

---

# Cisco UCS Architecture and Management: Comparison with Lenovo Flex System and XClarity

February 2016

---

## What You Will Learn

Recently, IBM moved its x86-architecture Flex System solutions over to Lenovo as a response to shrinking profit margins and little market differentiation. Lenovo will market the x86-architecture Flex System paired with its own XClarity management software. IBM will continue to sell IBM POWER systems.

Part of the reason that IBM handed over its x86-based sales is because in less than six years, Cisco has moved IBM out of its market position both in North America and worldwide. Cisco launched the Cisco Unified Computing System™ (Cisco UCS®) in 2009 to directly address the problems introduced by traditional blade server architectures from HP, IBM, and Dell. The promise of greater efficiency through shared power, cooling, networking, and management resources was undermined by these traditional server companies because they created blade servers that reproduced all the complexity of a rack (redundant Ethernet switches, Fibre Channel switches, and management modules) in every blade chassis. Because Cisco did not have to support a traditional blade chassis design, the company was able to create a single unified system that supports both blade and rack servers in a single management domain. Cisco reduced the complexity entailed in the use of separate IP, storage, and management networks by combining all three on a single unified fabric. Cisco® fabric extender technology eliminates the need for blade-chassis-resident and top-of-rack (ToR) switches by eliminating another layer of switching in favor of simple, low-power-consumption devices. The market has responded to Cisco UCS by making it one of the fastest growing technologies in history. Cisco UCS remains an apex of server innovation. This document describes the technology and philosophy behind the Cisco UCS platform that Lenovo and other vendors are trying to emulate.

## Contents

Cisco Philosophy .....	3
<b>Architectural Comparison</b> .....	<b>4</b>
System Resiliency .....	4
Blade-to-Blade (East-West) Traffic Performance .....	6
<b>Management</b> .....	<b>7</b>
Cisco UCS Service Profiles.....	8
Lenovo XClarity Patterns .....	10
Power Management.....	11
Open Management Architecture Versus Vendor Lock-in .....	11
<b>Conclusion</b> .....	<b>11</b>
<b>For More Information</b> .....	<b>12</b>

# Cisco UCS Architecture and Management:

White Paper  
February 2016

## Comparison with Lenovo Flex System and XClarity



### Cisco Philosophy

Cisco Unified Computing System™ (Cisco UCS®) continues Cisco's long history of innovation in delivering integrated systems based on industry standards and using the network as the foundation. From the beginning, Cisco's goal was to radically simplify data centers. The Cisco Nexus® Family of switches, with support for unified fabric and virtualization, began the unified computing phase of the Cisco® Unified Data Center strategy. Cisco then took an integrated approach to computing that unifies computing, networking, virtualization, and storage access resources in a single management domain. This Cisco innovation is at the core of the revolutionary Cisco UCS unified fabric and provides the platform for unified networking, computing, storage access, virtualization, and management.

Until the advent of Cisco UCS, integration of components into systems was a manual, time-consuming, error-prone activity that resulted in higher costs, configuration drift, and slower time to revenue. In creating a single converged system, Cisco used the concept of unified management to create a self-aware, self-integrating system in which every component can be configured through software rather than through manual processes. Cisco UCS service profiles capture every parameter needed to deploy a server (including personality, connectivity, configuration, and firmware), thus allowing servers to be configured for a particular application in minutes. The capability to provision any server in the system for any workload is called hardware-state abstraction, sometimes referred to as stateless computing. Cisco UCS unified management is supported through Cisco UCS Manager, an embedded management system that is available through an intuitive GUI, command-line interface (CLI), and XML API, making it easy for a wide range of ecosystem partners to integrate high-level processes (such as software provisioning) with the system's automated hardware configuration. Cisco UCS runs in the system's fabric interconnects and does not incur any additional licensing costs or the complexity of external management servers.

## Cisco UCS

- Integrated, unified management provides a single point of control for the entire system of up to 160 servers per domain at no cost, and up to 800 servers in multiple domains at no cost.
- Cisco UCS fabric interconnects are built on the standards-based Cisco Nexus 5500 platform.
- Cisco UCS unified fabric connects IP, storage, and management networks through a single set of cables.
- The system integrates with any upstream switching architecture as a single system.

The Cisco approach, in which server identity, connectivity, and settings are dynamically configured to meet application workload requirements, contrasts with the Lenovo Flex System design, which requires customers to purchase, configure, and manage two networking switches (assuming redundancy and the use of a converged network) for every 14 blades. If anything is misconfigured, the time and staff required to repurpose existing blades that are connected to the wrong upstream switches or repurpose an existing blade chassis that has the wrong LAN or SAN networking modules can delay application deployment for days.

Although Lenovo downplays the significance of a network-centric approach, IBM and Lenovo have been relying on third-party chassis-resident networking switches for years. With Lenovo Flex System x86, you must purchase, configure, and maintain up to six modules in a single chassis for local networking and management; up to four networking switches, and two chassis management modules—significantly increasing the number of components and management points in systems.

## Architectural Comparison

Cisco UCS is the only integrated infrastructure that reduces the overall amount of hardware and combines both blade and rack servers on a single unified fabric and management domain (Figure 1). Cisco's approach eliminates management and networking devices in every chassis, reducing the cost of powering, cooling, configuring, managing, monitoring, and maintaining the infrastructure. Cisco UCS places all management functions and configuration information in the fully redundant and highly available Cisco UCS Manager. Cisco SingleConnect Technology, with its wire-once capability, allows customers to scale their data centers easily, quickly, and efficiently without requiring reevaluation of networking infrastructure every time servers are added. With Cisco UCS, the network is established once, with no changes necessary as it scales to 160 servers per domain (and multiple domains with up to 10,000 servers can be managed with Cisco UCS Central Software). By aggregating management and connectivity in the fabric interconnects, every server in the domain is automatically connected northbound to the LAN or SAN without time-consuming and risky reconfiguration at the chassis and server levels.

Lenovo could have redesigned its blade server architecture with the purchase of the Flex Chassis platform from IBM. Instead of unifying networking and management, however, Lenovo replicated this traditional blade chassis architecture, with blades in the front and networking switches and management modules in the back. The minimum redundant configuration of a Flex System requires customers to purchase a pair of networking switches for every chassis. To meet this requirement, customers end up overprovisioning and overpurchasing hardware and port licenses. As a result of this approach, customers must purchase, configure, maintain, power, and cool one switch for every seven blades. Figure 2 illustrates the hardware consolidation and efficiency of a fully redundant Cisco UCS unified fabric compared to the traditional approach.

## System Resiliency

Cisco UCS delivers a level of availability not present in traditional designs. Complementing the system's active-active network configuration, Cisco UCS virtual interface cards (VICs) implement hardware-based fabric failover. This failover provides automatic connection movement (pinning) from one upstream fabric to

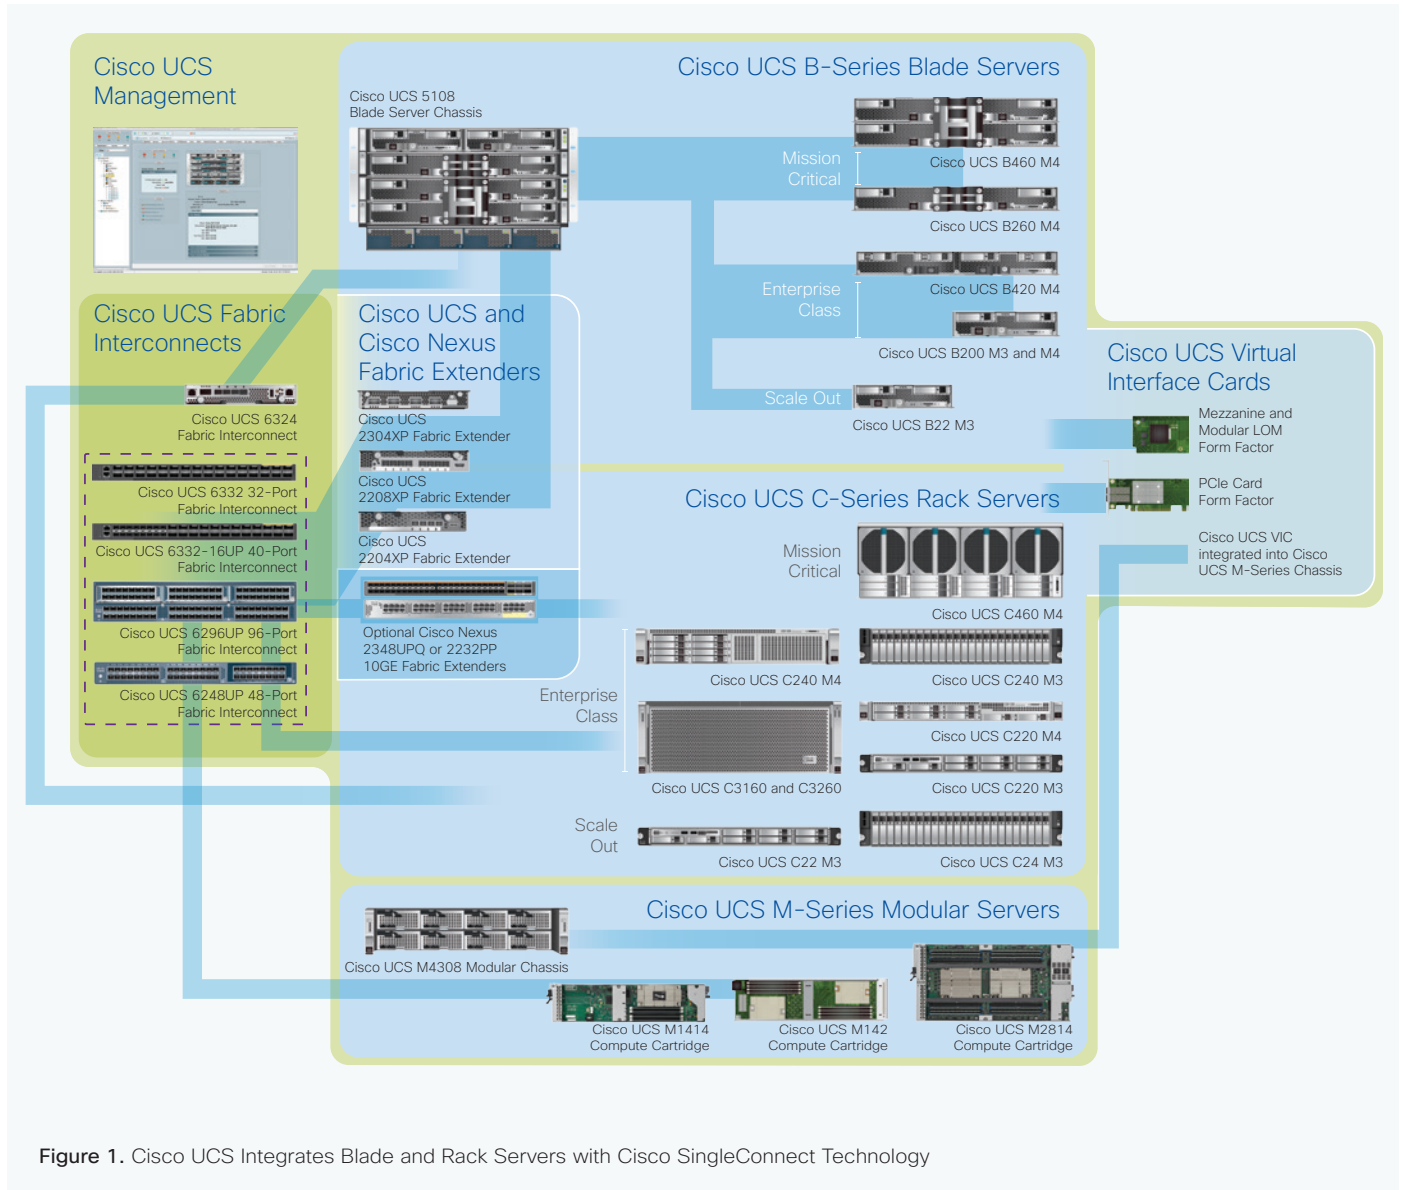
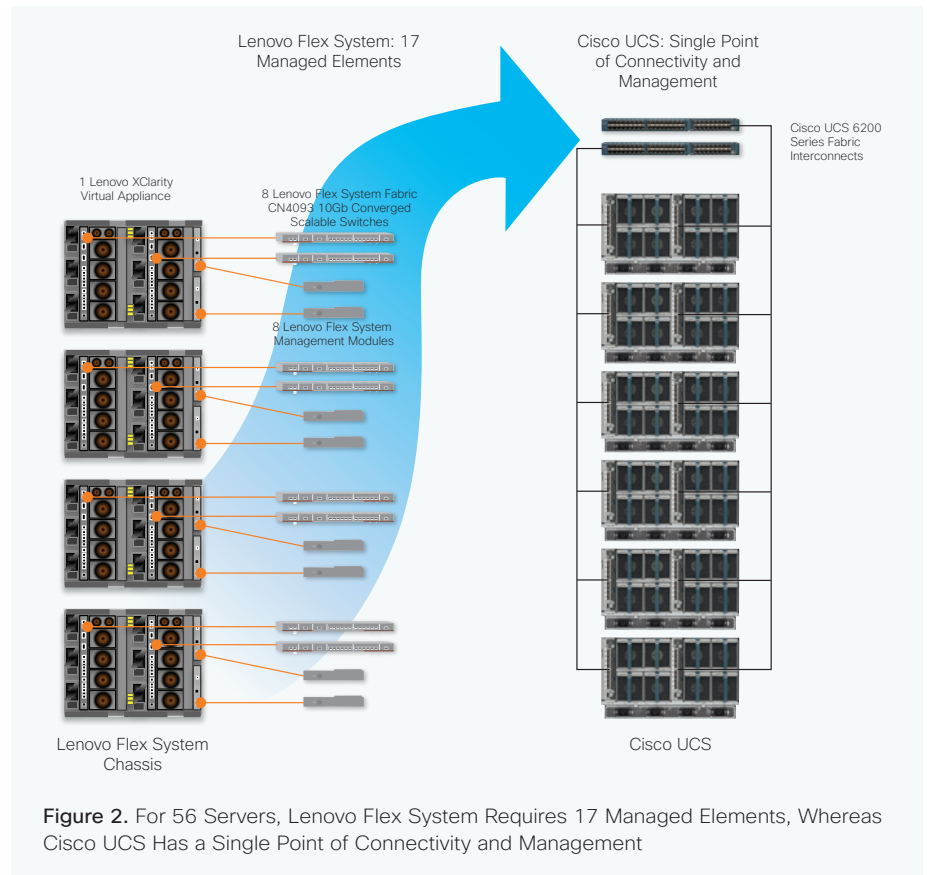


Figure 1. Cisco UCS Integrates Blade and Rack Servers with Cisco SingleConnect Technology

another by synchronizing adapter settings across both fabrics, without the need for OS-based network interface card (NIC) teaming or for any involvement by the operating system. Cisco UCS unified fabric provides an outstanding level of resiliency in comparison to traditional blade server environments.

Lenovo inherited the shortcomings of the IBM BladeCenter chassis. The Lenovo Flex System, basically copies the architecture that HP and Dell have been deploying for many years. It will be interesting to see if Lenovo, traditionally a PC company, will add value to Flex System other than to reduce its capabilities to lower prices.

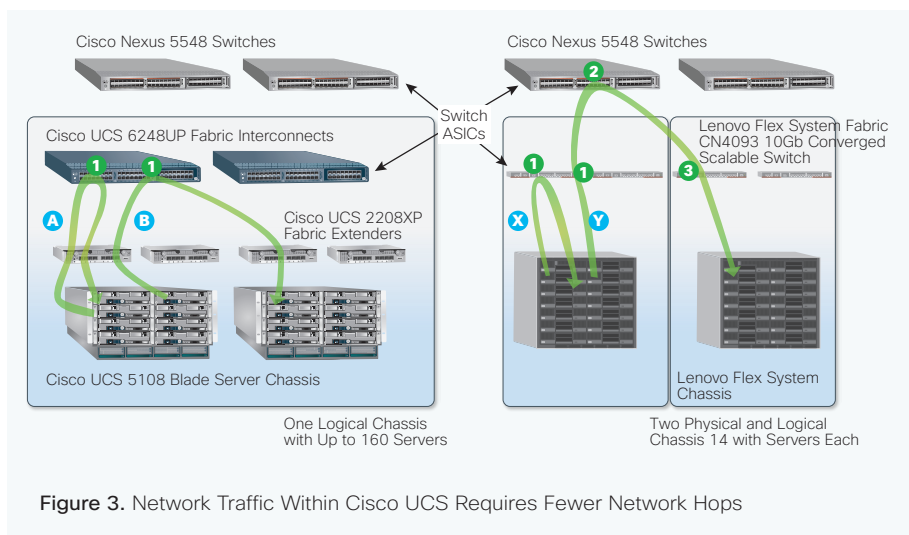


### Blade-to-Blade (East-West) Traffic Performance

IBM in the past has incorrectly stated that Cisco UCS has slower fabric performance than Flex System, without providing any supporting data. Cisco's exhaustive testing demonstrates that Cisco UCS decisively outperforms Lenovo Flex System in both east-west blade traffic latency and blade-to-blade virtual machine migration performance times.

Test results are available in the performance brief [Cisco UCS Outperforms IBM Flex System Blades on East-West Latency at \[http://www.cisco.com/en/US/prod/collateral/ps10265/le\\\_40202\\\_ibmlatencypb-130717.pdf\]\(http://www.cisco.com/en/US/prod/collateral/ps10265/le\_40202\_ibmlatencypb-130717.pdf\)](http://www.cisco.com/en/US/prod/collateral/ps10265/le_40202_ibmlatencypb-130717.pdf).

Cisco tested the Cisco UCS 5108 Blade Server Chassis with Cisco UCS 6200 Series Fabric Interconnects and compared this setup with the Lenovo Flex System with Lenovo Flex System Fabric CN4093 virtual fabric switch modules. Latency measurements were collected for two primary use cases: traffic between two servers within a single chassis and across multiple chassis. Cisco UCS demonstrated significantly lower latency when traffic spanned multiple chassis, and lower server-to-server latency within a single chassis as packet sizes increased. Cisco UCS performed better because of Cisco network innovations embodied in custom application-specific integrated circuits (ASICs), helping ensure only one network hop between any two servers (rack or blade) in the same management domain (see paths A and B in Figure 3). With the Flex System, the two use cases take either one or three network hops, significantly affecting application transaction times depending on where various components are located (paths X and Y in Figure 3).



Cisco also recorded faster virtual machine migration times over a wide array of sizing and load conditions compared to the Lenovo virtual fabric. Cisco tested migration within a single chassis and across multiple chassis and found Lenovo Flex System migration times to be up to 92 percent higher than those for Cisco UCS, with Cisco UCS averaging 22 percent faster migration. This data was collected and averaged from hundreds of controlled test samples with identically configured servers.

In addition to being a leader in enterprise networking, Cisco continues to demonstrate industry-leading performance, with over 100 world-record results on industry-standard benchmarks. More important, the east-west data results demonstrate the exceptional capability of the Cisco Unified Fabric design and the substantial gains that lead to better application performance.

## Management

Cisco UCS Manager provides a single point of connectivity and management for all components in Cisco UCS, including both blade and rack servers. Cisco UCS Manager is embedded in a pair of Cisco UCS 6200 Series Fabric Interconnects in a highly available, active-standby clustered configuration running alongside the system's active-active data paths. This approach provides an important advantage over Lenovo and other traditional architectures: Cisco UCS Manager is a fully redundant management engine that is ready the moment the system receives power—without the need for special clustering software or additional licensing fees.

Cisco UCS uses model-based management in which Cisco UCS Manager discovers and inventories all system components and incorporates them into an object model. Server configuration occurs as a side effect of manipulation of the object model through the Cisco UCS Manager GUI, CLI, or XML API. This approach contrasts with the traditional approach, in which separate element managers configure every component separately and (hopefully) accurately. At best,

traditional approaches support configuration through scripting. In contrast, Cisco UCS Manager orchestrates and automates server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, auditing, and statistics collection processes. Model-based management allows subject-matter experts to define policies for configuring specific types of servers. These policies can be embodied in Cisco UCS service profile templates, which can be used to generate Cisco UCS service profiles that can configure one or hundreds of servers in minutes. Cisco UCS service profiles allow administrators to apply a different service profile to any server, providing extreme flexibility to respond to business workload peaks. This approach is the foundation of stateless computing, in which servers are ready to be configured for any workload on demand rather than needing to be purchased and configured for a specific workload, with significant barriers to repurposing.

A single Cisco UCS management domain consists of up to 160 rack or blade servers and two Cisco fabric interconnects. For multidomain management, Cisco UCS Central Software can manage up to 10,000 servers in local or geographically dispersed data centers (providing global service profiles, statistics aggregation, and aggregated inventory).

Using Cisco UCS Director, you can automate and orchestrate the entire integrated infrastructure, including the deployment and management of Cisco UCS, your choice of industry-leading storage, and your choice of virtualization software. This level of unified management allows you to achieve greater efficiency, agility, and scalability in your server operations while reducing complexity and risk.

Lenovo XClarity is a virtual appliance that runs a software overlay stack on top of a Linux operating system. Lenovo requires node-based or chassis-based software licensing to provide infrastructure systems management. Administrators must choose between two versions of the XClarity software: Administrator or Pro. Each XClarity instance can manage up to 20 chassis or an equivalent number of rack servers (280 total nodes). XClarity is not redundant, so in a failure situation, a new XClarity instance must be manually restored or recovered to bring systems management back online. Additionally, no XClarity aggregation tool currently exists for shared policy management across multiple domains. When the twenty-first chassis is added, an entirely new XClarity domain must be licensed, set up, and configured. For example, a single management domain in an XClarity system of 160 servers requires a nonredundant appliance instance and licensing for each of the 160 devices.

**Table 1** provides a comparison of Cisco UCS and XClarity management features.

	Cisco UCS	Lenovo
No additional licensing fees required	Yes	No
Highly available	Yes	No
A profile can span server generations	Yes	No
A profile can span both rack and blade servers	Yes	No
A profile can span servers with different I/O	Yes	No

### Cisco UCS Service Profiles

A feature unique to Cisco UCS Manager is the use of service profiles to provision and manage Cisco UCS blade and rack servers and configure their I/O properties in a single management domain. Cisco UCS service profiles benefit both virtualized and bare-metal environments when workloads need to be moved from one server to another. Service profiles make it easy to change the hardware resources assigned to a workload, or to take a server offline for maintenance and to substitute another server in its place. Cisco UCS service profiles can be used to increase the mobility of workloads on bare-metal servers. They also can be used in conjunction with virtualization clusters to bring new resources online easily, complementing existing virtual machine mobility.



A Cisco UCS service profile consists of a software definition of the server identity, configuration, and associated network and storage connectivity that the server requires to support a workload. When a service profile is associated with a server, Cisco UCS Manager automatically configures the server, RAID controller, BIOS, I/O adapters, blade settings, firmware (if necessary), and fabric interconnects to match the configuration specified in the service profile (Figure 4). Service profiles improve IT productivity and business agility by defining the server state based on the application workload rather than making the workload fit narrowly defined servers. With service profiles, you can provision infrastructure in minutes instead of days, shifting the IT department's focus from maintenance to strategic initiatives. Cisco UCS service profiles enable preprovisioning of servers, enabling you to configure new servers and associated network and storage access settings even before the servers are physically deployed. The policies coordinate and automate element management at every layer of the hardware stack, including RAID levels, BIOS settings, firmware settings, server identities, adapter settings, virtual LAN (VLAN) and virtual SAN (VSAN) settings, network quality of service (QoS), and data center connectivity.

Cisco UCS service profile templates can be used to simplify the creation of new service profiles, helping ensure consistent policies throughout the system for a given service or application. A service profile is a description of a logical server, and the profile and the physical server have a one-to-one relationship, whereas a service

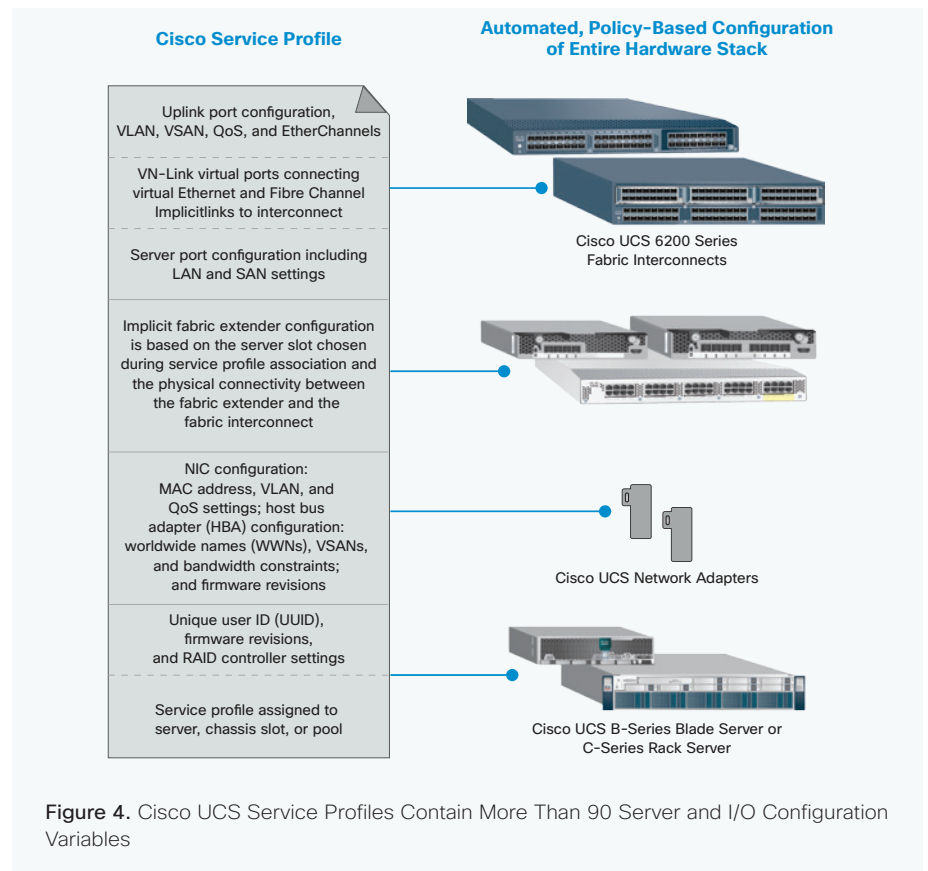


Figure 4. Cisco UCS Service Profiles Contain More Than 90 Server and I/O Configuration Variables

profile template can be used to define multiple servers. The template approach makes it just as easy to configure one server or hundreds of servers with perhaps thousands of virtual machines. This automation reduces the number of manual steps needed, helping reduce the opportunities for human error, improving consistency, and further reducing server and network deployment times. Competitors simply cannot duplicate the power, depth, and breadth of Cisco service profiles and the radical simplicity they bring to infrastructure management.

### Lenovo XClarity Patterns

Lenovo responded to the advancements embodied in Cisco UCS service profiles by implementing Lenovo Patterns for configuration of chassis and servers to prepare them for deployment. Patterns fall short of service profiles because of their limited capabilities and cumbersome real-world usability. Patterns script common configuration commands and impose a variety of prerequisites and requirements on the environment. For example, whereas a specific service profile can be applied to any server that meets the requirements of the service profile, Lenovo Patterns can be applied only to servers of the same type (blade or rack, similar I/O adapters and generations, etc.). In the end, XClarity users are left with software that merely scripts common commands in a “fire and forget” fashion. Unlike the model-based Cisco UCS Manager, scripted software has difficulty adapting to changing hardware environments, incorporating dynamic firmware into profiles, and maintaining statelessness.

XClarity has a number of limitations:

- **Firmware enablement is a manual process**, not tied to the Pattern itself. Firmware updates can be applied in XClarity, but separately from the pattern.
- **Only Lenovo Flex System blades and Lenovo System X M5 rack servers are supported.** IBM BladeCenter and System P and Z and older System X servers are not supported.
- **Scope is limited.** Lenovo Patterns do not adapt to the server and apply only to the same hardware and model type (for example, a Lenovo x240 blade pattern can be applied only to another Lenovo x240 blade).
- **XClarity can be costly and complex.** Lenovo Patterns require an XClarity software license for every managed chassis or server node.
- **No high-availability capability is available.**

Cisco invented Cisco UCS service profiles and built the entire system on the concept of state abstraction. Whether customers are running a bare-metal or virtualized environment or any combination of the two, they can gain the advantages of Cisco UCS service profiles. These service profiles have revolutionized computing, and competitors are challenged to try to replicate the increased productivity that automated configuration provides. Cisco’s approach has been so successful because every element of the system was designed from the beginning to have its configuration set through software. Cisco service profiles do not require licensing, because they are a feature of Cisco UCS Manager. Tens of thousands of satisfied Cisco UCS customers around the world enjoy the benefits that Cisco service profiles bring to a variety of production environments.

### Power Management

Lenovo XClarity Administrator does not support the capability to set or manage power caps. Power-related features are limited to monitoring and basic power management: power on and off the host and operating system. All other actions require the user to connect to the Lenovo Chassis Management Module (CMM) for the chassis and Integrated Management Module (IMM) for the server. XClarity allows a user to launch a CMM web user interface to manage power caps. An IMM2 standard features-on-demand (FoD) license is required to enable power capping.

Unlike Cisco UCS, Lenovo does not use power groups. Cisco UCS power groups allow users to define a group mapped to a power distribution circuit, with power balanced among all the resources in the group. With this approach, underused resources can share their power allocations with other resources in the group. To get the same capability, Lenovo users must add another management layer, such as Tivoli Endpoint Manager or Cisco EnergyWise® software.

### Open Management Architecture Versus Vendor Lock-in

Cisco UCS Manager provides centralized management capabilities that serve as the system's central nervous system. Cisco UCS Manager provides flexible role- and policy-based management. Cisco UCS Manager maintains a model of the system that forms the single source of truth about all connected components. This model can easily be exported to configuration management databases (CMDBs) for use in ITIL processes.

Cisco UCS Manager provides system visibility to higher-level systems management and lifecycle tools from independent software vendors (ISVs), including Microsoft, BMC, CA, HP, and IBM. ISVs and in-house developers can use the Cisco UCS XML API to further customize and automate the system according to their unique requirements by using the Cisco goUCS Developer's Toolkit. In addition to supporting standards including Simple Network Management Protocol (SNMP) and Intelligent Platform Management Interface (IPMI), the Cisco UCS XML API enables secure export of all Cisco UCS Manager commands and data through the API.

Currently Lenovo XClarity has no cloud software stack and cannot provide a single-source solution like Cisco UCS Director.

### Conclusion

Cisco UCS represents an advanced architecture based on Cisco innovation that uses a high-bandwidth, low-latency unified fabric for network and storage connectivity and server management. By using fabric interconnects as the single point of connectivity and management for the entire system, Cisco developed a single converged system that scales easily and efficiently while providing a superior level of fabric redundancy not found in Lenovo or other traditional architectures.

Lenovo Flex System has no single management convergence point and requires every 14-blade enclosure to be separately cabled, configured, and maintained. Although Lenovo markets Lenovo XClarity virtual appliance as a single point of management, in reality it provides no redundancy; it is the result of bundling several software tools together, and it offers only limited scripting capabilities for

For more information about the Cisco UCS Manager XML API please visit <http://developer.cisco.com/web/unifiedcomputing/goucs>.

rudimentary automation. The management of all the components that make up a solution is not integrated, which increases operating costs and TCO.

In contrast, Cisco UCS Manager is embedded in a fully redundant fabric configuration out of the box. Cisco UCS service profiles provide capabilities to the entire infrastructure, including robust features such as adaptive policies (rack or blade), firmware control, and BIOS and bandwidth policies for adapters. By making Cisco UCS Manager functions available through an XML API, Cisco introduced programmatic control over the entire Cisco UCS infrastructure, with Cisco UCS Central Software and Cisco UCS Director, enabling organizations to better automate data center setup, provisioning, and management—starting at the bare-metal system and OS and extending all the way to the applications and even the cloud for both physical and virtual environments. Designed for simplified management, Cisco UCS comes with robust manageability tools that help your team deliver a positive impact on your business through the data center.

## For More Information

- For more information about Cisco UCS, visit <http://www.cisco.com/go/ucs>.
- For more information about Cisco UCS Manager, visit <http://www.cisco.com/en/US/products/ps10281/index.html>.
- For more information about Cisco UCS Central Software, visit <http://www.cisco.com/en/US/products/ps12502/index.html>.
- For more information about the VersaStack solution combining the innovation of Cisco UCS Integrated Infrastructure with the efficiency of the IBM Storwize storage system, visit <http://www.cisco.com/go/versastack>
- For more information about Cisco UCS latency and virtual machine migration performance compared to Lenovo Flex System, please contact your Cisco sales representative, call 1-866-428-9596 (United States or Canada), or visit <http://www.cisco.com/web/ordering/root/index.html>.



**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

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

**Europe Headquarters**  
Cisco Systems International 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](http://www.cisco.com/go/offices).