



Installation Guide (version 2.0.45)

Scale Computing Storage Cluster Installation Guide

Table of Contents

1. How a Scale Storage Cluster Works	1
For NFS and CIFS implementations:	2
2. Installing a Scale Storage Cluster	4
Provisioning IP Addresses for a Scale Storage Cluster	4
About Private IP Addresses	4
About Public IP Addresses	4
About Virtual IP Addresses (CIFS and NFS only)	5
Example IP Addressing	5
Configuring DNS for a Scale Storage Cluster	5
Before Configuring	6
Racking and Cabling the Nodes	7
Racking the Nodes	7
Cabling the Nodes	9
Node Configuration	10
Configure the First Node in the Cluster	10
Configure Additional Nodes	14
Initialize the Cluster	15
Completing Cluster Configuration	17

List of Figures

1.1. Network Cable and IP Configuration	1
1.2. Cloud of Storage Resources	2
1.3. Nodes Functioning as a Cohesive Unit	3
2.1. Inner slide assemblies	7
2.2. Securing the slide assemblies	8
2.3. Sliding the chassis into place	9
2.4. Chassis in place	9
2.5. Public and Private Ports	9
2.6. Node Console Menu	11
2.7. Configure Node	12
2.8. Applying the Configuration	13
2.9. Logging Out of the Node Console Menu	14
2.10. Configuring Additional Nodes	15
2.11. Initializing the Cluster	16
2.12. A Configured Node	17

Chapter 1. How a Scale Storage Cluster Works

Before you start the installation of the Scale Storage Cluster it is important to understand how the system is architected to ensure a fast, smooth installation.

Scale's Storage cluster is deployed as a collection of 1U systems connected together using Gigabit Ethernet (GigE). Each 1U system or "storage node" is an off-the-shelf storage server paired with robust supercomputing storage software. The power of clustering is the reliability and performance that is gained by having multiple devices acting as one unified storage system. There is no master control unit. Each device can manage itself and the rest of the cluster, even if configurations change. Also, each node in the cluster serves as an entry point into the data store, which provides for a high level of aggregate bandwidth available to end clients.

To create cluster functionality, DNS plays a key role in NFS and CIFS installations and it is necessary that you have a DNS server capable of setting up round-robin addressing. DNS is not necessary for iSCSI-only implementations.

Each node in a cluster has 2 or more GigE ports. At least one port is for "private" traffic which operates as the backplane for the cluster. The private network uses a dedicated GigE switch or VLAN. The other port is for "public" traffic. The public network is where other clients connect to the storage cluster. Figure 1.1 illustrates basic network connections.

Figure 1.1. Network Cable and IP Configuration



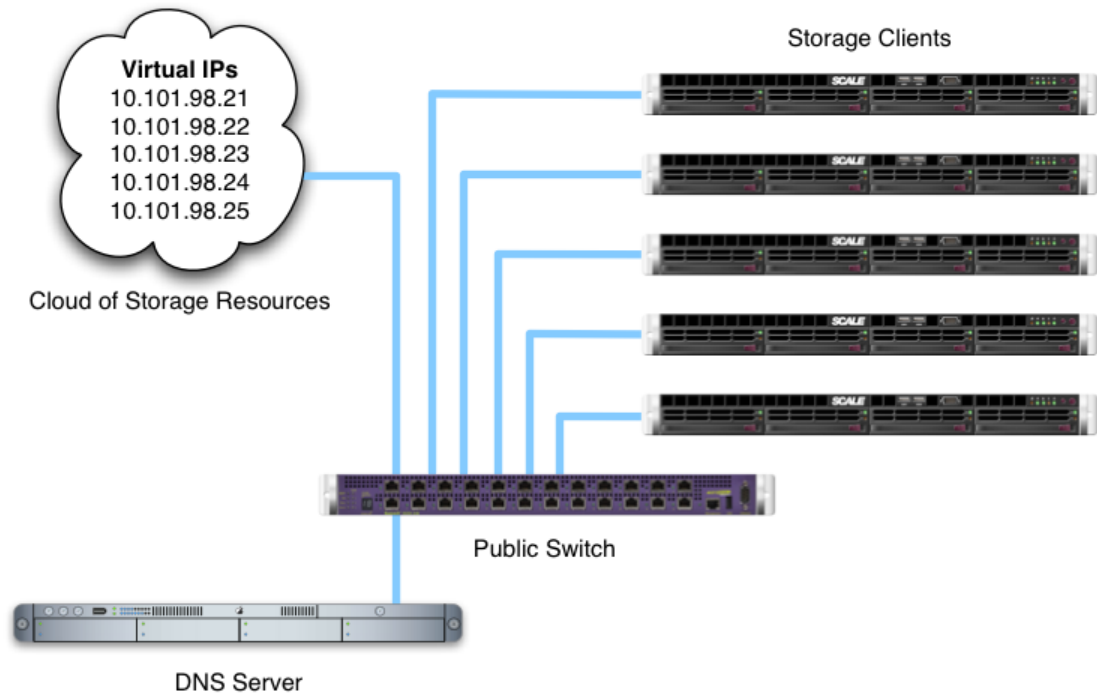
You will need to set aside up to 3 IP addresses for each node in the cluster. The private IP address serves only to identify the brick to the other storage nodes on the private network or VLAN. The private IPs can be provisioned on a completely separate subnet, and need not be routable to the public network. For example, if your corporate network uses IP addresses in the 10.101.98.0/16 network, you can provision the private IPs in the 192.168.98.0/16 network without conflict. The public IP address is the front-facing IP for management and iSCSI targets. The third IP address is a virtual IP address setup in DNS for only NFS and CIFS implementations. Please keep in mind the Scale Storage Cluster can run iSCSI, NFS and

CIFS all simultaneously on the same cluster. If you intend to only use the iSCSI functionality, you do not require the third, or virtual, IP address.

For NFS and CIFS implementations:

The DNS round-robin entry for NFS and CIFS deployments is key to ensure the Scale Storage cluster operates effectively. After configuration, you will have a cloud of storage resources that is accessed via the virtual IP addresses.

Figure 1.2. Cloud of Storage Resources



As nodes fail or are added, the virtual IP addresses fluidly move to other nodes around the cluster, enabling clients to maintain connection in the event of node failures. For example, as shown in the Figure 1.2, the cloud is made up of 4 individual nodes. Assuming all nodes are operational, each node corresponds to a virtual IP address.

Figure 1.3 illustrates an instance where one of the nodes is offline. As a result, one of the virtual IP addresses moves to an operational unit. There are 3 units functioning, one of which has 2 virtual IP addresses associated with it.

Figure 1.3. Nodes Functioning as a Cohesive Unit



Chapter 2. Installing a Scale Storage Cluster

Following is a list of tasks that create the Scale Storage Cluster. It is highly recommended you complete the tasks in the order shown below:

1. Provision IP addresses
2. Create a round-robin DNS entry (for NFS and CIFS)
3. Identify an external or internal time server the cluster can access
4. Rack and cable the cluster nodes
5. Initially configure the cluster nodes
6. Complete configuration via web interface

Provisioning IP Addresses for a Scale Storage Cluster

Before you provision IPs you must decide how you will deploy the storage cluster:

- iSCSI only
- iSCSI, NFS and CIFS
- NFS or CIFS only

Each node in the cluster will need three IP addresses. We refer to them as Private, Public and Virtual IP addresses. If you are using the cluster for only iSCSI, you only need 2 IP addresses for each node, as the virtual IP address is not used.

About Private IP Addresses

The private IP addresses do not need to be routable. They are used only for intra-cluster communication between nodes on a separate GigE switch or VLAN. All node-to-node traffic including data striping and mirroring occurs on this network.

In a typical installation, Scale Computing recommends setting the private IP addresses to a totally different IP address range than you normally use. For example, if you use 10.0.0.0/8 for internal traffic, use 192.168.0.0/16 for private IP addresses to ensure no unexpected routing issues.

The best practice for IP address assignment is to have the same final component for both the private and public IP addresses. For example, if a given node's assigned public IP address is 10.1.1.245, and the private network is 192.168.1.0/24, then use 192.168.1.245 for the private IP address.

About Public IP Addresses

The public IP addresses are the IP addresses of the nodes on your internal network. These are also the addresses that will be used by iSCSI initiators to connect to the cluster. In order for Scale Computing to

provide remote support, at least one of these addresses must be able to make an outgoing ssh connection over the Internet on TCP port 22.

About Virtual IP Addresses (CIFS and NFS only)

In addition to the public IP addresses, you must reserve one additional IP address per node in the cluster, to act as a virtual IP address. You must create a round-robin DNS entry that points to all of these IP addresses. The virtual IP addresses are assigned to the cluster through the web interface, and are managed by the cluster. These IP addresses are migrated from node to node in the event of a failure or other node-related issues. You do not need virtual IP addresses in iSCSI-only deployments.

Example IP Addressing

Following is an example of IP addresses for a five node cluster.

Public IP Addresses

node01 A 10.101.98.1

node02 A 10.101.98.2

node03 A 10.101.98.3

node04 A 10.101.98.4

node05 A 10.101.98.5

Private IP Addresses

node01 A 192.168.98.1

node02 A 192.168.98.2

node03 A 192.168.98.3

node04 A 192.168.98.4

node05 A 192.168.98.5

Virtual IP Addresses

cluster A 10.101.98.21

cluster A 10.101.98.22

cluster A 10.101.98.23

cluster A 10.101.98.24

cluster A 10.101.98.25

Configuring DNS for a Scale Storage Cluster

This is for implementations needing to use NFS and CIFS protocols. If you plan on only deploying iSCSI, skip to the section called “Racking and Cabling the Nodes”.

Before Configuring

Before configuring, you must have the following information:

- Private and Public IP addresses for each node (there is a third IP, but that will be identified later in the installation)
- The list of virtual IP addresses to be assigned to this cluster
- subnet masks for both the public and private networks
- The IP address of your internal DNS server
- A gateway or router IP address

Start the configuration by setting up a round robin DNS entry. Round robin DNS is a technique in which load balancing is performed by a DNS server instead of a dedicated machine. A DNS round-robin record has more than one IP address value. When a request is made to the DNS server which serves this record, the answer it gives alternates for each request. For instance, if you had three web servers between which you wish to distribute requests, you could setup your DNS zone as one of the following examples demonstrates.

In most NFS and CIFS implementations with the storage cluster, round robin distributes load evenly over the nodes, effectively aggregating the available bandwidth of the entire cluster.

Configuring DNS round robin using Windows

Log into Windows DNS manager and create an 'A' record name associated with each Virtual IP address you set aside.

Configuring DNS round robin using BIND

Below is an example BIND configuration file that includes multiple 'A' records related to the cluster virtual IPs. Each line starting with **cluster** is an 'A' record pointing to a unique IP address. As BIND responds to DNS requests to resolve **cluster**, it will rotate among the available IP addresses.

```

$ORIGIN .
$TTL 86400           ; 1 day
example.com         IN SOA  ns1.example.com. hostadmin.example.com. (
                        2009011601           ; serial
                        3600                 ; refresh (1 hour)
                        1800                 ; retry (30 minutes)
                        604800              ; expire (1 week)
                        86400               ; minimum (1 day)
                        )
                        NS      ns1.example.com.
$ORIGIN scale.example.com.
ns1                 A      192.168.1.65
node01              A      10.101.98.1
node02              A      10.101.98.2
node03              A      10.101.98.3
node04              A      10.101.98.4

```

node05	A	10.101.98.5
cluster	A	10.101.98.21
cluster	A	10.101.98.22
cluster	A	10.101.98.23
cluster	A	10.101.98.24
cluster	A	10.101.98.25

Racking and Cabling the Nodes

Racking the Nodes

Before you start

The Rail Packaging includes:


- One pair of inner slides, to be installed on the chassis
- One pair of outer slides, to be installed in the rack
- Two pairs of short brackets, to be used on the front side of the outer slides. Note: One pair of short brackets includes screw threads; the other pair does not include screw threads. Use only the pair that fit into your rack.
- One pair of long brackets, to be used on the rear side of the outer slides

You will need a Philips-head screwdriver.

Assembling the Inner-Slides of the Rack

1. Locate the right inner slide, which is the slide that will be used on the right side of chassis when facing the front panel of the chassis.
2. Align the 4 square holes on the right inner slide against the hooks on the right side of the chassis as show below on the left.
3. Securely attach the slide to the chassis with two M4 Philips-head screws and repeat the steps 1-3 to install the left inner slide to the left side of the chassis.

Figure 2.1. Inner slide assemblies




Installing Outer Slides of the Rack

1. Measure the distance from the front rail of the rack to the rear rail of the rack.
2. Attach a short bracket to the rear side of the right outer slide, and a long bracket to the front side of the right outer slide as shown above on the right.
3. Adjust the short and long brackets to the proper distance so that the chassis can snugly fit into the rack.
4. Secure the slides to the cabinet with screws. Repeat steps 1-4 for the left outer slide.

Installing the Slide Assemblies to the Rack

To install the slide assemblies:

Figure 2.2. Securing the slide assemblies



1. After you have installed the short and long brackets to the outer slides, you are ready to install the whole slide assemblies—the outer slides with short and long brackets attached—to the rack. See *Assembling the Inner-Slides of the Rack*.
2. Use M5 screws and washers to secure the slide assemblies into the rack, as shown in Figure 2.2.

Installing the Chassis into the Rack

Figure 2.3. Sliding the chassis into place





Figure 2.4. Chassis in place



1. Push the inner slides, which are attached to the chassis, into the grooves of the outer slide assemblies that are installed in the rack, as shown in Figure 2.3
2. Push the chassis all the way to the back of the outer slide assemblies, as shown in Figure 2.4

Cabling the Nodes

Figure 2.5. Public and Private Ports



1. After you install the nodes in the rack, connect the nodes to a power outlet. NOTE: Scale Computing recommends that you distribute the nodes between several power outlets that are connected to two or more Uninterruptible Power Source units (UPS) to ensure that the nodes continue to receive electricity.
2. Connect network cables to the nodes. Each node has at least two GigE ports labeled public and private, as shown in Figure 2.5.
3. Configure the private network—yellow ports and cables—to its own dedicated GigE switch or VLAN.
4. Connect the public port—blue port and cables—to the same switch or network that any network devices use to connect to the storage cluster.

Node Configuration

Prior to configuring the individual nodes, be sure to verify network and cabling setup has been completed. The Scale configuration software confirms connectivity between nodes, as well as the ability to connect to a gateway before allowing configuration to continue.

In order to configure your Scale Cluster, you will need to have serial or graphical console access to the physical nodes. Serial access is **VT100** at **115,200** baud. Any standard VGA monitor and USB or PS/2 keyboard will also work as a graphical console.

Setting up a Scale cluster can be broken down into the following steps:

- Configure the first node in the cluster.
- Configure all additional nodes in the cluster, pointing these nodes to the first node.
- Initialize the entire cluster on the last node.

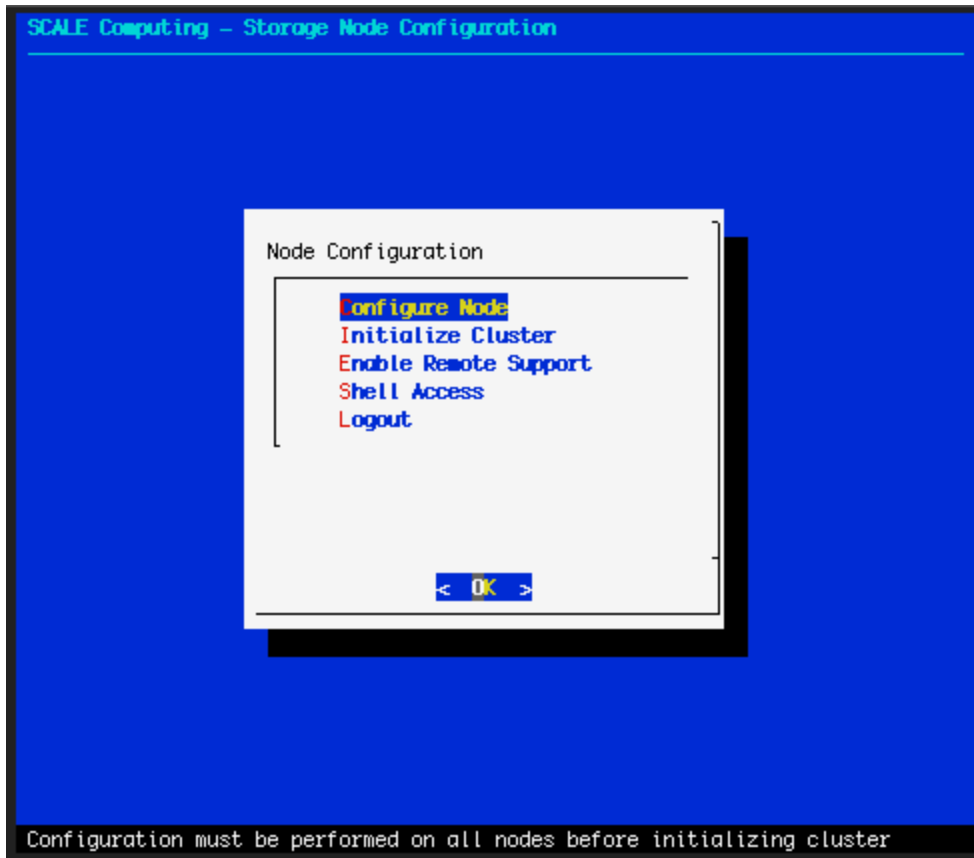
Configure the First Node in the Cluster

Log into the first node with the following credentials:

user: admin

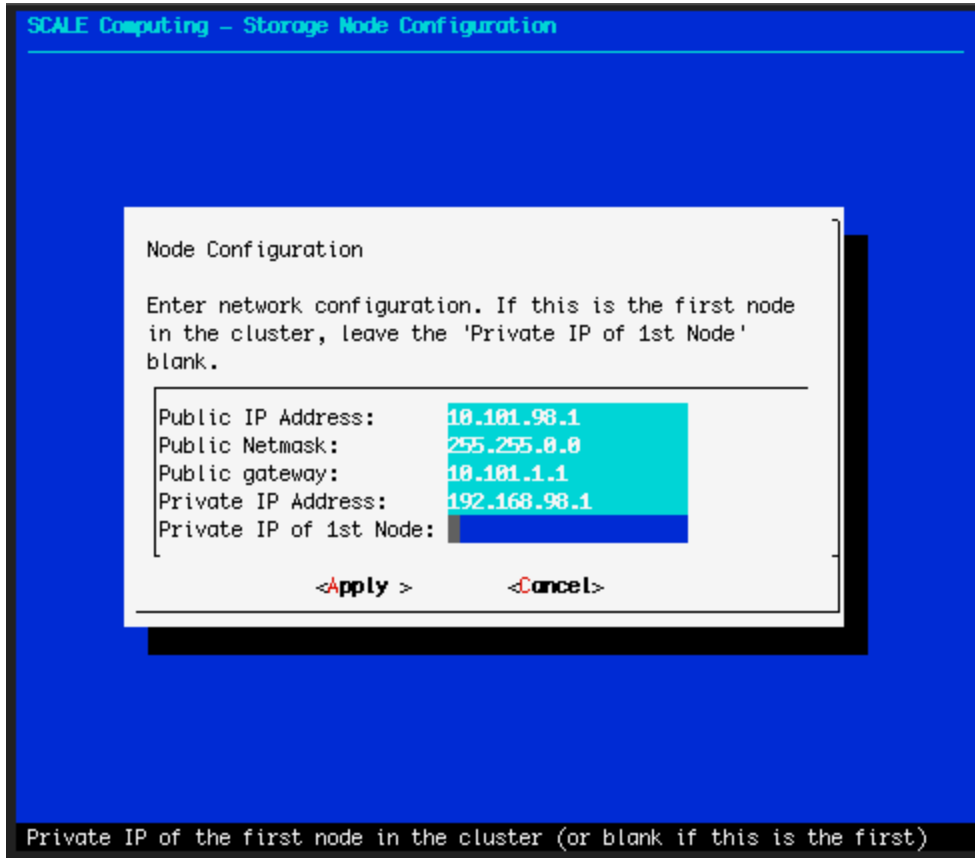
password: scale

You will be presented with the following graphical menu:


Figure 2.6. Node Console Menu

Select the first option in the menu: Configure Node. You will be presented with a form for inputting the networking information for the first node. Use the arrow keys to move between the fields. Press **tab** to cycle between the entry fields and the **<Apply>** and **<Cancel>** buttons. For the first node in the cluster leave the last field (**Private IP of 1st Node**) blank.

Figure 2.7. Configure Node



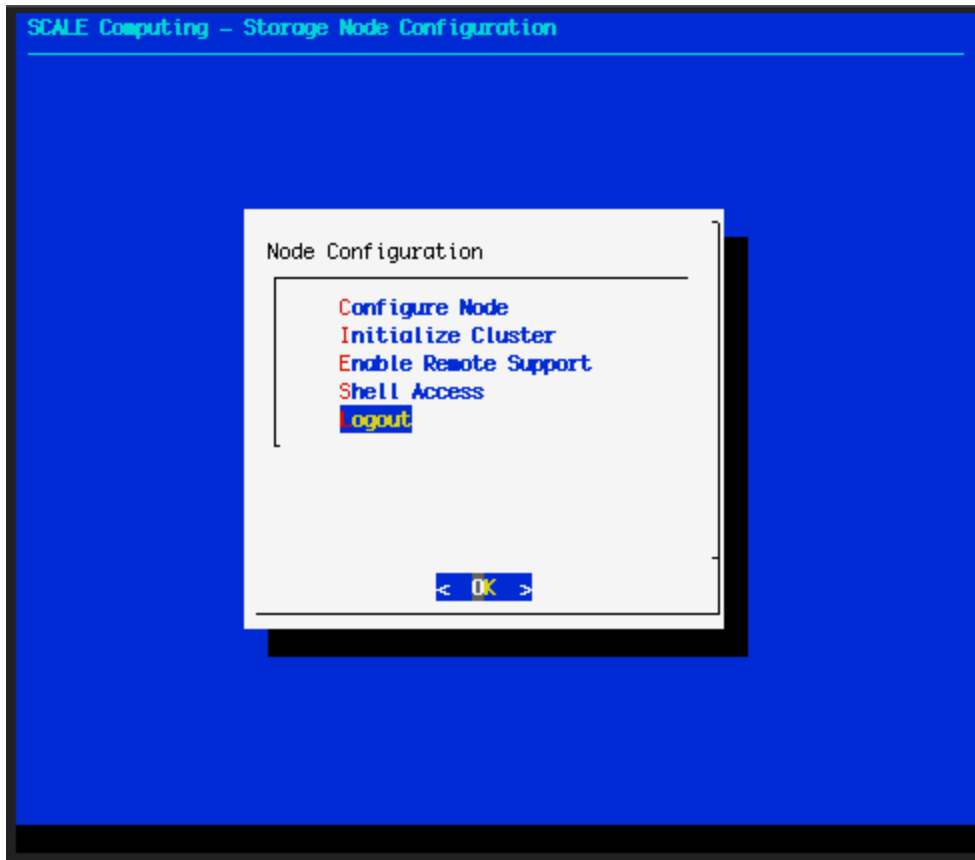
Press **tab** to cycle to the **<Apply>** button and hit return. The node will perform initial network configuration, and run internal checks to verify it is able to communicate with the network. You will see the following screen while it is performing these checks:

Figure 2.8. Applying the Configuration

If there are any errors in the configuration you will be presented with a dialog reporting the error encountered. You can return to the configuration screen, or optionally continue past the error. It is not recommended to continue past errors reported by the configuration routines. Continuing with the install may result in a misconfigured cluster node, and require Scale technical support staff to resolve.

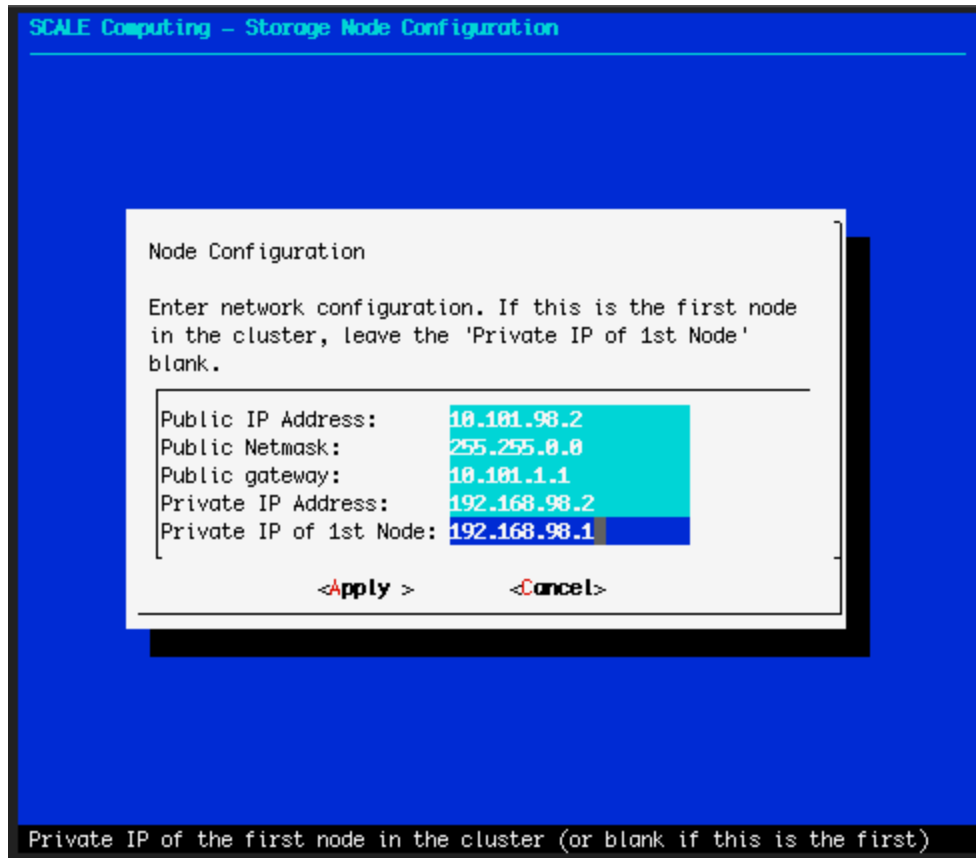
Once configuration has completed it will return to the main menu. You can now logout of the administrative interface.

Figure 2.9. Logging Out of the Node Console Menu



Configure Additional Nodes

Once configuration has completed on the first node, move your console to the second node in the cluster, logging in with the same credentials. Select **Configure Node** from the menu, and provide the networking formation for this node. For this node, instead of leaving the last entry blank, enter the private IP address of the 1st node in the cluster.

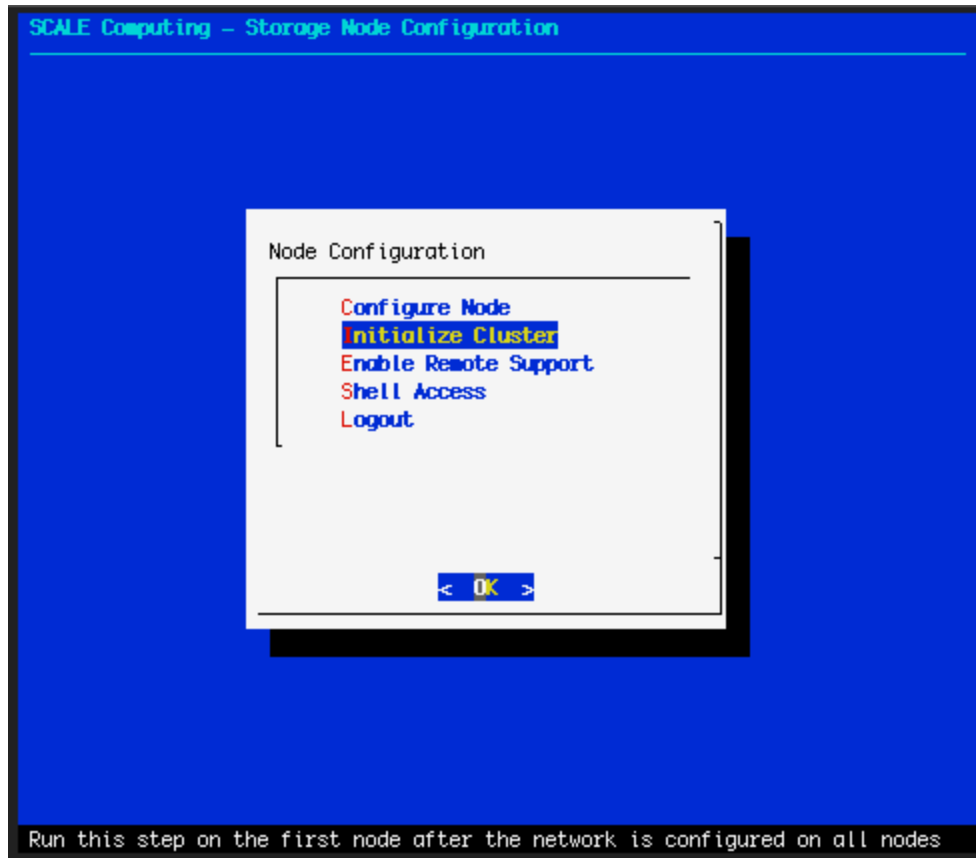
Figure 2.10. Configuring Additional Nodes

After applying the settings, the configuration software will verify network settings and communicate with the first node in the cluster. Once the configuration settings have been successfully applied, logout and connect to the next node in the cluster. Perform the same configuration on all remaining nodes in the cluster, providing each with its unique public and private IP addresses, and pointing them back to the first node in the cluster. Make sure each node has completed applying its network configuration before proceeding to the next node.

Initialize the Cluster

Once you have completed configuring all the nodes in the cluster, you are ready to perform the cluster-wide initialization. Return to the first node in the cluster and log back in as the admin user. From the main menu select the **Initialize Cluster** menu option.

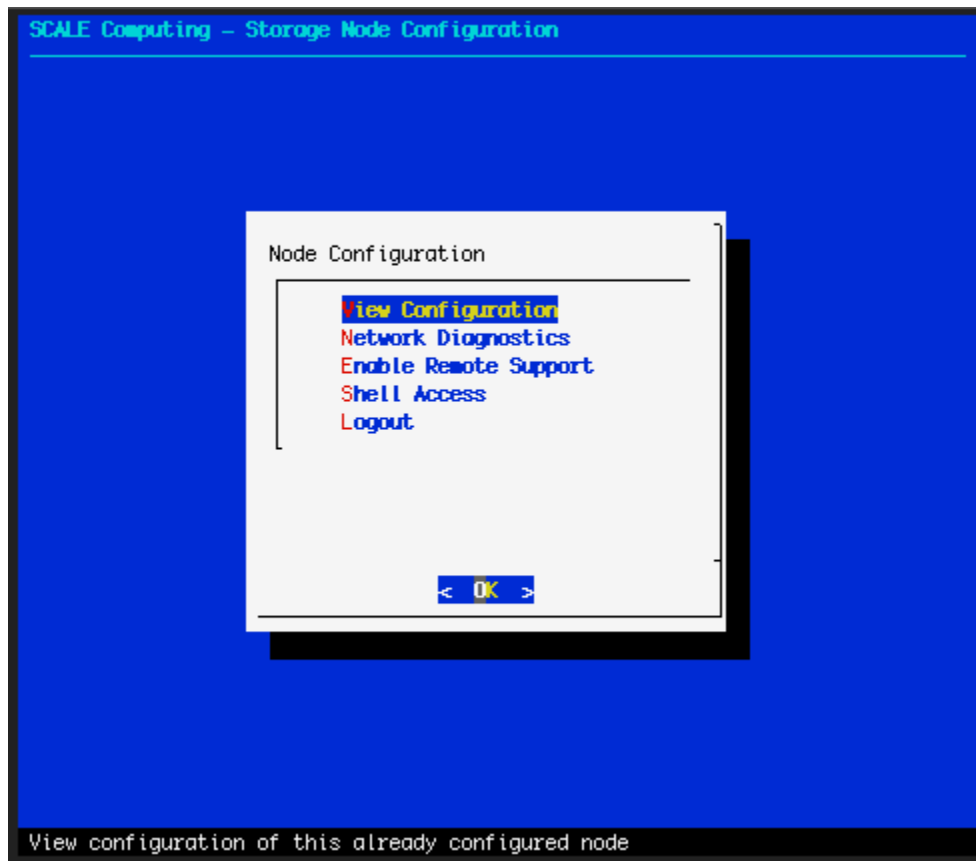
Figure 2.11. Initializing the Cluster



This will perform the whole-cluster initialization, and you will see disk activity on all nodes in the cluster. Depending on the size of your cluster, this initialization can take several minutes.

After configuration is complete you will be returned to the main menu. You will notice the menu options have changed for a configured cluster. You can use the administrative interface to check a node's configuration settings, perform network diagnostics, as well as enable remote support.

Figure 2.12. A Configured Node



Completing Cluster Configuration

Once you have completed setting up each node and run the cluster initialization you are finished with the console-based configuration. All remaining configuration is performed in the web-based interface. See the Scale Storage Cluster User Guide for more information.