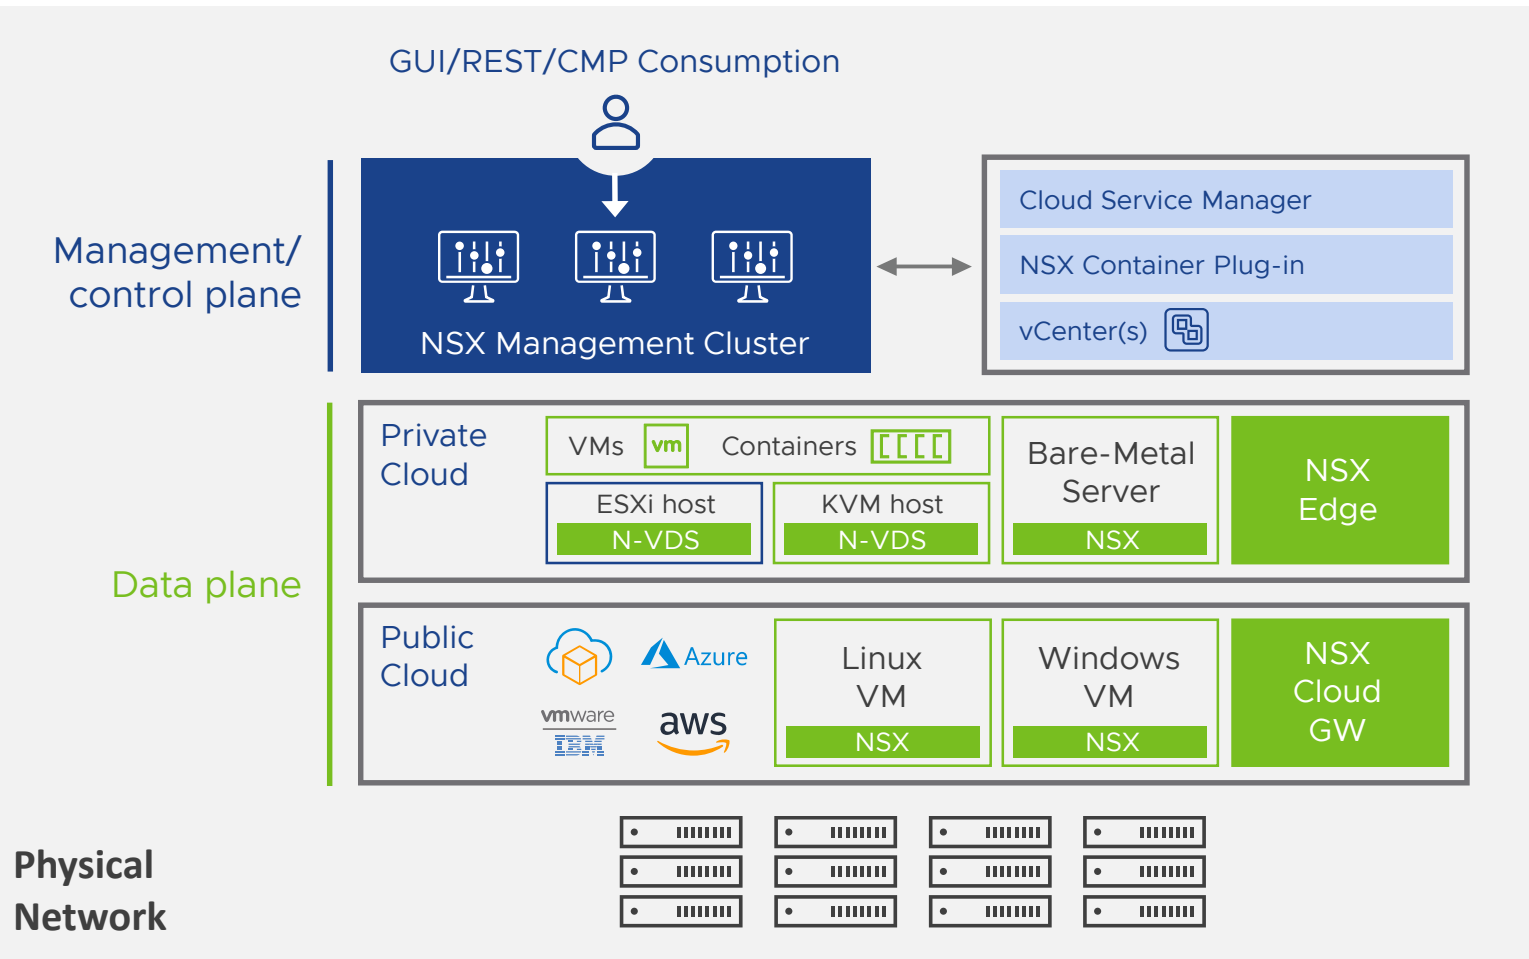


# NSX-T Architecture

Overview of the NSX-T Management,  
Control, and Data Plane

# NSX-T Datacenter Components

## Management, control and data planes



## Management/control plane

NSX-T utilizes a multi-tiered networking stack.

The NSX Management Cluster is a 3 node high availability cluster. This cluster consists of a converged Management (Policy) & Control Plane cluster services.

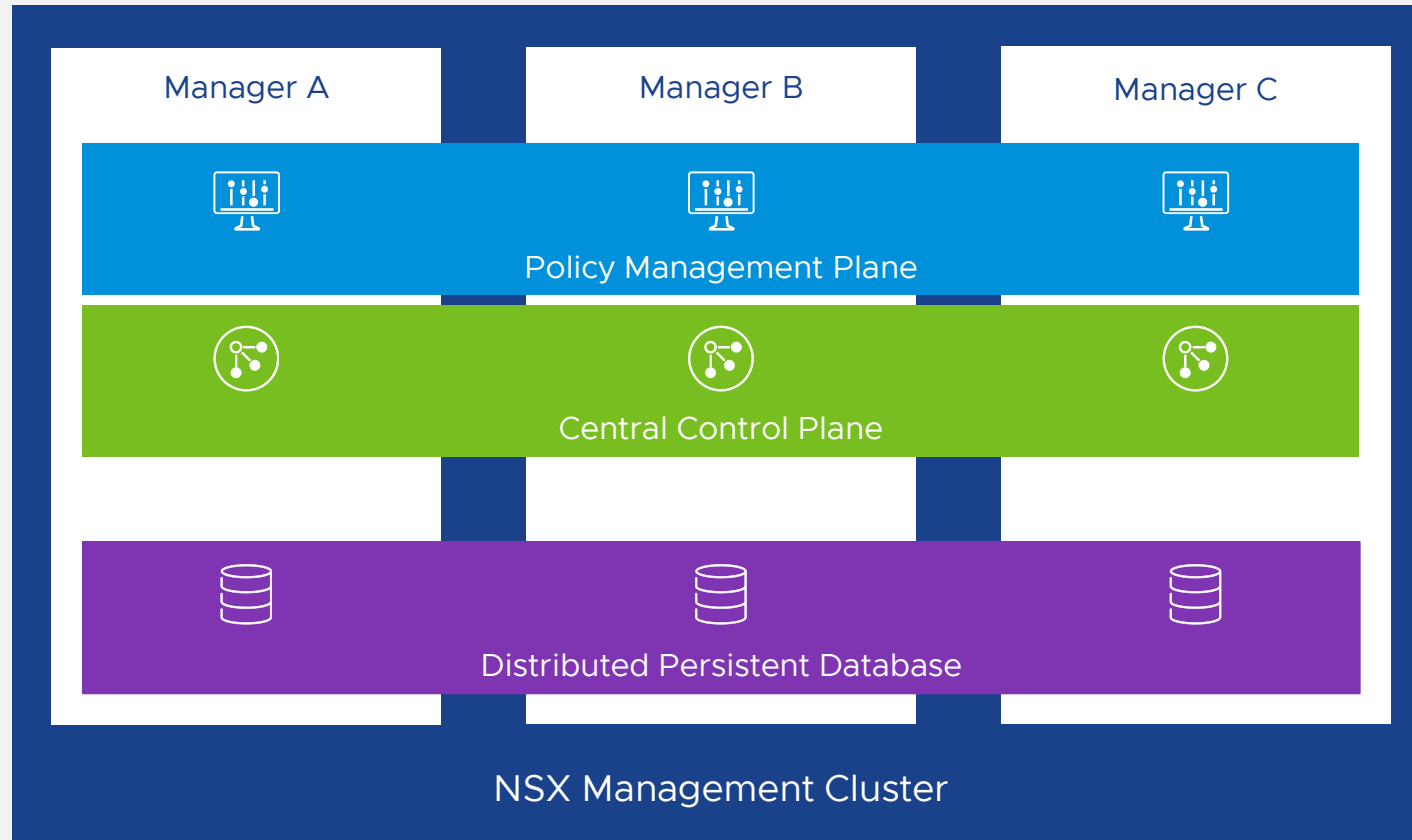
## Distributed data plane

The NSX-T distributed data plane connects and hosts workloads across an entire enterprise utilizing heterogeneous hypervisor support and multiple Public cloud connectivity. Supports hosting the most diverse array of application frameworks – VMs, containers, micro-services, bare-metal, etc.

Implements distributed switching, routing and firewalling.

# NSX-T Manager – A Converged Appliance

Merging policy, management, and central control services on a cluster of nodes



## Feature

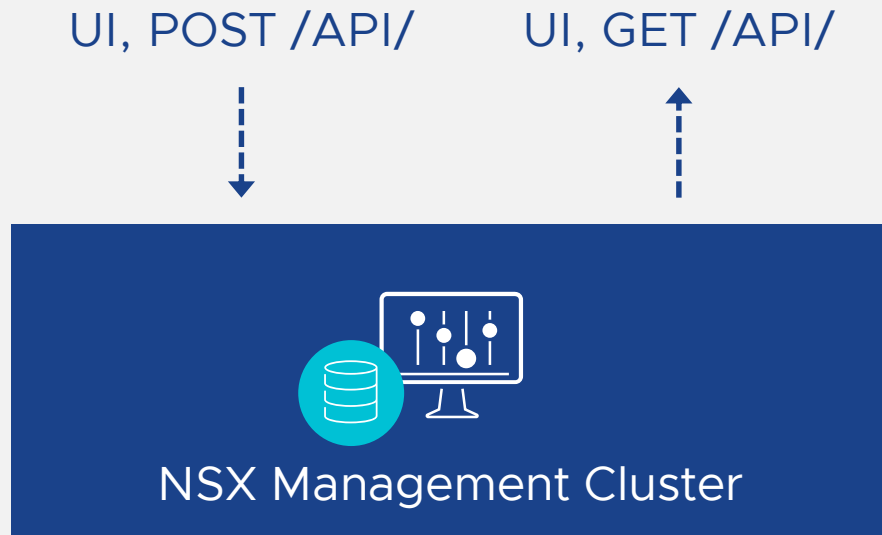
The NSX management nodes each contain a Management plane, a central control plane, a policy role and a replicated desired state datastore.

The NSX Management Cluster provides availability of all management services and increased performance.

The converged appliance allows for easier operations with less systems to monitor and maintain.

# NSX-T Manager

## Functional characteristics of the NSX-T Manager

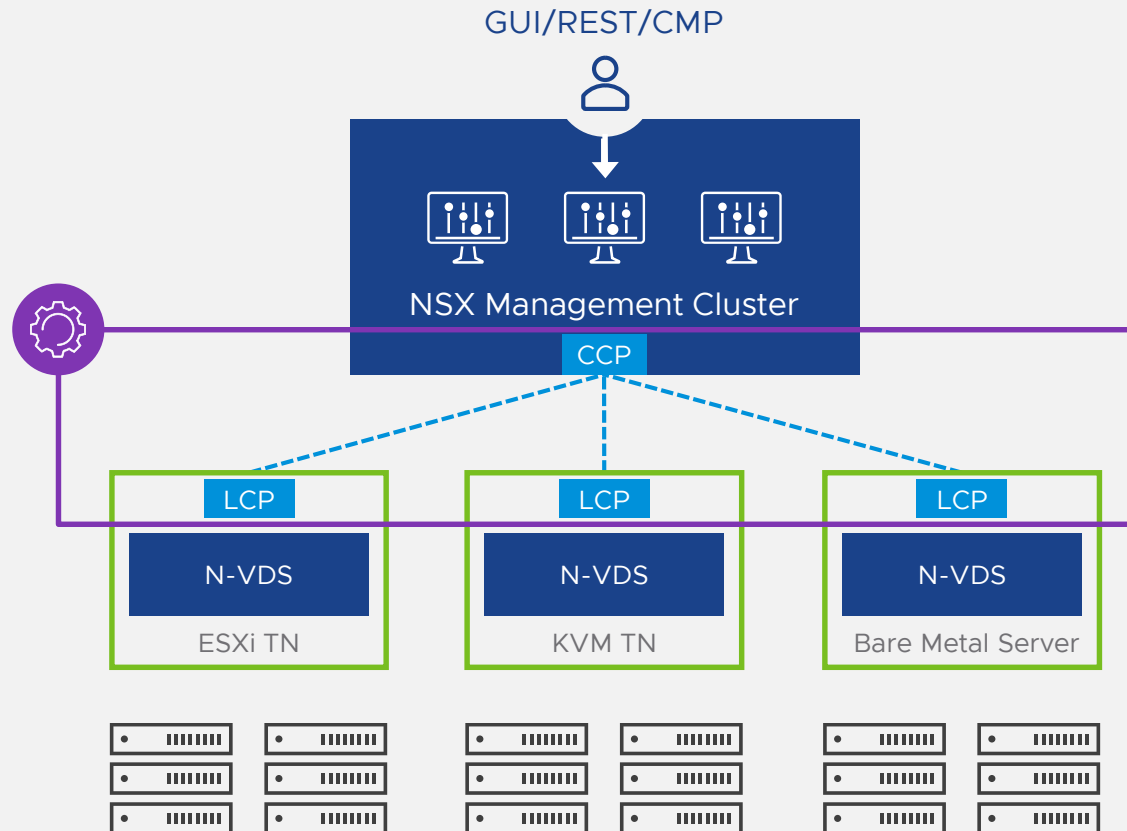


## The NSX-T Manager

- Maintains connectivity to all nodes in the system
- Provides entry point to the system via UI or API
- Handles user queries
- Persists user configuration and the desired configuration
- Validates the stores data state
- Maintains and propagates the dynamic state

# NSX-T Management and Control Plane

## Centralized & local control



Computes all ephemeral runtime states based on configuration from the management plane

Pushes stateless configuration to forwarding engines on transport nodes

Disseminates topology information reported by the data plane elements

The control plane is distributed between

- Central Control Plane (CCP) in the Manager Cluster and
- Local Control Plane (LCP) agents on the hosts

# NSX-T Manager Clustering with Virtual IP

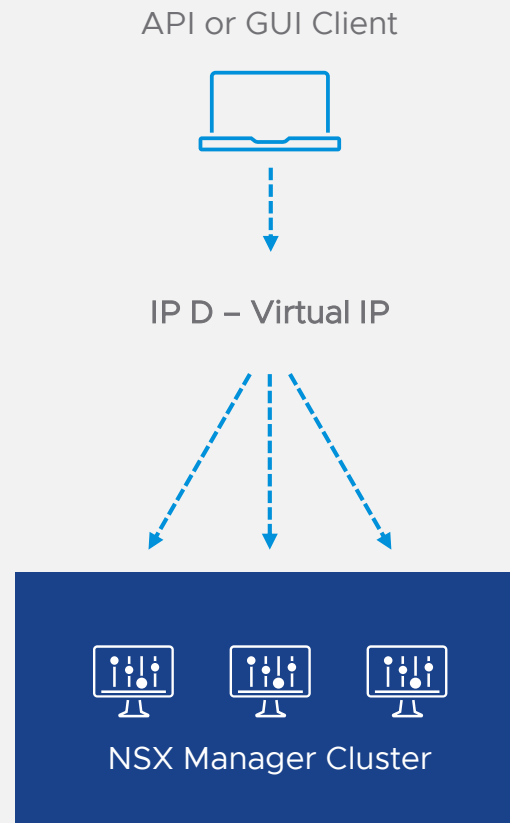
A highly available management plane for GUI and API

The cluster Virtual IP is assumed by one Manager called the leader.

All cluster nodes must be in the same subnet.

GARP is used when the Manager with the Virtual IP fails.

The cluster IP is used for the north bound operations. The south bound connectivity to the hosts uses the physical IP of each node.



## Feature

Cluster of three NSX managers

API and GUI available on all managers

Replicated desired state datastore

## Benefit

High availability of the NSX UI and API

Reduces the likelihood of failures of NSX operations

Provides API and GUI clients with multiple endpoints or a single VIP for availability

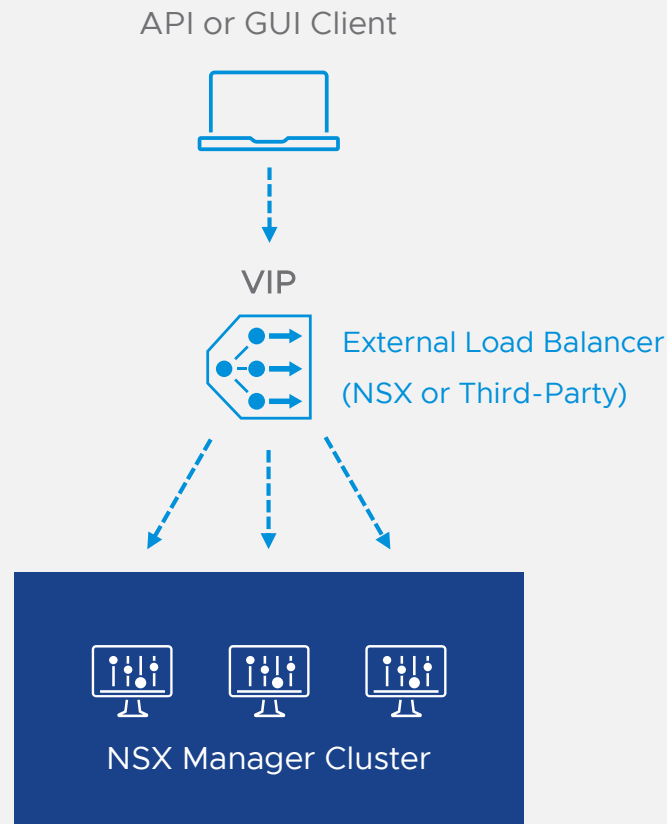
# NSX-T Manager Clustering with Load Balancer

A highly available management plane for GUI and API

All Nodes Active

VIP Load Balances to Multiple Managers

Managers can be in different subnets



## Feature

Use any external load balancer including the NSX Load Balancer

All NSX Managers within the cluster are available for responses from API and GUI issued requests.

## Benefit

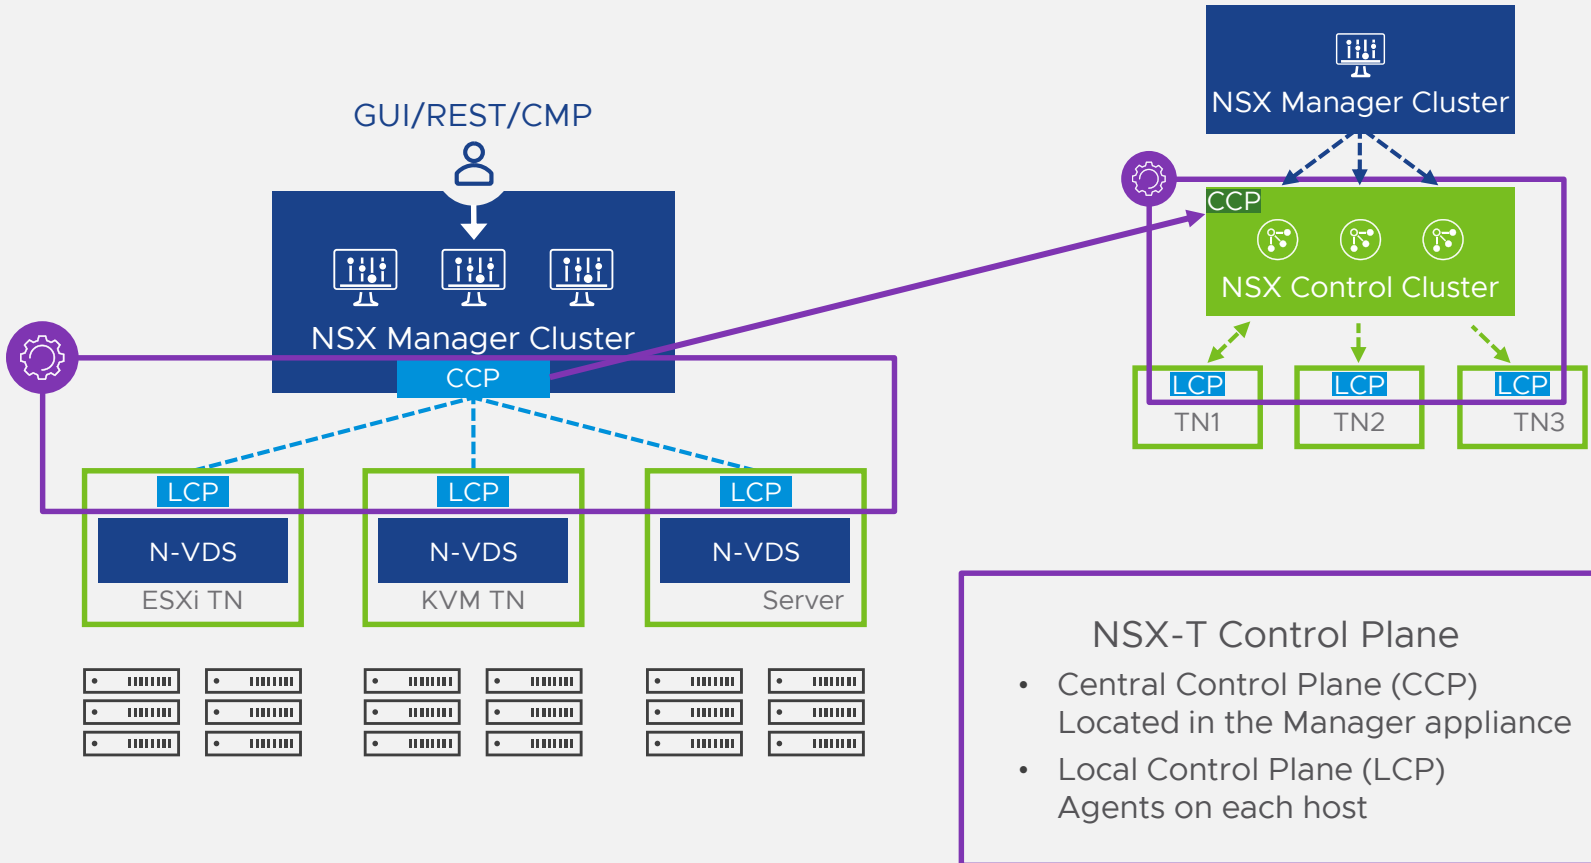
Enables high availability of the NSX UI and API and the NSX Manager.

Reduces the likelihood of failures of operation of NSX.

Provides API and GUI clients with multiple endpoints or single VIP for availability

# NSX-T Control Plane

## Centralized & Local Control Plane



NSX-T Central Control Plane (CCP) computes all ephemeral runtime states based on configuration from the management plane.

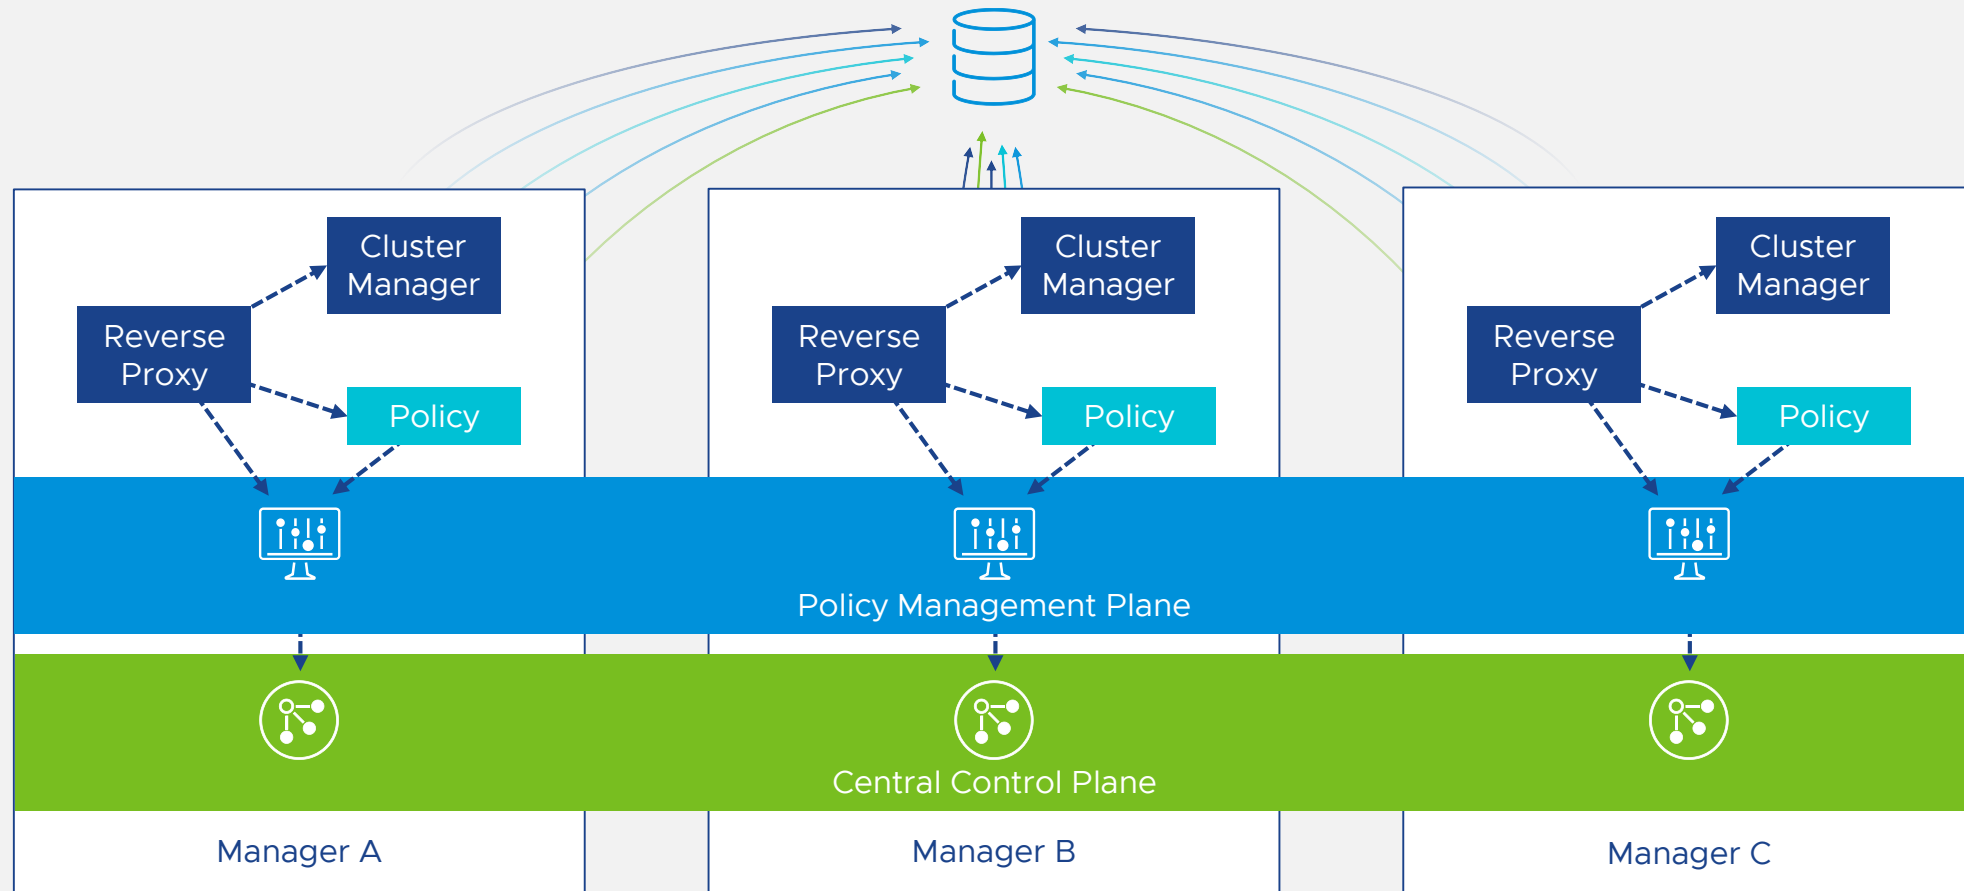
CCP Pushes stateless configuration to the transport nodes. The Local Control Plane (LCP) agents receive the configuration and push it into the data plane of the transport node.

Disseminates topology information reported by the data plane elements



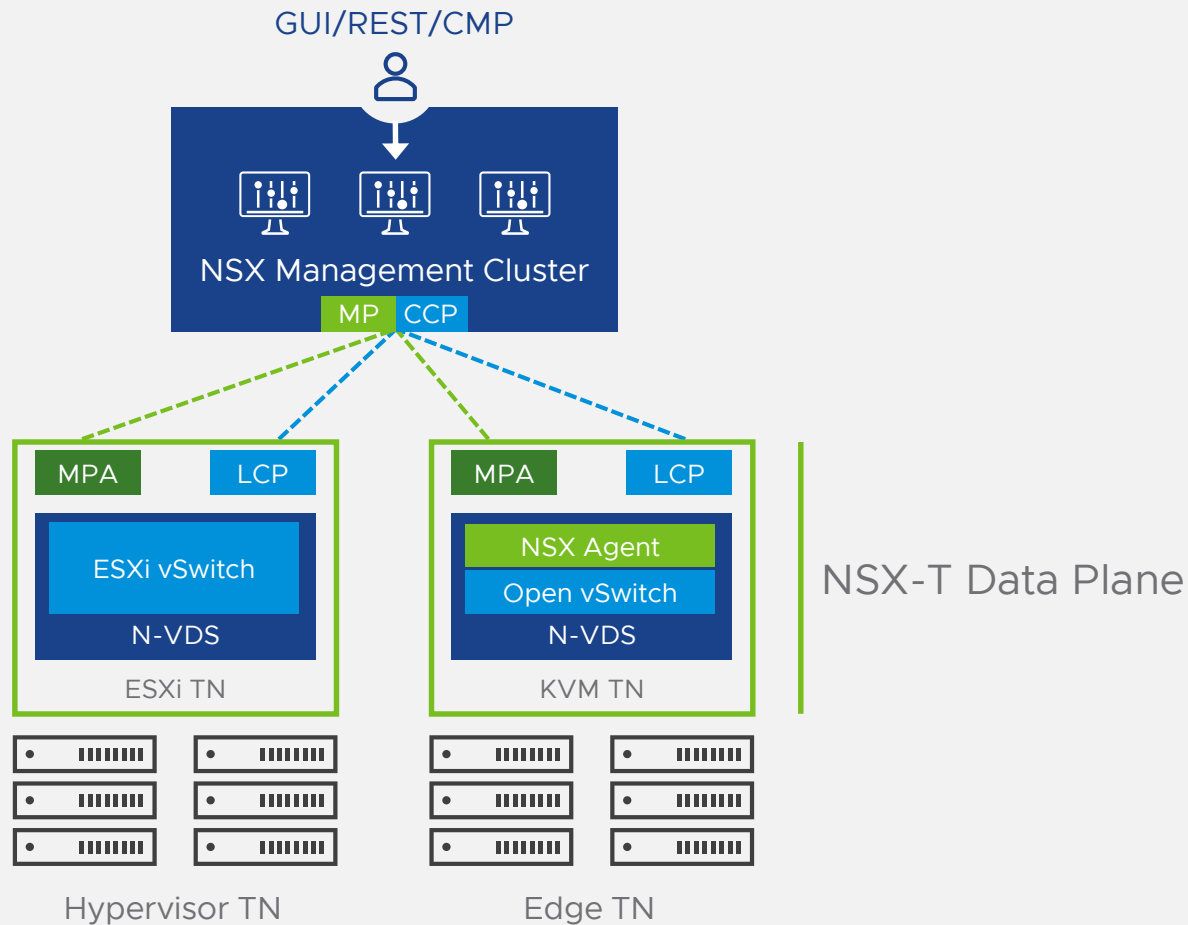
# NSX-T Management Cluster

High Level of the NSX-T Management architecture



# NSX-T Data Plane

## Transport Nodes



The Data Plane on each node is a high performance engine for logical switching, routing, and distributed firewall. The data plane encapsulates and decapsulates packets for the overlay network.

The data plane of each host contains the following elements:

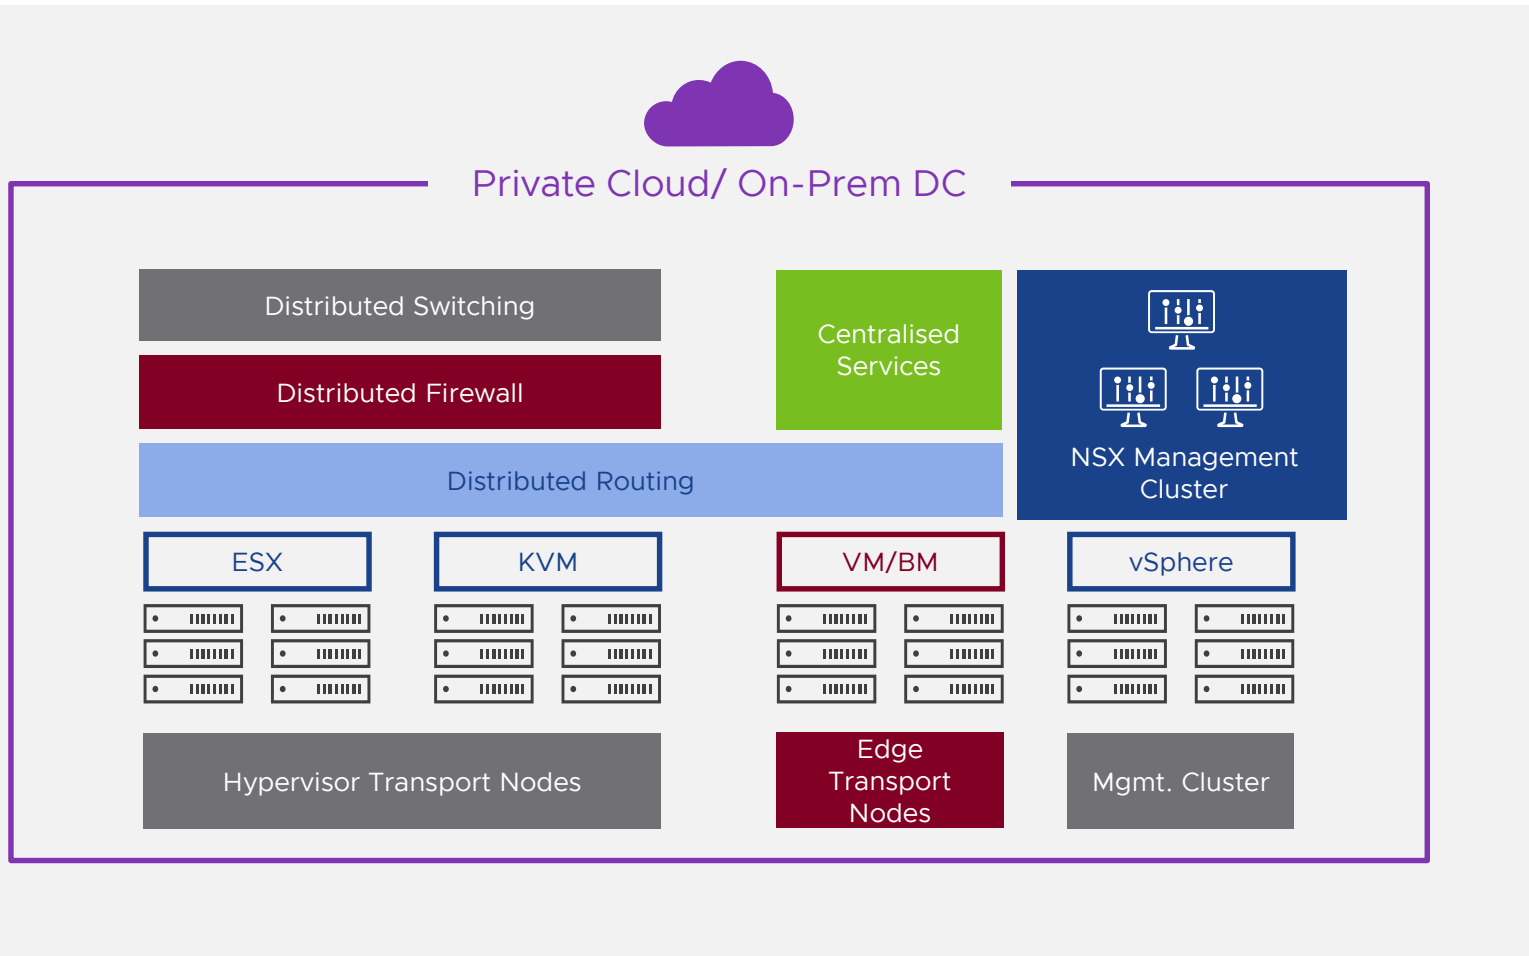
- Local Control Plane (LCP) agent
- Management Plane Agent (MPA)
- NSX Virtual Distributed Switch

The N-VDS is based on either:

- ESXi vSwitch for ESXi
- Open vSwitch (OVS) for KVM

# NSX in the Data Center

## NSX functional groups deployment



NSX-T best practices for deployment is a simplified structure of functional groups usually deployed as host clusters.

The **Management Cluster** contains NSX manager, vCenter, automation tools, day two operational tools and dependent shared services (LDAP, DNS, NTP, etc).

The **Compute cluster** high performance data plane kernel hosts the workloads. Communication is East/West for the workloads and connecting to the Edge.

The **Edge Cluster** provides north/south routing for the external to the NSX Overlay. The edge nodes are either VM or bare-metal form factors. The Edge cluster hosts the centralized (stateful) services such as Firewall, Load Balancer, IPAM, VPN, etc.



Thank You