Technical Report

Ethernet Storage Best Practices for Clustered Data ONTAP Configurations

Mike Worthen, NetApp
February 2014 | TR-4182

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

This technical report describes the implementation of clustered Data ONTAP® network configurations. It provides common clustered Data ONTAP network deployment scenarios and networking best practice recommendations as they pertain to a clustered Data ONTAP environment. A thorough understanding of the networking components of a clustered Data ONTAP environment is vital to successful implementations.

This report should be used as a reference guide ONLY. It is NOT a replacement for product documentation or specific clustered Data ONTAP technical reports or end-to-end clustered Data ONTAP operational recommendations or cluster planning guides. In addition, any solution specific guides will supersede the information contained in this document, double-check all related references.

2 Overview

There are different types of cluster configurations that can be implemented, all of which utilize various networking concepts and features.

- Single-node cluster (Figure 1): In a single-node cluster, settings such as Ethernet flow control and Transmission Control Protocol (TCP) options still need to be properly configured. Also, if the configuration will be upgraded from a single node, the best practice recommendation is to install the necessary components for the upgrade and expansion during the initial implementation. For example, install NICs for cluster interconnectivity. This could save a reboot or two when the need to move to a highly available solution presents itself.

Important Note
In a single node configuration there aren't other nodes in the cluster to replicate important cluster and node configuration information to. As such, accommodations will need to be made to ensure there is always a backup copy located off of the single node itself. For more information see the section "Backing up and restoring cluster configurations" in the System Administration Guide for Cluster Administrators on the support.netapp.com site.

- Two-node switchless cluster (Figure 2): In a two-node switchless cluster configuration, settings such as flow control, TCP options, and MTU size also need to be configured appropriately. If there is a potential this configuration will be moved to a multinode switched configuration at a later time, keep in mind this can be done nondisruptively. Also, the lack of a switch in this configuration does not change the port requirements per platform. Refer to the Hardware Universe for port requirements and follow the same guidance whether it's a switched or switchless configuration.
- Multinode switched cluster (Figure 3): In a multinode switched cluster, the customer will gain all the benefits that clustered Data ONTAP offers: the nondisruptive capabilities, the highly available capabilities, and the performance capabilities. This is the more complex solution of the three to configure, but it is also the one with the greatest return.

Figure 1) Single node cluster.

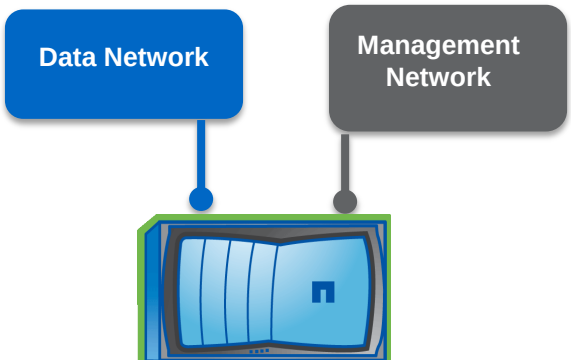

Figure 2) Two-node switchless cluster.

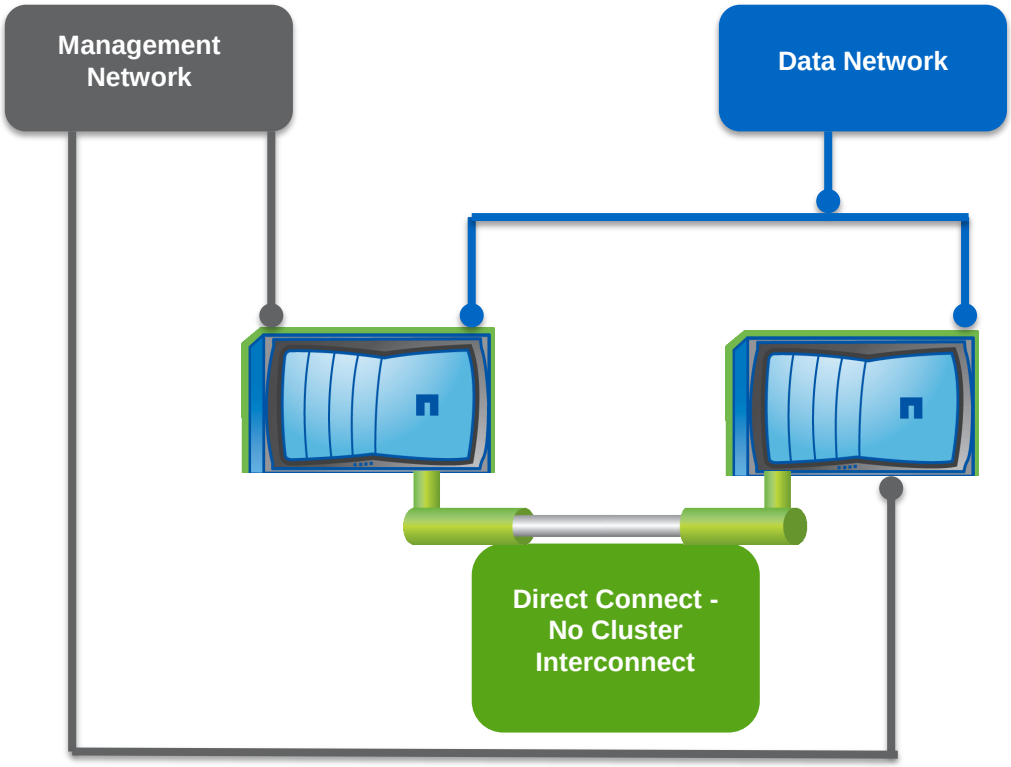

Figure 3) Multinode switched cluster

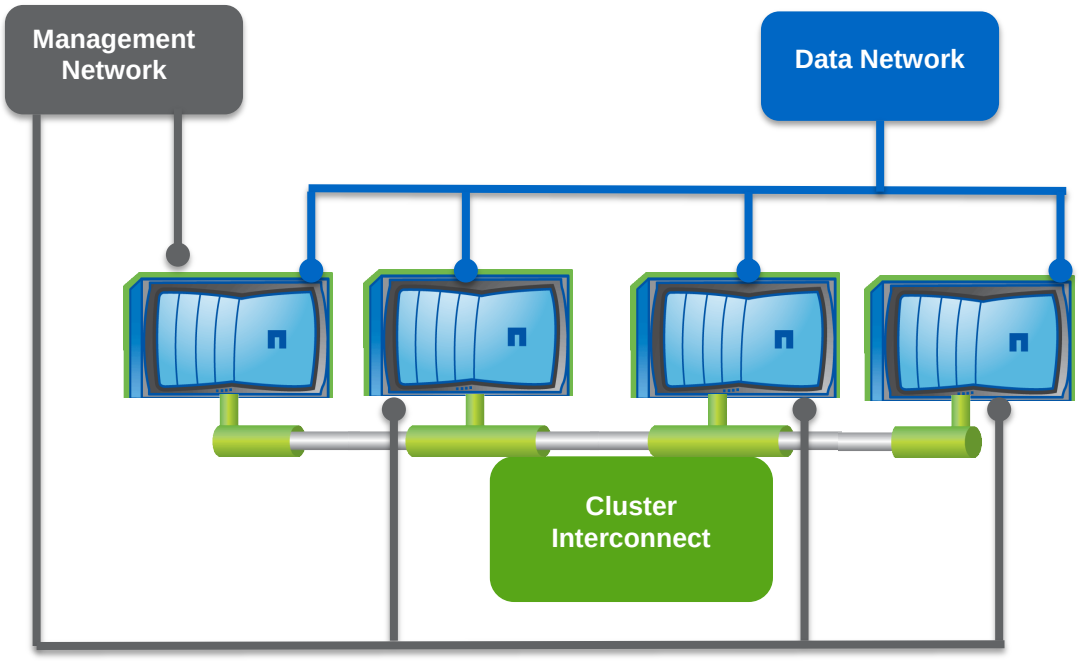

- Visit the Cluster Management and Interconnect Switches page for additional information regarding the cluster interconnect requirements.

2.1 Setting Up the Cluster

Several software and hardware prerequisites are required for initially setting up and configuring a clustered Data ONTAP implementation. We will begin with discussing the port types available and then go through the logical interfaces that will use the different port types to send and receive input/output requests to and from the cluster.

- Ports
  - Physical ports: Will be used for various functions and can have different types of configurations.
  - Virtual ports: VLANs and interface groups (IFGRPs) make up the options for virtual ports.
- Logical Interfaces (LIFs)
  - Data
  - Cluster
  - Cluster management
  - Node management
  - Intercluster

2.2 Ports

There are different types of ports in clustered Data ONTAP: physical and virtual. The different port types are used throughout the cluster in different configurations but are the building blocks that logical interfaces (LIFs, which are described in more detail in section 2.3) will use to allow the sending and receiving of data.

2.2.1 Physical Ports

Physical ports can be used individually or in combination when configuring virtual ports. If you use several physical ports together for an interface group it is important to remember to configure all relevant port settings in the same way (including MTU size and flow control). However, these settings will also be relevant if you exclusively use physical ports in a configuration in which only failover groups are in play; consistency in the settings and consistency in the configurations is needed.

2.2.2 VLANs

A Virtual Local Area Network (VLAN) subdivides a physical network into distinct broadcast domains. As a result, traffic is completely isolated between VLANs unless a router (Layer 3) is used to connect the networks. Complete isolation is one of the primary reasons to use VLANs from a security perspective. In clustered Data ONTAP VLANs subdivide a physical port into several separate virtual ports allowing for one of the key components of our secure multi-tenant messaging.

2.2.3 Interface Groups

IFGRPs can be configured to add an additional layer of redundancy and functionality to a clustered Data ONTAP environment. They can also be used in conjunction with a failover group, which would help protect against Layer 2 and Layer 3 Ethernet failures. Below are characteristics listed for each type as well as situations in which it would be advisable to use each.

Best Practice - IFGRPs

- The best practice recommendation when creating an interface group is: if it is physically possible (if there are enough slots, NICs, and so on), is to create the interface group using ports from different NICs but verify that they are the same model/chipset and have the same speed, functionality, and so on. This is critical in maintaining consistency in the ifgrp in the event of a port failure. By maintaining consistency with port aggregation and by spreading the ifgrp over NICs in different slots you decrease the chances of a slot being responsible for taking offline all the ports in an ifgrp. **For example, if performance is a top priority in the environment, do not create an ifgrp with 10GbE and 1GbE interfaces in the same ifgrp, regardless of the model of the interface.**
- The network interfaces and the switch ports that are members of an dynamic multimode (LACP) ifgrp **MUST** be set to use the same speed, duplex, and flow control settings. However, it is a best practice recommendation to follow the same practices if creating any of the different ifgrps below.

Note - IFGRPs

There are two methods to achieve path redundancy if using iSCSI in clustered Data ONTAP: by using ifgrps to aggregate more than one physical port in partnership with an LACP-enabled switch, or by configuring hosts to use MPIO over multiple distinct physical links.

Both of these methods allow a storage controller and host to use aggregated bandwidth and both can survive the failure of one of the paths from host to storage controller.

However, MPIO is already a requirement for using block storage with clustered Data ONTAP, and using MPIO has the further advantage of no additional switch configuration or port trunking configuration being required. Also, using an ifgrp for path management when using iSCSI is not supported although using an ifgrp as a port for an iSCSI LIF is supported.

Different Types of Interface Groups

2.2.3.1 Single Mode

A single mode interface group is an active-passive configuration (one port will sit idly waiting for the active port to fail) and it cannot aggregate bandwidth. The switch configuration on a single mode ifgrp is not continuously verified (speed, duplex) other than basic connectivity. As an effect of this, if the VLAN connectivity encompassing both ports is severed (e.g. that VLAN is not configured on an alternate switch-to-switch trunk line becoming active), unexpected behavior may result; in particular during reboots/cluster failover (either the "wrong" port becomes active or no port becomes active). Trying to enable the proper port will only work if the node port was fully operational and in the proper state to begin with. Due to its limited capabilities, as a best practice recommendation NetApp advises not using this type. To achieve the same level of redundancy, you could use failover groups (see section 3.1) or one of the two ifgrps listed below.

2.2.3.2 Static Multimode

A static multimode interface group might be used if you want to utilize all the ports in the group to simultaneously service connections. It does differ from the type of aggregation that happens in a dynamic multimode interface group (described in section 2.2.3.3) in that no negotiation or auto detection happens within the group in regard to the ports. A port will be sending data when the node detects a link regardless of the state of the connecting port on the switch side.

Note

For all data ports, and in particular single mode and static multimode interface groups, it is recommended to disable spanning tree on the adjacent switch ports, or set all the port-specific timers of the spanning tree to the minimum attainable. This reduces any loss of service time due to the network not properly accepting and forwarding storage data packets.

2.2.3.3 Dynamic Multimode (LACP)

A dynamic multimode interface group might be used to aggregate bandwidth of more than one port. LACP monitors the ports on an ongoing basis to determine the aggregation capability of the various ports and continuously provides the maximum level of aggregation capability achievable between a given pair of devices.

However, all the interfaces in the group will be active, will share the same MAC address, and will handle load balancing outbound traffic. But this does not mean a single host will achieve larger bandwidth, exceeding the capabilities of any of the constituent connections. For example, adding four 10GbE ports to a dynamic multimode interface group will not result in one 40GbE link for one host. This is due to the way the aggregation of the ports in the interface group is handled by both the switch and the node. A recommended best practice is to use this type of interface group so that you are able to take advantage of all the performance and resiliency functionality the interface group algorithm has to offer.

2.2.3.4 Load Balancing for Multimode IFGRPs

Four distinct load-balancing modes are available.

- Port: Use this distribution method for best load balancing results. However, it lends itself less well to troubleshooting, since the TCP/UDP port of a packet is also used to determine the physical port used to send a particular packet. It has also been reported that switches operating in particular modes (mapping MAC/IP/Port) may exhibit lower than expected performance in this mode.
- MAC: Only useful when the IFGRP shares the same VLAN with the clients having access to the storage. If any storage traffic traverses a router or firewall, do not use this type of load balancing as the MAC address for every outgoing IP frame will be the MAC address of the router which will result in only one interface in the IFGRP being used.
- IP: Second-best load distribution method, since the IP addresses of both sender (LIF) and client are used to deterministically select the particular physical link that a packet traverses. Although deterministic in the selection of a port, the balancing is performed using an advanced hash function. This has been found to work under a wide variety of circumstances, but particular selections of IP addresses might still lead to unequal load distribution.
- Sequential: Nondeterministic load balancing. Under specific circumstances, this type of load balancing can cause performance issues due to high overhead to the switch (potential constant remapping of MAC/IP/Port) or out-of-order delivery of individual packets destined for a client (because of this it's not supported by the IEEE LACP specification. Due to the potential of less than favorable load balancing and the potential for out or order delivery it is not recommended to use this type.

Note

Remember, the load balancing in an IFGRP happens on the outbound traffic, not inbound. So, when a response is being sent back to a requester, the load balancing algorithm comes into play to determine which "path" is the optimal to send the response back on.

2.3 Logical Interfaces

Logical interfaces (LIFs) are created as an abstraction on top of the physical (physical ports) or virtual interface (VLANs or IFGRPs) layer. IP-based LIFs for NAS or iSCSI are assigned IP addresses, and FC-based LIFs are assigned WWPNs.

2.3.1 Data

The data LIF is used for data traffic (NFS, CIFS, FC, iSCSI). Although you use a data LIF for either NAS or SAN traffic, you cannot use the same data LIF for both. The NAS data LIF can fail over or migrate to other data ports throughout the cluster if configured to do so via a failover group. Also, as a very important distinction, NAS data LIFs will migrate; SAN data LIFs **(including iSCSI)** do not migrate but will instead use ALUA and MPIO processes on the initiators to handle path failures.

Note: Make certain that failover groups and the LIFs assigned to them are configured correctly, meaning that you should configure the failover groups to use ports in the same subnet/VLAN and verify that LIFs are assigned to the correct failover groups. If ports from different subnets are used in the same failover group or if LIFs aren't assigned to the correct failover groups and a failover occurs, it will result in loss of network connectivity that will result in the loss of data availability.

2.3.2 Cluster

The cluster LIF can only be configured on 10GbE or 1GbE (1GbE can be used on the FAS2040 and FAS2220 platforms) ports of type cluster and can only failover to cluster ports on the same node.

It is used for operations such as:

- Volume moves
- To synchronize cluster/node configuration and metadata among the nodes in the cluster (this is a very important communication aspect since it keeps nodes in the cluster in quorum)
- Access data on multiple nodes in the cluster.
- Intracluster data replication

Visit the Cluster Management and Interconnect Switches page for additional information.

2.3.3 Cluster Management

The cluster management LIF is used to manage the cluster. It can only reside on and failover to data ports but can fail over to any data port on any of the nodes in the cluster. As such, ensure this LIF is assigned to a correctly configured failover group.

2.3.4 Node Management

A node management LIF exists on every node in the cluster and is used for processes such as AutoSupport, SNMP, NTP, DNS and other node specific management traffic (for a complete list of process refer to the "Clustered Data ONTAP 8.2 Network Management Guide"). It can also be used to manage the node directly in the cluster for system maintenance purposes. It can fail over to other data or node management ports on the same node only. As such, ensure this LIF is assigned to a correctly configured failover group.

2.3.5 Intercluster

The intercluster LIF is used for peering from cluster to cluster. These are node specific; they can only use or failover to intercluster or data ports on the same node. At least one intercluster LIF is required per node for replication between clusters. However, for the sake of redundancy the best practice recommendation

is to configure a failover group or host the LIF on an IFGRP port. Maintain consistent settings between the intercluster LIFs (same MTUs, flow control, tcp options, and so on).

Table 1) Additional information regarding LIFs

LIF Type	Function	Minimum Required	Maximum Allowed
Node management	Used for system maintenance of a specific node, SNMP, NTP, and ASUP™ tool	1 per node	1 per port/subnet
Cluster management	Management interface for the entire cluster	1 per cluster	N/A
Cluster	Used for intracluster traffic	2 per node	2 - 4 per node depending on the platform. Check Hardware Universe for specifications
Data	Associated with a Storage Virtual Machine and used for data protocols and protocol services (NIS, LDAP, AD, WINS, DNS)	1 per Storage Virtual Machine	128 per node in HA configuration 256 per node in non-HA
Intercluster	Used for intercluster communication, such as setting up cluster peers and SnapMirror® traffic	1 per node if cluster peering is enabled	N/A

3 Storage Virtual Machine Networking

3.1 Failover Groups

In the event of a failure, LIFs need to be migrated in a coordinated manner. When a LIF is created it is assigned to a system-defined failover group by default. However, the behavior of the system-defined failover group may not be sufficient for every different type of environment.

A failover group contains a set of network ports on one or more nodes. A failover group can have cluster management, node management, intercluster, and NAS data LIFs assigned to it. As mentioned previously in this document, SAN LIFs don't failover so they don't utilize failover groups. The network ports that are present in the failover group define the failover targets for the LIF. The best practice recommendation for LIFs capable of utilizing failover groups should always be to assign those LIFs to an appropriate failover group. Also, make double-checking a best practice by making certain all ports in the failover group are part of the same subnet. Failure to determine that ports are in the same subnet and failure to assign LIFs to appropriate failover groups will result in loss of connectivity to data.

There are currently three different types of failover groups. Below, each is described with details on when you might want to use each.

- User Defined: This is the type of failover group to use from a recommended best practice perspective. You can configure it to provide all configuration requirements of any environment due to its very flexible functionality. Functionality such as:
  - If multiple subnets exist—the system-defined or clusterwide groups will not keep LIFs on their own subnets.

- You want to logically group a certain type of interface (for example, 10GbE-based LIFs only failover to other 10GbE ports).
- For LIFs configured on top of VLAN ports and you want to be certain the LIFs move to port(s) that can communicate with the other devices that are members of that same VLAN configuration.

- Cluster-Wide: This type is automatically created during setup and it cannot be modified nor does it need to be. It includes all data ports in the cluster by default and is the default failover group for cluster management LIFs. As long as the network is flat (that is, there are no subnets), it will successfully control failover of the LIFs that are assigned to it.
- System Defined: This type is also automatically created during setup, cannot be modified, and will control failing over all data LIFs by default. As with the cluster-wide type of group, system defined is useful as long as the network is flat.

Note: **(see Figure 4): System-defined groups will only contain ports from a maximum of two nodes: ports from one node of an HA pair combined with ports from a node of a different HA pair. This decreases the chance of complete loss of connectivity if some type of network issues causes one node in an HA pair to fail followed by the second node in the same HA pair.**

Figure 4) Default behavior for the system-defined failover group beginning with cDOT 8.2.

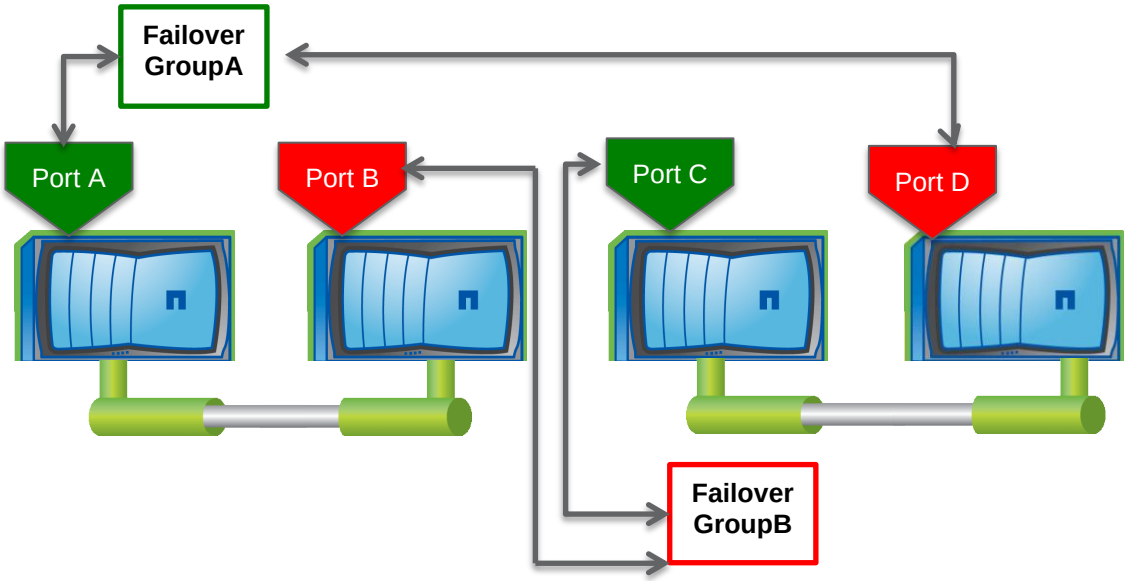

FAILOVER GROUP - EXAMPLES FOR DOUBLE-CHECKING CONFIGURATIONS

Pre clustered Data ONTAP 8.2: Listed below is an example of an incorrectly configured ifgrp; with explanations of why it is incorrectly configured and the steps to remedy it.

1. ontaptme-rtp::> network interface failover-groups create -failover-group scon_test1 -node ontaptme-rtp-01 -port e0a

   (network interface failover-groups create)

2. ontaptme-rtp::> network interface failover-groups create -failover-group scon_test1

   -node ontaptme-rtp-01 -port e0b

   (network interface failover-groups create)

3. ontaptme-rtp::> network interface modify -vserver scon_test1 -lif scon_test_lif1

   -failover-group scon_test1

The failover group is not configured correctly as it is assigned to use the system-defined failover group (look at the "Use Failover Group" field).

4. ontaptme-rtp::> network interface show -vserver scon_test1 -lif scon_test_lif1

   (network interface show)

   Abbreviated output……

   Abbreviated output……

   Use Failover Group: system-defined

   Abbreviated output……

   Failover Group Name: scon_test1

The following steps can be executed to properly configure the failover group.

5. ontaptme-rtp::> network interface modify -vserver scon_test1 -lif scon_test_lif1 -use-failover-group enabled -failover-group scon_test1

   (network interface modify)

6. ontaptme-rtp::> network interface show -vserver scon_test1 -lif scon_test_lif1
   (network interface show)

   Abbreviated output……

   Failover Policy: nextavail

   Abbreviated output……

   Use Failover Group: enabled

   Abbreviated output……

   Failover Group Name: scon_test1

To identify any NAS data LIFs that don't have failover groups configured.

7. ontaptme-rtp::> network interface show -failover-group system-defined -data-protocol nfs|cifs

Note
Beginning with clustered Data ONTAP 8.2 and later, the "Use Failover Group" field no longer exists so the two fields to be concerned with will be "Failover Policy" and "Failover Group Name". There are different LIF/failover group/VLAN/IFGRP configurations possible in a clustered Data ONTAP environment. The best practice recommendation is to utilize the configuration in Figure 5. This configuration takes advantage of the cluster-wide failover capabilities of failover groups, the port aggregation functionality of interface groups, and the security aspects of VLANs. For more examples refer to the "Clustered Data ONTAP 8.2 Network Management Guide".

Figure 5) Best practice LIF/failover/VLAN/IFGRP configuration.

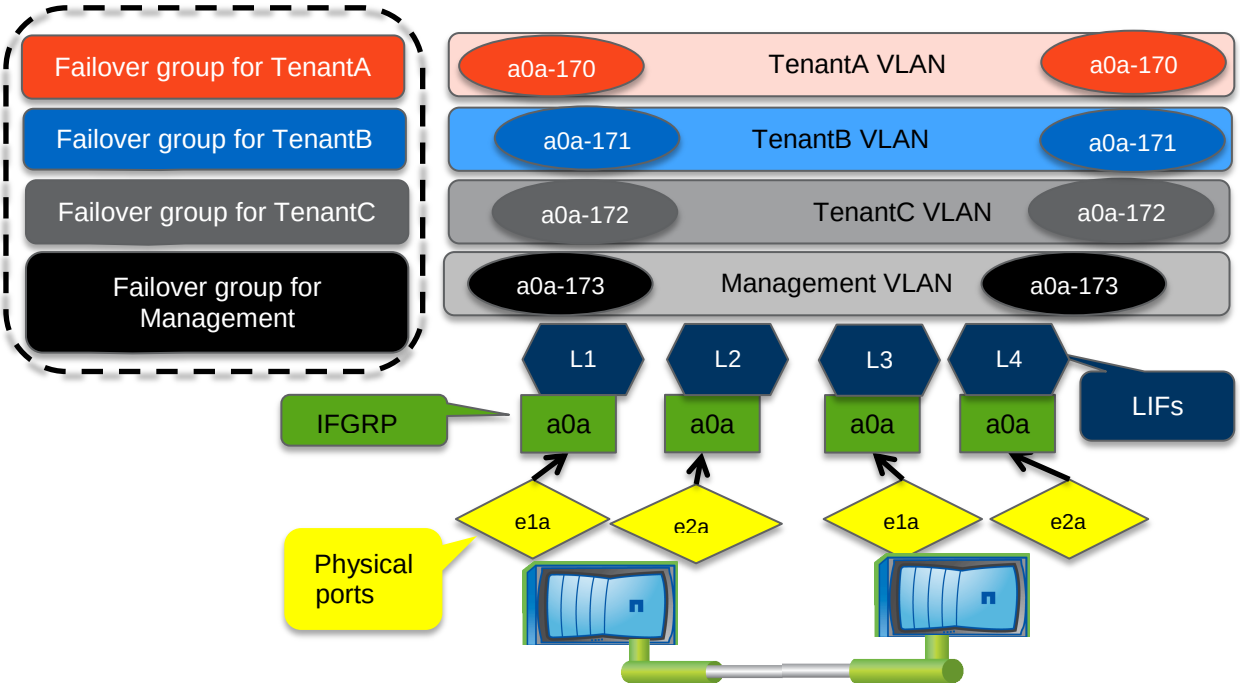

3.2 Design Considerations

3.2.1 DNS Load Balancing

DNS load balancing in its different forms takes advantage of different types of algorithms to determine the optimal LIFs to return resolution requests. Two different types of DNS Load Balancing are currently supported with clustered Data ONTAP; Zoning Based and Round Robin. Using Zoning Based (or commonly referred to as On-Box), you can create a DNS load-balancing zone on the Storage Virtual Machine that returns a lower loaded LIF based on the network traffic and the availability of the port resources (CPU usage, throughput, open connections, and so on). Using Round Robin based DNS Load Balancing (or commonly referred to as Off-Box) you will be using a round robin based algorithm to determine what resources respond to resolution requests. Also, for either option, only new NAS mounts will be affected. Any existing mounts/shares would need to be remounted to take advantage of the load balancing benefits. For details regarding setting up either of the options below in an NFS environment using different types of authentication please see TR-4067 Secure Unified Authentication with NetApp Storage Systems beginning on page 23.

3.2.1.1 Zoning Based/On-Box

In this configuration there may be a measurable increase in performance with certain workloads. However, it is a bit more complex to set up and manage. Although it is a bit more complex to manage, after the delegation zone is added to the site wide DNS server and the data LIF IP addresses are added to it (**NOTE: you are adding data LIF IP's belonging to each individual Storage Virtual Machine, not cluster mgmt., cluster, node mgmt.)**, the Storage Administrator could manage it from then on; which could decrease dependency on the Network team. It currently works with the following protocols – NFSv3, NFSv4, NFSv4.1, and SMB2.0

You are delegating queries to the DNS server (NAMED) running on each node inside the cluster, which then passes it to the individual Storage Virtual Machines; at that point clustered Data ONTAP calculates the load of each data LIF automatically and then passes the query to the appropriately balanced LIF. In the example below, the NFS client would query the site wide DNS server for resolution of the share it wants to mount from Storage Virtual Machine (SVM). The site wide DNS server would pass the query via the delegation zone CLUS.NETAPP.COM to NAMED running on the cluster, which would resolve access, based on the appropriately balanced data LIF in the SVM. Keep in mind, if an environment will have many SVM's, you will have to account for each of the data LIF's of each of the SVM's being added to the delegation zone on the site wide DNS server.

Figure 6) DNS Load Balancing – Zoning Based

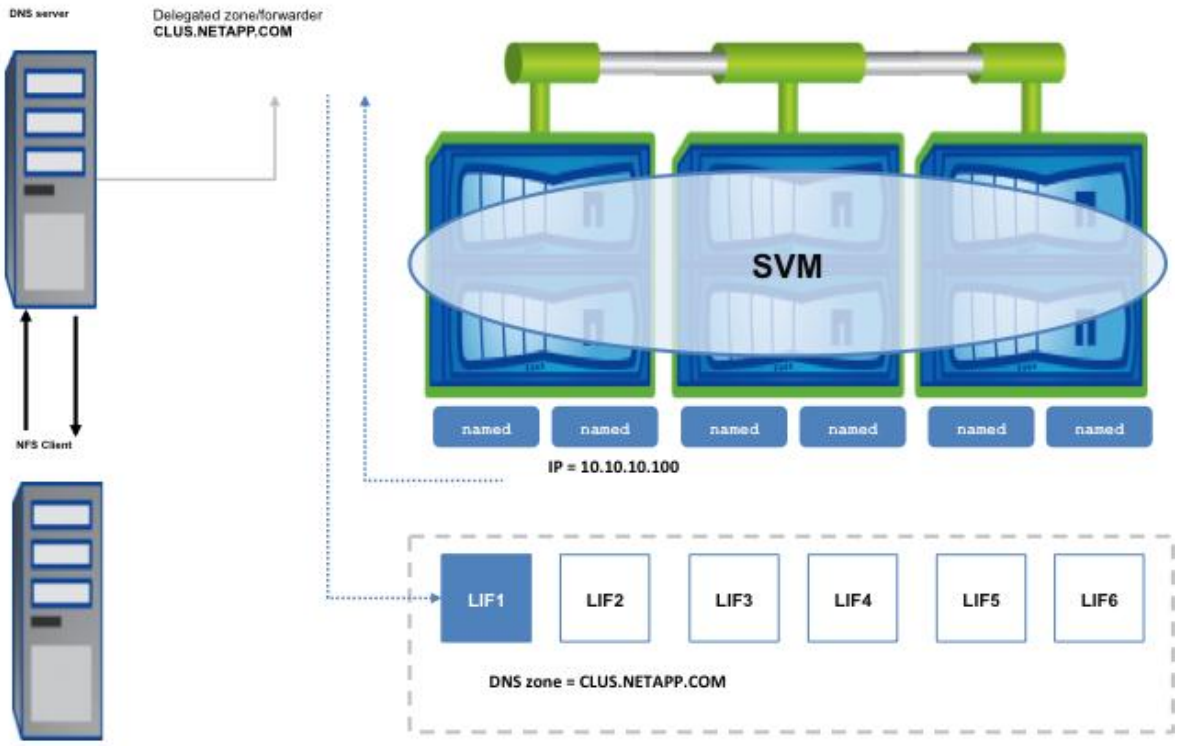

3.2.1.2 Round Robin/Off-Box

In this configuration, the setup and management is a bit easier compared to the solution in section 3.2.1.1 above however, issues such as client and DNS hierarchy caching could cause resource bottlenecks. Also, the Round Robin algorithm may not be as effective as the algorithm used by clustered Data ONTAP. To the contrary though, it may be a good solution if the environment is uniformly configured (similar type servers, evenly distributed connections). It might also be a better solution if for example; the same team manages the DNS and Storage environments.

In the example below each Data LIF in each Storage Virtual Machine residing in the cluster has a DNS "A" record created with the same name. The NFS client makes a request for name resolution to the Site Wide DNS server. Site Wide DNS server resolves the request to an IP address via Round Robin algorithm. Site Wide DNS server responds to client with chosen IP address. Keep in mind, similarly to the solution in section 3.2.1.1, if an environment will have many Storage Virtual Machine's, you will have to account for each of the data LIF's of each of the Storage Virtual Machine's being added to the site wide DNS server.

Figure 7) DNS Load Balancing – Round Robin

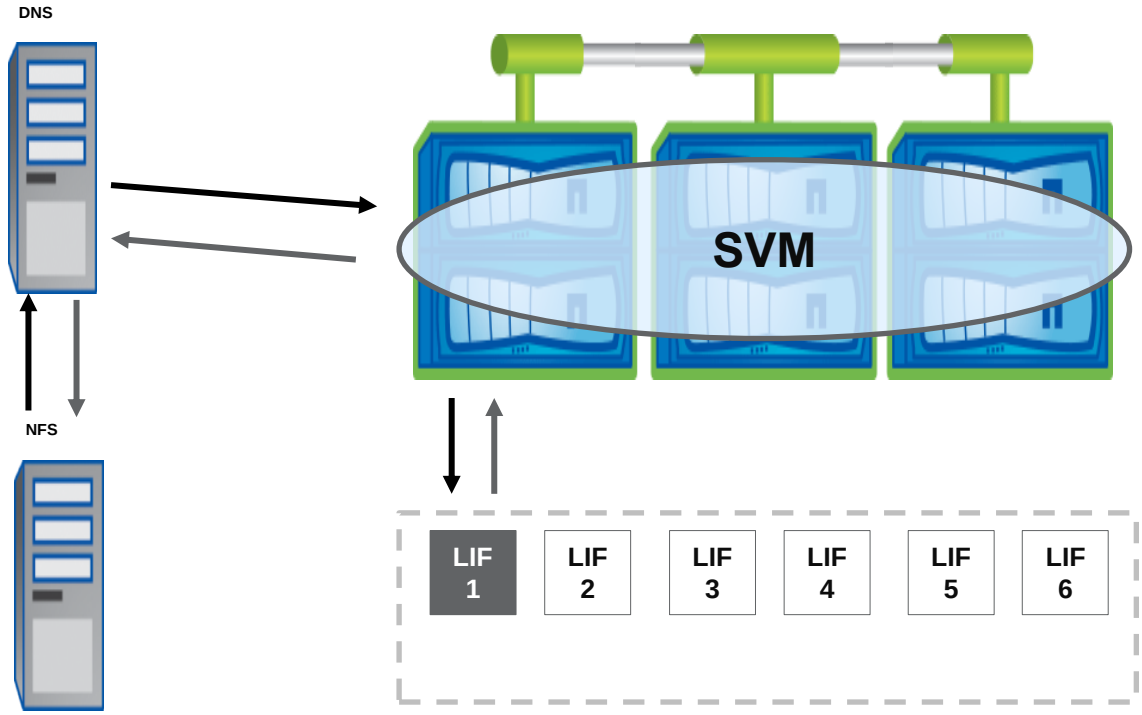

3.2.2 Automatic LIF Rebalancing

Automatic LIF Rebalancing can be used to allow clustered Data ONTAP to dynamically migrate LIFs residing on over utilized ports to ports with lower utilization.

- Only works with NFSv3. If the connection is moved, the I/O request will resume once the connection is reestablished.
- If any other NAS traffic exists with NFSv3 on a LIF, it negates auto rebalancing.
- Follows the failover group rules. Make certain the failover groups are configured correctly so failovers occur as expected. **Specifically, best practice requirement, make certain the ports that make up the failover groups are part of the same subnet. This way the LIFs will not lose connectivity if they are moved around or rebalanced within the failover group.**
- Calculates a weighted, balanced distribution of load across the ports. Is automatically assigned to the LIFs based on statistics on the current node and port resources.

4 Performance Considerations

4.1 Ethernet Flow Control

Ethernet flow control from a NetApp perspective can be thought of as the mechanics that allow the receiving party of a connection to control the rate of the sending party. With that said, due to limitations with buffer designs on switches in the industry today, the best practice recommendation is to disable flow control throughout the network (including host ports, switch ports, and all node ports). Allow the upper layer protocols to handle congestion control as needed.

Note: When creating or configuring an interface on a NetApp® controller, the default for flow control settings is set to be "on" for both send and receive. You will need to change the settings "send off" and "receive off" using the `network port modify` command.

Note: With the first generation network interface (NIC) cards it is not possible to disable flow control. With these types of cards NetApp recommends leaving flow control enabled on both the NIC port and the switch port the NIC is connected to.

4.2 Jumbo Frames

Jumbo frames are Ethernet frames with more than 1,500 bytes maximum transmission unit (MTU) of payload. In a clustered Data ONTAP environment, ports with a role type of cluster **must** be set to an MTU size of 9,000. The cluster may operate but will do so at a suboptimal level if the cluster ports are set to a 1,500 MTU size. Also, keep in mind that if the MTU is changed while the cluster port is active, the NIC will reset and connections will be dropped; this could have a very detrimental effect on the cluster.

Note: When a port is configured as type cluster, during initial setup the clustered Data ONTAP setup wizard will set it to 9,000 MTUs.

4.3 Topology

- To avoid undesired performance degradation during a physical port failover, do not configure ifgrp's comprised of different link speeds. For example, avoid creating an ifgrp and mixing a 10GbE and a 1GbE interfaces to it. Instead create the interface groups with like interfaces (refer to section 2.2.3 and figure 8 for additional information).
- When adding additional interfaces to an existing ifgrp, review the entire configuration in the environment from end-to-end. For example, if adding two additional 10GbE interfaces to an existing ifgrp containing two 10GbE interfaces (four interfaces total or 40Gigabits aggregated); review the port group configuration on the upstream switches and network infrastructure so that there will be adequate bandwidth to prevent a possible bottleneck at the port group level on the switches (refer to figure 9).
- Verify that all interfaces and network ports participating in the environment are correctly negotiating (auto, half, full, etc) as expected.
- Use the following command on the nodes in the cluster to check the settings for the ports for each node:
  - ***"network port show"***

    For the switches in the environment use commands such as the one below to check relevant port settings on the switches (this is an example only, refer to the documentation for your specific switch for the exact syntax):
  - ***"show interface ethernet 1/1"***
- Do not use interface e0M (or any other port designated for node management) for any type of protocol traffic other than node management. This interface should be exclusively for node management and administrative access.

Figure 8) DO NOT mix port speeds in interface groups.

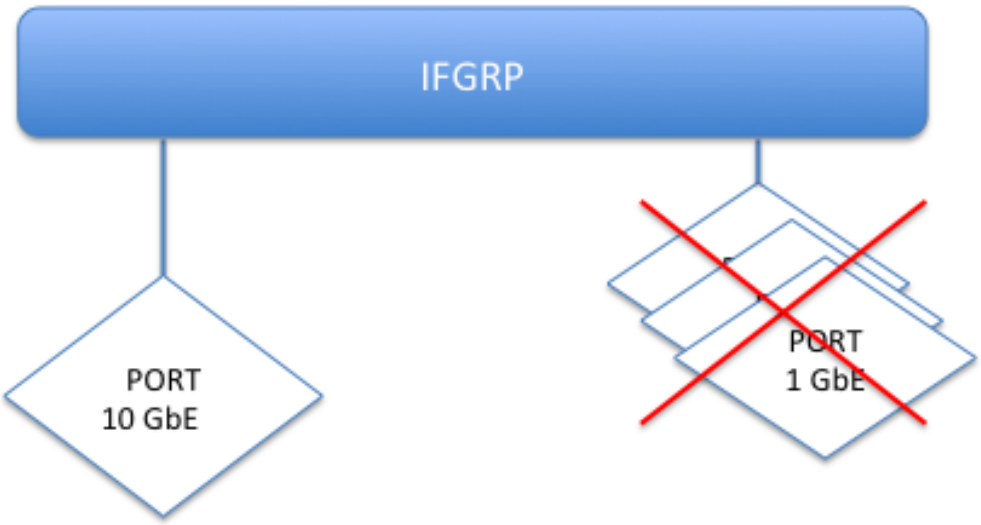

Figure 9) Review topology end to end if adding additional resources at any one point.

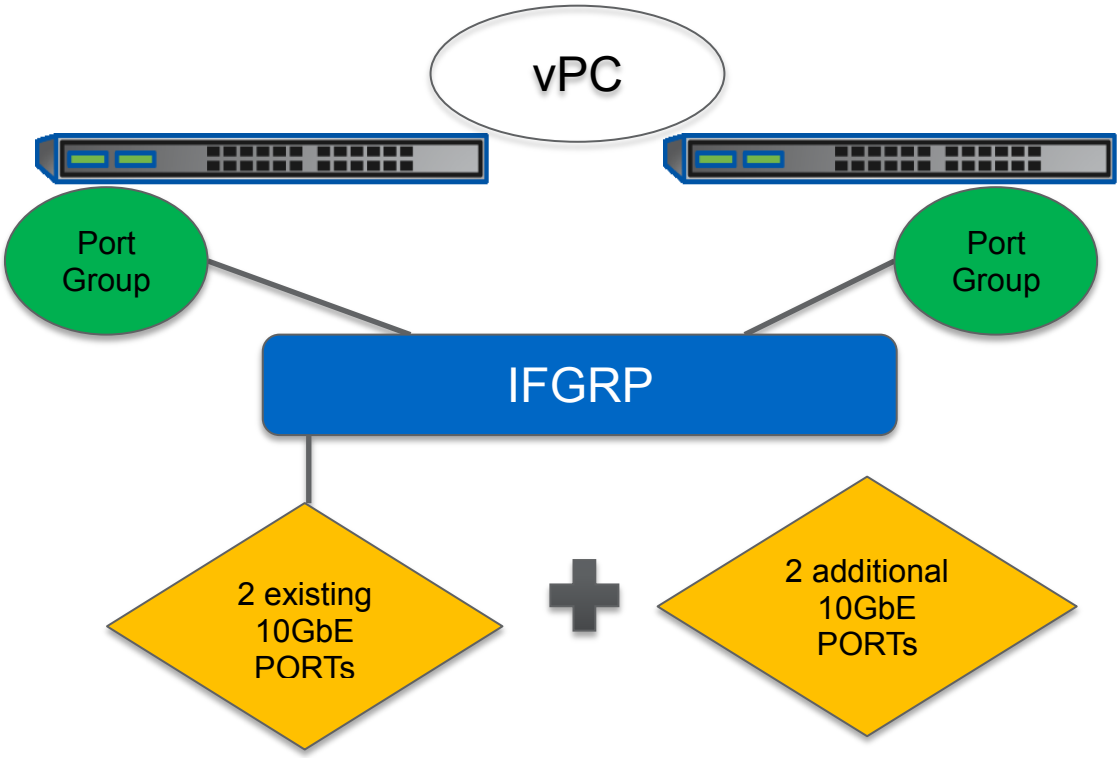

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

www.netapp.com

© 2014 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, ASUP, and SnapMirror are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4182-0214