

COHESITY

Version 1.2

February 2024

Optimal Network Designs with Cohesity

Table of Contents

Introduction.....	3
Key Concepts	4
Five Common Configurations.....	6
Simple Topology	6
Standard Topology.....	7
Layered Topology	8
Cisco FEX Topology	9
Layered Bond Mode 4 (LACP) Topology	10
Terminology.....	11
Conclusion.....	12
Your Feedback.....	13
About the Authors.....	13
Document Version History.....	13

Figures

Figure 1: Cohesity on a Simple Topology.....	6
Figure 2: Cohesity on a Standard Topology	7
Figure 3: Cohesity on a Layered Topology.....	8
Figure 4: Cohesity on a Cisco FEX Topology.....	9
Figure 5: Cohesity on a Layered Bond Mode 4 (LACP) Topology.....	10

Tables

Table 1: Key Concepts.....	4
Table 2: Terminology.....	11

Introduction

Cohesity provides a hyper-converged secondary storage solution that converges backups, file and object services, test/dev, and analytics into a single, web-scale platform. The Cohesity platform can span both private and public clouds and enables replication and archival across enterprise LANs and WANs. As with any hyper-converged platform, Cohesity achieves highly efficient storage by integrating network, storage, and computing into a single data platform. However, integrating into a production network can be very involved and complex, requiring changes to routers, switches, and configurations.

A seamless integration into such a complex environment requires careful consideration for networking plans, to ensure the performance and availability of the Cohesity cluster. In this paper, we explore the most common network configurations and present the best practices to design an optimal Cohesity cluster.

This document covers:

- Key Concepts
- Five Common Configurations
- Terminology

Key Concepts

Before we explore the different configurations, there are several core concepts we should review in the Cohesity context:

Table 1: Key Concepts

CONCEPT	COHESITY CONTEXT
Ethernet	A Cohesity cluster relies on a high-bandwidth, high-throughput Ethernet network that is resilient against bursts, congestion, and failures from network ports, switches, and the corresponding links.
Connectivity	The ideal connectivity provides a reliable, single hop between nodes without introducing latencies or bottlenecks due to oversubscription (for example, when multiple links fan-in into a single link).
Network	All nodes of a Cohesity cluster need to be part of a single, layer-2 network. The IP addresses need to belong to a single subnet.
Bond Mode	Bonding is the creation of a single bonded interface by combining two or more Ethernet interfaces for high availability or performance reasons. Bond Mode 1 is configured by default on the Cohesity nodes.
IP Address Allocation	<p>Configuration of a Cohesity cluster requires careful consideration of IP address allocations. IP addresses are required for each node, and optionally, an IP address is required for VIP. There must be at least one VIP for the entire cluster and one IPMI IP per node.</p> <p>The best practice calls for a VIP corresponding to each Cohesity node.</p> <p>Cohesity does not support DHCP and requires statically configured IP addresses at Cluster creation and Node addition.</p>
Nodes	<p>Cohesity clusters are deployed as a minimum of 3 nodes.</p> <p>In the current generation of Cohesity hardware appliances, like the C2000 and the C3000 series, there are two 10GbE interfaces available on each node for networking. One 10Gb interface should be connected to switch A and another 10Gb interface to switch B.</p>

East-West Traffic	In a Cohesity cluster, node-to-node traffic involves traffic traversing between the Cohesity nodes. This traffic includes node rebalancing, auto-healing operations, metadata replication, and replication for the purpose of data redundancy — due either to Replication Factors (RF) or Erasure Coding (EC).
North-South Traffic	Traffic in a Cohesity cluster also involves, for example, the actual backup, file services, or replication traffic from outside the cluster (North-South traffic).
TCP/IP Connections	TCP/IP connections are established between all nodes to facilitate replication and metadata transfers. TCP/IP provides protocol-level support in case of switch overloads (due to oversubscription) and retransmissions (due to packet drops).

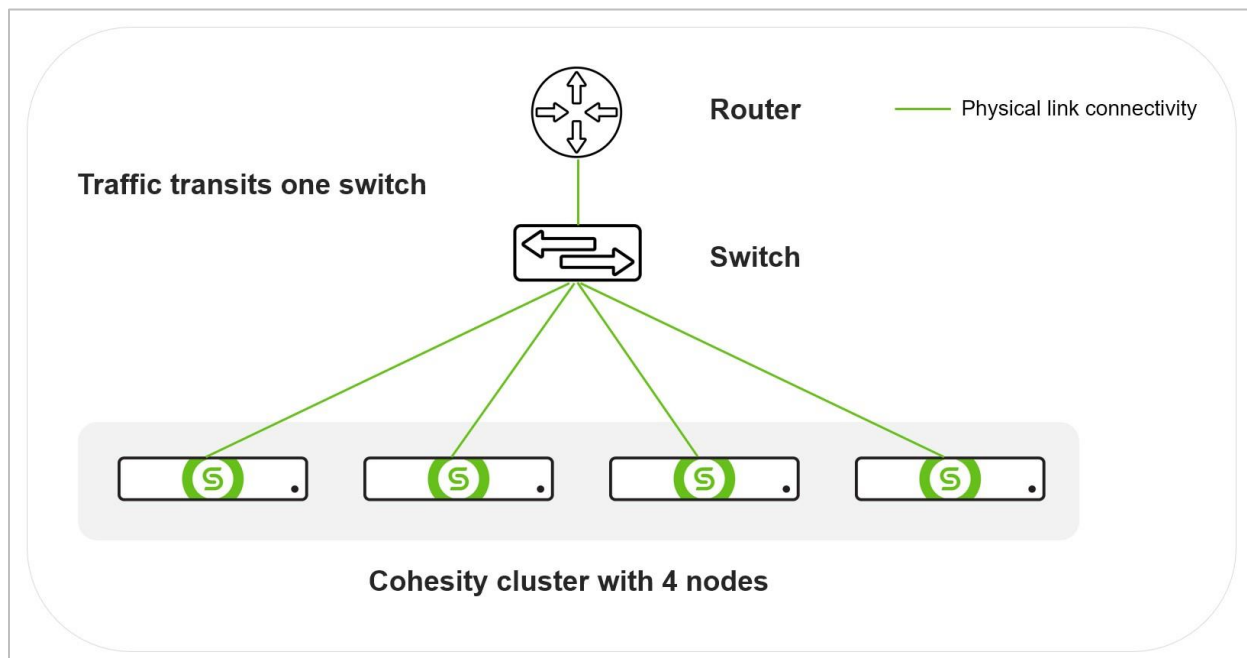
Five Common Configurations

The following sections provide the most common configurations encountered in production network environments.

Simple Topology

A simple topology is common for proof-of-concept testing and other non-production use cases.

Figure 1: Cohesity on a Simple Topology



NOTE: A single line in the diagram represents two network links per node that are bonded.

Characteristics of a simple network topology:

- A single path through a single switch from one Node to another Node.
- A single, separate path from every node to the router (North-South traffic).
- This enables a line-rate performance (for example, 10GbE) from one node to another.
- There is no inherent choke point, and all traffic is non-blocking.
- East-West inter-node traffic is deterministic and constant.
- North-South traffic is deterministic and constant.
- Ethernet communication happens directly between two MAC addresses and does not change.

Before choosing a simple topology, consider:

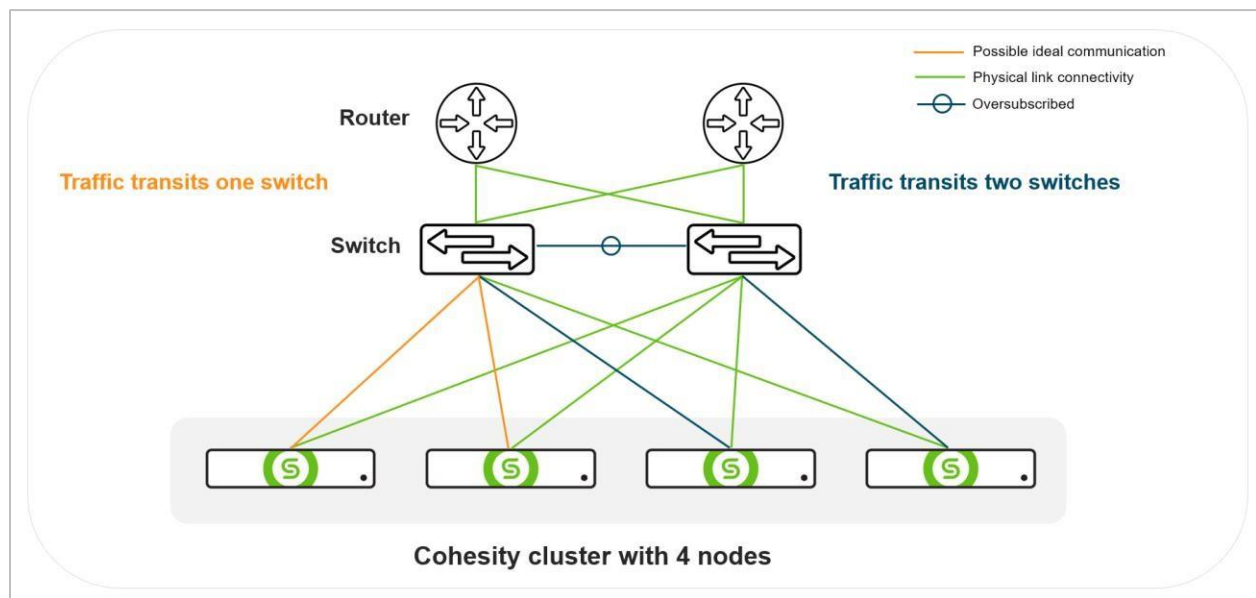
- No switch redundancy. A switch failure can lead to catastrophic failure that can make the whole cluster unreachable.
- There is a possibility of contention with other traffic. If other devices on the production network are connected to the same switch, it can lead to the switch becoming loaded and can lead to Ethernet drops and TCP retransmissions between nodes.

Best Practice: Default Bond Mode 1 enables active backup so that only one interface is in use, and the second interface is used when the first fails.

Standard Topology

A standard topology introduces a second switch for redundancy, which is achieved by connecting the two Ethernet interfaces to two separate switches.

Figure 2: Cohesity on a Standard Topology



Characteristics of a standard network topology:

- Provides Node NIC and network switch redundancy.
- All nodes might not forward traffic to the same switch, therefore forcing some traffic to transit the peer link between switches. This might introduce drops, delay, and latency.
- No special configuration is required on the network switches. (All node interfaces are required to be in the same VLAN)

Before choosing a standard topology, consider:

- Oversubscription or head-of-line blocking can occur where multiple links must traverse through the single inter-switch peer link. For example, multiple 10GbE links could go into a single 10GbE link.

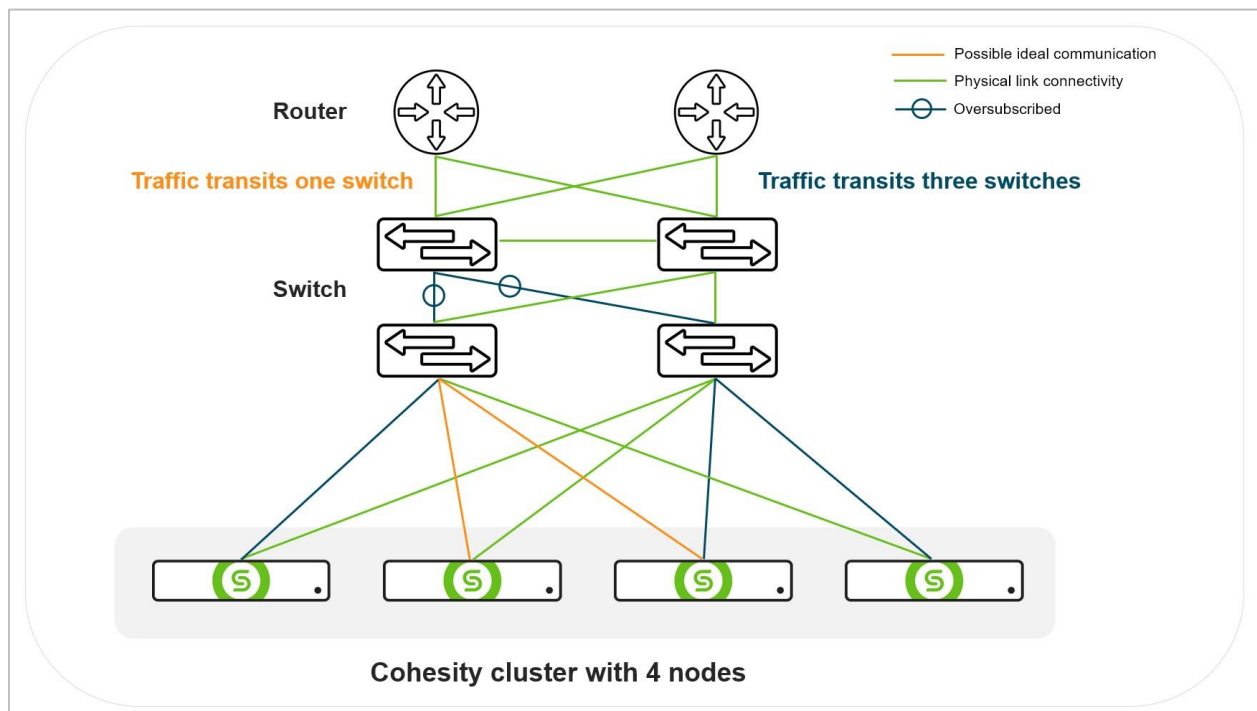
- Oversubscription can lead to frame drops. Frame drops are expensive and can lead to performance degradation because of packet retransmissions.
- Scaling becomes a problem as more switches are added.

Best Practice: For best results, configure Bond Mode 4 on the nodes with MLAG/VPC at the cluster switch layer.

Layered Topology

Some production environments have a layered topology to support many devices. This is like the standard topology, except that nodes may communicate over multiple switch layers to reach other nodes within the cluster.

Figure 3: Cohesity on a Layered Topology



Characteristics of a layered network topology:

- The peer link might be established only on the top-most layered switches, which means some traffic might have to flow through all the layers, and back down, before reaching its destination.
- Helps with redundancy because multiple switches are used.
- The peer links at the first-hop switches are sometimes removed to simplify the spanning tree topology.

Before choosing a layered topology, consider:

- This can lead to oversubscription if multiple paths traverse the same inter-switch links.

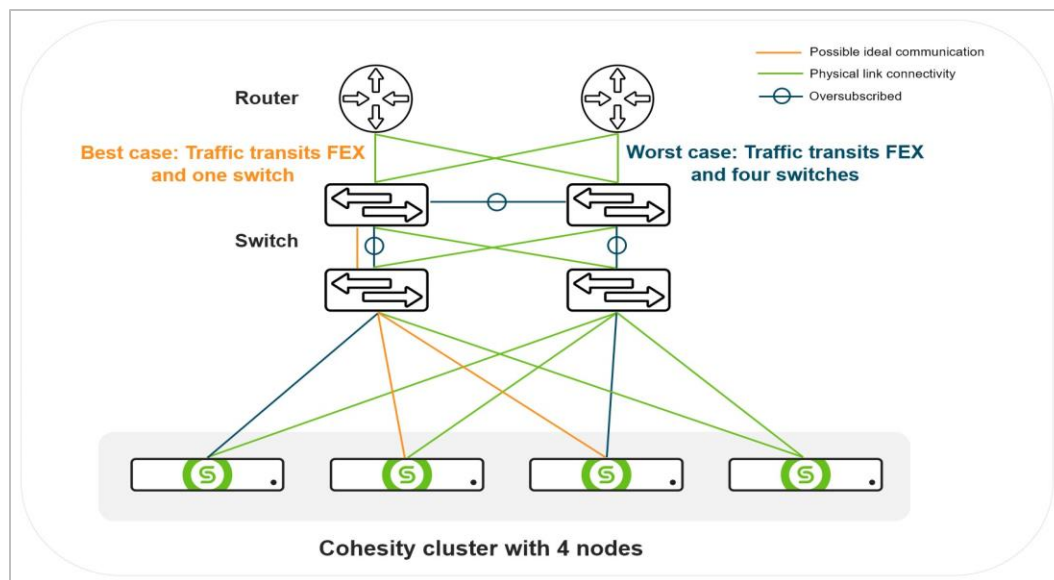
- This can lead to some nodes using a single hop and some using multiple hops to reach the other nodes, resulting in increased latencies.

Best Practice: Connect all Cohesity nodes directly to a pair of switches. The most ideal configuration is Bond Mode 4 with MLAG/VPC at the cluster switch layer.

Cisco FEX Topology

Production environments can include Cisco Fabric Extender Technology (FEX) switches that have already been installed for a legacy environment involving many network devices.

Figure 4: Cohesity on a Cisco FEX Topology



Characteristics of an FEX network topology:

- A Cisco FEX topology involves using intermediate FEX devices to increase port counts and improve fabric reach.

Before choosing an FEX topology, consider:

- An FEX is designed for north south traffic that must transit the uplink anyways.
- An FEX lacks the capability to perform any local switching.
- Can lead to traversing multiple hops between nodes.
- Even the use of Bond Mode 4 might not reduce oversubscription.
- For east-west traffic, as the number of paths, latency, and delays increase, the possibility of packet drops increases.
- Increased traffic congestion at the switch and router layer.

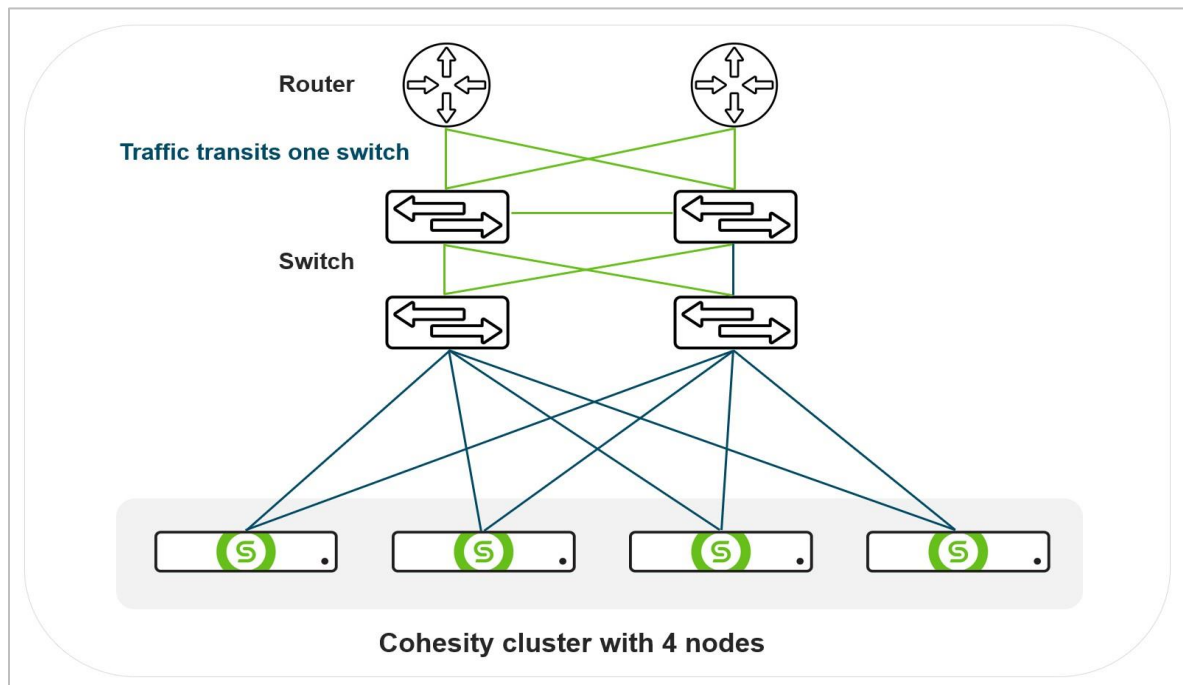
Best Practice: Cohesity supports Bond Modes 1 (active backup, the default) and 4 (802.3ad, aka LACP). Bond Mode 4 is useful with the LACP use case, where multiple ports on the switch might be configured to do LACP, while Bond Mode 1 can be used in the other use cases.

Layered Bond Mode 4 (LACP) Topology

A Bond Mode 4 topology currently provides the optimal solution for cluster connectivity. It solves the problem of oversubscription by having the traffic directly traverse only one intermediate switch instead of several intermediate switches. This topology involves configuring nodes for LACP, and a special configuration on the switches.

Administrators need to configure a pair of switches to perform either MLAG (Multi-Chassis Link Aggregation) or VPC (Virtual Port Channel).

Figure 5: Cohesity on a Layered Bond Mode 4 (LACP) Topology



Characteristics of a Bond Mode 4 (LACP) network topology:

- Cohesity nodes must be configured for Bond Mode 4 (802.3ad).
- Connected network switches must support LACP.
- Connection to two different switches requires software support from the switch vendor. For example, Cisco provides Virtual Port Channel (VPC) while Arista provides Multi-Chassis Link Aggregation (MLAG).

Before choosing a Bond Mode 4 topology, consider:

- Requires careful configuration and testing from the network team.

Best Practice: This topology represents the best configuration. Configure Bond Mode 4 on the nodes, connecting each interface on a node to a different switch, and configure the switches to enable VPC or MLAG.

Terminology

Table 2: Terminology

TERM	DEFINITION
Bond	A Linux virtual network interface used to aggregate physical network interfaces for the purpose of failover, increased throughput, or both. Bonds are numbered such that each Bond Mode <number> implies a certain way to aggregate network ports. Some bond types require special configurations on the switch that involves some initial negotiation on connectivity to make the bonding functionality successful.
Bond Mode 1	Combines network interfaces in an active-backup configuration, which can be useful in cases where the active one fails. In those cases, the backup interface can take over and continue serving traffic.
Bond Mode 4	Combines network interfaces in an active-active configuration, in accordance with the IEEE 802.3ad specification.
Cluster	A 3-node (minimum) deployment of Cohesity physical nodes.
FEX	A Cisco Nexus 2000 Series Fabric Extender is like a switch that helps extend the network fabric.
LACP	Link Aggregation Control Protocol, as defined in the IEEE 802.3ad specification.
MLAG and VPC	Multi Chassis Link Aggregation (Arista) and Virtual Port Channel (Cisco) enable the switches to learn about the interfaces that the other connects to.
NIC	Network Interface Card. Indicates a single network link on the node used to connect to other devices.
Node	A server running Cohesity software.
Virtual IPs (VIPs)	We recommend you use one virtual IP for each cluster node.

Conclusion

Cohesity provides a few different options for network topologies. It is possible to configure the internal networking for optimal performance without sacrificing efficiency. While configuring the network, administrators need to consider both the existing environment and their future needs.

We recommend that you deploy a layered topology using multiple switches and multiple NIC ports going from each node to each switch. If MLAG or VPC support is available on the switches, a Layered Bond Mode 4 topology represents the best configuration and solves the network oversubscription problem, while simultaneously leveraging both the switches to provide additional redundancy.

Your Feedback

Was this document helpful? [Send us your feedback!](#)

About the Authors

Barnaby DiAnni is a Principal SRE at Cohesity.

Other essential contributors include:

- Arvind Jagannath, Product Management
- Hashim Zargar, Senior Technical Solutions Engineer

Document Version History

VERSION	DATE	DOCUMENT HISTORY
1.2	Feb 2024	Updating the note under Simple Topology section.
1.1	July 2021	Cohesity Rebranding Updates
1.0	Nov 2018	First release

ABOUT COHESITY

[Cohesity](#) radically simplifies data management. We make it easy to protect, manage, and derive value from data -- across the data center, edge, and cloud. We offer a full suite of services consolidated on one multicloud data platform: backup and recovery, disaster recovery, file and object services, dev/test, and data compliance, security, and analytics -- reducing complexity and eliminating [mass data fragmentation](#). Cohesity can be delivered as a service, self-managed, or provided by a Cohesity-powered partner.

Visit our [website](#) and [blog](#), follow us on [Twitter](#) and [LinkedIn](#) and like us on [Facebook](#).

© 2024. Cohesity, Inc. All Rights Reserved. The information supplied herein is the confidential and proprietary information of Cohesity and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Cohesity software and services. Find the terms of Cohesity licenses at www.cohesity.com/agreements.

Cohesity, the Cohesity logo, SnapTree, SpanFS, DataPlatform, DataProtect, Helios, the Helios logo, DataGovern, SiteContinuity, DataHawk, and other Cohesity marks are trademarks or registered trademarks of Cohesity, Inc. in the US and/or internationally. Other company and product names may be trademarks of the respective companies with which they are associated. This material (a) is intended to provide you information about Cohesity and our business and products; (b) was believed to be true and accurate at the time it was written, but is subject to change without notice; and (c) is provided on an "AS IS" basis. Cohesity disclaims all express or implied conditions, representations, warranties of any kind.