

Executing Performance Tests

User's Guide



Version 6.0

Copyright

Copyright © 2004, 2005, 2006, 2007, 2008, 2009 Reflective Solutions Ltd.
All rights reserved.

Restricted Rights Legend

This software and documentation is subject to and made available only pursuant to the terms and conditions of the Reflective Solutions Ltd. license agreement and may be used or copied only in accordance with the terms of that agreement. It is against the law to copy the software except as specifically allowed in the agreement. This document may not, in whole or in part, be copied, photocopied, reproduced, translated or reduced to any electronic medium or machine-readable form without prior consent, in writing, from Reflective Solutions Ltd.

Information in this document is subject to change without notice and does not represent a commitment on the part of Reflective Solutions Ltd.

THE SOFTWARE AND DOCUMENTATION ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER REFLECTIVE SOLUTIONS LTD. DOES NOT WARRANT, GUARANTEE OR MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE, OF THE SOFTWARE OR WRITTEN MATERIAL IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY OR OTHERWISE.

Trademarks or Service Marks

jTune™ and StressTester™ are registered trademarks of Reflective Solutions Ltd.

All other trademarks are the property of their respective companies.

Contents

Introduction	4
Designing a Performance Test	5
Test Goal	5
Information Requirements	5
Injection Locations	6
User Journeys	6
Resource Monitors	7
Location Load Profiles	7
Considering Injection Locations	8
Run Components Workspace	9
Run Components Navigation Tree	9
Right-Click Menus	10
Injectors	12
Creating an Injector Configuration	12
Advanced Settings	13
SSL	14
Network Settings	14
Process Settings	15
Debug	16
Monitor Agents	17
Creating Test Run Configurations	18
User Journey Configuration	19
Resource Monitor Configuration	20
Starting Run Components	22
Starting Injectors	22
Starting Monitor Agents	22
Executing a Performance Test	24

Introduction

This user guide details the steps that you should consider in order to design and execute a performance test using StressTester™.

The design of the performance test is the last stage in ensuring that the load created during the test will be realistic when compared to the real World.

The sections of this guide cover everything that you should consider when designing a test, and then how you can configure and execute the designed test using StressTester™.

Designing a Performance Test

The following sections detail the areas that you should consider when designing your performance test.

Test Goal

The key thing to consider before configuring and executing a performance test is:

[What is the performance test's purpose?](#)

Typical performance test purposes include the following:

Load Test: what response times will the users receive at a certain load and transaction mix?

Stress Test: what is the maximum number of users that can be supported before the system breaks?

In this scenario, “break” may refer to the application failing, an unacceptable percentage of requests returning errors, or unacceptable response times occurring.

Spike Test: how does the application respond to a sudden short burst of high load?

Stability Test: are response times and system resource usage the same if a steady load runs for a long period of time (for example, for 24 hours)?

Information Requirements

Once the purpose of the test is established, the main reason for executing a performance test – to obtain information about the performance profile of the application under test – can be addressed.

[What is the information the tester/team/IT department/business requires from the performance test?](#)

In what format should the information be provided (graphs, data tables or both)?

The list of potential recipients of the information is provided on purpose; StressTester™ can be used to provide information to many different audiences from the same performance test due to its flexible and customizable analysis capabilities.

Please refer to the [Analysing Performance Test Results](#) user guide for more details of these capabilities.

Injection Locations

The next consideration is to determine from where the load should be generated.

StressTester™ has been designed to support distributed load simulation from multiple locations with little extra configuration effort required.

It is also licensed in such a manner that once you have purchased / rented a certain load, you can distribute this across as many injection locations as you like.

This flexibility, coupled with other StressTester™ features, allows you to simulate the correct number of users, at the correct locations, running correct User Journeys – truly realistic testing.

User Journeys

An obvious design consideration is which User Journeys will run during the performance test.

In addition, if load is being simulated from multiple locations, you should consider which User Journeys will run at which locations.

For example, it may be that your web site's purchase transaction is run from many places, but the restricted back office transaction is only run from your head office.

Resource Monitors

Depending on the purpose and information requirements of the performance test, you may wish to monitor the components of the system under test in order to be able to diagnose any performance and/or scalability problems that may be found.

If monitoring is needed, you need to consider the Resource Monitors you will require, and which aspects of the system they should monitor.

Please refer to the [Monitoring Underlying System Resources](#) user guide for more details on StressTester™'s system monitoring capabilities.

Location Load Profiles

The final consideration is:

[What load should initially be simulated, for each User Journey, at each location?](#)

In addition, if you intend to execute a stress test and therefore increase the load over time:

[How should these loads increase during the performance test?](#)

Once again, it is important that this information is as correct as possible to ensure the performance test is realistic and therefore the information obtained of use to the intended audience.

Considering Injection Locations

The StressTester™ component that simulates users interacting with the application under test is called an “Injector”.

Injectors can be thought of as synonymous with locations – typically one Injector will be installed at each location from which load will be simulated.

The granularity of “location” is dependent on the requirements of the performance test.

It could be that an Injector is installed on each floor of an office building to check the network components on each floor can handle that floor’s load.


Alternatively, locations may be of a coarser granularity, for example: New York, Sydney, Berlin and London.

At each location identified, the Injector software needs to be installed and configured.

Please refer to the [Installation](#) user guide for details of how to install Injectors.

Run Components Workspace

All tasks related to configuring the components used to execute a performance test and creating the test profile – known as the Test Run Configuration – are defined in the Run Components workspace.



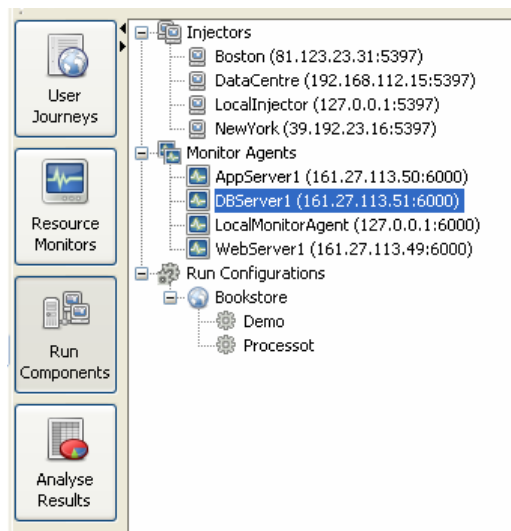
To access this workspace, click the “Run Components” button on the left hand side of the StressTester™ user interface.


Run Components Navigation Tree


When the workspace is entered, the workspace navigation tree and introduction page are displayed.


Note: It is highly recommended that you read the introduction page the first time you use the workspace – it contains valuable information.

By clicking a node in the navigation tree, the introduction page will be changed to a data panel - allowing the selected node’s information to be viewed and updated (StressTester™ license permitting).



In the screen shot shown above, there are four Injectors (represented by the  icon) – indicating the different locations from which load can be generated.

There are also four Monitor Agents (represented by the  icon). Each is responsible for retrieving platform metrics (CPU utilization, memory occupancy etc.) from one or more infrastructure components e.g. web server, load balancer, database server, etc.

In addition, each application has a node under the Run Configurations node, and by expanding a node, you can see the Test Run Configurations that are defined for the application (represented by the  icon).

Right-Click Menus

The majority of navigation tree nodes within the Run Components workspace have context sensitive right-click menus – providing easy access to the relevant StressTester™ functionality.

The sections below detail the different functionality available on the different tree nodes.

“Injectors” Node

Menu Item	Description
Create Injector	Creates a new Injector configuration.

Injector

Menu Item	Description
Delete	Deletes the Injector configuration.

“Monitor Agents” Node

Menu Item	Description
Create Monitor Agent	Creates a new Monitor Agent configuration.

Monitor Agent

Menu Item	Description
Delete	Deletes the Monitor Agent configuration.
Clone	Duplicates the Monitor Agent configuration in a new configuration.

Application

Menu Item	Description
Create Run Configuration	Creates a new Test Run Configuration.

Test Run Configuration

Menu Item	Description
Execute	Start a performance test as specified by the Test Run Configuration.
Delete	Deletes the Test Run Configuration and all results relating to it. Note: When you delete a Test Run Configuration you will no longer be able to view test results created using the configuration.

Injectors

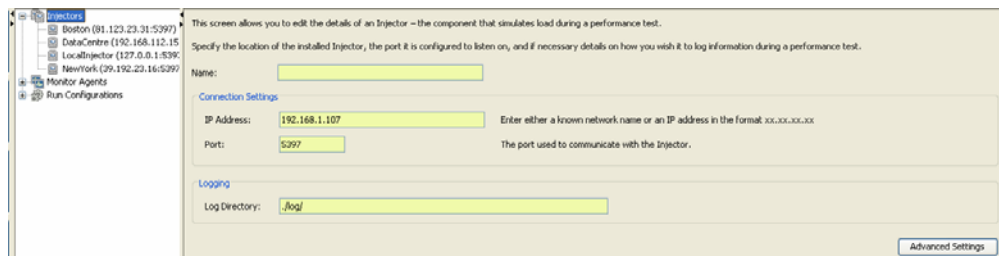
As stated earlier in this guide, an Injector is synonymous with a location from which simulated load should be generated.

Creating an Injector Configuration

You only need to define an Injector once within a StressTester™ installation. After this, it can participate in as many Test Run Configurations as you like.

To define an Injector within StressTester™, right-click on the Injectors node in the Run Components workspace and then select the “Create Injector” option.

A screen similar to the one shown below will be displayed.



The fields on the screen are as follows:

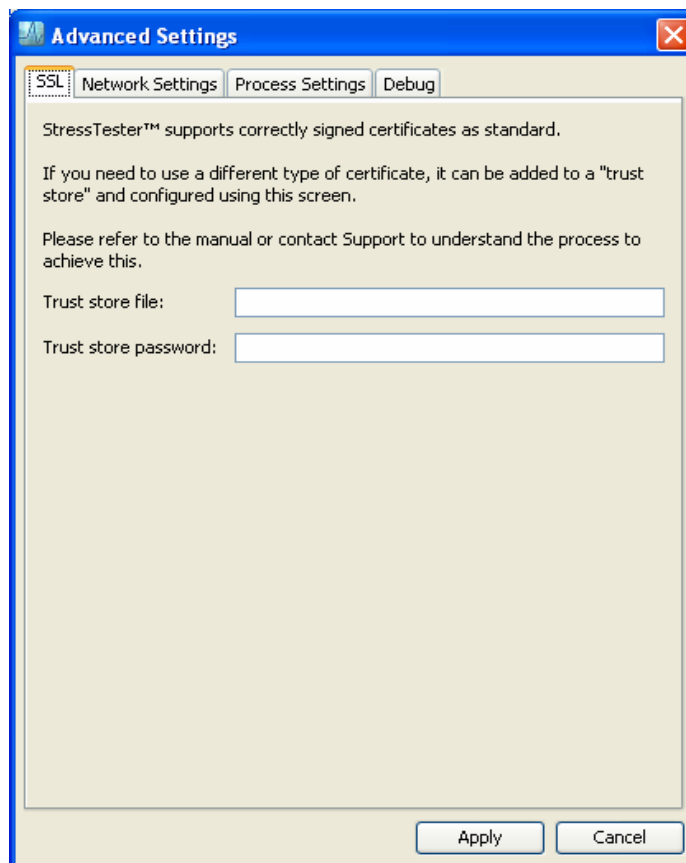
Field	Description
Name	A unique name to give to this Injector. It is recommended names reflect the location of the Injector.
IP Address	The IP address (in the format xx.xx.xx.xx) of the machine hosting the Injector.
Port	A port that can be used to send commands to the Injector. The default is 5937 – this is the port reserved by the IANA ¹ for StressTester™ Injectors.

¹ Internet Assigned Numbers Authority.

Field	Description
Log Directory	The path of a directory (either relative to the starting location of the Injector, or an explicit path name) within which this Injector will write log and error files.

Advanced Settings

By clicking the “Advanced Settings” button on the Injector properties screen, you can access additional configuration fields.



The fields on the different tabs on this pop-up are described in the sections below.

SSL

Secure communications between the server and client require the client to trust the server to be what it claims to be.

This is achieved using 'server certificates' that identify the server and are signed by a trusted third-party authority.

StressTester™ supports by default all certificates correctly signed.

However, if it is necessary to use a certificate that is not signed in such a way, this tab can be used to describe the file containing the certificate to be used, and the password used to read the file.

The table below describes the fields on this tab.

Field	Description
Trust store file	The file containing the certificate to be trusted by StressTester™.
Trust store password	The password used to access the file.

Please contact Reflective Solutions Support – support@reflective.com – for details on how to configure the “trust store”.

Network Settings

This tab allows you to configure details of a proxy that StressTester™ should use to access the application under test.

The table below describes the fields on this tab.

Field	Description
Proxy Address	The IP address of the proxy to use.
Proxy Port	The port on which the proxy 'listens'.

Process Settings

StressTester™ Injectors are implemented in Java. This means an Injector runs on one or more Java Virtual Machines (JVMs).

By default, StressTester™ Injectors use the host machine's default Java implementation, and a new Injector process is spawned for every 200 users simulated.

This tab allows these and other default parameters to be overridden.

Note: These are advanced parameters and should only be updated by experienced users with a detailed knowledge of JVMs, or when you are instructed to do so by Reflective Solutions Support.

The table below describes the fields on this tab.

Field	Description
Injector Process JVM	The full path to the Java command that should be used to run the Injector process(es). By default, the host machine's default Java implementation will be used.
Supplementary Classpath	Additional items that should be placed on the Injector process' Java class path.
Memory: Initially Assigned	The amount of memory that the JVM should initially request from the operating system when each Injector process starts. By default, Java's standard setting is used. This varies depending on the Java implementation.
Memory: Maximum Limit	The maximum amount of memory that the JVM running the Injector process can request during a performance test. By default, Java's standard setting is used. This varies depending on the Java implementation.
Supplementary Arguments	Java runtime arguments that will be passed to the JVM when the Injector process is started.

Field	Description
Maximum Threads	<p data-bbox="732 222 1390 352">In order to avoid JVM or operating system threading issues, a StressTester™ Injector will create a clone of itself when it has allocated a set number of threads.</p> <p data-bbox="732 394 1390 457">In this context, a Java thread is synonymous with a simulated user.</p> <p data-bbox="732 499 1390 562">Adjusting this field changes when clone processes are created.</p> <p data-bbox="732 604 1390 699">Note: The original Injector process manages all its clones and ensures they are killed at the end of a performance test.</p>

Debug

Under the instruction of Reflective Solutions Support, you may be asked to configure debug parameters in order that more information about a reported issue may be captured.

When you have updated the Injector properties, click the “Save” toolbar button to persist the changes.

Monitor Agents

As described earlier in this guide, a Monitor Agent is responsible for monitoring the system under test during a performance test.

A Monitor Agent can be installed either local, or remote, to the system it is monitoring.

Monitor Agents can obtain information on either hardware or software components. So, for example, the CPU utilization as well as the number of sessions running inside a database may be monitored.

There is no limit to the number or type of resources that may be monitored.

Please refer to the [Monitoring Underlying System Resources](#) user guide for details on how to install, configure and start Monitor Agents.

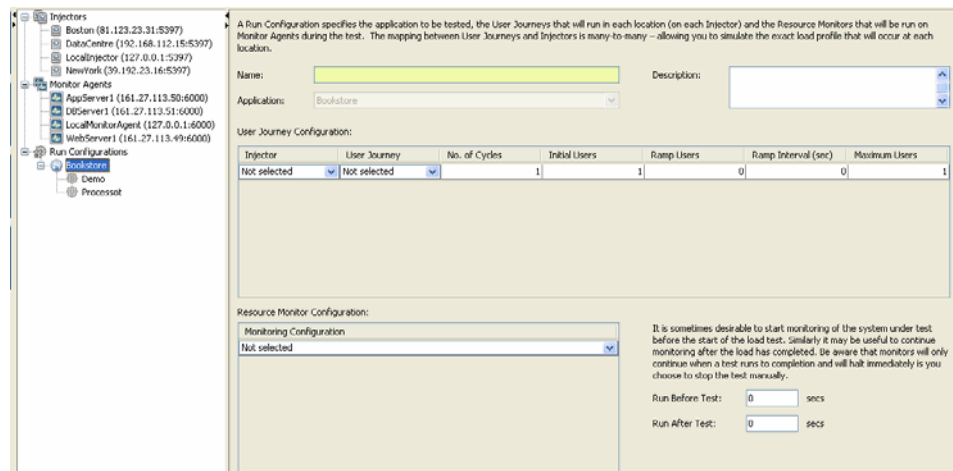
Creating Test Run Configurations

When you have designed your performance test, configured the Injectors and Monitoring Configurations that will be used during the test, you can create a Test Run Configuration within StressTester™.

Note: please refer to the [Monitoring Underlying System Resources](#) user guide for details on how to create Monitoring Configurations.

To create a new Test Run Configuration, right-click on the relevant application node and then select the “Create Run Configuration” option.

A screen similar to that shown below will be displayed.



The fields on the screen are as follows:

Field	Description
Name	A unique name for the Test Run Configuration.
Description	An optional free text description of the Test Run Configuration. It is recommended you use this field to record the purpose of the configuration i.e. the goal of the performance test.

Field	Description
Application	The application that will be tested by the Test Run Configuration. This field is read-only on this screen.
User Journey Configuration	Please see the User Journey Configuration section below for details of this table.
Resource Monitor Configuration	Please see the Resource Monitoring Configuration section below for details of the table and fields relating to resource monitoring.
Run Before Test	
Run After Test	

User Journey Configuration

This table within the Test Run Configuration screen allows you to define which User Journeys will be run from which locations, and how the number of simulated users for a User Journey, at a location, will be ramped.

Note: The exact numbers of Injectors you can use, the number of users you can simulate, etc. may be restricted by your StressTester™ license.

The fields in the table are described below.

Field	Description
Injector	The Injector (location) to use to simulate the User Journey.
User Journey	The User Journey to be executed on the Injector. Note: The relationship between Injector and User Journey is many-to-many: an Injector can run many different User Journeys, and a User Journey can run in many different locations.
No. of Cycles	The number of iterations of the User Journey each simulated user should complete during the performance test. A value of '0' indicates that the simulated users should continue to run until you stop the performance test (by pressing the stop button in the Analyze Results workspace).

Field	Description
Initial Users	The initial number of users that should be simulated for the selected User Journey running on the selected Injector.
Ramp Users	The number of simulated users that should be added to those already running every time the number of users is ramped.
Ramp Interval	The time period in seconds between subsequent additions of Ramp Users to those users already running.
Maximum Users	The maximum number of simulated users that should run the selected User Journey on the selected Injector. When this number of simulated users is reached, no further ramping occurs.

Resource Monitor Configuration

This table within the Test Run Configuration screen allows you to select which Resource Monitors will measure underlying system metrics during the performance test.

The field in this table is described below.

Field	Description
Monitoring Configuration	The Monitoring Configuration to execute when running a performance test based upon this Run Configuration. Multiple Monitoring Configurations can be selected – allowing many aspects of the system under test to be simultaneously monitored.

Two additional fields are provided to control the starting and stopping of the monitoring.

Field	Description
Run Before Test	A value entered into this field indicates that the resource monitoring should start in advance of the execution of the User Journeys. The objective is to ensure that the resource monitors will capture baseline information before the load is applied. Enter the time required (in seconds).
Run After Test	A value entered into this field indicates that the resource monitoring should continue once the execution of the User Journeys has stopped. The objective is to ensure that the resource monitors will continue to capture information after the load has been removed. Enter the time required (in seconds)

When you have updated the Test Run Configuration properties, click the “Save” toolbar button.

Starting Run Components

You are now ready to start the processes that are defined within the Test Run Configuration.

These processes are either Injectors or Monitor Agents.

Starting Injectors

To start an Injector, navigate to the directory in which StressTester™ is installed and then run:

```
injector.cmd or injector.sh
```

(depending on your operating system)

Each of these takes a single parameter – the name of the Injector as defined within StressTester™.

For example, to start an Injector that has been named “Boston” on a Microsoft Windows machine, you would run the command:

```
injector.cmd "Boston"
```

Note: Injectors, once started, are long running processes that will only stop when you kill the process. This means that once an Injector has been started, it does not need to be started again until you either kill the Injector or restart the machine.

Starting Monitor Agents

To start a Monitor Agent (the process that executes Resource Monitors), navigate to the directory in which you installed the Monitor Agent and then run:

```
monitor_agent.bat or monitor_agent.sh
```

(depending on your operating system).

The command does not require any parameters but the Monitor Agent properties should have been correctly configured in StressTester™ prior to it being started.

If unsure as to whether your Monitor Agent is configured within StressTester™, please refer to the [Monitoring Underlying System Resources](#) user guide for details on how to install and configure Monitor Agents.

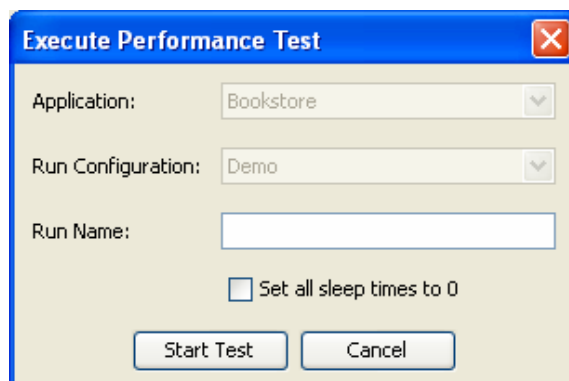
Note: Like Injectors, Monitor Agents are long running processes that will only stop when you kill the process or restart the machine.

Executing a Performance Test

Once you have defined a Test Run Configuration and started the components it requires, you can execute a performance test.

There are two ways to start a performance test: either right-click on a Test Run Configuration navigation tree node and choose the “Execute” option, or click the “Execute Test Run” button on the StressTester™ user interface toolbar.

Whichever you choose, you will be presented with the following pop-up screen.



The fields on the screen are as follows:

Field	Description
Application	The name of the application to be tested. This field will be read-only if the execution was invoked by right-clicking the Test Run Configuration navigation tree node.
Run Configuration	The Test Run Configuration that defines the performance test to be executed. This field will be read-only if the execution was invoked by right-clicking the Test Run Configuration navigation tree node.
Run Name	A unique name for the performance test.

Field	Description
Set all sleep times to 0	<p>If you want to execute a performance test in which StressTester™ will ignore configured sleep (inter-step wait or “think”) times, check this box.</p> <p>However, ensure that this box is unchecked when executing real performance tests – to correctly simulate the timings that would occur in the real world.</p>

When you have completed the screen fields, click the “Start Test” button to start the performance test.

At this point, the StressTester™ user interface will perform the following:

- ✓ Communicate with all the run components involved in the performance test to ensure they are running and ready to execute the test.

If a component is not running or ready, you will see an appropriate error message and the performance test will not be executed.

- ✓ Notify all the components of the performance test that is to be executed, in order that each can understand what it needs to do in the performance test
- ✓ Switch the user interface display to the Analyze Results workspace with the current test selected in that workspace’s navigation tree

To understand how to use the analysis screen, and how to further analyse the performance test results, please refer to the [Analyzing Performance Test Results](#) user guide.