

# Why VXLAN is Awesome, and Why it May Not Be

✖ If you're not already familiar with VXLAN, you should get to know it. VXLAN (acronym for Virtual eXtensible Local Area Network) is fast becoming the hottest new discussion item for Cloud service design.

## What is VXLAN?

VXLAN technology overlays a virtualized Layer 2 on top of a Layer 3 network allowing the extension of Layer 2 traffic between physical segments. While this has been achieved to a degree with VLANs there is a limit of 4096 VLAN IDs per network and there are many challenges with subnet design due to the limit of Layer 3 boundaries.

While 4096 VLAN IDs seems like a rather large number, if you are a Cloud provider or SMB you will be approaching that boundary far earlier than you may have realized.

## Understanding the Broadcast Domain

If we go back to our early studies of network fundamentals and subnet design in switched networks we remember that each subnet has a range which contains the usable addresses, a subnet address, a broadcast address and a gateway address.

The broadcast address receives and transmits the datagrams destined for the segment to all hosts on the segment. This means one simple thing: The boundary of the network is the subnet.

## Why do I care about Broadcast Domain

The MAC address (Layer 2) and IP address (Layer 3) have real boundaries. Let's say that you have an application that runs on a server which may failover to a secondary host. The ideal way for the application to continue to roll on without interruption is if the destination host has the same network address.

Right away we have hit an issue. We will most likely have to failover to a host on a different physical network and thus a different network address outside of the broadcast domain. We have accepted this as a limitation up to now and we do lots of application logic and leverage DNS and other tools to transition applications between hosts on different networks.

## VXLAN to the Rescue

By virtualizing Layer 2 VXLAN can use a single Multicast group across multiple logical networks which allows us to bridge datacenters without changing address or gateway because the MAC and IP info are encapsulated into the Multicast traffic and can cross previously unreachable boundaries.

It is possible to have your network across physical locations using a broad address space by creating large subnets, but if you do this the protocol mashup and incredible amount of broadcast traffic would be untenable. Just ask your network designer if it's a good idea, and when they finish wiping the tears of laughter away you can tell them you've found a way around it ☐

VXLAN lets us extend the address space, logical isolation and more importantly remove the limitations of the physical capability of the current network hardware. This in a nutshell is the heart of SDN (Software Defined Networking).

That's a rather simple explanation of it, so I would encourage everyone to read up and get familiar with what VXLAN does at the core and how it works to get a full understanding. SDN and the SDDC (Software Defined Data Center) will become part of your arsenal of tools very soon I'm sure.



Source:

[http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/white\\_paper\\_c11-685115.html](http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/white_paper_c11-685115.html)

## Why VXLAN Won't Work for Everyone

While VXLAN is an important step towards true Cloud extensibility it also isn't the magic bullet. The first thing that many think is "Why don't I just run this in my regular LAN?". That is a great question.

There are a number of issues with VXLAN and how it affects your current application infrastructure. You most likely have software today that depends on logical segmentation to work. VMware SRM is a great example of this. Today, SRM is designed for the destination recovery machine be on a different logical segment. VMware vSphere Replication uses the same concept. It requires that the source address is unavailable to bring the failover device online. In a VXLAN, single logical subnet design these products won't work as they need to.

Stretched Clusters are the ideal use case for VXLAN, but with everything that they gain in functionality, there are pitfalls as mentioned with products like SRM. The challenge is that your application needs to be able to fully support the features that are offered with VXLAN. It may seem like this should be a no-brainer but you may be surprised by how the application acts when it's moving between nodes.

## Hybrid is the new Standard

VXLAN is new, exciting, and should be a consideration in your design. It will not replace what you do today. No more than IPv6 can replace IPv4. What it can do is augment your current design and have specific application within your environment to deal with a particular use case.



Think of it as someone discovering that they like pepper. That doesn't mean that they should put pepper on everything they eat all of a sudden. We need to take the new capabilities and look for the appropriate use case and then apply as needed.

As Application Designers, Systems Architects, Systems Admins and System Developers we need to do our best to be aware of every tool that is available to understand where each may be appropriate.

Infrastructure design should definitely involve an understanding of the Law of the Instrument as defined by Abraham Maslow. The often paraphrased Maslow stated: "I suppose it is tempting, if the only tool you have is a hammer, to treat everything as if it were a nail."

We have a whole array of options available to us, so build your recipe to include what you need, and feel free to mix and match as needed. A little bit of SDN, some VXLAN, a little bit of OpenStack (this is a whole new and awesome topic) and just flavour to taste!

If you find that after your analysis that you don't have a way to integrate VXLAN into your environment, that isn't a bad thing. But the key is to know that it is available as an option.