CHAPTER 16

Configuring LACP in vSphere Distributed Switch

In this chapter we will discuss the importance of Link Aggregation Control Protocol (LACP) in the vSphere Distributed Switch (vDS) environment. LACP plays a role in enhancing network performance and reliability by combining network connections. By understanding how to configure and utilize LACP within vDS, you will be well equipped to optimize your VMware infrastructure for increased throughput and redundancy.

16.1 What Is Link Aggregation Control Protocol (LACP)

Before diving into the usage of a feature, it's essential to understand its purpose and functionality. In this section, we explore LACP, also known as Ether-Channel, Ethernet trunk, port channel, LACP, vPC, and Multi-Link Trunking. It's important to note that the availability and configuration of LACP may vary depending on the specific switch vendor and model you are using.

L. Patrão, *VMware vSphere Essentials*, https://doi.org/10.1007/979-8-8688-0208-9_16

Link aggregation combines multiple physical network links into a single logical link, providing higher bandwidth and redundancy in case of link failures. Two common link aggregation methods are used: EtherChannel and LACP (IEEE 802.3ad).

Link aggregation, which enables multiple physical links between network devices to function as a single logical link, can be achieved through both trunking and LACP. However, they differ in their approach:

- **Trunking**: In trunk mode, individual ports within a link aggregation group are statically configured to form a trunk. The Switch treats these ports as a logical interface, facilitating traffic exchange between switches or network devices. Trunking doesn't require negotiation between devices and operates independently of any specific protocol.
- **LACP**: LACP is a standard protocol that allows two devices to dynamically negotiate and form a link aggregation group. LACP utilizes control frames for negotiation and management, providing advanced load-balancing techniques beyond trunking. With LACP, switches can automatically detect and configure link aggregation with other LACP-enabled switches.

In summary, while both trunking and LACP achieve link aggregation, trunking relies on static configuration without protocol negotiation. On the other hand, LACP is a dynamic, protocol-based approach that enables the automatic formation and management of link aggregation groups between devices.

Understanding these concepts will help you to proceed with configuring and using LACP in the vSphere Distributed Switch.

Link aggregation concepts:

Link aggregation is the process of combining multiple physical network links into a single logical link, providing higher bandwidth and redundancy in case of link failures. There are two common methods used for link aggregation:

- **EtherChannel**: EtherChannel is a proprietary technology offered by Cisco that allows for up to eight physical Ethernet links to be combined into a single logical link, providing fast connectivity and failover protection between switches, routers, and servers.
- **LACP or IEEE 802.3ad**: LACP is an IEEE standard included in the 802.3ad specification that uses a Link Aggregation Control Protocol to bundle several physical ports into a single logical channel automatically. When two devices support LACP, they can negotiate the creation of a link aggregation group and the load-balancing method used to distribute traffic across the links.

While both technologies achieve the same goal of link aggregation, they have some differences. EtherChannel is a Cisco-proprietary technology that can be used only with Cisco equipment. In contrast, LACP is an open standard that can be used with any vendor's equipment that supports it. Additionally, EtherChannel provides greater flexibility in load balancing, while LACP mandates a specific load-balancing algorithm.

Figure 16-1 illustrates the design of a vDS with LACP using Link Aggregation Group (LAG) configuration and the connection to physical network interfaces and switch ports. In this scenario, we have three ESXi hosts connected to the Management vDS. Each vmnic (virtual machine network interface card) from the ESXi hosts is connected to corresponding physical switch ports and configured with LACP at the switch level.

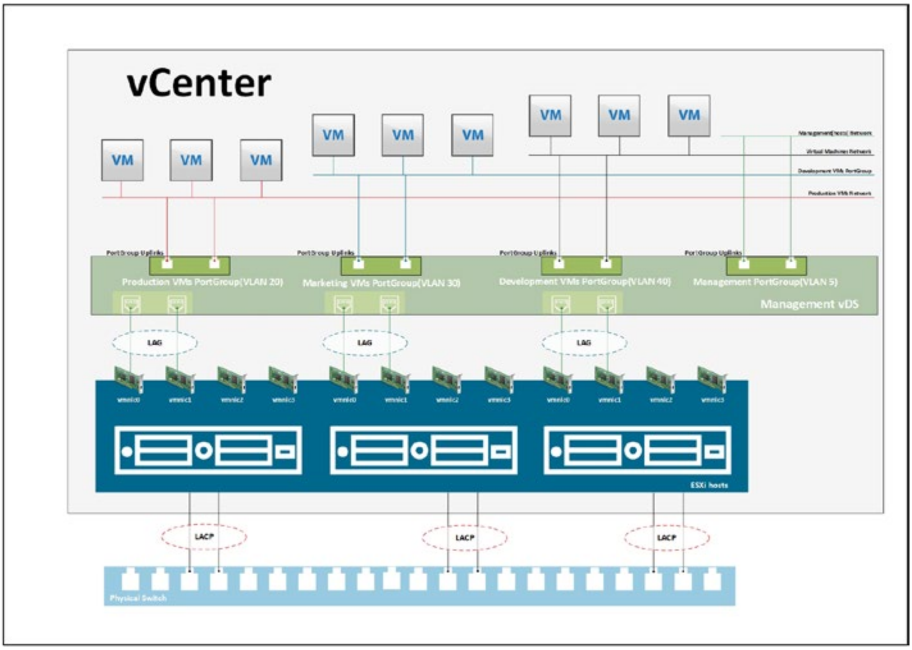

Figure 16-1. *(LACP 1 of 1)*

This configuration exemplifies implementing LACP in your environment, showcasing how the vDS, ESXi hosts, and physical switches are interconnected and configured to utilize LACP for improved network performance and redundancy.

A Link Aggregation Group (LAG) is a logical grouping of multiple physical network links or ports combined to form a single high-bandwidth connection. LAGs are commonly used to increase network capacity, enhance redundancy, and improve overall network performance.

The diagram in Figure 16-1 illustrates the concept of a LAG. The vDS has three ESXi hosts, each of which has multiple physical network interface cards (NICs) connected to the physical switch. Instead of treating these individual NICs as separate connections, you can create a LAG by combining them into a single logical link. For example, you can configure LACP on the physical switch and the vDS to enable dynamic negotiation and management of the LAG. LACP will establish a LAG between the physical switch and the vDS by bundling the multiple NICs together.

Once the LAG is formed, virtual machines and other network resources in your vSphere environment can utilize the combined bandwidth and redundancy provided by the LAG. This means traffic can be distributed across the individual NICs within the LAG, allowing for increased throughput and improved resilience.

By using LAGs, you can effectively utilize multiple physical links as a single logical link, providing higher bandwidth and fault tolerance for the VMs networks.

16.2 Requirements and Limitations of LACP in vSphere

In order to effectively manage your network, it is crucial to understand the requirements and limitations associated with implementing the Link Aggregation Control Protocol in vSphere vDS. It's important to consider how LACP will integrate with your existing network infrastructure, including hardware and software compatibility as the configuration of physical NICs. By being aware of these limitations, you can optimize your network planning to fully leverage LACP's benefits within your vSphere environment.

The following restrictions and limitations apply when using LACP in vSphere vDS:

- LACP requires a vSphere Enterprise Plus license for the vDS feature.
- An ESXi host supports NIC teaming only on a single physical switch or a stack of switches.
- Link aggregation is not supported when using different trunked switches. To enable link aggregation, the

Switch must be configured to perform 802.3ad link aggregation in static mode ON, while the virtual Switch should have its *Load balancing* method set to **Route based on IP hash** (demonstrated later in this chapter).

- Enabling either **Route based on IP hash** without 802.3ad aggregation or vice versa causes disruptions in networking. Therefore, it is recommended first to change the virtual Switch. This results in the service console being unavailable. Still, the physical switch management interface remains accessible, enabling aggregation on the involved switch ports and restoring networking.
- Do not use LACP for iSCSI software multipathing. iSCSI software multipathing requires just one uplink per VMkernel, and link aggregation gives it more than one.
- Do not use beacon probing with IP HASH load balancing.
- Do not configure standby or unused uplinks with IP HASH load balancing.
- VMware supports only one EtherChannel bond per Virtual Standard Switch (vSS).
- ESXi supports LACP on vDS only.
- In vSphere Distributed Switch 5.5 and later, all load-balancing algorithms of LACP are supported.

You need to ensure that the load-balancing algorithm used in ESXi matches the load-balancing algorithm implemented on the physical switch is essential. For inquiries regarding the specific load-balancing algorithm employed by the physical Switch, please consult the physical switch vendor's documentation.

Caution Due to potential network disruption, changes to link aggregation should be done during a maintenance window.

As with any networking change, there is a chance for network disruption, so a maintenance period is recommended for changes. This is especially true on a vDS because vCenter owns the vDS, and the hosts alone cannot change the vDS if the connection to vCenter is lost.

Enabling LACP can complicate vCenter or host management recovery in production-down scenarios because the LACP connection may need to be broken to move back to a Standard Switch if necessary (since LACP is not supported on a Standard Switch).

Limitations:

- vSphere Distributed Switches are the only supported switch type for LACP configuration.
- LACP cannot be used for software iSCSI multipathing.
- Host Profiles do not include LACP configuration settings.
- LACP is not supported within guest operating systems, including nested ESXi hosts.
- LACP cannot be used together with the ESXi Dump Collector.

Note The management port must be connected to a vSphere Standard Switch to use the ESXi Dump Collector.

- Port Mirroring cannot be used with LACP to mirror LACPDU packets used for negotiation and control.

- The teaming health check does not function for LAG ports, as the LACP protocol ensures the health of individual LAG ports. However, VLAN and MTU health checks can still be performed on LAG ports.
- Enhanced LACP support is limited to a single LAG per distributed port (dvPortGroup) or port group to handle the traffic.
- Up to 64 LAGs can be created on a distributed switch, with each host supporting up to 64 LAGs.
    - The actual number of usable LAGs depends on the capabilities of the physical environment and the virtual network topology.
    - For example, if the physical switch allows a maximum of four ports in an LACP port channel, you can connect up to four physical NICs per host to a LAG.
    - LACP is currently not supported with Single Root I/O Virtualization (SR-IOV, discussed in Chapter 26).

16.3 Configure LACP in vSphere Distributed Switch

Let's begin by configuring LACP in the vDS we created in the previous chapter about vDS. We will create a vDS named Vembu-LAG with a port group called LACP-LAG vDS.

Note Means that user can see different names that we have already seen in previous chapters, since this part needs to be done in physical lab, not nested. As stated above LACP is not supported in nested environments.

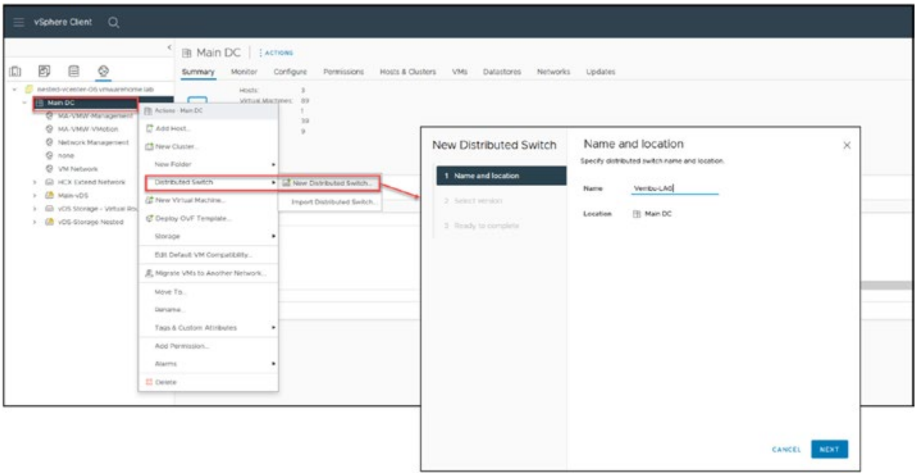

Figure 16-2. (LACP in vSphere vDS 1 of 13)

Although I will demonstrate the steps using our nested environment (the same we have been using since the beginning of this book), it's important to note that the actual LACP and physical configurations are performed in our physical VMware environment, which has been expressly set up with LACP for the purpose of this chapter.

After the vDS and port group are created, select the vDS you just created, click the **Configure** tab, select the **LACP** option, and then click **NEW**, as shown in Figure 16-3. Create a LAG called **LAG-Vembu**.

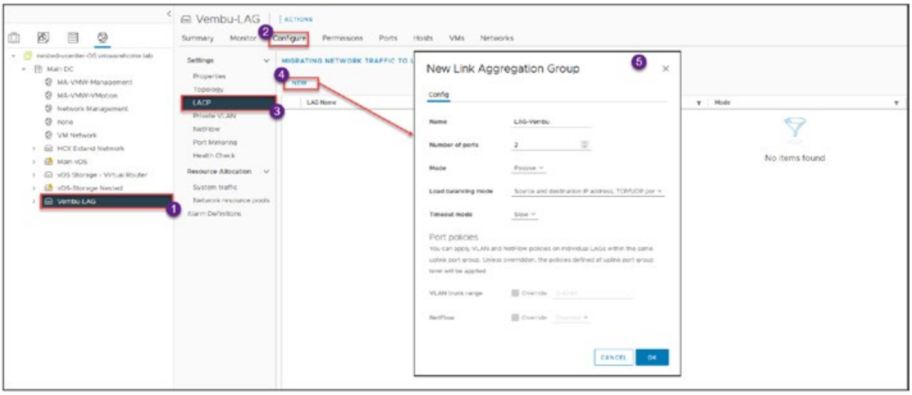

Figure 16-3. *(LACP in vSphere vDS 2 of 13)*

Set the *Number of ports* option to **2** to match the number of uplinks in your vDS.

Link Aggregation Group mode:

The vSphere Distributed Switch (vDS) LACP mode option is a configuration setting that determines how the vDS handles LACP negotiation with physical switches.

There are two LACP mode options available in vSphere vDS:

- **Active LACP mode**: In Active LACP mode, the vDS initiates the LACP negotiation with the physical switch. The vDS sends LACP packets to the switch and waits for a response. If the switch responds, the vDS forms a LAG with the switch.

- **Passive LACP mode**: In Passive LACP mode, the vDS does not initiate LACP negotiation with the physical switch. Instead, it listens for LACP packets from the switch. If it receives a packet, it will respond and form a LAG with the switch.

The choice of LACP mode depends on your network configuration and requirements. Active LACP mode is recommended if you have control over the physical switch and want to ensure that the vDS initiates the LACP negotiation. This allows you to configure specific settings, such as the LAG load-balancing algorithm, on both the vDS and the physical switch.

On the other hand, Passive LACP mode may be more appropriate if you connect the vDS to a third-party switch or an external service provider's network. In this scenario, you may not control the switch configuration, so Passive LACP mode allows the vDS to adapt to the switch's negotiation behavior.

Overall, the LACP mode option in vSphere vDS allows you to configure the LACP negotiation behavior to match your network environment, ensuring optimal performance and reliability.

In our example, Passive mode is used in the vDS LAG and Active mode is used in the physical switch.

The *Load balancing mode* option in vSphere vDS (see Figure 16-3) determines how traffic is distributed across the physical uplinks within a LAG. It controls the algorithm used to determine which uplink is used for transmitting traffic from virtual machines. Several *Load balancing mode* options are available in vSphere vDS, but since we will be selecting the **Route based on the IP hash** option in our vDS, select the **Source and destination IP address** (default) here.

This load-balancing mode is helpful for traffic that flows between the same source and destination IP addresses. The source and destination IP addresses are used to create a hash value in this mode. The hash value is then used to determine the traffic's uplink.

Moving down to the *Port policies* section of Figure 16-3, we are not using port policies in our LAG example, but I'll quickly explain how port policies work.

In vSphere, a LAG combines multiple physical network connections into a single logical connection, increasing bandwidth and redundancy. LAG port policies are a set of rules that define how the LAG is configured for each port.

When configuring LAG port policies, you can specify VLAN policies and NetFlow policies for individual LAGs within the same uplink port group. This allows you to apply different policies to different LAGs within the same group based on their specific requirements.

By default, the policies defined at the uplink port group level will be applied to all LAGs in the group. However, if necessary, you can override these policies for individual LAGs by specifying different policies at the LAG level. For example, you may want to apply a different VLAN policy to a particular LAG to isolate traffic from other LAGs in the same port group. Alternatively, you might want to apply a different NetFlow policy to a specific LAG to monitor traffic flow more closely.

Overall, LAG port policies provide a flexible and granular approach to configuring LAGs in vSphere, allowing you to tailor your network configuration to meet the specific needs of your environment.

Now that we have a vDS, a port group, and a LAG, we must set our Teaming from the vDS uplink to our LAG created in the previous step. Right-click the LACP-LAG vDS, click **Edit Settings**, and then select **Teaming and failover**, as shown in Figure 16-4. We need to move down the vDS uplinks to unused and move the LAG to active.

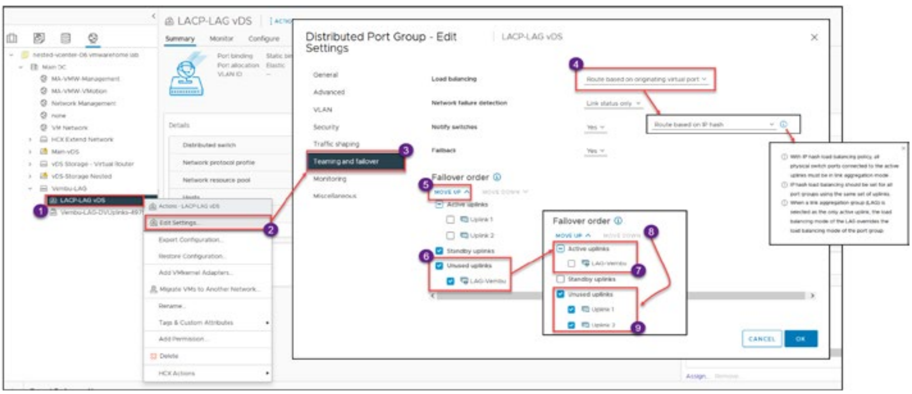

Figure 16-4. *(LACP in vSphere vDS 3 of 13)*

As previously explained, when using LACP, you need to set the vDS *Load balancing* field to **Route based on IP hash**, as shown in Figure 16-4.

Figure 16-5 shows the final result after finishing our vDS reconfiguring to use LAG.

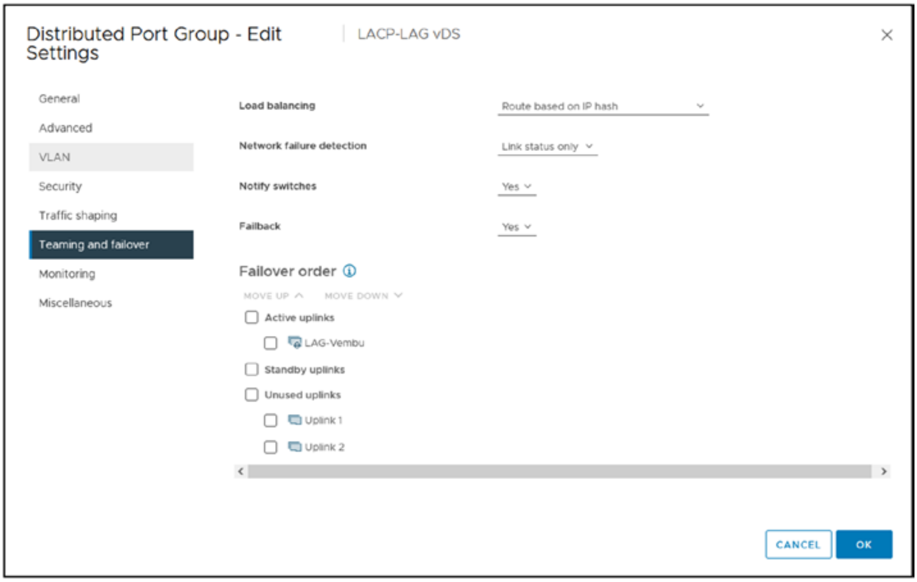

Figure 16-5. *(LACP in vSphere vDS 4 of 13)*

Now that we have our LAG configured in our new vDS, we will add ESXi hosts to our vDS.

Right-click the vDS and click **Add and Manage Hosts** in the context menu. See Figure 16-6. Since this is a new vDS, select the **Add hosts** option in the wizard that opens; If this was to be done in a vDS that already exists, in this case we need to create one, so extra steps are needed, then you would need to use the **Manage host networking** option (as described in Chapter 15 in the "Migrate from a Standard vSwitch to vSphere Distributed Switch" section). Click **Next**.

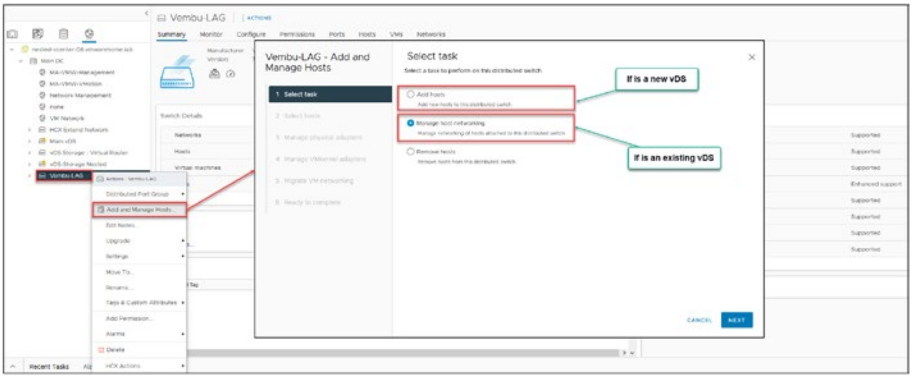

Figure 16-6. *(LACP in vSphere vDS 5 of 13)*

In step 2 of the wizard, select the ESXi hosts to add to this vDS, as shown in Figure 16-7. Click **Next**.

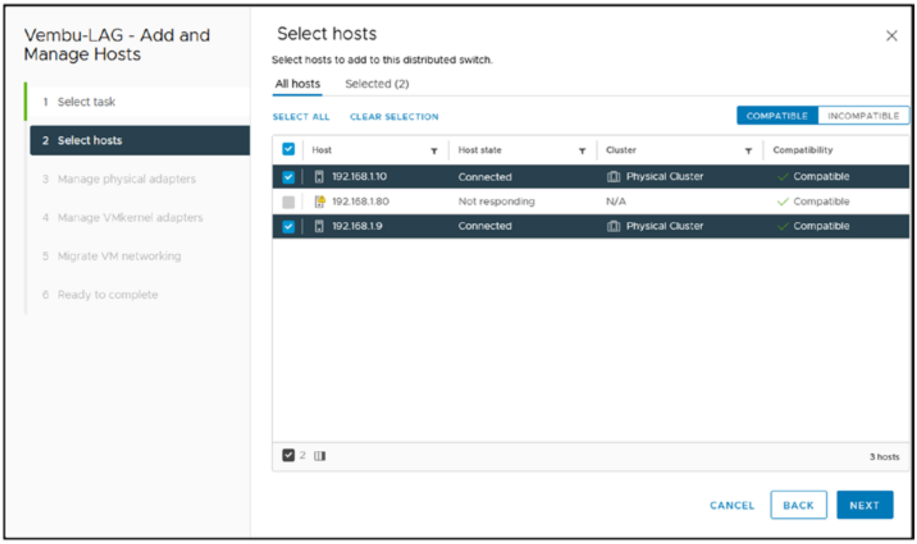

Figure 16-7. *(LACP in vSphere vDS 6 of 13)*

Next, select the vmnics connected to the physical switch ports and configured with LACP—in our case, `vmnic1` and `vmnic7`, as shown in Figure 16-8. Click **Next**.

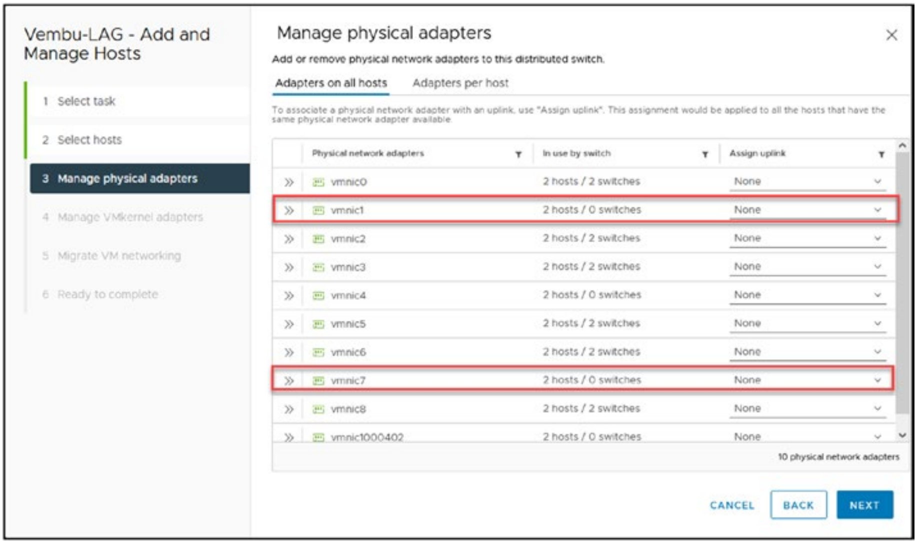

Figure 16-8. *(LACP in vSphere vDS 7 of 13)*

Note Don't forget that all your ESXi hosts' network interface connections need to be set as LACP in your physical switch.

In step 3 of the wizard, bind the vmnics with the LAG that you previously created, as shown in Figure 16-9. Click **Next**.

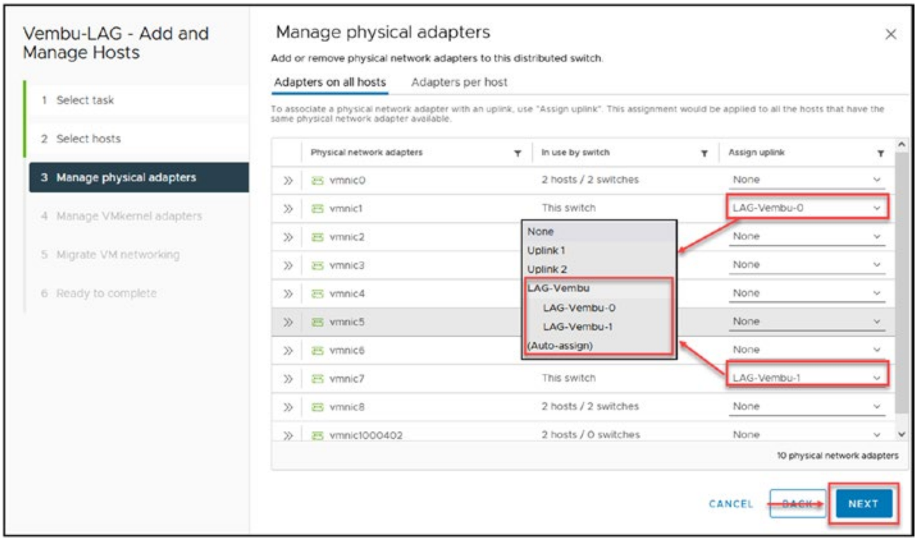

Figure 16-9. *(LACP in vSphere vDS 8 of 13)*

Step 4, shown in Figure 16-10, is about migrations of the VMkernel networks. In the next option we will skip this step, since this is only if you have VMkernel with LAG enabled, since is not the case we skip it. Click **Next**.

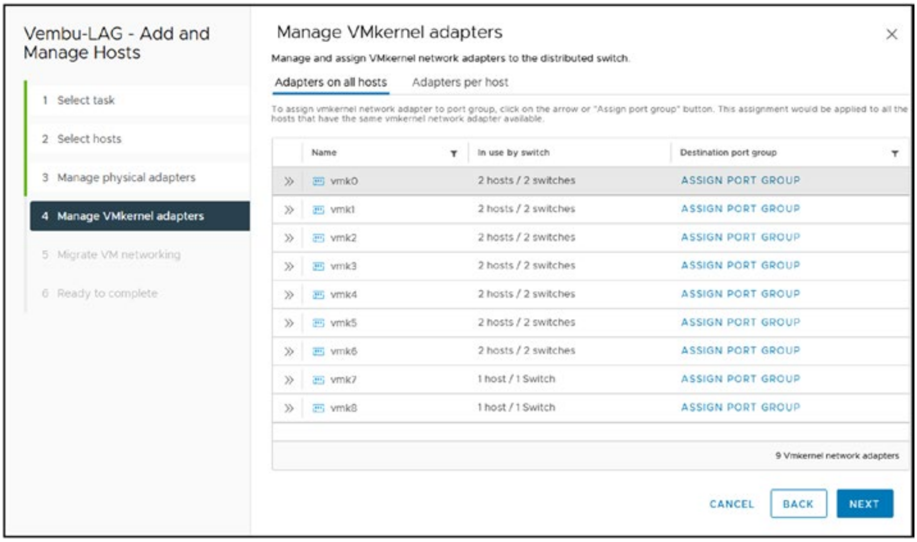

Figure 16-10. *(LACP in vSphere vDS 9 of 13)*

In step 5 of the wizard, you migrate existing VMs connected to this vDS to use the LAG port group. As shown in Figure 16-11, check the **Migrate virtual machine networking** check box and click the **Configure per virtual machine** tab.

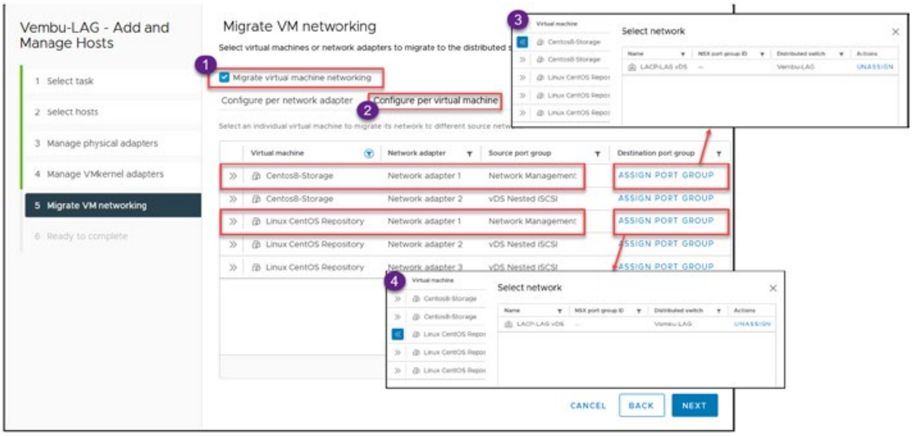

Figure 16-11. *(LACP in vSphere vDS 10 of 13)*

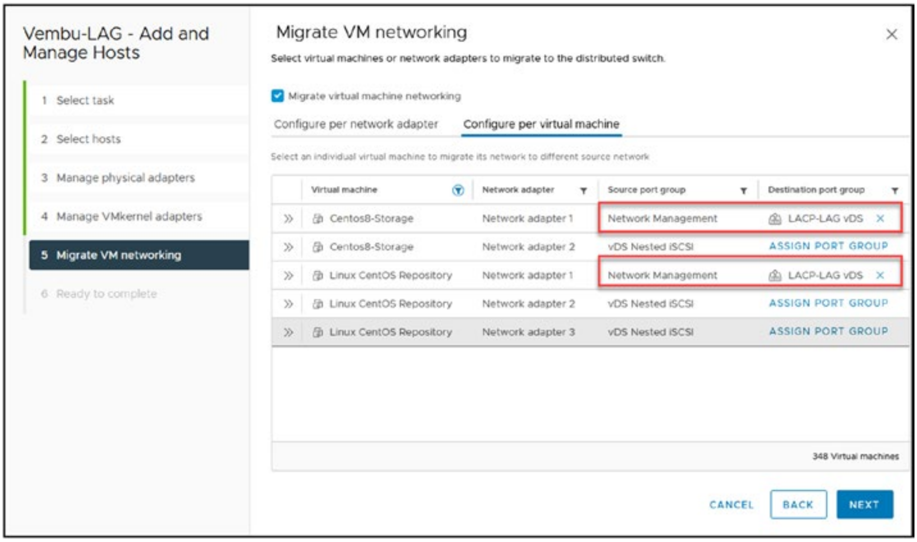

Figure 16-12. *(LACP in vSphere vDS 11 of 13)*

Note If your VMs have more than one virtual network adapter and are in different networks (one using LAG and the other not), but all are in the same network (for example, `nic1` is management and `nic2` is production), then you would use the *Configure per network adapter* tab.

In our test lab design, all VMs have more than one network adapter, but we want to change only the VM network management. Select the VM and the network adapter, click **Assign Port Group**, and **Assign**. Then click **Next**.

After VMs are connected to LAG, and will use an LACP network. Click **Next**.

In the final wizard step, shown in Figure 16-13, verify the details and click **Finish** to complete the port group configuration for utilizing LAG.

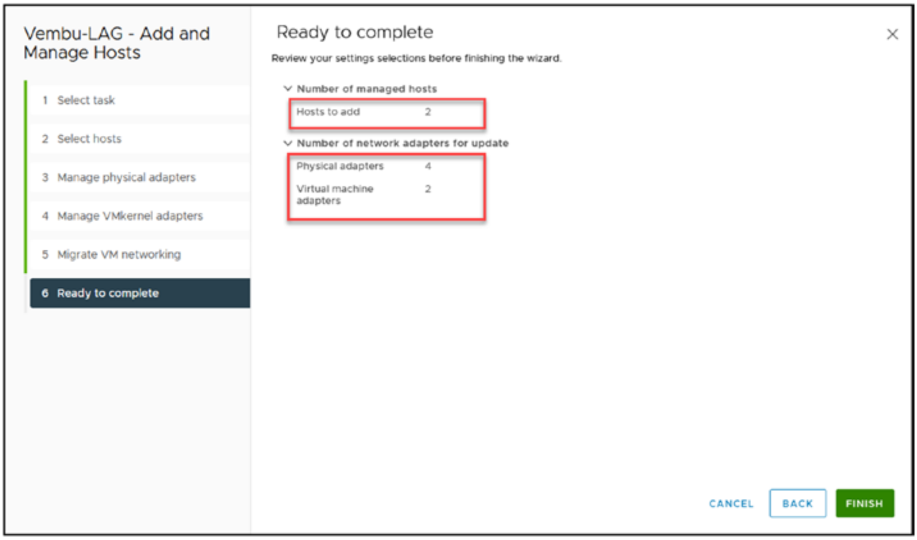

Figure 16-13. *(LACP in vSphere vDS 12 of 13)*

With the completion of all LACP configurations, you have successfully executed the following steps:

- Create a vDS for LACP.
- Create a LAG (link aggregation group).
- Set up the port group with LAG.
- Associate the LAG uplinks with vmnics.
- Transfer the VMs' networks to the LAG port group.

Now we can check the ESXi host console to see if the host uses LACP and is enabled on the physical switch. To do so, running the following command:

```
esxcli network vswitch dvs vmware lacp status get
```

If we see the physical switch and the physical ports number, as in Figure 16-14, then LACP is active and working.

Figure 16-14. *(LACP in vSphere vDS 13 of 13)*

This command shows us the configuration of our vDS with the LAG:

```
esxcli network vswitch vmware lacp config get
```

With the information, we have the guarantee that our LACP is adequately configured and active on both sides (in the vDS and the physical switch).

16.4 Migrate vDS VMkernel Network to vDS LACP

When enabling LACP (LAG) in our vDS and using LACP on our VMkernel, we need to consider the timeout or even losing communication while LACP is enabled.

In the previous section, we already checked how the LACP runs on both sides (physical and virtual), meaning that the LACP is working on your vSphere with LAG. But to migrate a VMkernel (that could be using Active/Passive Teaming or even Active/Active Teaming), one physical adapter must be done each time. If not, you could lose connection to your ESXi host or to your Datastores if using NFS.

Note Before starting the migration, if you have two new physical adapters (vmnics) configuring a vDS with LACP, this migration is more straightforward. But if you plan to use vmnics that are in use, the migration must be done one vmnic at a time.

Migrate VMkernel and vmnics to a new vDS with LACP using the same vmnics.

As previously explained, when migrating to an LACP but using the same physical network adapters, the process must be done differently, one vmnic at a time. If we enable LACP on both ports in the physical switch, we will automatically lose the connection in the Virtual Switch (a Standard Switch or a vDS) because we didn't change the settings in the vSphere. If we are doing this for a management VMkernel, then we don't have access to the ESXi host to make these changes.

Whenever there is a networking modification, it is advisable to allocate a maintenance window due to the potential for network disruptions. This is particularly crucial when working with a vDS since the vDS is controlled by vCenter, and individual hosts cannot make alterations to the vDS without a connection to vCenter.

Implementing LACP can introduce complications for restoring vCenter or host management functionality during critical situations, since the LACP connection may need to be terminated to revert to a Standard Switch if required (as LACP is not compatible with a Standard Switch).

How do we do this?

Note Since LACP is not supported in nested environments, we will not use the nested environment that we are using throughout the book. We will use a physical environment for these tasks.

First, this process is recommended with the ESXi in maintenance mode.

Figure 16-15 shows our network vDS for the physical environment: two ESXi hosts with a management vDS.

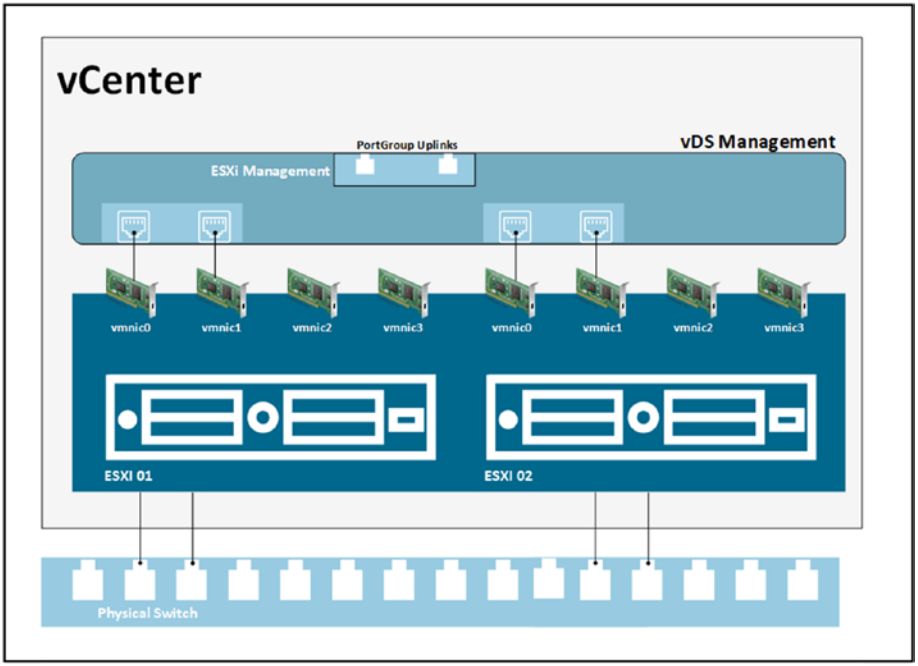

Figure 16-15. *(Migrate VMkernel LACP 1 of 11)*

As mentioned, we will use `vmnic0` and `vmnic1` for this migration. So our first step is to shut down the port in the physical switch where `vmnic0` is connected to the ESXi 01.

Note I will use my HPE ProCurve switch for the physical switch example. Don't forget that other switch vendors' commands are different to perform these tasks.

The following command are to shutdown the physical port in the Switch.

There were no number in my initial writing, only bullet points. But each number here is a step a command that needs to be done in physical switch, nothing related to what you see in the next image. Next image is a consequence of the steps, vmnic down. As is explained in the first paragraph.

This was done in a HPE physical switch, that is not discuss in this book, reader can have different devices, he just need to see what happen when we shutdown the port that is show in the Figure 16-16:

1. `conf` (enter the configuration shell).
2. `interface 13 disable` (shut down port 13 where vmnic0 is connected).

   After the port is disabled, you should see it in the vDS. See Figure 16-16.

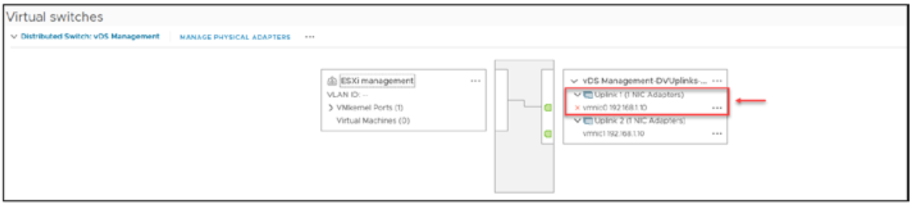

Figure 16-16. *(Migrate VMkernel LACP 2 of 11)*

3. `interface 13 LACP` (add port 13 or `vmnic0` to the LACP/EtherChannel configuration).

   In the vDS Management, create a LAG. See Figure 16-17.

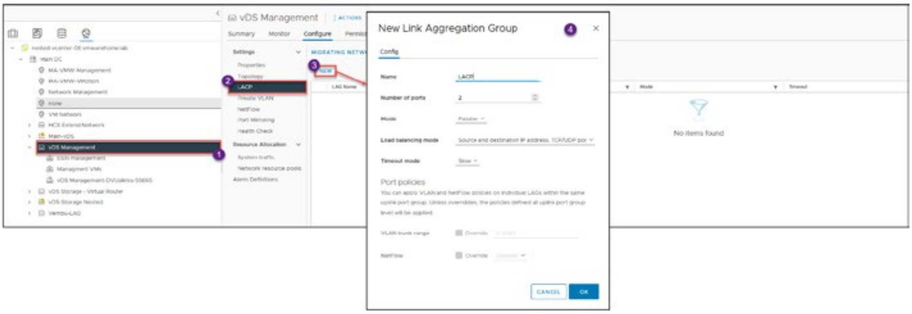

Figure 16-17. *(Migrate VMkernel LACP 3 of 11)*

Now let’s us add the `vmnic0` to the LAG uplinks.

On the vDS Management, select **Manage host networking** and then click **Next**. See Figure 16-18.

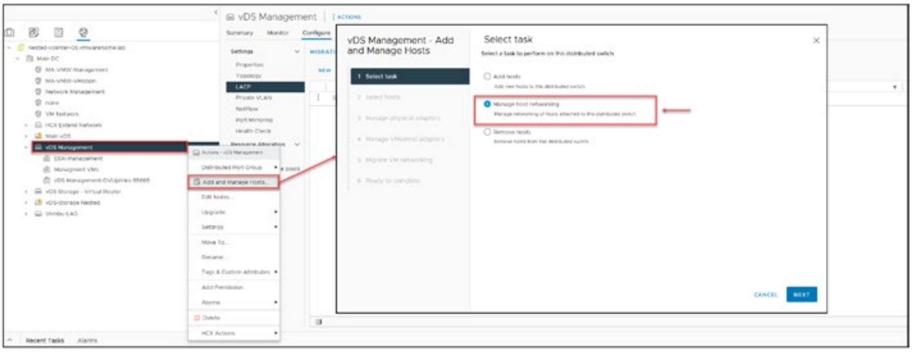

Figure 16-18. *(Migrate VMkernel LACP 4 of 11)*

Then select the ESXi host on which you disabled the vmnic, in our case ESXi 01, and click **Next**. See Figure 16-19.

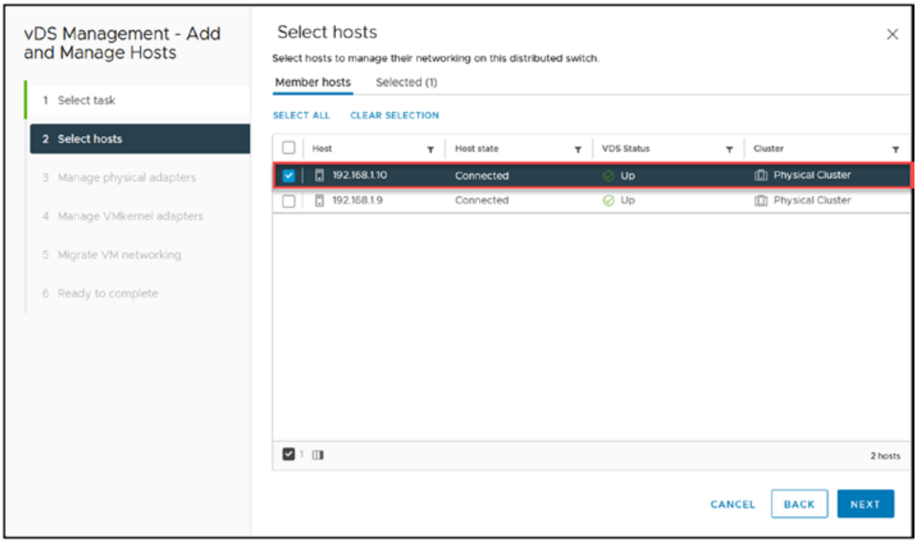

Figure 16-19. *(Migrate VMkernel LACP 5 of 11)*

Next, select `vmnic0` and that interface to the LAG uplink (the vDS LACP). See Figure 16-20.

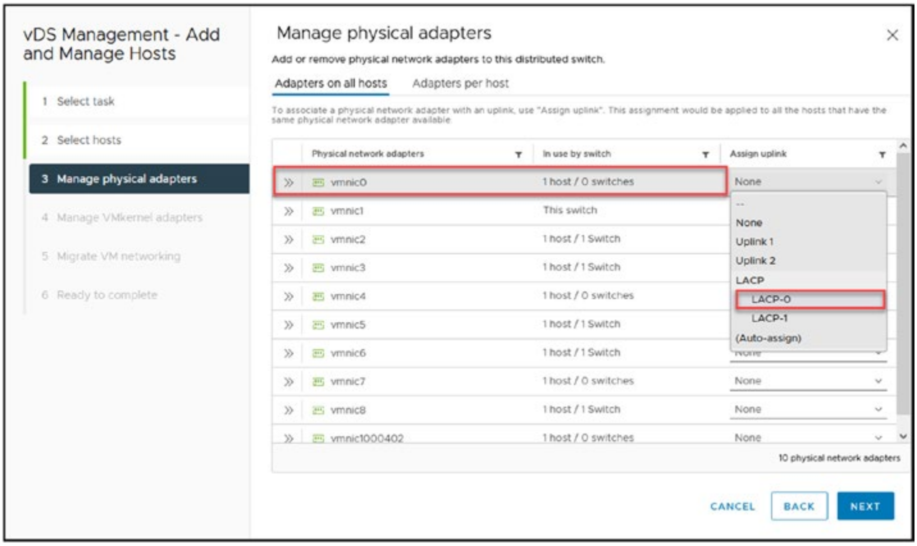

Figure 16-20. *(Migrate VMkernel LACP 6 of 11)*

Do not change anything else: click **Next** until the final wizard step and then click **Finish**.

Return to the physical switch and reenable the port 13 (`vmnic0`).

4. `interface 13 enable`

Now edit the vDS LACP, and in the *Teaming and failover* section, move the LACP group to the *Active uplinks* and move the Uplinks to the *Unused uplinks*. See Figure 16-21.

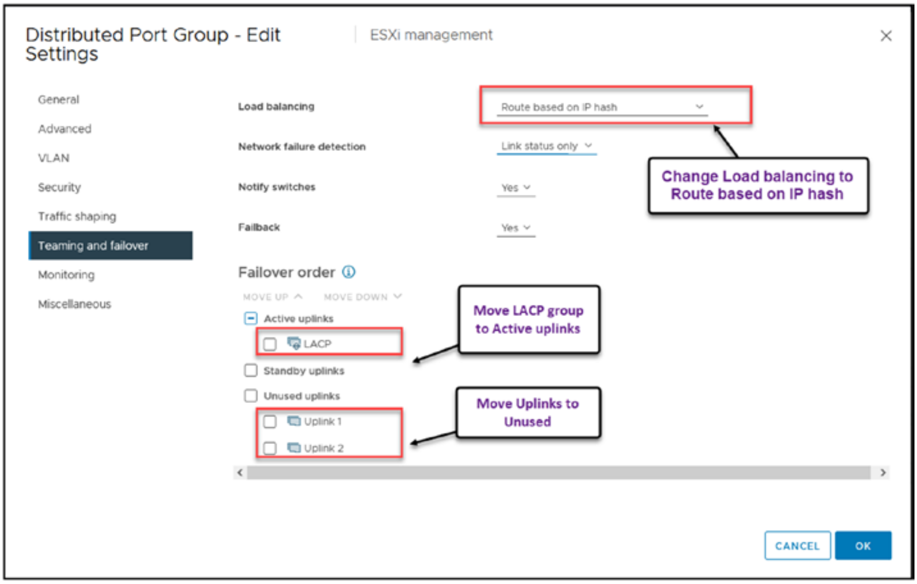

Figure 16-21. *(Migrate VMkernel LACP 7 of 11)*

Note While doing these changes, ensure the ESXi host doesn't lose its connection.

In the Figure 16-22, we show the process that we created above witht the commends in your physical switch.Moving from a Standalone port, to an LACP port, by disabling ports and enabling the ports with LACP.

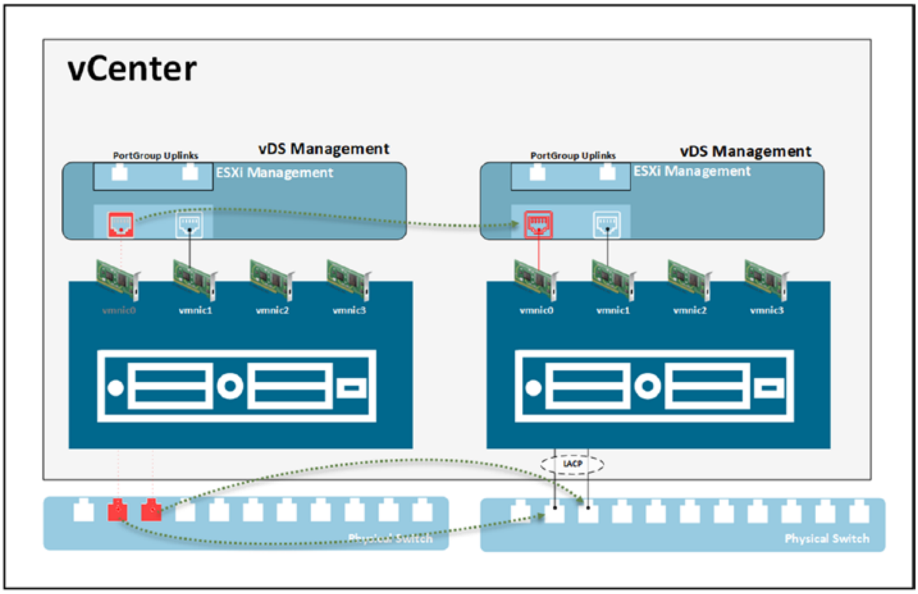

Figure 16-22. (Migrate VMkernel LACP 8 of 11)

Now let's disable the second port (`vmnic1`)

5. `interface 12 disable` (shut down port 12 where `vmnic1` is connected)

As you can see in Figure 16-23, `vmnic1` is down in vCenter.

Figure 16-23. *(Migrate VMkernel LACP 9 of 11)*

Next, following the same procedure previously described, move the `vmnic1` to the LAG link. See Figure 16-24.

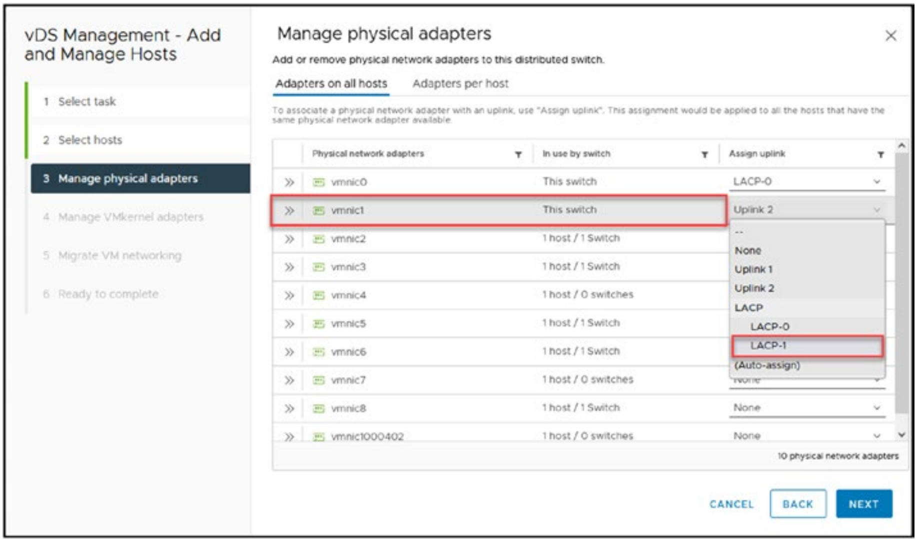

Figure 16-24. *(Migrate VMkernel LACP 10 of 11)*

Again, do not change anything else; click **Next** to proceed through the wizard steps and then click **Finish** in the last step.

Finally, add port 12 (`vmnic1`) to LACP in the physical switch.

6. `interface 12 LACP` (add port 12 or `vmnic1` to LACP/EtherChannel configuration).
   To finalize the VMkernel migration, reenable the port for the `vmnic1`.

7. `interface 12 enable`
   We have migrated our VMkernel to a vDS with LACP without losing any connection to our ESXi host.
   We did this process for ESXi 01, but you also need to do this for all ESXi hosts in the Cluster that are added to this vDS Management.
   For a final check to see if LACP is running on your ESXi hosts, run the following commands (see Figure 16-25):

   - `esxcli network vswitch vmware lacp config get`
   - `esxcli network vswitch dvs vmware lacp status get`

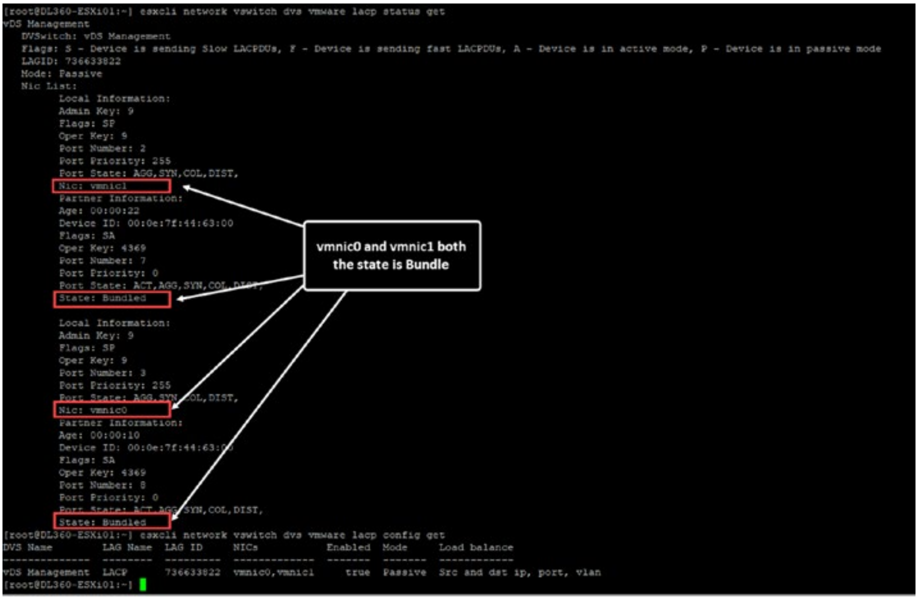

Figure 16-25. *(Migrate VMkernel LACP 11 of 11)*

16.5 Summary

This chapter explored the intricacies of configuring LACP in vSphere Distributed Switch and explained its core concepts, requirements, and the necessary steps for implementation and migration. In Chapter 16, we shift our focus to gaining a deeper understanding of advanced configurations in vSphere Distributed Switch. Our next topic is the utilization of VLANs within vDS, which plays a role in improving network segmentation and performance. Furthermore, we will also explore the process of migrating from a Standard vSwitch to vDS—a step for network advancement in VMware environments. This phase aims to enhance your expertise in managing networks by leveraging the capabilities of vDS within VMware vSphere.