NEW HORIZON COLLEGE OF ENGINEERING

DEPARTMENT OF INFORMATION SCIENCE & ENGINEERING

2015 - 16

INFOTECH PATRIKA - A Half- Year Publication

VOLUME 1 ISSUE 1

About the Department:

Information science and Engineering department focuses on current Information Technology trends, and Domain Specific Applications. The program facilitates the evolution of skills in students to help them attain a higher degree of knowledge, global competency and excellence, for the betterment of the society. The Department of Information science and Engineering at NHCE was established in the year of 2001 and offers graduate and PhD programs. The four year B.E degree equip the students to meet day- to- day Technological advancements of the ever dynamic IT field through adept training on various subjects of curriculum of Information Science and engineering and beyond. The department offers B.E program through autonomous scheme from the year 2015. The department has a total intake of over 380 students with a very good team of highly qualified and talented faculty members including Professors, Associate Professors and Assistant Professors.

Information Science and Engineering course at New Horizon College of Engineering is designed to meet industry standard and cope up with the emerging technology. There is a great emphasis on holistic learning to help the students to make significant contributions at all levels and to meet the expectations of stakeholders. The department is well known for its research excellence in various competitive areas of Information Science. Students are made to involve vigorously in research activities. The department provides industry collaborated courses for the students.

Emerging Technologies in Computer Networks

Data center networking is progressing through a major shift that aims to simplify and automate the provisioning of network resources. We take a look at the next generation networking technologies and tools from some of the top enterprise vendors.

During the last decade, the majority of data center networks have been built using a very common design that involves the implementation of a two to three layer modeled solution. This solution works quite well, which is why it has been implemented for so long, but change is on the horizon. The recent introduction of software defined networking (SDN), network overlay technologies, network virtualization (NV), and a general focus on deploying and upgrading to more operationally efficient systems has lit a fire under many companies, forcing them to evolve their data center networking strategies.

As we dive into data center networking technologies and tools, this article will review some of the main network designs that have been used over the last 15 years as well as some of the emerging technologies that are currently being developed and deployed to evolve the standard methods of implementation. We will also take a look at some of the available tools and the top data center networking vendors.

A Short History Of Data Center Networking

As many a network engineer will discuss, the basic design used within most networks (and not just in data centers) over the last 20 years has been one that involves the implementation of a two to three layer architecture (Fat Tree). The Fat Tree architecture typically involved the implementation of server racks placed at the access/edge layer, which was then connected to each other via an aggregation/distribution layer, which in turn was interconnected together at a core layer.

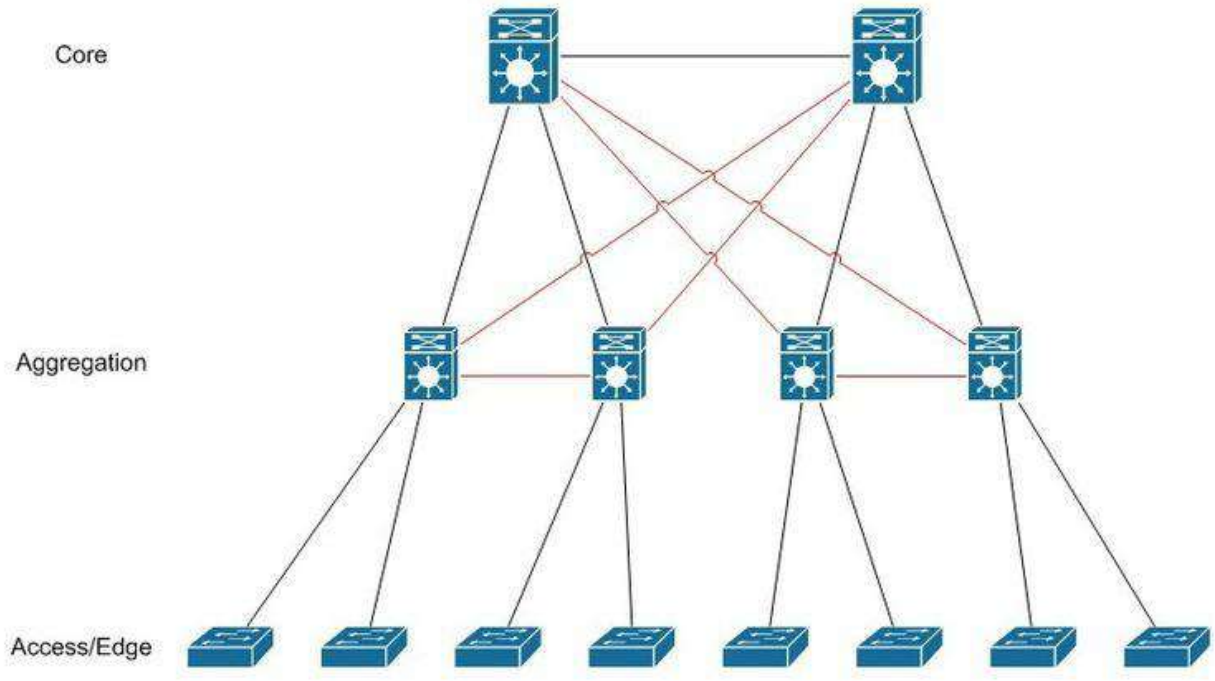

Figure 1: Fat Tree (Red Links are Optional)

This design was originally developed with the idea that most traffic would be going from the access layer to the core layer and back (North - South). The problem however, in many modern networks is that this pattern assumption is incorrect in the data center. Modern data centers have many different resources (compute, storage, etc) that are interconnected with one or many different virtualization technologies. This evolution has changed the traffic patterns so that much more access to access layer (East - West) traffic occurs than up through the aggregation and core layers; this causes many different potential bottlenecks to form and for the network to quickly become inefficient.

On top of these traffic changes, the network has been going through a evolutionary shift in the way that resources are used and managed. 20 years ago if an organization was going to deploy a web server, they would buy a physical server (or blade) and install in in their data center. Now this is done typically without any additional hardware by utilizing virtualized compute and storage resources. However, the problem up until recently has been that while the compute and storage resources have evolved to a point where they could be quickly provisioned from a central location, the network resources required to make the changes could take considerably more time to provision (hours or days, not minutes).

- Ms. Sushma D S

Next Generation Data Center Networking

The next generation of data center networking is still in the process of being determined, but there are a few main camps that have the momentum, whether they combine together or

go in completely different directions is the current big unknown question. One thing that is a common objective is to reduce the amount of time it takes to provision network resources on a network and to automate as much of the individual Command Line Interface (CLI) provisioning as possible.

Software Defined Networking (SDN)

The first of camps that exist revolve around the concept of Software Defined Networking (SDN). SDN itself is broadly defined but basically involves the de-coupling of the control and data planes of the networking equipment (typically switches). What this means is that while it is common for protocols like OSPF (Open Shortest Path First), IS-IS (Intermediate System to Intermediate System) or BGP (Border Gateway Protocol) to be used inside a modern data center the implementation of SDN would remove this routing intelligence (control plane) from the individual devices and move it to a central controller. The actual forwarding of the data (data plane) would remain the responsibility of the individual device, but its forwarding tables would be controlled and modified from a central location.

There are a number of new terms and protocols that have been introduced and implemented for SDN. One of the most common terms that was initially written about quite frequently was OpenFlow. OpenFlow is a "Southbound" control protocol that is used to communicate between the controller and the individual devices.

The main organization that is pushing OpenFlow as well as Open SDN is the Open Networking Foundation (ONF). ONF includes a large membership of most of the major network vendors, which support the general idea of bringing SDN into the mainstream. The key point to highlight is in the openness of the protocols developed with and by ONF members; this is going to be a potential deciding factor in many organizations' selection of vendors. Some vendors support the idea of Open SDN, but still have invested in other non-open technologies in the interest of maintaining their existing customers. Sometimes this method is successful as these specific vendors have decided to develop in other directions that end up with a superior implementation, and sometimes it doesn't; this is the fundamental question being answered by decision makers in today's data centers.

- Ms. Pavithra

SDN And The Evolution Of Existing Topologies

One of the biggest questions that many organizations will need to answer is whether they wish to jump on the SDN bandwagon and whether this decision is even worthy of their specific implementation. In small to medium sized networks, the cost of changing over to SDN may not be worth the cost of implementation in the short or long term. Where the ideas behind SDN really shine are on larger networks and especially in large- to massive-

sized data center networks. Another shift that is coming along at the same time as the SDN debate is whether the legacy topologies are designed for their specific traffic requirements.

In data center networks this really isn't much of a question, it is a fact that the traffic patterns within data centers have changed fundamentally from a North-South traffic pattern to a East-West traffic pattern. There's also the fact that excessive and/or wasted capacity is also becoming less of an acceptable condition and is being more focused on as a source of potential savings.

With this, there is a slow evolution from the legacy Fat Tree topology to a newer Clos style of topology. Many forward thinking organizations are well on their way to changing the way their data center networks are connected by completely altering the way that devices are interconnected.

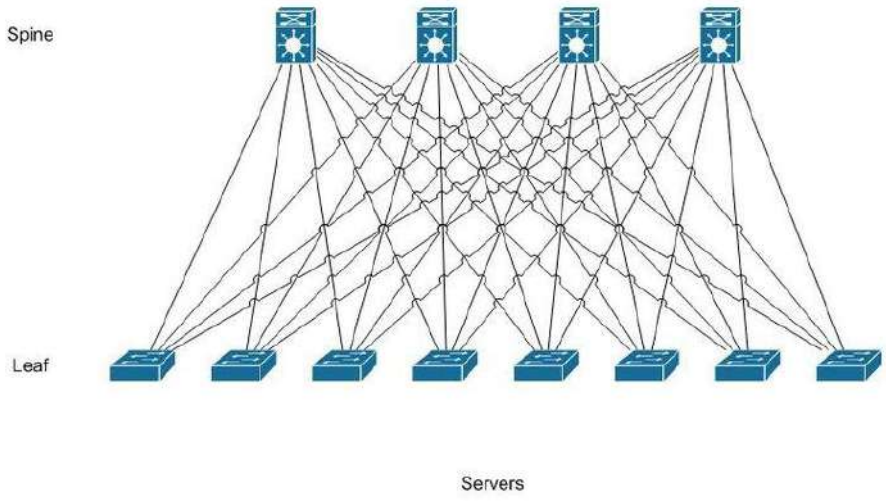

Figure 2: Clos Topology

Overall, most organizations will end up with a hybrid of these topologies, depending on their specific requirements and how their current networks are interconnected. There are variations of these two topology types that allow existing Fat Tree topologies to implement some of the Clos topology advantages without needing to completely alter the way things are currently interconnected.

- Mr. Anto Malvin

Overlay Technologies

Some of the technologies that have gotten the most traction over the last several years have revolved around their ability to "overlay" on top of existing networks regardless of their physical or logical design. Some of the common technologies that are being implemented include Virtual Extensible LAN (VXLAN), Network Virtualization Using Generic Routing

Encapsulation (NVGRE), Transparent Interconnection of Lots of Links (TRILL), Shortest Path Bridging (SPB), and Location/Identifier Separation Protocol (LISP). Each of these has their own advantages and disadvantages and are implemented more or less depending on the specific vendor that is selected. The selection of which one is better than the other is not a settled argument.

One thing is for sure, data center networks, along with networking in general, are eventually going to evolve from their existing best practices. While data centres do require a specific use case, they are also a place where networking will need to evolve the fastest and with it will be an intense time of technology review. The technologies that best fit the solutions of the majority will eventually bleed out and be used on other parts of the network.

Now that we have an overview of the existing and evolving data center networking technologies, we'll take a closer look at some of the specific networking tools. In the coming weeks we'll review some of the top data center networking solutions from vendors including Arista Networks, Cisco, Dell and HPE.

Now that we have covered the high level history of how data center networking has evolved over the last few decades and discussed some of the current directions that it is taking in modern data centers, let's review solutions from four of the top data center networking vendors.

- Mr. Akash S

Arista Networks: CloudVision SDN

Arista Networks is currently one of the top data center networking vendors and a strong player in this space, mainly because of the company's advanced prediction of how networking will be changing in the future. Arista's solutions are based on a number of different switching platform options, each of which run a copy of Arista's Extensible Operating System (EOS), which is the centerpiece of the company's solutions. EOS is built with complete openness in mind and has an open API that can be interconnected with most other platforms, including OpenStack, VMware, F5 and many others. EOS is modular in nature and has a number of advantages over traditional network operating systems because the different running processes can be independently modified, configured and upgraded without bringing down the other processes running on a device.

Arista's CloudVision SDN solution takes advantage of its open EOS software; it utilizes the flexible software and programmability to support a number of different plug-in solutions, including those from VMware, Microsoft, Dell, Palo Alto and Rackspace. CloudVision can be implemented in a number of ways depending on the specific size and needs of an organization.

Support includes four different principle topology designs and each of these topology options can be deployed using a number of Arista's supported switches depending on the needs of the organization:

- Spline -- a combination of Leaf and Spine Clos elements
- Clos using Layer 2 and Multichassis Link Aggregation (MLAG)
- Clos using Virtual Extensible LAN (VXLAN)
- Clos using Equal-cost multi-path routing (ECMP)

While Arista may be a newer player compared to some of the other vendors we will discuss below, it has a strong relationship with a number of different partners and an advantage in software-centric solutions, which looks to be the direction that networking is heading in the near future. This in itself should be enough for any organization to take a look at their available options and weigh them against existing deployed equipment and the future direction of the organization's software integrated network.

- Mr. Rahul Yadav

Cisco Systems: ACI

Cisco Systems has been around for a long time and the company is and has been the leader in many networking categories for years, and data center networking is no exception. Because of Cisco's history in so many different parts of the network, it has a greater number of available equipment options across all network sizes, from SMB to enterprise. However, a higher number of equipment options can be both an advantage and a disadvantage when trying to make a purchasing decision; when there are so many available solutions, the "best" option can be subjective and differ considerably.

Cisco offers documentation on a number of different Cisco Validated Designs (CVDs), including several that are focused on the data center. Cisco's Nexus line of switches are focused on the data center and are central to its medium and larger sized solutions. Cisco's data center networking solutions include the Prime - Data Center Network Manager, which provides a management solution for both the Nexus-based solutions and Storage Attached Networking (SAN) technologies. It also has the capability to integrate with other orchestration products like OpenStack and VMware's vCloud Director.

Unlike some of its competitors, Cisco was late in the Software Defined Networking (SDN) market but has made it a mission to quickly regain pace by following a few different paths. The company has offered the Cisco Extensible Network Controller (XNC), which is based on the OpenDaylight SDN solution and has developed the Application Centric Infrastructure (ACI). ACI is the solution Cisco is primarily pushing, which is centered around its Application Policy Infrastructure Controller (APIC) and the Nexus 9000 series of switches. The Nexus 9000 series are able to operate in both "normal" and in ACI mode allowing organizations

thinking about implementing a partial or full SDN solution in the future to deploy equipment that supports both and can be configured to the specific requirements. ACI has an expansive partner ecosystem allowing a number of different solutions to be supported; this wouldn't be possible without a well documented Application Program Interface (API).

Since Cisco has so many different options it should obviously be part of most organization's selection process. The company does however have one major disadvantage, it tends to be the most expensive solution, which can lead many companies to select other options.

- Ms. Deepthi Bhat

Dell Data Center Networking

Dell is one of the top networking vendors in the data center and is attempting to gear up and gain on the lead of its competitors. One of Dell's main advantages is that it is well known in a number of different sectors and because of this it has both name recognition and reputation for being flexible and competent at the enterprise level. This also provides the company the ability to tie in solutions from a number of different sectors into a single offer to increase savings.

Dell has two different primary data center networking product lines. The first runs Dell's proprietary network operating system and the second is designed to be open and run a number of supported operating systems from companies like Big Switch, Cumulus, IP Infusion, Midokura and Pluribus Networks. Both of these solutions can be deployed in a number of different configurations including both Fat-tree and Clos-based topologies. Dell's open-based systems are designed to be worked into specifically engineered third party solutions where the switch will be plugged into a topology solution designed by each respective vendor.

Dell's networking equipment is designed and licensed in a way that allows for equipment and features to be added on an as needed basis and into a number of different design topologies. Core layer topologies are focused on the use of the Z-series fabric switches, which offer support for both "normal" mode operations and support for SDN implementations. Dell's SDN support includes OpenFlow and VXLAN and is extended with its open networking support for third party vendors that have more integrated SDN implementations.

Because Dell is so well established in the server space this can and should be a big factor in implementation decisions. While its in-house implementations are not as expansive as their networking-centric competitors, they should be a considered a selection in any new or existing data center implementations.

- Mr. Nandiprasad G

HPE Data Center Networking

HP Enterprise Co. (HPE) was formerly the enterprise focused division of Hewlett-Packard (HP), but as of November 2015 has been split off as a separate company. HPE offers a number of different enterprise focused solutions across a number of different platforms and service areas. HPE is currently one of the top players in the data center networking market and has been a strong networking player for a number of years. HPE offers a number of different individual networking devices, including routers and switches as well as flexible networking fabric solutions.

Because of HPE's product diversity, each solution is highly contingent on the specific implementation and organization. For smaller data centers, a network of connected and designed individual components may offer enough capacity and meet the organization's requirements. For companies that are looking to build a large-scale Fat-Tree or Clos-based data center networking topology, HPE offers a number of different high density switched solutions, principally the FlexFabric line of switches. And for data centers that are looking into the future and have bought into SDN, HPE has a complete turnkey solution with its Synergy platform. Through HP OneView it is also possible to build a network that supports a more traditional infrastructure as well as SDN.

Currently, HPE seems to be putting a lot of attention on its newer Synergy platform, this along with its support for a "compostable infrastructure" will allow organizations to provision Compute, Storage and Networking dynamically via an integrated API. This is HPE's fist offering where the networking would be integrated into a turnkey solution. It's also a great example of where SDN is really taking networking in the future. The Synergy platform can be integrated into a number of different third-party partner solutions, including OpenStack, Docker, Chef, Puppet and VMware, to name a few.

- Mr. Nagendra Devadiga