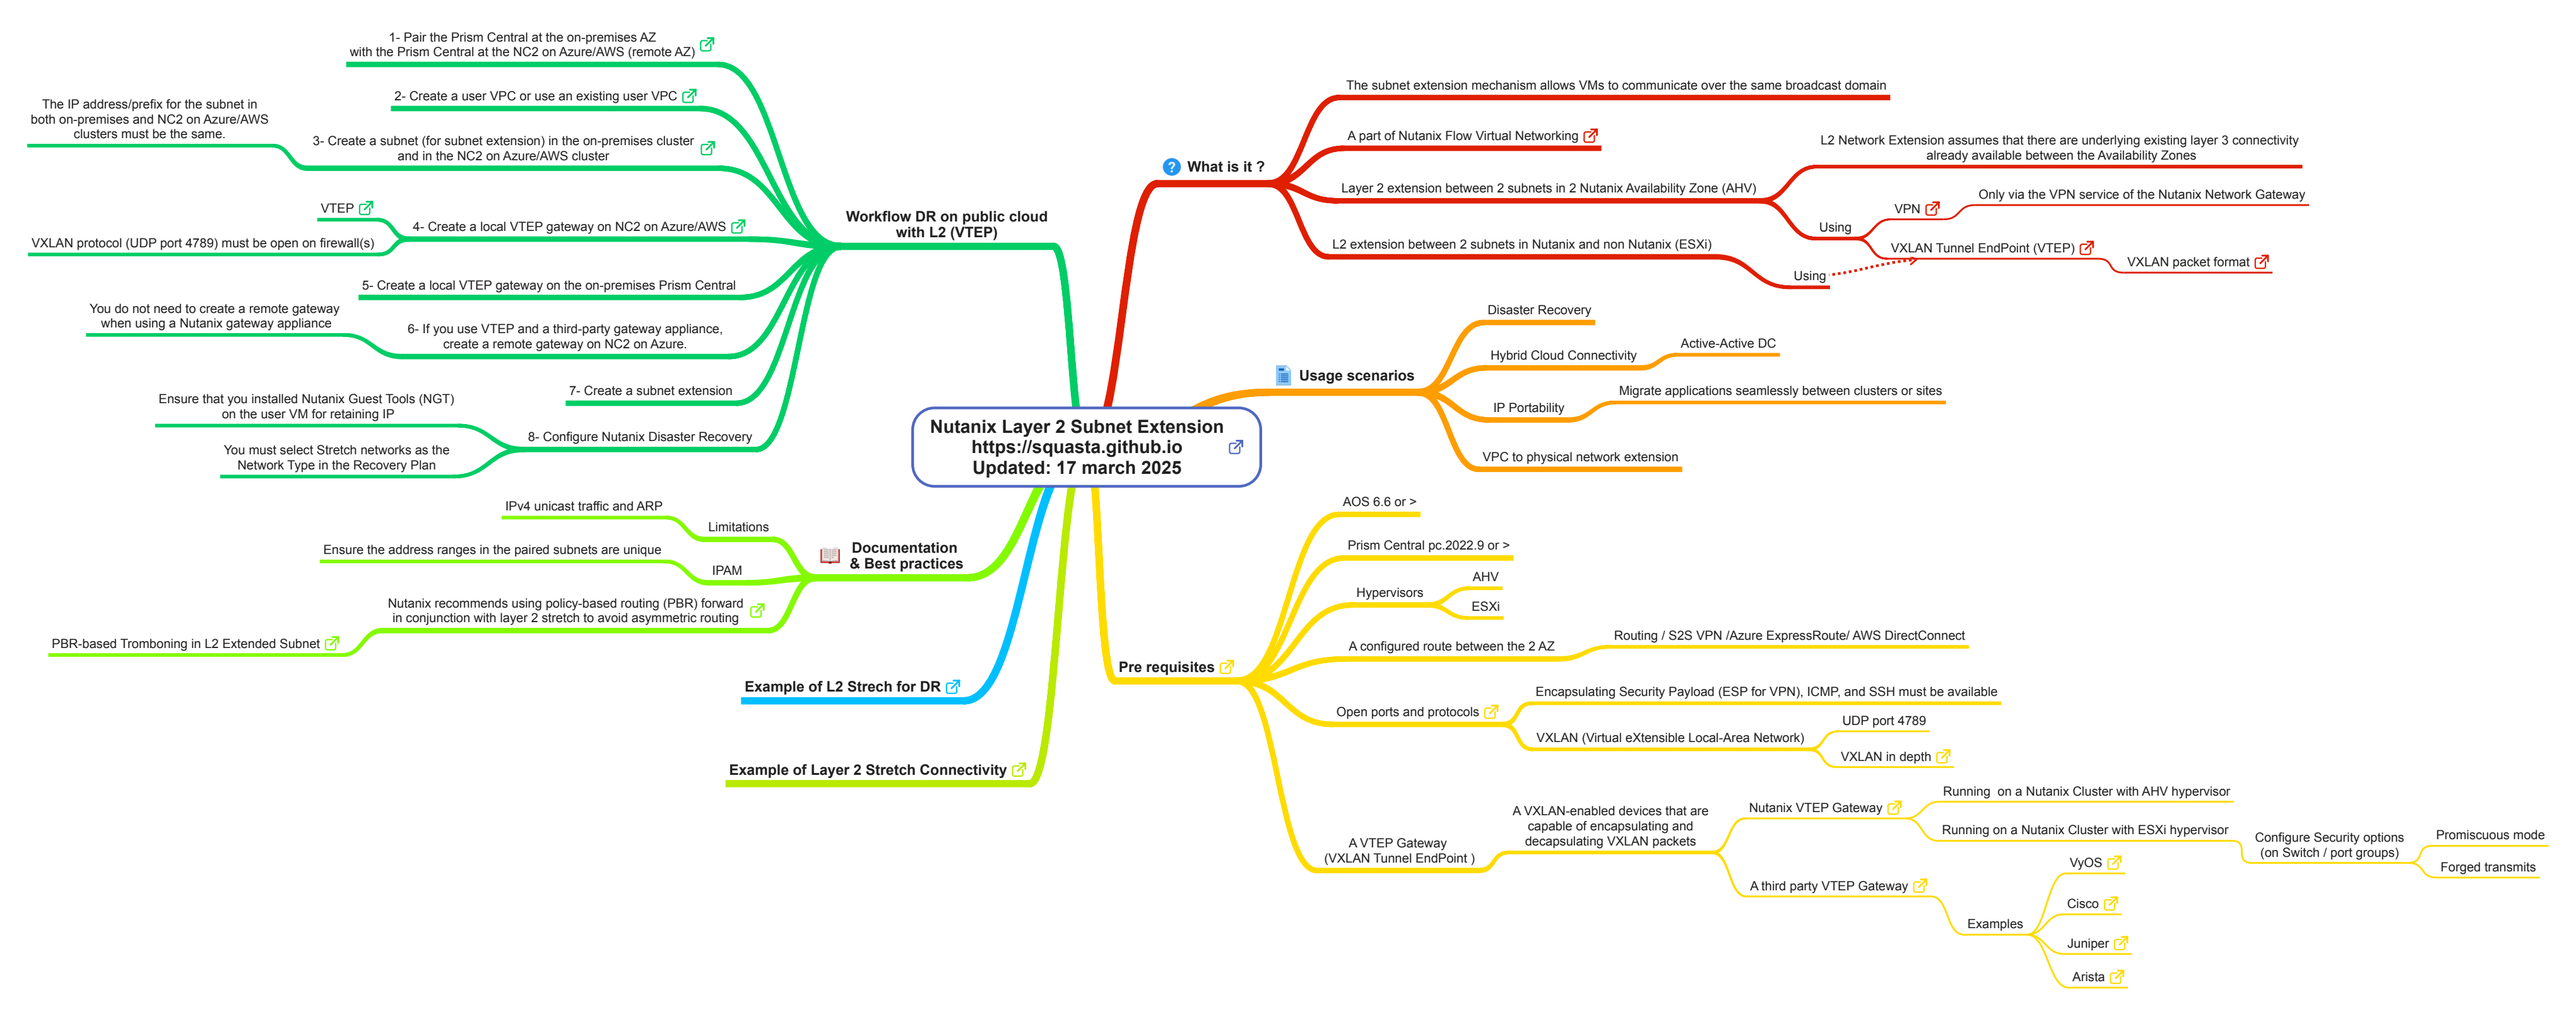
Nutanix Layer 2 Subnet Extension
https://squasta.github.io
Updated: 17 march 2025
What is it ?
The subnet extension mechanism allows VMs to communicate over the same broadcast domain
A part of Nutanix Flow Virtual Networking
Layer 2 extension between 2 subnets in 2 Nutanix Availability Zone (AHV)
L2 Network Extension assumes that there are underlying existing layer 3 connectivity already available between the Availability Zones
Using
VPN
Only via the VPN service of the Nutanix Network Gateway
VXLAN Tunnel EndPoint (VTEP)
VXLAN packet format
L2 extension between 2 subnets in Nutanix and non Nutanix (ESXi)
Using
Usage scenarios
Disaster Recovery
Hybrid Cloud Connectivity
Active-Active DC
IP Portability
Migrate applications seamlessly between clusters or sites
VPC to physical network extension
Pre requisites
AOS 6.6 or >
Prism Central pc.2022.9 or >
Hypervisors
AHV
ESXi
A configured route between the 2 AZ
Routing / S2S VPN /Azure ExpressRoute/ AWS DirectConnect
Open ports and protocols
Encapsulating Security Payload (ESP for VPN), ICMP, and SSH must be available
VXLAN (Virtual eXtensible Local-Area Network)
UDP port 4789
VXLAN in depth
A VTEP Gateway
(VXLAN Tunnel EndPoint )
A VXLAN-enabled devices that are capable of encapsulating and decapsulating VXLAN packets
Nutanix VTEP Gateway
Running on a Nutanix Cluster with AHV hypervisor
Running on a Nutanix Cluster with ESXi hypervisor
Configure Security options (on Switch / port groups)
Promiscuous mode
Forged transmits
A third party VTEP Gateway
Examples
VyOS
Cisco
Juniper
Arista
Workflow DR on public cloud with L2 (VTEP)
1- Pair the Prism Central at the on-premises AZ with the Prism Central at the NC2 on Azure/AWS (remote AZ)
2- Create a user VPC or use an existing user VPC
3- Create a subnet (for subnet extension) in the on-premises cluster and in the NC2 on Azure/AWS cluster
The IP address/prefix for the subnet in both on-premises and NC2 on Azure/AWS clusters must be the same.
4- Create a local VTEP gateway on NC2 on Azure/AWS
VTEP
VXLAN protocol (UDP port 4789) must be open on firewall(s)
5- Create a local VTEP gateway on the on-premises Prism Central
6- If you use VTEP and a third-party gateway appliance, create a remote gateway on NC2 on Azure.
You do not need to create a remote gateway when using a Nutanix gateway appliance
7- Create a subnet extension
8- Configure Nutanix Disaster Recovery
Ensure that you installed Nutanix Guest Tools (NGT) on the user VM for retaining IP
You must select Stretch networks as the Network Type in the Recovery Plan
Documentation & Best practices
Limitations
IPv4 unicast traffic and ARP
IPAM
Ensure the address ranges in the paired subnets are unique
Nutanix recommends using policy-based routing (PBR) forward in conjunction with layer 2 stretch to avoid asymmetric routing
PBR-based Tromboning in L2 Extended Subnet
Example of L2 Strech for DR
Example of Layer 2 Stretch Connectivity