White Paper

Cisco UCS Architecture and Management: Comparison with IBM Flex System

Contents

Cisco UCS Architecture and Management: Comparison with IBM Flex System

White Paper
September 2013

What You Will Learn

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 better efficiency through shared power, cooling, networking, and management was undermined by 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 by 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, and this document describes the technology and philosophy behind the Cisco UCS that IBM and other vendors are trying to emulate.

Cisco Philosophy

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.

The Cisco approach, in which server identity, connectivity, and settings are dynamically configured to meet application workload requirements, contrasts with the IBM 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 IBM downplays the significance of a network-centric approach, the company has been relying on third-party chassis-resident networking switches for years. IBM requires customers to purchase, configure, and maintain up to seven modules in a single chassis for local networking and management: up to four networking switches, two chassis management modules, and one IBM Flex System Manager (FSM), significantly increasing the number of components and management points in its systems.

This approach aligns with a business model that includes fee-based management software and fee-based services. IBM PureSystems solutions encumber the customer with ongoing management software licensing, while at the same time IBM promotes costly IBM Power Systems software and services contracts. According to IBM's public 2012 annual report, 80.6 percent of total revenue for its fiscal year 2012 came from IBM's Software and Global Services businesses. IBM profits significantly through this type of software and services revenue in the form of continuous customer payments and annuities. IBM's 2012 annual report notes that:

> "Approximately two-thirds of external software segment revenue is annuity based, coming from recurring license charges and ongoing post-contract support.
>
> "Approximately 60 percent of external Global Services segment revenue is annuity based, coming primarily from outsourcing and maintenance arrangements."

Source: http://www.ibm.com/annualreport/2012/bin/assets/2012_ibm_annual.pdf

Architectural Comparison

Cisco UCS is the only converged system that reduces 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 Unified Fabric, with its wire-once capability, allows customers to scale their data centers easily, quickly, and efficiently without requiring a 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 (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.

IBM could have redesigned its blade server architecture with the IBM Flex Chassis platform. Instead of unifying networking and management, however, IBM replicated its traditional blade chassis architecture with blades in the front and networking switches and management modules in the back. IBM's minimum redundant

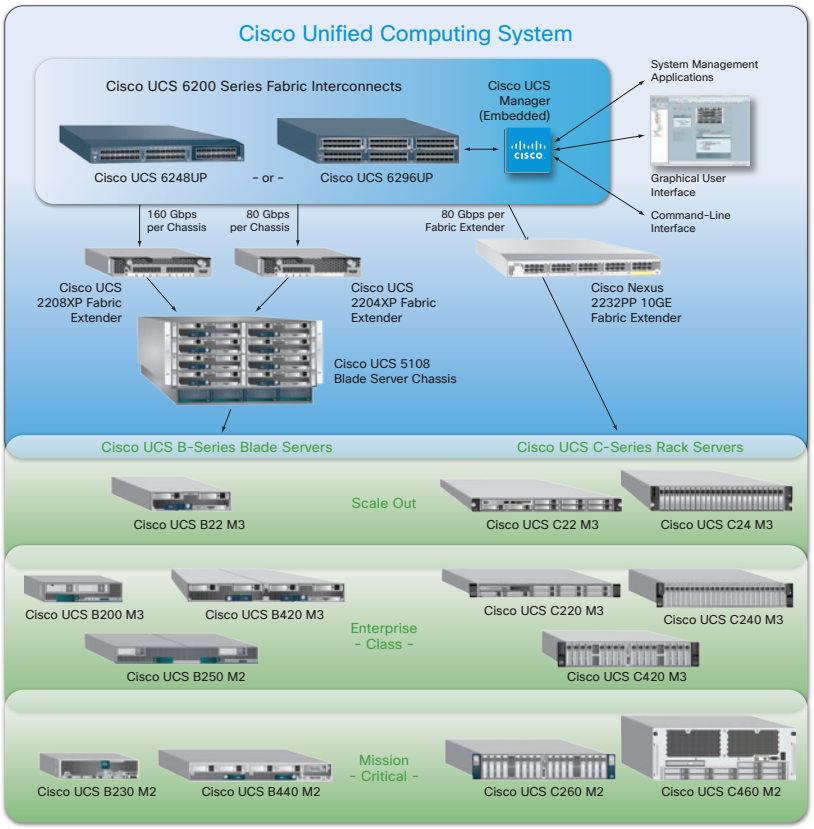

Figure 1. Cisco UCS Integrates Blade and Rack Servers with a Unified Fabric

configuration 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 implemented by IBM.

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.

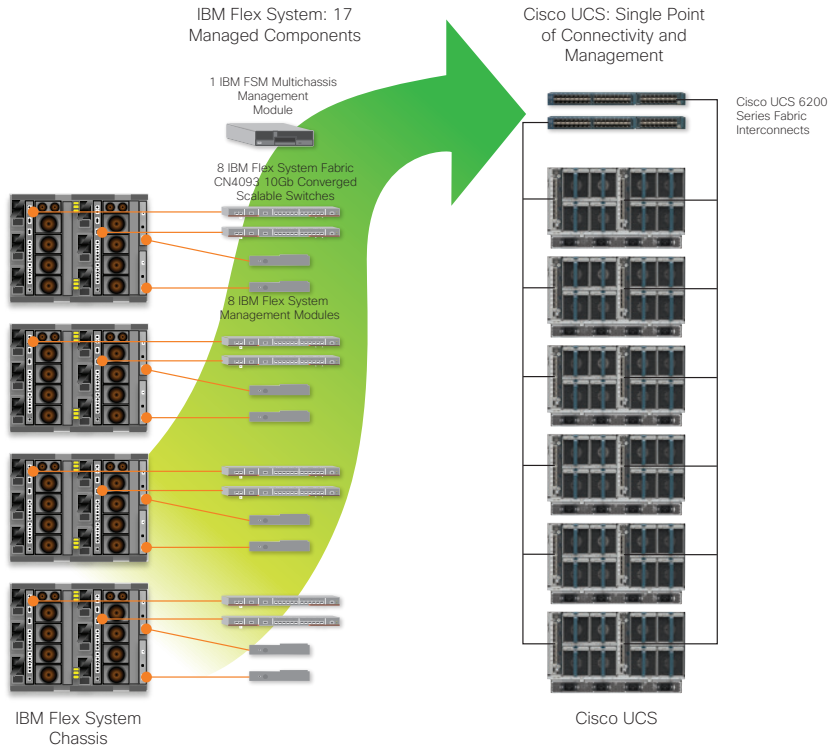

Figure 2. For 56 Servers, IBM Flex System Requires 17 Managed Components, Whereas Cisco UCS Has a Single Point of Connectivity and Management

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 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.

IBM also demonstrates the shortcomings its previous IBM BladeCenter chassis by undoing all the mainframe-inspired hardware from the previous generation. In IBM

Flex System, IBM no longer uses blowers, instead using off-the-shelf fans. IBM no longer endorses an active backplane. IBM no longer defends compartmentalized cooling as superior, opting instead for a dual-zone cooling design. With IBM Flex System, IBM basically copied an architecture that HP and Dell have been deploying for many years.

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

IBM has incorrectly stated that Cisco UCS has greater latency and slower fabric performance than its solution without providing any supporting data. Cisco's exhaustive testing demonstrates that Cisco UCS decisively outperforms IBM 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.

Cisco tested the Cisco UCS 5108 Blade Server Chassis with Cisco UCS 6200 Series Fabric Interconnects and compared this setup with the IBM Flex System with IBM 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 demonstrates significantly lower latency when traffic spans multiple chassis, and lower server-to-server latency within a single chassis as packet sizes increase. 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 IBM 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 IBM virtual fabric. Cisco tested migration within a single chassis and across multiple chassis and found IBM migration times to be up to 92 percent higher than those for Cisco UCS, with Cisco UCS averaging 22 percent faster. 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 70 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 IBM 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.

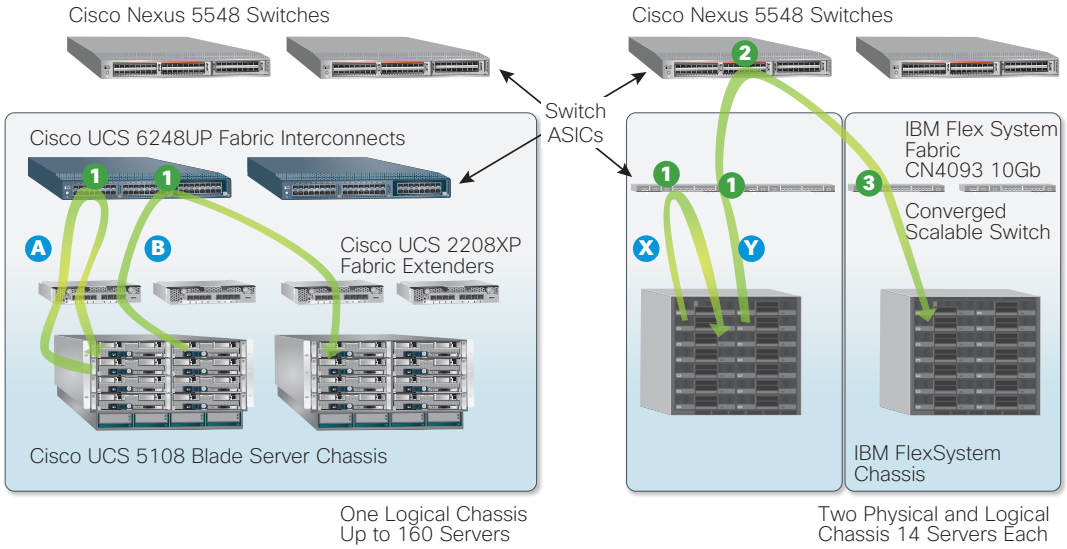

Figure 3. Network Traffic Within Cisco UCS Requires Fewer Network Hops

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). By enabling automation of processes, Cisco UCS unified management allows data center managers to achieve greater efficiency, agility, and scalability in their server operations while reducing complexity and risk.

IBM continues to use a model in which customers license management software and deploy a separate blade server for multichassis management. IBM FSM is a locked appliance running a proprietary grouping of IBM director-inspired software on the Linux operating system. In addition to requiring the purchase of IBM FSM, IBM's licensing model requires standard or advanced software licensing for every IBM Flex System chassis to implement shared management. Each IBM FSM has the capability of managing up to 16 IBM Flex System chassis and is not redundant (as of the date of this document). Therefore, in a failure situation, a new IBM FSM must be manually replaced and configured to restore the system. Additionally, no IBM FSM aggregation tool currently exists today for shared policy management across multiple domains. When customers add a seventeenth IBM chassis or move to a different location, an entirely new IBM FSM domain must be purchased. Cisco UCS has up to 160 servers in a single management domain, with no slots used for management systems or software licensing requirements. Up to 10,000 servers can be managed across multiple Cisco UCS domains through Cisco UCS Central Software. For instance, a single management domain in an IBM Flex System of 160 servers requires a nonredundant IBM FSM system and IBM FSM software licensing for each chassis, for a total cost of US$163,799. The cost for each chassis in a 56-server example is illustrated in Figure 4.

IBM Flex System Management

- No aggregated IBM FSM domain management
- No redundant IBM FSM management option
- Each domain requires one IBM FSM blade (US$12,599) and licensing for every chassis (3-year cost of US$12,600)
- Example: Total cost of US$62,999 for IBM FSM management and licensing for four chassis

(Prices are list prices obtained from ibm.com as of September 17, 2013.)

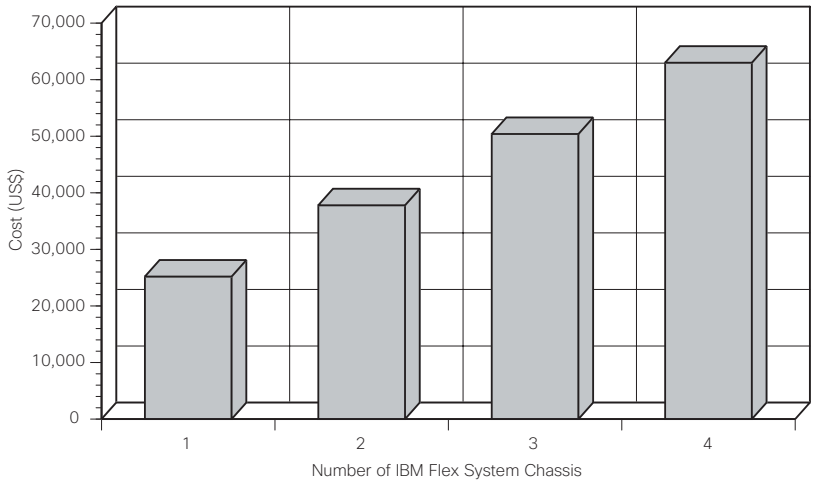

Figure 4. IBM FSM Server Plus 3-Year Software Licensing Costs

Cisco UCS Service Profiles

A feature unique to Cisco UCS Manager is the use of Cisco UCS 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. Cisco UCS 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 5). Cisco UCS 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 Cisco service profiles, infrastructure can be provisioned in minutes instead of days, shifting the IT department's focus from maintenance to strategic initiatives. Cisco UCS service profiles enable preprovisioning of servers, making it possible 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.

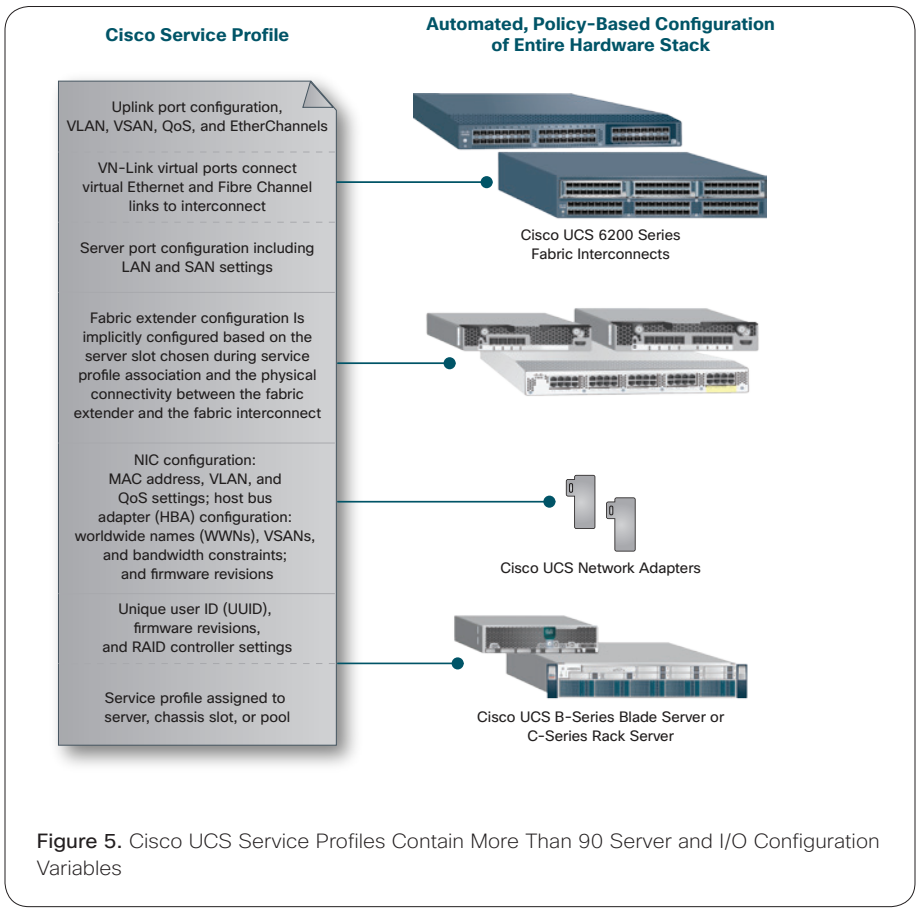

Figure 5. Cisco UCS Service Profiles Contain More Than 90 Server and I/O Configuration Variables

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 there is a one-to-one relationship between the profile and the physical server, whereas a service 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.

IBM Patterns

IBM responded to the advancements embodied in Cisco UCS service profiles by implementing IBM Patterns for configuration and server deployment. IBM Patterns fall short of Cisco UCS service profiles because of their limited capabilities and cumbersome real-world usability. IBM Patterns script common configuration commands and impose a variety of prerequisites and requirements on the environment. For example, whereas a specific Cisco UCS service profile can be applied to any server meeting the requirements of the service profile, IBM Patterns can be applied only to the same server type. IBM environments must use additional software to manage firmware. This restriction is representative of the limitations of a separate management system that is not unified; customers 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.

Limitations include:

- No firmware management (separate web interface): There is no single point of hardware abstraction.
- Support for IBM Flex System blades only: IBM BladeCenter and System X, P, and Z are not supported.
- Limited scope: IBM Patterns do not adapt to the server and apply only to the same hardware and model type (for example, an IBM x240 blade pattern can be applied only to another IBM x240 blade).
- Cost and complexity: IBM Patterns requires an IBM FSM appliance and software licensing for every managed chassis.
- No high-availability capability available today: A second IBM FSM is required for clustering.

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. Cisco UCS 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.

Open Management Architecture Compared to 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.

For more information about the Cisco goUCS Automation Tool please visit http://developer.cisco.com/web/unifiedcomputing/goucs

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.

Imagine developing a cloud orchestration layer that has to interface with only a single API to provision and manage all computing elements. This capability is far beyond the reach of IBM's single-vendor approach, which requires customers to purchase additional IBM SmartCloud software and IBM services to support hybrid cloud environments. Cisco offers an open management interface that increases choice, including cloud and management tools developed by Cisco and third-party, open source, and user-developed ecosystems.

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 IBM or other traditional architectures.

IBM Flex System has no single management convergence point and requires every 14-blade enclosure to be separately cabled, configured, and maintained. Although IBM markets the IBM FSM 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 rudimentary automation. IBM requires customers to purchase management hardware, software licensing, and support for every chassis.

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 that include 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, making it possible 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, please visit http://www.cisco.com/go/ucs.
- For more information about Cisco UCS Manager, please visit http://www.cisco.com/en/US/products/ps10281/index.html.
- For more information about Cisco UCS Central Software, please visit http://www.cisco.com/en/US/products/ps12502/index.html.
- For more information about Cisco UCS latency and virtual machine migration performance compared to IBM 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**.

LE-40501-00 09/13