WHITE PAPER

Introduction to Cisco UCS & Converged Infrastructure

Understanding the concepts of stateless hardware and converged infrastructure for internal personnel without direct hardware experience

Contents

Introduction

This guide intends to introduce Pure employees to the basics of converged infrastructure and the Cisco Unified Computing System (UCS) platform.

The Reasons for Converged Infrastructure

When customers are planning for hardware to purchase for their datacenter, they will be looking at the following three architecture choices:

- Composable infrastructure - think of this as fully DIY, where all components must be acquired and configured independently to work together (whether by customer or partner)
- Converged infrastructure - integrated and connected components, which provide a more "building block" hardware approach, tested and documented by the vendors whose products comprise the solution
- Hyper-converged infrastructure - specific pre-built components including compute/memory/storage, requiring that all resources scale in a linear fashion and each component consumes shared resources to run some "controller" software

The decision for utilizing any of these may have many different factors and implications for the customers, including things like:

Relevant to Business Decision Maker (BDM):

- Costs can vary widely moving across the 3 types of architectures, with some being more expensive up front, and others being more expensive over the life of the platform
- Varying power, cooling & network requirements which can be specific to vendor or model of the components used
- The requirements of each architecture choice fall to the customer to determine and manage for the lifecycle of the equipment
- There can be a significant amount of time required for the deployment of equipment to bring it into production
- There is an additional period of time where re-architecture or replacement of equipment overlaps and requires additional management and infrastructure resources

Relevant to Technical Decision Maker (TDM):

- Number of management interfaces can be overwhelming, up to single management interface per piece of equipment
- Number of cables for power & network can quickly get to be a management issue within racks
- Interoperability of all software versions can range from completely customer managed & supported to being strictly defined by vendors
- Potentially disrupting to production if incompatible versions are used, support can be denied if customers do not adhere to guidelines or install/upgrade to non-tested versions

Within composable infrastructure, most customers tried to decrease overhead through a scale-up approach - meaning that they increase the capacity of a server by increasing the available CPU, memory, and IO devices.

If a customer is not completely virtualized, or they do not want to be limited to fewer endpoints which are most costly to handle extra components, then they adopt a scale-out approach, which aligns more with converged & hyperconverged infrastructures.

The desires to ease the management overhead, reduce power/cooling/cost requirements, and increase virtualization & simplification of infrastructure components are some of the factors that really drove development of both converged & hyperconverged infrastructure.

Converged infrastructure moves us away from the traditional standalone servers of the past, and most vendors tend to favor blade servers as the core of this configuration.

Short history & Strengths of the Cisco UCS platform

Utilizing blade servers gives customers the ability to gain much more utilization within a virtualized environment, or more easily re-purpose these compute resources in an environment which also runs bare-metal instances.

Within a converged infrastructure environment leveraging a scale-out model for compute resources, it is easier to incorporate shared data storage capable of large capacity and high IO performance for mixed workloads. Converged infrastructure can also consume shared storage while not drastically increasing the number of physical connections to allow connectivity to every compute resource.

Many blade server platforms had significant limitations around the inability to host PCI cards, the requirement for "mezzanine" cards in proprietary form factors to provide networking via basic built-in network connections, and very basic platform management controllers built with closed access user interfaces.

The Cisco Unified Computing System, which was originally known as the codename of "Project California," was built to reduce or eliminate the limitations in server deployments.

Cisco decided that the best place to start building a scalable data center architecture was with the natural aggregation point that exists: the network.

Cisco's original definition for Unified Computing was:

> *"Unified Computing unifies network virtualization, storage virtualization, and server virtualization into one, within open industry standard technologies and with the network as the platform."*

Cisco UCS Architecture Components

Fabric Interconnects

The components that make up a UCS environment start at the top-level with the fabric interconnects, also known as FIs (pronounced as "F-I"). These fabric interconnects serve as the physical connection point for the entire UCS architecture, including the servers and the unified fabric – encapsulating both ethernet and fiber channel networks. The fabric interconnects are the central connectivity point where all servers, both blades & rack mount, are connected together, and from these fabric interconnects there will be uplinks to the data center networks, both ethernet and fiber channel.

The FIs are run in a highly available pair that provides redundancy for the management functions and an active-active data plane for the network connectivity. The connectivity provided by the FI ports has evolved from a maximum of 10Gig Ethernet and 8GB fiber channel at release in 2009, to now support up to 100Gig Ethernet and 32GB fiber channel on the more recent hardware. The FIs are based on the same switching technology as the Cisco Nexus platform, but they run a different operating system.

UCS Manager & Intersight

The specific version of the NX-OS operating system which runs on these fabric interconnect management switches hosts UCS Manager, which serves as the logical controller for the UCS architecture. Cisco UCS was designed to provide a single point of management for multiple chassis or rack-mount servers which are controlled through its internal data management engine. This engine controls all configuration and state data for the servers & components and provides access to control the platform through external interfaces including GUI, CLI and API.

UCS is also built around the concept of stateless hardware, where the entire logical configuration of a server is separated from the hardware. The complete server identity includes things like BIOS/boot configuration, storage & network connectivity through virtual interfaces, and even firmware versions; these configurations and interfaces are all applied to the hardware when the server profile is associated/assigned to it.

This management software uses managed object model to provide the ability for administrators to control all the UCS hardware through pools, policies/templates and an organizational structure. Pools allow for a standardized identifier for things like UUID, MAC addresses, and WWNN/WWPNs. Policies allow for settings to be specified in a way that they will remain consistent, and these policies can be utilized by templates which define all settings for any profile which is derived from the template.

The organization structure provided by Cisco UCS is very beneficial for large-scale server administration, as it allows for the use of defined templates for all use cases, where any sub-organizations (think of these like folders) allow for overriding of pools of identifiers along with policy settings. In addition, role-based access control can also be assigned to the organizational structure, which allows for access being granted to administrators to control their own resources & configurations, without overriding the higher-level templates & settings.

The most recent change to the management for Cisco UCS comes in the form of Cisco Intersight. Intersight is a cloud-based SaaS platform which elevates from managing all templates, policies and profiles within the maximum of 160 servers under single pair of FIs to being able to define the same components but apply them across any number of UCS environments. In addition, Intersight also allows for additional integrations with many third-party producs, such as Pure Storage. Pure produced the first third-party integration for the Intersight platform, and more details can be found in this (insert link) document focused on Intersight along with the new X-Series platform.

Servers – Chassis/Blades or Rack-mount

When we get down to the compute resources themselves, there are 2 options for form factor for consideration. Most customers who run Cisco UCS will prefer blade servers, which are known as the B-Series, but there is also a strong use case for the rack-mount C-Series servers as well.

Within UCS, the base component for blade servers is the 5108 chassis. This is a 6U chassis which has 8 blade slots in a two horizontal x four vertical slot layout. Each chassis has redundant power supplies and can host two fabric extenders which provide connectivity from the fabric interconnects down into the chassis backplane, which connects to each blade itself.

For the blade servers, there are options for dual-processor models which fill a single blade slot, all the way up to quad-processor models which take up 4 blade slots. Each blade option is capable of housing multiple terabytes of memory, along with capabilities for Intel Optane persistent memory, and even GPUs.

For a more traditional hardware model, some customers utilize the C-series rack-mount servers within a UCS environment. These are 1RU, 2RU or 4RU servers which can also be leveraged as standalone servers, but primarily are used where there is a need for additional PCIe slots not found in the blade servers and for some specific applications where heavy GPU acceleration is needed for the workloads.

Each server, whether blade or rack-mount, follows a numbering scheme that tells you the server's generation. The server model number ends with "B," followed by a number that denotes the server hardware generation. For example, a B200M4 blade is a fourth generation B-series blade server, whereas C220M5 denotes a fifth-generation rack-mount C-series server.

Virtual Interface Cards (VIC)

For all UCS managed servers, the network connectivity at the host level is provided by the virtual interface cards. These cards are converged network adapters, which means that they can present both ethernet (NIC) and fiber channel (HBA) connectivity to the hosts where they are installed, along with all management traffic.

These cards are powered by Cisco ASICs, which means that they can provide advanced network functionality, even to the virtualized interfaces the expose to the hosts, while still having reduced needs for power consumption and space.

The only significant difference to be aware of with the VICs between blade and rack-mount servers is that you will only plug cables directly into the rack-mount servers. Blade servers connect to the internal chassis backplane– meaning you plug cables into the chassis through the fabric extender, not each blade.

Overall, any hardware component within a UCS environment, from the blades or rack mount servers to the fabric extenders providing network connectivity to chassis, to even the fabric interconnects running the environment themselves can all be expanded or upgraded with no downtime (minus the case of taking servers offline naturally). This same statement holds true for software upgrades to add in new features and support for newer server models, as well as the hardware within the UCS architecture.

Customers can add capacity with new chassis/blades or rack-mount servers. In addition, they can upgrade network connectivity by replacing the fabric extenders. Finally, they could upgrade the system to new capabilities and uplink network speeds by replacing the fabric interconnects. All of these changes are possible without taking the entire platform offline – the platform only has decreased redundancy during the FI upgrade.

Benefits of converged infrastructure with Pure Storage

The real value to many customers is not with only the Cisco UCS architecture itself, as blade servers are severely limited on storage capabilities, having only two hard drive slots per single chassis slot blades. Blade servers in general are a better solution for customers when leveraging boot from SAN, and this is even more the case with the logic behind the "stateless computing" model of Cisco UCS.

Since the entire identity of a server profile can be associated with any blade within the UCS domain (if it has compatible hardware), this means that the virtual HBAs presented to the host, the boot from SAN targets, and the installed operating system can effectively be migrated throughout a UCS environment.

To associate or dis-associate a server profile from a blade server takes about 20 minutes, but that is not an unacceptable amount of time to transfer the entire identity of a server from one blade to another. This method ends up being the quickest & easiest way to upgrade the capabilities of a server in many cases, when the identity can be moved from older to newer server hardware with only a minor downtime to move the identity. Many customers perform quick upgrades of servers this way, and then the older hardware can be utilized as test/dev systems – and this is all with only needing to slot another blade.

However, this functionality all relies on having storage which lives outside of the servers themselves. As a result, there has been a significant market for converged infrastructure solutions where storage vendors have partnered with Cisco. In these instances, each has jointly developed a CI platform – these are led by Pure Storage & NetApp as the market leaders, but also include storage arrays from EMC, IBM, and Hitachi.

Based on limitations of the underlying storage, or fear of support issues to vendors due to unvalidated or untested versions of both software & hardware running within the systems, led many of the converged infrastructure solutions to be centered around an RCM - a Release Certification Matrix.

This means that everything from the hardware – specific server models, ethernet & fiber channel switches – to the software – including the driver firmware on the servers, software versions on the ethernet & fiber channel switches, the software running on the FIs, and even the specific versions of hypervisors or operating systems installed on the servers – are all restricted to a limited list of versions that had been certified within your specific platform.

When it came time for you to try and upgrade your platform, primarily driven by a new version of a hypervisor or operating system due to a business need for new features or capabilities, you were required to engage with vendor support to initiate a project for your upgrade. With most converged infrastructure platforms, part of the benefit or curse of the entire platform is that you have vendor support for the platform as unit, but this means that you are required to follow their restrictions and requirements - or you risk losing your support if things go wrong.

For anyone who did not have to deal with the pain of lifecycle management and a somewhat long-term project of upgrades run through a new RCM version, you should be thankful. Many times, getting from initial request to vendor support to having your system upgraded could take weeks; I had many platform upgrades that took up to 6 weeks of near constant engagement from the customer side.

When we talk about converged infrastructure with Pure Storage, meaning our FlashStack platform, customers get simplicity and flexibility, not a rigid set of rules they must comply with to maintain vendor support. As a result, they get the best performance while allowing them to be concerned with operations focused on their workloads, not on the underlying infrastructure.

Conclusion

Converged infrastructure is about trying to help balance the needs of customers to scale their resources, manage them easily, and not face the risk of building everything in a DIY fashion. There are factors and costs associated with deciding on the type of data center architecture, and these will differ slightly for business or technical decision-makers. The Cisco UCS platform addressed the complexity of a modular data center infrastructure that could scale and reduce or eliminate limitations in server deployments; this architecture was built upon the natural aggregation point – the network. This converged infrastructure platform is comprised of many components which collectively allow for a stateless software-driven architecture, which leverages policies and templates to simplify the management of hardware that scales to meet any needs from the smallest to largest customers.

In my opinion, when we combine the Cisco UCS platform with Pure Storage, we now have the best data center platform available to customers. That stamenent is based on my 15 years of building and managing data centers and colo facilities, and architecting massive hardware footprints covering composable (DIY) environments, to every converged platform on the market, and some of the leading hyperconverged options. After all my experience in the field with building and managing those options daily for years, there is only one platform that I believed in enough to come work on – and that is FlashStack.

Reach out to me for any questions or to find additional information, visit the FlashStack page under the Pure Portfolio on Sales Central.

About the Author

Joe Houghes is a Senior Solutions Architect in the Portfolio Solutions team within Pure Storage, focused on solutions on the FlashStack platform and automation and integration. He has experience from over 20 years in Information Technology across various customer/vendor organizations with architecture and operations. His expertise covers computing, networking, storage, virtualization, business continuity and disaster recovery, cloud computing technologies, and automation and integration across many applications and vendor platforms.

Pure Storage, Inc.
650 Castro Street, #400
Mountain View, CA 94041

[Publication number and date here]