Ignition Redundancy

Ignition redundancy supports 2-node systems, this means there are two copies of the Ignition Gateway. One node (the Master Gateway) is considered the **master node** and the other (the Backup Gateway) is the **backup node**. Both nodes share the same **state** or configuration. In other words, all projects, Gateway settings, and so on are shared between the nodes. The master node manages the configuration then replicates it to the backup node.

When you have the redundant system in place, you can get detailed status information by going to the **Status** tab of the Gateway webpage as shown here and looking at the system's map:

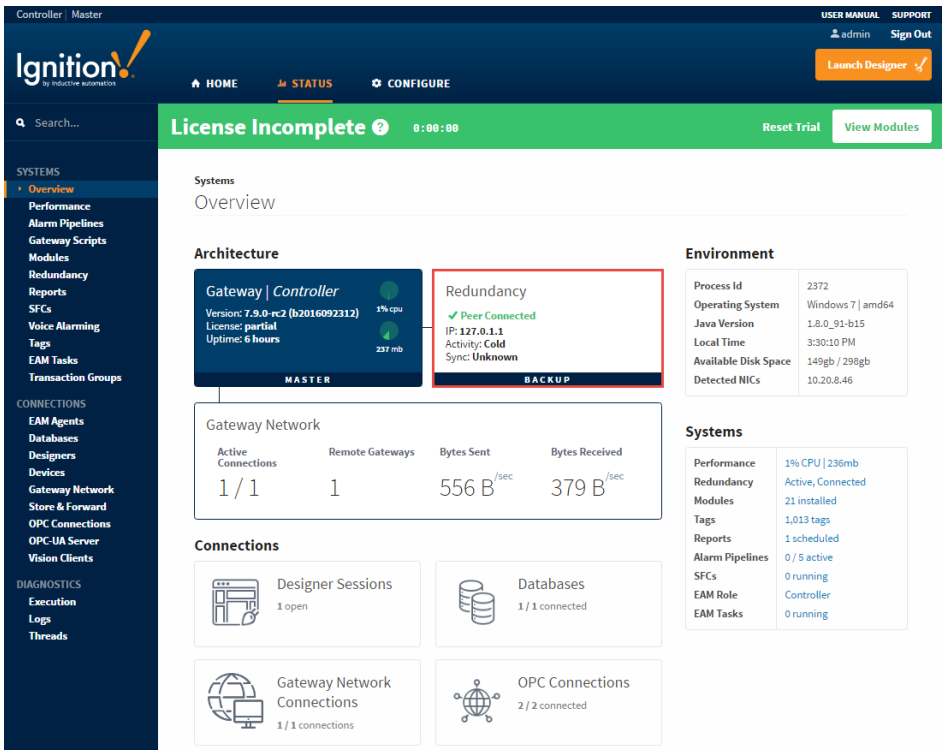

On this page

...

INDUCTIVE UNIVERSITY

How Redundancy Works

Watch the Video

Knowledge Base

Updating or Patching a Redundant Ignition Pair

Learn about updating redundant servers and how to make the process a success.

Link to Knowledge Base Article

Node Communication

The master and backup nodes communicate over TCP/IP. Therefore, they must be able to see each other over the network, through any firewalls that might be in place. By default, all communication goes from the backup to the master node on **port 8750**. Therefore, that port must allow TCP listening on the master machine. You can change the port from the Gateway webpage in the **Redundancy** settings page.

Configuration Synchronization

The master node maintains the official version of the system configuration. You must make all changes to the system on the master Gateway, the backup Gateway does not allow you to edit properties. Similarly, Designer only connects to the master node.

When changes are made on the master, they are queued up to be sent to the backup node. When the backup connects, it retrieves these updates, or downloads a full system backup if it is too far out of date.

If the master node has modules that aren't present on the backup, they are sent across. Both types of backup transfers, **data only** and **full**, will trigger the Gateway to perform a soft reboot.

Runtime State Synchronization

Information that is only relevant to the running state, such as current alarm states, is shared between nodes on a differential basis so that the backup can take over with the same state that the master had.

On first connection or if the backup node falls too far out of sync, a full state transfer is performed. This information is light-weight and does not trigger a Gateway restart.

Status Monitoring

Once connected, the nodes begin monitoring each other for liveliness and configuration changes. While the master is up, the backup runs according to the **stand by activity level** in the settings.

When the master cannot be contacted by the backup for the specified amount of time, it is determined to be down and the backup assumes responsibility. When the master becomes available again, responsibility is dictated by the recovery mode and the master either takes over immediately or waits for user interaction.

System Activity

When a node is active, it runs fully, connecting to any configured OPC servers, and communicating with devices. When it is not active, its activity level is dictated by the settings, either **warm** or **cold**.

In **warm** standby, the system runs as if it were active, with the exception of logging data or writing to devices, allowing for faster fail-over. In **cold** standby, the system connects to all OPC servers but does not subscribe to tag values. The Ignition OPC UA server does not communicate with any device but third party OPC UA servers may still have device connections. This allows the system to standby without putting additional load on the devices and network. Fail-over takes slightly longer, as tags must be subscribed and initialized.

Historical Logging

Historical data presents a unique challenge when working with redundancy because it is never possible for the backup node to know whether the master is truly down or simply unreachable. If the master was running but unreachable due to a network failure, the backup node becomes active and begins to log history at the same time as the master, who is still active.

In some cases this is OK because the immediate availability of the data is more important than the fact that duplicate entries are logged. But in other cases, it's desirable to avoid duplicates, even at the cost of not having the data available until information about the master state is available.

Ignition redundancy provides for both of these cases, with the **backup history level**, which can be either **Partial** or **Full**.

- In **Full** mode, the backup node logs data directly to the database.
- In **Partial** mode, however, all historical data is cached until a connection is reestablished with the master. At that time, the backup and master communicate about the uptime of the master, and only the data that was collected while the master was truly down is forwarded to the database.

Client Fail-over

All Vision clients connect to the active node. When this system fails and is no longer available, they automatically re-target to the other node. The reconnection and session establishment procedures are handled automatically, but the user is notified that they have been transferred to a different node so that they can notify the system administrator that the system may need attention.

In This Section ...