Initial Configuration

Tractor's flexibility and power stems from its ability to be customized. While there are dozens of tuning paramters and opportunities to write custom plug-ins, Tractor has been designed so that no editing of configuration files is necessary to be operational.

If customizations are required, a directory which is readable by the engine's process owner *must first be created* to store the configuration files. Because the engine uses /var/spool/tractor as its default data directory, /var/spool/tractor/config would be a reasonable choice for the configuration directory.

Configuration files are written in JSON.

Key Parameters

There are a few key parameters that administrators should review while setting up a new system in order to ensure that the default settings are desirable. All of the following parameters can be overridden in your site's copy of the tractor.config file. **NOTE: the default values are usually correct for most studios.**

EngineOwner names a specific run-time owner for the engine process. System services start tractor-engine as root in order to acquire protected ports such as port 80; however, the process subsequently reduces its privilege level to EngineOwner. Studios are encouraged to create a special tractor user for this purpose.

ListenerPort specifies the port on which the engine listens for blade, dashboard, and other client connections. The engine is a form of web server and listens on the usual HTTP port 80 by default, which is the simplest configuration for users and site administration. An alternate port number is usually only necessary in cases where there is already another web server running on port 80, or when the engine is not started as root and cannot acquire 80.

EngineDiscovery specifies how the engine is to be discovered by blades and other client programs. Clients will attempt to connect to the hostname tractor-engine by default, or an alternative hostname given by either the TRACTOR_ENGINE environment variable or each tool's --engine=NAME command line parameter. When that hostname cannot be resolved then the EngineDiscovery scheme is tried.

LicenseLocation specifies a path to a license file, or a license server name. This setting *is usually unnecessary* since Tractor will read thepixar. license file in the parent directory of the Tractor installation, which is shared by all Pixar applications.

MaxConcurrentDispatch specifies the maximum number of Tractor license seats that will be consumed by the engine process. A custom setting can be useful when several test engines are running from the same license server. Using the default value, zero, causes tractor-engine to acquire *all* available seats at start-up.

TractorDataDirectory specifies an EngineOwner-writable location where job and user data is stored. By default this is /var/spool/tractor. A different location may be selected for testing, due to constraints in file system permissions, or to point to volumes with larger or faster storage. Some location local to the engine host is usually the best choice. Ensure the data directory is first created and made readable and writable by EngineOwner; this is necessary anyway if the configuration directory is to be a subdirectory, such as with /var/spool/tractor/config.

Example

The following represents an example tractor.config file that would be stored in the site-determined configuration directory. In this example, a site has set up a user named tractor, using its home directory to store data, and listening on a higher port. An in-house license server is used, and only 10 concurrent dispatches will be allowed by the engine.

```
{
  "LicenseLocation": "my-license-server@9010",
  "EngineOwner": "tractor",
  "ListenerPort": 8000,
  "MaxConcurrentDispatch": 10,
  "TractorDataDirectory": "/ssd/tractor/data"
}
```

Further Reading

The full set of engine-specific configuration parameters are described in the Engine Configuration section. This and other site customization documentation can be found in the general Configuration section.