DACETEK

FlexPod for Remote Office and Data Center

Design Guide

Version 1.0
Published: 6th May 2024

Table of Contents

Introduction

FlexPod is a converged infrastructure solution which is jointly developed by Cisco and NetApp. A converged infrastructure essentially combines compute, storage, and networking resources into a single, pre-configured system. This simplifies deployment, management, and potentially reduces costs. These systems are highly flexible and suitable for small-medium businesses and large enterprise data centers. It enables a flexible infrastructure solution that encompasses many different components and software versions.

FlexPod has been offered in two options:

- **FlexPod Datacenter:** This is a comprehensive converged infrastructure solution designed for large-scale deployments and demanding workloads. It offers a high degree of scalability and flexibility, allowing you to add or remove resources as needed. FlexPod Datacenter is ideal for organizations with complex IT environments and significant data storage requirements.
- **FlexPod Express:** This is a more streamlined and cost-effective solution targeted towards smaller and mid-sized businesses. FlexPod Express comes in pre-defined configurations, making it easier to deploy and manage, especially for organizations with limited IT staff. It prioritizes affordability and simplicity over the ultimate level of scalability offered by FlexPod Datacenter

In FlexPod, compute and network components are provided by Cisco and storage by NetApp. This modular approach enables independent deployment, and scaling of individual stack and thus provides better ROI.

Some of the benefits of FlexPod are:

- Highly available and scalable platform with flexible architecture that supports various deployment models.
- Can run virtualized, non-virtualized workload, containerized and AI/ML workloads on a single stack.
- Co-operative support model
- Simple to deploy, consume, and manage architecture, which saves time and resources required to research, procure, and integrate off-the-shelf components.

- Support for configuration, monitoring, automation and orchestration, and workload optimization using Cisco Intersight
- Cloud-based monitoring of all the components in the solution
- Validated policy driven modular design.

FlexPod Components

FlexPod is a pre-engineered solution equipped with Cisco UCS Compute integrated with the Cisco Nexus networking platform and NetApp Data Storage housed in a single unit.

Compute: The Cisco Unified Computing System (UCS) provides a modular, easy-to-use solution for powering modern workloads. Customers can choose between X-Series, B-series Blade, and C-series Rack Server models with an option to customize between Intel and AMD processors based on their workloads. The compute node is equipped with adequate PCIe, mLOM ports to interface adapters such as Virtual Interface Card (VIC), GPUs, etc., and offers room for further extension. For data center networks, UCS can be clubbed with Fabric Interconnects and Fabric Extenders to expand connectivity. Management Tools such as UCSM and Cisco Intersight can be used to tune the processor configurations accurately with ease including Boot options, BIOS settings, vNIC placement, Storage RAID setting for Local Boot, etc., to ensure modern applications run efficiently on both Bare Metal and Virtualized deployments.

Networking: Cisco Nexus is a highly modular, software-driven data center switching solution offering both spine/leaf and top-of-the-rack deployments. The unified ports support 10/25GE, 8/16/32G FC, RDMA over converged Ethernet (RoCE), and IP storage to reduce operational costs. Built on Cisco silicon, Nexus can operate in two modes: NXOS and Cisco ACI, thus simplifying hybrid cloud deployments. Nexus platform provides a rich feature set including comprehensive protocol support for L3, VXLAN EVPN overlay fabrics, Segment Routing (SR-MPLS), and state-of-the-art traffic engineering making it the right fit for next-generation data center and cloud networks. Customers can take advantage of features like Virtual Port Channel (vPC) to improve port efficiency via traffic load-balancing and build redundant networks via vPC Peer Link.

Storage Networking (SAN): Cisco MDS (Multilayer Fabric Switch) is the highly reliable, flexible, switch for SAN connectivity. It provides multiprotocol and multiservice convergence. Based on the workload requirement suitable MDS switches can be opted.

Storage: Powered by NetApp ONTAP, a proprietary storage OS, NetApp Flash Storage system offers industry's best storage efficiency, latency, and scalability for any workload. The A-series, and C-series platforms under the All Flash Array (AFF) umbrella support the classic NFS, iSCSI to the newest NVMe/FC, NVMe/TCP data access protocols. For a dedicated block storage experience, customers can opt for the All-Flash SAN Array (ASA) portfolio of products. All data storage platforms can be coupled with Data Management services for running resilient backup/recovery, Disaster Recovery, Business Continuity, Ransomware Protection, and cold data tiering to the cloud use cases. To ensure a smooth transition into the cloud, NetApp Cloud Volumes ONTAP is built for all the major hyperscalers that help manage, and protect data storage in the cloud. For modern containerized applications in the Hybrid Cloud, this modular approach enables independent development, deployment, and scaling of individual services, fostering agility and accelerating time-to-market. Microservices promote fault isolation, allowing failures in one service to be contained without impacting the entire application. They also facilitate diversity in technology, enabling teams to choose the most suitable technologies and frameworks for each service.

FlexPod: An ideal platform for various workloads

FlexPod is a proven solution and highly suitable to run various workloads and its flexibility allows you to run different environments on the same rack. We can deploy different environments with different infrastructure requirements such as vSphere, Kubernetes, VDI etc. on a single FlexPod. FlexPod is highly capable of running Modern Apps, Artificial Intelligence & Machine Learning, Virtual Desktop Infrastructure, Enterprise Databases, Healthcare, SAP, etc.

With the disaggregated architecture, customers can purchase exactly the infrastructure they need to start with their current application requirements and can scale up by adding more resources to the FlexPod system or scale out by adding more FlexPod instances. All components can be scaled as needed, allowing the FlexPod to meet the exact customer needs. The blade and rack servers can be added in the same rack and can be managed by UCSM or IMM. Similarly, the storage can be expanded by adding extra disk shelves to the same controller.

Design Considerations

Designing a FlexPod for a highly available production-grade environment is very crucial. There are several critical decisions that need to be considered to ensure that the cluster meets the performance, reliability, security, and scalability requirements of the applications it hosts. Let us discuss some important design decisions that an enterprise should consider.

A FlexPod configuration has the following minimum hardware requirements:

- A pair of Cisco Nexus switches in a redundant configuration. This configuration can consist of two redundant switches from the Cisco Nexus 7000, or 9000 Series. The two switches should be of the same model and should be configured in the same mode of operation. Nexus switches operate in two modes – NXOS and ACI.

  If you are deploying an ACI architecture, you must observe the following additional requirements:

  - Deploy the Cisco Nexus 9000 Series Switches in a leaf-spine topology.
  - Use at least three Cisco Application Policy Infrastructure Controllers (APICs).
- Pair of Cisco UCS 6400, 6500 Series Fabric Interconnects in a redundant configuration.

- Cisco UCS Servers:

  - If the solution uses the latest X series chassis then one X9508 Chassis with two x210c M6/M7 blade servers and two Intelligent Fabric Module (IFM) or for B-Series Servers, one Cisco UCS 5108 chassis, two B-Series blade servers with two I/O modules (IOMs).
  - If the solution uses only C-Series Servers, two Cisco UCS C-Series rack servers.
  - For larger deployments of Cisco UCS C-Series rack servers, you can connect them to Fabric Interconnects directly.

- A pair of MDS switches for SAN environment, few Nexus switches do support Ethernet as well FC but in case if there is specific requirement for dedicated SAN switches then MDS switches can be connected to FlexPod environment. For example if you want to use NetApp’s ASA for the block environment then an MDS switch can be used.

- NetApp storage controllers in a high-availability (HA) pair configuration:

  This configuration can consist of any supported NetApp FAS, AFF, or ASA-series storage controllers. See the NetApp Hardware Universe application for a current list of supported FAS, AFF, and ASA controller models.

  - The HA configuration requires two redundant interfaces per controller for data access; the interfaces can be FCoE, FC, NFS, iSCSI, NVMe/FC, NVMe/TCP.
  - Verify the onboard ports to be used for data access, or else order a networking card to support the requirement.
  - For ONTAP clusters with more than two nodes, a pair of cluster interconnect switches are required.
- Few controllers come with internal storage, if ordering a controller that does not support storage then one NetApp disk shelf with any supported disk type. See the Shelves tab of the NetApp Hardware Universe for a current list of supported disk shelf models.

When utilizing Cisco UCS Servers with shared storage, it is recommended to configure boot from SAN and store the boot partitions on remote storage. This enables administrators to take full advantage of the stateless nature of Cisco UCS Server Profiles for hardware flexibility across the server hardware and overall portability of server identity. In case of any server failure the replacement becomes very easy and can be done with the help of a few clicks. Boot from SAN also removes the need to populate local server storage thereby reducing cost and administrative overhead.

Cisco UCS supports different Virtual Interface Cards (VIC) that support multiple protocols and come with different speed options. The VIC card can handle different traffic and provides traffic isolation with the help of VLANs. For each type of traffic, we can create a separate VLAN on UCS.
In a standard design, we can create multiple vNICs for ESXi hosts and assign them to specific virtual and distributed switches. VIC will provide redundancy and traffic load balancing as it connects to both fabric interconnects.

NetApp ONTAP provides unified capabilities of handling multiple storage protocols. A single HA pair of NetApp controllers can be used for multiple storage protocols such as NFS, CIFS, FC, iSCSI, etc. ONTAP systems are inherently designed for multi-tenancy. With ONTAP, you use Storage Virtual Machines (SVMs) to serve data to clients and hosts. A SVM is a logical entity that abstracts physical storage resources. ONTAP offers a wide range of storage efficiency technologies that reduce the amount of physical hardware or cloud storage consumed by your data and also yield significant improvements to system performance - including faster reads of data, faster copies of datasets, and faster VM provisioning.

Deployment

Deployment of a FlexPod can be done manually or via IaC tools like Ansible. For the compute part the configuration can be done either via UCSM or IMM. The latest X series servers can only be configured via IMM.

Conclusion

FlexPod has been running in the data centers for more than a decade and has proven its capabilities. With the flexible architecture, scaling of FlexPod is just simple and one can start with the minimum setup and can scale based on future requirements.

With the latest innovation from Cisco and NetApp, FlexPod is an ideal solution for building a Hybrid cloud environment. It enables agility, scalability, resilience, and efficiency, ensuring that applications can handle large volumes of traffic and adapt to changing requirements effectively. By carefully considering and making the appropriate design decisions, one can build a highly available data center for their environment. To enable yourselves to make the right design decisions and implement it with confidence and trust, contact us at info@dacetek.com.