

Cisco UCS M81KR Virtual Interface Card: Design a Flexible and Scalable Data Center

What You Will Learn

With traditional servers, I/O is inflexible and rigid. Changing the I/O configuration generally requires physical intervention and results in application downtime. The rapid adoption of virtualization brings with it additional configuration complexity. The Cisco® UCS M81KR Virtual Interface Card (VIC) helps address these challenges by creating more I/O connections to a given server while limiting the number of required physical connections to the server.

Cisco developed the Cisco UCS M81KR to provide I/O flexibility within the Cisco Unified Computing System™. This document discusses how to use the Cisco UCS M81KR's capability to virtualize I/O, in combination with other Cisco Unified Computing System features such as a wire-once cabling model and service profiles, to make moving or rehosting a physical or virtualized environment a straightforward and easily scalable operation.

Challenge

Data centers face increasing challenges resulting from two related trends: an increasing overall number of transactions, and accelerating adoption of virtualization. These trends are creating unprecedented demand for I/O accessibility and vastly improved I/O management within every server, both physical and virtual. This growing demand is affecting overall data center economics and operations in three fundamental ways:

- Need to scale applications while containing physical costs: Most existing servers are deployed with multiple network interface cards (NICs) and host bus adapters (HBAs) that all need to be powered, cooled, and stocked for sparing. Adding new servers to keep up with transaction demands also leads to cable and switch proliferation. Looking forward, I/O needs to scale with the increasing performance of processor and memory technologies. In a virtual environment, I/O also needs to scale to accommodate virtual machine mobility.
- Need to simplify dynamic operations: End users expect their applications to be available on demand, which forces IT to repurpose and reprovision servers with increasing frequency. Proliferation of components and devices—from servers and adapters to switches and cables—creates more management complexity and also more points of potential failure. Since management paradigms are typically device based and have generally not kept pace with the scale and complexity of physical implementations, troubleshooting becomes increasingly challenging. IT departments often separate server, SAN, and LAN teams, making synchronization of servers with I/O difficult.
- Need to optimize I/O for virtualized environments: Currently, it is difficult to apply consistent network policies (security, quality of service [QoS], etc.) as virtual servers move around, or to provide consistent management and troubleshooting mechanisms between a physical server and a virtual server.

Cisco Solution

Cisco addresses the I/O challenges faced by customers today with the Cisco UCS M81KR (Figure 1).

Figure 1. Cisco UCS M81KR Virtual Interface Card



The Cisco UCS M81KR is a virtualization-optimized Ethernet and Fibre Channel over Ethernet (FCoE) mezzanine card designed for use with Cisco UCS B-Series Blade Servers. The VIC is a dual-port 10 Gigabit Ethernet adapter that supports up to 128 Peripheral Component Interconnect Express (PCIe) standards-compliant virtual interfaces that can be dynamically configured so that both their interface type (NIC or HBA) and identity (MAC address and worldwide name [WWN]) are established using just-in-time provisioning. Virtual interfaces are administratively defined and are instantiated on the server at the time a service profile is assigned. Virtual interfaces are recognized by the BIOS, OS, and hypervisors as regular PCIe devices.

The architecture of the Cisco UCS M81KR allows presentation of the virtualized adapter both inward toward the CPU and outward toward the network from the perspective of the physical connection to the fabric. Because the existence, identity, and policy are applied from the service profile to the adapter, these devices are known to the BIOS prior to the operating system boot. Normal scans of the PCI bus by the system are not only unaffected, but the information is persistent across multiple system boots and migrations of the service profile to other blades within the Cisco Unified Computing System. This behavior is consistent with the stateless model that is a fundamental part of the Cisco Unified Computing System.

IT administrators can define policies for the individual virtual interfaces directly, with no restrictions on subgroupings of these interfaces. Cisco VN-Link technology enables the unique identification of the virtual adapters and presents them as logical interfaces (which are logically associated with the virtual adapters) on the Cisco UCS 6100 Series Fabric Interconnects. Cisco VN-Link technology also allows administrators to configure policy groups and to include the logical interfaces in this grouping.

Within virtualized environments, the Cisco UCS M81KR offers close integration with VMware vCenter. As virtual NICs (vNICs) on given virtual machines are added, moved, or deleted, the port group is controlled through VMware vCenter. These port groups are not configured on each VMware ESX Server, but defined as port profiles within the Cisco Unified Computing System (VLAN, access control list [ACL], etc.). This approach allows an abstraction of traditional networking configuration away from the VMware ESX configuration tasks (no need to configure networking on the VMware ESX).

Solution Benefits

The Cisco UCS M81KR offers a number of benefits to address the challenges of today's data center.

Scale Applications and Contain Costs

The Cisco UCS M81KR offers a high-performance unified I/O architecture that enables a wire-once deployment model. It dramatically reduces the number of adapters and switches that must be purchased, cabled, configured, powered, cooled, and secured. Instead of using multiple different types of adapters, switches, and management tools, customers can use a single 10-Gbps converged network adapter (CNA) and unified fabric to simplify the

infrastructure. This design makes the Cisco Unified Computing System highly energy efficient. The new Cisco VIC in combination with the innovative Cisco UCS 2104XP Fabric Extender significantly reduces cabling.

The Cisco UCS M81KR is a high-performance, dual-port 10-Gbps adapter that delivers more than 600,000 I/O operations per second (IOPS) with very low latency, helping eliminate any bottlenecks for I/O-intensive applications. The Cisco adapter also enables great flexibility in the number of virtual interfaces a user can create to meet the growing connectivity needs of applications. To scale, IT administrators need a way to create multiple distinct interfaces to a given server, while containing the number of physical connections. Further, each I/O interface needs QoS capabilities to prioritize the different types of traffic. The Cisco UCS M81KR achieves this goal by allowing users to apply QoS policies to each interface. This capability in particular helps optimize I/O for a virtualized server in which virtual machine mobility is important.

Simplify Dynamic Operations

IT infrastructure needs to adapt to changing business needs. The Cisco Unified Computing System abstracts the characteristics of the physical infrastructure and manages them through service profiles in Cisco UCS Manager. For example, an IT manager in a financial institution may use data center servers as virtual desktop interface (VDI) servers during the day, and at night repurpose them to run analytics useful for the next day. In this example, VDI servers may need two HBAs and six NICs, and analytic servers may need only two NICs. The Cisco UCS M81KR is the only adapter on the market that can create different combinations of NICs and HBAs dynamically to enable a truly stateless computing infrastructure. The IT manager can easily provision and configure these adapters using a service profile in Cisco UCS Manager. With the ability to present more adapters to a server directly, IT no longer needs to install, configure, and certify vendor-provided adapter software for every OS and hypervisor to support VLAN tagging on the host.

Resilience is very important to an agile and flexible data center. Most customers use some form of NIC teaming software that needs to be provided by the NIC vendor for every OS and hypervisor. NIC teaming also requires certification for every application environment. The Cisco UCS M81KR offers fabric failover, which enables interface failover at a physical level without involving the OS or hypervisor or certification overhead.

Cisco UCS Manager also simplifies and facilitates collaboration between the server and network groups. The network administrator can define network profiles that can be used by the server administrator to define the service profile. In a virtualized environment, the Cisco UCS M81KR integrates tightly with VMware vCenter, the VMware management tool. IT can define a port profile for a vNIC in Cisco UCS Manager, and the server administrator can apply it to a virtual machine through VMware vCenter in a transparent and collaborative way.

Optimize I/O for Virtualized Environments

A typical virtualized server with a software switch in the hypervisor often connects to multiple distinct LAN and SAN interfaces to provide separate connectivity for VMware VMkernel, service console, and virtual machine production traffic and shared SAN storage. In such a deployment, the use of four to eight NICs and two or more HBAs is common. The Cisco UCS M81KR is the only adapter on the market today that allows a user to create these from a single dual-port 10-Gbps adapter and apply network policies to each. Hence, it helps customers consolidate multiple NICs and HBAs (and the associated Ethernet and storage ports) and save costs.

The Cisco UCS M81KR can create up to 128 virtual adapters and map them to different virtual machines in a VMware deployment using Cisco VN-Link technology in hardware. VN-Link technology, jointly developed by Cisco and VMware, enables the network to be virtual machine aware. It provides policy-based virtual machine connectivity, enables policies to persist after VMware VMotion migration, and maintains a consistent operational model. VN-Link, implemented in hardware, connects virtual machines directly to virtual interfaces on the fabric interconnect and helps free CPU cycles that are normally used to provide networking functions on the server, resulting in better virtual server performance.

Since a consistent operational model between physical and virtual servers is necessary to facilitate virtualization automation, the Cisco UCS M81KR uses VN-Link capabilities to manage and operate interfaces on a virtual server in a way that closely follows the operational model for interfaces on a physical server. Management of the Cisco UCS M81KR is tightly integrated with VMware vCenter. Network policies and configurations, called port profiles, can be created in Cisco UCS Manager by a network administrator, exported to VMware vCenter, and applied to a virtual machine by a server or virtualization administrator using VMware vCenter. This capability helps customers greatly simplify management of adapters and policies.

The Cisco UCS M81KR has built-in architectural support for technology that allows each virtual machine to directly access the adapter hardware, bypassing the hypervisor completely. Known as hypervisor bypass technology, this capability relieves some of the computing burden on the hypervisor and further improves performance.

Use Cases

The Cisco UCS M81KR can provide significant benefits for all application environments. The capabilities offered by the Cisco UCS M81KR address the following scenarios especially well:

- **VMware virtualization:** The Cisco UCS M81KR helps reduce the number of physical adapters, simplify management, increase hypervisor performance, and facilitate collaboration between groups.
- **Database deployments:** The Cisco UCS M81KR helps reduce the number of physical adapters, simplify management, and scale performance for demanding applications such as databases.
- **Infrastructure as a service (IaaS):** The Cisco UCS M81KR helps reduce the number of physical adapters, repurpose servers dynamically, and reliably scale performance for the various services.

Conclusion

The Cisco UCS M81KR Virtual Interface Card provides many new capabilities through its innovative technology. With it, the user can:

- Reduce physical infrastructure by consolidating and virtualizing multiple NICs and HBAs
- Simplify operations
- Facilitate collaboration between IT groups
- Enable a stateless and agile infrastructure
- Provide a scalable and high-performance I/O architecture
- Optimize I/O for a virtual environment

For More Information

For more information about the Cisco UCS M81KR, please visit

http://www.cisco.com/en/US/prod/collateral/ps10265/ps10280/data_sheet_c78-525049.html.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems (transitions) BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, CCSP, Cisco IaaS, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Nexus Connect, Cisco Prime, Cisco Sessionless, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Hio for Coda, Hio Mind, HioShare (Design), Hio Ultra, Hio Video, Hio Video (Design), Indent Broadband, and Welcome to the Human Network are trademarks. Changing the Way We Work, Live, Play and Learn, Cisco Capital, Cisco Capital (Design), Cisco Finance (Stylized), Cisco Store, Flip Gift Card, and One Million Acts of Green are service marks, and Access Register, Alronet, All built, AsyncOS, Bringing the Meeting to You, Catalyst, CCDA, CCDE, CCIE, CCIP, CCNA, CCNP, CCS, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Lumina, Cisco Nexus, Cisco Prime, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Continuum, Fireport, IronSwitch, Event Center, Isdler, Know Me, Learning, GetMeat, LYNX, OS, iPhone, IronPort, the IronPort logo, LaserLink, LightStream, Linksys, MeetingPlace, MeetingPlace Online Sound, MGX, Networks, Networking Academy, PCNow, PX, PowerKEY, PowerPanel, PowerTV, PowerTV (Design), PowerVu, Prisma, ProConnect, ROBA, Solderless, SMI/Thin, Spectrum Expert, StackWise, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0910)