hardware/software platform. With NPAR, HPE now brings some of these same capabilities to the rack and tower server environment. Like Virtual Connect, NPAR allows a single adapter port to be presented as four separate adapter functions to the server Operating System. Also, like Virtual Connect, NPAR also allows the administrator to set QoS for each of the virtualized functions. However, Virtual Connect manages QoS across multiple servers, and orchestrates network management across multiple adapters, switches, and servers, where NPAR management is limited to managing a single adapter within a single server at a time. NPAR for HPE ProLiant rack and tower servers does provide a significant benefit for consolidation of server to network connectivity and bandwidth utilization within a given server.

FAT PIPES, MULTIPLE FLOWS

Network Partitioning begins by configuring a single 10GbE port to represent as many as four separate partitions or physical functions. Each partition is an actual PCI Express® (PCIe®) function that appears to the system ROM, OS, or virtualization OS as a discrete physical NIC with its own software driver, and each partition behaves as an independent NIC port (see Figure 1).

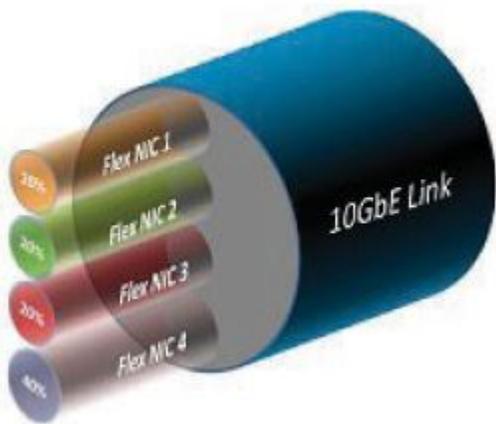


Figure 1. Physical Port Partitioned into Four Functions

BANDWIDTH PROVISIONING ON DEMAND

Traditionally, system administrators may oversubscribe bandwidth for shared connections to handle demand surges by the same applications. Instead, administrators can take advantage of the Network Partitioning capability to handle such demands. Furthermore, systems administrators can configure the weighting of each partition to provide increased bandwidth presence when an application requires it (see Figure 2). Each virtual machine (VM) assigned to a specific partition with equal weights can acquire as much free bandwidth as is available, while incrementally yielding bandwidth back as demand increases on the other partitions.

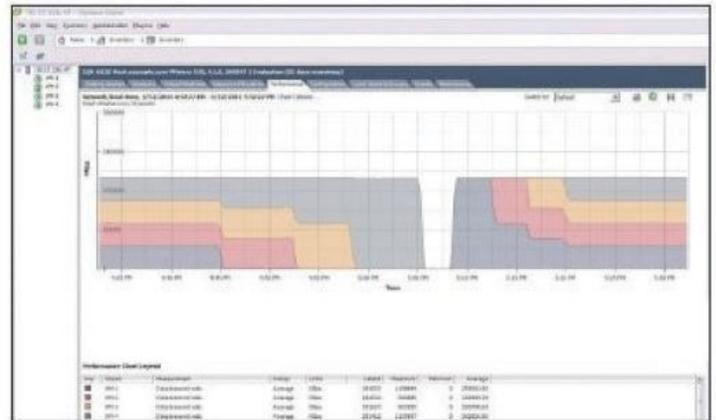


Figure 2. Variable Bandwidth Weighting

NPAR CONFIGURATIONS

Each NIC partition is presented with a dedicated PCI function that can be enabled or disabled by administrators. With this PCI function, partitions behave as normal multi- port PCIe Ethernet devices from a host system and OS perspective—requiring no changes to the OS. After the driver is loaded, each enabled partition is essentially a dedicated network connection or Ethernet and storage controller device. Standard device configurations can then be applied to the partition.

Each partition can support networking functions such as large send offload, Transparent Packet Aggregation (TPA), and multiqueue receiveside scaling, along with storage Host Bus Adapter (HBA) features such as Internet SCSI (iSCSI) Host Bus Adapter and Fibre Channel over Ethernet (FCoE) Host Bus Adapter. Administrators can configure a partition to run iSCSI or FCoE.iSCSI or FCoE. All features can be enabled or disabled on each partition as necessary.

Administrators can provision HPE FlexFabric 530/630 Series NICs, HPE Synergy 2820C/3820C adapters, and HPE StoreFabric CN1100R CN1100R-T Converged Network Adapters quickly and easily using Comprehensive Configuration Management (CCM) utility which is part of the adapter firmware. Administrators can enable and configure NPAR using this CCM, which is invoked during the server boot process.

Figure 3 shows the CCM configuration screen for setting NPAR parameters for function 2 after NPAR has been enabled with four functions on Port 0 of the adapter.

STREAMLINE THE NETWORK WITH NPAR

Network Partitioning helps simplify the data center and network and storage infrastructure in several ways. For example, when connecting servers to LANs and SANs, administrators may have to use many cables. Network Partitioning provides an alternative—consolidating Ethernet (and iSCSI or Fibre Channel when using Converged or FlexFabric adapters) connections onto a significantly reduced numbers of ports. The HPE FlexFabric 530/630 Series, HPE Synergy 2820C/3820C adapters, and HPE StoreFabric CN1100R/CN1100R-T Converged Network Adapters from Marvell support this functionality.

Like switches, Network Partitioning reduces the number of cables without adding workloads on the network—but Network Partitioning utilizes the existing adapter and eliminates the need to add switches and cables. Network Partitioning is designed to offer the following benefits over standard connectivity:

- **Reduced Network Sprawl:** A single dual-port HPE FlexFabric 530/630 Series adapters, HPE Synergy 2820C/3820C adapters, or HPE StoreFabric CN1100R/CN1100R-T Converged Network Adapters can replace up to 8 connections in a single server—consolidating infrastructure onto much less equipment, simplifying management, reducing costs, and using less energy and facility space than when deploying multiple GbE devices needed to achieve the same goal.
- **Maximized Network Scalability:** Having a reduced number of network devices and cables allows IT organizations and data centers to easily scale their networks and add servers and network devices to meet growing IT demands.
- **Simplified Administration:** Network Partitioning can also save time and labor by reducing the number of physical adapters required, which simplifies management for IT administrators and enhances their self-sufficiency. They can add or replace network cards or move workloads from one partition to another within minutes.
- **Optimized Resource Allocation:** Network Partitioning can play an important role in helping optimize bandwidth allocation and improving utilization of that allocation in both virtualized and non-virtualized environments. Virtualization can reduce the number of physical servers but can increase I/O demand for each server. Using 10/20GbE connectivity with Network Partitioning

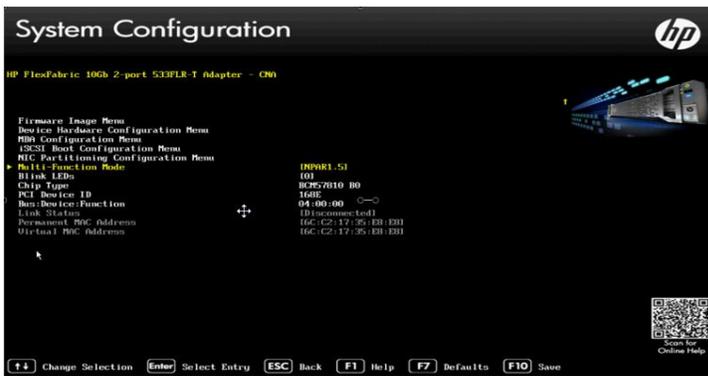


Figure 3. CCM NPAR Configuration Screen

Figure 4 shows the view of the NPAR-enabled Server adapter from Microsoft Windows® device manager. Note that there are eight 10GbE adapters being presented to the operating system by the HPE 534FLB adapter. CCM supports both UEFI or legacy boot options.

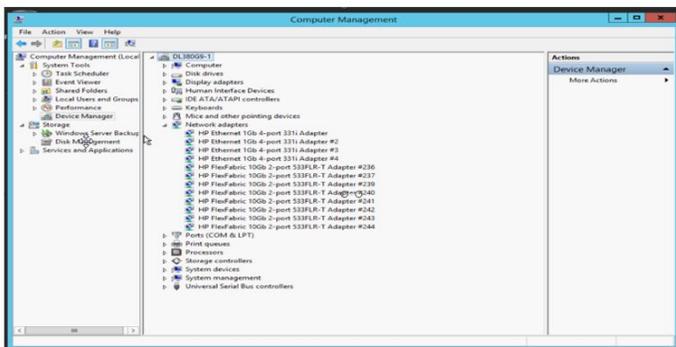


Figure 4. Windows Device Manager NPAR Enabled

addresses this problem by providing eight functions on a single integrated network adapter built into the server, so the server can efficiently handle bandwidth requirements but allow for plugging in additional devices to meet growing bandwidth requirements. Network Partitioning also allows system administrators to fine-tune the amount of bandwidth for each adapter port as required by the application and to set the bandwidth for each physical function (PF).

- **Improving CPU Utilization:** By shifting the network virtualization task from the OS to the network adapter, CPU cycles are made available to be used for other critical tasks, such as server virtualization or application processing. While system administrators who are virtualizing servers can use the OS to create multiple virtual network connections, this approach requires CPU resources. Using NPAR offloads that responsibility from the OS and processor to the adapter itself, freeing up CPU resources to run more applications or manage more virtual machines. In HPE and Marvell internal tests, NPAR has been shown to reduce server-side CPU utilization by up to 35 percent when operating a single adapter versus two or more teamed adapters.

By providing all of these benefits for 10/20GbE networks, NPAR helps optimize enterprise data centers.

BOOSTING NETWORK EFFICIENCY WITH NPAR

The Network Partitioning approach helps make migrating to today's powerful 10/20GbE networks an easy and compelling option for many organizations. 10/20GbE networks provide organizations an opportunity to achieve the performance and efficiency capabilities these networks can offer and reduce total cost of operation. Additionally, 10/20GbE networks can deliver flexibility that enhances the efficiency of HPE rack and tower servers and storage. NPAR delivers consolidation and provisioning capabilities for rack and tower servers, including the HPE ProLiant Gen8 servers, that are similar to HPE's Flex-10, Flex-20 and FlexFabric capabilities deployed in HPE BladeSystem server and Virtual Connect environments.

NPAR FEATURES IN HPE PROLIANT SERVERS

Administrators deploying HPE ProLiant Gen8, Gen9, or Gen10 servers paired with HPE FlexFabric 530/630 Series, HPE Synergy 2820C/3820C adapters, and HPE StoreFabric CN1100R/CN1100R-T Converged Network Adapters from Cavium can take advantage of several features and functions provided by Network Partitioning:

- Support for up to eight partitions per device and up to four partitions per port
- Support for server and virtualization OSs—Microsoft Windows, VMware®, and Linux® operating systems
- No OS or BIOS changes required
- Support for industry-standard 10 Gigabit Ethernet (10GbE) switches
- NIC control of the transmit flow rate from the server
- Flexible and dynamic bandwidth allocation

For HPE ProLiant and BladeSystem Server customers, Network Partitioning available in the HPE FlexFabric 530/630 Series, HPE Synergy 2820C/3820C adapters, and HPE StoreFabric CN1100R/CN1100R-T Converged Network Adapters from Cavium provides a powerful means for consolidating NICs, optimizing network connectivity, and enabling bandwidth control.

RESOURCES

HPE and Marvell Learning Site: www.marvell.com/hpe



ABOUT MARVELL: Marvell first revolutionized the digital storage industry by moving information at speeds never thought possible. Today, that same breakthrough innovation remains at the heart of the company's storage, processing, networking, security and connectivity solutions. With leading intellectual property and deep system-level knowledge, Marvell semiconductor solutions continue to transform the enterprise, cloud, automotive, industrial, and consumer markets. For more information, visit www.marvell.com.