

Scale Computing Storage Cluster Installation Guide

**Scale Computing
2121 S. El Camino Real
San Mateo, CA**

CHAPTER 1	How a Scale Cluster Works.....	1
	Cluster Structure.....	1
	LAN and Backplane IP Addresses	1
	Maintaining Cluster Availability	4
	Special Requirements for NFS and CIFS Implementations.....	5
	Revision History.....	7
CHAPTER 2	Installing a Scale Cluster	9
	Provisioning IP Addresses for a Scale Cluster.....	9
	About Backplane IP Addresses	10
	About LAN IP Addresses.....	10
	About Virtual IP Addresses (CIFS and NFS Only).....	11
	Example IP Addressing	11
	Configuring DNS for a Scale Storage Cluster.....	12
	Before Configuring	12
	Configuring DNS Round Robin Using Windows	12
	Configuring DNS Round Robin Using BIND	13
	Racking and Cabling the Nodes	14
	Racking the Nodes	14
	Rail Kit Contents	15
	Putting the Front and Back Chassis Slides on the Chassis.....	18
	Removing Front Slides from the Long Brackets.....	19
	Installing Front Slides on the Chassis	19
	Installing Back Slides on the Chassis	21
	Putting the Back Rack Rails, Long Brackets, and Short Brackets Together	23
	Installing the Chassis into the Rack	25
	Cabling the Nodes.....	25
	Node Configuration	26
	Configure the First Node in the Cluster	27
	Configure Additional Nodes	31
	Initialize the Cluster	32
	Completing Cluster Configuration	34
	Revision History.....	35

List of Figures

Cluster Configuration for Regular Nodes	3
Cluster Configuration for High Availability Nodes	4
Cloud of Storage Resources	6
Nodes Functioning as a Cohesive Unit	7
Short Brackets - Threaded and Unthreaded	15
Front Chassis Slides	16
Back Chassis Slides.	16
Back Rack Rails	17
Long Brackets	17
Panhead Screw	18
Flathead Screws.	18
Removing Front Slides from the Long Brackets.	19
Chassis Hooks	20
Orientation of Front Slides on Chassis	21
Front and Back Slides on Chassis	22
Chassis Sides with Front and Back Slides Installed, Side Views	22
Sliding the Long Bracket into the Back Rack Rail	23
Positioning Short Bracket for Rack.	24
Sliding the Chassis into Place	25
Backplane and LAN Ports.	26
Node Console Menu	28
Configure Node	29
Applying the Configuration	30
Logging Out of the Node Console Menu	31
Configuring Additional Nodes	32
Initializing the Cluster.	33
A Configured Node	34

How a Scale Cluster Works

Before you start the installation of the Scale cluster it is important to understand how the system is architected to ensure an easy, smooth installation. The key concepts you need to know prior to installing your Scale cluster are discussed in the following sections:

- **Cluster Structure**
- **LAN and Backplane IP Addresses**
- **Special Requirements for NFS and CIFS Implementations**

Cluster Structure

The Scale 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 Scale's robust supercomputing storage software.

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.

LAN and Backplane IP Addresses

The Scale cluster uses a private subnetwork as a backplane to connect nodes to one another. This subnetwork is only accessible by nodes in your cluster and enables them to communicate their status to one another and perform data striping and mirroring. To ensure proper function-

ality of your cluster, the selected IPs for this subnetwork should be comprised of IP addresses that do not overlap with any others within your company.

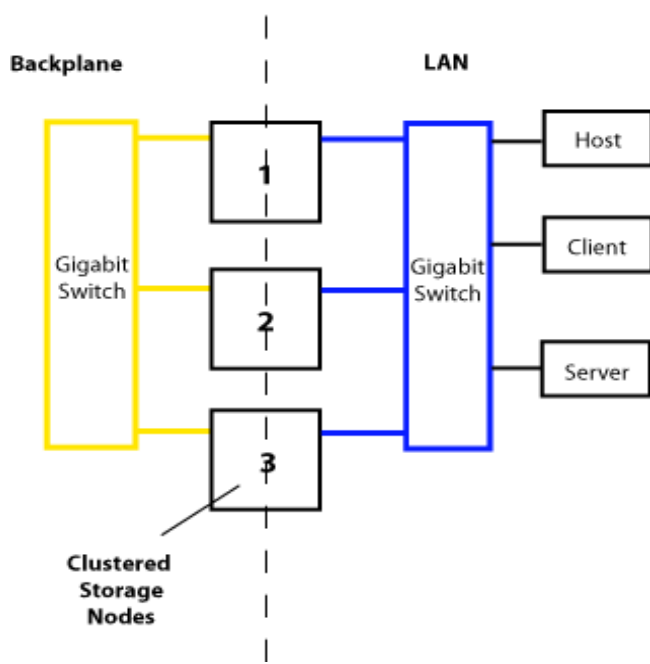
NOTE: The IP addresses you select are PERMANENT. You will not be able to change them once you set them. Careful planning prior to choosing these addresses is recommended.

When connecting a node to your cluster, use the port(s) labelled **Backplane** to construct the private subnetwork that will act as your cluster's backplane. Use the port(s) labelled **LAN** to connect your cluster to the rest of your company's network.

You will need to set aside up to 3 IP addresses for each node in the cluster. These addresses are the backplane IP address, so that your node can communicate with other nodes in the cluster, the LAN address so that your node can talk to the rest of your network, and the virtual IP address, which is setup in DNS for NFS and CIFS implementations. The Scale cluster can run iSCSI, NFS, and CIFS simultaneously on the same cluster. If you intend to only use the iSCSI functionality, you do not require the third IP address.

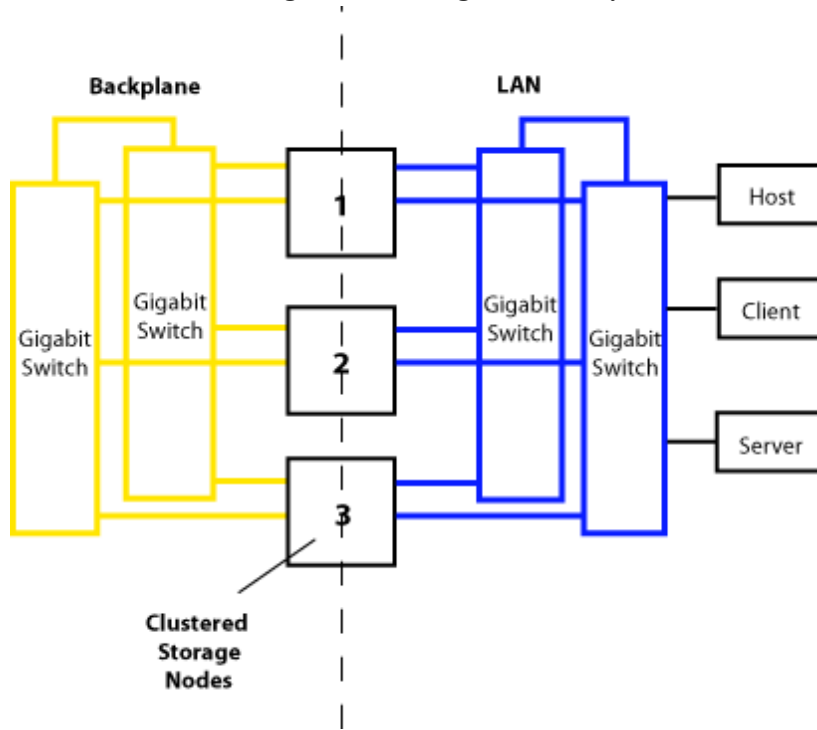
Configuration for a cluster comprised of regular nodes (a regular node has one **Backplane** port and one **LAN** port) would look like the example cluster shown in **Figure 1-1, Cluster Configuration for Regular Nodes**.

FIGURE 1-1. Cluster Configuration for Regular Nodes



Configuration for a cluster comprised of high availability (HA) nodes (a HA node has two **Backplane** ports and two **LAN** ports) looks like the example shown in **Figure 1-2, Cluster Configuration for High Availability Nodes**.

FIGURE 1-2. Cluster Configuration for High Availability Nodes



In each configuration example, the network used to form the backplane for your cluster is kept completely separate from the rest of your network. This concept is key to configuring an effective, functional cluster.

Maintaining Cluster Availability

Your Scale cluster can continue operation without issue if you have drive failure on a single node. However, if drives fail on more than one node, your cluster's filesystem unmounts in order to protect your data from corruption.

You can help ensure cluster availability by doing the following:

- Responding immediately if a node indicates a drive is down.

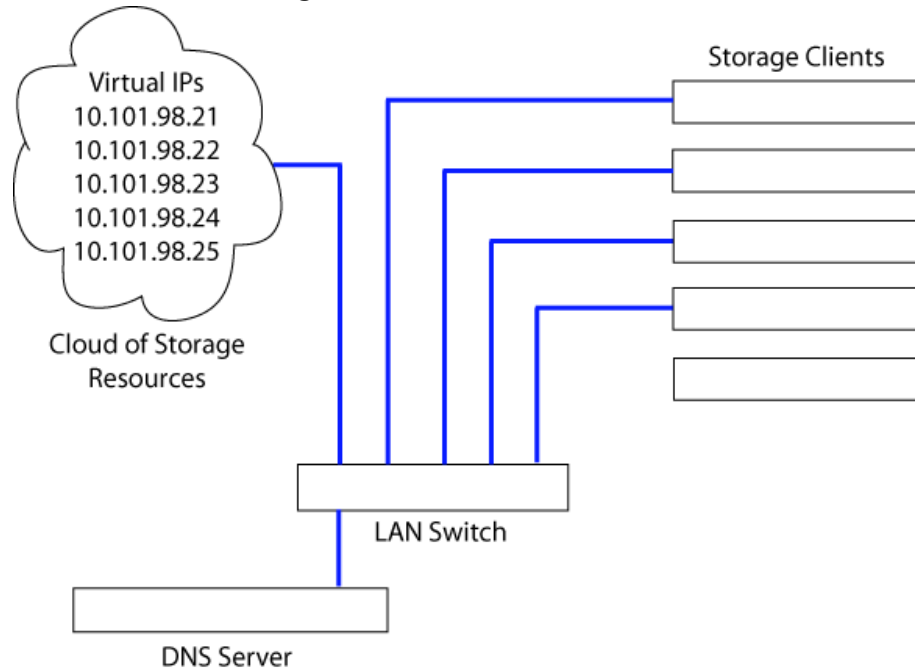
- Avoiding disconnecting nodes from the private network acting as your cluster's backplane.
- Ensuring all cables are securely plugged in for all nodes in your cluster.

It is not recommended that you move nodes around your cluster by unplugging them from the backplane and then plugging them in again. If you determine that you must move nodes by unplugging them from the backplane, be aware that you must perform this procedure one node at a time. This is because an unplugged node registers as a failed node on your cluster. If an additional node fails or has a drive fail at the same time, the filesystem unmounts. When you plug a node back in, you must wait until it is fully online and available before unplugging another node and moving it. Otherwise, the filesystem unmounts.

Special Requirements for NFS and CIFS Implementations

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

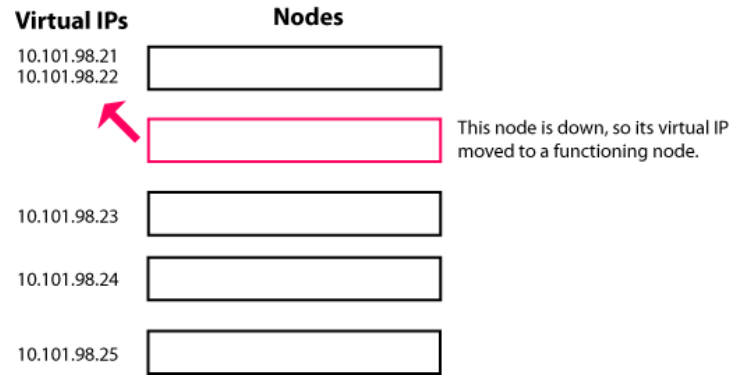
FIGURE 1-3. Cloud of Storage Resources



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 a node failure. For example, as shown in **Figure 1-3, Cloud of Storage Resources**, the cloud is made up of 5 individual nodes. When all nodes are operational, each node corresponds to a virtual IP address.

Figure 1-4, Nodes Functioning as a Cohesive Unit 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 4 units functioning, one of which has 2 virtual IP addresses associated with it.

FIGURE 1-4. Nodes Functioning as a Cohesive Unit



Revision History

This section contains information describing how this chapter has been revised.

Release 2.3.1:

- No changes this release.

Installing a Scale Cluster

When installing your Scale cluster, it is highly recommended you complete the following tasks in the order listed:

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

The tasks required to create the Scale cluster are discussed in the following sections:

- **Provisioning IP Addresses for a Scale Cluster**
- **Configuring DNS for a Scale Storage Cluster**
- **Racking and Cabling the Nodes**
- **Completing Cluster Configuration**

Provisioning IP Addresses for a Scale Cluster

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

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

Each node in the cluster needs three IP addresses. In Scale Computing's documentation, these three addresses are referred to as backplane, LAN, and virtual IP addresses. Backplane IP addresses refer to IP addresses assigned to a private subnetwork that acts as your cluster's backplane. LAN IP addresses refer to IP addresses used to connect to your company's network. Virtual IP addresses are used for set-ups that include the use of NFS, CIFS, or NFS and CIFS. If you are using the cluster for iSCSI only, you do not require virtual IP addresses for your nodes.

About Backplane IP Addresses

Backplane IP addresses are addresses in a private subnetwork used solely by your Scale cluster's nodes. These addresses are part of the non-routable subnetwork acting as your cluster's backplane. They are used only for intra-cluster communication between nodes, data striping, and mirroring.

In a typical installation, Scale Computing recommends setting the backplane IP addresses to an IP address range that does not overlap with any other address ranges assigned in your network.

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

About LAN IP Addresses

The LAN IP addresses are the IP addresses of the nodes on your internal network. These are also the addresses that are 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 LAN 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.

- **LAN 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
- **Backplane 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 using NFS and CIFS protocols. If you plan on only deploying iSCSI, skip to section [Racking and Cabling the Nodes](#).

Before Configuring

Before configuring, you must have the following information:

- Backplane and LAN IP addresses for each node (there is a third IP, but that is identified later in the installation)
- The list of virtual IP addresses to be assigned to this cluster
- Subnet masks for both the LAN and backplane 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.

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 rotates among the available IP addresses.

```
<![CDATA[
$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

This section discusses racking and cabling nodes in two sections:

- **Racking the Nodes**
- **Cabling the Nodes**

Racking the Nodes

This section discusses how to rack nodes in your cluster in the following subsections:


- **Rail Kit Contents**
- **Putting the Front and Back Chassis Slides on the Chassis**
- **Putting the Back Rack Rails, Long Brackets, and Short Brackets Together**
- **Installing the Slide Assemblies to the Rack**
- **Installing the Chassis into the Rack**

Rail Kit Contents

The rail kit accompanying your Scale cluster includes:

- **Three pairs of short brackets**, as shown in **Figure 2-1, Short Brackets - Threaded and Unthreaded**. Two pairs are unthreaded, one pair is threaded. It is very likely you will only need one pair. Choose threaded/unthreaded based on what is most appropriate for your rack.

FIGURE 2-1. Short Brackets - Threaded and Unthreaded



-
- **One pair of front chassis slides**, as shown in **Figure 2-2, Front Chassis Slides**. These will be installed on the front of the chassis. These pieces arrive inserted into the long bracket.

FIGURE 2-2. Front Chassis Slides




-
- **One pair of back chassis slides**, as shown in **Figure 2-3, Back Chassis Slides**. These are mounted on the back of the chassis.

FIGURE 2-3. Back Chassis Slides




-
- **One pair of back rack rails**, as shown in **Figure 2-4, Back Rack Rails**. These will be installed on the back of the rack.

FIGURE 2-4. Back Rack Rails




-
- **One pair of long brackets**, as shown in **Figure 2-5, Long Brackets**. These are mounted on the rack.

FIGURE 2-5. Long Brackets




- Eight panhead screws, as shown in **Figure 2-6, Panhead Screw**.

FIGURE 2-6. Panhead Screw



- Eight small flathead screws and eight large flathead screws, as shown in **Figure 2-7, Flat-head Screws**.

FIGURE 2-7. Flathead Screws



- Eight washers.

Putting the Front and Back Chassis Slides on the Chassis


In order to slide your node in and out of the rack, you must add slides to the sides of the chassis. The pieces you need for this task are:

- Front Chassis Slides (2)
- Back Chassis Slides (2)
- Node you want to add slides to
- Panhead screws (6)
- Phillips-head Screwdriver

Removing Front Slides from the Long Brackets

The front chassis slides often arrive inserted into the long brackets. If this is the case, you need to remove them from the long brackets prior to installation on the chassis. To do this, push the front chassis slide forward in the long bracket (towards the end with a lip) until it catches on the front of the long bracket. Push up on the black plastic piece on the front chassis slides to release them from the long brackets, as shown in **Figure 2-8, Removing Front Slides from the Long Brackets**.

FIGURE 2-8. Removing Front Slides from the Long Brackets



Installing Front Slides on the Chassis

To install the front slides on the chassis of your node, do the following:

- 1 Orient one front slide so the black plastic catch is visible when you look at the side of the chassis.

- 2 Align the 6 square holes on one of the front slides against the hooks on one side of the chassis and slide the hooks through the holes on the slide. It should look like the side of the chassis shown in **Figure 2-9, Chassis Hooks**.
- 3 Securely attach the slide to the chassis with one panhead screw. The slide should be against the side so that the loop is at the front of the chassis.
- 4 When you repeat steps 1-3 for the other side of the chassis, the black plastic catch on the other front slide will be oriented so that it is facing upside down when compared to the other slide, as shown in **Figure 2-10, Orientation of Front Slides on Chassis**.

FIGURE 2-9. Chassis Hooks





FIGURE 2-10. Orientation of Front Slides on Chassis



Installing Back Slides on the Chassis

Installing back slides on the chassis is similar to installing front slides. To install a back slide on the chassis, do the following:

- 1 Orient the back slide so the black plastic catch faces the same direction as it does on the front slide. If you have not already installed the front slides do so before continuing.
- 2 Line up the 4 rectangular holes on the back slide with the hooks on the side of the chassis and slide it in place. Ensure that the black plastic catch faces out from the chassis so you can see it.
- 3 If you have installed the front and back slides correctly, they will appear as shown in **Figure 2-11, Front and Back Slides on Chassis** and **Figure 2-12, Chassis Sides with Front and Back Slides Installed, Side Views**.

FIGURE 2-11. Front and Back Slides on Chassis





FIGURE 2-12. Chassis Sides with Front and Back Slides Installed, Side Views



Putting the Back Rack Rails, Long Brackets, and Short Brackets Together


The back rack rails and the long bracket combine to create an outer slide assembly that matches the length of your rack. For this task you need:

- Both back rack rails
- Both long brackets
- Two short brackets
- Two panhead screws
- Phillips head screwdriver
- Measuring tape (optional)

When the back rack rail, long brackets, and short brackets are all pieced together, they are referred to as the outer slide assembly. Do the following to put together the back rack rail, long brackets, and short brackets:

- 1 Slide the long bracket into the back rack rail as shown in **Figure 2-13, Sliding the Long Bracket into the Back Rack Rail**. Ensure the long bracket is oriented so the word **Front** and the arrow next to it are pointing away from the back rack rail.


FIGURE 2-13. Sliding the Long Bracket into the Back Rack Rail



- 2 Slide the long bracket until it is flush against the end of the back rack rail. This enables you to push the peg on the back rack rail into the groove on the long bracket.
- 3 Arrange the back rack rail so it fits snugly against one of the rack columns. The grooves in the back rack rail and the long bracket should be pointing in towards the rack columns on the other side.
- 4 Take one of the short brackets and position it so it is flush against the backside of the long bracket.

- 5 Slide the short bracket until it is flush against the back of the front rack column, as shown in **Figure 2-14, Positioning Short Bracket for Rack.**

FIGURE 2-14. Positioning Short Bracket for Rack




- 6 When you have correctly sized the rail and brackets, secure them using panhead screws.
- 7 Secure the short bracket to the long bracket using one panhead screw through the groove in the short bracket and through the threaded hole. Attach the short bracket so that the head of the panhead screw is on the backside of the long bracket.
- 8 Secure the long bracket inside the back rack rail with a panhead screw using one of the threaded holes on the grooved side of the back rack rail.
- 9 Now attach the combined long bracket, back rack rail, and short bracket (outer slide assembly) to the rack columns using flathead screws. It is recommended that you use two screws for each.
- 10 Repeat these steps for the other side.

Installing the Chassis into the Rack

Push the front and back slides, which are attached to the chassis, into the grooves of the outer slide assemblies installed in the rack, as shown in **Figure 2-15, Sliding the Chassis into Place**.

FIGURE 2-15. Sliding the Chassis into Place



Push the chassis all the way to the back of the outer slide assemblies. Repeat these racking steps for additional nodes.


Cabling the Nodes

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.

Connect network cables to the nodes. Each node has at least two GigE ports labeled **LAN** and **Backplane**, as shown in **Figure 2-16, Backplane and LAN Ports**. If you are using a high availability (HA) node, such as an R1, R2, or R4, these models offer two places to plug-in power sources, and two **LAN** and two **Backplane** ports. Be sure to maximize redundancy for power by plugging in to two separate power sources. Maximize redundancy for your ports by plugging into different switches (up to four different switches).

FIGURE 2-16. Backplane and LAN Ports



Connect the yellow (**Backplane**) port and cables to its own dedicated GigE switch or VLAN.

Connect the blue (**LAN**) port and cables to the same switch or network that any network devices use to connect to the Scale 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 must 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 also works 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 Additional Nodes**
- **Initialize the Cluster**

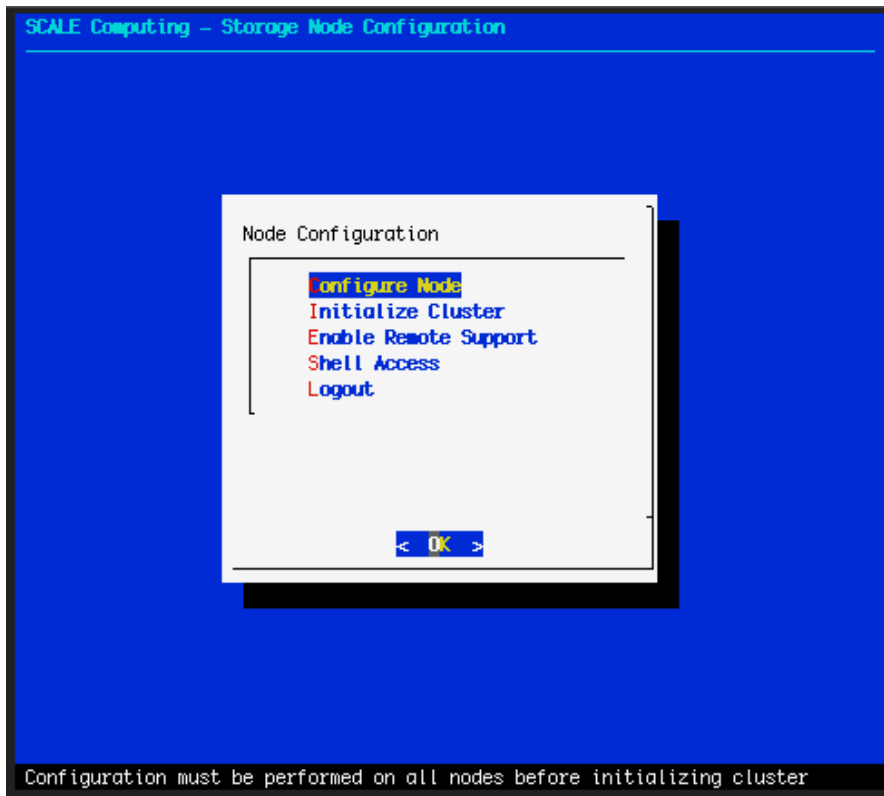
Configure the First Node in the Cluster

Log into the first node with the following credentials:

- user: admin
- password: scale


You are presented with the graphical menu displayed in **Figure 2-17, Node Console Menu**:

FIGURE 2-17. Node Console Menu




Select the first option in the menu: **Configure Node**. You are 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 (Backplane IP of 1st Node) blank as shown in **Figure 2-18, Configure Node**.

FIGURE 2-18. Configure Node



Press **tab** to cycle to the **Apply** button and hit **return**. The node performs initial network configuration and runs internal checks to verify it is able to communicate with the network. You see the screen displayed in **Figure 2-19, Applying the Configuration** while it is performing these checks:

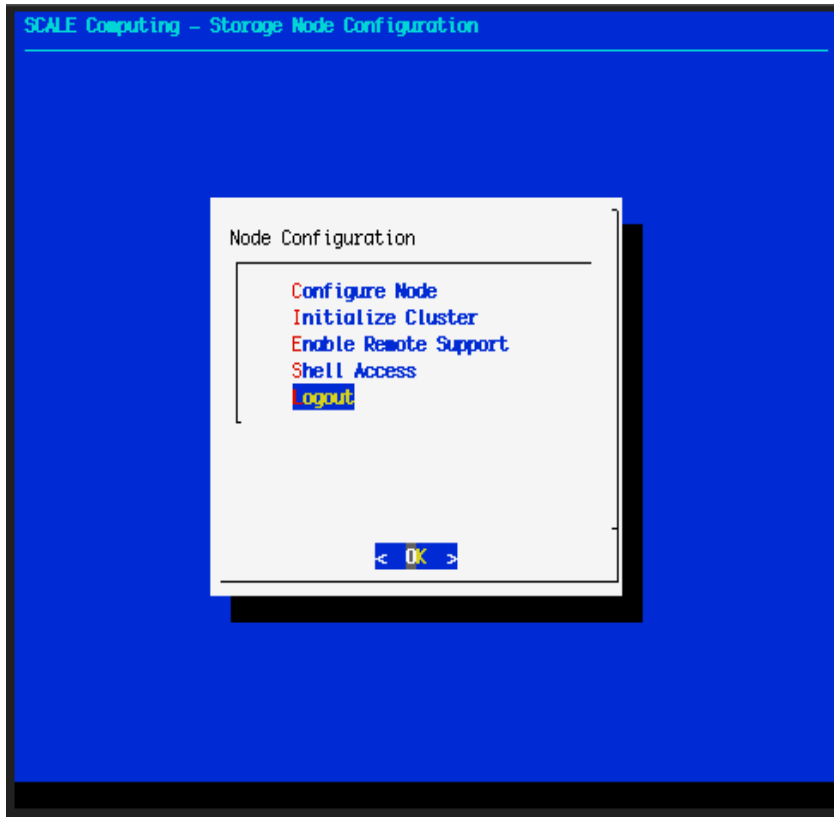
FIGURE 2-19. Applying the Configuration



If there are any errors in the configuration you are 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 completes you return to the main menu as shown in **Figure 2-20, Logging Out of the Node Console Menu**. You can now log out of the administrative interface.

FIGURE 2-20. 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 backplane IP address of the 1st node in the cluster as shown in **Figure 2-21, Configuring Additional Nodes**.

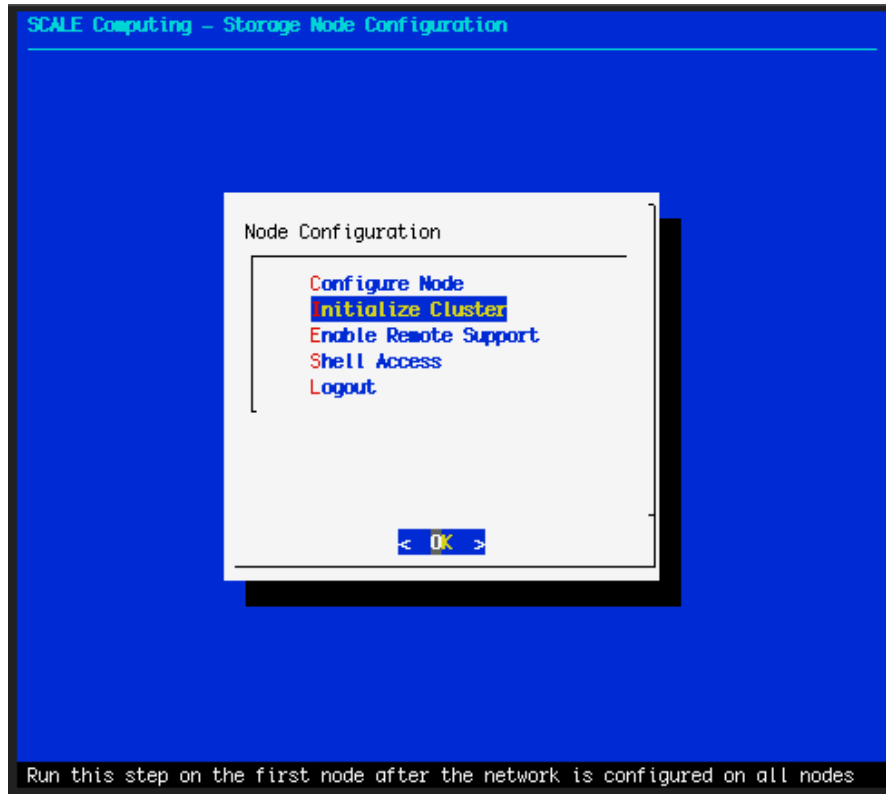
FIGURE 2-21. Configuring Additional Nodes

After applying the settings, the configuration software verifies network settings and communicates 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 LAN and backplane 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 as shown in [Figure 2-22, Initializing the Cluster](#).

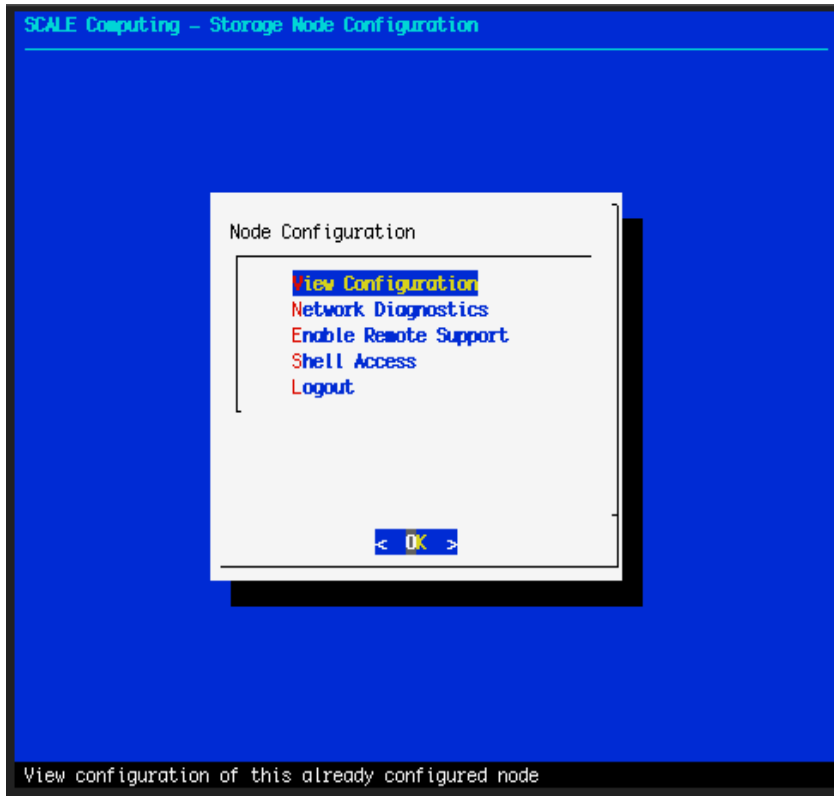
FIGURE 2-22. Initializing the Cluster



This performs the whole-cluster initialization, and you 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 return to the main menu. The menu options are slightly different 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 as shown in **Figure 2-23, A Configured Node**.

FIGURE 2-23. A Configured Node



Completing Cluster Configuration

Once you complete 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.

Revision History

This section contains information describing how this chapter has been revised.

Release 2.3.1:

- Updated the **Racking the Nodes** section, adding additional sections and images.

