

fe-safe 2017

HIGH-PERFORMANCE COMPUTING WITH *fe-safe*



Contents

1	INTRODUCTION TO DISTRIBUTED MEMORY PROCESSING (DMP)	5
1.1	ABOUT <i>FE-SAFE</i>	5
1.2	ABOUT DISTRIBUTED MEMORY PROCESSING	5
1.3	HOW TO USE THIS MANUAL	6
2	INSTALLATION	7
2.1	MASTER NODE (<i>FE-SAFE</i>) INSTALLATION	7
2.2	ANALYSIS SERVER NODE INSTALLATION	7
2.3	ANALYSIS NODE ADMINISTRATION UTILITY INSTALLATION	11
3	DMP LICENSING	13
4	DISTRIBUTED ANALYSIS CONFIGURATION	15
4.1	MASTER NODE SERVER LIST	15
4.2	USING THE ANALYSIS NODE ADMINISTRATION UTILITY	16
4.3	EDITING THE MASTER NODE SERVER LIST	18
5	MASTER NODE BEHAVIOUR	21
6	ANALYSIS SERVER NODE BEHAVIOUR	23
6.1	SERVER DAEMON AND SERVER DAEMON CONTROLLER	23
6.2	USING THE SERVER DAEMON	24
6.3	USING THE SERVER DAEMON CONTROLLER	24

Trademarks

fe-safe, Abaqus, Isight, Tosca, the 3DS logo, and SIMULIA are commercial trademarks or registered trademarks of Dassault Systèmes or its subsidiaries in the United States and/or other countries. Use of any Dassault Systèmes or its subsidiaries trademarks is subject to their express written approval. Other company, product, and service names may be trademarks or service marks of their respective owners.

Legal Notices

fe-safe and this documentation may be used or reproduced only in accordance with the terms of the software license agreement signed by the customer, or, absent such an agreement, the then current software license agreement to which the documentation relates.

This documentation and the software described in this documentation are subject to change without prior notice.

Dassault Systèmes and its subsidiaries shall not be responsible for the consequences of any errors or omissions that may appear in this documentation.

© Dassault Systèmes Simulia Corp, 2016.

Third-Party Copyright Notices

Certain portions of *fe-safe* contain elements subject to copyright owned by the entities listed below.

© Battelle
© Endurica LLC
© Amec Foster Wheeler Nuclear UK Limited

fe-safe Licensed Programs may include open source software components. Source code for these components is available if required by the license.

The open source software components are grouped under the applicable licensing terms. Where required, links to common license terms are included below.

IP Asset Name	IP Asset Version	Copyright Notice
Under BSD 2-Clause		
UnZip (from Info-ZIP)	2.4	Copyright (c) 1990-2009 Info-ZIP. All rights reserved.
Under BSD 3-Clause		
Qt Solutions	2.6	Copyright (c) 2014 Digia Plc and/or its subsidiary(-ies) All rights reserved.

1 Introduction to Distributed Memory Processing (DMP)

1.1 About *fe-safe*

fe-safe is a powerful, comprehensive and easy-to-use suite of fatigue analysis software for Finite Element models. It is used alongside commercial FEA software to calculate:

- where fatigue cracks will occur;
- when fatigue cracks will initiate;
- the factors of safety on working stresses (for rapid optimisation);
- the probability of survival at different service lives (the 'warranty claim' curve).

Results are presented as contour plots which can be plotted using standard FE viewers, *fe-safe* has direct interfaces to the leading FEA suites.

1.2 About Distributed Memory Processing

Distributed Memory Processing (DMP) is a framework designed to enable distributing analysis jobs to multiple analysis solvers in a network cluster in order to reduce the analysis time. It currently features the following capabilities:

- Automatic load balancing between the analysis nodes
- Fail safe functionality - detecting if an Analysis Node has gone off-line
- The licensing is controlled by a Master Node
- Performance is proportionally scalable

Built into the *fe-safe* application is the Master Node which coordinates distributed analyses; it distributes the analysis to remote Analysis Server Nodes which perform the analysis calculations, monitors the progress of distributed analyses and collects and merges the analysis results. The local machine (the Master Node) can also take part in the analysis using its internal solver.

Models are automatically split into variable sized chunks of items for analysis, depending on the processing speed of individual Analysis Nodes. To minimise the effect of network latency the size of the chunks is optimised to take a minimum of 1 second to analyse. The framework ensures that there are always chunks of data ready for an Analysis Server Node; this is done by queuing chunks on both the Analysis Server Node and the Master Node.

DMP can be used to minimise memory allocation problems with analysis of large models and large, complex load histories. This is due to the fact that the Analysis Server Node does not need to load model or group information, and the Master Node does not require memory for load superimposition or other resource-consuming processes.

The effectiveness of using distributed processing is dependent on the size of the model and the amount of data that is common between analysed FEA elements and nodes. For example a large model with many nodes using a complex superimposition with few datasets will scale better than stress/strain dataset sequences.

Note: *If it is quicker for the Master Node to analyse an FEA node than it is to transfer its the data to an Analysis Server Node (installed in comparable hardware), perform the calculations and then transfer it back to the Master Node, the analysis time will not scale favourably.*

1.3 How to Use This Manual

Chapter 2 (Installation) provides an overview of the actions required in order to install an Analysis Server Node and the Analysis Noda Administration Utility. Licensing is explained in Chapter 0. Chapter 4 describes configuration of distributed analysis and the use of the Master Node and the Analysis Node Administration Utility. Chapter 5 describes individual characteristics of the Master Node itself and Chapter 6 deals with the Analysis Server Node.

Users new to *fe-safe*

Because this manual assumes some familiarity with the *fe-safe* application, it will be necessary to learn a little about the main program first. Work through some of the tutorials in the *fe-safe User Manual*, including at least one demonstrating the use of data from your preferred FEA software, then return here.

Experienced users of *fe-safe*

Experienced users are most likely to refer to chapters 2-6 as a working reference.

2 Installation

2.1 Master Node (*fe-safe*) Installation

The Master Node is integrated into *fe-safe*; see Section 2 of the *fe-safe* Installation and Licensing Guide for installation instructions.

2.2 Analysis Server Node Installation

Operation of the Analysis Server Node is handled by the STL Server Daemon which ensures that the Analysis Server process is always running. The STL Server Daemon is installed as a SYSTEM service on Windows, which allows automatic restarting when the machine reboots. Therefore installation of the Analysis Server Node should be done by a system administrator.

Note: *Although it is possible to run multiple Analysis Server Node services from differing versions of fe-safe on a single machine, it is recommended that this be avoided if at all possible.*

Should an administrator wish to run multiple instances of the Analysis Server Node, strict attention should be made to assigning non-conflicting port numbers that are configured at install time (Figure 5 - Directory for temporary files) and through the parameters in the Service Daemon (Using the Server Daemon Controller 6.3).

These port numbers should then be communicated to the operator of the Host machine when configuring the Master Node configuration (Section 15) to ensure that the correct Analysis Server nodes are used for the software version that is being run as the Master Node.

Installation of an Analysis Server Node is similar to installing *fe-safe* (the installer is launched and the first step is to click Continue in the lower right corner, see Figure 1 - The Installation Welcome Screen).



Figure 1 - The Installation Welcome Screen

- I. During product selection select *fe-safe* Analysis Server Node, see Figure 2 - Installing an Analysis Server Node.

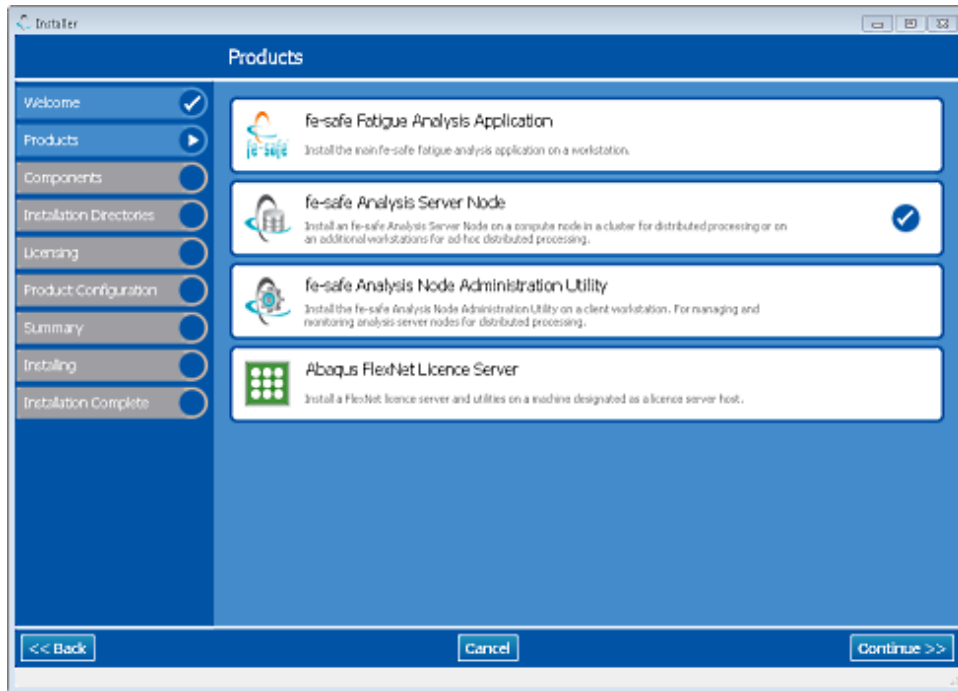


Figure 2 - Installing an Analysis Server Node

- II. Choose if the documentation is to be installed, see Figure 3 – Optional Components.

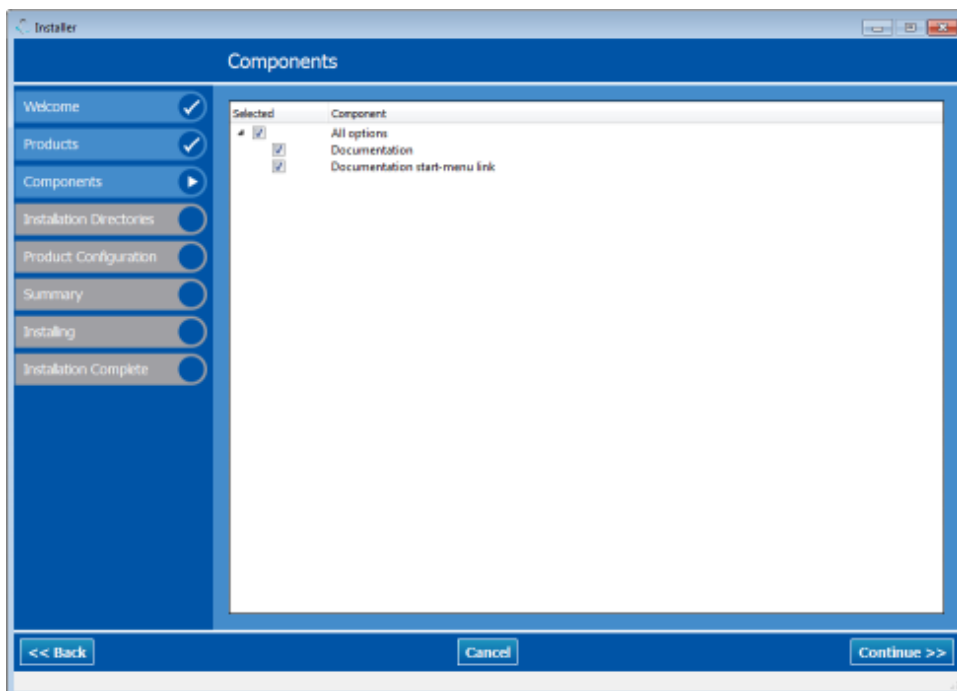


Figure 3 – Optional Components

- III. Select the installation directory, see Figure 4 - Analysis Server Node Installation Directory. If the Analysis Server Node is to be installed together with an existing *fe-safe* application a separate product directory should be used, though using a common root directory is recommended.

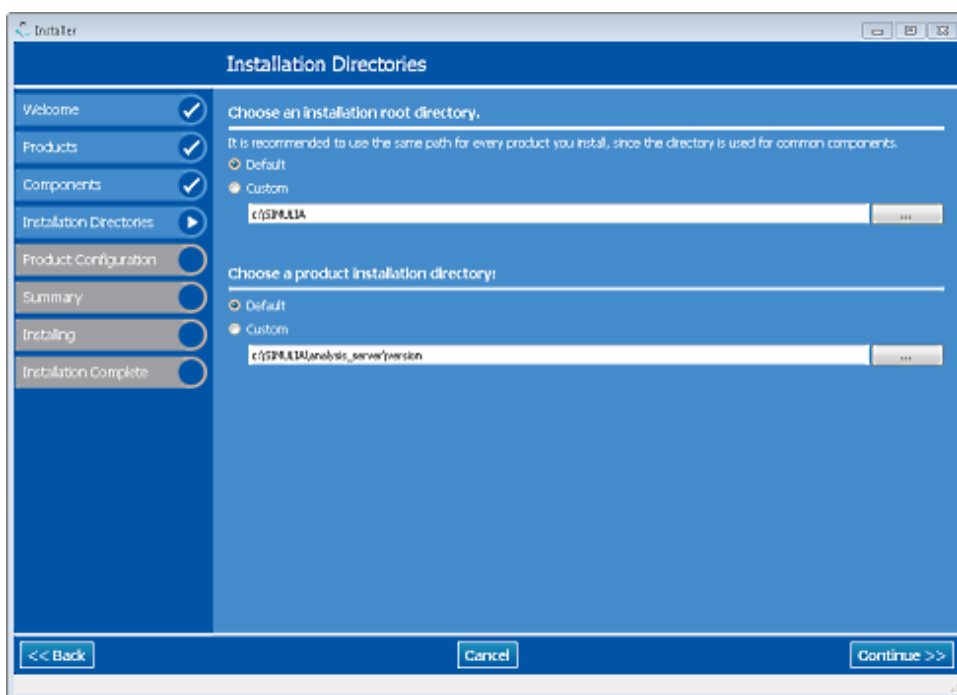


Figure 4 - Analysis Server Node Installation Directory

A directory for temporary working files needs to be specified, and a network port must be specified for the Analysis Server Node to listen on; the default is **18666**, see Figure 5 - Directory for temporary files

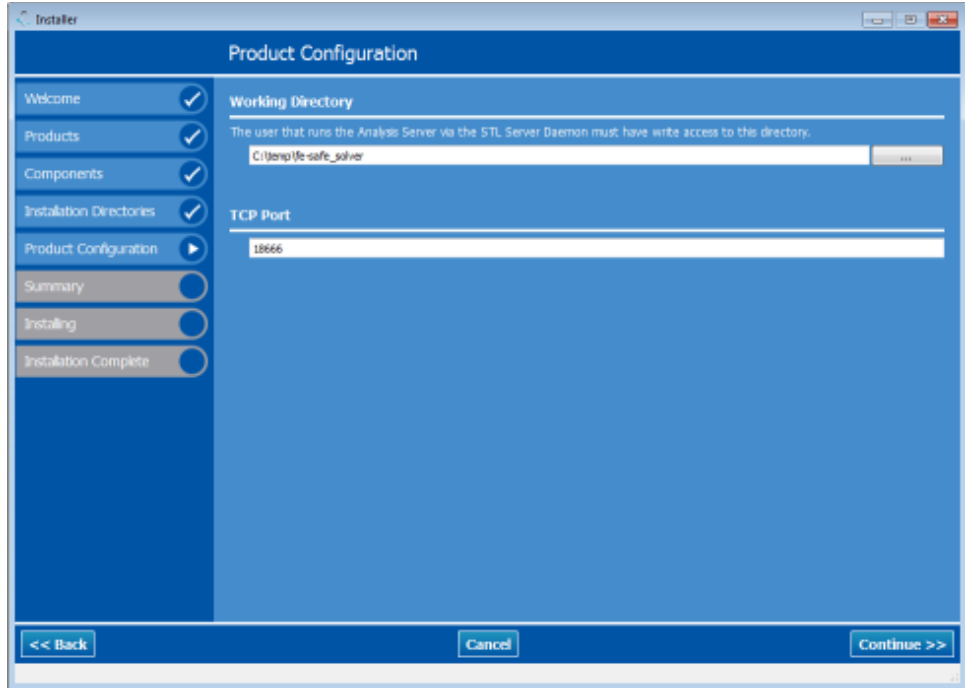


Figure 5 - Directory for temporary files

The installer will try to configure the Windows Firewall to allow access to the specified port, the Analysis Server Node is added to the Windows Firewall list of exceptions. This may not apply to third-party firewalls.

Note: This will have no effect if the use of exceptions is turned off, a check should be made after installation. The Analysis Node Administration Utility can then be used to see if a remote Analysis Server Node can be accessed.

- IV. A summary of the installation selections will then be shown and then the installation process will commence. When the installation process is finished licence code for the machine will be shown, see Chapter 3 for the description of licensing.

2.3 Analysis Node Administration Utility Installation

To install the Analysis Node Administration Utility, launch the installer and select the Analysis Node Administration Utility on the product page, see Figure 6 – Installing the Node Administration Utility. After accepting the licence agreement, a product installation location needs to be specified, see Figure 7 – Node Administration Utility Installation Directory.

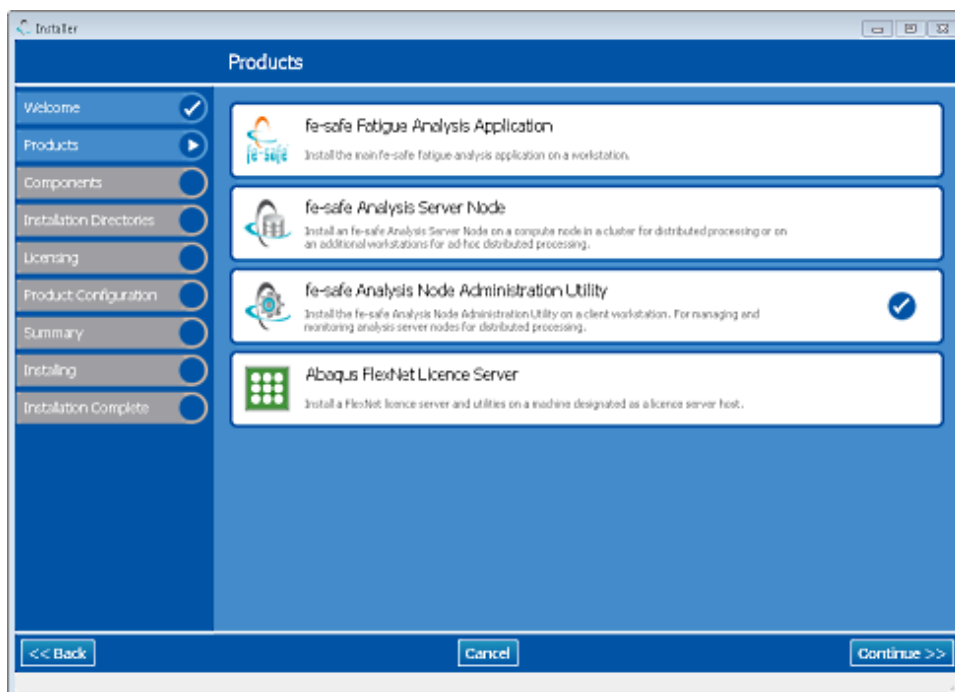


Figure 6 – Installing the Node Administration Utility

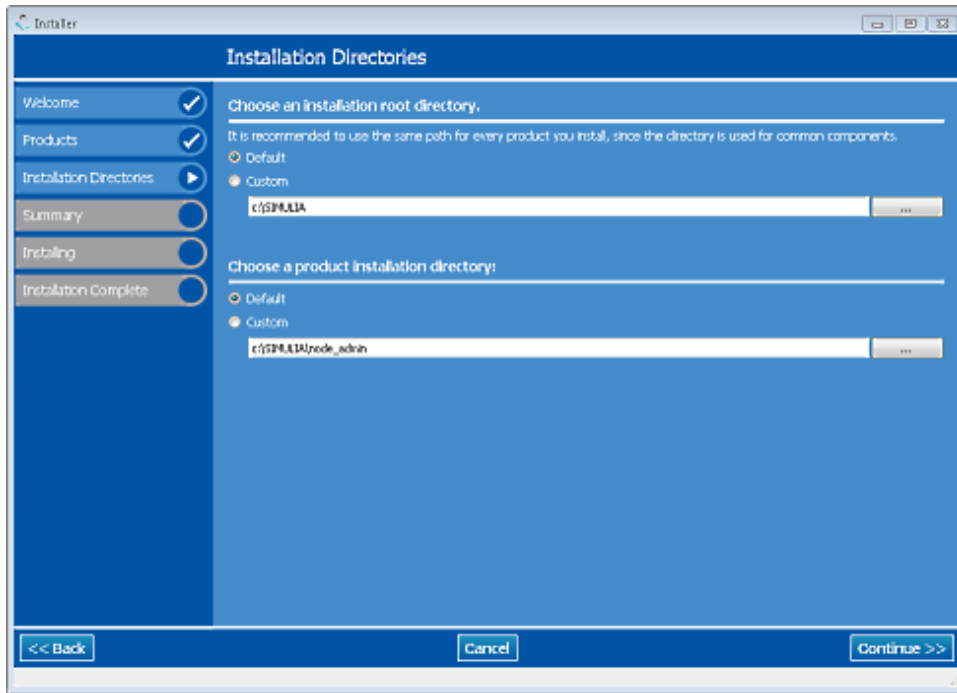


Figure 7 – Node Administration Utility Installation Directory

A summary of the installation selections will then be shown and then the installation process will commence. When the installation process is finished licence code for the machine will be shown, see Chapter 3 for the description of licensing.

Note: as well as a stand-alone application, the Analysis Node Administration Utility is embedded in the fe-safe installation and can be accessed via the menu option **FEA Fatigue >> Distributed Analysis Servers...**, see Figure 8 – Accessing the Node Administration Utility in fe-safe.




Figure 8 – Accessing the Node Administration Utility in fe-safe

3 DMP Licensing

To use the distributed analysis capability a DMP licence token is required. This will enable the distributed processing capability and will allow distribution of analyses based on the number of available *fe-safe* Client and/or Analysis tokens.

When running a distributed analysis, the Master Node will attempt to check out Analysis tokens for all configured analysis nodes. If there are insufficient tokens for a module, the number of Analysis Server Nodes that will be used is reduced, e.g. for an *fe-safe*/TURBOLife analysis on a 10-node cluster, if there are 6 *fe-safe*/Analysis tokens and 4 *fe-safe*/TURBOLife tokens available, only 4 of the Analysis Server Nodes will be used and only 4 *fe-safe*/Analysis tokens will be checked out.

If the Master Node is not configured to use its internal solver, the Master Node will allocate its *fe-safe*/Analysis token to an Analysis Server.

4 Distributed Analysis Configuration

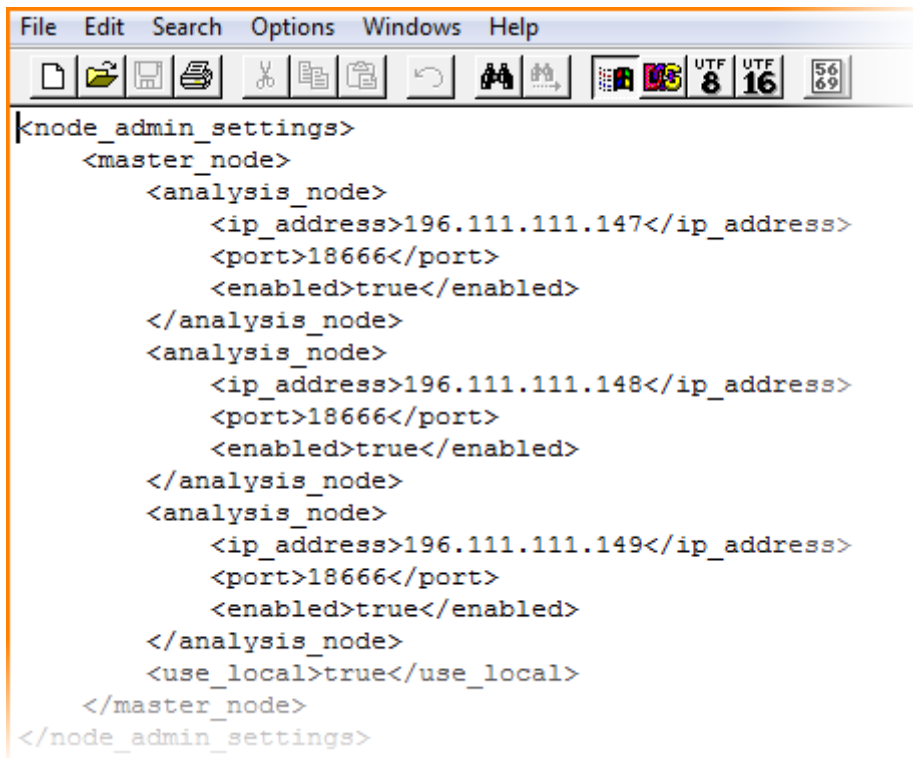
Analysis Server Nodes are run on remote machines and can be part of a cluster along with the Master Node, or where the Master Node is remote to the cluster, or in an ad-hoc network. The machines used for the Master Node and Analysis Servers do not need to be configured identically, nor need the hardware or even the operating systems be the same.

At installation time the port used for the Analysis Server Nodes needs to be specified (refer to Section 2.6 of this manual). The port and the IP addresses of the machines they are running on are used in the Analysis Node Administration Utility to create a list of usable servers.

Note: *Once a DMP configuration has been specified fe-safe will use it for all future analyses until otherwise specified. A DMP configuration file can also be specified using a macro.*

4.1 Master Node Server List

By default the list of Analysis Servers used by the Master Node is stored in a file named `master_node_settings.xml` in the user settings directory. To apply a common configuration to all users the configuration file should be stored in the `/setup` subdirectory of the *fe-safe* installation. To protect the DMP configuration the common `master_node_settings.xml` file can be configured as read-only with limited write access. An example of a file with 3 Analysis Server Nodes is shown in Figure 9 - A sample master node settings file below.



```

File Edit Search Options Windows Help
[Icons: New, Open, Save, Print, Undo, Redo, Copy, Paste, Find, Home, End, Refresh, Zoom, Font, Encoding, Line Numbers, Column Guide]
UTF 8 UTF 16 56 89
<node_admin_settings>
  <master_node>
    <analysis_node>
      <ip_address>196.111.111.147</ip_address>
      <port>18666</port>
      <enabled>true</enabled>
    </analysis_node>
    <analysis_node>
      <ip_address>196.111.111.148</ip_address>
      <port>18666</port>
      <enabled>true</enabled>
    </analysis_node>
    <analysis_node>
      <ip_address>196.111.111.149</ip_address>
      <port>18666</port>
      <enabled>true</enabled>
    </analysis_node>
    <use_local>true</use_local>
  </master_node>
</node_admin_settings>

```

Figure 9 - A sample master node settings file

If *fe-safe* is unable to locate the `master_node_settings.xml`, the default configuration will be used which is to use only the embedded solver. At the start of any analysis, if the DMP configuration is invalid or the analysis is not suitable for distributed solving this default setting is used instead, but **only** for that analysis.

4.2 Using the Analysis Node Administration Utility

The Analysis Node Administration Utility has two uses: to create and edit a list of Analysis Server Nodes for a Master Node to use, and to monitor Analysis Servers. The GUI is shown in Figure 10 - The Node Administration GUI below.




Figure 10 - The Node Administration GUI

Each server in the server list can be enabled or disabled separately, allowing subsets of Analysis Server Nodes to be used from one analysis to the next without having to add or remove them from the list. There are three ways to add new Analysis Servers, see Figure 11 below:




Figure 11 - Adding Server Nodes to the List

- Using an individual IP address and port number.
- Using a range of IP addresses, e.g. 192.168.0.10 to 192.168.0.14. It is assumed that all Analysis Servers in the range were configured to use the same port number.
- Using a subnet mask. The Master Node will try to establish contact with all computers on the subnet using the specified port number and automatically add all Analysis Servers found to the list. This option may require more time to complete depending on the number of computers present on the subnet.

Clicking the **Add to list** button will populate the Servers list, the list must then be saved before the configuration can be used for a fatigue analysis. Additionally, use of the internal solver of the Master Node can be controlled with the **Enable internal processing** option. Figure 12 - A Sample Analysis Server Node List below shows an example of a configured Analysis Server Node list.




Figure 12 - A Sample Analysis Server Node List

The list contains four analysis nodes, of which one has been disabled (first column). The second column lists IP addresses (or machine names) and port numbers of all configured nodes, the third column shows their activity status and the last their connection status. Only the first node in the list is connected, the other nodes have either not yet been connected or are waiting to establish contact with the master node, the internal solver will also be used in the analysis. If the analysis was started at this point, *fe-safe* analysis would be executed on 2 nodes: 1 local node (enable internal processing) + 1 remote node.

Note: After adding new nodes to an existing configuration it may be necessary to restart the existing nodes in order to reset the configuration and ensure that it is ready to use.

4.3 Editing the Master Node Server List

The list of Analysis Server Nodes used by the Master Node is stored in a file named `master_node_settings.xml`, the default configuration is to use only the embedded solver, see Chapter 0 above. The servers list can be edited in three different ways:

- 1) From the *fe-safe* menu option **FEA Fatigue >> Distributed Analysis Servers...** This starts the Analysis Node Administration Utility which will automatically allow editing of the `master_node_settings.xml` file.
- 2) A stand-alone version of the Analysis Node Administration Utility can be used to load, edit and save the servers configuration list.

- 3) Manually editing the `master_node_settings.xml` file in a text editor and then loading it using the Analysis Node Administration Utility.

Note: If the `fe-safe master_node_settings.xml` is modified outside of the application, its contents will automatically be reread. However, DMP configuration changes will only ever take effect the next time an analysis is started – a running analysis will not change the servers it is using.

Enabling or disabling nodes can be done in two ways:

- 1) Using the Analysis Node Administration Utility, right-click the relevant server entry and select the required option, see Figure 13 below,




Figure 13 - Enabling and disabling Analysis Server Nodes using the Admin Utility

- 2) Edit the `master_nodes_settings.xml` file manually as shown in Figure 14 below:

```
<root>
  <master_node>
    <analysis_node>
      <ip_address>196.111.111.147</ip_address>
      <port>18666</port>
      <enabled>true</enabled> ← enabled
    </analysis_node>
    <analysis_node>
      <ip_address>196.111.111.148</ip_address>
      <port>18666</port>
      <enabled>>false</enabled> ← disabled
    </analysis_node>
    <analysis_node>
      <ip_address>196.111.111.140</ip_address>

```

Figure 14 - Enabling and disabling Analysis Server Nodes in the .xml file

Note: Analysis progress will be displayed in the `fe-safe` message log. If the Master Node is experiencing difficulties or interruptions in communicating with the Analysis Server Nodes this will be reported in the log. The analysis will continue, but the solution time will suffer. For more details see Chapter 5.

5 Master Node Behaviour

A DMP analysis can be run through the fe-safe GUI or by running fe-safe from the command line. The machine that initiates an analysis is referred to as the Master Node. Each connection to an Analysis Server Node is handled concurrently to prevent latency or time-outs.

- If the Master Node cannot connect to *any* of the specified Analysis Server Nodes, the analysis cannot begin and will be aborted immediately.
- If an analysis is running and the Master Node is terminated, the analysis will be aborted immediately.
- If an analysis is running and all Analysis Server Nodes become disconnected and a time-out of at least 15 seconds is reached, the analysis will be aborted.
- If an analysis is running and one or more Analysis Server Node becomes disconnected the analysis will continue and the Master Node will attempt to reconnect to the Analysis Server Node(s) for approximately 2 minutes. If the disconnected Node(s) reconnects within that time the analysis will continue where it was interrupted, otherwise the Master Node will distribute the analysis data to other nodes.
- The usual time-out when trying to connect to an Analysis Server Node is 5 seconds. A time-out of 10 seconds is imposed on all remote calls including sending, remote processing and receiving a response.
- When the Master Node is not uploading analysis data or downloading results, it polls the Analysis Server Nodes approximately every second to determine if the Analysis Server has encountered connection issues when communicating with the Master Node.

6 Analysis Server Node Behaviour

Analysis Server Nodes can be installed and run on cluster nodes or on remote machines in an ad-hoc network. The Analysis Server handles analysis requests and solves portions of the analysis allocated to it by the Master Node.

The Analysis Server is hosted by the STL Server Daemon which ensures that the Analysis Server process is always running. The STL Server Daemon is installed as a SYSTEM service on Windows which allows automatic restarting when the machine reboots, therefore installation of the Analysis Server should be done by a system administrator.

On Linux, the STL Server Daemon starts in the background when it is first installed, but to restart it after a machine reboot, the script `start_server_daemon.sh` must be called. Adding the script to individual user startup or login scripts is **not** recommended, as the user would have to log in for the STL Server Daemon and Analysis Server processes to be started and the call to the script will require the same privileges as it did at installation.

The Analysis Server uses a working directory to save temporary working files and analysis and connection logs. To modify the location of the working directory after installation, use the **Server Daemon Controller** and modify the **Server Parameters** as required (see Chapter 6.3).

An analysis log file named `analysisNode.log` can be found in the Analysis Server working directory. This file records connections to Master Node(s), progress of analyses and any diagnostic information if requested.

Cached data for an analysis is saved into a cache subdirectory in the Analysis Server working directory (e.g. `\fe-safe_solver\cache\1\...\current_chunk_file_name.raw`).

Each analysis is provided with a different subdirectory named by the ID assigned to the analysis by the Analysis Server process. The cached files are only needed while the Analysis Server process is running and are deleted once the analysis has completed and all chunks transferred back to the Master Node.

Note: *If an Analysis Node is shown as green in the Servers list (enabled and connected), see Figure 12, but time-outs are being reported in the Master Node log file at the beginning of an analysis, it is possible that the Analysis Server Node needs to be restarted as described in Chapter 6.3.*

6.1 Server Daemon and Server Daemon Controller

The Server Daemon is common to both the Analysis Server and the STL Licence Server. It runs as a service and maintains any registered servers. By default it is installed in a subdirectory `common/server_daemon/` under the installation root directory specified during the main installation, e.g. `C:\SIMULIA` or `/usr/SIMULIA`.

6.2 Using the Server Daemon

After installation, the `st_service(.exe)` process should be running, check using the windows task manager or by issuing the `top` or `ps -ef | grep st_service` command in a Linux terminal.

Use the scripts `stop_server_daemon.bat` or `./stop_server_daemon.sh` to stop and unregister the service and `start_server_daemon.bat` or `./start_server_daemon.sh` to register and start it. Once registered on Windows, the service can also be stopped and started in the services list (where it is labelled as **STL Server Daemon**).

For testing purposes the Server Daemon can be run as a standalone process using a `st_service -exec` command in a windows command prompt. In this environment, simply use **Ctrl+C** to terminate it.

6.3 Using the Server Daemon Controller

To see the current Server Daemon configuration, launch the Server Daemon Controller (`st_service_controller.exe`) to display the current settings. By default it is installed in the `.\common\server_daemon\` folder of your *fe-safe* installation directory and can be accessed as follows:

- On Windows
 - 1) Go to **Start >> All Programs >> fe-safe.x >> Start or Stop Analysis Server**
 - 2) Double-click `st_service_controller.exe` in its installation location.
- On Linux
 - 1) Open a terminal window and issue the following command:
`./../common/server_daemon/st_server_controller`
 - 2) Single-click `st_service_controller.exe` in its installation location.

Note: On Windows platforms with a virtual store (e.g. Windows 7) ensure that the controller is started 'as administrator' (right-click on the Start menu option or the file and select 'Run as administrator'), otherwise the changes will be made in the virtual registry and therefore NOT affect the Server Daemon configuration.

The Server Daemon Controller window is shown in Figure 15 below. If the controller is launched and does not detect that the Server Daemon is running, the user will be prompted to install it. Once the Server Daemon Controller is running, the hosted servers are listed in the drop-down list at the top of the dialogue, see Figure 15 below.

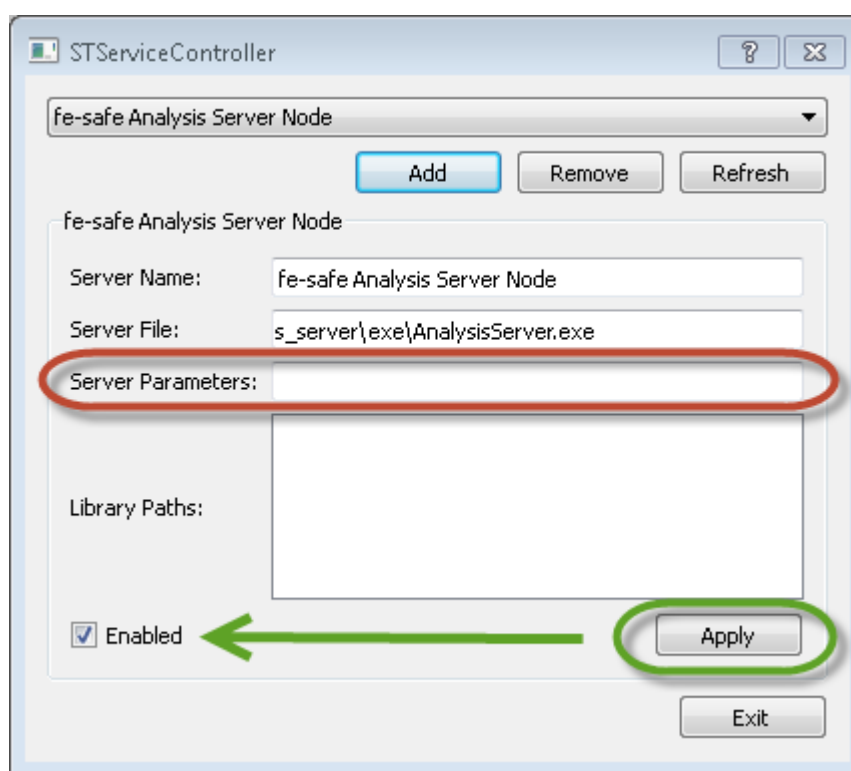


Figure 15 - The Server Daemon Controller GUI

To turn the Analysis Server on or off, select or deselect the **Enabled** checkbox, and click the **Apply** button to confirm the change. This is the only way to start or stop an Analysis Server, if the process is terminated unexpectedly (e.g. using a Task Manager) the Server Daemon will immediately start a new instance of the Analysis Server.

If the **Server Parameters** are not set, they should be configured as follows:

-p <port_number> -w <WORKING_DIR>

The dialogue should then look similar to Figure 16.

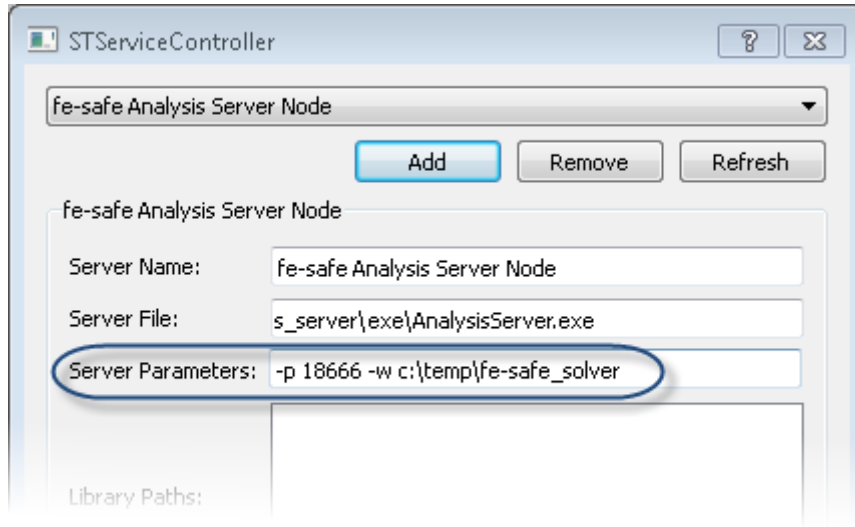


Figure 16 - Port and Write Parameters in the GUI

These parameters are used to configure the Analysis Server Node; the network port to be used (<port_number>, default is **18666**) and the temporary working folder (<WORKING_DIR>, e.g. c:\temp\fe-safe_solver). If the path contains spaces double quotes must be used to enclose it. If the working path does not exist when the Analysis Server starts it will be created automatically. If the Server Daemon and its servers are **NOT** to be run with **administrator** privileges, the controller must be run under the same account as the Server Daemon (and thus its guest servers). If the Server Daemon and its servers **will** be run with **administrator** privileges, the controller will need the same privileges, though not necessarily with the same account. If the controller cannot access the administrator settings, a warning will be given indicating this in case the application has been run with insufficient privileges.