



Feature Description LoadMaster Clustering

8 January 2024

Copyright

Visit the following page online to see Progress Software Corporation's current Product Documentation Copyright Notice/Trademark Legend: [Product Documentation Copyright Notice & Trademarks | Progress](#)

Table of Contents

Chapter 1: Introduction.	4
Document Purpose.	5
Intended Audience.	5
 Chapter 2: LoadMaster Clustering.	 6
Prerequisites.	6
Restrictions.	7
Clustering Overview.	8
How it Works.	9
Configuring Clustering.	9
Create a New Cluster and Add Nodes to it.	9
Remove a Node LoadMaster from the Cluster.	14
Changing the Cluster Parameters.	15
Backing Up and Restoring the Configuration.	16
Cluster-Wide Software Updates.	17
Clustering WUI Options.	18
General WUI Options.	19
Cluster Control WUI Options.	20
Cluster Interface Options.	23
Cluster-Wide Software Update WUI Options.	24
Application Program Interface (API) Commands.	24
Troubleshooting.	24
 Chapter 3: References.	 26

Introduction

Introduction

Clustering involves deployment of multiple LoadMaster instances (nodes) as a single management and control domain where each node is actively passing traffic. If any of these nodes fail - the remaining nodes continue to provide service to the remote client, providing resilience.

Clustering supports scaling - LoadMaster nodes can be added to or removed from a cluster so that capacity can be dynamically adjusted to match business requirements. Each additional LoadMaster node which is added to the cluster can provide a linear increase in throughput and performance. The maximum number of nodes that can be added to a cluster is defined by your license agreement.

An example use case is - if a website experiences higher than usual traffic during a holiday period and the amount of traffic cannot be predicted in advance - often the capacity is over-provisioned to handle these requests that come in spikes. One of the simplest solutions is to deploy LoadMasters in a cluster and then add new nodes as the capacity requirements (concurrent connections and performance requirements) increase.

Related Links

- [Document Purpose](#)
- [Intended Audience](#)

Document Purpose

Document Purpose

The purpose of this document is to provide information on the LoadMaster clustering feature and step-by-step instructions on how to set up and configure clustering.

Intended Audience

Intended Audience

This document is intended to be read by anyone who is interested in finding out more about LoadMaster clustering.

LoadMaster Clustering

LoadMaster Clustering

Refer to the following sections for details on LoadMaster clustering.

Related Links

- [Prerequisites](#)
- [Restrictions](#)
- [Clustering Overview](#)
- [Configuring Clustering](#)
- [Cluster-Wide Software Updates](#)
- [Clustering WUI Options](#)
- [Application Program Interface \(API\) Commands](#)
- [Troubleshooting](#)

Prerequisites

Prerequisites

There are some prerequisites to aware of before clustering LoadMasters:

- A minimum of three LoadMasters are needed to set up clustering. However, four LoadMasters are recommended.

- All of the LoadMasters in the cluster must be physically identical in terms of the hardware used. Clustering also works with Virtual LoadMasters but the Virtual Machines must have the same resources assigned; this includes CPU, memory and Hard Disk Drive (HDD). Once an empty machine is brought into the cluster, all parameters from the other machines (network interfaces, and so on), are copied over onto the new machine. Local addresses for network interfaces must then be set.
- Requirements for clustering in a virtual environment (VMware, Hyper-V, and so on) are the same as those for High Availability (HA). In these environments:
 -
 - Ensure that Media Access Control (MAC) spoofing is allowed and enabled
 - Ensure that multicast traffic flow is working in both directions between the devices
- A 2-arm configuration is required.

Restrictions

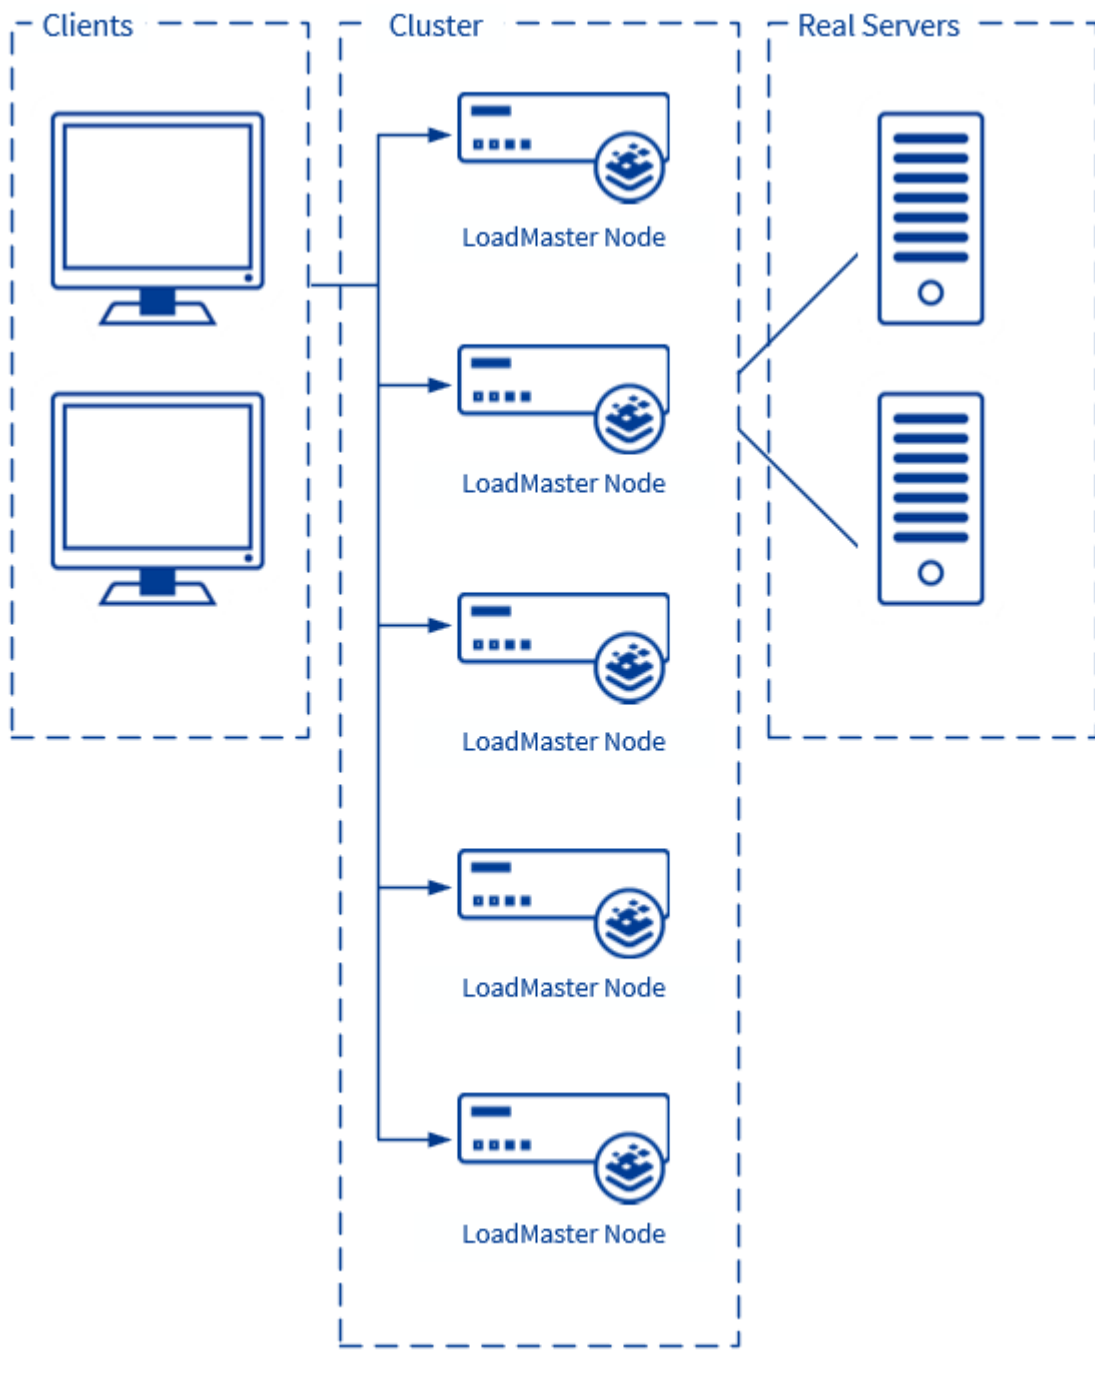
Restrictions

Most functionality that is available in a single LoadMaster is available when a LoadMaster environment is clustered. The only restrictions are listed below:

- All Virtual Services must operate at Layer 7
- Transparency cannot be enabled
- Subnet originating addresses are always enabled
- Non-local Real Servers are not supported
- HA mode cannot be used
- VXLAN and IPsec tunneling are not supported
- A 2-arm configuration is required
- It is only possible to have a Virtual Service on the primary interface. Virtual Services on a second interface do not work.
- GEO does not currently work properly when configured on a LoadMaster in clustering mode (as opposed to High Availability (HA) mode). If you want to use GEO with LoadMasters configured into clustering mode, GEO must be running as a standalone GEO Virtual Machine or configured on a LoadMaster that is not in clustering mode.

Clustering Overview

Clustering Overview



With LoadMaster clustering, the load-balancing capability can be extended as needed by adding additional LoadMasters for the same Virtual Service. Traffic which is intended for a failed node is either passed on to an existing node or load-balanced across the remaining nodes.

The LoadMasters all work in parallel. If a LoadMaster node fails, traffic is redistributed amongst the remaining nodes. Health checking is performed on each of the LoadMasters.

Most administrative changes should be made on the shared IP address interface. These changes are then propagated to the LoadMasters in less than a second. Local administration of the individual LoadMasters can also be performed by accessing the WUI of the units directly. The WUI of the local LoadMasters contains limited configuration options.

Related Links

- [How it Works](#)

How it Works

How it Works

When traffic is forwarded to the cluster, the admin node uses source-based persistency to pass the connections among the child nodes. When the child node receives the packet, it processes the requests on the Virtual Services and sends those requests to the back-end Real Servers. Once the response is received from back-end Real Servers, the child node forwards the received response directly back to the client. This is done to efficiently allow for resource pooling of all nodes in the cluster. For this reason, you cannot configure transparency.

Configuring Clustering

Configuring Clustering

If clustering is enabled on the LoadMaster, the **Cluster Control** menu item is available under **System Configuration** in the main menu of the LoadMaster WUI. If clustering functionality is not enabled and you would like to use it, please contact your Progress Kemp representative.

Related Links

- [Create a New Cluster and Add Nodes to it](#)
- [Remove a Node LoadMaster from the Cluster](#)
- [Changing the Cluster Parameters](#)
- [Backing Up and Restoring the Configuration](#)

Create a New Cluster and Add Nodes to it

Create a New Cluster and Add Nodes to it

To set up a cluster and add LoadMaster nodes to it, follow the steps in the sections below, in order.

Create a New Cluster

To create a new cluster, follow the steps below on a LoadMaster WUI:

1. In the main menu of the LoadMaster WUI, go to **System Configuration > HA and Clustering**.

Confirm

☐ HA Mode

An HA configuration requires two LoadMasters, only one of which is active and processing traffic at any time. The other passive unit continuously monitors the health of the active unit and will begin serving traffic when the active unit becomes unavailable. Once you configure HA mode, clustering options will be unavailable.

☐ Clustering

A Clustering configuration requires the following:

1. At least three LoadMasters (four or more are recommended). All LoadMasters in a cluster actively process traffic.
2. All hardware LoadMasters must be the same model. Virtual LoadMasters must have the same CPU, RAM and disk storage assigned. You cannot mix hardware and virtual LoadMasters in a cluster.
3. All LoadMasters should be set to use factory-default settings, with the exception of networking.

Once you configure clustering, HA mode options will be unavailable.

Confirm

Cancel

2. This screen describes both **HA Mode** and **Clustering**. Select **Clustering** and click **Confirm**.

Convert to Cluster

Create a new Cluster

Create New Cluster

Add this LoadMaster to an existing cluster

Add to Cluster

3. Click **Create New Cluster**.

Convert to Cluster

Cluster Shared Address [Create a New Cluster](#)

4. Enter the desired shared IP address for the cluster.

Note: This is the address that is used to access the shared WUI where configuration changes should be made.

5. Click **Create a New Cluster**.

Reboot

Rebooting and switching to the Shared Address to finish the conversion to Cluster mode

Please reconnect to 10.154.11.91

[Continue](#)

6. The LoadMaster will then reboot. A prompt will appear asking to reconnect to the shared IP address.
7. Click **Continue**.

When the LoadMaster has rebooted, you are brought to the shared IP address.

The cluster is now set up. The next step is to add a LoadMaster to the cluster.

Add a Node LoadMaster to the Cluster

To add a node LoadMaster to the cluster, the node LoadMaster needs to become available and then have its IP address added in the shared IP address WUI. To do this, follow the steps below:

1. In the node WUI, go to **System Configuration > Cluster Control** in the main menu.

Convert to Cluster

[Create a new Cluster](#) [Create New Cluster](#)
[Add this LoadMaster to an existing cluster](#) [Add to Cluster](#)

2. Click **Add to Cluster**.

Note: The LoadMaster is now available to be added to the cluster.

Waiting to Join Cluster. Please Wait...



- 3. Open the WUI of the shared IP address.
- 4. In the main menu, go to **System Configuration > Cluster Control**.

Current Cluster Configuration

ID Address	Status	Operation
1 10.154.11.90	Admin	<div>DisableDelete</div>

IP Address10.154.0.0Add New Node

- 5. Enter the IP address of the node LoadMaster in the **Address** text box.
- 6. Click **Add New Node**.

Rebooting to finish the conversion to cluster mode

Continue

Note: This must be done while the node LoadMaster is available to join the cluster.

- 7. The node LoadMaster needs to reboot to finish the conversion to cluster mode. Click **Continue**.

Current Cluster Configuration

ID Address	Status	Operation
1 10.154.11.90	Admin	<div>DisableDelete</div>
2 10.154.11.80	Disabled	<div>EnableDelete</div>

IP Address10.154.0.0Add New Node

- 8. By default, the new node is disabled. Click **Enable** to enable the node in the cluster.

Note: When a node is enabled, it will not be immediately brought into rotation. It will only come online after it has been up for 30 seconds. The page may need to be refreshed to display any state changes. If there are no Virtual Services in the node, the node is in a **Disabled** state.

9. Click **OK**.

The cluster is now configured and a node has been added. To add more nodes, repeat the steps in the Add a Node LoadMaster to the Cluster section.

When all nodes have been added to the cluster, other configuration changes can be made - such as creating and configuring Virtual Services. These changes are propagated from the shared IP address to the LoadMaster nodes.

Configure the Network Interface(s)

Some configuration changes are needed on the interfaces. Refer to the sections below for further details.

Set the Cluster Shared IP Address

The cluster shared IP address needs to be set on any interfaces that are used to handle traffic. To do this, follow the steps below:

1. In the main menu of the shared IP address LoadMaster WUI, go to **System Configuration > Network Setup**.
2. Click the relevant interface link.

Network Interface 0

Interface Address (address/prefix)	10.154.11.70/16	Set Address
Cluster Shared IP address	10.154.11.90	Set Shared address
Use for Cluster checks	<input checked="" type="checkbox"/>	
Use for Cluster Updates	<input checked="" type="checkbox"/>	
Use for GEO Responses and Requests	<input checked="" type="checkbox"/>	
Link Status	Speed: 10000Mb/s, Full Duplex	Automatic <input type="button" value="Force Link"/>
MTU:	1500	Set MTU
Additional addresses (address/prefix)		Add Address

[VLAN Configuration](#)
[Interface Bonding](#)

3. Enter the **Cluster Shared IP address** and click **Set Shared address**.

Note: GEO does not currently work properly when configured on a LoadMaster in clustering mode (as opposed to High Availability (HA) mode). If you want to use GEO with LoadMasters configured into clustering mode, GEO must be running as a standalone GEO Virtual Machine or configured on a LoadMaster that is not in clustering mode.

Enable the Use for Cluster Checks Option

The **Use for Cluster checks** option must be enabled on at least one interface. To enable this option, follow the steps below:

1. In the main menu of the LoadMaster WUI, expand **System Configuration**.
2. Click the relevant interface.

Network Interface 0

Interface Address (address/prefix)	10.154.11.70/16	Set Address
Cluster Shared IP address	10.154.11.90	Set Shared address
Use for Cluster checks	<input checked="" type="checkbox"/>	
Use for Cluster Updates	<input checked="" type="checkbox"/>	
Use for GEO Responses and Requests	<input checked="" type="checkbox"/>	
Link Status	Speed: 10000Mb/s, Full Duplex	Automatic <input type="button" value="Force Link"/>
MTU:	1500	Set MTU
Additional addresses (address/prefix)		Add Address

VLAN Configuration Interface Bonding

3. Tick the **Use for Cluster Checks** check box.

Enable the Use for Default Gateway Option

The **Use for Default Gateway** option needs to be enabled on each of the LoadMasters.

To do this, follow the steps below:

1. In the main menu of a node LoadMaster, expand **System Configuration**.
2. Click the relevant interface.

Network Interface 0

Interface Address (address/prefix)	10.154.11.70/16	Set Address
Cluster Shared IP address	10.154.11.90	Set Shared address
Use for Cluster checks	<input checked="" type="checkbox"/>	
Use for Cluster Updates	<input checked="" type="checkbox"/>	
Use for GEO Responses and Requests	<input checked="" type="checkbox"/>	
Link Status	Speed: 10000Mb/s, Full Duplex	Automatic <input type="button" value="Force Link"/>
MTU:	1500	Set MTU
Additional addresses (address/prefix)		Add Address

VLAN Configuration Interface Bonding



3. Tick the **Use for Default Gateway** check box.

Remove a Node LoadMaster from the Cluster

Remove a Node LoadMaster from the Cluster

It is possible to remove a node LoadMaster from the cluster which will mark the node as unavailable for new traffic (that is, a Draining state) to facilitate the removal from the cluster. To do this, follow the steps below in the shared IP address WUI:

1. In the main menu, go to **System Configuration > Cluster Control**.

ID	Address	Status	Operation
1	10.154.11.90	 Admin	<button>Disable</button> <button>Delete</button>
2	10.154.11.80	 Up	<button>Disable</button> <button>Delete</button>

- Click **Disable** on the node to be removed.

Note: Before a node can be deleted, it must first be disabled. A node cannot be deleted if it is handling traffic.

- The node will change to a **Draining** state. Wait for the node to finish draining and change to a **Disabled** state.

Note: The page will need to be refreshed to see the new state.

Note: When a node is in the **Draining** state, the connections that are still being served by the node are allowed to continue for the amount of time specified in the **Node Drain Time** text box. For more information, refer to the [Cluster Parameters](#) section. No new connections are handled by the node during this time.

- Click **Delete** to remove the node from the cluster.
- Click **OK** to confirm the removal.
- Click **OK** to the confirmation message.

When a node is deleted it becomes a regular single LoadMaster instance. If the LoadMaster is later added back in to the cluster, any configuration changes that have been made in the shared IP address will propagate to the node LoadMaster.

Changing the Cluster Parameters

Changing the Cluster Parameters

The cluster Virtual ID and node drain time can be updated, if needed. Refer to the sections below for further information.

Updating the Cluster Virtual ID

When using multiple clusters or LoadMaster HA systems on the same network, the virtual ID identifies each cluster so that there are no potential unwanted interactions. To change the cluster ID, follow the steps below in the shared IP address WUI:

- In the main menu, go to **System Configuration > Cluster Control**.
- Click the **Show Options** button.

Cluster Parameters

Cluster Virtual ID	<input type="text" value="2"/>	Set Cluster Virtual ID (Valid Values: 1-255)
Node Drain Time	<input type="text" value="10"/>	Set Node Drain Time (Valid Values: 1-600)

1. Change the **Cluster Virtual ID** value.

Note: Valid values range from 1 to 255.

1. Click **Set Cluster Virtual ID**.

Updating the Node Drain Time

When a node is disabled, the connections that are still being served by the node are allowed to continue for the amount of seconds specified in the **Node Drain Time** text box. No new connections are handled by the node during this time. The **Node Drain Time** is set to **10** seconds by default. To change the **Node Drain Time**, follow the steps below in the shared IP address WUI:

1. In the main menu, go to **System Configuration > Cluster Control**.
2. Click **Show Options**.

Cluster Parameters

Cluster Virtual ID	<input type="text" value="1"/>	Set Cluster Virtual ID (Valid Values: 1-255)
Node Drain Time	<input type="text" value="20"/>	Set Node Drain Time (Valid Values: 1-600)

1. Change the value in the **Node Drain Time** text box.

Note: Valid values range between 1 and 600 (seconds).

1. Click **Set Node Drain Time**.

Backing Up and Restoring the Configuration

Backing Up and Restoring the Configuration

Backups can be taken and restored as usual in the **System Configuration > System Administration > Backup/Restore** section of the LoadMaster WUI. Backups and restores should only be performed on the shared IP address because the changes made there are propagated to the other nodes. When restoring, the individual LoadMasters are rebooted, one at a time.

There are certain settings that need to be set in the Virtual Services for them to work with clustering:

- Transparency must be disabled

- Subnet Originating Requests must be enabled
- All Virtual Services must operate at Layer 7.
- High Availability (HA) mode cannot be used

As a result of these restrictions, there may be issues with restoring a backup taken from a non-cluster LoadMaster on to a cluster LoadMaster and vice versa, because the settings may be incompatible.

When restoring a configuration, you specify what parts of the configuration should be restored:

- The LoadMaster base configuration only
- The Virtual Service configuration only
- The GEO configuration only
- The ESP SSO configuration only
- A combination of the Virtual Service, GEO, ESP, and/or LoadMaster base configuration

The base configuration contains the information about the basic configuration of the LoadMaster, that is, the IP addresses of the various interfaces and the keyboard and time zone settings.

The Virtual Service configuration contains only the settings relating to the Virtual Services and the Real Servers.

The GEO configuration contains only the settings relating to the GEO configuration.

Note: GEO does not currently work properly when configured on a LoadMaster in clustering mode (as opposed to High Availability (HA) mode). If you want to use GEO with LoadMasters configured into clustering mode, GEO must be running as a standalone GEO Virtual Machine or configured on a LoadMaster that is not in clustering mode.

The ESP SSO configuration stores the SSO domains, LDAP endpoints and SSO custom image sets. This does not restore the Virtual Service settings - use the **VS Configuration** option to restore those.

Cluster-Wide Software Updates

Cluster-Wide Software Updates

The firmware on all LoadMasters in a cluster can be updated using the shared IP address.

Note: Software updates require the LoadMasters to be rebooted. Functionality is maintained but existing connections will be lost for a short period of time.

To perform a cluster-wide software update, follow the steps below in the WUI of the shared IP address:

1. Get a LoadMaster firmware patch file.

Note: You can download firmware patches from the [Downloads page](#).

2. In the main menu, go to **System Configuration > Update Software**.

Update LoadMaster Software

Software Update File: No file chosen

3. Click **Choose File**.
4. Browse to and select the firmware update patch file.
5. Click **Update Cluster**.
6. Click **OK** to continue.

Transferring and Validating Software Update. Please Wait...



7. Wait for the software update to be transferred and validated.
8. Click **OK** to install the software update.
9. A message will appear saying that the new software has been installed and the **Cluster Control** screen is displayed.
10. Each of the non-admin nodes need to be rebooted in order to activate the software. Click **Reboot** to reboot a node.
11. Click **OK** to the pop-up message.
12. When all non-admin nodes have been rebooted, the admin node needs to be rebooted. Click **Reboot** in the admin node row to reboot the admin node.

Clustering WUI Options

Clustering WUI Options

Refer to the following sections for details on the clustering options in the LoadMaster WUI.

Related Links

- [General WUI Options](#)
- [Cluster Control WUI Options](#)
- [Cluster Interface Options](#)
- [Cluster-Wide Software Update WUI Options](#)

General WUI Options

General WUI Options

- ▼ System Configuration
 - Network Setup
 - Cluster Control
 - System Administration
 - Logging Options
 - Miscellaneous Options

If clustering functionality is enabled on a non-node LoadMaster, a **Cluster Control** menu option should be available in the WUI within the **System Configuration** section.

Convert to Cluster

Create a new Cluster	Create New Cluster
Add this LoadMaster to an existing cluster	Add to Cluster

Before setting up clustering, clicking the **Cluster Control** menu item will give the option to either create a new cluster or add this LoadMaster to a cluster.

Local Home

▼ Local Administration

- ▼ Interfaces
 - eth0 ▸ eth1
 - Virtual LAN ▼
- Host & DNS Configuration
- User Management
- Default Gateway
- Update License
- System Reboot
- Update Software
- Backup/Restore
- Date/Time
- WUI Settings
- Log Files
- Extended Log Files
- Backup/Restore Certs.
- Statistics
- Historical Stats

Home

- Virtual Services
- Global Balancing
- Statistics
- Real Servers
- Rules & Checking
- Certificates
- System Configuration

When logging in to the cluster, use the shared IP address to view and set the full functionality of the cluster nodes. Logging in to the local IP address of any of the nodes displays different menu options (see menus above). Logging in to one of the nodes directly is usually reserved for maintenance.

Cluster Control WUI Options

Cluster Control WUI Options

Convert to Cluster

Create a new Cluster	Create New Cluster
Add this LoadMaster to an existing cluster	Add to Cluster

Create New Cluster: If setting up a new cluster, click this button.

Add to Cluster: Add this LoadMaster to an already existing cluster.

Convert to Cluster

Cluster Shared Address	10.154.11.90	Set Shared Address
------------------------	--------------	--------------------

Cancel Create a Cluster

When the **Create New Cluster** button is clicked, the screen above will appear which prompts to set the shared IP address of the cluster. The shared IP address is the address which is used to administer the cluster.


Rebooting and switching to the Shared Address to finish the conversion to Cluster mode

Please reconnect to 10.154.11.90

Continue

When the **Create a Cluster** button is clicked, the LoadMaster reboots. A message will appear asking to reconnect to the shared IP address that was just set.



Current Cluster Configuration

ID Address	Status	Operation
1 10.154.11.90	 Admin	<button>Disable</button> <button>Delete</button>

IP Address Add New Node

After creating a cluster, the **Cluster Control** screen in the WUI of the shared IP address will allow the addition of LoadMaster nodes into the cluster.

Note: A LoadMaster can only be added to a cluster when the cluster is available and waiting to join the cluster. Refer to the [Add a Node LoadMaster to the Cluster](#) section for further information and steps.

ID Address	Status	Operation
1 10.154.11.90	 Admin	<button>Disable</button> <button>Delete</button>
2 10.154.11.80	 Up	<button>Disable</button> <button>Delete</button>

The **Cluster Control** screen, in the shared IP address WUI, displays details for each of the nodes in the cluster.

Show Options: Clicking the **Show Options** button will display the **Cluster Parameters** section which contains two additional fields which can be used to set the **Cluster Virtual ID** and **Node Drain Time**. For further information, refer to the Cluster Parameters section.

ID: The cluster ID.

Address: The IP address of the LoadMaster node. If a second IP address appears in brackets after the first one - the second IP address is the IP address of the interface port. The IP address and status text is coloured depending on the status:

- Blue: The node is the admin node.
- Grey: The node is disabled.
- Green and red: The node is starting.
- Green: The node is up.
- Red: The node is down.

Status: The status of the node. The possible statuses are:

- **Admin:** The node is the primary control node.
- **Up:** The node is up.

- **Down:** The node is down.
- **Drain stopping:** The node has been disabled and the connections are being shut down in an orderly fashion. Drain stopping lasts for 10 seconds by default. This can be updated by changing the **Node Drain Time** value on the **Cluster Control** screen. For more information, refer to the [Updating the Node Drain Time](#) section.
- **Starting:** The node is starting
- **Disabled:** The node is disabled - connections will not be sent to that node. If there are no Virtual Services in the node, the node is in a **Disabled** state.

Operation: The different operations that can be performed in relation to the nodes:

- **Add new node:** Add a new node with the specified IP address to the cluster.
- **Disable:** Disable the node. Nodes that are disabled will first go through drain stopping. During the drain stopping time, the connections are shut down in an orderly fashion. After the drain, the node is disabled and no traffic is directed to that node.
- **Enable:** Enable the node. When a node comes up, it will not be immediately brought into rotation. It will only come online after it has been up for 30 seconds.
- **Delete:** Delete a node from the cluster. When a node is deleted it becomes a regular single LoadMaster instance. If the LoadMaster is later added back in to the cluster, any configuration changes that have been made in the shared IP address will propagate to the node LoadMaster.

Cluster Parameters

Cluster Parameters

Cluster Virtual ID	<input type="text" value="1"/>	Set Cluster Virtual ID (Valid Values: 1-255)
Node Drain Time	<input type="text" value="10"/>	Set Node Drain Time (Valid Values: 1-600)

When the **Show Options** button is clicked, the **Cluster Parameters** section appears. This section contains two additional WUI options - **Cluster Virtual ID** and **Node Drain Time**.

Cluster Virtual ID

When using multiple clusters or LoadMaster HA systems on the same network, the virtual ID identifies each cluster so that there are no potential unwanted interactions. The cluster virtual ID is set to **1** by default, but it can be changed if required. Valid IDs range from 1 to 255. Changes made to an admin Loadmaster propagate across all nodes in the cluster.

Node Drain Time

When a node is disabled, the connections that are still being served by the node are allowed to continue for the amount of seconds specified in the **Node Drain Time** text box. No new connections are handled by the node during this time. The **Node Drain Time** is set to **10** seconds by default, but it can be changed if required. Valid values range from 1 to 600 (seconds).

During the drain time the status changes to Draining until the specified drain time elapses.

When the drain time has elapsed the status changes to disabled.

Cluster Interface Options

Cluster Interface Options

There are a couple of fields relating to clustering in the **Network Interface** management screen.

Network Interface 0

Interface Address (address[/prefix])	<input type="text" value="10.154.11.70/16"/>	Set Address
Cluster Shared IP address	<input type="text" value="10.154.11.90"/>	Set Shared address
Use for Cluster checks	<input checked="" type="checkbox"/>	
Use for Cluster Updates	<input checked="" type="checkbox"/>	
Use for GEO Responses and Requests	<input checked="" type="checkbox"/>	
Link Status	Speed: 10000Mb/s, Full Duplex	<input type="text" value="Automatic"/> Force Link
	MTU: <input type="text" value="1500"/>	Set MTU
Additional addresses (address[/prefix])	<input type="text"/>	Add Address

[VLAN Configuration](#)
[Interface Bonding](#)

To get to this screen, go to **System Configuration > Network Setup** in the main menu of the LoadMaster WUI and click one of the interfaces. The cluster-related options are described below.

Cluster Shared IP address

Specify the shared IP address which can be used to access the cluster. This is also used as the default source address when using Server NAT.

Use for Cluster checks

Use this option to enable cluster health checking between the admin LoadMaster and the nodes. At least one interface must be enabled.

Use for Cluster Updates

Select this option to use this interface to synchronize the configuration with all of the other nodes in the cluster.

Cluster-Wide Software Update WUI Options

Cluster-Wide Software Update WUI Options

Update LoadMaster Software

Software Update File: No file chosen

The firmware on all LoadMasters in a cluster can be updated using the shared IP address.

This can be done in **System Configuration > System Administration > Update Software** by using the **Update Cluster** option.

After the firmware update patch has been uploaded, a new **Reboot** button will appear on the **Cluster Control** screen. When you click **Reboot**, a pop-up message will appear saying that the node is rebooting.

For more information, including step-by-step instructions on how to perform a cluster-wide software update, refer to the [Cluster-Wide Software Updates](#) section.

Application Program Interface (API) Commands

Application Program Interface (API) Commands

Progress Kemp supports APIs that can be used to allow remote applications to access the LoadMaster in a simple and consistent manner:

- RESTful API
- PowerShell API

Clustering can be configured using API commands. For details on each of the commands that can be used, refer to the relevant API Interface Description document on the Progress Kemp documentation page: www.kemptechnologies.com/documentation.

Troubleshooting

Troubleshooting

Logs are recorded to track certain clustering events, such as when a node goes offline. Viewing the logs can help to troubleshoot problems which may occur. To view the logs relating to clustering, in the LoadMaster WUI go to **System Configuration > Logging Options > System Log Files**. The logs relating to clustering are available in the **Warning Message File** and the **System Message File**.

Some examples of log entries relating to clustering are listed and described below.

In the following log example, there was a worker node and admin node. The admin node failed and the worker node became the admin node. There were then no worker nodes left because there was only one previously, which is why the last line says **Cluster has no functioning nodes**:

```
2023-01-24T13:40:06+00:00 LM03 ucarp: [CRITICAL] Partner has failed, becoming
ACTIVE
2023-01-24T13:40:06+00:00 LM03 logger: HA Switching to Active
2023-01-24T13:40:07+00:00 LM03 l4d: Cluster has no functioning nodes.
```

The following log example shows a node going offline. This could be because the node was rebooted.

lb100 l4d: Node 3 now deadlb100 l4d: Cluster running ononly 3 of 4 nodes

The following log example indicates that a node has gone offline.

lb100 l4d: Cluster running on 3 of 4 nodes (1 disabled)

The following log example indicates that a node is alive, for example a disabled node has been enabled.

lb100 l4d: Node 3 now alive

References

References

Unless otherwise specified, the following documents can be found at <https://docs.progress.com/>.

PowerShell, Interface Description

High Availability (HA), Feature Description