



HRG Assessment: Cisco's Unified Computing System (UCS)

The emergence of converged systems represents an innovation in packaging. These systems bring together or package network, compute, storage, and system level software into a single easy to consume, implement, and manage IT infrastructure solution

Today with converged systems available from vendors like Cisco, Dell, HP, IBM, and VCE customers are finally able to gain control of their IT infrastructure and costs.

The rapid success of Cisco's Unified Computing System (UCS) has made Cisco a leader in converged systems both in terms of revenue and market share. Cisco UCS meets customer requirements for a cost effective solution that is easy to buy, install, implement, manage and scale.

Cisco's approach to meeting customer requirements involves Cisco UCS Manager, Cisco B-Series blade and C-Series rack servers, Cisco Fabric Interconnects and Fabric Extenders, Cisco partner ecosystems solutions, and Cisco UCS Central Software. Using a modular "fabric centric" approach customers can easily and rapidly scale their corporate IT infrastructure from 100s to 1000s of blades all managed through a single management software console. This rapid scale out capability positions customers to most effectively reap the benefits of their preferred Cisco ecosystem partner based Cloud Computing solution.

CXOs and IT professionals are challenged daily by the increased complexity of business environments and the need to control the cost of operations. Increasingly IT budgets are consumed by management, administration, and indirect non-technology costs.

Customers increasingly see Converged Systems as a good way to meet these challenges. Cisco's Unified Computing System is an easy to configure, easy to scale, easy to manage, converged computing platform designed specifically to meet many of today's business and IT challenges. With Cisco UCS Cisco's B-Series blades and C-Series rack mount servers can be integrated, managed, and viewed as a single unified compute resource.

Cisco UCS

The Cisco approach is a modular “fabric centric” architecture built on top of Intel Xeon and XML API technologies in combination with Cisco Networking expertise and Cisco’s 10GB converged fabric which comprise this solutions base level or lowest common denominator building block. The software (UCS Manager) and hardware (UCS Fabric Interconnects) are sold as an integrated and inseparable package. Cisco embeds their UCS Manager software in the Cisco UCS Fabric Interconnect switches. Cisco UCS Manager, a device manager, only manages Cisco Blades, Rack Mount Servers, and UCS certified components. All Cisco Blades and Rack mount servers are Intel based. A single Cisco UCS manager domain can scale out to manage up to 160 UCS servers.

Cisco provides a rich XML API for developers and system level software, tools, and utilities providers to integrate their offerings with Cisco UCS Manager. BMC and CA use this API to work with Cisco UCS Manager to stand up, provision, and manage UCS Blade and Rack mount servers as well as the Virtual Machines on those servers. Cisco UCS Manager manages Cisco Switches, Blades, Rack Mount servers, CNAs, and in chassis FEX (Fabric Extenders). Since the introduction of Cisco UCS in 2009 Cisco and its partners have been continually evolving their XML API based integrations to enhance the manageability of the UCS platform.

Basic Cisco UCS Architecture

The principal Cisco UCS benefit for customers is a significant reduction of time and money spent on software and hardware deliberation, selection, installation, implementation, and administration.

The Cisco UCS 5108 Blade Server Chassis is six rack units (6RU) high and can mount in any industry-standard 42 U rack. Each 42 U rack can fit 2 Fabric Interconnect switches and six Cisco UCS Blade chassis. A Cisco 5108 chassis accommodates 2 FEX or Fabric Extenders with eight half-width Cisco UCS B-Series Blade Servers, four full-width, or a mix of these form factors. Cisco servers are only available with Intel processors and the UCS manager only manages Cisco UCS certified hardware.

Each of the 2 FEX or pass through switches (they do not route traffic) in each Cisco UCS Blade chassis is connected to the south side server ports on the Fabric Interconnect. Two Fabric Extenders (FEX) are required for each UCS 5108 chassis for availability and fail over purposes. These FEX or pass through switches are connected to the south side ports on one of the top of rack Fabric Interconnect switches. Each FEX has either 4 or 8 north bound 10 GB uplink ports depending on the model.

Each Fabric Extender connects to a different top of rack Fabric Interconnect (FI) to ensure that there are redundant 10 GB Cisco converged fabric paths in the event that either one of the FEX or one of the Fabric Interconnect switches fails. Cisco UCS provides multiprotocol support for FCOE and iSCSI on the Cisco 10GB converged fabric. Under normal conditions UCS manager runs primarily on one FI, with the other FI mirroring all the activity via a heartbeat connection. If the primary FI goes down, UCS Manager takes over from the secondary FI with no major disruption of service. It is only when one of the interconnects is taken out of service for maintenance or goes down that all the traffic in the system flows through one Fabric Interconnect creating a temporary potential for increased latency at the switch.

Availability

The Cisco UCS Blade System UCS 5108 chassis based solution requires 2 identical Top of Rack Fabric Interconnect switches in order to provide redundancy and availability at the Layer 2 Fabric Interconnect Switch level. The same number of uplink ports on each FEX must be cable connected to the south side

ports on each of the Fabric Interconnects such that FEX 1 will connect only to FI 1 and FEX 2 will only connect to FI 2 in order to preserve system availability and fail over capability throughout the Fabric path.

For the UCS Blade system to deliver its maximum throughput, all uplink ports on each of the Fabric Extenders (FEX) in the UCS 5108 chassis must be connected to one or the other of the top of rack Fabric Interconnect switches. Currently HRG does not recommend any converged system for extremely latency sensitive HPC class workloads that require continual real-time collaboration and messaging between compute nodes. When considering an IT infrastructure that is heterogeneous, and based on open standards, customers need to understand ahead of time how a converged system like Cisco UCS will be integrated and managed as part of that environment. Although Cisco in chassis and in rack node to node connectivity is not as fast as infiniband many cluster oriented workloads can be run effectively on UCS.

HRG rates the Cisco UCS recommended configuration at an **AEC 3** level of availability where the system supports business functions that require uninterrupted computing services, either during essential time periods, or during most hours of the day and most days of the week throughout the year. This means that the user stays on-line. However, the current transaction may need restarting and users may experience some performance degradation. Go to <http://www.hrgresearch.com/High%20Availability.html#> to learn more.

Latency

Cisco UCS provides predictable levels of latency and performance regardless of where a workload or blade server is located as long as they are connected to the same fabric interconnect, often referred to as a “UCS instance.” A single UCS instance can span several racks (up to 160 servers) all connected to one pair of FI’s. Within a single Cisco UCS instance or domain it takes a predictable amount of time (latency) for each blade to be accessed through the Fabric Interconnect. In those cases where customers are running latency sensitive HPC type workloads that require extremely low levels of messaging latency this type of solution may not be a best fit and real-world benchmarking of customer workloads on a Cisco UCS system is recommended to ensure that the system will perform as expected.

Inside the Cisco UCS 5108 Blade System Chassis

Each FEX connects to 8 UCS B series blade servers through the 5108 chassis mid-plane. The 10 GB converged FCoE Fabric is brought to each FEX through cables plugged into 8 of the south side ports on each of the UCS Fabric Interconnects thereby extending the Cisco 10 GB FCoE converged network fabric as well as the out of band UCS manager bi-directional system level communications from the top of rack switch to the in-chassis Fabric Extender and then from there via the 5108 chassis mid-plane connecting to the B Series Blades.

Cisco UCS efficiently distributes workloads across chassis of blades or racks of rack mount servers to effectively meet workload specific inter node messaging latency and SLA requirements. Furthermore, easy reconfiguration of the environment to meet changing workload demands can be facilitated through the use of pre-established customer specific Cisco US service profiles.

Cisco UCS Central Software

Cisco UCS Central software, a virtual machine image, manages numerous Cisco UCS Manager domains, can be run on a hypervisor and can, according to Cisco, be scaled up to manage 10,000 Cisco

UCS servers from a single unified UCS Manager console. Cisco UCS Central integrates and consolidates multiple Cisco UCS Manager domains and utilizes UCS Manager to provide global configuration capabilities for resource pools, policies, and firmware thereby simplifying global policy compliance across distributed data centers giving administrators a "define once, deploy many times" workflow capability for the entire compute infrastructure.

Cisco UCS Central brings inventory and fault information across multiple Cisco UCS domains together in one place facilitating service assurance for the entire Cisco UCS infrastructure. UCS central exposes a rich XML API like Cisco UCS Manager for the integration of customer and partner system management software solutions.

Cisco UCS Manager

Cisco UCS Manager's embedded device management software manages the software and hardware components of the Cisco Unified Computing System™ across multiple chassis and virtual machines through a Java based GUI, a CLI (command-line interface), or an XML (Extensible Markup Language) Application Programming Interface (API). Service Profiles in the UCS Manager application can be used to set up and configure stateless Intel Xeon based Cisco Blades, Rack mount servers, and virtual machines. UCS manager by itself does not configure VM's. However, when UCS Manager is integrated using the available XML API with either VMware, Microsoft, or Openstack VM configuration can be accomplished from within those platforms. VMs can be deployed from within UCS Director when deploying Service Profile settings. These Service Profiles can be 'moved' with a virtual machine when it is moved using VMware's VMotion as in the case of a server failover or when reallocating capacity to satisfy changing workload requirements on the fly.

Service Profiles

Cisco UCS Manager Service Profiles are created by server, network, and storage administrators and stored on the UCS Fabric Interconnect in an object based data store. Cisco UCS Manager discovers UCS devices that are added, moved, or removed from the UCS system. This information, added to the UCS Manager's inventory (a light weight CMDB), is saved on the Fabric Interconnect switch. Cisco reports that a number of their customers are using the XML API to connect UCS Manager to their CMDB (BMC is one of Cisco's partners providing enterprise class CMDB technology) in order to facilitate change tracking. UCS Manager then uses this information when deploying Service Profiles to newly discovered resources. When a Service Profile is deployed UCS Manager configures the server, adapters, fabric extenders, fabric interconnects, NICs, HBAs, LAN, and SAN switches. Service Profiles can also be used to enable Virtual Network Link (VN-Link) capabilities for VN-Link supported hypervisors.

SingleConnect Technology

SingleConnect Technology, an umbrella term, is Cisco's new name for VN link and all VM FEX related functionality. However, the name of each discrete bundled technology remains the same. SingleConnect Technology is the term Cisco uses to describe various approaches to providing physical network type visibility to non-physical or non-directly attached devices. For example, virtual machines, virtual interfaces on a remote interface card, or physical interfaces on a non-switching remote device. SingleConnect benefits include Unified Fabric / Ethernet, multiprotocol storage (FCOE, iSCSI) and improved support for mixed rack and blade form factors.

Hypervisors

Cisco UCS supports the VMware ESX, ESXi, Microsoft Hyper-V, KVM, OVM, and Xen hypervisors. Cisco's implementation of VMware virtualization ESX and ESXi run directly on the UCS system hardware, without additional software, providing hypervisor functionality to host guest operating systems such as Windows or Linux on the physical server.

Partner Focused

Cisco UCS Manager when used in conjunction with partner provided software solutions can stand up, provision, and manage UCS Blade and Rack mount servers as well as the Virtual Machines on those servers. Cisco does not sell system level management and monitoring software instead BMC, EMC, CA, IBM and others provide this functionality. Cisco UCS Manager by itself does not do bare metal operating system installations or application software installations. Cisco UCS provides this functionality through close integration with additional stand-alone software like Altiris, BMC BladeLogic and Cisco's own Cisco Server Provisioner (acquired via the Tidal acquisition).

The XML API for the UCS Manager application can be used by 3rd party management tools. Using this API Data Center Management software from BMC, CA, EMC, and IBM can provision and decommission servers based on demand. Currently, only BMC and EMC use this API to this extent allowing Cisco UCS compute pods or islands of computing to be integrated into a broader Converged Data Center environment. Cisco UCS Manager Ecosystem Partners include BMC, CA, Canonical, Citrix, Compuware, HP, IBM, Infovista, Microsoft, NetApp, Nimsoft, Openstack, Oracle, Redhat, Sciencelogic, SolarWinds, Splunk, Suse, Symantec, VMware, and Zenoss. Those partners offering the tightest level of integration with Cisco' UCS are EMC, BMC, CA, and NetApp.

Storage Connectivity

Cisco UCS Manager preserves traditional Ethernet and Fibre Channel connectivity to LAN and SAN environments North of the Fabric Interconnect. Cisco relies on other vendors like NetApp and EMC for storage and storage management. Cisco provides highly customizable XML APIs that developers and system level software, tools, and utilities providers use to integrate their offerings with Cisco UCS Manager. Storage tools can export the definitions of pools of storage resources through these APIs to the UCS Manager for use in allocating storage resources to UCS servers using the UCS Service Profiling capability.

When running in End Host Mode Cisco UCS Manager allows the use of FCoE to connect a storage array directly to the Fabric Interconnect by pinning a North side Fabric Interconnect port to a VSAN. While operating in End Host Mode it is not possible to perform either LUN (Logical Unit Number) masking nor do normal storage Zoning as would be the case if the Fabric Interconnect Top of Rack switch was operating in Layer 2 switch mode with STP enabled. When configuring the Fabric Interconnect to run as a Layer 2 switch running STP the switch based CPU is utilized to generate, send, and receive BPDUs thereby negligibly increasing messaging latency.

Native Fibre Channel (FC) connectivity is available on the North or upstream side of the Fabric Interconnect (FI) switch with the purchase of an appropriate expansion module. Cisco's current generation FI uses 16 port modules with unified port technology so that individual upstream ports can be

configured as either native FC or Ethernet, providing customers with enhanced flexibility in meeting changing workload requirements. However, there is no native Fibre Channel connectivity available within the Cisco UCS 5108 chassis or within the rack that contains the chassis. Customers currently running native Fibre Channel for SAN connectivity from individual rack mount or blade servers will need to migrate to the converged 10 GB FCoE Cisco fabric. For most customers this should not pose a problem.

End Host Mode

UCS Fabric Interconnects are configured to initially “power up” into what Cisco calls “End-Host Mode” which uses pinning logic to pin a server uplink to a specific VLAN or VSAN virtualized device. However, in those cases where a customer wants a standard layer 2 Spanning Tree Protocol (STP) switch behavior on its up link ports this requires a configuration change after which the FI switch requires rebooting for it to come back up and on-line in standard switch mode running the STP protocol.

End Host Mode was implemented by Cisco in order to remove the requirement to use spanning tree protocol to make routing decisions. By not using STP the FI switch does not have to consume CPU resources required to do STP calculations and send or receive BPDUs.

End Host Mode assigns a link, for example, linking or pinning server A to uplink #1. It is as if the server had a hardwired connection from this specific server port to this specific uplink. It is worth noting that even with the implementation of End Host Mode that the F.I. switch will use MAC address learning and behave as a Layer 2 Switch for local traffic. This means that for one blade to communicate with any other blade in the same chassis or another blade in another chassis in the same physical rack that traffic has to be routed from the first blade up to the Top of Rack F.I. switch and then from that switch it is routed to the second blade. In this way the Cisco UCS architecture inadvertently introduces some additional latency in the case of blade to blade and chassis to chassis messaging which for the majority of workloads will have negligible if any impact.

The current version of Cisco UCS Manager when configured to run in End Host Mode allows the use of FCoE to connect a storage array directly to the Fabric Interconnect. In this configuration Cisco End Host Mode is used to pin a north side Fabric Interconnect port to a VSAN based storage array.

Conclusion and Recommendation

It is HRG's opinion that Cisco UCS is a good fit for many general business workloads. In those cases involving mission critical workloads where reduced transactional latency is a requirement HRG recommends real world benchmarking of customer workloads on a Cisco UCS system prior to purchase. Cisco's Intel centric approach to the Blade market is a simplified, easy to understand, and easy to scale solution.

Harvard Research Group is an information technology market research and consulting company. The company provides highly focused market research and consulting services to vendors and users of computer hardware, software, and services. For more information please contact Harvard Research Group:

Harvard Research Group

Harvard, MA 01451 USA

Tel. (978) 456-3939

e-mail: hrg@hrgresearch.com

<http://www.hrgresearch.com>