Telstra Private Cloud

User Guide

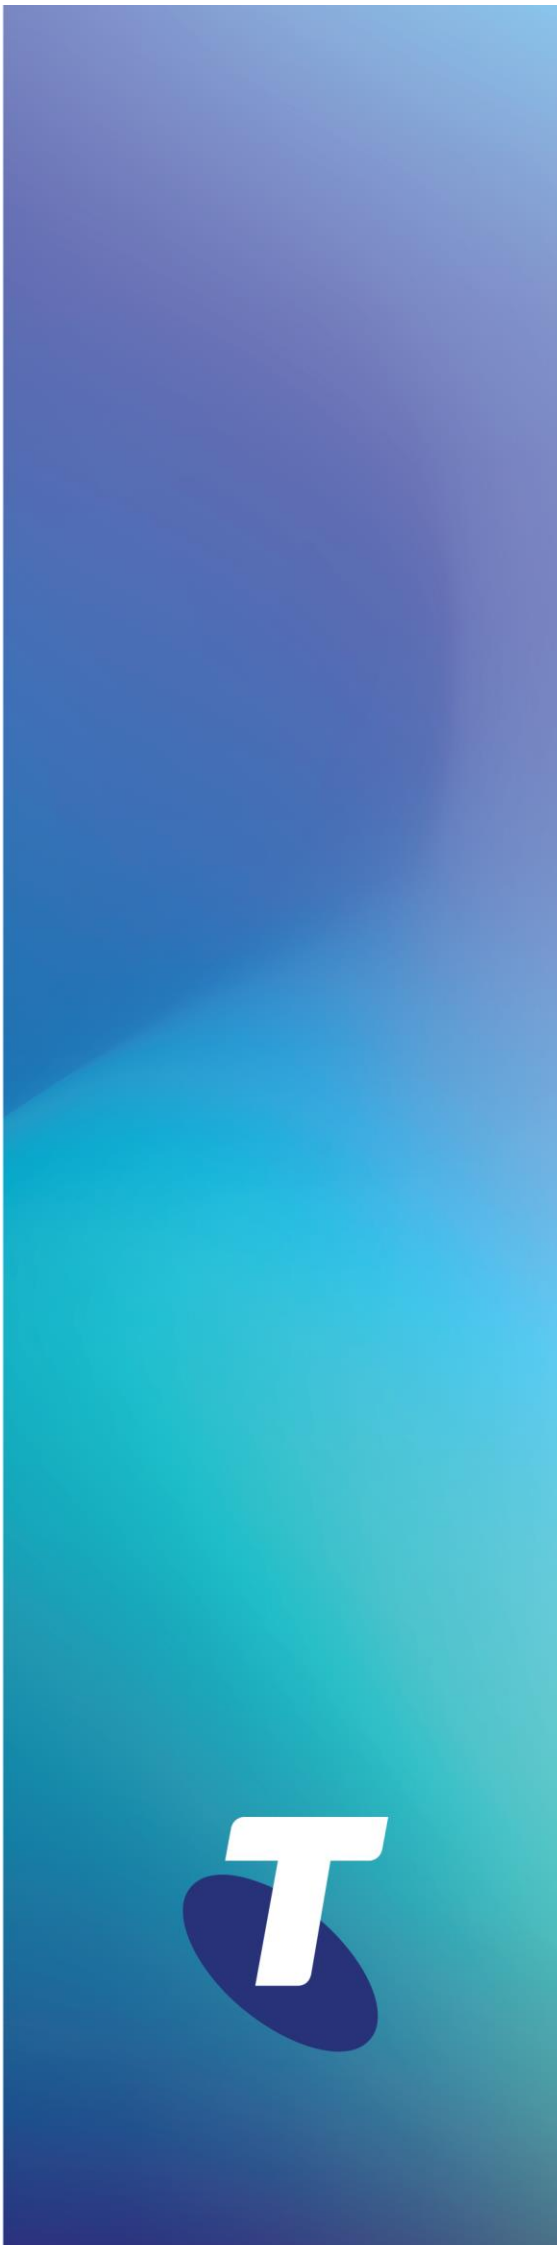

Welcome to the Telstra Private Cloud: User Guide

About Telstra Private Cloud Guides

Telstra has produced several guides to help you order, manage and configure Telstra Private Cloud (TPC):

Telstra Private Cloud: Administration Guide	Explains TPC's operating environment and provides detailed instructions and tips for you to order, access and manage your TPC virtual data centre (vDC) and its physical resources. It discusses: • General TPC concepts • TPC-related tasks using Telstra Cloud Sight • Submitting orders for tenancy resources and other administrative requests using Telstra's TPC vCenter Plug-in
Telstra Private Cloud: User Guide	Discusses TPC terminology as well as technical and topological concepts for your vDC. It also covers important configuration tasks completed using these VMware tools: • vCenter Portal • NSX-T Manager Portal
Telstra Private Cloud: Quick Reference Guides	Each guide lays out the sequence of tasks required to complete a common provisioning activity, such as: • How to add a virtual server to a public network • How to add a virtual server with new storage.

Which Guide Do I Need?

You can use the following table to help direct you to the correct guide. You can also find additional information on telstra.com.

I need information about…	Use this guide first
Purchasing a TPC tenancy (vDC)	Telstra Private Cloud: Administration Guide
Adding TPC administrators for vCenter or NSX-T	Telstra Private Cloud: Administration Guide
Providing SSL VPN Access to TPC administrators	Telstra Private Cloud: Administration Guide
Adding hosts to my TPC vDC	Telstra Private Cloud: Administration Guide
Adding storage to my TPC hosts	Telstra Private Cloud: Administration Guide
Requesting additional public IP addresses for my vDC	Telstra Private Cloud: Administration Guide
vDC topologies and connections	Telstra Private Cloud: User Guide
Using NSX-T to add, connect or modify logical network entities (switches, gateways, routers, distributed firewall, etc)	Telstra Private Cloud: User Guide
Using vCenter to create VMs, DRS groups, VM-Host affinity rules	Telstra Private Cloud: User Guide
The portals used and steps required to complete common or complex jobs	Telstra Private Cloud: Quick Reference Guides

For information on Telstra's older private cloud products such as Virtual Server (Dedicated) Gen 1, Gen2 or Gen 2+, refer to their documentation on telstra.com.

Telstra Private Cloud User Guide, v1.0, June 2022

Terminology Used in TPC Guides

Job	Equivalent to a Use Case in Telstra Private Cloud. A Job consists of a set of Tasks that collectively achieve a configuration goal for a customer. We cover common or complex jobs in TPC Quick Reference Guides
Task	One self-contained component of a Job. A given Task might occur within many different Jobs.
Procedure	The set of steps that make up a task.

Example Job: Add a Publicly-Reachable Web Server to a Tenancy Using a New VM on New Host with New Storage

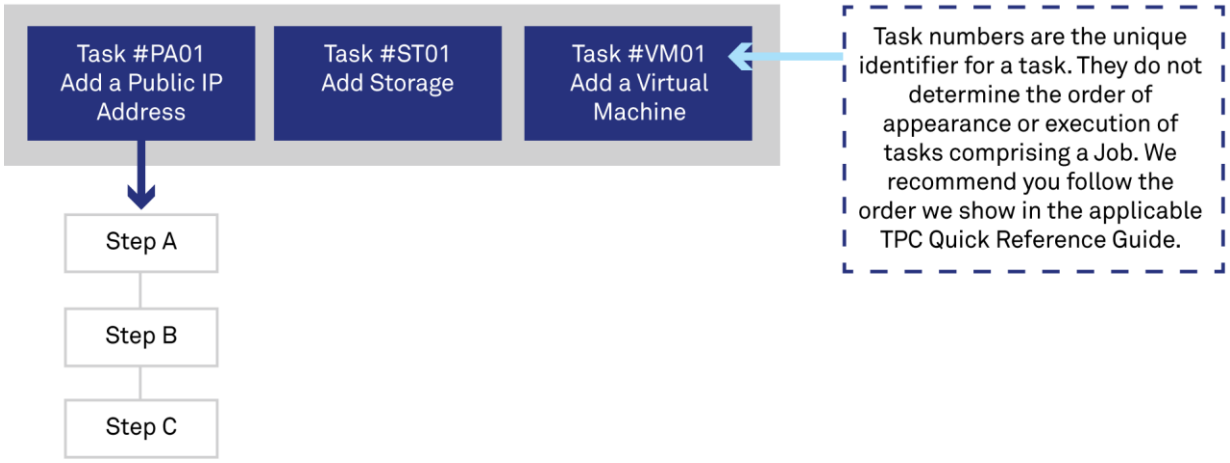

What's Inside

Chapter 1: External Interconnects to Your vDC 7
Types of External Interconnects 7
Public Interconnect 8
Private Interconnect 9
Public Interconnect 10
Identifying Your Public Tier-0 Gateway 10
Addresses on the Public Interconnect 11
Public IP Address Fees 12
Obtaining Usable Public IP Addresses 12
Public Interconnect Routing 12
Security Considerations for the Public Interconnect 14
Private Interconnect 15
Identifying Your Private Tier-0 Gateway 15
Private Interconnect Addressing 16
Private Interconnect Routing 17
Advertising Additional vDC Routes to Next IP 17
Chapter 2: vDC Topologies 19
Basic Topologies 21
Public 21
Private 23
Complex Topologies 25
Chapter 3: NSX-T Manager Tasks 27
Task #NS01: Add a Segment 27
Task #NS02: Modify an Edge Node 31
Task #NS03: Add/Modify a Tier-1 Gateway (Logical Router) 34
Task #NS04: Add/Modify NSX-T Gateway Firewall 37
Task #NS05: Add/Modify NSX-T Distributed Firewall 41
Task #NS06: Add a Layer-2 VPN 44
Task #NS07: Add a Load Balancer 47
Task #NS08: Add a NAT Rule 50

Chapter 4: vCenter Tasks 52
Task #VM01: Create a Virtual Machine 52
Task #VM02: Create a VM DRS Group 54
Task #VM03: Create a Host DRS Group 57
Task #VM04: Create a VM-Host Affinity Rule 60
Task #VM05: Create a VM-VM Affinity Role 63
Chapter 5: Guidelines for Common Jobs 66
Job: Add a VM to a Public Network 66
Job: Add a VM to a Private Network 67
Chapter 6: Resource Sizing Considerations 69
TPC Host Characteristics 69
Rightsizing Your VMs 70
Host-VM CPU Contention 72
RAM Consumption 73
Use VMware Tools 74
Resource Locking and HA/DRS 75
Power-on Priority 76
Storage 76
Software Licences 77
Chapter 7: Other Considerations 78
Integrating Advanced Management Tools and Software 78
Planning Migration to or from TPC 78

Chapter 1: External Interconnects to Your vDC

Using vCenter, NSX-T and Telstra's vCenter Plug-in, you can define, arrange and connect the resources within your vDC in your preferred topology, as well as define and implement the firewall rules to protect them.

You cannot physically build or supply your own externally-facing links from your vDC. Rather, Telstra has pre-installed fast, highly reliable connections from our TPC facilities to our world-class network infrastructure. After we hand over your vDC, you can complete self-service actions to activate and/or secure its external connections.

Types of External Interconnects

Your vDC will typically communicate with one or more outside networks or services, such as:

- The Internet
- A Next IP VPN
- Other Telstra value-added products and infrastructures and/or
- Third-party public cloud providers.

Traffic will cross an external interconnect as it leaves or enters your vDC. There are two external interconnects relevant to TPC, each containing specific links and physical and logical devices. Depending on which outside networks are involved, your traffic will use one or the other, called a ***Public Interconnect*** and ***Private Interconnect*** respectively. Your vDC may have only a Public Interconnect or may have both.

Public Interconnect

If you want to communicate between your vDC and any destination via the Internet, your traffic will use a Public Interconnect. The Public Interconnect joins your vDC to the Internet over dual-redundant, high-speed uplinks to Telstra Internet Direct (TID).

When we provision your TPC tenancy, we create a logical device called a ***Public Tier-0 Gateway*** along with its corresponding Edge nodes in your vDC. You will be able to see the Public Tier-0 Gateway and modify it using NSX-T when you receive management access to your tenancy.

Figure 1: Public Interconnect

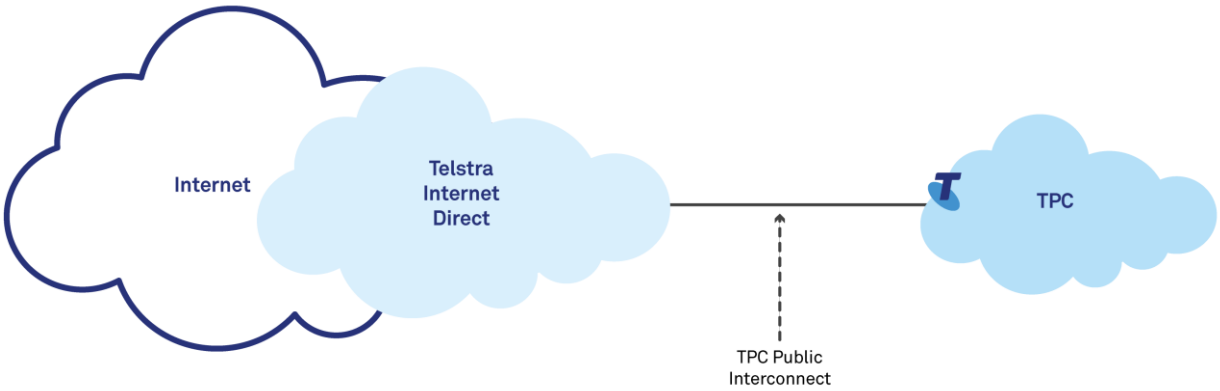

Private Interconnect

If you want to communicate between your vDC and sites on or reached through your Next IP VPN, you will need a Private Interconnect. The Private Interconnect joins your vDC to an auxiliary network service called ***Telstra Cloud Connector***. In addition to Next IP, a Cloud Connector can enable private connectivity between your vDC and supported public cloud services like AWS and Azure.

If you don't already have a Cloud Connector, you need to order one and link it to your Next IP VPN and to your TPC vDC before you can use your Private Interconnect. If you do already have a Cloud Connector, you need only to link it to your TPC vDC to be able to use it. We discuss how to order and link Telstra Cloud Connector to your vDC in the Telstra Private Cloud: Administration Guide, available here.

When we provision your TPC tenancy, we create a logical device called a ***Private Tier-0 Gateway*** along with its corresponding Edge nodes in your vDC. You will be able to see the Private Tier-0 Gateway and modify it using NSX-T when you receive management access to your tenancy. Later, when you link your Cloud Connector to TPC, we modify the configuration of your Private Tier-0 Gateway to activate the Private Interconnect and make your Next IP VPN and your vDC visible to each other.

Figure 2: Private Interconnect

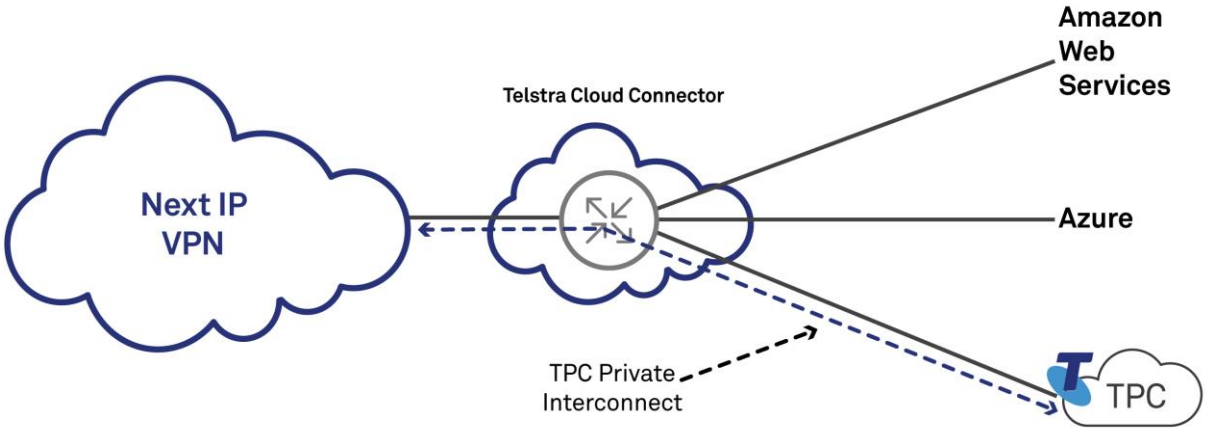

Public Interconnect

Every TPC tenancy receives a Public Interconnect running between TID and your vDC. It consists of a pair of VLANs over a redundant switching fabric, each running from an Edge node in your vDC to a Telstra Public Router. Those Edge nodes are then associated with your Public Tier-0 Gateway.

At the time we hand over your vDC, your Public Interconnect is active but you will not be able to use it until:

1. You obtain a range of public IP addresses from Telstra
2. We complete subsequent provisioning activity in the Telstra Public Routers, and
3. You configure your vDC to use the public addresses we assigned to you.

The range we provide is specific to the TPC vDC for which you ordered it. You cannot use it anywhere outside of that vDC. Moreover, you cannot use your own public addresses in your vDC, nor use Telstra-supplied ranges obtained from our other products (including from TID itself).

Identifying Your Public Tier-0 Gateway

When you log in to NSX-T with the *networkadmin* username, you can see and modify resources in NSX-T, including the Public Tier-0 Gateway. Telstra encourages you to avoid changing its configuration unless necessary to support your downstream topology, because:

- Changing the configuration risks introducing accidental or unanticipated side-effects on your external connectivity, and
- Telstra employs automation to provision moves/adds/ changes (MACs) you have ordered through the Telstra vCenter Plug-in. If you have adjusted your Public Tier-0 Gateway configuration in the past, our provisioning systems may unexpectedly reset your changes while completing a later MAC. You may wish to develop a plan and process beforehand to capture the details of your changes so you can re-apply them later if needed.

When administering NSX-T, you can recognise the Public Tier-0 Gateway by its name, which is always:

`t0-pub01`

A key piece of text to notice in the name is `pub`, which indicates this is the Public Tier-0 Gateway rather than the Private Tier-0 Gateway.

To provide your vDC with resilient operation, we always configure the Public Tier-0 Gateway in a HA active/standby configuration.

Addresses on the Public Interconnect

The Public Tier-0 Gateway and Telstra Public Routers at each end of the VLANs that comprise the Public Interconnect use IP addresses Telstra has selected from a range that is reserved by IANA specifically for carrier use. The range is contained in RFC 6598 and drawn from 100.64.0.0/10. You will be able to see the addresses and masks we have used for your Public Tier-0 Gateway in your NSX-T portal.

While not identical to the RFC 1918 private addressing ranges you are probably familiar with, the carrier-specific range has similar characteristics because it can be used by any carrier and is not globally routable nor advertisable on public links. We chose to use this range for the Public Interconnect because we expect it to be compatible with your existing network.

Figure 3: Addresses on the Public Interconnect

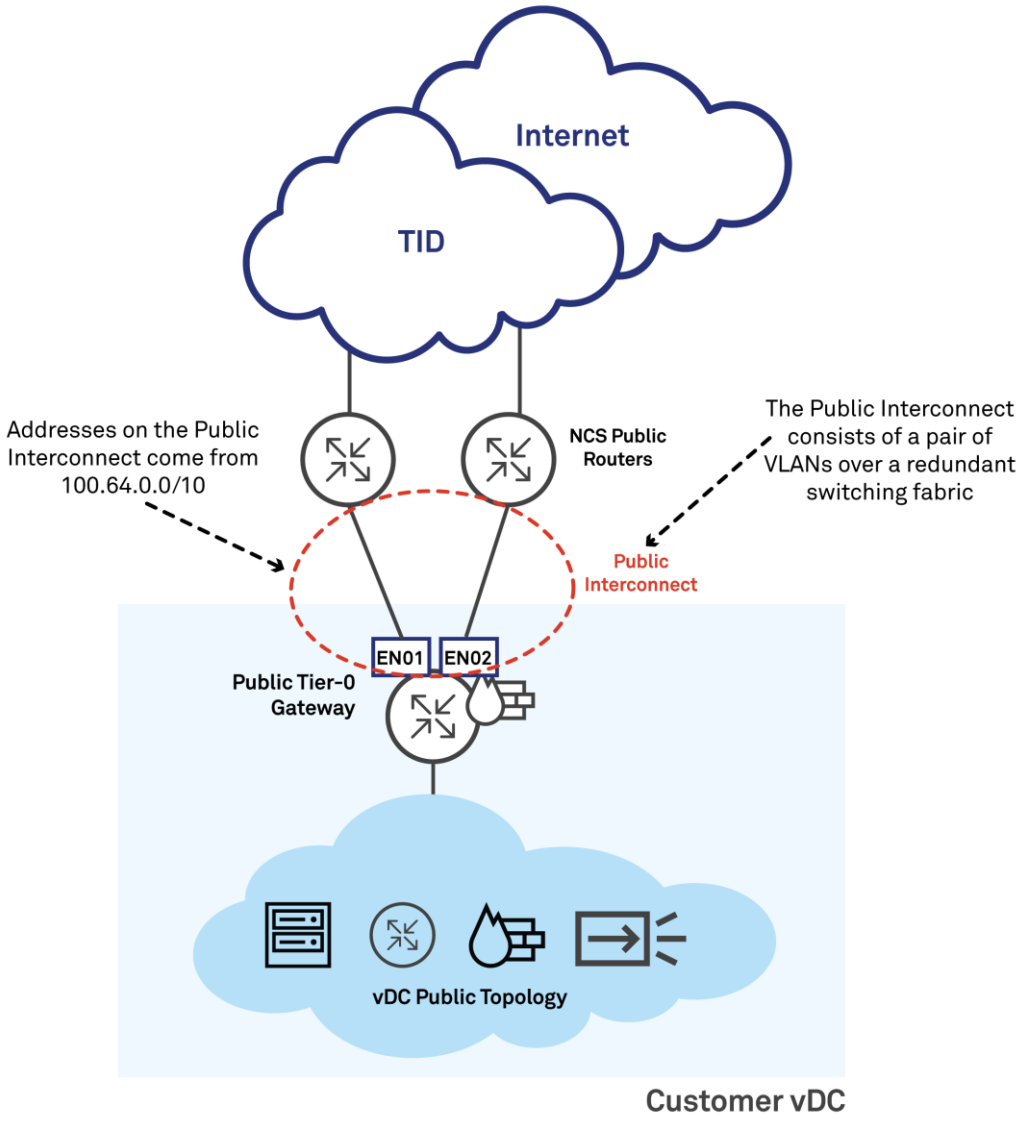

Public IP Address Fees

All public IP addresses you order for your vDC are subject to monthly fees (refer to the TPC Pricing Guide for more information on pricing).

TPC includes five complimentary addresses with each vDC. Each month, Telstra will credit your TPC account with the amount charged for 5 addresses (ie. equivalent to the number of usable addresses in one /29 subnet in TPC). If you order a larger subnet, or more than one, TPC will bill you for the difference based on the number of usable addresses involved.

Obtaining Usable Public IP Addresses

TPC does not automatically provision your complimentary public IP addresses. At any time after we build your vDC and hand it over to you, you can use the TPC vCenter Plug-in to request a subnet (or more than one of them) of the necessary size. Once TPC automation has reserved your subnet, configured the routing of that subnet towards your vDC and advised you of the subnet's details, you can commence using it in your vDC and configure any required intra-vDC routing.

We explain how to order public IP address ranges for TPC in Task #PA01 in the Telstra Private Cloud: Administration Guide, available here.

Public Interconnect Routing

The Public Tier-0 Gateway and Telstra Public Routers statically route traffic between them. We configure your Public Tier-0 Gateway with a default route pointing at the Telstra Public Routers. We also configure your vDC Edge nodes to use Bi-directional Forwarding Detection (BFD) to quickly identify path disruptions and re-direct traffic flows away from a failed link.

After you order a Public IP Address Range for your vDC using the Telstra vCenter Plug-in, our automated provisioning system will reserve a suitable range and then configure our Telstra Public Routers with a route to it via your Public Tier-0 Gateway. After provisioning completes, you can configure your Public Tier-0 Gateway and/or vDC topology to appropriate forward and/or use the range and make appropriate adjustments to your vDC security policy.

If you subsequently order another range, during provisioning we will configure another static route in the Telstra Public Routers pointing at your Public Tier-0 Gateway.

Figure 4: Public Interconnect Routing

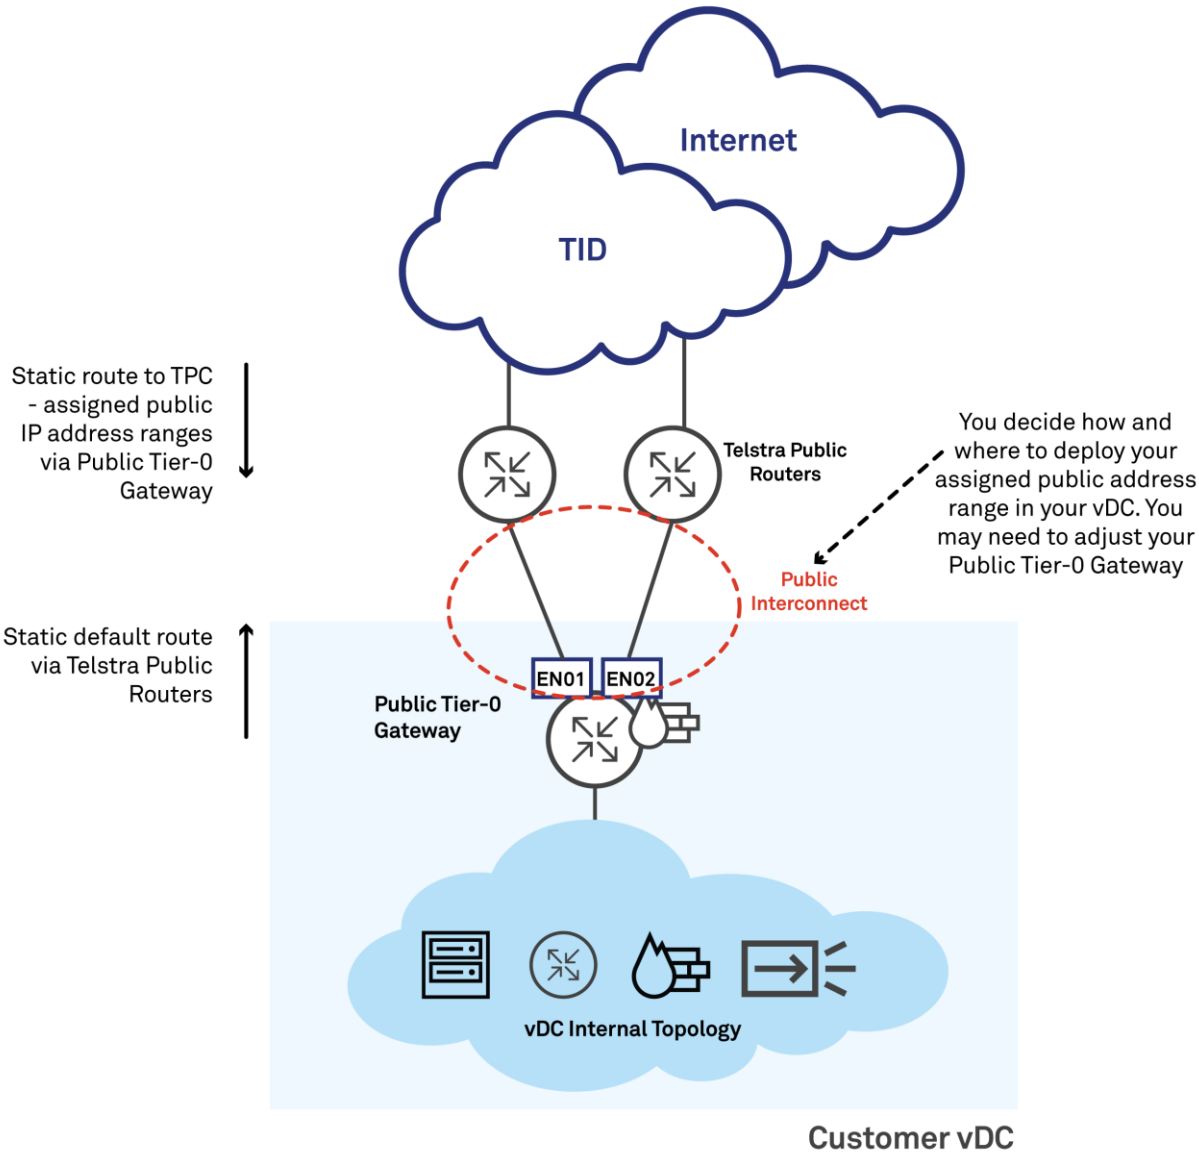

Security Considerations for the Public Interconnect

It is up to you to implement appropriate firewall rules in your vDC to protect it from threats. This includes your Public Tier-0 Gateway and all resources reachable in your vDC behind it. You must assume the Public Interconnect will behave like any open Internet access service.

You can use the inbuilt NSX-T Firewall or a VM containing your preferred virtual firewall appliance to configure and apply security in your vDC.

What is the NSX-T Firewall?

The NSX-T Firewall provides logical security mechanisms for your vDC. It broadly categorises its policy application into two parts:

- Security between workloads and applications: known as the Distributed Firewall (DFW), this deals with 'east-west', stateful firewall inspection (SFI) security. Conceptually there is only one DFW in each tenancy, which we automatically enable for you in every host cluster during provisioning. Because it runs in a distributed fashion in each hypervisor's kernel, the DFW will intelligently apply your security policy at an interface level on every VM in your vDC
- Perimeter security: known as a Gateway Firewall, this refers to 'north-south' security that offers stateful inspection at the edge of the vDC.

You can review your rules using the Security Overview Dashboard within the NSX-T portal. You determine where to apply your rules and their precedence when you define or modify them in NSX-T.

Within this guide, we will typically use the term NSX Firewall to cover both firewall types unless we need to discuss a specific or unique characteristic of one or the other.

Private Interconnect

The Private Interconnect runs between your Private Tier-0 Gateway and your Next IP VPN over a Telstra Cloud Connector. It consists of four parallel VLANs each addressed using /30 subnets drawn from the /24 private IP range you supplied when you ordered your original vDC. The private interconnect runs eBGP to exchange routes with your Next IP VPN.

Telstra initialises the Private Tier-0 Gateway when we provision your vDC, but at that point the vDC does not have a Private Interconnect. Immediately following provisioning, the Private Tier-0 Gateway only provides the path for management traffic between your dedicated versions of vCenter and NSX-T and your vDC's hosts.

Our provisioning tools build the Private Interconnect and apply the addresses to the endpoints after you link your Cloud Connector to your vDC. At the vDC, those addresses sit on the Private Tier-0 Gateway, each applied to an interface contained in one of a pair of Edge nodes that land the VLANs.

Identifying Your Private Tier-0 Gateway

When you log in to NSX-T with the *networkadmin* username, you can see and modify resources in NSX-T, including the Private Tier-0 Gateway. Telstra encourages you to avoid changing its configuration unless necessary to support your downstream topology, because:

- Changing the configuration risks introducing accidental or unanticipated side-effects on your external connectivity, and
- Telstra employs automation to provision moves/adds/ changes (MACs) you have ordered through the Telstra vCenter Plug-in. If you have adjusted your Private Tier-0 Gateway configuration in the past, our provisioning systems may unexpectedly reset your changes while completing a later MAC.

When administering NSX-T, you can recognise the Private Tier-0 Gateway by its name, which is always:

`t0-prv01`

A key piece of text to notice in the name is `prv`, which indicates this is the Private Tier-0 Gateway rather than the Public Tier-0 Gateway.

To provide your vDC with resilient operation, we always configure the Private Tier-0 Gateway in a HA active/standby configuration.

Private Interconnect Addressing

When we build your vDC, we subdivide the /24 range you provided when you ordered your vDC and use it to address various resources and interfaces in your tenancy. If you later link your vDC to Cloud Connector to reach Next IP, we use four /30 subnets from that range to address the VLANs comprising the Private Interconnect.

Figure 5: Private Interconnect Addressing

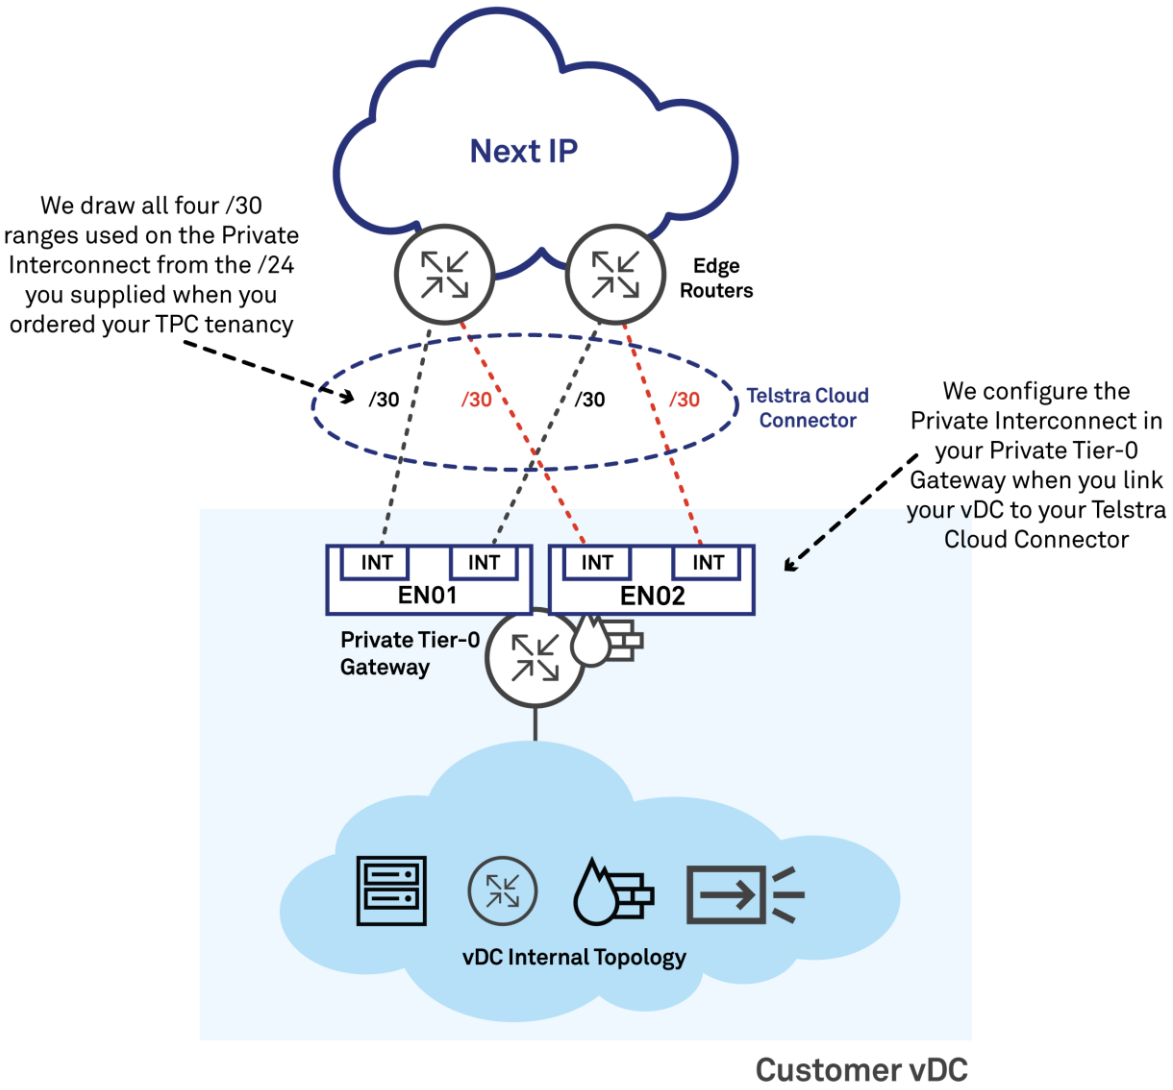

Private Interconnect Routing

The Private Tier-0 Gateway and Next IP VPN routers dynamically route traffic between them using eBGP. We configure your Private Tier-0 Gateway with a BGP ASN you supplied when you linked your Cloud Connector to your vDC. The Private Interconnect also uses Bi-Directional Forwarding Detection (BFD) to rapidly detect path failures and adjust routing.

The Edge nodes for the Private Tier-0 Gateway operate as a pair of NSX-T Service Routers (SRs) in an active/standby high-availability mode. Even though they operate in an active/standby arrangement, our configuration causes both Service Routers (ie. the active SR and the standby SR) to regularly exchange routing information with the Next IP VPN routers. In order to ensure that only the active SR actually handles traffic to and from Next IP under normal conditions, the standby SR prepends three additional copies of its ASN in the AS-PATH attribute. If the active SR or its connections should fail, the standby SR will become active.

The Private Tier-0 Gateway resides on a host in your vDC. You can see it and inspect or change its configuration, but you need to be careful not to make a mistake or try to configure a feature or setting we do not support. This is because you could accidentally disrupt:

- The connectivity between your vDC and your Next IP VPN, or
- Management access to your vDC using TPC's SSL VPN.

Advertising Additional vDC Routes to Next IP

When we build your vDC, we subdivide the /24 range you provided when you ordered your vDC and use it to address various resources and interfaces in your tenancy. If you later link your vDC to Cloud Connector to reach Next IP, we use four /30 subnets from that range to address the VLANs comprising the Private Interconnect. You do not need to take any further action to advertise the /24 range in your network.

However, depending on your vDC's topology, you are likely to use one or more additional private ranges within your vDC for resources you later define and operate, like VMs, Tier-1 routers and so on. In order to reach these resources from your Next IP VPN, you will need to advertise their subnets over the Private Interconnect using BGP.

To do this, you can configure your Private Tier-0 Gateway to advertise your vDC's private IP address ranges to their peering Next IP edge routers, which will advertise them into the rest of the Next IP VPN.

Telstra's infrastructure does not filter outbound advertisements from the Dedicated Private Tier-0 Gateway to the Next IP routers. In principle you can advertise any ranges you like, including the default route and public IP ranges, and Next IP will accept and propagate them.

However, you must be mindful of the impacts on upstream networks. For example, while Next IP will usually accept a default route (0.0.0.0/0), advertising it from TPC might cause unexpected behaviour if your VPN Internet gateway is meant to sit elsewhere. Next IP also restricts certain other loopback and link-local ranges.

Figure 6: Advertising Additional vDC Routes to Next IP

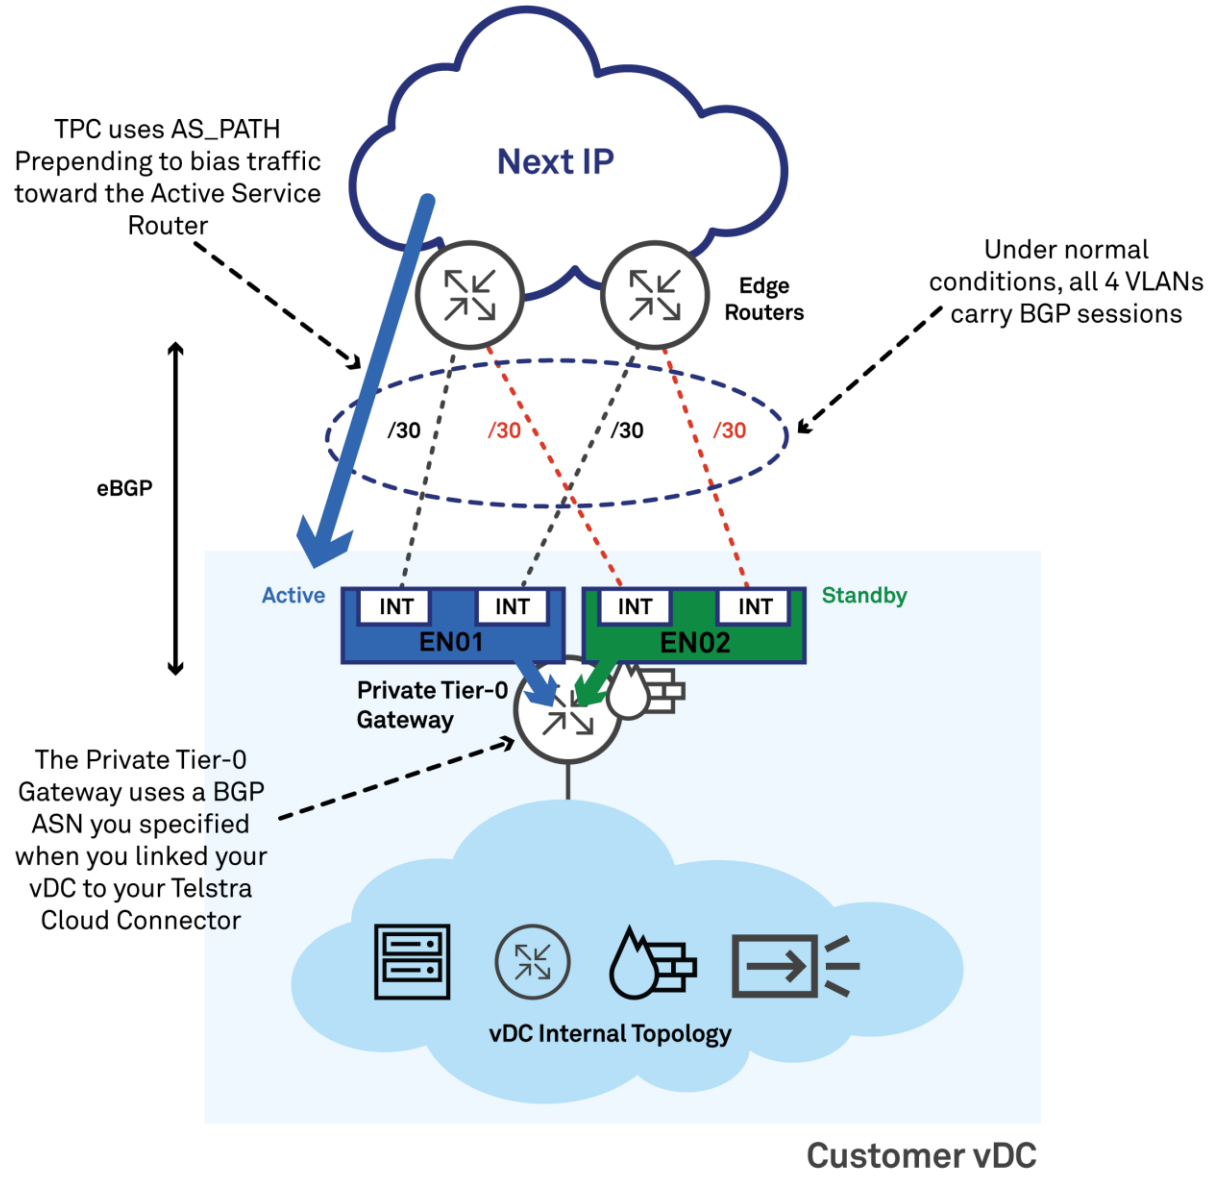

Chapter 2: vDC Topologies

You can configure your vDC's internal topology to suit your needs. Telstra does not explicitly restrict the number of VMs or certain other resources you deploy in your vDC, but you should observe VMware's specifications[1], Telstra's product capabilities and rules, and the practical limits of your physical and logical resources.

While TPC allows you to define, arrange and connect the logical resources inside your vDC as you see fit, Telstra does impose certain limitations to meet these special circumstances:

1. Restrictions necessary to protect the general integrity of our customers' tenancies and our infrastructure
2. Routing rules that control how your vDC connects to external networks through the Public and Private Interconnects
3. Administrative account conditions (the number of accounts, password rules, etc.) in vCenter and NSX-T.

In Telstra's experience over many years with CSX-based services, our customers tend to divide their vDCs into public and private topologies, logically linked but monitored using inbuilt or third-party virtual firewalls, load-balancers, VPN concentrators and/or other virtual devices that support operational needs. For TPC, that could result in a conceptual vDC layout that looks something like Figure 7.

[1] Refer to the VMWare Configuration Maximums tool for further information and support

Figure 7: vDC Topologies

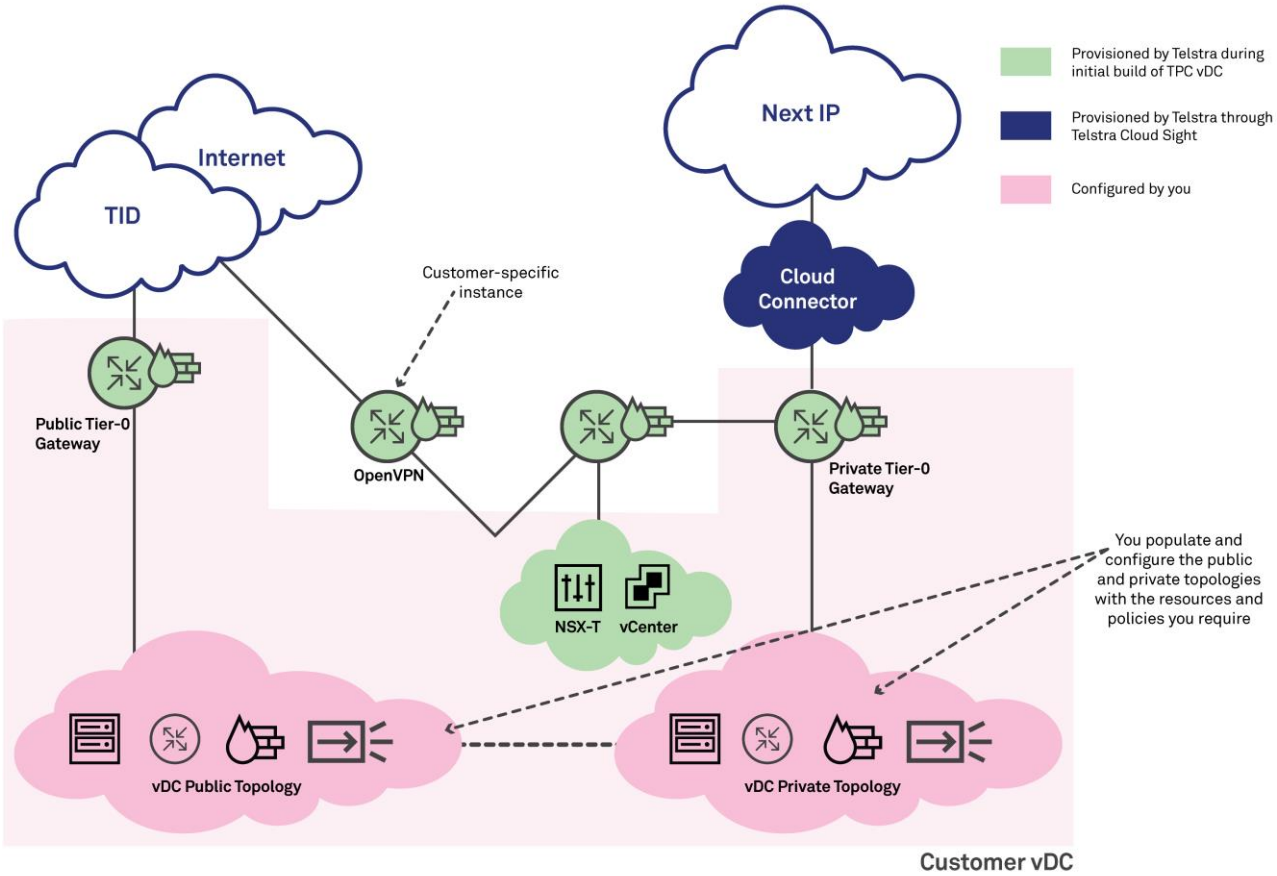

It is up to you to carefully consider the security, performance, resiliency and cost impacts of your topology. For example, if you make a resource reachable from both the Public Interconnect and Private Interconnect, you have created a path to your vDC and downstream private networks from the Internet, so you will need to implement effective security measures to protect them. Or, if you decide to request a larger range of public IP addresses for your vDC than you really need, you might needlessly incur higher fees.

Basic Topologies

Public

A basic public topology is one that exclusively communicates with external parties over the Public Interconnect. It is up to you to determine the appropriate public network topology and addressing plan.

In order to make your VMs visible to the Internet over the Public Interconnect, you will need an appropriate range of public IP addresses to assign to them, either directly to a logical interface or by configuring NAT in an intermediate gateway. You can request an additional range(s) from /29 up to /26 by submitting a request using Telstra's vCenter Plug-in.[2]

Telstra will configure the Public Tier-0 Gateway with a default route pointing to the Telstra Public Routers when we build your vDC. We will also configure the Telstra Public Routers with a route to each public IP address range assigned to your vDC following your order submission. We show the standard routing configuration for the Public Interconnect in.

Figure 8: Basic Public Topology

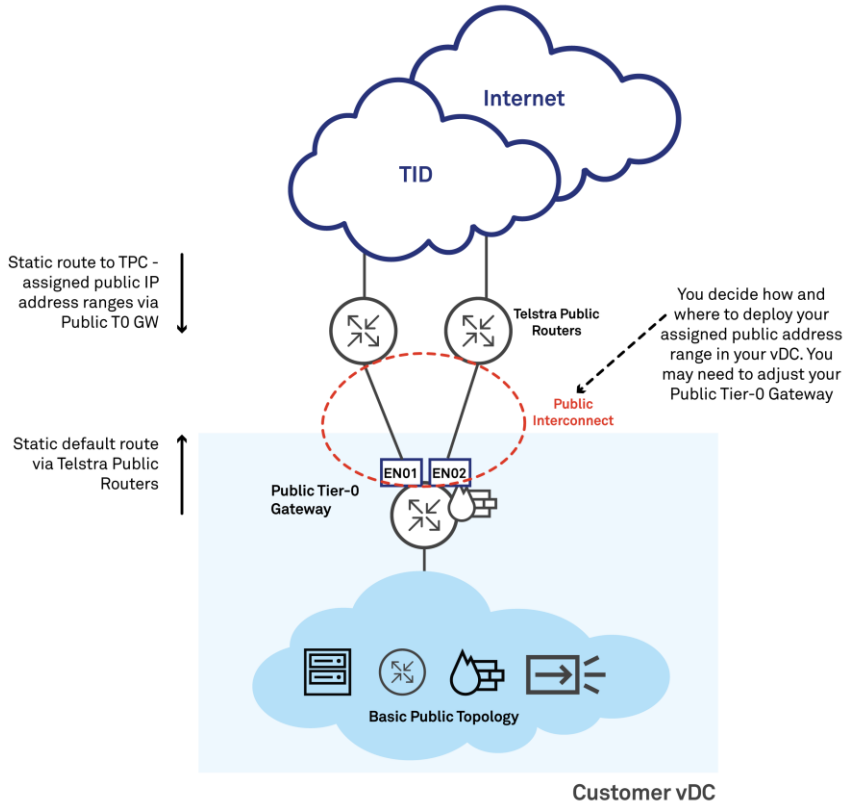

[2] As we discuss in the Telstra Private Cloud: Administration Guide (available here), you will receive five complimentary public IP addresses with your vDC through a credit offset on your TPC bill. In order to obtain and use them, you need to order them (along with any additional addresses you may need) through the Telstra vCenter Plug-in

Using Public Address Ranges Assigned to My vDC

When you request a public addressing range for your vDC, you are free to use the addresses from that range in ways that meet your needs for your public topology and complies with networking practices and rules. This includes one or more of the following options:

- Applying the addresses directly to interfaces on VMs in your public topology
- Using one or more addresses for load balancing in NSX-T
- Employing source and/or destination NAT in your gateway(s) to allow Internet breakout or server reachability further into your vDC
- Further subnetting your range to divide and apply it across different topological segments
- Forwarding that range or a subnet further into your public topology.

Example

In Figure 9 we show a basic public topology for example 'Customer A'. It shows the Public Tier-0 Gateway and a small number of Internet-facing servers sharing a common LAN segment. The servers each use a public IP address drawn from a /29 range Customer A has requested through the Telstra Plug-in. After we assigned the public IP range to the vDC in response to Customer A's request, we configured the static route in the Telstra Public Routers, pointing at the Public Tier-0 Gateway.

Customer A has configured the inbuilt NSX-T Distributed Firewall to help secure the servers. It is entirely up to Customer A to determine whether this level of security is sufficient for their needs or further measures are necessary or advisable. For example, Customer A could supplement the NSX-T DFW with a third-party, VM-based next generation firewall (NGFW) appliance to implement more advanced security functions. A third-party cloud-based service might be another option. Telstra offers a number of comprehensive security solutions that may satisfy this requirement, some of which deploy logical security appliances while others are cloud-based.

In order to build this topology, Customer A will:

- Use the vSphere Telstra Plug-in to request a /29 public IP address range. (Refer to the Telstra Private Cloud: Administration Guide for more information.)
- Create a Logical Segment and attach it to Customer A's Public Tier-0 Gateway
- Configure one of the public reachable addresses on the Logical Segment interface to the Public Tier-0 Gateway. This address will be the default gateway for the servers
- Configure a pair of VMs to act as the publicly reachable application servers, connect them to the Logical Segment, and assign one of the public addresses to each of them
- Suitably configure the NSX-T Distributed Firewall
- Configure any other security or applications necessary to achieve Customer A's objectives for the vDC.

Figure 9: Basic Public Topology

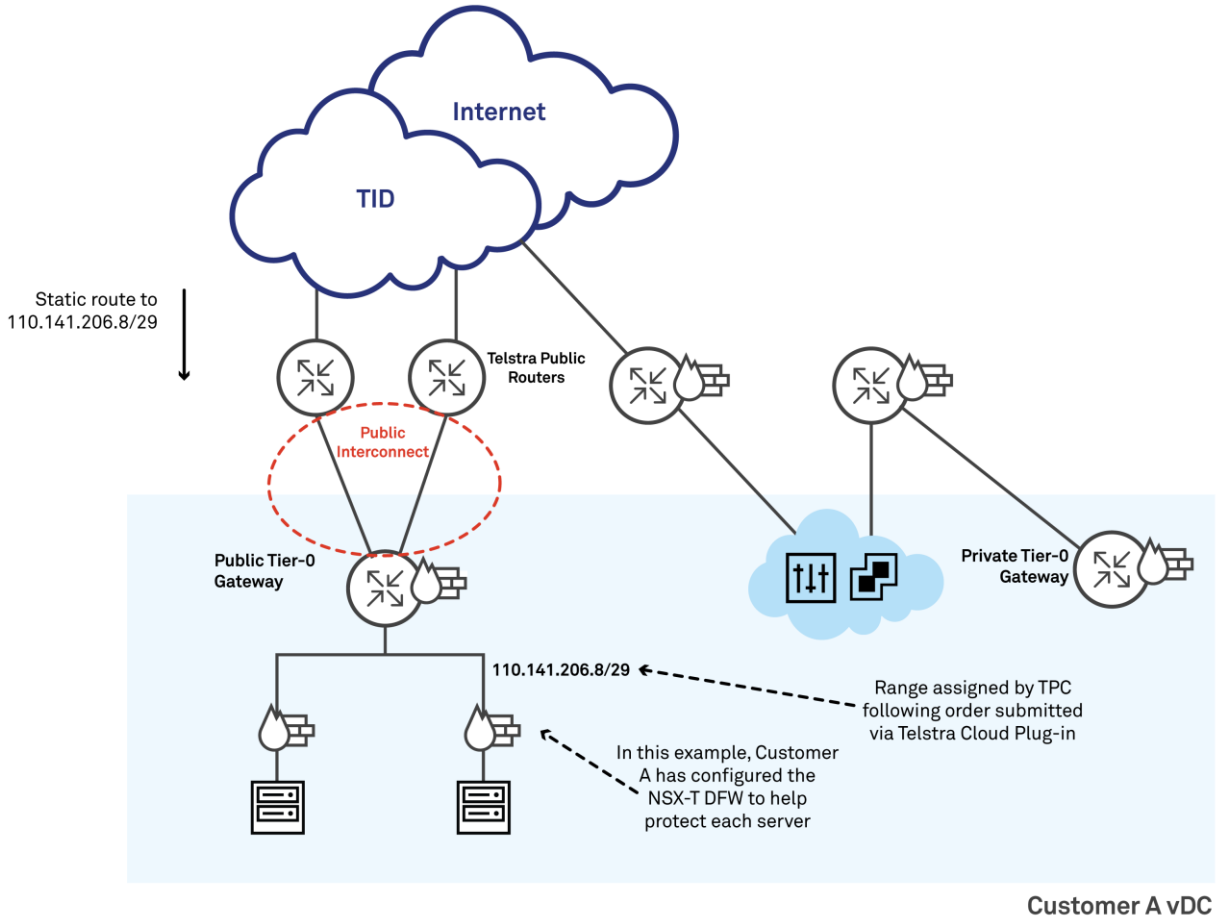

Private

Most basic private topologies will exclusively use addresses from the RFC 1918 private ranges but this is not their defining characteristic. Rather, a basic private topology is one that exclusively communicates with external parties over the Private Interconnect.

You are responsible for providing all addresses you employ in a basic private topology. You must ensure they are compatible with your wider private network, and you decide which ranges are advertised to Cloud Gateway or your Next IP VPN.

If you try to use a public range you did not obtain through the Telstra vCenter Plug-in, you must remember that while you can communicate with addresses from that range over the *Private* Interconnect, it will not subsequently work with a *Public* Interconnect. All addresses used with the Public Interconnect must be obtained from TPC, which requires you to submit an order in the Telstra vCenter Plug-in.

In Figure 10 we show a basic private topology for Customer B. This customer chose 192.168.1.0/24 for the private servers in the vDC because it is compatible with their wider internal network. Customer B will configure the Private Tier-0 Gateway to advertise 192.168.1.0/24 to Next IP via BGP.

Figure 10: Basic Private Topology

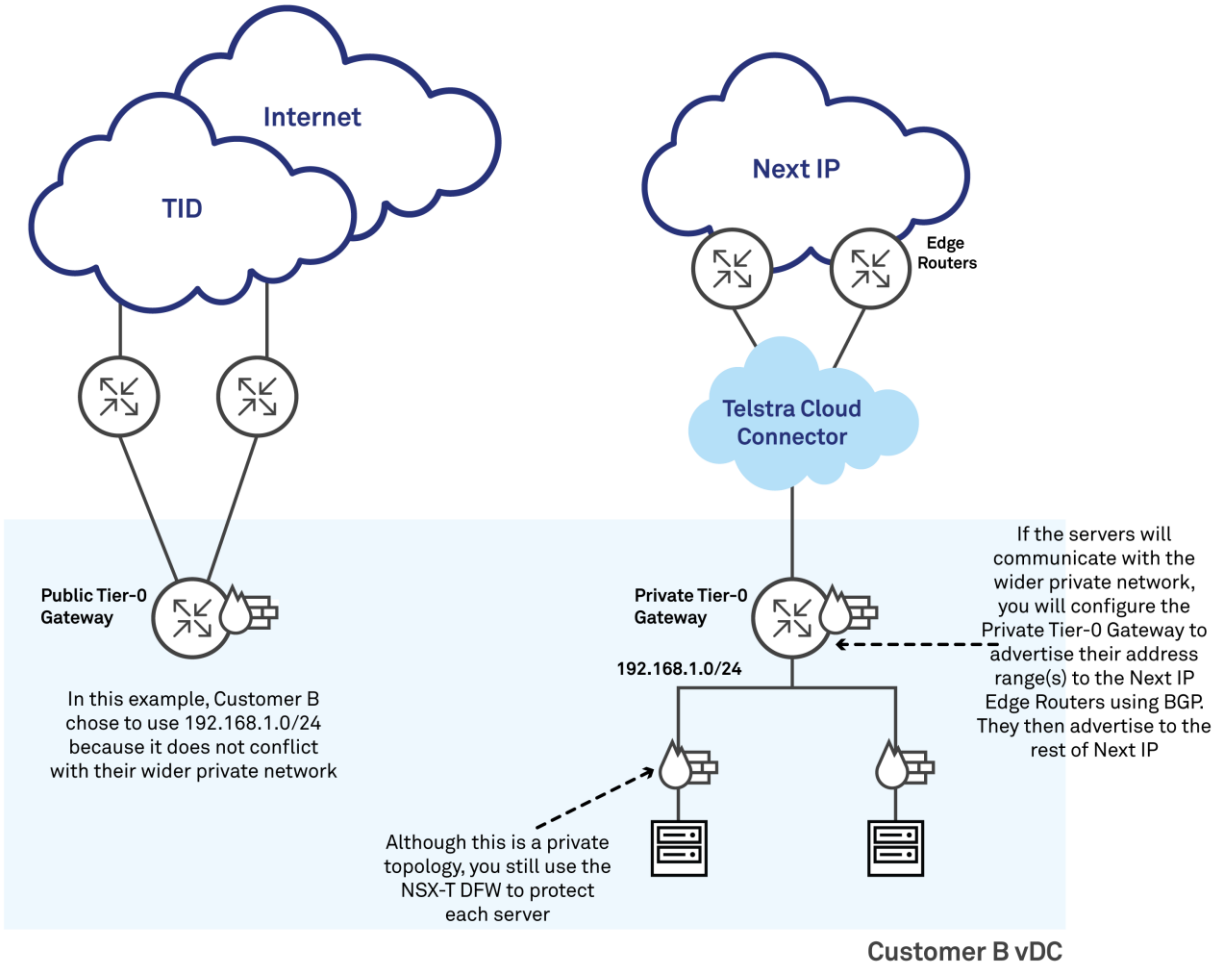

Complex Topologies

You might need to use an advanced design in your vDC, involving public and private access to various parts of the topology but in a secure and controlled fashion. That is, some servers and other resources are visible to the Public Interconnect and some are visible to the Private Interconnect. Moreover, certain servers or resources might be visible to both, but via different interfaces and paths.

We show an example of a complex topology for 'Customer D' in Figure 11. Customer D has divided the publicly visible topology into multiple logical segments using tiers and zones. Customer D has:

- Defined logical switches to build segments to arrange VMs and other resources in common zones and tiers
- Extensively used the NSX-T DFW on VM interfaces to apply stateful firewall rules locally on every server
- Added a Tier-1 Gateway to load balance traffic to certain servers
- Implemented further VMs running third-party NGFW software for advanced security to protect the vDC as traffic enters from the Public Interconnect, and to separate and control communication between that area of the vDC and the rest of the customer's private networks.

Regardless of the design you choose to build into a complex topology, always remember that:

1. Telstra must assign all public IP addresses that can be reached in your vDC through the Public Interconnect. If you use addresses in your public network topology that we did not assign to you, the Telstra Public Routers will not be able to route traffic to them
2. It is up to you to construct a topology that can successfully and securely communicate between:
    - Your Public Tier-0 Gateway and the VMs that are visible to the Internet
    - Your Private Tier-0 Gateway and the VMs the are visible to your private networks
    - VMs that need to be able to reach each other
3. You remain responsible for the logical security of your vDC. This includes the configuration of appropriate rules within the NSX-T DFW as well as additional security necessary to protect your vDC and any downstream systems and networks.

Figure 11: Complex Topology

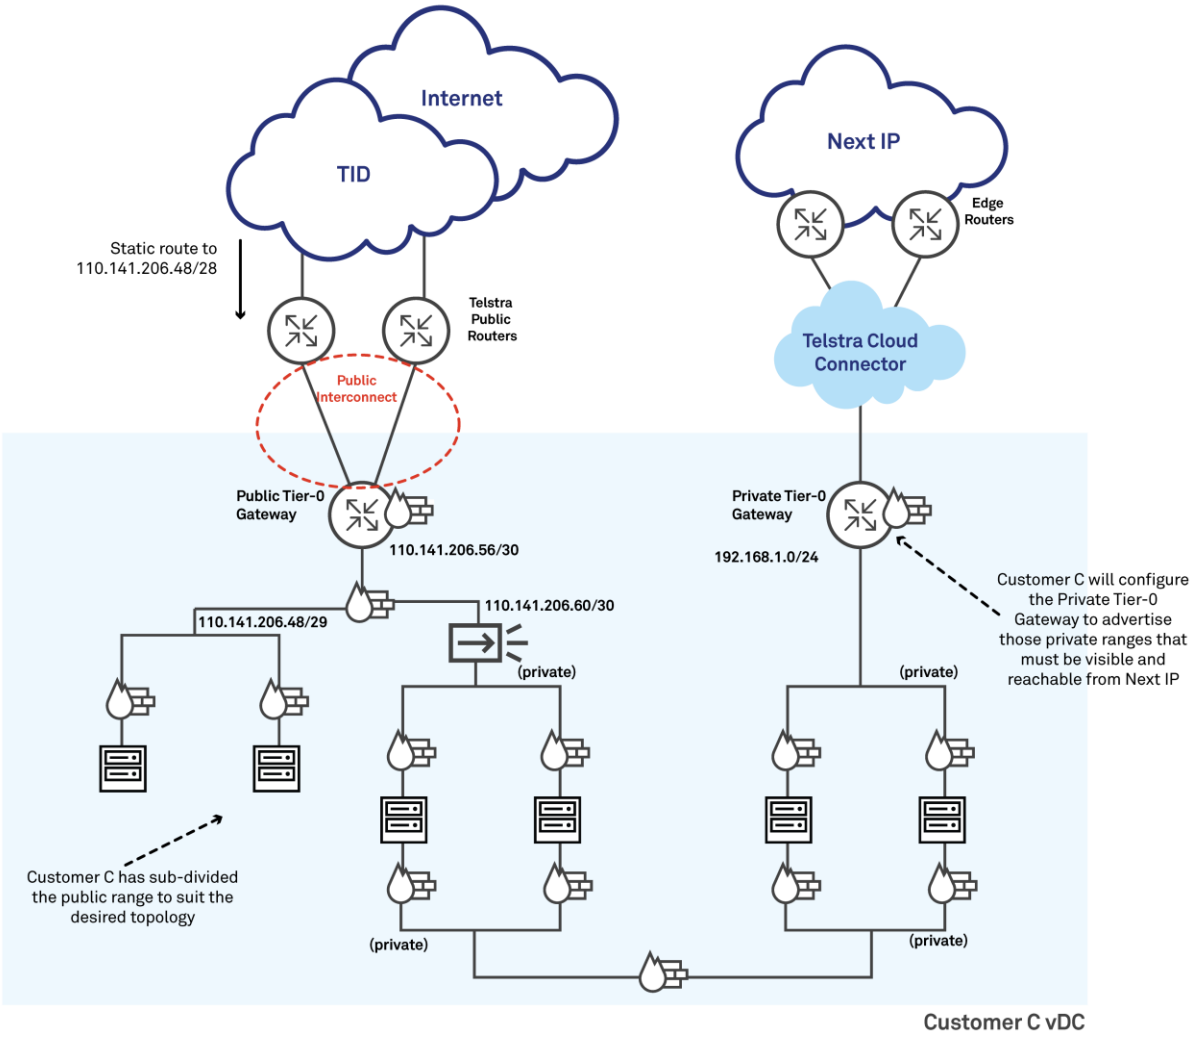

Chapter 3: NSX-T Manager Tasks

Task #NS01: Add a Segment

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☐ vSphere Admin ☐ vSphere Read-Only ☑ NSX-T Admin	☐ Telstra Cloud Sight ☐ vSphere (Native) ☐ Telstra Cloud plug-in ☑ NSX-T

Purpose

To create a logical Layer-2 broadcast domain to allow communication between:

- VMs
- Logical Routers (LRs) and/or
- Gateway interfaces.

Generally, you will create and configure your new Segment before adding or configuring the logical devices that will connect to it.

Overview

As you construct your vDC topology, you may need additional ***Segments***. You can use them in a flat, multi-tier or multi-zone topology, or as a stub connected to a Tier-0 or Tier-1 Gateway, providing connectivity to external public or private IP address ranges.

We show an example of a complex multi-tier topology in Figure 12. This example uses several Segments to provide connectivity between resources in different levels of the topology.

Figure 12: Using Segments to Connect Resources

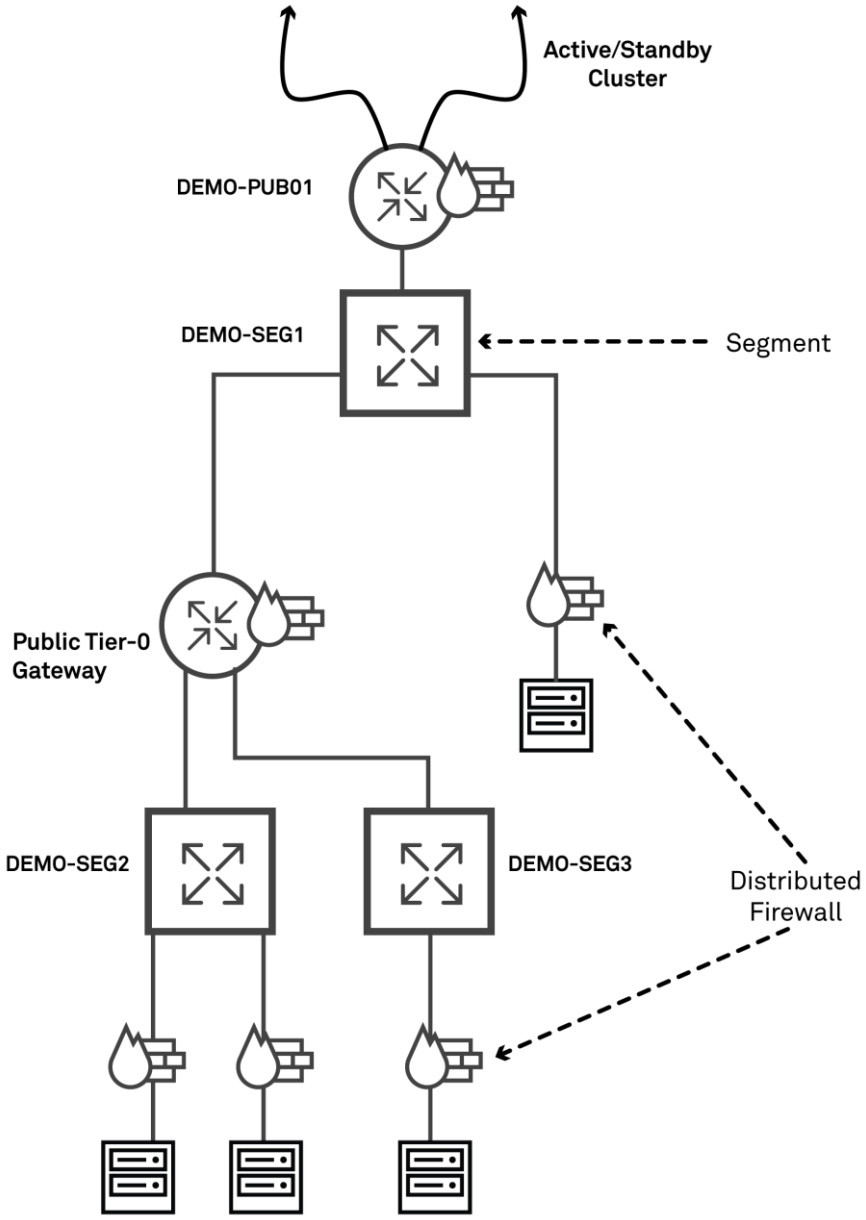

In NSX-T, Segments come in two types:

1. VLAN-backed: employs a traditional VLAN running between hosts over the physical infrastructure, or
2. Overlay-backed: runs through GENEVE tunnels between hosts, carried by the overlay network.

You can only create Overlay-backed Segments within your vDC.

Procedure

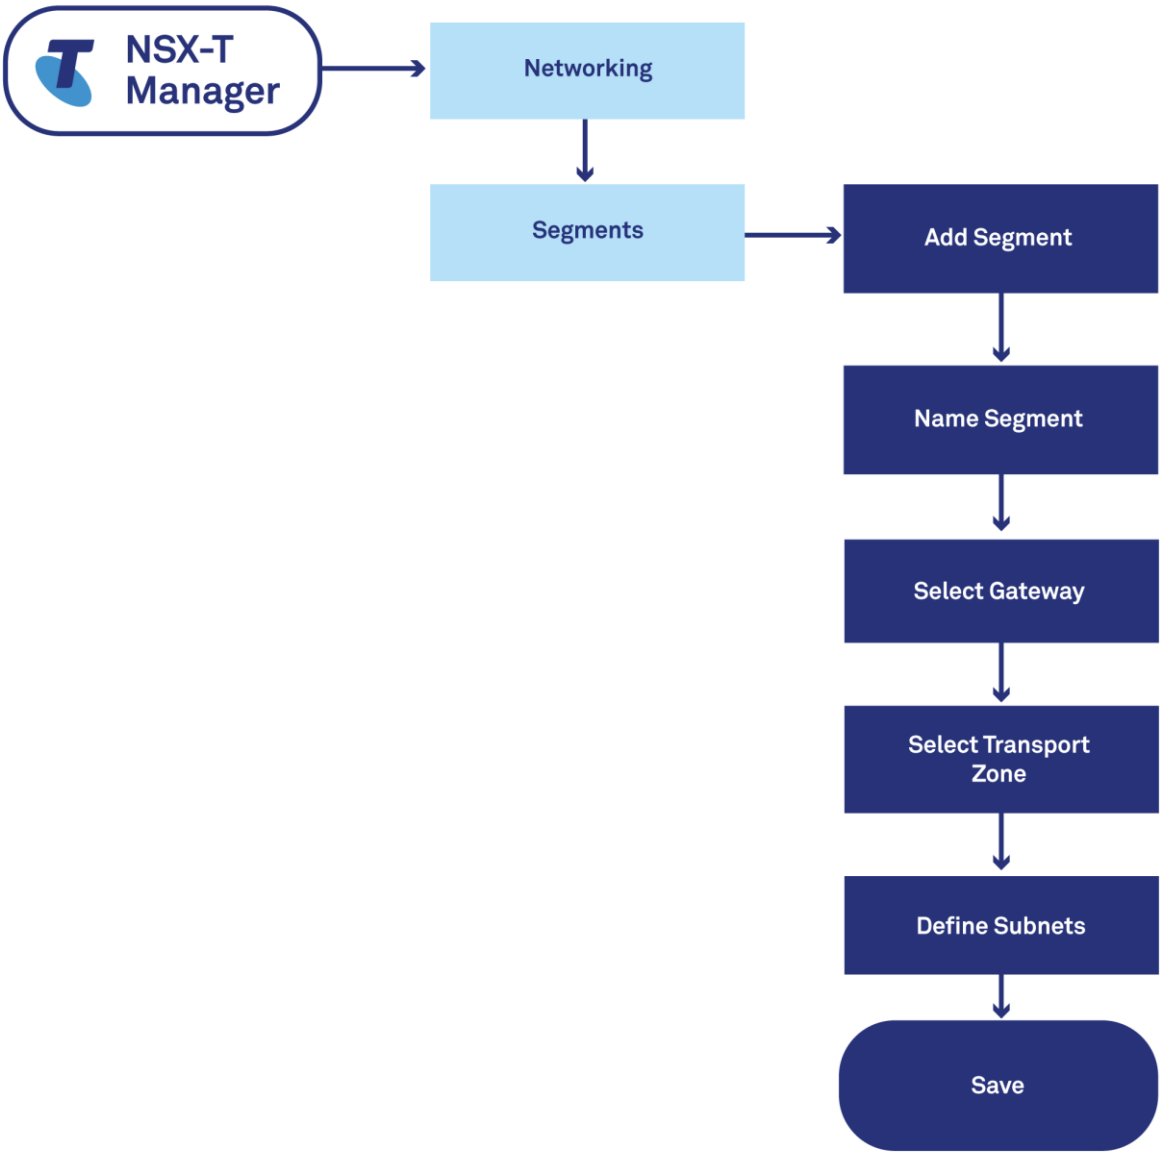
NSX-T Manager
Networking
Segments
Add Segment
Name Segment
Select Gateway
Select Transport Zone
Define Subnets
Save

Configuration Tips

You can only administer your Segments through NSX-T Manager

When you create a Segment or edit its configuration, please observe these recommendations:

1	You may name your new Segment according to your own standards. However, you should not re-name or change the provisioned configuration of any component Telstra has created in your TPC vDC on your behalf. When next invoked, Telstra's automated provisioning systems will reset those entities to standard templates and discard any of your alterations that conflict with our settings.
2	Except for fields that require mandatory input, we recommend that you accept the default configuration settings because they will be suitable in most cases. You can usually amend them later if required
3	You must select ***nsx-overlay-transportzone | Overlay*** as the Segment's Transport Zone. We created this Transport Zone when we provisioned your vDC
4	After you create your Segment, you will need to provision and connect the appropriate network components (using NSX-T) or virtual machines using the vSphere vCenter portal.

Information Resources

You can learn more about Segments by referring to these VMware documents:

- Segment Concepts
- How to Add a Segment (VMware Docs)
- NSX-T Reference Design Guide 3-0 | VMware
- NSX-T Data Center Quick Start Guide

Task #NS02: Modify an Edge Node

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☐ vSphere Admin ☐ vSphere Read-Only ☑ NSX-T Admin	☐ Telstra Cloud Sight ☐ vSphere (Native) ☐ Telstra Cloud plug-in ☑ NSX-T

Purpose

To configure an Edge node to deliver one or more specific functions for your vDC, such as routing, security, load balancing, NAT, VPN, DHCP and/or DNS.

Overview

An NSX-T Edge node can undertake several valuable functions:

- Static and dynamic routing
- Load balancing
- NAT
- A VPN concentrator
- DHCP server or relay
- DNS server.

Edge nodes support HA active/standby configurations running on dual VM-based appliances.

Standard Deployment in vDC

Edge nodes are integral components to TPC. When we provision your vDC, we create both your Private Tier-0 Gateway and Public Tier-0 Gateway as Edge nodes in a high availability cluster. All four are visible in the NSX-T Manager (they are also visible in vSphere as Virtual Machines).

You cannot create additional Edge nodes in your vDC. While you can adjust some of the features of Edge nodes we have built for you, you should remember that our standard settings comply with provisioning templates that we will reapply to your vDC from time to time. Consequently, if you change an Edge node's name or any of our standard settings, you will find them reset at some future point. Moreover, if you make changes detrimental to the operation of your vDC, you may incur additional fees if you need to raise service requests for Telstra to resolve issues those changes have caused.

Procedure

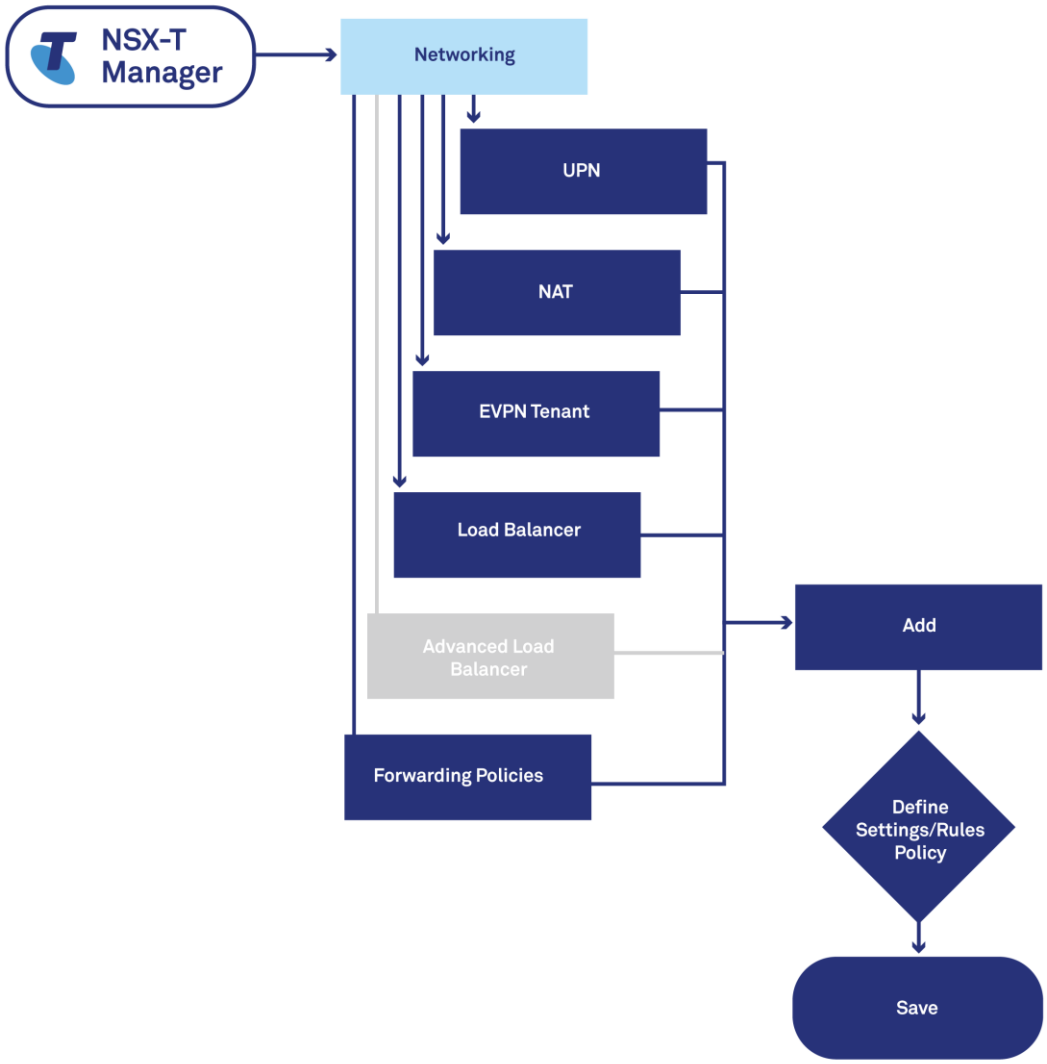

Configuration Tips

When you edit an existing Tier-0 Gateway configuration, please observe these recommendations:

1. Where you choose to activate a specific function, you will find the default values are quite suitable most of the time. Possible exceptions are:
   - Where you wish to configure specific routes and nodes to achieve complex vDC operations.
   - When the function of the Edge Node is not in line with the standard functions of that device
2. You can enable a default firewall policy
3. Telstra employs a ***Quad Large*** VM for each Public and Private Tier-0 Gateway when we build them during vDC provisioning
4. You are free to incorporate HA active/standby VM appliances underneath your Tier-0 Gateway
5. When you nominate a datastore to hold the VM system files, ensure you choose one of the correct type (ie. Standard or Performance tier) and with sufficient space to hold your files
6. If you plan to connect the Tier-0 Gateway to a new Virtual Machine, create the Segment and network configurations before updating the Edge node, Tier-0 Gateway or the VMs in the vSphere portal.

Task #NS03: Add/Modify a Tier-1 Gateway (Logical Router)

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☐ vSphere Admin ☐ vSphere Read-Only ☑ NSX-T Admin	☐ Telstra Cloud Sight ☐ vSphere (Native) ☐ Telstra Cloud plug-in ☑ NSX-T

Purpose

To build a new Tier-1 Gateway in your vDC to:

1. Uplink to a Tier-0 Gateway for north-south routing into downstream Segments
2. Perform east-west traffic routing among VMs spread over two or more Segments, and/or
3. Provide stateful services for traffic, such as firewall, load balancing and NAT.

Overview

You cannot create a Tier-0 Gateway for your vDC, although you can adjust the functionality of, and some settings for, those Tier-0 Gateways we initially built for you. If you wish to create additional Logical Routers in TPC, they will always be Tier-1 Gateways.

Unlike the case for an Edge node, you do not choose the size of the underlying VM appliance hosting a Tier-1 Gateway, although if you will associate it with a load balancer, you can elect to allocate more resources to it.

Procedure

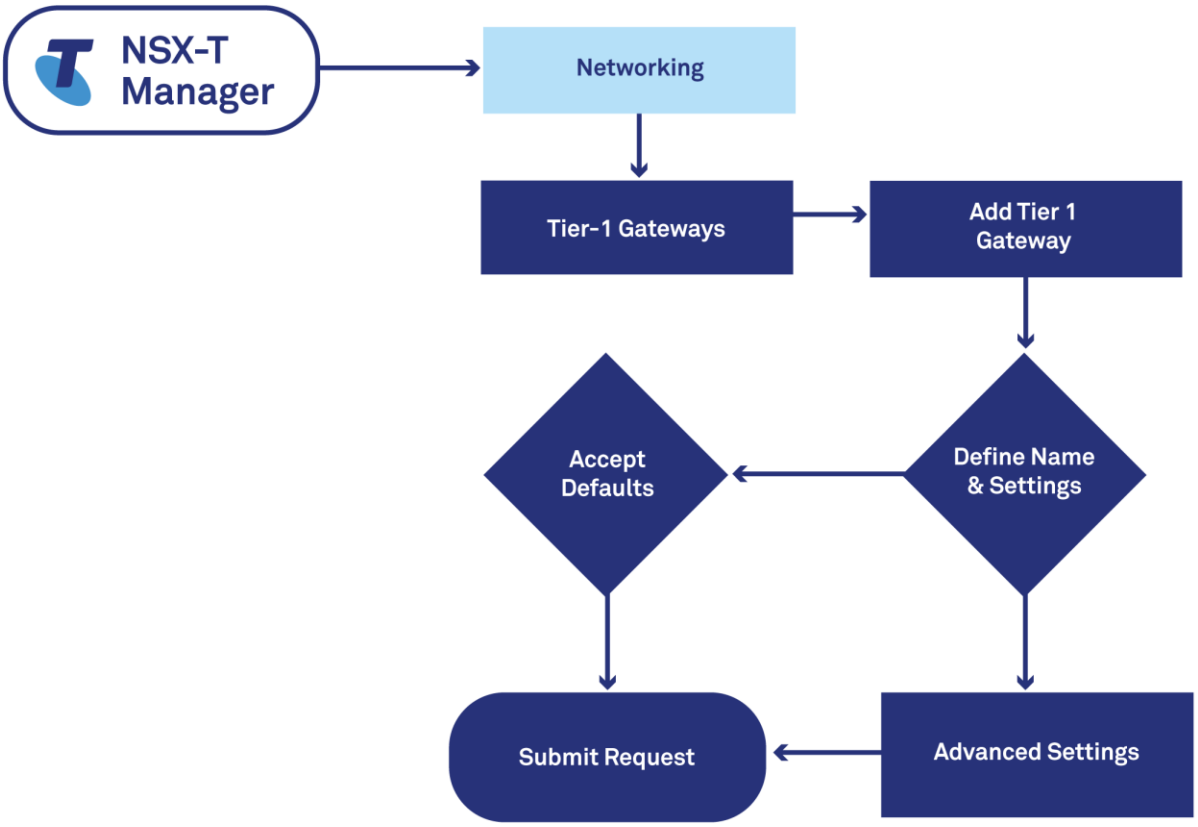

Configuration Tips

When you add or modify your Tier-1 Gateway, please observe these recommendations:

1	You may name your new Gateway according to your own standards
2	You will find the default values for configuration parameters are quite suitable in most cases
3	You can incorporate HA active/standby Tier-1 Gateways but remember that each new device will consume vDC resources to operate
4	If you plan to connect the Tier-1 Gateway to a new Segment, you will need to create that Segment and configure its settings appropriately before you create the Tier-1 Gateway

Information Resources

You can learn more about Tier-1 Gateways and how to use them in TPC vDCs by referring to these VMware links:

- Tier-0 Logical Router
- Tier-1 Logical Routers
- NSX-T Key Concepts

Task #NS04: Add/Modify NSX-T Gateway Firewall

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☐ vSphere Admin ☐ vSphere Read-Only ☑ NSX-T Admin	☐ Telstra Cloud Sight ☐ vSphere (Native) ☐ Telstra Cloud plug-in ☑ NSX-T

Purpose

To adjust the perimeter firewall policies to manage traffic flowing into and out of your vDC.

Overview

A Gateway Firewall is a perimeter firewall that resides on any Tier-0 Gateway or Tier-1 Gateway in your vDC.

When Telstra provisions your vDC, the initial configuration for your Public and Private Tier-0 Gateways will contain a generic Gateway Firewall policy. Because you have administrative access to your vDC, you will be able to add or modify the Gateway Firewall on both Tier-0 Gateways.

While you can make the changes you see fit, Telstra recommends you exercise caution. You should carefully evaluate your security policies because NSX-T Manager will not check your rules for consistency or efficacy, nor their effects on the security of your vDC's resources.

Moreover, Telstra uses standardised templates to configure devices we provision on your behalf, and we periodically re-apply those templates to your vDC. Consequently, you should avoid changes to names and settings that we configure for you.

If you wish, you are able to supplement your vDC with further security products and services that you purchase from us or acquire through other means.

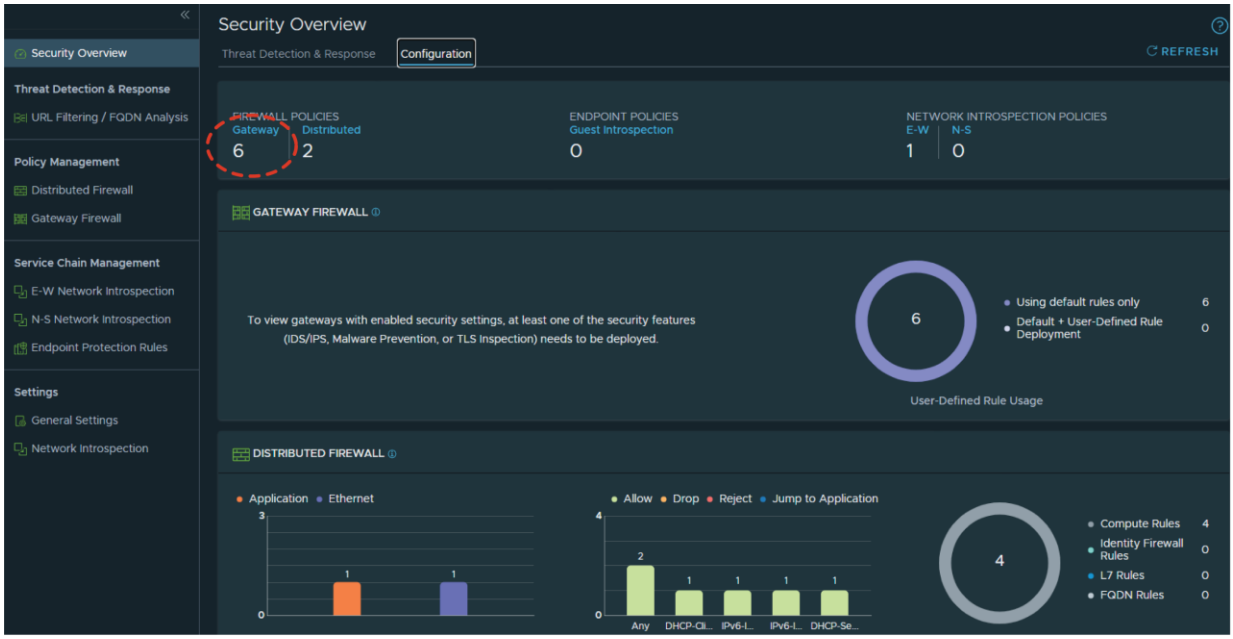
Security Overview
Security Overview
Threat Detection & Response
URL Filtering / FQDN Analysis
Policy Management
Distributed Firewall
Gateway Firewall
Service Chain Management
E-W Network Introspection
N-S Network Introspection
Endpoint Protection Rules
Settings
General Settings
Network Introspection
Threat Detection & Response
Configuration
REFRESH
FIREWALL POLICIES
Gateway
Distributed
6
2
ENDPOINT POLICIES
Guest Introspection
0
NETWORK INTROSPECTION POLICIES
E-W
N-S
1
0
GATEWAY FIREWALL
To view gateways with enabled security settings, at least one of the security features (IDS/IPS, Malware Prevention, or TLS Inspection) needs to be deployed.
6
Using default rules only 6
Default + User-Defined Rule Deployment 0
User-Defined Rule Usage
DISTRIBUTED FIREWALL
Application
Ethernet
3
0
1
1
Allow
Drop
Reject
Jump to Application
4
0
2
1
1
1
1
Any
DHCP-Cli...
IPv6-I...
IPv6-I...
DHCP-Se...
4
Compute Rules 4
Identity Firewall Rules 0
L7 Rules 0
FQDN Rules 0

Procedure

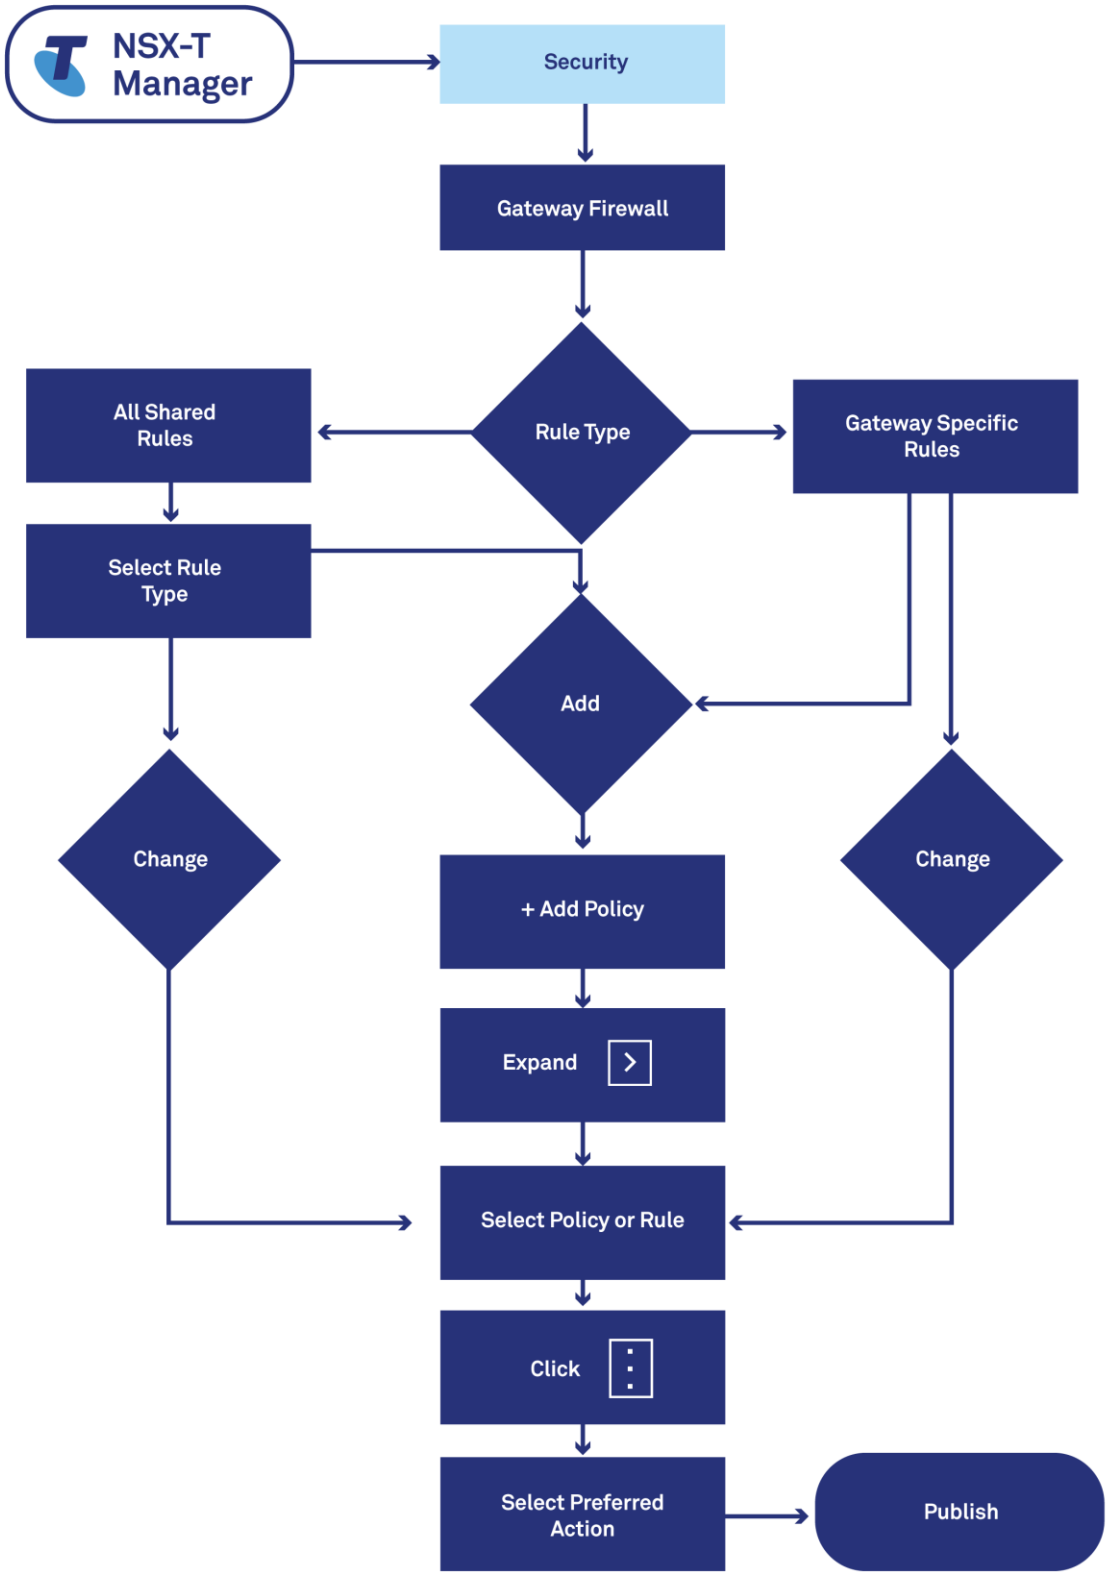
NSX-T Manager
Security
Gateway Firewall
Rule Type
All Shared Rules
Gateway Specific Rules
Select Rule Type
Add
Change
Change
+ Add Policy
Expand
Select Policy or Rule
Click
Select Preferred Action
Publish

Configuration Tips

When we hand it over to you, your vDC will contain a Public Tier-0 Gateway and Private Tier-0 Gateway each using a generic configuration, consisting of:

- Any custom rules required for the vDC to connect to your private MPLS network and existing resources
- Any custom rules required for the vDC to connect to your public internet connection.
- Default (catch-all) rules we have added for your private topology and resources
- Default (catch-all) rules we have added for your public topology and resources
- Default rules required by NSX-T.

You have a significant degree of configuration flexibility over your Gateway Firewalls. But while powerful and adaptable, you must apply sufficient rigour and care when you define objects, groups and policy rules. Otherwise, you risk introducing policy conflicts that may admit undesirable traffic to your vDC or block legitimate communication.

Information Resources

You can learn more about the NSX-T Gateway Firewalls by referring to these VMware documents:

- Logical Firewall
- Edge Firewall
- Working with Firewall Rule Sections
- Working with Firewall Rules
- The NSX-T Gateway Firewall Secures Physical Servers

Task #NS05: Add/Modify NSX-T Distributed Firewall

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☐ vSphere Admin ☐ vSphere Read-Only ☑ NSX-T Admin	☐ Telstra Cloud Sight ☐ vSphere (Native) ☐ Telstra Cloud plug-in ☑ NSX-T

Purpose

To apply policies and rules to traffic flowing through a segment, or into and out of your vDC.

Overview

You can define an NSX-T Distributed Firewall (DFW) at the vNIC level of your hosts and resources, enabling you to granularly monitor and manage traffic as it flows through your vDC.

Procedure

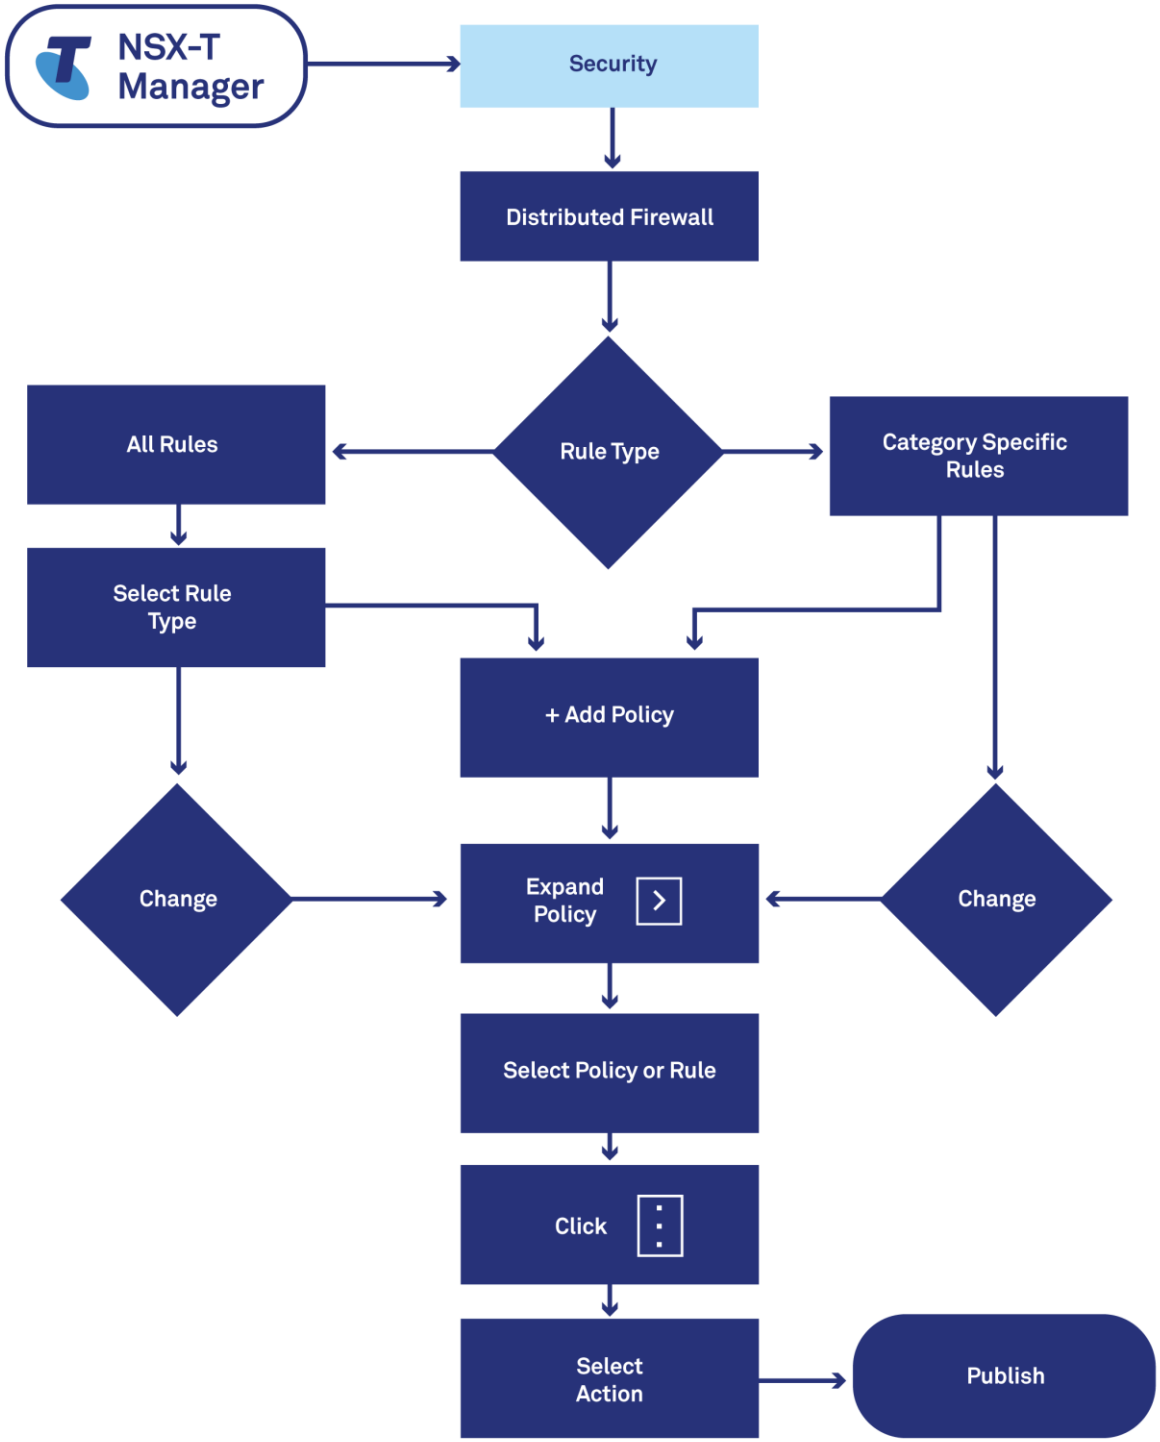
NSX-T Manager
Security
Distributed Firewall
Rule Type
All Rules
Category Specific Rules
Select Rule Type
+ Add Policy
Change
Expand Policy
Change
Select Policy or Rule
Click
Select Action
Publish

Configuration Tips

The DFW for your vDC will contain a small set of generic protocol rules when we hand it over to you after provisioning.

You can apply DFW security policies to manage traffic between one or more VMs within a Segment. You need to create and configure your Segment(s) prior to creating the DFW policy or rule. While the DFW requires the vNIC information to be available when you create the policy or rule, once applied they will operate independently of the location of the VM, meaning that if vMotion moves the VM later, the DFW polices will follow it, reducing outages and interruptions.

Information Resources

You can learn more about the NSX-T Distributed Firewall refer to these VMware documents:

- Distributed Firewall
- Edge Firewall
- Working with Firewall Rule Sections
- Working with Firewall Rules
- Add a Distributed Firewall

Task #NS06: Add a Layer-2 VPN

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☐ vSphere Admin ☐ vSphere Read-Only ☑ NSX-T Admin	☐ Telstra Cloud Sight ☐ vSphere (Native) ☐ Telstra Cloud plug-in ☑ NSX-T

Purpose

To create a site-to-site IPsec tunnel into your vDC, through which you can extend Layer 2 networks (VNIs or VLANs) across multiple sites on the same broadcast domain.

Overview

You can use Layer-2 VPN services to cross any public or private network, or a mix of them depending on the location of your facilities and your routing arrangement. This could include:

- The Internet
- Your Next IP VPN
- Separated networks within your vDC, and/or
- Another vDC.[3]

The L2 VPN runs within an IPsec tunnel. You must create and configure the IPsec service before you attempt to deploy the L2 VPN service. The IPsec and L2 VPN settings will both need to use the same Tier-0 Gateway or Tier-1 Gateway to ensure the connection and the direction of traffic flow is identical.

[3] Refer to the Telstra Private Cloud: Administration Guide (available here) for further advice regarding routing between vDCs owned by the same organisation

Procedure

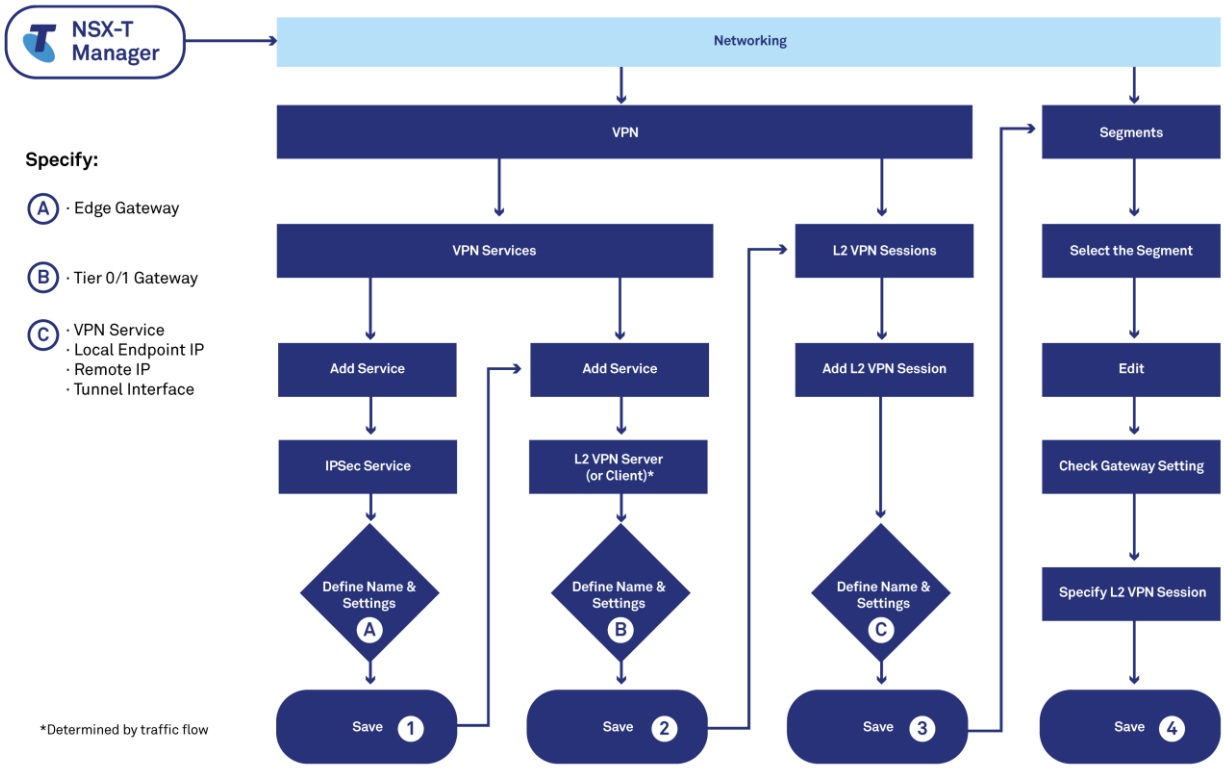

Configuration Tips

The L2 VPN feature is available only for NSX-T Data Center. It does not offer any third-party interoperability.

You can configure a L2 VPN on either a Tier-0 Gateway or Tier-1 Gateway across your vDC. However, each Tier-0 Gateway or Tier-1 Gateway can host only one L2 VPN service.

You can extend L2 networks from VLAN to VNI, VLAN to VLAN, and VNI to VNI.

If you plan to use L2 VPN through your Public Tier-0 Gateway, remember that you need to configure your tunnel endpoint with a TPC-compatible public IP address. You cannot supply your own addresses for use in your vDC, but must acquire them from us. Refer to Task #PA01 in the Telstra Private Cloud: Administration Guide (available here) if you need to request an additional public IP address range from Telstra.

A L2 VPN needs specific configuration settings in the logical devices that support it, such as Edge nodes and Logical Switches. Misconfigurations may cause looping and duplicate packets. Refer to VMware publications for L2 VPN best practices.

Information Resources

You can learn more about the Layer 2 (L2) VPNs and their configuration by referring to these VMware documents:

- L2 VPN Overview
- L2 VPN Best Practices
- L2 VPN Over SSL
- L2 VPN Over IPSec
- Understanding Layer 2 VPN

Task #NS07: Add a Load Balancer

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☐ vSphere Admin ☐ vSphere Read-Only ☑ NSX-T Admin	☐ Telstra Cloud Sight ☐ vSphere (Native) ☐ Telstra Cloud plug-in ☑ NSX-T

Purpose

To configure load balancing using the existing Public or Private Tier-0 Gateways.

Overview

Your Public Tier-0 Gateway and Private Tier-0 Gateway each support NSX-T load balancing configurations for both layer-4 and layer-7 parameters. The contrasting approaches offer these advantages:

- Layer-4 load balancing is implemented using TCP and UDP. It is fast because it does not buffer the whole request or stub the connection. Instead, it processes each packet header and sends it directly to the selected server
- Layer-7 load balancing is socket-based, so it creates back-to-back sessions. It receives and processes the whole request, allowing advanced traffic manipulation and DDOS mitigation. This is the default mode for TCP, HTTP and HTTPS virtual servers.

Procedure

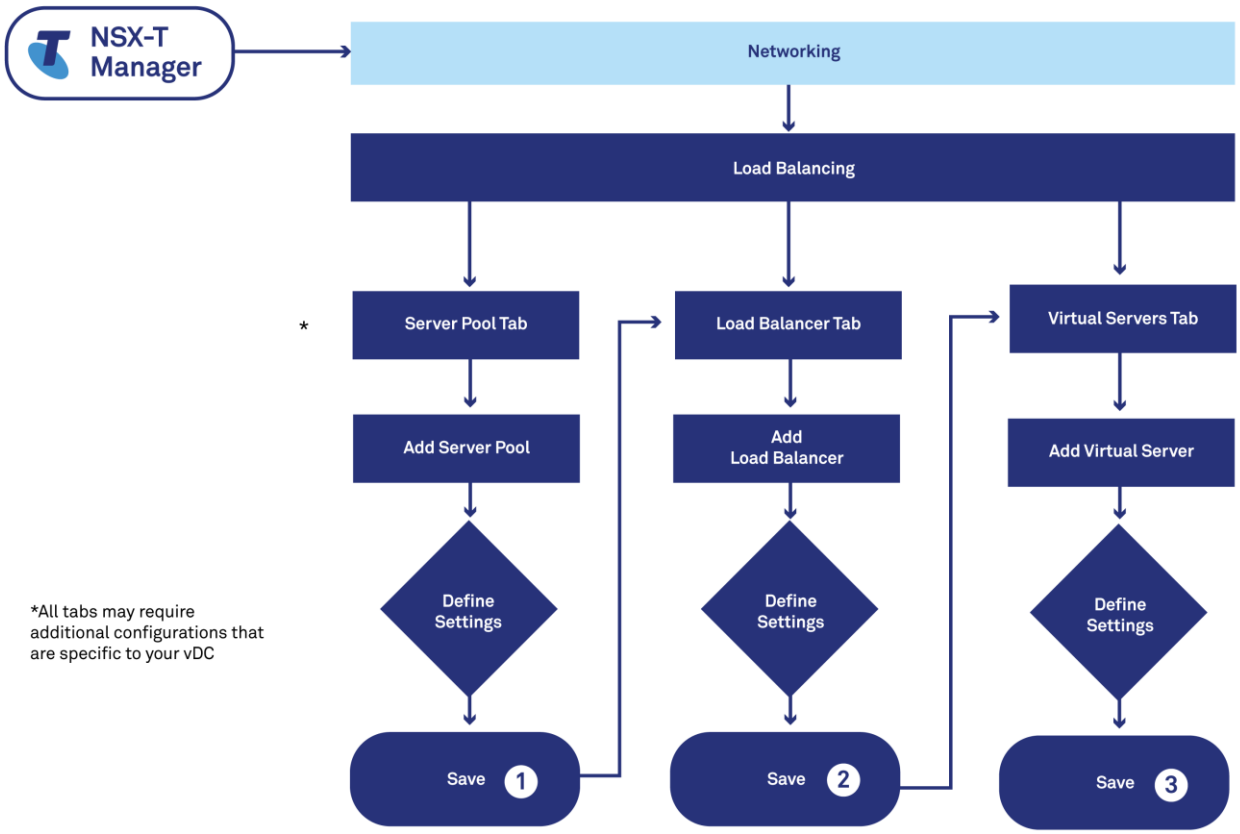

Configuration Tips

Within the new vDC you can configure load balancing by modifying your existing Public Tier-0 Gateway and/or Private Tier-0 Gateway. You can also add one when you deploy a Tier-1 Gateway.

In previous versions, our provisioning policy recommended deploying each Public or Private load balancer as a separate device, but with TPC the NSX-T Manager provides access to modify the configuration of the existing Dedicated Public or Dedicated Private Tier-0 Gateway.

Because you cannot create an additional Tier-0 Gateway within your vDC and Telstra configures the default Tier-0 Gateways in a HA Cluster, ensure you consider the impact of the processing load on the underlying VM and host if you activate this feature. For example, you may need to alter the VM's specifications or adjust other settings on the Tier-0 Gateway.

Note: An option for an 'Advanced Load Balancer' is visible in NSX-T Manager. This feature is not yet available in our initial platform release.

Information Resources

You can learn more about NSX-T load balancing by referring to these VMware Docs:

- Logical Load Balancer
- Setting Up Load Balancing
- Managing Service Monitors
- Managing Server Pools
- Managing Virtual Servers
- Managing Application Rules
- Scenarios for NSX-T Load Balancer Configuration

Task #NS08: Add a NAT Rule

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☐ vSphere Admin ☐ vSphere Read-Only ☑ NSX-T Admin	☐ Telstra Cloud Sight ☐ vSphere (Native) ☐ Telstra Cloud plug-in ☑ NSX-T

Purpose

To configure Network Address Translation (NAT) rules for your vDC.

Overview

Network Address Translation (NAT) rules are best used to map a specific IP address to another as a dedicated destination within your vdc.

You can apply NAT to your Tier-0 Gateways or Tier-1 Gateways. NSX-T supports Source NAT (SNAT), Destination NAT (DNAT) or Reflexive NAT (RNAT). SNAT and DNAT are stateful NAT techniques, while Reflexive NAT is stateless. Consequently, SNAT and DNAT can work with gateways running in Active/Standby HA mode, but only RNAT supports Active/Active HA mode.

Typically, you will use the various options for different purposes:

- Source NAT: to change the source address of outbound packets so that they appear to be originating from a specific IP address. Supports one-to-one and many-to-one modes
- Destination NAT: to overwrite the destination address of inbound packets with an address for a different target
- Reflexive NAT: to apply address translation to both inbound and outbound traffic as they pass through a gateway. It modifies the destination address of inbound packets and the source address of outbound packets.

Procedure

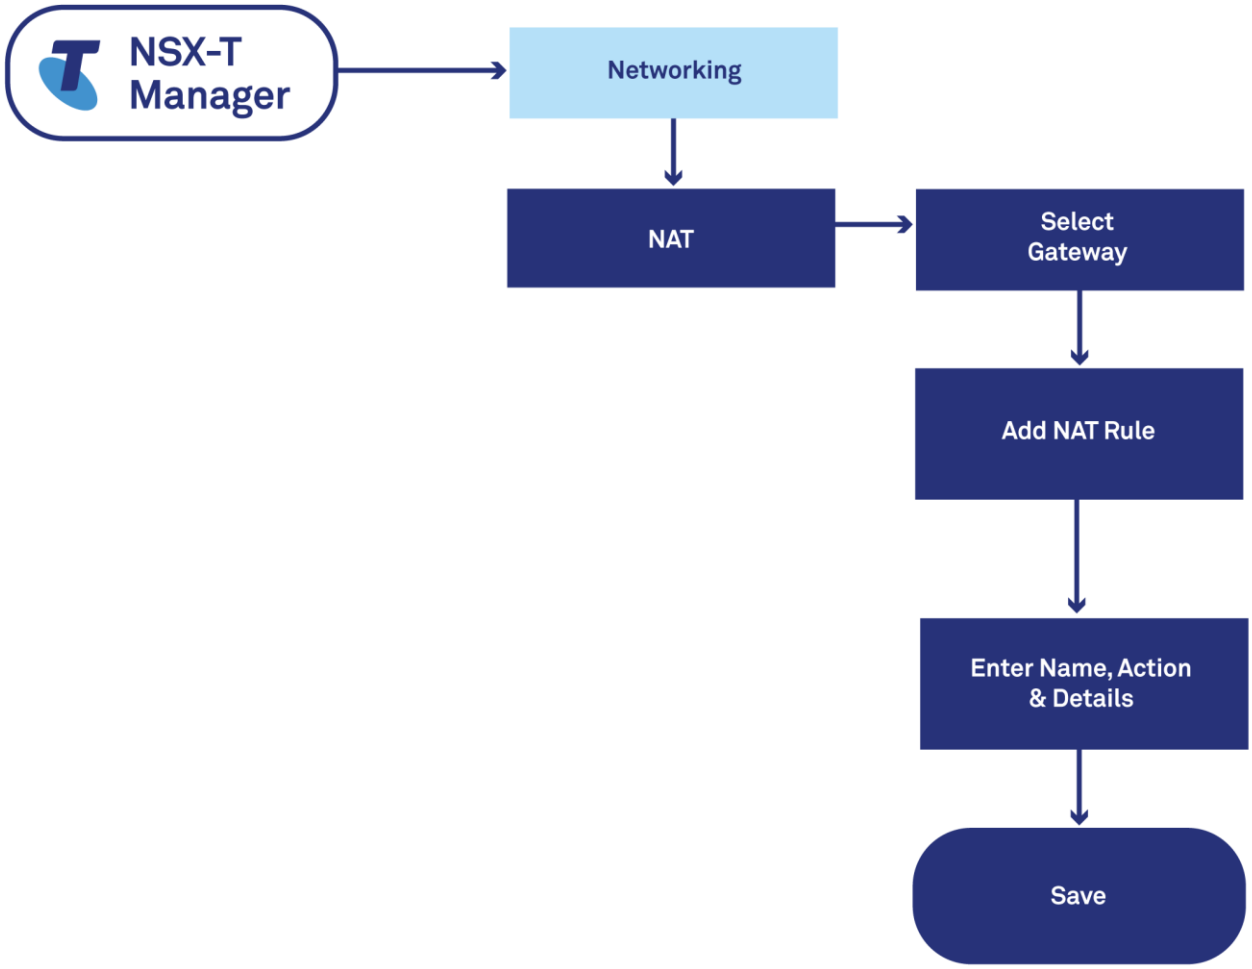

Configuration Tips

DNAT is not supported on a Tier-1 Gateway that also has IPsec VPN policies configured.

Determining the appropriate NAT Rules across the vDC will assist in defining both the Gateway Firewall rules and routing expected within the network.

Information Resources

You can learn more about NSX-T NAT rules by referring to:

- Configure NAT on a Gateway
- Configuring NAT services in NSX-T
- NSX-T: Configure Network Address Translation (NAT)

Chapter 4: vCenter Tasks

Task #VM01: Create a Virtual Machine

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☑ vSphere Admin ☐ vSphere Read-Only ☐ NSX-T Admin	☐ Telstra Cloud Sight ☑ vSphere (Native) ☐ Telstra Cloud plug-in ☐ NSX-T

Purpose

To create a VM for a new virtual server or virtual appliance.

Overview

You can create, modify and remove VMs from the hosts and clusters in your vDC directly using your vSphere client. Each VM can host one of several different types of resources, such as a virtual server or a virtual appliance hosting a Tier-1 Gateway or Load Balancer.

When you create a VM, you can define it with customised specifications or base it on another entity that is identical or similar to what you need:

- Use a custom definition: where no other VMs in your vDC have the requirements you need (eg. OS or hardware configuration)
- Use a pre-configured VM definition: where you can export an existing VM or virtual appliance, or use a vApp stored as an OVF file
- Use a VM template: where you create a master copy of a VM and then deploy multiple copies of it
- Clone an existing VM: you can do this many times, either directly or by first cloning the existing VM to a template.

vSphere includes various ‘wizards’ that guide you through each process. The wizard will ask you to specify a name for your new VM and a storage location.

Procedure

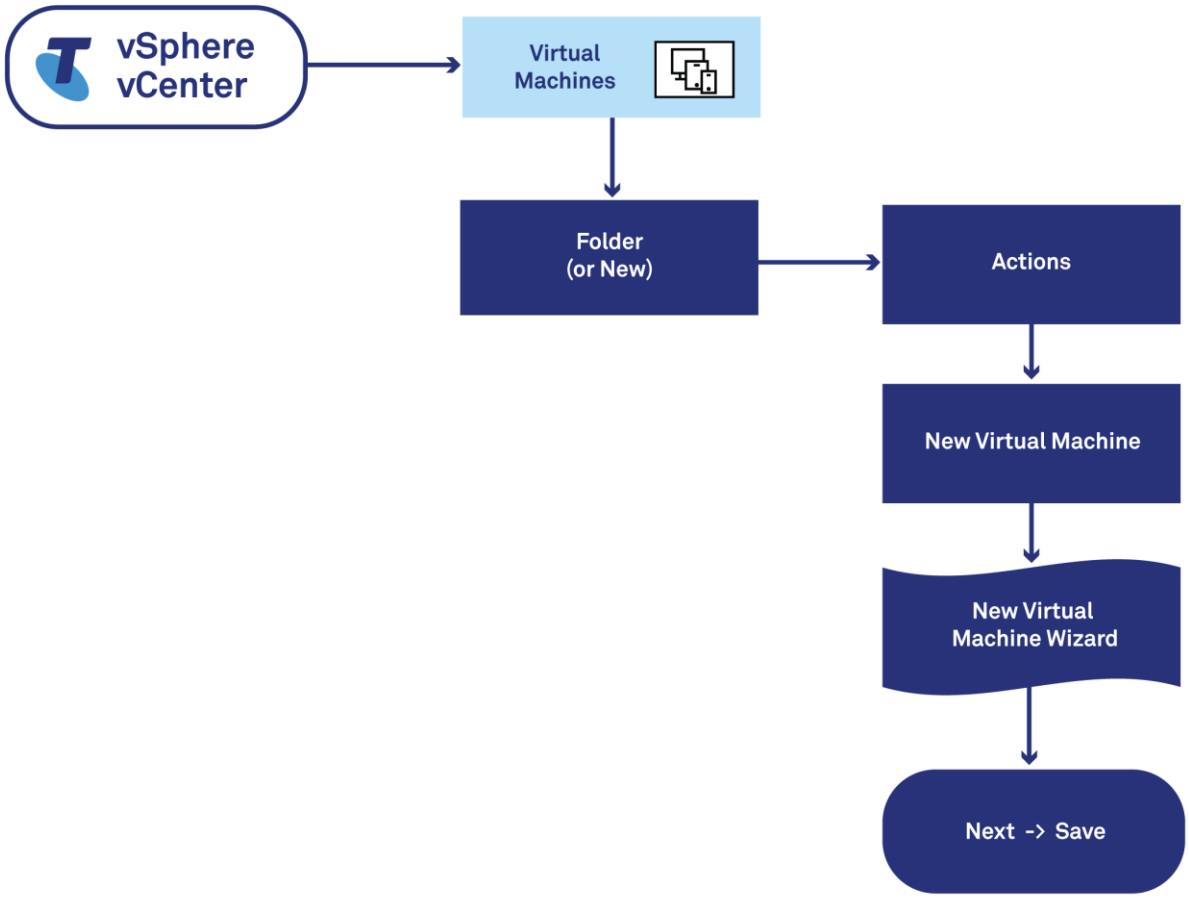

Configuration Tips

You set the specifications for your VMs at the time you create it. VMware provides documentation that can help you to determine the best size for your VM and how to use vMotion with it, as well as wizards you can invoke from the vSphere client.

You will need to prepare the folder you intend to hold your new VM in advance, typically using the "New Folder" wizard.

Information Resources

You can learn more about VMs and VMWare 'Best Practices' by referring to these VMware documents:

- Deploying Virtual Machines
- Configuring Virtual Machine Hardware
- Configuring Virtual Machine Options

Task #VM02: Create a VM DRS Group

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☑ vSphere Admin ☐ vSphere Read-Only ☐ NSX-T Admin	☐ Telstra Cloud Sight ☑ vSphere (Native) ☐ Telstra Cloud plug-in ☐ NSX-T

Purpose

To add one or more VMs to a VM Distributed Resource Scheduler (DRS) Group for later use in DRS VM-Host affinity rules.

Overview

You can create DRS VM-Host affinity rules to influence the distribution of VMs across certain hosts. These rules can cause a VM to prefer to be placed on a certain host, or to avoid it.

Procedure

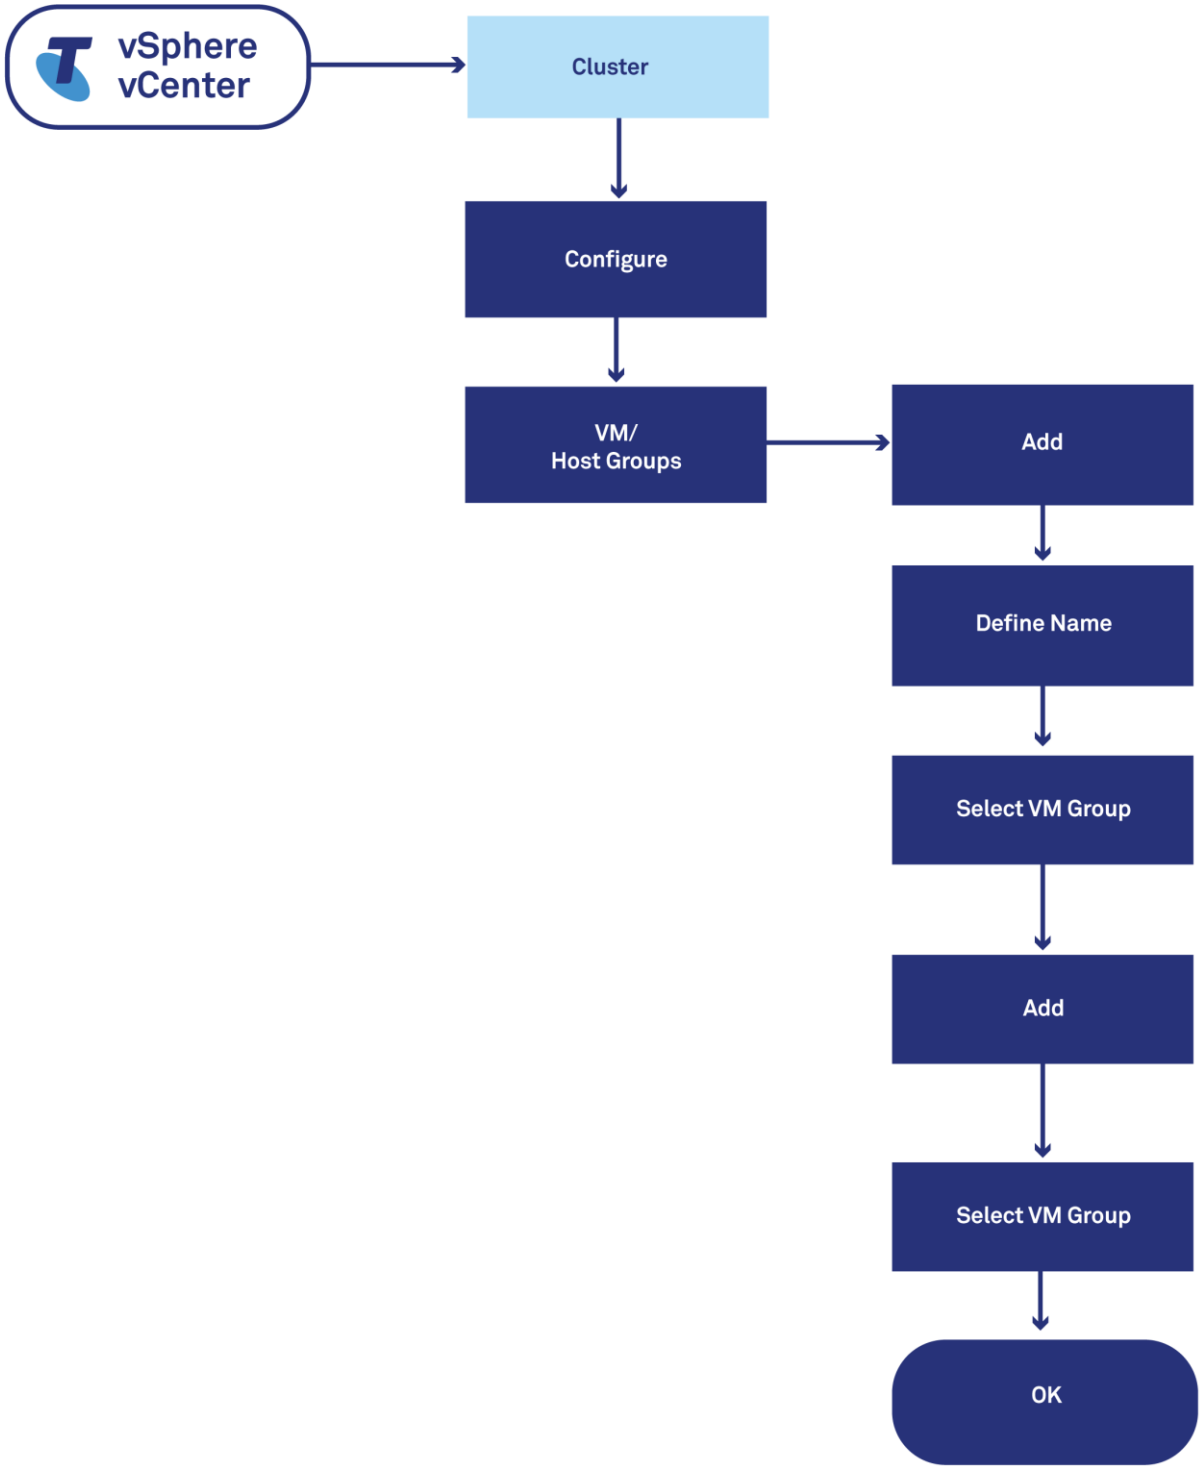
vSphere
vCenter
Cluster
Configure
VM/
Host Groups
Add
Define Name
Select VM Group
Add
Select VM Group
OK

Configuration Tips

When you create a VM DRS Group, please observe these points and recommendations:

- Each VM DRS Group will consist of one or more VMs and draw its VM membership from those configured in a single host cluster
- If you create VM DRS Groups with overlapping memberships, you must be mindful of the way in which VMware resolves conflicts with their corresponding affinity rules
- Any given VM can be a member of more than one VM DRS Group at a time, though the VMs must reside on the same cluster
- Removing the VM from the cluster will also cancel its membership of the DRS Group. DRS Group membership is not automatically reinstated, even if you should later move the VM back into that cluster.

Information Resources

You can learn more about DRS Groups and VM-Host affinity rules and how to use them in TPC by referring to these VMware documents:

- Using DRS Affinity Rules
- Create a Virtual Machine DRS Group
- VM-Host Affinity Rules
- Using vSphere DRS Groups and VM-Host Affinity Rules with WSFC Virtual Machines
- DRS Cluster Validity

Task #VM03: Create a Host DRS Group

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☑ vSphere Admin ☐ vSphere Read-Only ☐ NSX-T Admin	☐ Telstra Cloud Sight ☑ vSphere (Native) ☐ Telstra Cloud plug-in ☐ NSX-T

Purpose

To add one or more hosts to a Host DRS Group for later use in DRS VM-Host affinity rules.

Overview

You can create DRS VM-Host affinity rules to influence the distribution of VMs across certain hosts. These rules can cause a VM to prefer to be placed on a certain host, or to avoid it.

If you plan to create DRS VM-Host affinity rules that influence the distribution of VMs across certain hosts, you must first arrange those respective hosts into one or more Host DRS Groups. Each Host DRS Group will consist of one or more hosts and draw its host membership from those configured in a single host cluster.

If you only plan to create DRS VM-VM affinity rules (ie. between individual VMs rather than between VMs and hosts) then you do not need to complete this task.

Procedure

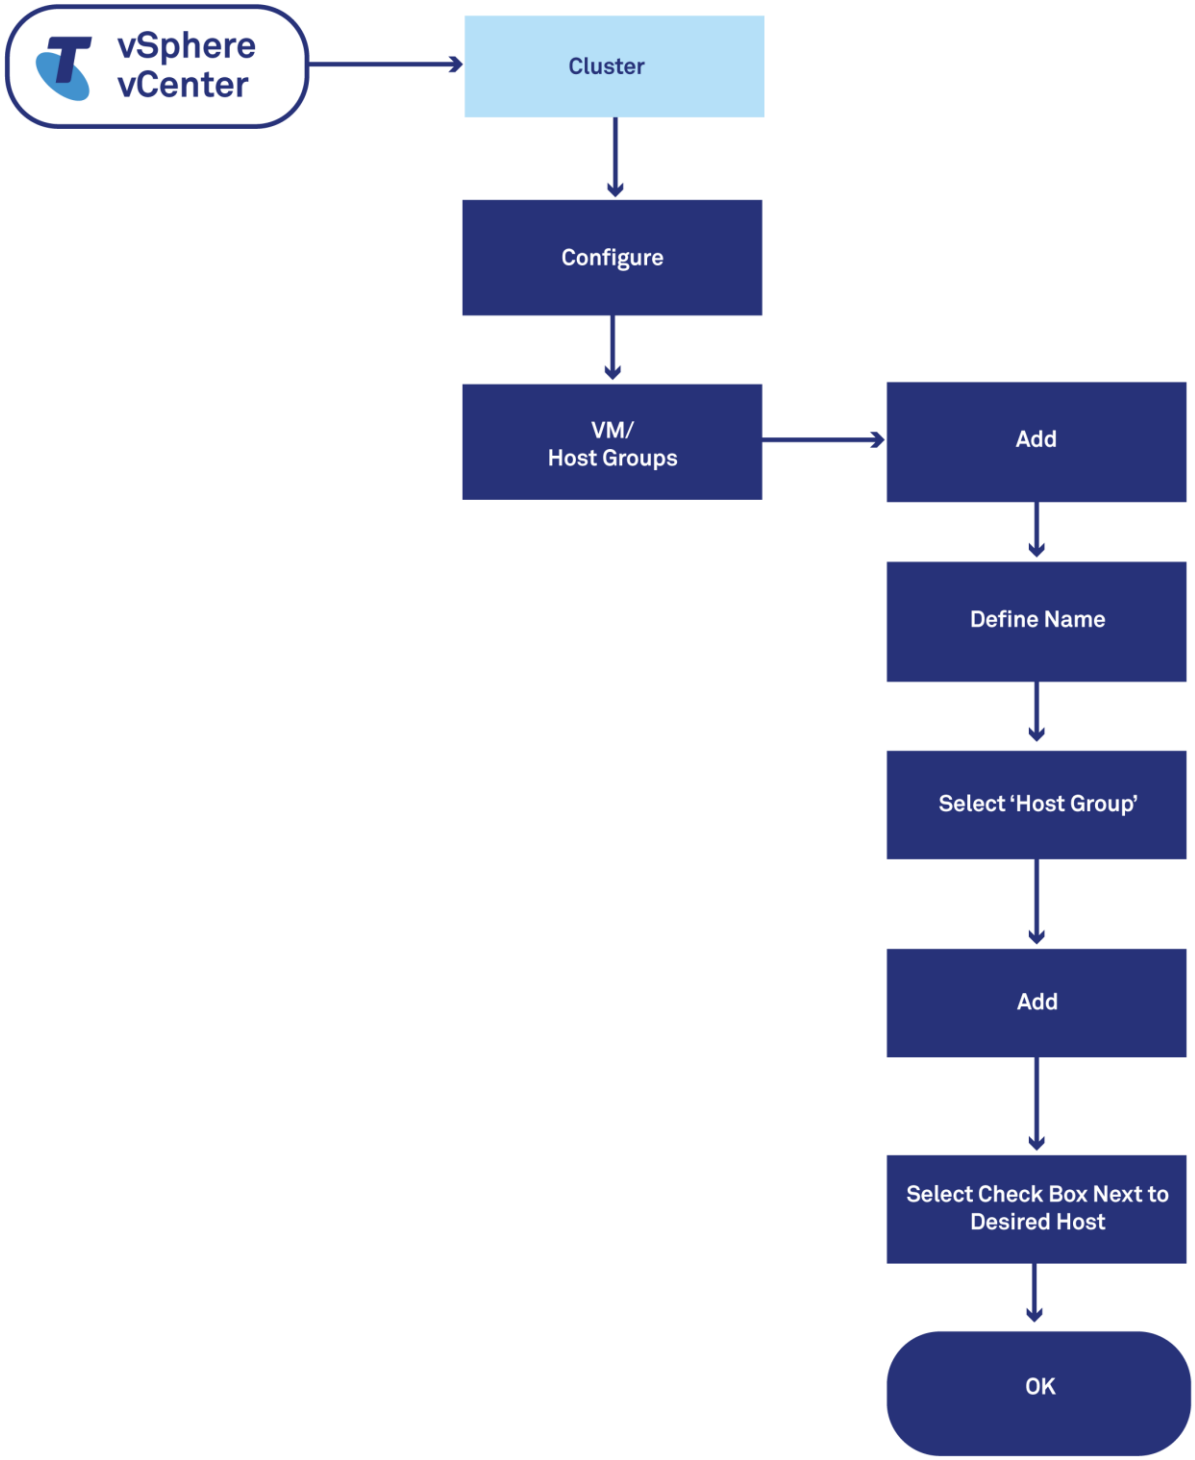
vSphere
vCenter
Cluster
Configure
VM/
Host Groups
Add
Define Name
Select 'Host Group'
Add
Select Check Box Next to
Desired Host
OK

Configuration Tips

When you create a Host DRS Group, please observe these points and recommendations:

- If you create Host DRS Groups with overlapping memberships, you must be mindful of the way in which VMware resolves conflicts with their corresponding affinity rules.

Information Resources

You can learn more about DRS Groups and VM-Host affinity rules and how to use them in TPC by referring to these VMware documents:

- Using DRS Affinity Rules
- Create a Host DRS Group
- VM-Host Affinity Rules
- DRS Cluster Validity

Task #VM04: Create a VM-Host Affinity Rule

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☑ vSphere Admin ☐ vSphere Read-Only ☐ NSX-T Admin	☐ Telstra Cloud Sight ☑ vSphere (Native) ☐ Telstra Cloud plug-in ☐ NSX-T

Purpose

To configure a DRS affinity rule that specifies a relationship between a DRS Group of VMs and a DRS Group of hosts.

Overview

This task refers specifically to VM-Host rules. You can categorise each DRS VM-Host affinity (or anti-affinity) rule as either *required* ('must') or *preferential* ('should').

These four common conditions are available for VM-Host Affinity rules:

1. Must run on hosts in group: VMs in VM Group X must run on hosts in Host Group A
2. Should run on hosts in group: VMs in VM Group X should run on hosts in Host Group A but are not required to
3. Must not run on hosts in group: VMs in VM Group X must never run on a host in Host Group A
4. Should not run on hosts in group: VMs in VM Group X should not run on a host in Host Group A, but they might.

Procedure

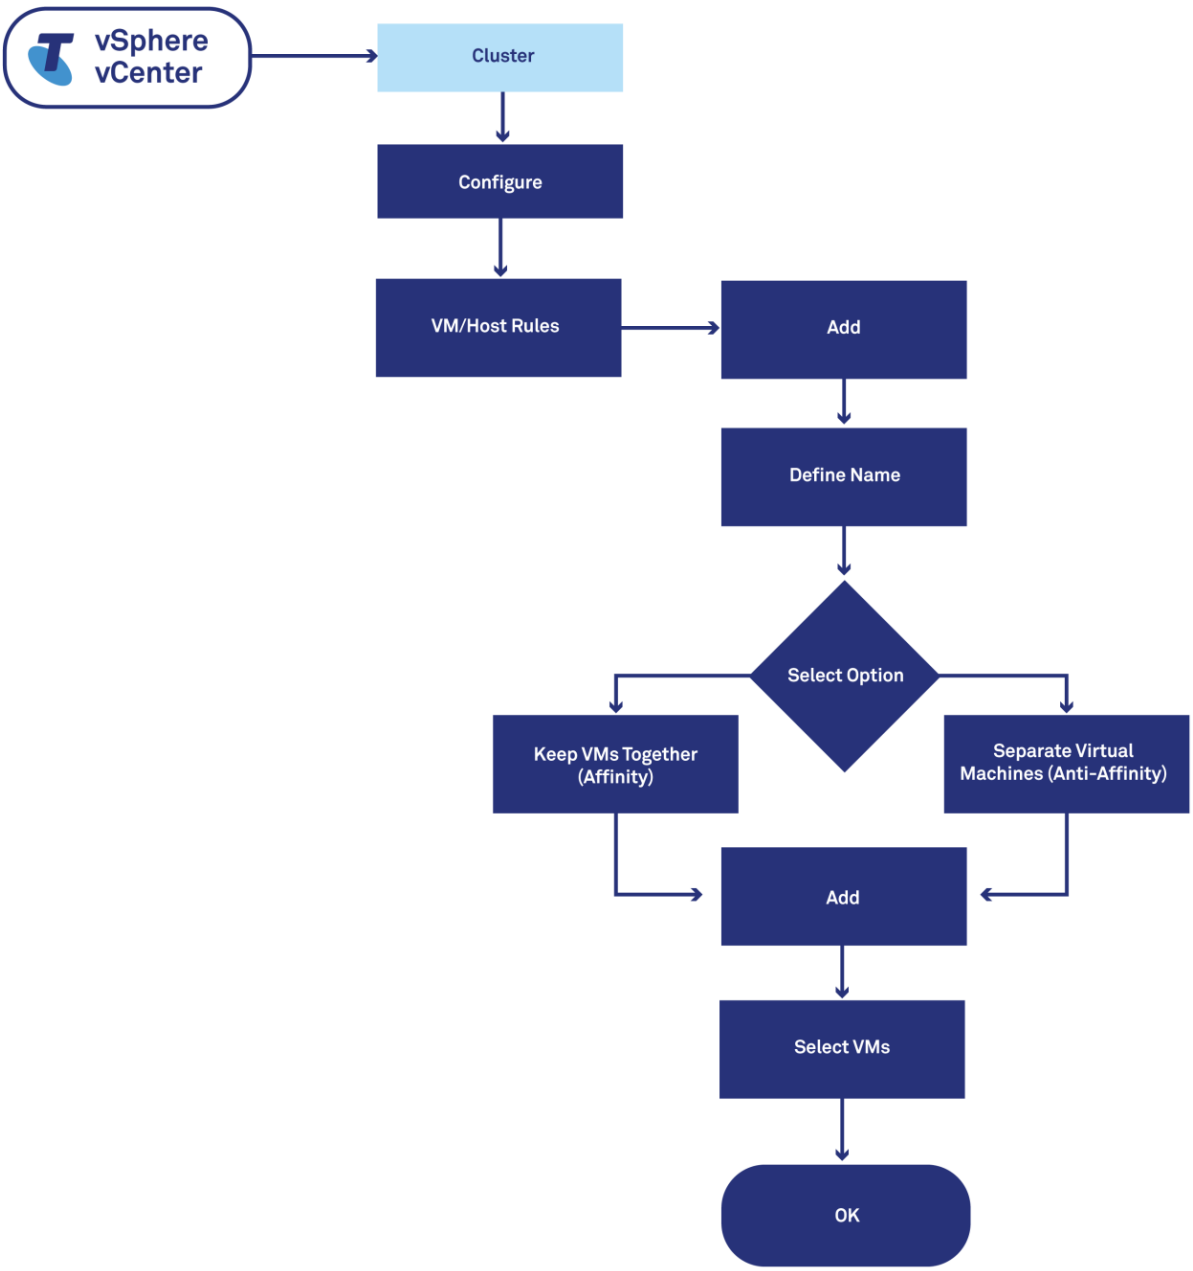
vSphere
vCenter
Cluster
Configure
VM/Host Rules
Add
Define Name
Select Option
Keep VMs Together (Affinity)
Separate Virtual Machines (Anti-Affinity)
Add
Select VMs
OK

Configuration Tips

DRS VM-Host affinity (or anti-affinity) rules apply to VMs and hosts in groups. You need to create your designated groups of each before you configure VM-Host affinity or anti-affinity rules. Refer to:

- Task #VM02: Create a VM DRS Group
- Task #VM03: Create a Host DRS Group

When DRS distributes VMs across a cluster of hosts, it will attempt to honour all your affinity (or anti-affinity) rules. However, sometimes it cannot fulfil the affinity conditions because of conflicting rules or a lack of available resources. This is called an *affinity violation*. vSphere reports affinity violations and the reason(s) they have occurred under 'Faults' in the DRS monitoring panel.

Information Resources

You can learn more about DRS Affinity rules and VMWare 'Best Practices by referring to these VMware documents:

- Create a VM-Host Affinity Rule
- Using DRS Affinity Rules
- VM-Host Affinity Rules

Task #VM05: Create a VM-VM Affinity Role

Process	Required User Type	Tool
☐ Telstra Provisioned ☐ Optional Task via Engagement ☑ Customer Configured	☐ TCS Tenancy Owner / Admin / Manager ☐ TCS Workspace Admin Owner / Admin ☐ TCS Cloud Service Admin / Manager ☐ TCS Cloud Connector Admin / Manager ☑ vSphere Admin ☐ vSphere Read-Only ☐ NSX-T Admin	☐ Telstra Cloud Sight ☑ vSphere (Native) ☐ Telstra Cloud plug-in ☐ NSX-T

Purpose

To configure a DRS affinity rule that determines whether an individual VM should run on the same or different hosts as other VMs in a DRS cluster.

Overview

This task refers specifically to VM-VM affinity rules.

Procedure

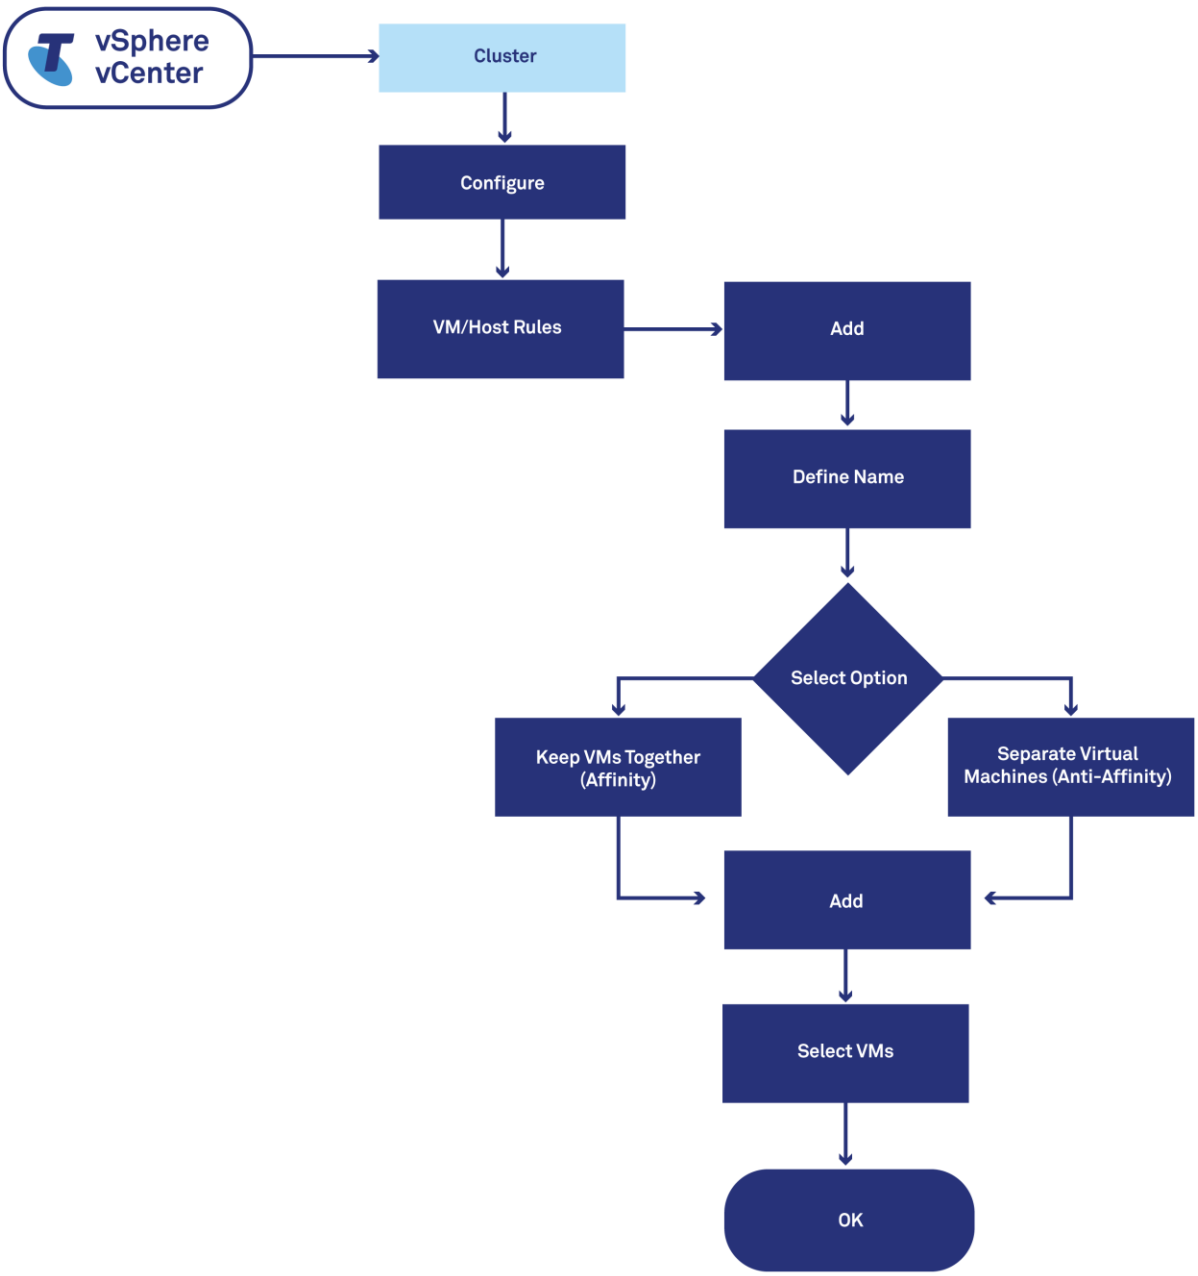
vSphere
vCenter
Cluster
Configure
VM/Host Rules
Add
Define Name
Select Option
Keep VMs Together (Affinity)
Separate Virtual Machines (Anti-Affinity)
Add
Select VMs
OK

Configuration Tips

You can nominate the VMs for your DRS VM-VM affinity or anti-affinity rule as you configure it. You do not need to define VM Groups in advance.

When DRS distributes VMs across a cluster of hosts, it will attempt to honour all your affinity rules. However, sometimes it cannot fulfil the affinity conditions because of conflicting rules or a lack of available resources. This is called an *affinity violation*. vSphere reports affinity violations and the reason(s) they have occurred under 'Faults' in the DRS monitoring panel.

Information Resources

You can learn more about DRS Affinity and how to use it in Telstra Private Cloud by referring to these VMware documents:

- Using DRS Affinity Rules
- VM-VM Affinity Rules

Chapter 5: Guidelines for Common Jobs

Telstra has produced Quick Reference Guides to help you navigate certain administration jobs in TPC. We have published those QRGs here.

In this section, we have re-produced images of the QRGs for particular jobs that new or less experienced customers routinely ask us about. We have included hyperlinks directly to key tasks covered in this User Guide to help you easily review and follow their steps if you are not already confident about them.

Job: Add a VM to a Public Network

Most commonly, you will add a virtual machine to a public network when you want that VM to be able to communicate with parties reached over the Internet.

If the meaning of a 'public network' is not clear to you, refer to Chapter 1: External Interconnects to Your vDC found earlier in this guide before continuing.

Figure 13: Quick Reference Guide: Add a VM to a Public Network

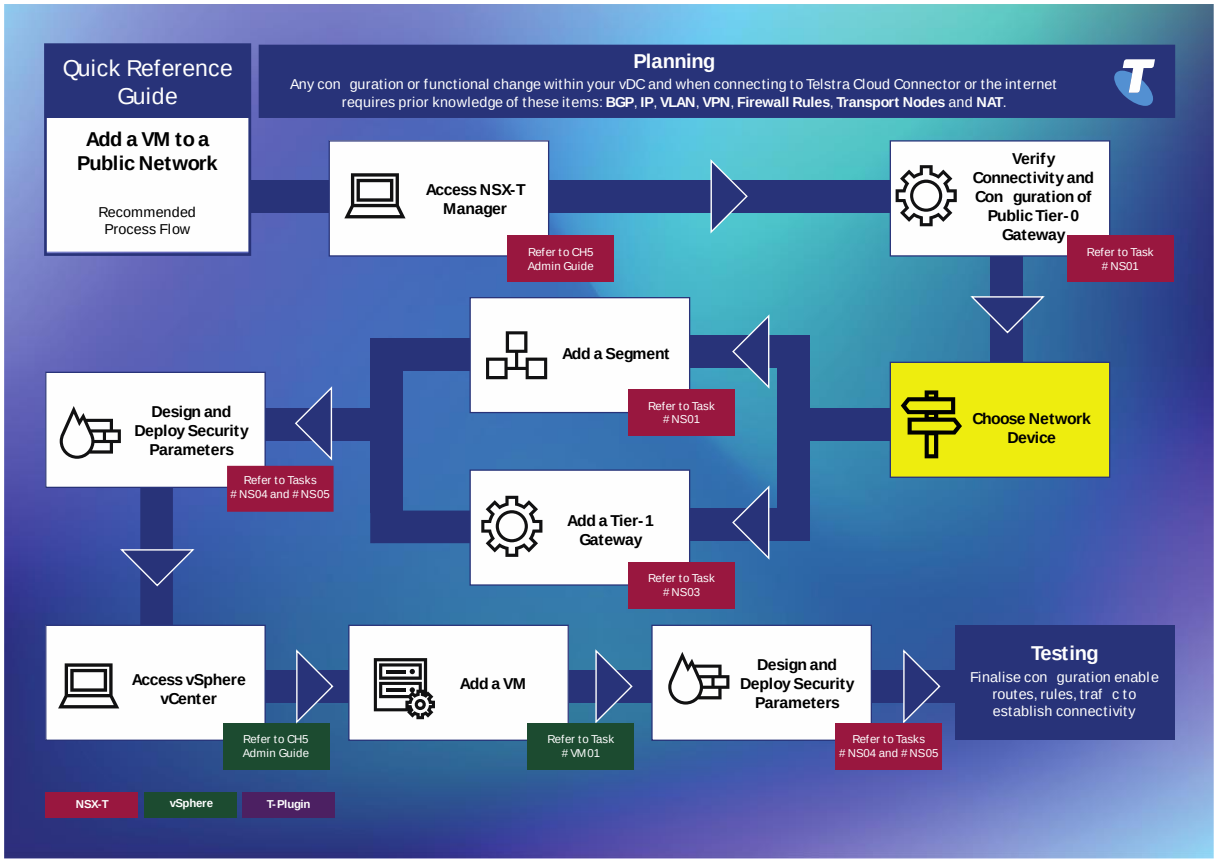

You can follow these links to the designated tasks shown above in Figure 13 to read more about them:

- Task #NS01: Add a Segment
- Task #NS03: Add/Modify a Tier-1 Gateway (Logical Router)
- Task #NS04: Add/Modify NSX-T Gateway Firewall
- Task #NS05: Add/Modify NSX-T Distributed Firewall
- Task #VM01: Create a Virtual Machine

Job: Add a VM to a Private Network

You will add a virtual machine to a private network when you do not want that VM to be able to communicate with parties reached over the Internet. Furthermore, you are likely to want that VM to be able to communicate with parties reached over your Private Interconnect to your Next IP VPN or other resources available through your Cloud Connector.

If the meaning of a 'private network' is not clear to you, refer to Chapter 2: vDC Topologies found earlier in this guide before continuing.

Figure 14: Quick Reference Guide: Add a VM to a Private Network

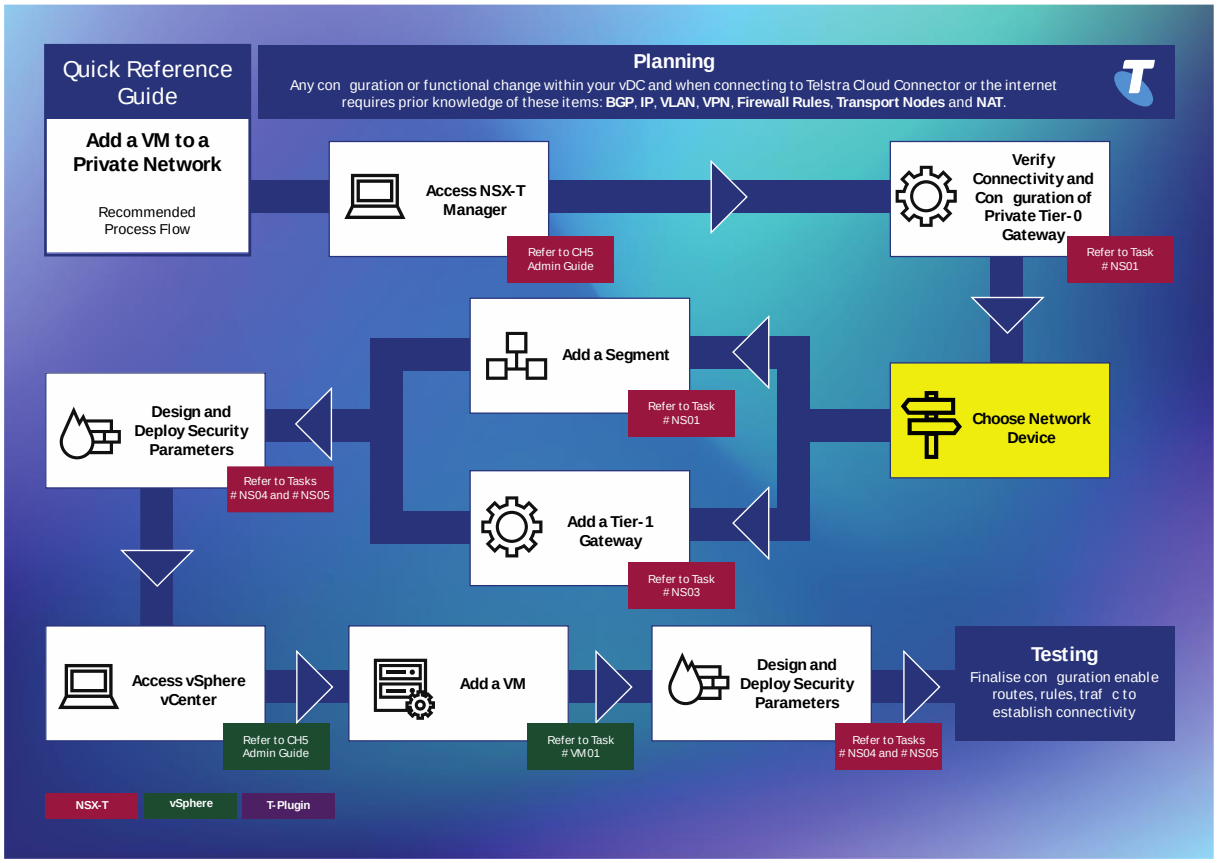

You can follow these links to the designated tasks shown above in Figure 14 to read more about them:

- Task #NS01: Add a Segment
- Task #NS03: Add/Modify a Tier-1 Gateway (Logical Router)
- Task #NS04: Add/Modify NSX-T Gateway Firewall
- Task #NS05: Add/Modify NSX-T Distributed Firewall
- Task #VM01: Create a Virtual Machine

Chapter 6: Resource Sizing Considerations

Once you assume control of your vDC, you need to allocate your vDC's resources to your workloads. Unless you have very simple needs, it is likely to require careful planning initially, and then ongoing monitoring and adjustments once you move into production or whenever you add or subtract resources or workloads. Moreover, resource allocation decisions are rarely subject to straightforward or firm rules, relying on operational judgements that balance several complementary and competing needs.

This section contains guidelines to help you effectively spread and manage the load in your vDC. They embody the advice of our technology partners as well as Telstra's experience running and hosting a substantial private cloud infrastructure for more than a decade. If you are new to TPC or don't have extensive experience with virtualisation, they provide important advice. Even if you are a sophisticated user of this technology, you may still find our guidelines provide a useful comparison against your established practices and metrics.

TPC Host Characteristics

Modern physical CPUs (pCPU), particularly those in data centre-based hosts, commonly contain multiple processing *cores*. Each core can focus on its own task at any one time, so the CPU can do multiple things simultaneously. In some pCPUs, a *physical core* can support hyperthreading, leading to a greater number of so-called *logical cores*. Some hosts can even contain multiple pCPUs (most often two or four) each served by its own RAM, further increasing processing power or capacity.

TPC offers different types of hosts, each differentiated from the others by its combination of pCPU count, number of cores and quantity of RAM. At the time Telstra launched TPC in 2022, there were three options:

1. 1 pCPU and 128 GB RAM: each pCPU has 20 physical cores and 40 logical cores
2. 1 pCPU and 256 GB RAM: each pCPU has 24 physical cores and 48 logical cores
3. 2 pCPUs and 512 GB RAM: each pCPU has 20 physical cores and 40 logical cores. In this case, each half of the RAM (256GB) serves one processor respectively.

When you define a virtual machine (VM) in vSphere, you will set the number of vCPUs and quantity of RAM it can consume. If the pCPU supports hyperthreading, ESXi automatically exploits it by scheduling each vCPU to a logical core; otherwise it will schedule a physical core.

Rightsizing Your VMs

If you assign more resources to a VM than it can use effectively, you may unintentionally harm the overall performance of your vDC. Software vendors typically recommend an optimal mix of resources, such as RAM and storage, to allow their applications to run well.

Where it matters, an application vendor may also nominate desirable CPU characteristics, such as the number of cores the application can use. You should observe these recommendations as you develop the VM's specifications and select its host to avoid performance issues or unnecessary expense.

Allocating Virtual CPUs

To benefit from using multiple vCPUs, your VM's applications and workloads must exhibit characteristics that can exploit them. Typically, this means:

1. If your VM hosts a single main application: the application employs *multi-threading* by dividing its operation into two or more units of (parallel) execution. The system can provide each thread its own vCPU, or
2. If your VM hosts multiple applications: the system can provide each application with its own vCPU (or more than one if it uses multi-threading).

If you configure your VMs to use multiple vCPUs, each one will trigger small resource overheads in the ESXi hypervisor, plus context switching penalties in the processor that generate real pCPU consumption on the host, even if some of those vCPUs sit unused. While ESXi includes features that can detect idle vCPUs and somewhat reduce their footprint, it cannot entirely erase the impact. Indeed, a heavily loaded system might amplify it.

Consequently, as a matter of standard practice, you should not over-allocate vCPUs to your VMs because it can reduce, rather than improve performance. Instead, try to understand the needs of each VM's applications and match them to your VM settings accordingly.

Over-allocating vCPUs to a VM can reduce performance. Understand your workload and determine the number of vCPUs it will actively use.

For more information, refer to this article from Intel

Turbo Boost

A pCPU with idle physical cores can shut down the inactive ones to reduce heat and save energy. It might also divert the saved energy to remaining active cores and increase their clocking rate, allowing applications to run faster. Intel calls this feature *Turbo Boost* (or some refinement thereof). Other vendors have their own names.

If you keep the number of vCPUs consumed by VMs to fewer than the number of cores in your pCPU, you encourage the conditions to induce Turbo Boost. This can be valuable when a key VM (or several of them) hosts a single threaded application that would benefit from a faster clock rate to increase processing speed.

Allocating fewer vCPUs than available cores can create the conditions for Turbo Boost to occur

Multi-Processor Hosts and Virtual CPUs

Some hosts contain multiple pCPUs to further increase their processing capacity. (Within the industry, CPUs are also called colloquially called *sockets*, after the plug-in housing in which some CPUs sit.) These hosts generally employ non-uniform memory access (NUMA) to help accelerate processing speed. NUMA localises some of the on-board RAM to one of the pCPUs, making it is faster to access. When necessary, the pCPU can draw data from other non-localised or shared memory, but that takes slightly longer. Consequently, the system operates best when an application's data sits in RAM local to the pCPU it is running on.

In the absence of multiple cores, a multi-processor host can still support multiple vCPUs simply by consuming an entire processor for each. But multiple cores are more likely, and each core will still inherit the speed benefits of localised RAM. When you need multiple vCPUs for a VM running on such a host, your configuration choices can influence whether the system draws all vCPUs from the same pCPU or spreads them.

Should your VM demand more vCPUs than one pCPU contains in cores, the system will necessarily distribute them over two (or more) pCPUs. In this case, if you can reasonably lower your vCPU demand (for example, can you split your VM?) you can then fit the VM on one pCPU. That said, while not impossible, VMs that need more than 32 vCPUs are unusual so the individual pCPUs on TPC's hosts have enough cores to hold all vCPUs for most VMs you are likely to use.

Provided your VM can fit on one pCPU and you don't configure it to override default behaviour, an ESXi system will always try to optimise performance by placing its vCPUs on one pCPU and allocate memory from localised RAM.

To lessen the performance impact on multi-processor hosts, you should configure your VM's settings to allow the system to co-locate its vCPUs on a single pCPU and use localised RAM

Host-VM CPU Contention

Even busy applications are unlikely to need 100 per cent of the host resources available to them all the time, which usually leaves spare capacity for other tasks. This is the premise of virtualisation. Indeed, one of the advantages of virtualisation is that one host can carry multiple workloads, resulting in financial, operational and environmental benefits to the enterprise.

You induce contention when you commit more resources to your VMs than the physical capacity of your infrastructure, particularly its CPU(s). This contention is measured as a ratio of total vCPUs (tallied across all VMs on a host) to pCPU cores. For example, 60 vCPUs assigned to VMs running on a host with 20 cores is a contention ratio of 3:1.

If you embrace contention, what level is appropriate? While there is no absolute rule, Telstra's experience suggests CPU contention ratios from 3:1 to 5:1 generally constitute a good range when the VMs aren't chronically busy.

A host-VM CPU contention ratio of 3:1 to 5:1 is a good starting point for many workloads

You can use various VMware operating statistics to help identify whether a particular host is dealing acceptably with CPU contention. In addition to our tips below, we recommend following VMware's documentation on these subjects to manage your vDC.

CPU Ready Time

CPU ready time measures the average amount of time a VM was ready to execute, but the system could not immediately find a host CPU (ie. a core) to accommodate it.

There are many public texts on this topic. They generally assert that a value of 50ms or less is normal, and that any value of greater than 1000ms may suggest problems. Higher numbers are certainly less desirable, but if your monitoring reveals a value of 1000ms or more, or you encounter an alarm, you should look at additional factors before you conclude there is an issue or begin to modify your installation.

The nature of your applications matters. For example, an application built for constant user I/O, such as virtual desktop software, may exhibit unacceptable performance at lower values than a multi-threaded DBMS. Evaluate whether your application performance is poor or if users are dissatisfied with a particular application and if not, you may not need to do anything more than monitor conditions.

If your application performance could be better, look at your host-CPU contention ratio. If it is aggressive, you may wish to move some of your workloads to another host to lower contention and evaluate the outcome.

In contrast, if contention is conservative, look at whether some VMs are particularly busy or over-consuming resources. If so, they could be classic 'noisy neighbours', therefore consider shifting this VM (or perhaps those it affects) to another host with sufficient spare capacity, and set up DRS rules to better influence the future automatic placement of these VMs.

You can also look at your vCPU allocations. Have you assumed an overly heavy load and subsequently provided more vCPUs to some VMs than they really need? When these VMs occupy the pCPU, they may not use their allocated cores productively yet consistently delay timely access for other VMs.

If your CPU ready time is higher than you expect, look for correlating factors before you assume it is problematic

RAM Consumption

It is normal to oversubscribe your host's RAM among its member VMs. Unless you have specifically *reserved* memory for a VM, it is generally *thin-provisioned*, meaning that the system incrementally commits chunks of specified RAM to VMs as demand rises. On a busy host depleted of uncommitted RAM, the system may begin to swap memory contents in and out of disk (called *paging*) to continue to honour new requests. This dramatically lowers performance and may even crash the host.

You can use your vSphere dashboard to monitor RAM consumption. While it will naturally rise and fall as VMs claim and release memory, you should ensure consumption remains below 90 per cent to prevent paging and poor performance, and the associated risks to ESXi itself.

If your monitoring reveals that you have approached 90 per cent consumption and you think it might increase further, you might need to take immediate action. For example:

- If you have reserved memory resources for one or more VMs, re-evaluate those reservations
- Use vMotion to move a VM (or several of them) to another host(s) in the cluster that has sufficient spare capacity
- Temporarily shut down a VM (or several of them) and/or

- Restart a VM that has incrementally increased its RAM consumption over time despite a fairly stable workload and/or which you suspect has a substantial 'memory leak'.[4]

Monitor your host's active RAM consumption to ensure it stays below 90 per cent. If you believe it will rise above that threshold, consider taking action quickly

RAM for Individual VMs

In principle, you can allocate 100 per cent of the host's RAM to one VM. If you do this and the VM later attempts to utilise all the memory you have nominally allowed, it will induce paging and the associated risks to system operation.

As was the case for vCPUs, you should understand the profile of your VM's applications to allocate sufficient RAM for efficient operation but not be unnecessarily generous.

Never allocate 100 per cent of host RAM to a single VM. Understand your applications so you don't allocate more RAM to a VM than it needs

Use VMware Tools

VMware Tools is an ancillary set of services and modules for better management of guest operating systems in VMs on hosts. It is a package installed on a VM after it is built, and often re-installed after the virtualisation software is updated. Among other features, VMware Tools installs key drivers for virtual devices (eg. optimum network adaptors) and intercedes in host-guest OS calls to enhance resource management.

While the use of VMware Tools is not mandatory in ESXi virtualisation, there are few if any contemporary cases where you would not use it (there are some older virtual appliances that were incompatible with VMware Tools, but they have become increasingly rare). Indeed, it is quite possible that you will experience instability and connectivity issues with your VMs if you do not install VMware Tools and keep it up to date.

Some modern OS releases incorporate aspects of VMware Tools utilities and drivers with their default distributions. Nevertheless, we recommend you install VMware Tools as a matter of practice.

Always install VMware Tools and keep it up to date

[4] Such leaks occur when a process continually requests blocks of memory from the system but 'forgets' to give them back when no longer used or needed. Over time, this can significantly exhaust the common pool of RAM

Resource Locking and HA/DRS

If you reserve (sometimes called *lock*) resources for your VM(s), you can unintentionally influence the system's abilities to manage resiliency and balance load.

ESXi can manage resource demands using *slot policies*, which influence VM *admission control*. VM admission control monitors the system's capacity to launch a new VM without threatening system resilience. It ensures that within a cluster, a certain number of hosts can fail with sufficient remaining resources to re-home active VMs to running hardware. The number of hosts that can fail is set by the HA policy (which is 1 host by default in TPC).

A slot is a logical representation of the memory and CPU required to support a VM in the cluster. It is based on the largest running VM. If the largest VM consumes relatively little memory and vCPUs, then your slots are smaller so you can nominally run more VMs in the cluster than when the largest VM is resource hungry. Whenever ESXi evaluates its HA position for the cluster, it holds some slots in reserve to re-home VMs from a failed host.

ESXi continually re-evaluates slot size. If the largest VM grows in consumption, the slot size grows with it, leading to less slots per host. The admission control function subsequently prevents the system launching new VMs earlier than it otherwise might. (However, admission control does not force-stop any VMs already launched.)

Resource locking can skew ESXi's admission control behaviour. If a VM reserves a large amount of memory and/or vCPUs, it could conceivably become the benchmark size of a slot. This will cause ESXi to assume a large slot size, with the commensurate impact on cluster capacity and resource efficiency.

For example, consider a cluster of hosts each containing 128 GB of RAM. You have defined a VM with 32 GB of locked (ie. reserved) memory. Because a slot is based on the largest VM and this VM has reserved all 32GB of memory, it artificially limits each host to four slots (ie. 128 divided by 32), some of which are then reserved to comply with HA policy. In turn, DRS decisions are affected (which need to observe HA reservations), as is your ability to launch more VMs.

Resource locking is a valid operational tactic. However, it can cause ESXi to base its slot policy and admission control judgements on worst-case, theoretical resource demands rather than actual ones.

Be mindful of the impact of resource locking on admission control behaviour, particularly if you are considering speculatively reserving a large amount of RAM for one VM. Do not lock resources to a VM unless you are sure it is necessary

Power-on Priority

Shortly after activating your cluster, your VMs may not yet have consumed a large amount of memory, so the combination of HA policy and admission control allows you to launch any additional VM you need. If your VMs' resource consumption later grows, so will the slot size, which may then hinder your ability to boot more VMs. In extreme cases where this situation arises and you did not control the launch order of your VMs, a VM running an application of lower business value could prevent the system from booting another VM containing a key application.

To avoid this example, you can configure *power-on priority* for certain VMs in the cluster. This ensures those VMs launch first, ensuring their access to the resources they need. You can then launch the other VMs later.

Use power-on priority to launch your key VMs first. You might also consider preventing certain VMs from auto-starting if they skew slot policy admission control and stop other VMs from booting

Storage

Several complementary questions arise when you allocate virtual disks to your VMs:

- Do you need to reserve your entire virtual disk immediately (*thick provisioning*) or can it partly exist only in theory, with the system allocating real storage later at the time of use (*thin provisioning*)?
- If you plan to use thick provisioning, do you also want the system to wipe the disk clean of previous data on the physical media during provisioning (*eager zeroing*) or wipe it just prior to actual use (*lazy zeroing*)?
- Should I use one big datastore for my virtual disk(s), or spread them over more?

Thin-provisioned disks are quicker to provision. Since many VMs will run quite acceptably with thin-provisioned resources, this should be your default choice in most cases. However, thin provisioning can also cost additional latency when the system later brings additional storage into service because the VM begins to use it. While momentary, if this will unacceptably affect a VM holding a high-performance workload, thick provisioning might be more appropriate.

Industry professionals generally consider eager-zeroed storage to be more secure than lazy-zeroed because it prevents unauthorised recovery of deleted data. Because it needs to process the storage ahead of time, eager-zeroed storage is necessarily also thick- provisioned.

In general, thin provisioning is appropriate for most enterprise workloads. However, if you are creating a VM to host a high-performance workload, consider providing it with thick-provisioned, eager zeroed storage

Datastore Size

Storage characteristics, particularly around performance, will often influence application behaviour and responsiveness. TPC arranges internal storage into ***tiers*** named ***Standard*** and ***Performance*** respectively. Standard tier storage suits general file, print and mixed applications, while the Performance tier can better fit demanding workloads like DBMS and analytics. An individual datastore always consists of one tier, but your vDC can contain multiple datastores to provide some of each.

TPC limits datastores to 8TB. This is generally large enough to hold one or more virtual disks of a substantial size without allowing them to become so big they are unwieldy to manage.

You can share the cluster's datastores among different member hosts, and a host can use virtual disks stored on different datastores. By spreading your virtual disks across multiple datastores, you can exploit the performance characteristics of the different storage tiers to meet the needs of each workload, as well as compartmentalise faults or issues that may affect one particular datastore.

If you use multiple datastores for your VMs you can select storage tiers that suit your workload and separate the storage fault domains

Software Licences

Some software vendors include licence conditions that force their customers to buy enough licences to cover every host on which it may run. Since HA/DRS will nominally move a VM onto any available host in the cluster, you may wish to set DRS to limit the pool of candidate hosts for affected VMs, to lower licence costs.

Telstra is a Microsoft Authorized Mobility Partner. If you have License Mobility, you can bring your existing Microsoft licences to TPC.[5]

[5] Carefully check the wording of your licence agreement

Chapter 7: Other Considerations

Integrating Advanced Management Tools and Software

If you are an administrator experienced with virtualised environments similar to TPC, you may wish to install and use advanced management tools and software in your vDC. Examples include Carbon Black, vRealize Operations and other products from VMware, as well as enterprise-level backup software such as Veeam.

Telstra does not prohibit these tools and you are welcome to use them, but with conditions:

a. *We do not promise they will work*: we do not routinely test or validate them with TPC, nor establish the precise conditions and settings they need to operate successfully

b. *They cannot require root access to the tenancy's resources*: restricted access is an important principle that underpins the protection of your applications and data from outsiders, including other tenancies

c. *They must operate within our standardised authorisations*: these authorisations apply to every tenancy. In other words, Telstra does not create bespoke or customised settings in your tenancy to allow a package to work.

Regarding the third point: if you identify an authorisation issue that may hinder your software's operation, please raise it with your Telstra account manager or authorised Telstra partner, who can notify our TPC product team for further investigation.

If you want to review a list of access control permissions for your vDC, refer to this menu hierarchy in your vSphere portal:

Access Control -> Roles -> Telstra Customer Admin Role

Commercial Considerations

You must carefully consider consequences of your configuration actions and those of your management tools and software. You should avoid changes that are detrimental to the vDC's connectivity, security or robustness, such as:

- Leaving the vDC isolated from external networks
- Creating a security vulnerability by a poor application of policy rules
- Altering workload policies that hamper the ability of VMware's DRS and/or HA to shift workloads to manage resource demands and withstand equipment failures.

Telstra may charge you a fee to reset resources in your tenancy to restore correct functioning.

Planning Migration to or from TPC

Telstra Private Cloud offers hosted private cloud infrastructure and employs a 'shared responsibility' management model:

- In concert with our strategic partners, Telstra manages and maintains the infrastructure platform
- Our customer (ie. you) configures all required compute and network features, and provides, installs and operates all IT applications.

Importantly, your management responsibilities also encapsulate all necessary planning and activities required to migrate your applications and data into your Telstra Private Cloud vDC after you acquire it, or out of TPC if you wish to leave. Telstra does not provide complimentary professional services to help you do this as a matter of course.

There are many technical aspects of migration you should consider and plan carefully before you commence. Table 1 summarises some of those considerations to help visualise the potential size and complexity of the task.

Table 1: Planning Migration

Category	Examples of Questions and Considerations to Address
Backups and archives	• Can I retain copies of past backup and archive files, and are they portable? • Which many backups and archives should I keep and/or move? • How might I export them? • How might I import the exports into my new cloud / tool?
Virtual machines and workloads	• How do I plan to export my VMDKs and related files? • How do I manage downtime across the transition? (SLAs for business-critical applications) • Are my OS and application licences portable, or do I need to acquire new ones? • Can I easily transfer my firewall rules and NIC configurations to a new installation?
Secure export	• What tunnelling mechanisms are available to link current and future platforms to secure in-flight data during export and import?
IP addresses	• Where will I obtain (ie. who will provide) replacement public IP addresses (if leaving TPC)? • Do my topology changes requiring planning for a re-arrangement of my private addressing as well? • How will I manage the address transition to avoid downtime for my users?
SSL VPN	• How do I preserve remote user access to their key applications across the transition? • Can I export a list of current users for import on the new platform?

Category	Examples of Questions and Considerations to Address
	• How do I manage passwords (since they are generally not reversible or exportable)? • The SSL VPN solution may change. How do I and my users deal with new tools and profiles?
Firewall rules	• How do I modify rules to account for new addressing plan? • Is the firewall technology different? If so, how do I translate the rules effectively to maintain my security posture and avoid gaps?
DRS rules	• If the underlying virtualisation technology changes, how do I capture and re-create my affinity/anti-affinity groups and rules?
Network configuration	• What internal network routes do I need to exchange with the provider or outside networks? • Can I re-create my topology in the new facility? • What load balancing and NAT functions are available to me?
Monitoring solutions	• Will my current self-managed monitoring solution work with the new tenancy?
Overlaid products	• Did / does my current TPC vDC host other overlay products from Telstra that support other areas of my business? Examples can include Digital Managed Services, Telstra Backup as a Service, Disaster Recovery as a Service and Gateway Protection Advanced.

In addition, you will need to consider the financial impacts of the migration such as:

- One-off costs generated by the movement of significant amounts of data
- Early termination charges (ETCs) if you are breaking a contract commitment
- Whether the aggregated cost of multiple cloud-hosted workloads is lower on hosted private cloud (eg. TPC) or hyperscaler-based public cloud.

Telstra Purple and authorised partners offer paid professional services to business and enterprise customers that need assistance with their TPC migrations. Refer to your account manager or authorised partner for more information.