

# Advanced: MQTT Modules in Redundant Ignition Environment

## Prerequisites

- Knowledge of Ignition and Module installation process: [Cirrus Link Module Installation](#)

## Summary

The Inductive Automation platform and MQTT modules can be resilient to failures when configured to use redundancy. Redundant Ignition systems can be set up and configured to act as failover backups for primary/master Ignition instances. This tutorial will provide step by step instructions for installing a set of Ignition systems with redundancy on the host/primary Ignition instance as well as redundancy on the MQTT enabled edge nodes. For this tutorial we will show how to set up a total of six Ignition systems. These will be:

- **Ignition Primary**
  - An Ignition system running as master with MQTT Distributor and MQTT Engine installed. This is what remote Ignition systems will send data to in normal operation.
- **Ignition Primary Backup**
  - An Ignition system running as backup with MQTT Distributor and MQTT Engine installed. This is what remote Ignition systems will send data to when Ignition Primary fails.
- **Ignition Edge 1**
  - An Ignition system running as master with MQTT Transmission installed. This will send data to Ignition Primary in normal operation. If Ignition Primary is in a failed state, this will send data to Ignition Primary Backup.
- **Ignition Edge 1 Backup**
  - An Ignition system running as backup with MQTT Transmission installed. This will send data to Ignition Primary in the event that Ignition Edge 1 fails. If Ignition Primary is in a failed state and Ignition Edge 1 is in a failed state, this will send data to Ignition Primary Backup.
- **Ignition Edge 2**
  - An Ignition system running as master with MQTT Transmission installed. This will send data to Ignition Primary in normal operation. If Ignition Primary is in a failed state, this will send data to Ignition Primary Backup.
- **Ignition Edge 2 Backup**
  - An Ignition system running as backup with MQTT Transmission installed. This will send data to Ignition Primary in the event that Ignition Edge 2 fails. If Ignition Primary is in a failed state and Ignition Edge 2 is in a failed state, this will send data to Ignition Primary Backup.

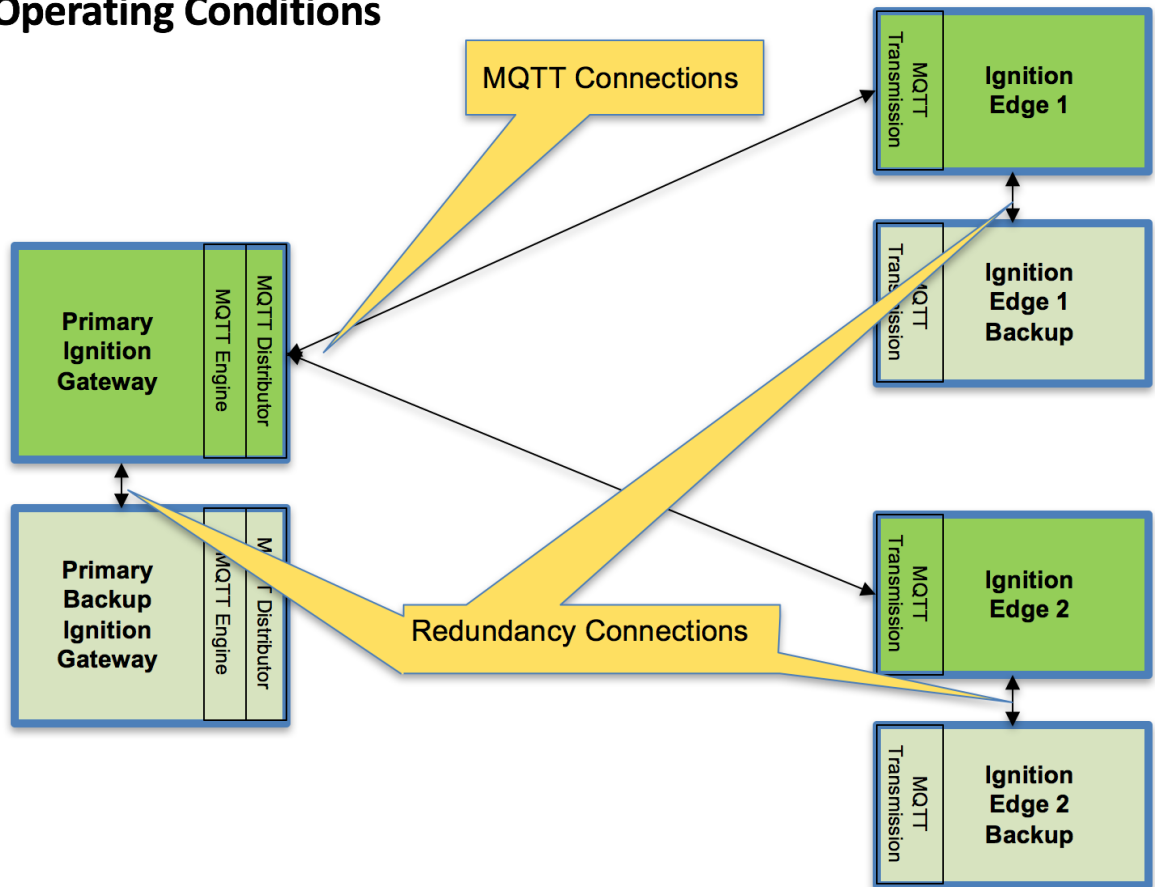
Additional Edge Nodes could be added to this infrastructure. It is also important to note that the Ignition Edge Nodes with MQTT Transmission could also be instances of Ignition Edge MQTT depending on your requirements (<https://inductiveautomation.com/whats-new-ignition-edge>). There are additional considerations when setting up a real world system using redundancy. These topics are not covered in this tutorial but should be taken into consideration.

- **Network paths**
  - It may make sense to have edge nodes support multiple network paths to the MQTT servers. For example, Ethernet, cellular, and satellite could all exist as supported network paths on a single Edge gateway. This will help ensure additional reliability by supporting failover of networks.
- **Primary Ignition and Primary Ignition backup placement**
  - This tutorial was created by modeling this exact environment using Amazon AWS EC2 instances in the cloud. Reliability could be improved by putting Ignition Primary and Ignition Primary Backup in different AWS availability zones or even different AWS regions. This would allow the primary Ignition with MQTT Distributor and MQTT Engine to continue to operate even in the case of AWS failures. In the case of on premises installations of Ignition these could be placed in different physical locations and/or on secondary networks.
- **MQTT Servers**
  - Additional MQTT Servers can be added so MQTT connections from remote edge nodes remain established. Additional Chariot MQTT Servers can be used to make the system more robust (<https://www.cirrus-link.com/iiot-mqtt-servers/>).
- **History enablement in MQTT Transmission**
  - MQTT Transmission supports caching of data in the case that it can not establish a connection to any of the configured MQTT Servers. Once a connection is reestablished, it will begin reporting and flush the stored historical values to prevent data loss in catastrophic failures.

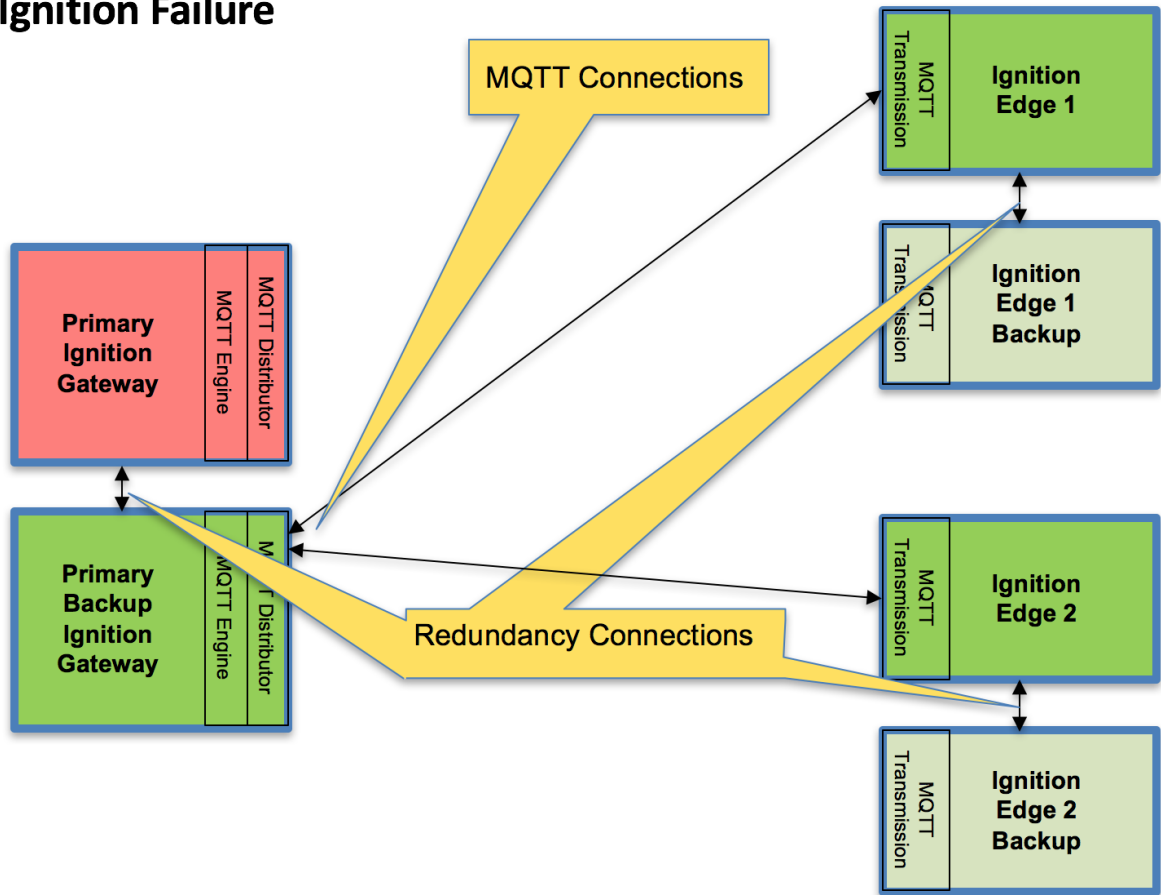
Upon completion of this tutorial you will have a functional system with redundancy/failover support for both remote edge nodes as well as the primary Ignition system that the remote edge nodes are reporting to.

## Architecture

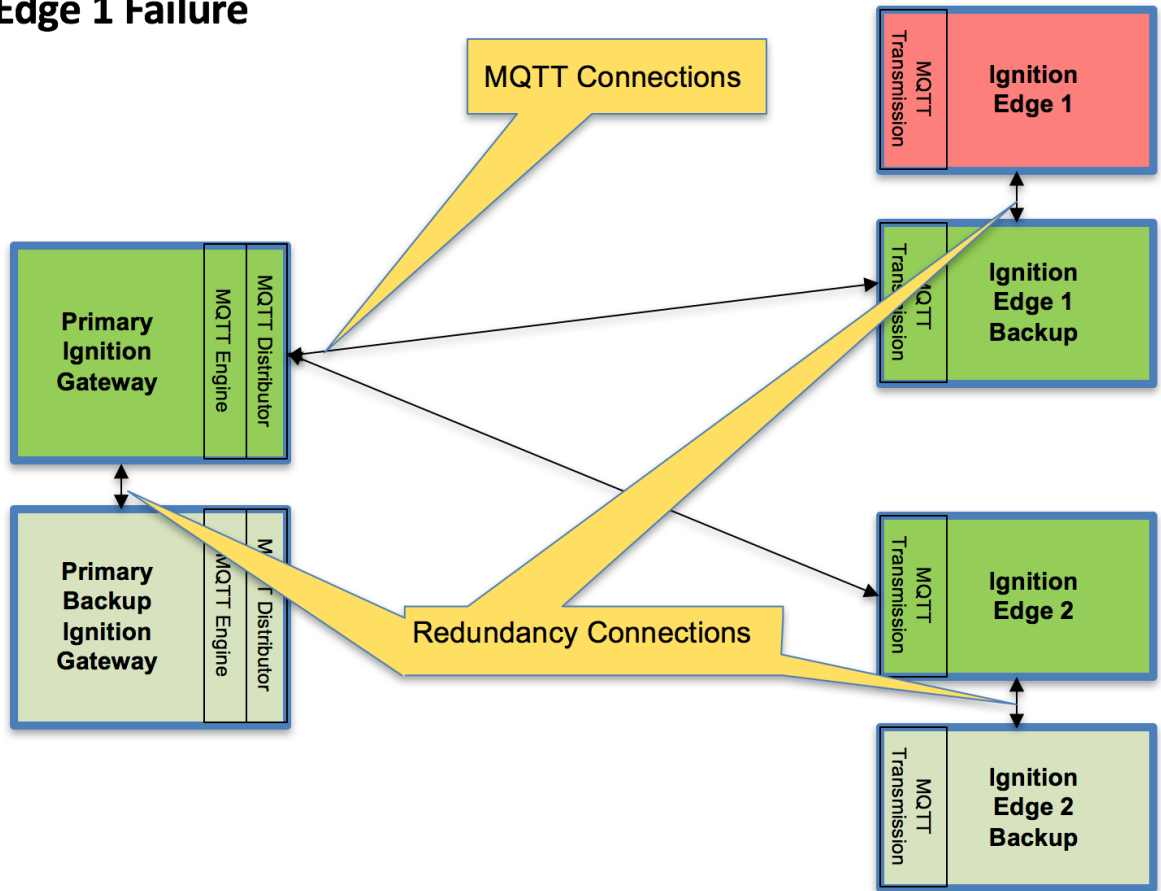
## Normal Operating Conditions



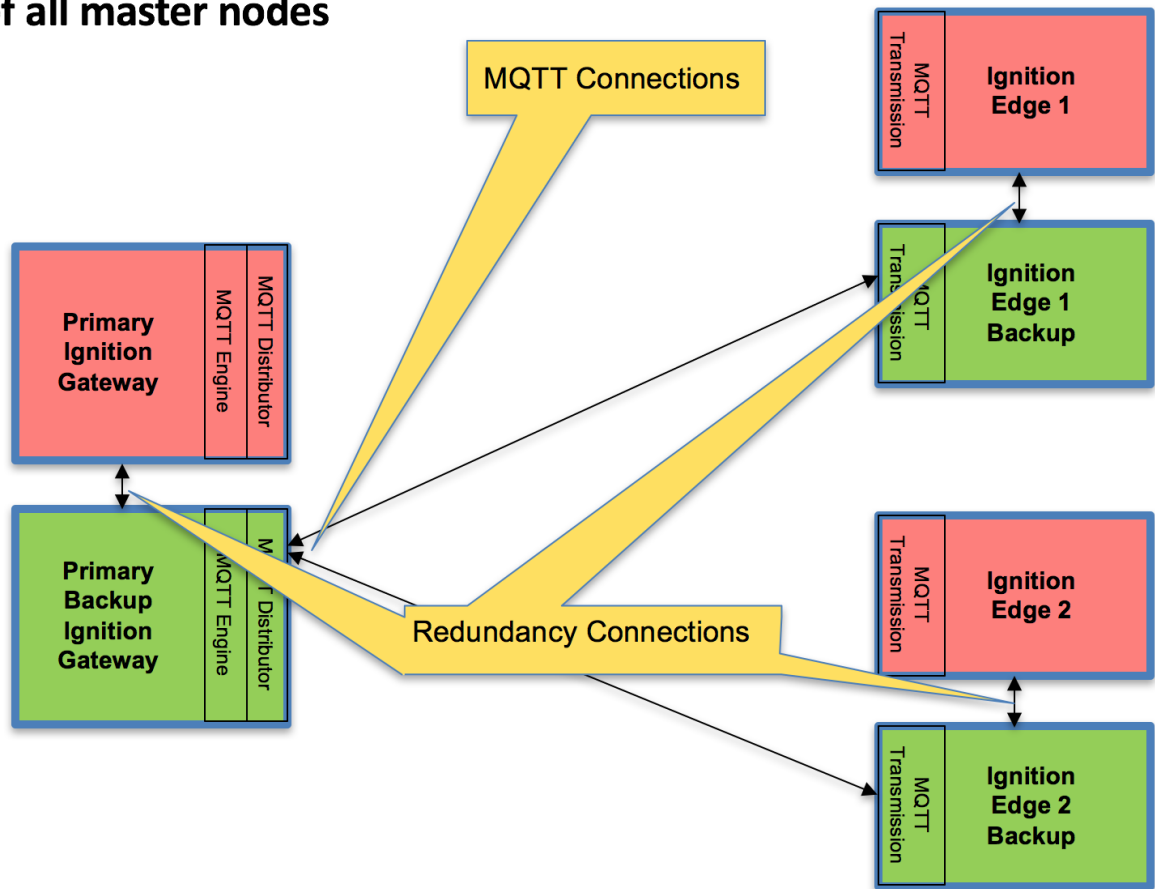
## Primary Ignition Failure



## Ignition Edge 1 Failure



## Failure of all master nodes



## Tutorial

### Step 1: Download and Install Ignition on Six Systems

Ignition is an Industrial Application Platform that can be used to create SCADA and HMI solutions. A fully functional Ignition system can be downloaded and run in trial mode.

Go to the Inductive Automation download page and download the desired Ignition installer for Windows, Linux or MacOS;  
<https://inductiveautomation.com/downloads/archive>.

Once the Ignition installer has been downloaded, follow the instructions provided by Inductive Automation to install and startup Ignition.

(Note: For this test infrastructure, MQTT Distributor will be installed as an Ignition module on both the Primary Ignition Gateway as well as the Primary Ignition Gateway Backup. Make sure to either turn off firewalls or at a minimum allow inbound connections to TCP/IP port 1883, as remote MQTT Clients will need to be able to establish a TCP/IP socket connection to these ports).

### Step 2: Download and Install the Cirrus Link MQTT Modules

Go to the Inductive Automation download page again and scroll down to the Third Party modules section. Find the Cirrus Link modules section and download the MQTT Distributor, MQTT Engine, MQTT Transmission modules.

<https://inductiveautomation.com/downloads/archive>. For each of the Ignition instance, install the following MQTT Modules.

- **Ignition Primary**
  - MQTT Distributor and MQTT Engine
- **Ignition Primary Backup**
  - MQTT Distributor and MQTT Engine
- **Ignition Edge 1**
  - MQTT Transmission
- **Ignition Edge 1 Backup**
  - MQTT Transmission
- **Ignition Edge 2**
  - MQTT Transmission
- **Ignition Edge 2 Backup**
  - MQTT Transmission

### Step 3: Configure the MQTT Modules

Once Ignition is installed, the MQTT Modules are installed, and everything is running we can configure the systems. Since we are going to have a backup for each master system, we only need to do most of the configuration for the master systems. Later, we can sync the configurations to the backups automatically. We'll start by configuring the modules and configure the redundancy settings in the next step.

- **Ignition Primary - MQTT Distributor**

- No modifications to the default parameters are required. However, it is important to make sure the Operation System allows inbound connections on port 1883 and there are no firewalls blocking inbound connections on this port from the remote edge nodes.

- **Ignition Primary - MQTT Engine**

- The only change from defaults is to set a Primary Host ID. MQTT uses Quality of Service (QOS) levels to ensure messages get delivered. However, this only ensures delivery between a single MQTT client and the MQTT server. In other words, it doesn't ensure delivery from one MQTT client to another MQTT client. Sparkplug introduces the notion of a Primary Host ID which is used to ensure client to client communications. The only requirement is that it match exactly on both the MQTT Engine and MQTT Transmission configurations.

MQTT Engine Settings

Servers Advanced Namespaces

Advanced Settings

Configuration

Enabled ☒ Enable or Disable the MQTT Engine (default: true)

Primary Host ID  The Primary Host ID to allow connecting clients to ensure they remain connected to this application (optional)

Set a Primary Host ID to some unique name.

- **Ignition Edge 1 and Ignition Edge 2 - MQTT Transmission** (Configure the same on both Ignition instances)

- As with the MQTT Engine configuration, the Primary Host ID must be set on the General tab as shown below.

MQTT Transmission Settings

General Servers Sets Transmitters

General Settings

Configuration

Enabled ☒ Enable or disable MQTT Transmission from connecting to the configured MQTT Servers

Primary Host ID  Primary Host ID of the backend application the MQTT clients in MQTT Transmission should remain connected to (optional)

Save Changes

Set a Primary Host ID to some unique name (matching the MQTT Engine configuration).

- Delete the existing default MQTT Transmission Server.

MQTT Transmission Settings

General Servers Sets Transmitters

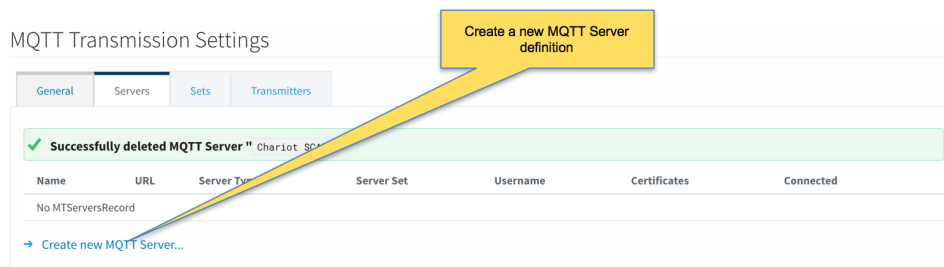
Name	URL	Server Type	Server Set	Username	Certificates	Connected	
Chariot SCADA	tcp://localhost:1883	MQTT Distributor	Default	admin		0 of 0	<a href="#">delete</a> <a href="#">edit</a>

[Create new MQTT Server...](#)

Delete the existing MQTT Transmission Server

- Create a new MQTT Server configuration by clicking the link below.

## MQTT Transmission Settings



The screenshot shows the 'Servers' tab in the MQTT Transmission Settings. A yellow callout bubble points to the 'Create new MQTT Server...' link at the bottom of the table.

**Create a new MQTT Server definition**

General Servers Sets Transmitters

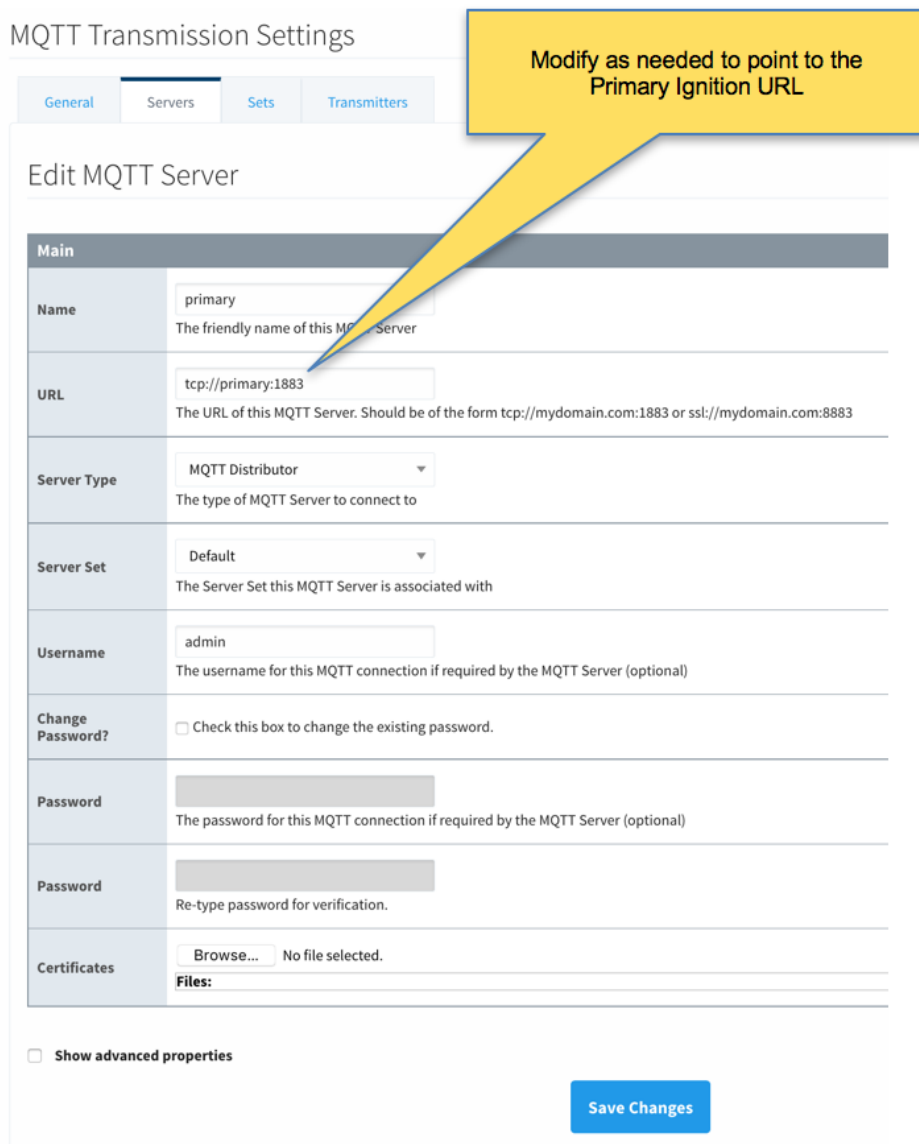
✓ Successfully deleted MQTT Server "Chariot SC..."

Name	URL	Server Type	Server Set	Username	Certificates	Connected
No MTServersRecord						

→ Create new MQTT Server...

- Configure MQTT Transmission to point to the Primary Ignition. Configure as shown below making sure to change the URL to reflect your network settings. For example, if your Ignition Primary is at 192.168.1.100 the MQTT Server URL would be: <tcp://192.168.1.100:1883>. After setting the parameters as shown below. Click the 'Save Changes' button at the bottom.

## MQTT Transmission Settings



The screenshot shows the 'Edit MQTT Server' form. A yellow callout bubble points to the 'URL' field.

General Servers Sets Transmitters

### Edit MQTT Server

**Main**

Name	primary	The friendly name of this MQTT Server
URL	tcp://primary:1883	The URL of this MQTT Server. Should be of the form tcp://mydomain.com:1883 or ssl://mydomain.com:8883
Server Type	MQTT Distributor	The type of MQTT Server to connect to
Server Set	Default	The Server Set this MQTT Server is associated with
Username	admin	The username for this MQTT connection if required by the MQTT Server (optional)
Change Password?	<input type="checkbox"/> Check this box to change the existing password.	
Password		The password for this MQTT connection if required by the MQTT Server (optional)
Password		Re-type password for verification.
Certificates	Browse... No file selected.	Files:

☐ Show advanced properties

Save Changes

**Modify as needed to point to the Primary Ignition URL**

- Verify the MQTT Server has been created and is shown in the list of MQTT Servers as shown below.

## MQTT Transmission Settings

General Servers Sets Transmitters

✓ Successfully created new MQTT Server "primary"

Name	URL	Server Type	Server Set	Username	Certificates	Connected	
primary	tcp://primary:1883	MQTT Distributor	Default	admin		0 of 0	delete edit

→ Create new MQTT Server

You should now see a new MQTT Server shown in the Servers Settings

- Repeat the process of creating a MQTT Server but instead point it to the Ignition Primary Backup MQTT Server. These are the parameters to use:
  - Name: primary-backup
  - URL: <tcp://primary-backup:1883>
    - Change 'primary-backup' in the URL to reflect the network address of the Ignition Primary Backup server.
  - Server Type: MQTT Distributor
  - Server Set: Default
  - Username: admin
  - Password: changeme
- When complete, verify both MQTT Servers appear in the list as shown below.

## MQTT Transmission Settings

General Servers Sets Transmitters

✓ Successfully created new MQTT Server "primary-backup"

Name	URL	Server Type	Server Set	Username	Certificates	Connected	
primary	tcp://primary:1883	MQTT Distributor	Default	admin		0 of 0	delete edit
primary-backup	tcp://primary-backup:1883	MQTT Distributor	Default	admin		0 of 0	delete edit

→ Create new MQTT Server...

- Finally, make sure to set up the same MQTT Transmission configuration in the Ignition Edge 2 instance.

## Step 4: Configure Redundancy

The following configuration shows all of the redundancy settings that were used in setting this environment up using Amazon's AWS EC2 instances (virtual machines). The configuration will vary based on your network configuration. Additional Ignition redundancy resources can be found at the following links:

<https://docs.inductiveautomation.com/display/DOC79/Setting+Up+Redundancy>  
[https://support.inductiveautomation.com/usermanuals/ignition/index.html?redundancy\\_settings.htm](https://support.inductiveautomation.com/usermanuals/ignition/index.html?redundancy_settings.htm)

### Ignition Primary

- Select Redundancy on the left navigation bar. Then set the Mode to 'Master' and set the Standby Activity left to 'Warm' as shown below.

Ignition! by inductive automation

HOME STATUS CONFIGURE

Search...

Trial Version 1:26:34 We're glad you're test driving our software. Have fun.

SYSTEM

- Overview
- Backup/Restore
- Licensing
- Modules
- Projects
- Redundancy
- Gateway Settings

NETWORKING

- Gateway Network
- Email Settings

SECURITY

- Auditing
- Users, Roles
- Service Security
- Security Zones

DATABASES

- Connections

### Redundancy and Network Configuration

#### Redundancy Settings

Mode	Master	Enable or disable redundancy, and specify this node's role. There should be one m
Standby Activity Level	Warm	How the node should run when it is not currently the active node. If 'cold', the no high level, reducing failover times.
Failover Timeout	10000	The time of inactivity, in milliseconds, before the backup assumes responsibility. (default: 10000)

- Set up the Redundancy Network Settings. The settings here are specific to your network setup. **On many LAN configurations none of these changes are required.** What is shown below was the configuration for setting up all of these components in Amazon's AWS EC2 instances. The changes were:
  - Uncheck 'Auto-detect network interface'
  - Set the 'Network Bind Interface' to the public IP address of the Ignition Primary EC2 instance. On a LAN this would be the primary network interface address of the Ignition Primary machine.
  - Uncheck the 'Autodetect HTTP Address' tickbox.



- Specify two explicit HTTP addresses for clients to use. These were the public IP addresses of the Ignition Primary and Ignition Primary Backup EC2 instances. On a LAN, these would be the primary network interface addresses of the Ignition Primary and Ignition Primary Backup machines. Also note the HTTP port is 8088 which is the default Ignition HTTP port.

Network Settings

Port: 8750  
The TCP port used for redundancy operations. Make sure that this port is not blocked by a firewall.  
(default: 8750)

Auto-detect network interface: ☒ If true, the system will automatically detect which network interface to use. Most commonly disabled on systems with multiple network cards, in order to explicitly specify which interface to use.  
(default: true)

Network Bind Interface: 34.201.169.174  
The IP address of the network interface to use for redundancy. Only used if "auto-detect" is turned off.

Autodetect HTTP Address: ☒ To specify an explicit HTTP address for clients to use, turn this off. Most users will leave autodetect on.  
(default: true)

Address	HTTP Port	HTTPS Port	
primary	8088	443	Remove
primary-backup	8088	443	Remove
Add New Address			
	80	443	Add

If autodetect HTTP address is false, clients will use these addresses. Usually only necessary in multi-homed or port-forwarded redundancy configurations.

- Set the Master Node Address. Note in the configuration below a hostname is being used. This should be the primary network interface address of the Ignition Primary Gateway.

Set the Master Node Address

Backup Node Settings

Master Node Address: primary  
The address of the master Ignition system.

Ping Rate: 1000  
The time, in milliseconds, between messages from the backup to the master.  
(default: 1000)

Reconnect Period: 10000  
How often, in milliseconds, to re-attempt connection when the backup node is not connected to the master.  
(default: 10000)

History Mode: Full  
How history is treated by the backup system. If Full, history will be stored normally, as it would be on the master system. If Partial, history will be cached until the master is available again and the backup node is able to determine the exact time that the master was down.

Save Changes

- Finally, click the 'Save Changes' button.
- Ignition Primary Backup**
  - Select Redundancy on the left navigation bar. Then set the Mode to 'Backup' and set the Standby Activity left to 'Warm' as shown below.

Select Redundancy on the left navigation bar

Redundancy and Network Configuration

Redundancy Settings

Mode: Backup  
Enable or disable redundancy, and specify this node's role. There should be one mas

Standby Activity Level: Warm  
How the node should run when it is not currently the active node. If cold, the node high level, reducing failover times.

Failover Timeout: 10000  
The time of inactivity, in milliseconds, before the backup assumes responsibility.  
(default: 10000)

- Set up the Redundancy Network Settings. The settings here are specific to your network setup. **On many LAN configurations none of these changes are required.** What is shown below was the configuration for setting up all of these components in Amazon's AWS EC2 instances. The changes were:
  - Uncheck 'Auto-detect network interface'
  - Set the 'Network Bind Interface' to the public IP address of the Ignition Primary Backup EC2 instance. On a LAN this would be the primary network interface address of the Ignition Primary Backup machine.
  - Uncheck the 'Autodetect HTTP Address' tickbox.
  - Specify two explicit HTTP addresses for clients to use. These were the public IP addresses of the Ignition Primary and Ignition Primary Backup EC2 instances. On a LAN, these would be the primary network interface addresses of the Ignition Primary and Ignition Primary Backup machines. Also note the HTTP port is 8088 which is the default Ignition HTTP port.

The screenshot shows the 'Network Settings' page in the Ignition Primary web interface. The left sidebar contains navigation links like Journal, Notification, On-Call Rosters, Schedules, SECs/GEM, Equipment, Module Settings, Simulator, User Manual, TAGS, History, Realtime, OPC-UA SERVER, Certificates, Devices, Settings, OPC CONNECTIONS, Servers, Quick Client, MOBILE, Settings, ENTERPRISE ADMINISTRATION, Agent Settings, and SEQUENTIAL FUNCTION. The main content area is titled 'Network Settings' and includes the following fields:

- Port:** 8750. The TCP port used for redundancy operations. Make sure that this port is not blocked by a firewall. (default: 8750)
- Auto-detect network interface?** ☐ If true, the system will automatically detect which network interface to use. Most commonly disabled on systems with multiple network cards, in order to explicitly specify which interface to use. (default: true)
- Network Bind Interface:** 54.211.69.177. The IP address of the network interface to use for redundancy. Only used if "auto-detect" is turned off.
- Autodetect HTTP Address:** ☐ To specify an explicit HTTP address for clients to use, turn this off. Most users will leave autodetect on. (default: true)
- HTTP Addresses:** A table with columns for Address, HTTP Port, and HTTPS Port. It lists 'primary' and 'primary-backup' both with HTTP Port 8088 and HTTPS Port 443. There is an 'Add New Address' button and an 'Add' button at the bottom.

At the bottom, a note states: "If autodetect HTTP address is false, clients will use these addresses. Usually only necessary in multi-homed or port-forwarded redundancy configurations."

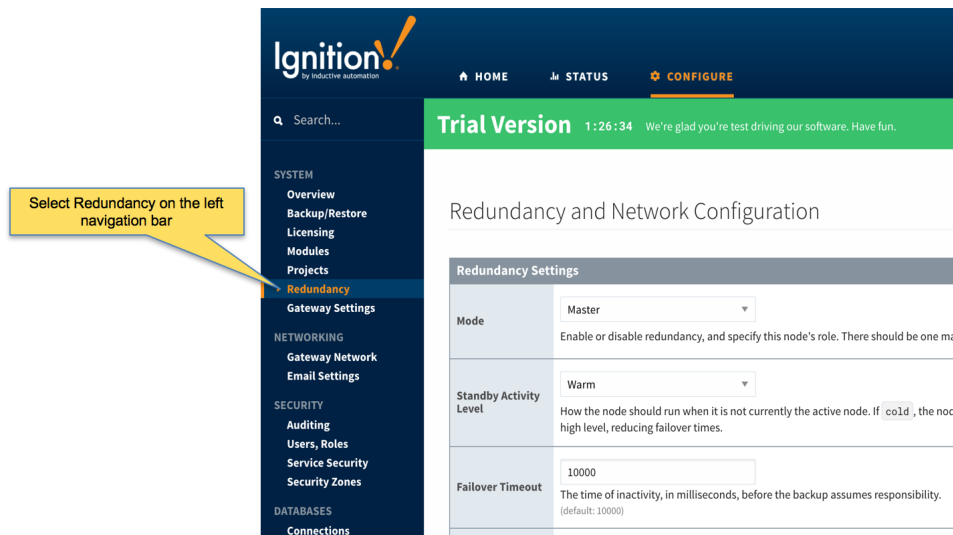
- Set the Master Node Address. Note in the configuration below a hostname is being used. This should be the primary network interface address of the Ignition Primary Gateway.

The screenshot shows the 'Backup Node Settings' page in the Ignition Primary web interface. A yellow callout bubble with the text "Set the Master Node Address" points to the 'Master Node Address' field, which is set to 'primary'. The page includes the following settings:

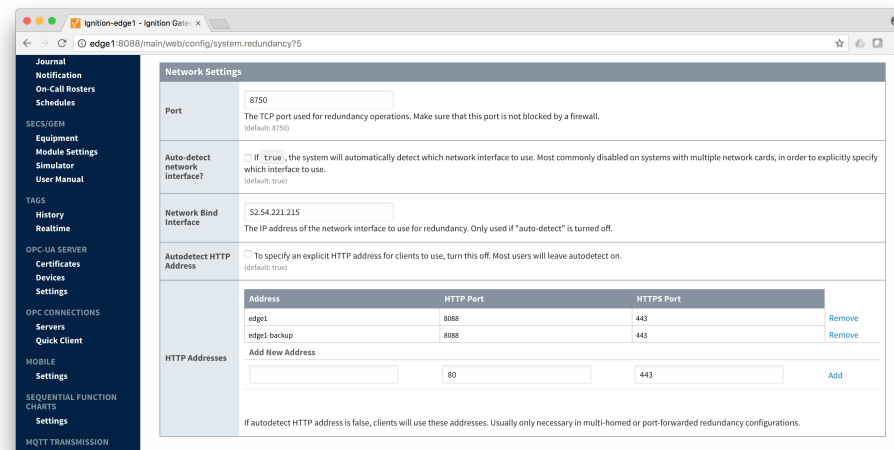
- Master Node Address:** primary. The address of the master Ignition system.
- Ping Rate:** 1000. The time, in milliseconds, between messages from the backup to the master. (default: 1000)
- Reconnect Period:** 10000. How often, in milliseconds, to re-attempt connection when the backup node is not connected to the master. (default: 10000)
- History Mode:** Full. How history is treated by the backup system. If Full, history will be stored normally, as it would be on the master system. If Partial, history will be cached until the master is available again and the backup node is able to determine the exact time that the master was down.

At the bottom right, there is a 'Save Changes' button. The Inductive Automation logo is visible at the bottom center.

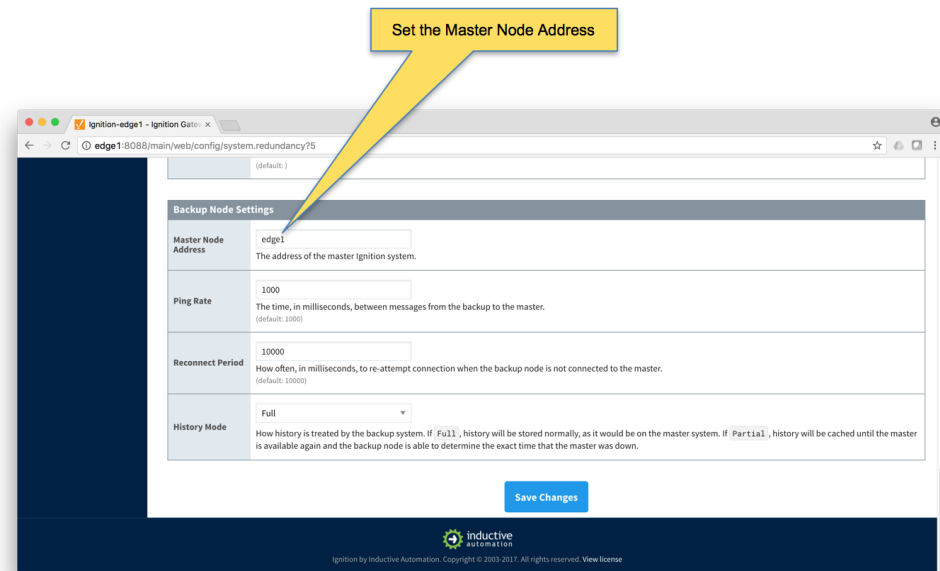
- Finally, click the 'Save Changes' button.
- **Ignition Edge 1**
  - Select Redundancy on the left navigation bar. Then set the Mode to 'Master' and set the Standby Activity left to 'Warm' as shown below.



- Set up the Redundancy Network Settings. The settings here are specific to your network setup. **On many LAN configurations none of these changes are required.** What is shown below was the configuration for setting up all of these components in Amazon's AWS EC2 instances. The changes were:
  - Uncheck 'Auto-detect network interface'
  - Set the 'Network Bind Interface' to the public IP address of the Ignition Edge 1 EC2 instance. On a LAN this would be the primary network interface address of the Ignition Edge 1 machine.
  - Uncheck the 'Autodetect HTTP Address' tickbox.
  - Specify two explicit HTTP addresses for clients to use. These were the public IP addresses of the Ignition Edge 1 and Ignition Edge 1 Backup EC2 instances. On a LAN, these would be the primary network interface addresses of the Ignition Edge 1 and Ignition Edge 1 Backup machines. Also note the HTTP port is 8088 which is the default Ignition HTTP port.



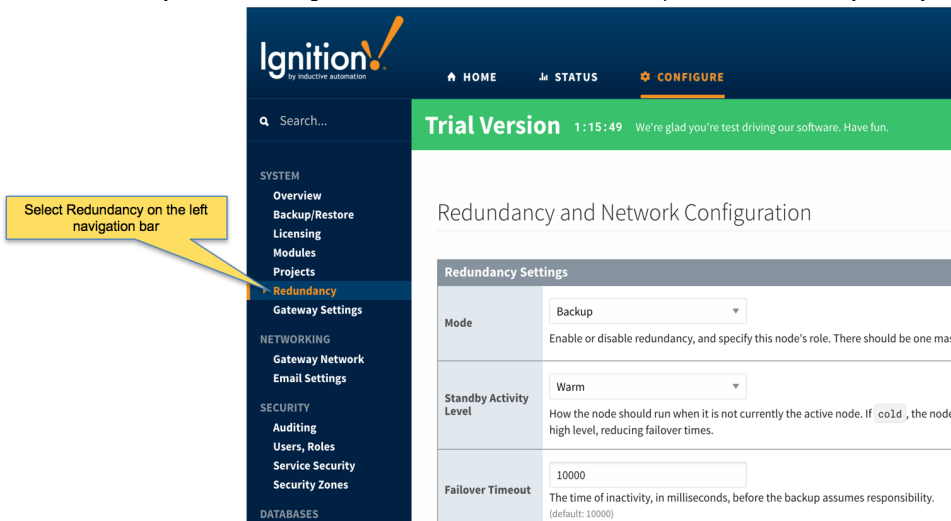
- Set the Master Node Address. Note in the configuration below a hostname is being used. This should be the primary network interface address of the Ignition Edge 1 Gateway.



- Finally, click the 'Save Changes' button.

- **Ignition Edge 1 Backup**

- Select Redundancy on the left navigation bar. Then set the Mode to 'Backup' and set the Standby Activity left to 'Warm' as shown below.



- Set up the Redundancy Network Settings. The settings here are specific to your network setup. **On many LAN configurations none of these changes are required.** What is shown below was the configuration for setting up all of these components in Amazon's AWS EC2 instances. The changes were:
  - Uncheck 'Auto-detect network interface'
  - Set the 'Network Bind Interface' to the public IP address of the Ignition Edge 1 Backup EC2 instance. On a LAN this would be the primary network interface address of the Ignition Edge 1 Backup machine.
  - Uncheck the 'Autodetect HTTP Address' tickbox.
  - Specify two explicit HTTP addresses for clients to use. These were the public IP addresses of the Ignition Edge 1 and Ignition Edge 1 Backup EC2 instances. On a LAN, these would be the primary network interface addresses of the Ignition Edge 1 and Ignition Edge 1 Backup machines. Also note the HTTP port is 8088 which is the default Ignition HTTP port.

The screenshot shows the 'Network Settings' page in the Ignition Edge 1 configuration interface. The left sidebar contains a navigation menu with categories like Journal, Notification, On-Call Rosters, Schedules, SECS/GEM, Equipment, Module Settings, Simulator, User Manual, TAGS, History, Realtime, OPC-UA SERVER, Certificates, Devices, Settings, OPC CONNECTIONS, Servers, Quick Client, MOBILE, Settings, SEQUENTIAL FUNCTION CHARTS, Settings, and MQTT TRANSMISSION.

The main content area is titled 'Network Settings' and includes the following fields:

- Port:** 8750 (The TCP port used for redundancy operations. Make sure that this port is not blocked by a firewall. (default: 8750))
- Auto-detect network interface?** ☐ If `true`, the system will automatically detect which network interface to use. Most commonly disabled on systems with multiple network cards, in order to explicitly specify which interface to use. (default: true)
- Network Bind Interface:** 34.229.100.212 (The IP address of the network interface to use for redundancy. Only used if "auto-detect" is turned off.)
- Autodetect HTTP Address:** ☐ To specify an explicit HTTP address for clients to use, turn this off. Most users will leave autodetect on. (default: true)
- HTTP Addresses:** A table with columns for Address, HTTP Port, and HTTPS Port. It lists 'edge1' and 'edge1-backup' with their respective ports (8088, 443). There is an 'Add New Address' button and an 'Add' button for a new entry.

At the bottom, a note states: "If autodetect HTTP address is false, clients will use these addresses. Usually only necessary in multi-homed or port-forwarded redundancy configurations."

- Set the Master Node Address. Note in the configuration below a hostname is being used. This should be the primary network interface address of the Ignition Edge 1 Gateway.

The screenshot shows the 'Backup Node Settings' page in the Ignition Edge 1 configuration interface. A yellow callout box with the text 'Set the Master Node Address' points to the 'Master Node Address' field, which contains the value 'edge1'.

The main content area is titled 'Backup Node Settings' and includes the following fields:

- Master Node Address:** edge1 (The address of the master Ignition system.)
- Ping Rate:** 1000 (The time, in milliseconds, between messages from the backup to the master. (default: 1000))
- Reconnect Period:** 10000 (How often, in milliseconds, to re-attempt connection when the backup node is not connected to the master. (default: 10000))
- History Mode:** Full (How history is treated by the backup system. If `Full`, history will be stored normally, as it would be on the master system. If `Partial`, history will be cached until the master is available again and the backup node is able to determine the exact time that the master was down.)

At the bottom right, there is a 'Save Changes' button.

- Finally, click the 'Save Changes' button.

- **Ignition Edge 2**

- Select Redundancy on the left navigation bar. Then set the Mode to 'Master' and set the Standby Activity left to 'Warm' as shown below.

The screenshot shows the Ignition Edge 2 configuration interface. The left sidebar contains a navigation menu with categories like SYSTEM, NETWORKING, SECURITY, and DATABASES. The 'Redundancy' option under the SYSTEM category is highlighted with a yellow callout box that says 'Select Redundancy on the left navigation bar'.

The main content area is titled 'Redundancy and Network Configuration' and includes the following fields:

- Mode:** Master (Enable or disable redundancy, and specify this node's role. There should be one master node.)
- Standby Activity Level:** Warm (How the node should run when it is not currently the active node. If `cold`, the node will be in a low power state, reducing failover times.)
- Failover Timeout:** 10000 (The time of inactivity, in milliseconds, before the backup assumes responsibility. (default: 10000))

- Set up the Redundancy Network Settings. The settings here are specific to your network setup. **On many LAN configurations none of these changes are required.** What is shown below was the configuration for setting up all of these components in Amazon's AWS EC2 instances. The changes were:
  - Uncheck 'Auto-detect network interface'
  - Set the 'Network Bind Interface' to the public IP address of the Ignition Edge 2 EC2 instance. On a LAN this would be the primary network interface address of the Ignition Edge 2 machine.
  - Uncheck the 'Autodetect HTTP Address' tickbox.
  - Specify two explicit HTTP addresses for clients to use. These were the public IP addresses of the Ignition Edge 2 and Ignition Edge 2 Backup EC2 instances. On a LAN, these would be the primary network interface addresses of the Ignition Edge 2 and Ignition Edge 2 Backup machines. Also note the HTTP port is 8088 which is the default Ignition HTTP port.

Network Settings

Port: 8750  
The TCP port used for redundancy operations. Make sure that this port is not blocked by a firewall.  
(default: 8750)

Auto-detect network interface? ☐ If 'true', the system will automatically detect which network interface to use. Most commonly disabled on systems with multiple network cards, in order to explicitly specify which interface to use.  
(default: true)

Network Bind Interface: 54.236.219.237  
The IP address of the network interface to use for redundancy. Only used if "auto-detect" is turned off.

Autodetect HTTP Address ☐ To specify an explicit HTTP address for clients to use, turn this off. Most users will leave autodetect on.  
(default: true)

Address	HTTP Port	HTTPS Port	
edge2	8088	443	Remove
edge2 backup	8088	443	Remove
Add New Address			Add

If autodetect HTTP address is false, clients will use these addresses. Usually only necessary in multi-homed or port-forwarded redundancy configurations.

- Set the Master Node Address. Note in the configuration below a hostname is being used. This should be the primary network interface address of the Ignition Edge 2 Gateway.

Set the Master Node Address

Backup Node Settings

Master Node Address: edge2  
The address of the master Ignition system.

Ping Rate: 1000  
The time, in milliseconds, between messages from the backup to the master.  
(default: 1000)

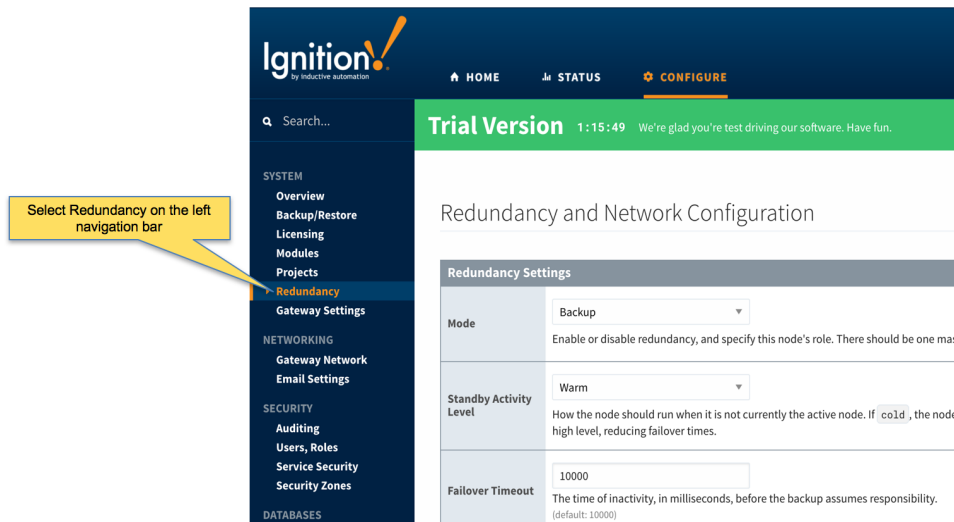
Reconnect Period: 10000  
How often, in milliseconds, to re-attempt connection when the backup node is not connected to the master.  
(default: 10000)

History Mode: Full  
How history is treated by the backup system. If 'Full', history will be stored normally, as it would be on the master system. If 'Partial', history will be cached until the master is available again and the backup node is able to determine the exact time that the master was down.

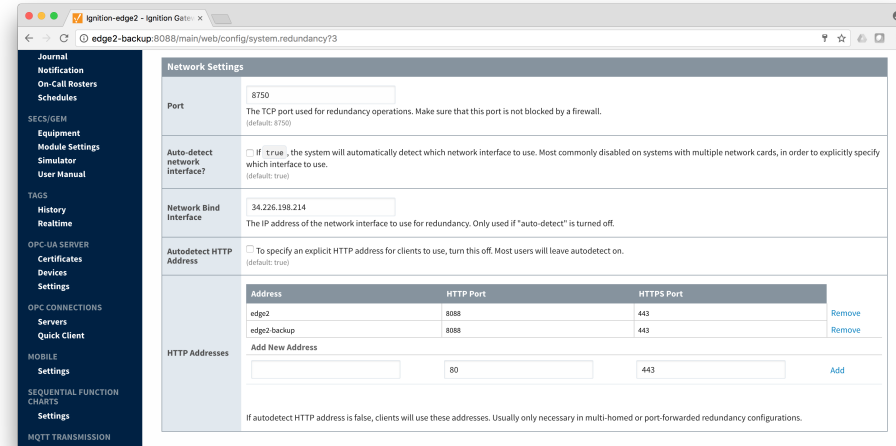
Save Changes

Inductive Automation  
Ignition by Inductive Automation. Copyright © 2003-2017. All rights reserved. View license

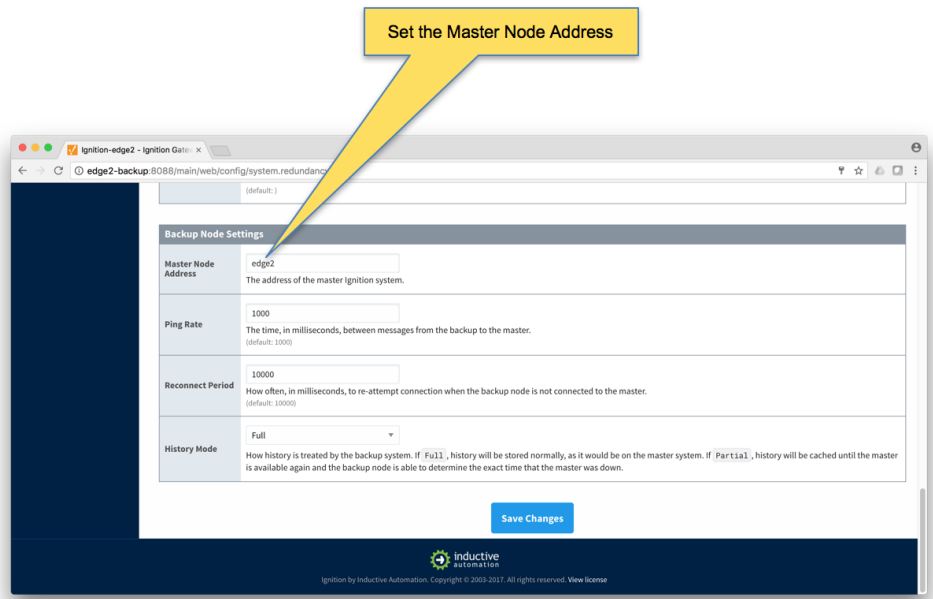
- Finally, click the 'Save Changes' button.
- Ignition Edge 2 Backup
  - Select Redundancy on the left navigation bar. Then set the Mode to 'Backup' and set the Standby Activity left to 'Warm' as shown below.



- Set up the Redundancy Network Settings. The settings here are specific to your network setup. **On many LAN configurations none of these changes are required.** What is shown below was the configuration for setting up all of these components in Amazon's AWS EC2 instances. The changes were:
  - Uncheck 'Auto-detect network interface'
  - Set the 'Network Bind Interface' to the public IP address of the Ignition Edge 2 Backup EC2 instance. On a LAN this would be the primary network interface address of the Ignition Edge 2 Backup machine.
  - Uncheck the 'Autodetect HTTP Address' checkbox.
  - Specify two explicit HTTP addresses for clients to use. These were the public IP addresses of the Ignition Edge 2 and Ignition Edge 2 Backup EC2 instances. On a LAN, these would be the primary network interface addresses of the Ignition Edge 2 and Ignition Edge 2 Backup machines. Also note the HTTP port is 8088 which is the default Ignition HTTP port.



- Set the Master Node Address. Note in the configuration below a hostname is being used. This should be the primary network interface address of the Ignition Edge 2 Gateway.



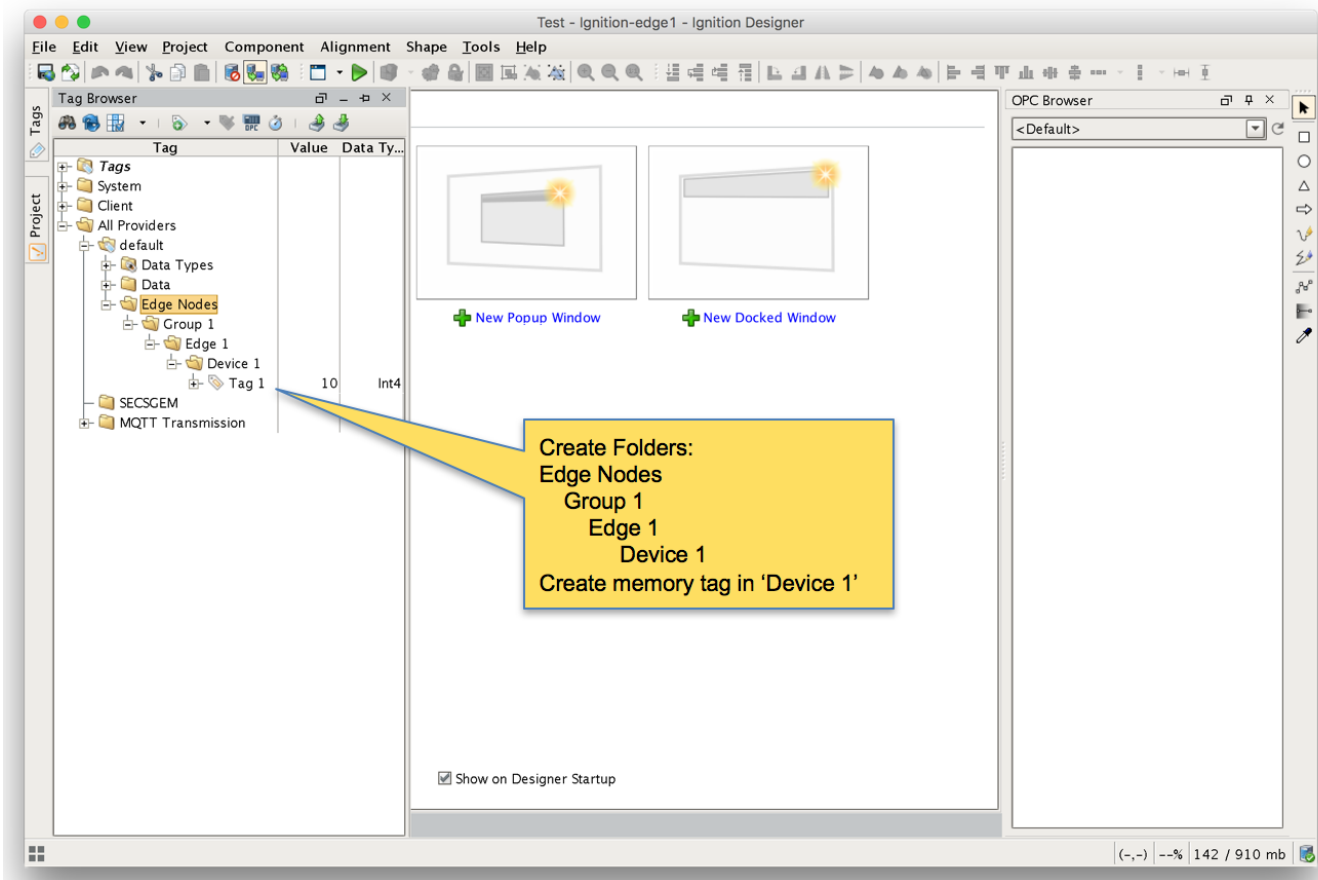
- Finally, click the 'Save Changes' button.

## Step 5: Create some tags in Edge 1 and Edge 2

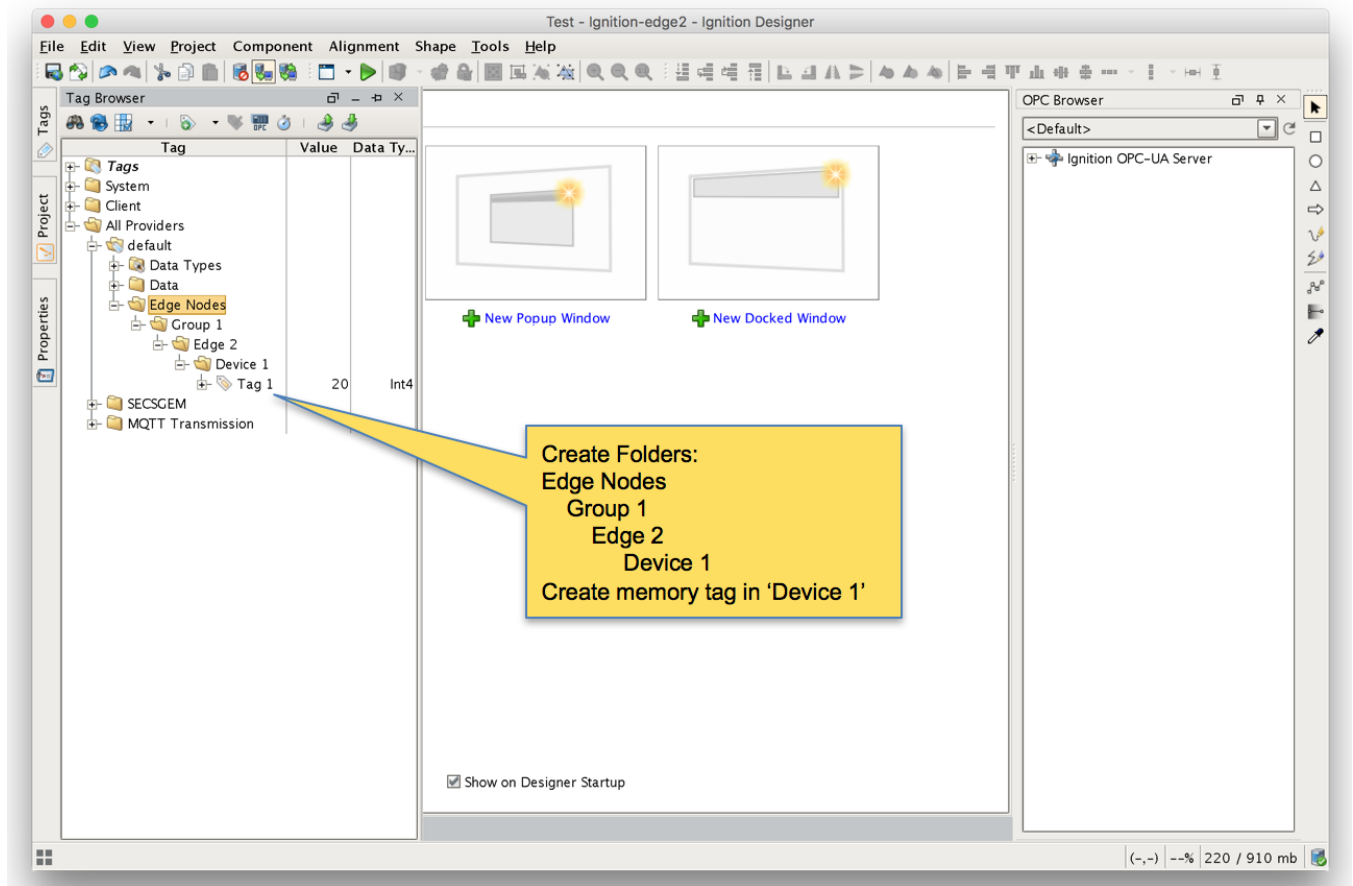
In this step we will use Ignition Designer in the Ignition Edge 1 and Edge 2 instances to create some tags. These tags will be used by MQTT Transmission and the Default Transmitter to push data to MQTT Distributor and MQTT Engine in the Ignition Primary instances.

- Using Ignition Designer on Edge 1, do the following:
  - Create a new project called 'Test'.
  - In the Tag Browser, confirm there is a folder called 'Edge Nodes'. If there is not, confirm MQTT Transmission is installed.
  - In the 'Edge Nodes' folder, create a folder called 'Group 1'.
  - In the 'Group 1' folder, create a folder called 'Edge 1'.
  - In the 'Edge 1' folder, create a folder called 'Device 1'.
  - In the 'Device 1' folder, create a Tag called 'Tag 1'.
  - At the end, you should see something similar to what is shown below.

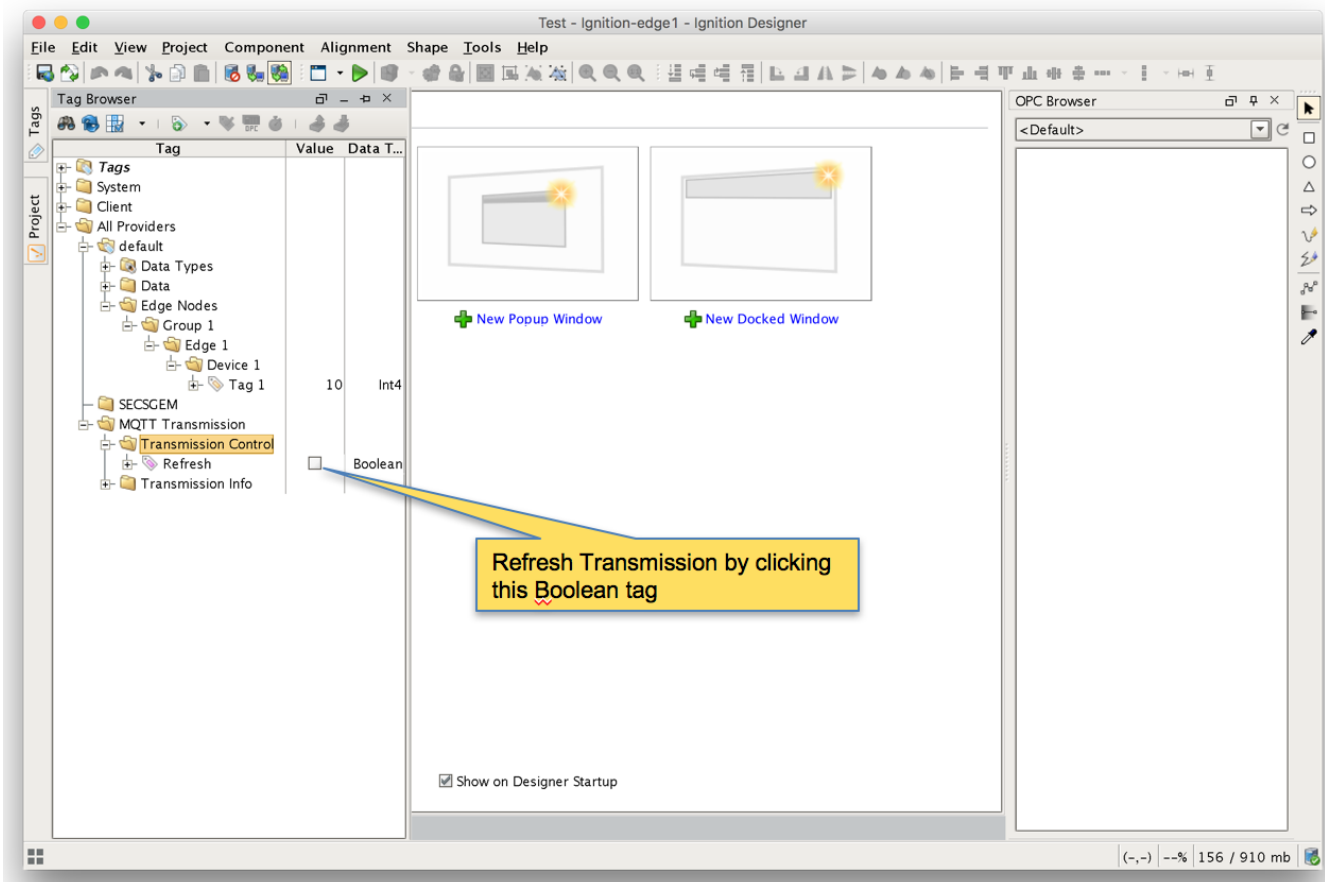




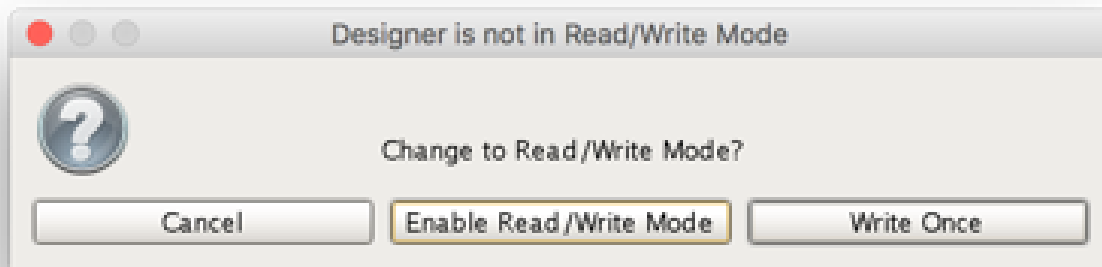
- Using Ignition Designer on Edge 2, do the following:
  - Create a new project called 'Test'.
  - In the Tag Browser, confirm there is a folder called 'Edge Nodes'. If there is not, confirm MQTT Transmission is installed.
  - In the 'Edge Nodes' folder, create a folder called 'Group 1'.
  - In the 'Group 1' folder, create a folder called 'Edge 2'.
  - In the 'Edge 2' folder, create a folder called 'Device 1'.
  - In the 'Device 1' folder, create a Tag called 'Tag 1'.
  - At the end, you should see something similar to what is shown below



- Finally, refresh the Transmission runtime. This is done by clicking the 'MQTT Transmission/Transmission Control/Refresh' Boolean tag.



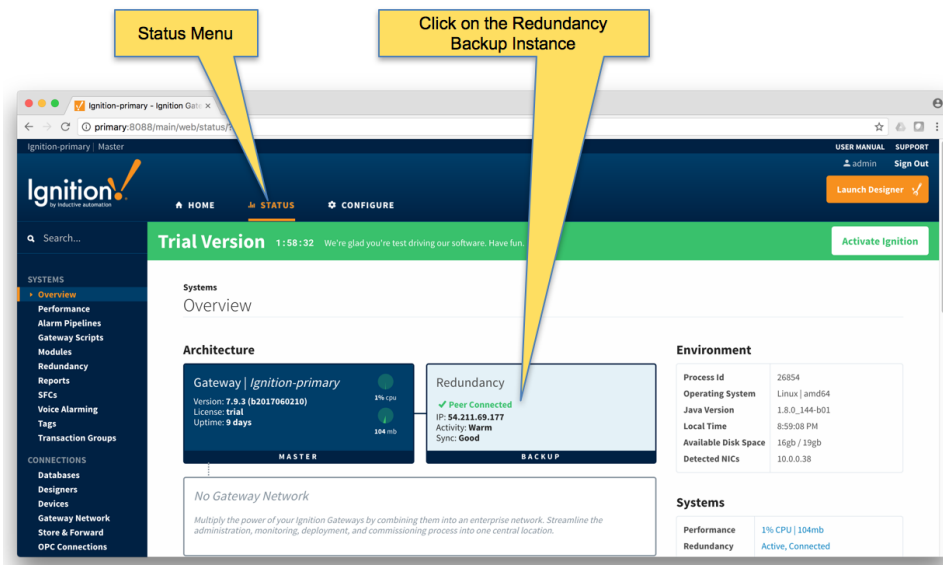
- After clicking the Boolean tag you may need to 'Enable Read/Write Mode'



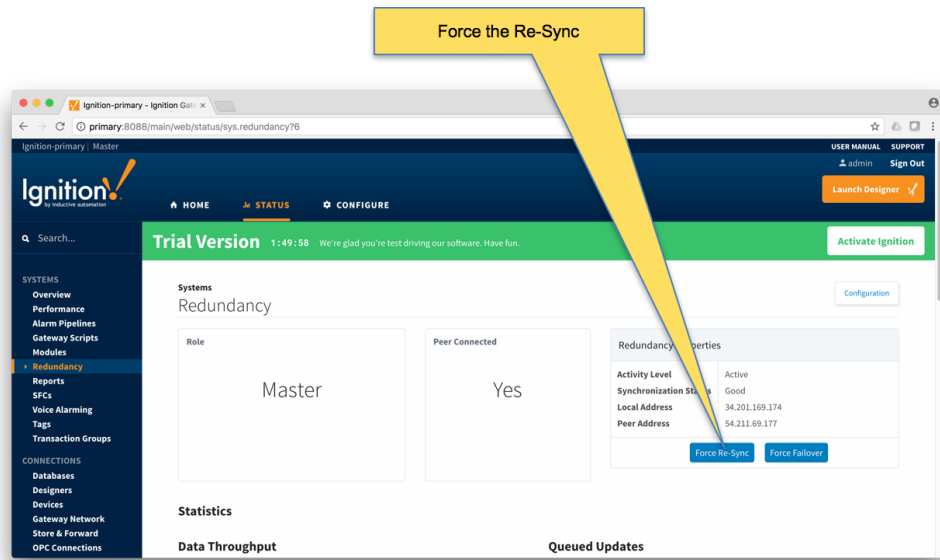
## Step 6: Force Sync of Backup Instances

The next step is to force the backup Ignition instances to receive their configurations from the master Ignition instances. This will happen automatically eventually but for expediency we're going to force the action.

- Repeat the following steps for the following Ignition instances: Primary, Edge 1, and Edge 2
  - Browse to the Status menu and then click in the Redundancy box as shown below.



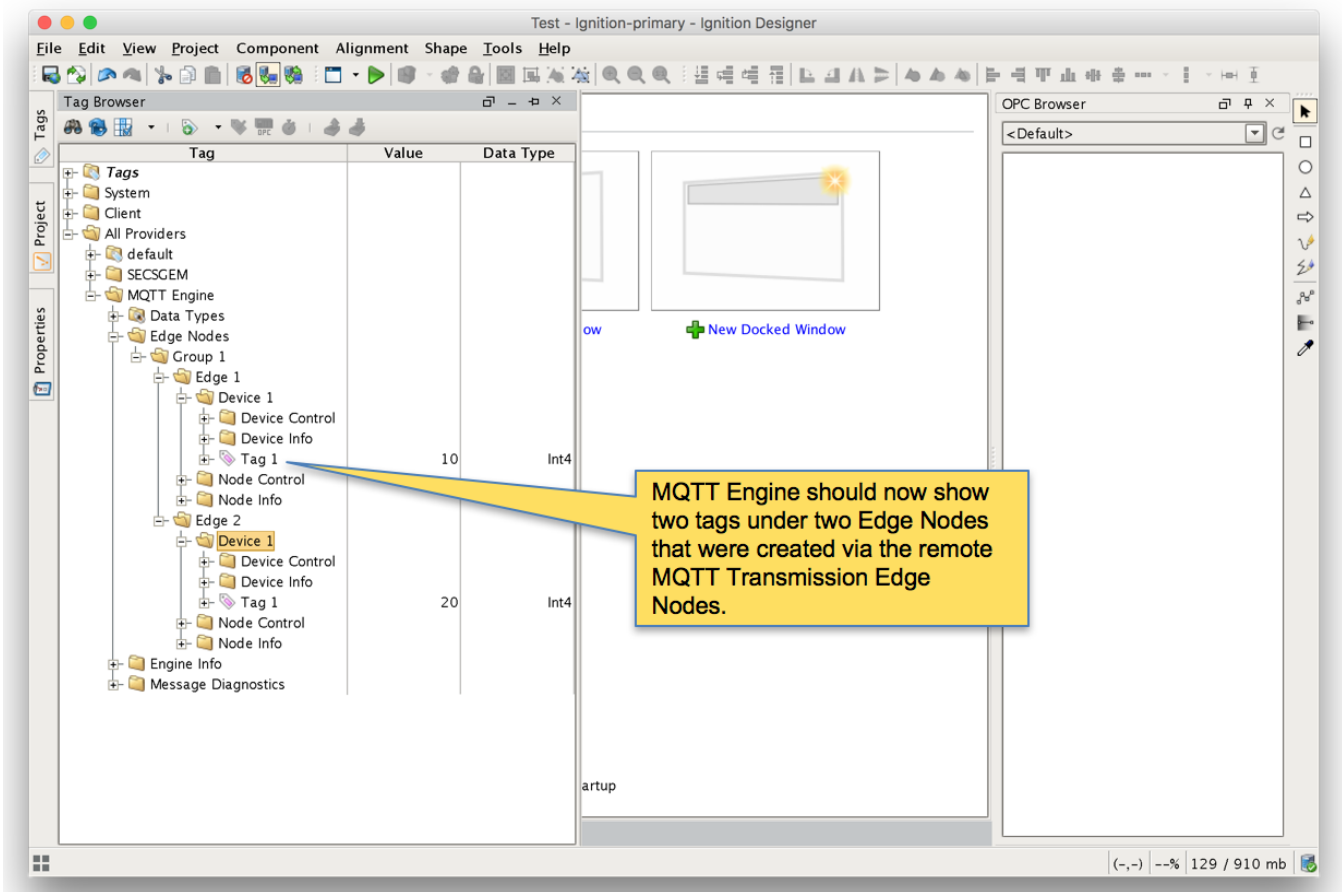
- Force the Re-sync via the button below. This will duplicate the master configuration to the backup Ignition instance it is connected to.



- Repeat the above Re-sync steps on the Edge 1 and Edge 2 Ignition instances.

## Step 7: Verify MQTT Engine is getting data from the MQTT Transmission Edge 1 and Edge 2

Open Ignition Designer on the Ignition Primary instance. Expand the MQTT Engine tag tree and validate the following tags are present. If they are present and not stale, they are properly connected.



## Step 8: Test the Redundancy

In order to test the redundancy, we need to make a few simple dashboards. It is important to note that this can not be tested with Ignition Designer alone. Designer can not be opened from an Ignition backup instance since projects get replicated to the backup instances. So, to show everything working, we'll make some very rudimentary dashboards.

- **Ignition Primary**

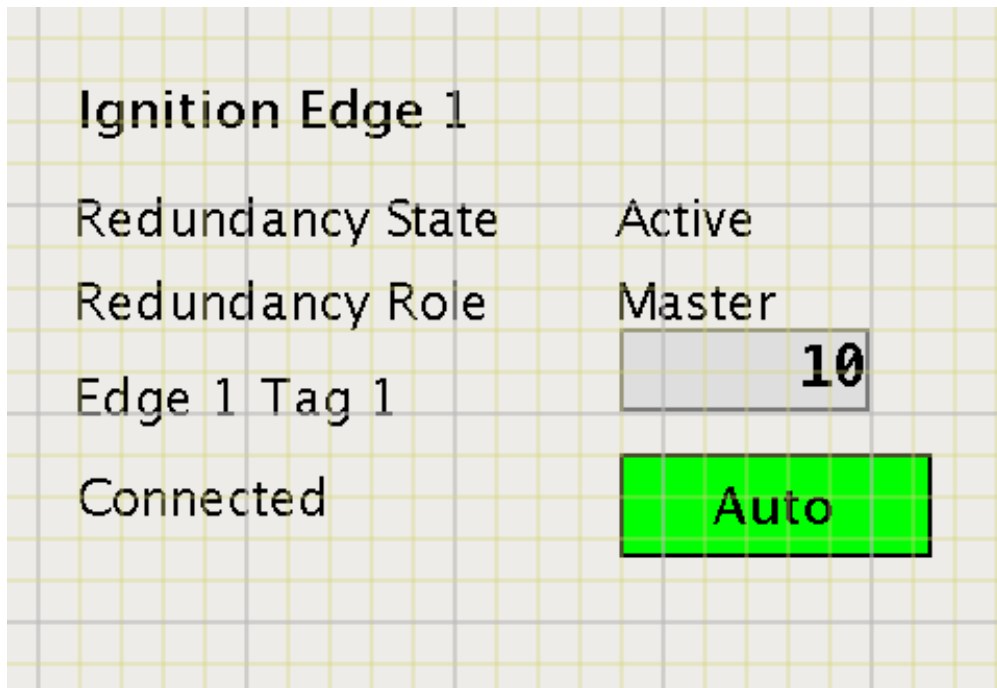
- Create the following widgets.

- Label - Ignition Primary
- Label - Redundancy State Label with Tag Path of "[MQTT Engine]/Engine Info/Redundancy State"
- Label - Redundancy Role Label with Tag Path of "[MQTT Engine]/Engine Info/Redundancy Role"
- Label - Edge 1 Tag 1 Label with Tag Path of "[MQTT Engine]/Edge Nodes/Group 1/Edge 1/Device 1/Tag 1"
- Label - Edge 2 Tag 1 Label with Tag Path of "[MQTT Engine]/Edge Nodes/Group 1/Edge 2/Device 1/Tag 1"
- Label - MQTT Engine Connected Multi-State Indicator with Tag Path of "[MQTT Engine]/Engine Info/MQTT Clients/Chariot SCADA/Online"
- Label - Edge 1 Connected Multi-State Indicator with Tag Path of "[MQTT Engine]/Edge Nodes/Group 1/Edge 1/Node Info/Online"
- Label - Edge 2 Connected Multi-State Indicator with Tag Path of "[MQTT Engine]/Edge Nodes/Group 1/Edge 2/Node Info/Online"

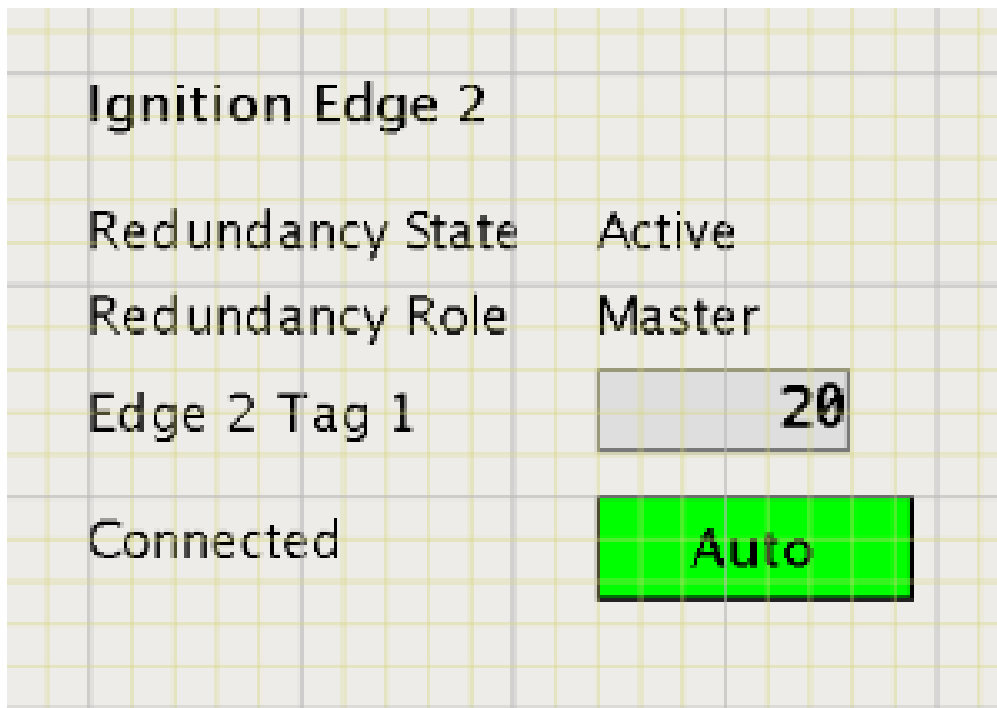
- When complete - it should look similar to what is shown below

Ignition Primary	
Redundancy State	Active
Redundancy Role	Master
Edge 1 Tag 1	10
Edge 2 Tag 1	20
MQTT Engine Connected	Auto
Edge 1 Connected	Auto
Edge 2 Connected	Auto

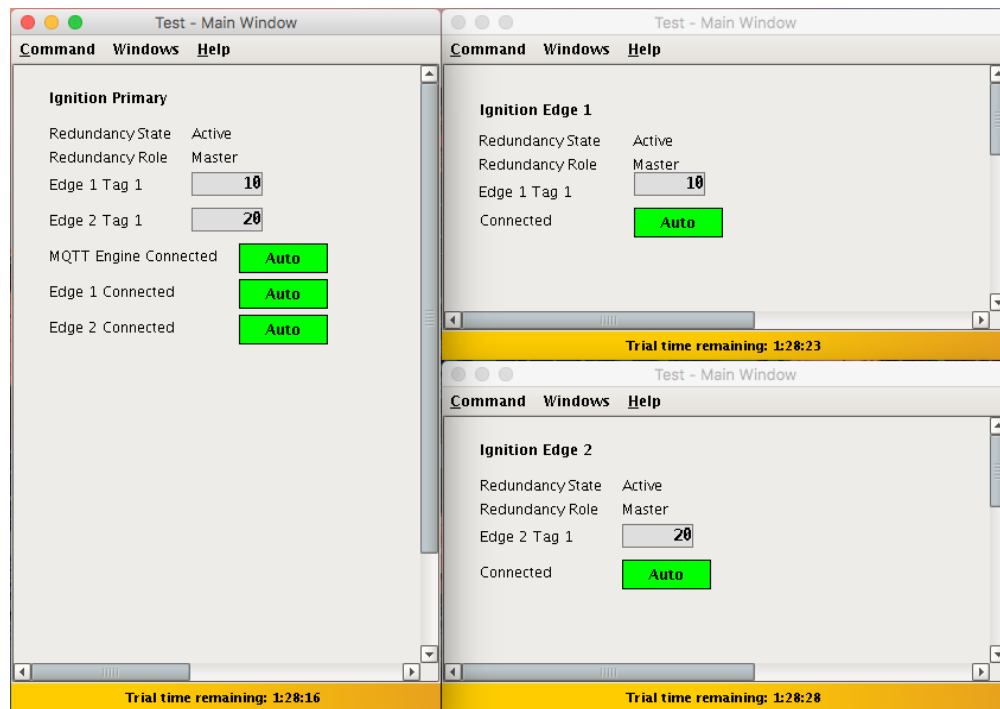
- Ignition Edge 1
  - Create the following widgets
    - Label - Ignition Edge 1
    - Label - Redundancy State Label with Tag Path of "[MQTT Transmission]/Transmission Info/Redundancy State"
    - Label - Redundancy Role Label with Tag Path of "[MQTT Transmission]/Transmission Info/Redundancy Role"
    - Label - Edge 1 Tag 1 Label with Tag Path of "[default]/Edge Nodes/Group 1/Edge 1/Device 1/Tag 1"
    - Label - Connected Multi-State Indicator with Tag Path of "[MQTT Transmission]/Transmission Info/MQTT Clients/Group 1-Edge 1/Online"



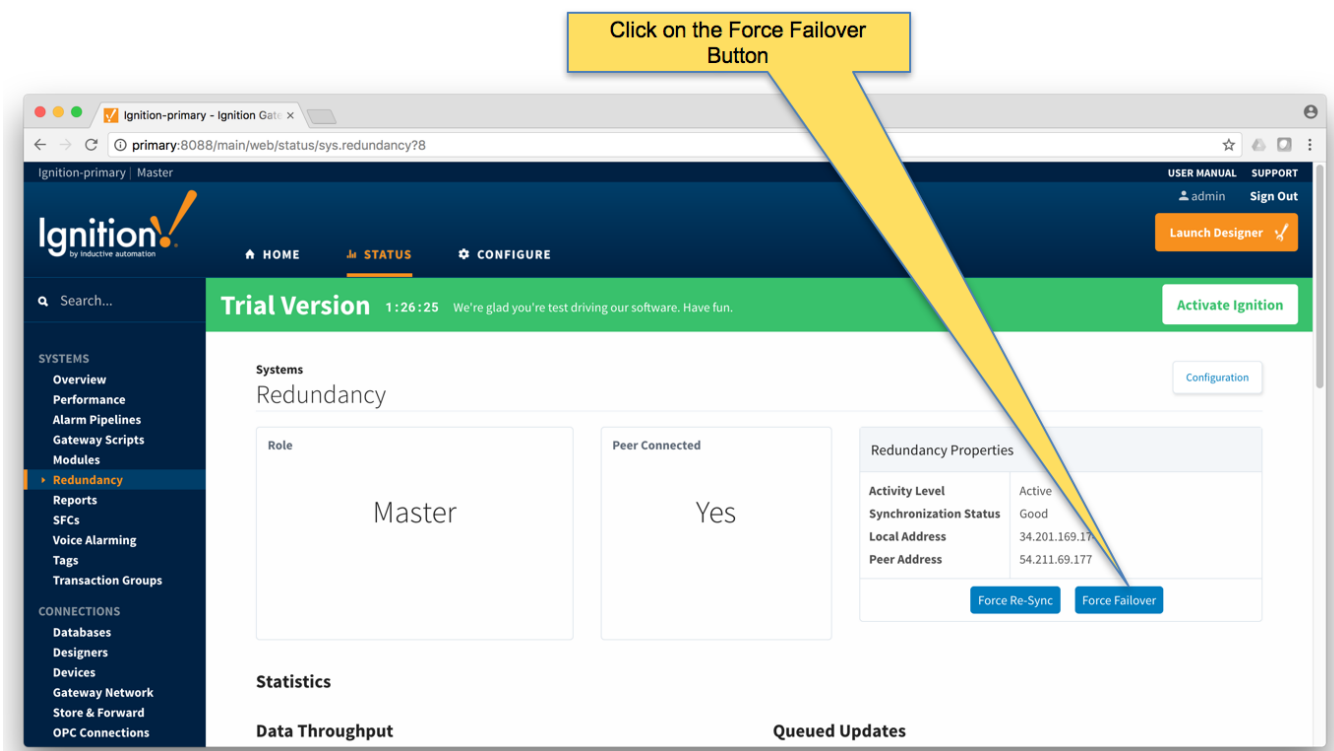
- **Ignition Edge 2**
  - Create the following widgets
    - Label - Ignition Edge 2
    - Label - Redundancy State Label with Tag Path of "[MQTT Transmission]/Transmission Info/Redundancy State"
    - Label - Redundancy Role Label with Tag Path of "[MQTT Transmission]/Transmission Info/Redundancy Role"
    - Label - Edge 2 Tag 1 Label with Tag Path of "[default]/Edge Nodes/Group 1/Edge 2/Device 1/Tag 1"
    - Label - Connected Multi-State Indicator with Tag Path of "[MQTT Transmission]/Transmission Info/MQTT Clients/Group 1-Edge 2/Online"



- Once all three dashboards have been created, save and publish them and close the Ignition Designer windows.
- Now open each of the Ignition client 'Test' projects. With everything running you should see three windows similar to the following.

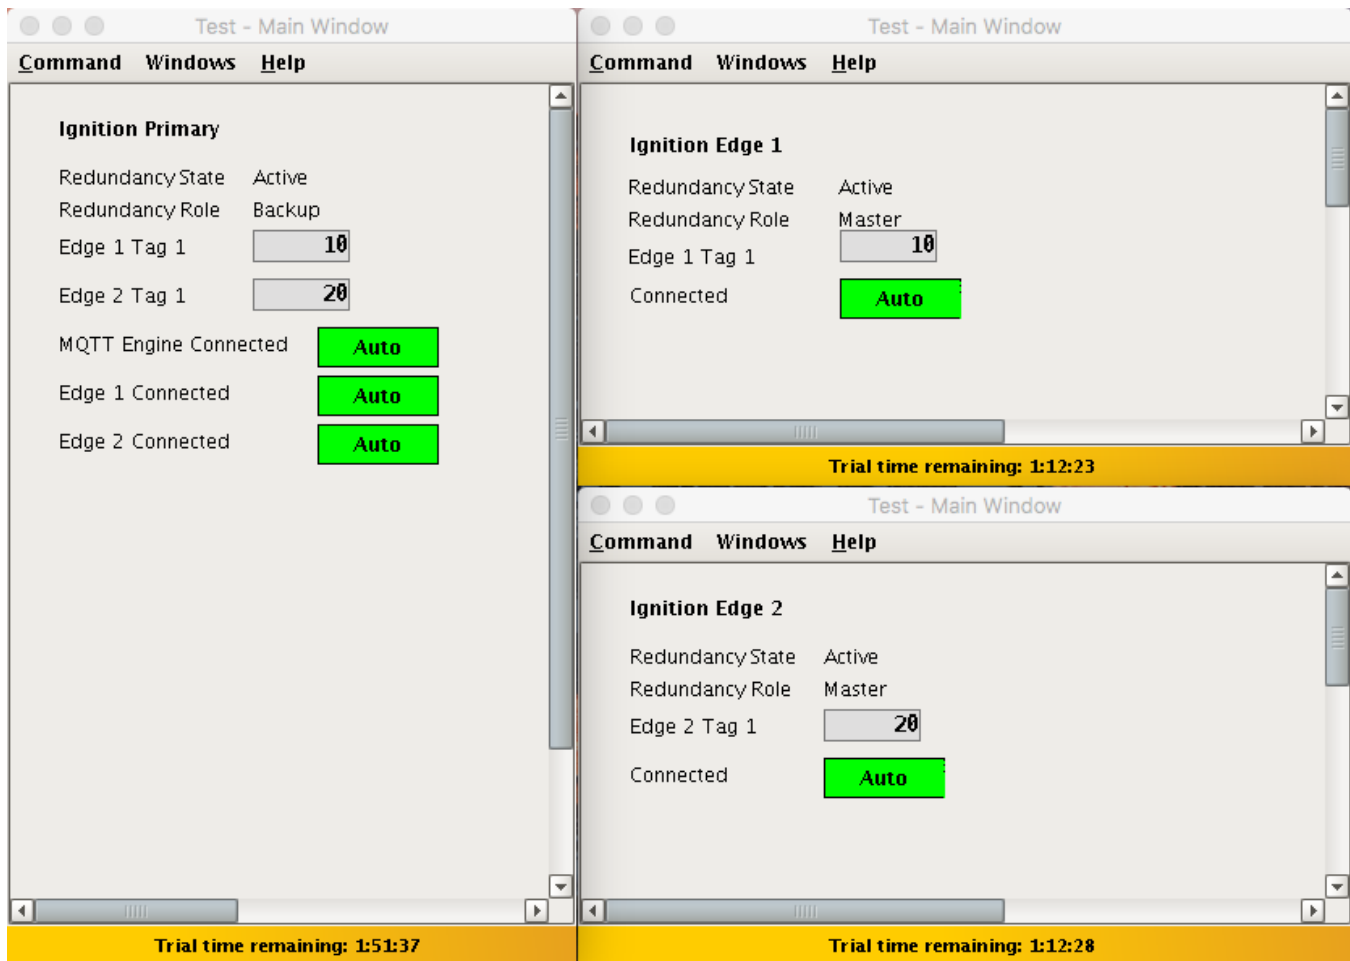


- At this point we can begin failing Ignition instances. From the Status Redundancy page we can use the 'Force Failover' button as shown below. Of course stopping the actual Ignition instance is another option.




- Primary Ignition Failure:** Failing the Ignition Primary instance will cause the following.
  - Ignition Primary will go down and be unreachable
    - This results in all MQTT connections being lost.
  - Ignition Primary Backup will come up and take the place of Ignition Primary
  - MQTT Engine will reconnect on Ignition Primary
  - The MQTT Transmission instances will reconnect to the new MQTT Server (MQTT Distributor) running in Ignition Primary Backup
  - This is all shown in the screenshot below of the Ignition projects. Note all connections are valid and the 'Redundancy Role' of Ignition Primary is now Backup.

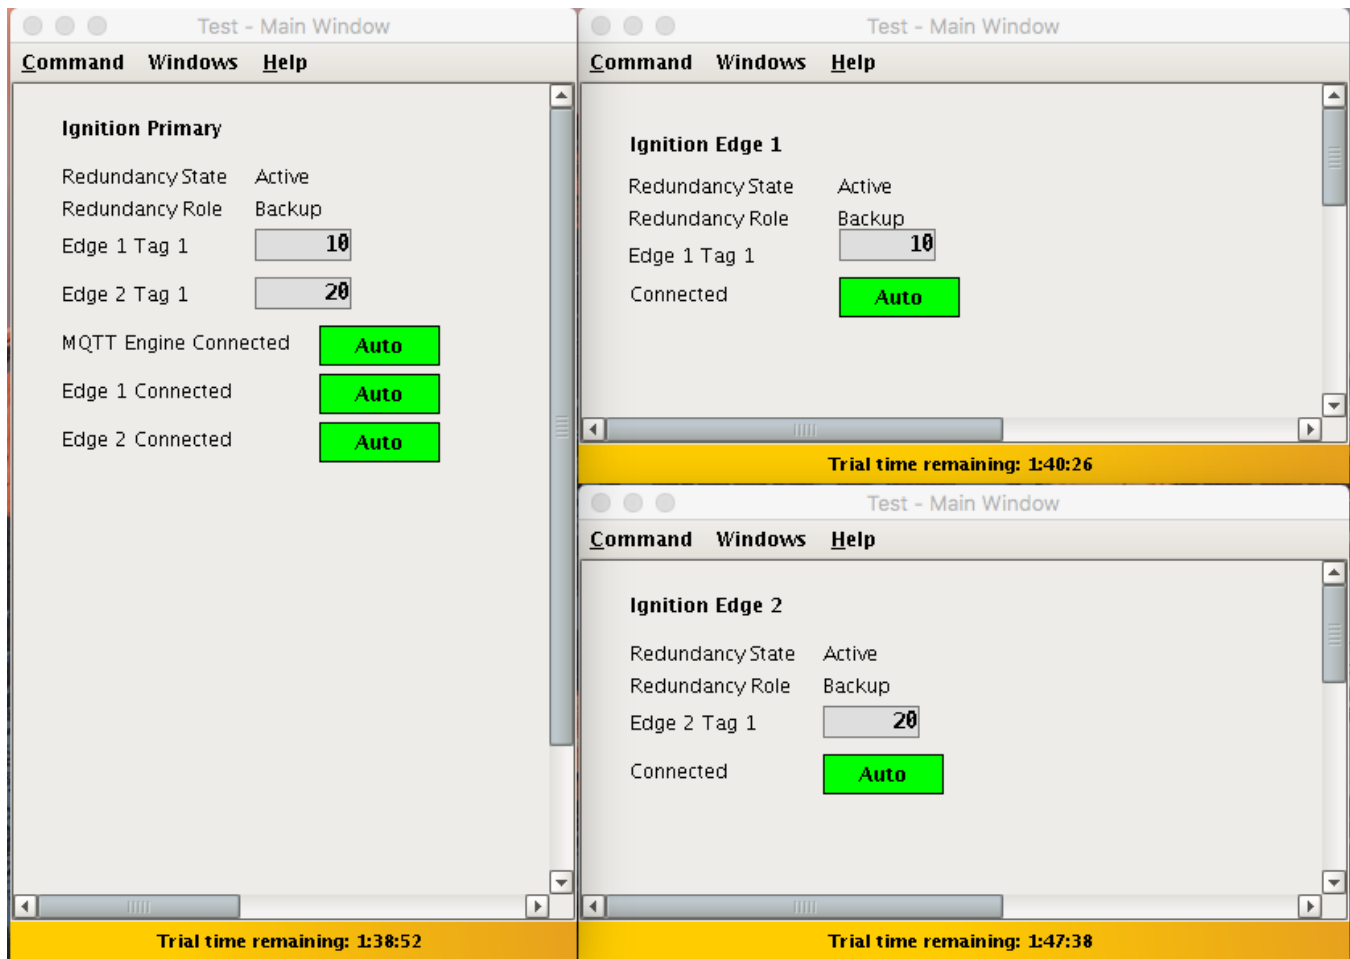




- **Edge 1 Failure:** Failing the Ignition Edge 1 instance will cause the following.
  - Ignition Edge 1 will go down and be unreachable
    - This results in the MQTT Connection between Ignition Edge 1 and Ignition Primary being lost.
  - Ignition Edge 1 Backup will come up and take the place of Ignition Edge 1.
  - The MQTT Transmission instance on Ignition Edge 1 Backup will connect to the MQTT Server (MQTT Distributor) running in Ignition Primary
  - This is all shown in the screenshot below of the Ignition projects. Note all connections are valid and the 'Redundancy Role' of Ignition Edge 1 is now Backup.



- **Failure of all Master Nodes:** Failing all master Ignition instances (Primary, Edge 1, and Edge 2) will cause the following.
  - Ignition Primary, Edge 1, and Edge 2 will all go down and be unreachable
    - This results in all MQTT connections being lost
  - Ignition Primary Backup, Edge 1 Backup, and Edge 2 Backup will all come up and start their MQTT services.
  - The new MQTT Transmission instances on Ignition Edge 1 Backup and Ignition Edge 2 Backup will connect to the new MQTT Server (MQTT Distributor) running in Ignition Primary Backup
  - This is all shown in the screenshot below of the Ignition projects. Note all connections are valid and the 'Redundancy Role' of all instances is now Backup.



To summarize, this tutorial shows how Ignition and the MQTT Modules can be used to create a resilient infrastructure that is able to withstand failures of machines and network connections within the architecture. As noted earlier, this tutorial shows the basic requirements of configuring failover support with Ignition and the MQTT Modules. This can be further improved with additional advanced concepts. Feel free to contact [sales@cirrus-link.com](mailto:sales@cirrus-link.com) for more information.