



Integral Information Solutions GmbH
Schickardstr 32 • D-71034 • Böblingen • Germany

FusionReactor 3.0.1

Help Documentation

Doc. Rev. 1, 17 November 2008

INTER **GRAL**

Trademarks and Warranties

FusionReactor, the FusionReactor logotype, and the Intergral logotype are trademarks of Intergral Information Solutions GmbH and may not be used without permission. Other trademarks are the property of their respective owners.

The FusionReactor software product is commercial software and may not be redistributed except with the express written agreement of Intergral Information Solutions GmbH. The software may only be used in accordance with the appropriate FusionReactor license agreement.

To the fullest extent applicable by law, Intergral Information Solutions hereby disclaims all warranties, including but not limited to the warranty of merchantability and the warranty of fitness for a particular purpose.

Feedback

We welcome feedback on all our products and publications. Please e-mail them to support@fusion-reactor.com and we will address them as quickly as possible.

Published in Germany

Copyright © 2005-2008 Intergral Information Solutions GmbH

All Rights Reserved

Table Of Contents

Welcome to FusionReactor Help	1
Getting Started.....	3
Introduction	3
Requests.....	5
Running Requests	5
Request History.....	6
Request Settings.....	7
Request Log.....	8
Request Capture	10
Enterprise	13
Enterprise Dashboard (Enterprise Version Only)	13
Manage Servers (Enterprise Version Only)	16
Manage Groups (Enterprise Version Only).....	17
Enterprise Settings (Enterprise Version Only).....	17
Using Monitored Server Scripts (Enterprise Version Only)	19
Local Monitoring: Monitoring Instances on the same computer as the Enterprise Monitor ...	23
Distributed Monitoring	24
Summary	24
Metrics	27
System Metrics.....	27
Longest Requests	27
Slow Requests	28
Longest JDBC	30
Metrics Settings.....	30
Resources.....	33
Request Activity Graph	33
Request Time Graph.....	33
JDBC Activity Graph	34
JDBC Time Graph.....	35
Memory Graph	35
CPU Graph.....	36
List All Threads	36
Resource Settings.....	37
Resource Log.....	37
Crash Protection.....	41
Crash Protection Restrictions	41
Crash Protection Settings	42
Crash Protection Log	45
JDBC	49
JDBC Settings.....	49
JDBC Stack Trace Filter	50
JDBC Log.....	50
Instances	53
Instance Manager (Enterprise Version Only).....	53
Compression.....	55
Compression Settings.....	55
MIME Type Restrictions	55
Exclude URLs	56
Content Filters	59
Search and Replace	59
Content Filter Restrictions.....	59
Filter Settings	60
FusionReactor	63

Enabling/Disable	63
FusionReactor Settings.....	63
FusionReactor Restrictions.....	64
Change Password.....	65
FusionReactor Log.....	65
About.....	66
Help.....	66
Logout.....	66
Request Detail	67
Stack Tracing	67
Request Detail Page	67
Analyzing Data.....	69
Importing and Graphing Data in Excel.....	69
Users	71
Types Of User Account.....	71
Managing The User Accounts.....	72
Regular Expressions	75
Regular Expressions and Restrictions	75
Regular Expressions and Content Filters	75
Installation.....	77
Installation Guide	77
License File - Uploading and Activation.....	77
default-web.xml File Modifications	78

Welcome to FusionReactor Help

Please pick a topic

<p><i>From the Table of Contents:</i></p> <p>Requests Running Requests Request History Request Settings <i>(Manager)</i> Request Log Request Capture</p> <p>Enterprise <i>(Enterprise)</i> Enterprise Dashboard Manage Servers <i>(Manager)</i> Manage Groups <i>(Manager)</i> Enterprise Settings <i>(Manager)</i> Using Monitored Server Scripts</p> <p>Metrics Request Metrics Longest Requests Slow Requests Longest JDBC Metrics Settings <i>(Manager)</i></p> <p>Resources Request Activity Graph Request Time Graph JDBC Activity Graph JDBC Time Graph Memory Graph CPU Graph List All Threads Resource Settings <i>(Manager)</i> Resource Log</p> <p>Crash Protection CP Restrictions <i>(Manager)</i> CP Settings <i>(Manager)</i> CP Log</p> <p>JDBC JDBC Settings <i>(Manager)</i> JDBC Stack Trace Filter <i>(Manager)</i> JDBC Log</p> <p>Instances <i>(Enterprise)</i> Instance Manager <i>(Administrator)</i></p> <p>Compression Compression Settings <i>(Manager)</i></p>	<p><i>Other Topics</i></p> <p>Getting Started Introduction</p> <p>Request Detail Stack Tracing Request Detail Page</p> <p>Analyzing Data Importing and Graphing Data in Excel</p> <p>Users Types Of User Account Managing The User Accounts</p> <p>Regular Expressions (RegEx) Using RegEx in Restrictions Using RegEx with Content Filters</p> <p>Installation Installation Guide License File - Uploading and Activation default-web.xml file modifications</p> <p>General <i>(from the www.fusion-reactor.com website)</i> What can you do with FusionReactor? Frequently Asked Questions Support Questions Contact</p> <p>Advanced Topics <i>(from the www.fusion-reactor.com website)</i> There are several features within FusionReactor which can become very in-depth and pretty complicated. For this reason we have a set of PDF documents which cover these topics in much greater detail: FusionReactor JDBC Driver Wrapper: User Guide (pdf) JDBC Tutorial (pdf) FRAPI Examples in ColdFusion (pdf) FusionReactor API (FRAPI) (JavaDoc) FusionReactor Enterprise Scripting (pdf) FusionReactor Installation Guide (pdf)</p> <p>Note:</p>
--	---

<p>MIME Type Restrictions <i>(Manager)</i> Exclude URLs <i>(Manager)</i></p> <p>Content Filters Search and Replace Filter Restrictions Filter Settings</p> <p>FusionReactor Enable/Disable <i>(Manager)</i> Settings <i>(Manager)</i> Restrictions <i>(Manager)</i> Change Password <i>(Manager)</i> Log <i>(Manager)</i> About Help Logout</p>	<p><i>(Manager)</i> - These pages are only available to Manager or Administrator level users. <i>(Administrator)</i> - These pages are only available to Administrator level users. <i>(Enterprise)</i> - These pages are only available on Enterprise versions of FusionReactor.</p>
---	---

Getting Started

Introduction

FusionReactor is a professional server monitor, designed for production environments and supporting a range of J2EE servers and Adobe servers (ColdFusion 6, 7 and 8, LiveCycle and Flex Data Services). As a pro-active monitor, it increases server stability and performance by continuously examining applications, servers and databases to identify, highlight and address issues before they impact critical processes.

Once you have installed FusionReactor you will want to access the FusionReactor Administrator.

There are several ways to do this:

The simplest method is to use the desktop or start menu icons which were created during installation.

Alternatively you can call the FusionReactor Administrator directly from a web browser.

There are two ways to do this, either using the FusionReactor built in web server with a URL like this:

- `http://127.0.0.1:8088/fusionreactor/`

or you can use the external webserver with a URL like this:

- `http://127.0.0.1/fusionreactor/fhtml.cfm`

The first page you will see when you access the FusionReactor will be a login page. Choose the User type you wish to login as and enter the password which you entered during installation and then click on the Login button. Note that if you did not enter passwords for the Manager or Observer users, then these accounts are disabled and you will be unable to login as these types of users.

Requests

Running Requests

Running Requests give you the ability to view and administer running requests on your server. From the FusionReactor Administrator, click "Running Requests" from the table of contents (Requests section) and you will see the Running Requests screen.

At the top of the screen you will see the refresh bar. You can use this to manually or automatically update the running requests listed below. For each request shown you will see the time at which it began, the IP address of the machine which created it, the page which was actually requested, the amount of time it has been running and the state of the memory when this request entered the system. Above the running requests you will also see the state of the memory right now.

The request table has the following columns:

Started	The time at which the request was received.
IP	The IP address of the caller.
ID	The request ID and the name of the thread which is servicing the request.
URL/Parameters	The URL path and parameters passed to the request. You can turn on or off the hostname in the URL by modifying the URL Format value on the Request Settings page.
Time (ms)	The number of milliseconds that the request has been running so far.
Memory (KB)	A snapshot of the state of system memory at the time the request was started. The following values are shown:Used memory as a percentage of the total available to the JVM, the actual used memory value in KB, and the margin of free memory available before the JVM must extend its own heap (if it has not already reached its limit).

The Running Requests page lists all currently running pages as well as showing the current memory values.

Max	The amount of memory available to the JVM (this is a hard upper limit).
Alloc	The amount of memory currently allocated by the JVM for its heap – this may grow until the 'Max' value is reached
Free	The amount of free memory from the currently-allocated heap
Used	The amount of memory which is actually in use. This is expressed as a percentage as well as in KB.

Queued Requests (due to Crash Protection)

If there are requests queued due to Crash Protection then they will be listed below the currently running requests. The queued requests are listed in the same format as the running requests (see above) but there is an extra column giving you the reason that the request was queued.

Sorting Requests

You can change the order in which requests are displayed on this page by clicking on any of the column titles. Clicking on a title for a second time will reverse the order on that column.

Managing Requests

Each request listed on this page has four icons next to it. They are:

Stack Trace	Clicking this button will take you to the Stack Trace page, allowing you to see the current Java stack trace of the request.
Request Detail	Clicking this button will take you to the Request Details page, allowing you to access all information associated with a request; e.g cookies, JDBC, headers and User Trace Markers.
Kill Request	Clicking this button allows you to manually kill any tracked request. You will be asked if you are sure. Click on OK to kill the request. WARNING: Although FusionReactor first attempts to signal the request to let it exit safely, if this times out or you skip the timeout, then the request will be killed in a very abrupt manner. Please be aware that in some very limited cases, this can cause the JVM to become unstable.
Add to CP Restrictions	If you click this button then you will be taken to the Crash Protection Restrictions page which will be filled out for this request.

As well as the icons, you can also Stack Trace or Kill All currently running requests by using the two buttons ("Stack Trace ALL" and "Kill ALL") above the request list.

WARNING: The "Kill ALL" button does not give the requests time to complete in their own time. Using this feature can cause the JVM to become unstable.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.

Request History

The Request History page shows a list of completed requests. Please be aware that FusionReactor only keeps a fixed amount of requests in memory. If the history size gets too long or you restart you application server then old requests will be removed. To change the size of the request history, change the History Size value on the Request Settings page. Note that you can still examine these older requests by looking into the Request Log.

The request history table has the following columns:

Finished	The time at which the request was completed.
Ret.Code/IP	The final HTTP return code of the request and the IP address of the caller. If the request was killed or if it triggered Crash Protection then you will see a note here letting you know what happened.
ID	The request ID and the name of the thread which is servicing the request.
URL/Parameters	The URL path and parameters passed to the request. You can turn on or off the hostname in the URL by modifying the URL Format value on the Request Settings page.
Time (ms)	The number of milliseconds that the request took to complete. If available then you will also see the amount of actual CPU time that this request

	used.
Memory (KB)	A snapshot of the state of system memory at the time the request was started. The following values are shown: Used memory as a percentage of the total available to the JVM, the actual used memory value in KB, and the margin of free memory available before the JVM must extend its own heap (if it has not already reached its limit).

Sorting Requests

You can change the order in which requests are displayed on this page by clicking on any of the column titles. Clicking on a title for a second time will reverse the order on that column.

Managing Requests

Each request listed on this page has two icons next to it. They are:

Request Detail	If you click this button then you will be taken to the Request Details page which gives you access to all information associated with a request. eg. Cookies, JDBC, Headers and User Trace Markers.
Add to CP Restrictions	If you click this button then you will be taken to the Crash Protection Restrictions page which will be filled out for this request.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.

Request Settings

The Request Settings page lets you modify various elements relating to the capture, storage and display of Request information. The page has the following items:

Request Formatting:

URL Format	This option changes how URLs are displayed throughout FusionReactor. "Full" includes the hostname, while "Short" just shows the path. Select "Full" if you are dealing with servers which host multiple websites.
------------	---

Request Logging:

Request Logging	If Request Logging is enabled then requests will be written to the Request Log as they arrive and complete.
View Size (KB)	This value indicates how much of the log file should be displayed when you view the Request Log from inside FusionReactor.
File Size (KB)	Specifies the size of each file in the Request Log rotating file set.
File Count	Specifies the number of files in the Request Log rotating file set.

Request Content Capture:

Request Capture	Activating this option instructs FusionReactor to store the complete Request and Response objects to disk.
-----------------	--

	WARNING: Turning on Request Content Capture will create multiple XML files for every request run on a server. This may use very large quantities of disk space quickly. Please be careful when enabling this feature and do not leave it enabled if you are not monitoring it.
Capture Path	Specifies the location into which captured requests and responses are stored.

History:

History Size	This value defines how many requests FusionReactor stores in memory. Once the Request History reaches this size, old requests will be removed when new ones arrive.
--------------	---

Stream Metrics:

Stream Metrics	Specifies whether FusionReactor should track statistics about how long it took to get a response back to the client. (Time to first byte, time to completion, bytes per second, etc.).
----------------	--

Flash Remoting Support:

AMF Decoding	If FusionReactor receives an Flash Remoting (AMF) request then you can use this option to have FusionReactor decode it. Decoded AMF can be viewed on the Request Details page.
--------------	--

Request Log

The Request Log page is a simple display page which shows the last entries of the Request Log file. The Request Logs (a set of rotating log files accessible over the Request Settings page) are simple space-separated files so they can be easily imported into database or spreadsheet applications. By default these files are located in the `/FusionReactor/instance/[your_instance_name]/log` directory. Note that the amount of data stored in the log files is typically far greater than the number of requests shown on either the Request History or Request Log screens. After a server restart the Request Log screen will show the last entries on the Request Log, even if these requests were run before the server restart. You can change the amount of data shown on the Request Log screen by changing the View Size on the Request Settings page.

The log file is formatted as follows:

Field	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The date on which this log entry was written.
Date/Time (ms)	3 [C]	This is the millisecond time at which this request was started or finished (depending upon the Request Status.)
Version	4 [D]	Version number of this log entry.
Server Startup Time (ms)	5 [E]	Startup time of the application server in milliseconds (ms) since epoch.
Request ID	6 [F]	The FusionReactor Request ID.
Request	7 [G]	The current state of the request. This will tell you if a request has

Status		started, finished, was killed, queued, etc...
CP Reason	8 [H]	The reason given by Crash Protection for the current action.
Thread ID	9 [I]	The name of the thread responsible for responding to this request.
Client IP Address	10 [J]	The IP address of the machine making the request.
Request Method	11 [K]	This will usually be "GET" or "POST"
Request URL	12 [L]	This is the requested URL. The format of this column will change depending on URL Format key which can be found on the Request Settings page.
Execution Time (ms)	13 [M]	The amount of milliseconds it took to complete the request. (For incomplete requests, this column will be 0.)
Used Memory (percentage)	14 [N]	The amount of memory (expressed as a percentage) which was used when this request started or finished.
Max Memory (KB)	15 [O]	The total amount of physical memory available to this instance.
Used Memory (KB)	16 [P]	The amount of memory which was used when this request started or finished.
Total Memory (KB)	17 [Q]	The amount of memory which was allocated by the instance when this request started or finished.
Free Memory (KB)	18 [R]	The amount of free memory (within the allocated block) when this request started or finished.
Query String	19 [S]	If the URL has any parameters then they will appear here.
Return Status Code	20 [T]	This is a HTTP return code such as 200 (OK,) 404 (Not found,) or 500 (Internal Server Error.) For "Started:" rows this column will be 200.
CPU Time (ms)	21 [U]	This is the amount of actual CPU time which this request required. (For incomplete requests, this column will be 0.)
AMF Request	22 [V]	If AMF decoding is enabled (Request Settings page) and there is AMF to decode, then the method names will appear here.
JSESSIONID	23 [W]	The J2EE Session Id for this request.
CFID	24 [X]	The CF Id for this request.
CFTOKEN	25 [Y]	The CFTOKEN for this request.
JDBC Query Count	26 [Z]	The number of JDBC queries run by this request at the time of this log.
JDBC Total Time (ms)	27 [AA]	The total amount of time spent running JDBC queries at the time of this log.
JDBC Total Execution Time (ms)	28 [AB]	The total amount of time spent by the database running JDBC queries at the time of this log.
JDBC Total Rows	29 [AC]	The total number of rows returned by JDBC queries at the time of this log.
Bytes Sent	30 [AD]	The amount of data which was sent back to the client.
Time to First Byte (ms)	31 [AE]	The number of milliseconds it took to deliver the first bit of data.

Time to Last Byte (ms)	32 [AF]	The number of milliseconds it took to deliver the complete content.
Time to Stream Open (ms)	33 [AG]	The number of milliseconds before the data stream was opened.
Time to Stream Close (ms)	34 [AH]	The number of milliseconds until the completed data stream was closed.

Deleting the Request Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the Request Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Also See

Importing and Graphing Data in Excel

Request Capture

Request Capture is the ability to store incoming requests and their associated responses to the file system.

When request capture is set to requests and responses you will end up with three files for each request. A request file, a body file and a response file. The header XML file holds all information sent to server by the client. The response XML file holds everything sent back (besides body) and the body file holds the actual returned content.

WARNING: Turning on Request Content Capture will create multiple XML files for every request run on a server. This may use very large quantities of disk space quickly. Please be careful when enabling this feature and do not leave it enabled if you are not monitoring it.

How do I set up Request Capture?

From the FusionReactor Administrator, click "Request Settings" from the table of contents (Requests section) and you will see the Request Settings screen. Scroll down until you reach a section named "Request Capture". From here you can control the request capture feature.]

Request Capture can be in three states. Disabled, Capture Requests only, or capture both requests and responses. The Capture Path is the directory into which the captured requests will be stored.

Why would I want to Capture Requests?

Request capture can be useful for a variety of reasons from technical to legal, these include:

Analysis

The FusionReactor stored requests hold much more information than you would normally expect to find in web logs and as such they have a much higher value for analyzing your application traffic.

Development

You could use stored requests in order to be able to reproduce specific load conditions or test cases.

Support

Anyone who has had to have a user describe the contents of their screen will instantly realize how much easier life would be if you had an exact copy of their response in easy access.

Enterprise

Enterprise Dashboard (Enterprise Version Only)

The Enterprise Dashboard lets you track the state of installations on one or more servers. When you install an Enterprise version of FusionReactor on a server it can be queried for its current state. The Enterprise Dashboard lets you instantly monitor the health of ALL of your server instances/machines from a single window. You immediately see the state of memory, CPU, DB load and average request time, in an easy to read graphical dashboard display.

By default you will see your install server in the Enterprise Dashboard but, using the Manage Servers page, you can add any other server which is running FusionReactor.

Once you have your machines set up, go to the Dashboard. The default view when bringing up the Enterprise Dashboard is to show all groups (which can be set up on the Manage Groups page) along with any servers which are currently ungrouped.

Group Icons

You can group machines together to help organize your dashboard. Each machine and group is represented by an icon which will change colour to indicate its current health. As well as this, clicking on a server or group will give you its most recent statistics. The dashboard also provides links to each monitored instance so it is a simple job to investigate a server further.

Clicking on a group icon will display all servers within that group. You will also see the Group Information Pane which will display statistics about the group and the servers within. By default the group icon will be displayed in blue, but if any grouped server is having problems then the colour of the icon will change and a small warning panel will be displayed to let you know what the problem is.

How Do I Add Machines?

There are three ways to add servers to your Enterprise Dashboard: Firstly, you can import servers from your Instance Manager. Secondly, you can get FusionReactor to scan its machine for servers, or you can simply insert the URL of your server into the Manage Servers page.

Server Icons

The server icon shows four bars. These bars each have a letter above them to indicate what is being displayed. These bars are:

M	Memory	This represents the amount of memory currently being used by this server.
C	CPU	This represents the amount of CPU time used by this instance since the last CPU value was measured. It is expressed as a percentage of the total CPU time (across all CPU cores) that was available during that period.
R	Request Time	This represents the average request time on this server over the last 60 seconds.
D	DB Time	This represents the average time spent on DB operations in the last 60 seconds.

Clicking on a server icon will display the Server Information Pane which will display statistics about the server. By default the server icon will be displayed in blue, but if there are any problems with the server then the colour of the icon will change and a small warning panel will be displayed to let you know what the problem is. Additionally, there are two small buttons on the Server icon:

- At the top right corner of the icon is a small button which can be used to flip a server between Online and Offline status. An offline server will be displayed in gray and will not cause warnings to appear within the Enterprise Dashboard.
- At the bottom right corner of the icon is a small arrow. If you click on this arrow then you will be taken directly to the instance of FusionReactor running on this server.

Warning Lights

When FusionReactor detects a problem on a monitored server it will light up one or more "Warning Lights" on the server cube. You can mouse over a warning light to see a tool tip explaining the problem in more detail. The following table describes the warning lights and their meanings:

CONN	Connection Lost	The connection (heartbeat) has been lost to the monitored server. The monitored server may be down.
SLOW	Slow Requests	The number of requests running longer than the slow page threshold (within the last 60 seconds) has exceeded the Warn or Critical values configured on the Enterprise Settings page.
LONG	Long Running Requests	The average request time (over the past 60 seconds) has exceeded the Warn or Critical value configured on the Enterprise Settings page.
MEM	Low Memory	The used memory percentage exceeded the Warn or Critical value configured on the Enterprise Settings page.
QTY	CP Request Protection	The number of triggered crash protections (for the request quantity limit) exceeded the Warn or Critical value configured on the Enterprise Settings page.
TIME	CP Timeout Protection	The number of triggered crash protections (for the request timeout limit) exceeded the Warn or Critical value configured on the Enterprise Settings page.
DB	JDBC Slow Requests	The average JDBC request time (over the past 60 seconds) has exceeded the Warn or Critical value configured on the Enterprise Settings page.
ACT	Please Activate License	You need to activate the license installed on the monitored server.

Group Information Pane

The group information pane shows you how many servers are within this group as well as showing statistics for all grouped servers. These statistics are:

Group Size	This shows the number of servers within this group.
Request Count	This shows you how many requests are running on all grouped servers along with the total number of requests which have been run.
Recent Slow Pages	This shows how many requests within the last 60 seconds have run longer than the Slow Page Threshold value.

Average Request Time	This is the average request time for all grouped servers for the last 60 seconds, and for the complete history size.
Average DB Time	This is the average time spent on database operations for all grouped servers.
Average Memory	This is the average used and total allocated memory for all grouped servers.

Additionally, you will see two graphs (Average Used Memory and Request Count.) These graphs show the same information displayed in the statistics section but they are presented visually for ease of use.

Server Information Pane

The server information pane shows you various statistics about your selected server. Statistics which are listed with blue titles are links. Click on one of these titles to be taken to a more detailed page within the instance of FusionReactor running on that server. If a server is currently having problems then the appropriate statistic will be highlighted to indicate where the problem is. The displayed statistics are:

FusionReactor Instance	This is the full ID of the instance of FusionReactor running on this server.
License	This shows the type of FusionReactor license running on this server.
Activation Status	This shows you if the license on this machine has been activated or not.
Crash Protection (Page Aborts)	
Request Quantity	This shows the number of pages which triggered Crash Protection due to request quantity in the last 60 seconds and in total.
Runtime Timeout	This shows the number of pages which triggered Crash Protection due to runtime timeout in the last 60 seconds and in total.
Low Memory	This shows the number of pages which triggered Crash Protection due to low memory in the last 60 seconds and in total.
Overview	
Server Uptime	This shows how long this server has been up.
Current Request Count	This is the number of requests currently running on this server.
Total Request Count	This is the total number of requests which have run on this server.
Average Request Time	This shows the average request time on this server over the last 60 seconds, and for the complete duration of it's request history.
Used Memory	This is the amount of memory currently used by this server (not the amount currently allocated.)
Total Memory	This is the total amount of memory available of this server.
Average Instance CPU	This shows the amount of CPU time used by this instance since the last CPU value was measured. It is expressed as a percentage of the total CPU time (across all CPU cores) that was available during that period.
Average DB Time	This shows the amount of time spent of database operations on this server.
Queue Size	This is the number of requests currently queued due to Crash Protection
Recent	
Recent Slow Pages	This shows how many requests within the last 60 seconds have run longer than the Slow Page Threshold value.

Additionally, you will see five graphs (Memory, Instance CPU, Average Request Time, Request Load, Request Count, JDBC Time and JDBC Load.) These graphs show the same information displayed in the statistics section but they are presented visually for ease of use. Both the Memory and Instance CPU graphs have arrow buttons by them. Clicking on this button will take you directly to the Memory or CPU graph page within the instance of FusionReactor which is running on this server.

Configuring the Enterprise Dashboard

Many elements within the Enterprise Dashboard can be modified. In order to make changes, please refer to the Enterprise Settings page. From there you can:

- Specify whether you should receive an email when servers start up and shut down.
- Specify what the maximum values should be on graphs and bars.
- Specify the warning and critical state threshold values for various statistics.

Manage Servers (Enterprise Version Only)

From this screen you can add new servers to the Enterprise Console or manage existing servers. It is also possible to directly add local instances to the Enterprise Dashboard from the Instance Manager. The form to add a new server has the following fields:

Name	This is the name which will appear in the Enterprise Dashboard under your server icon.
URL	This is the location of FusionReactor on the target server. An example URL will look something like this: <code>http://127.0.0.1:8088/fusionreactor</code>
Password	This is the password of FusionReactor on the target server.
Script	If you specify the path to a script on your server then it will be called when the server goes down or comes up. See the script section below for more information.
Status	You can choose if this new server is online or offline. An online server will collect statistics and display them in the Enterprise Dashboard and if there are any problems then you will see warnings. If a machine is offline then no information will be gathered and the icon will be displayed in gray. If you are adding a server which is currently turned off then you should set its status to offline. You can change its status later by returning to the Manage Servers page or by clicking the status button (top right hand corner of a server icon,) directly within the Enterprise Dashboard.
Groups	If you have any groups set up then they will be listed here. Simply select which groups you want this new server to appear within. You can add new groups at any time from the Manage Groups page.

Once you submit the form the page will be refreshed and your new server will be listed at the bottom of the page along with its current connect status.

Managing Existing Servers

Below the Add Server form you will see a list of all servers currently available in the Enterprise Dashboard. Each server has three icons next to it:

Modify Server	Clicking on this icon will take you to the Modify Server page. From this page you can change any of the information which is listed above.
---------------	--

Check Server	Clicking on this icon will check the connect status for a server. The screen will be refreshed and you will see a message telling you if there were any problems trying to access FusionReactor on this server.
Delete Server	Clicking on this icon will prompt you if you are sure you want to delete the server. If you OK the dialog then the server will be removed from the Existing Servers list and from the Enterprise Dashboard.

Clicking on the URL for any server in the Existing Servers list will take you directly to the instance of FusionReactor which is running on that server.

Shutdown / Startup Script

The script file which you can associate with a server will get called whenever that server goes down or comes back up again. It takes 5 parameters, they are:

Status	This parameter will have the value "DOWN" or "UP"
InstanceName	This is the instance name of the server in question (as defined on the Manage Servers page.)
InstanceIP	This is the IP address of the server in question.
PID	The Process ID of the Server. You could use this in order to terminate a server which had stopped responding. (If the PID is unknown then a value of -1 will be passed.)
LastSeen (ms)	This is the last time (expressed in milliseconds) that the server was last successfully contacted.

Manage Groups (Enterprise Version Only)

From this screen you can create or modify groups in order to better organize your servers within the Enterprise Dashboard. The Add Group form is very simple:

Name	This is the name of your new group which will appear in the Enterprise Dashboard and on the Manager Servers screen.
Servers	If you have any servers currently set up then they will be listed here. Simply select which servers should be added to your new group.

Managing Existing Groups

Below the Add Group form you will see a list of all existing groups. Each group has two icons next to it:

Modify Group	Click this icon to change a group name or to change the servers which are currently a part of this group.
Delete Group	Clicking on this icon will prompt you if you are sure. OK the dialog box to delete the group and refresh the page.

Enterprise Settings (Enterprise Version Only)

The Enterprise Settings page lets you modify various elements relating to the Enterprise Dashboard. The page has the following items:

Server Shutdown/Start Up Alerts:

Send Alert Email	With this option you can decide if you want notifications when a machine starts up and shuts down, or just when it shuts down if you prefer. To receive notifications you must first configure your email server on the FusionReactor Settings page.
------------------	--

Heartbeat Settings for Monitored Instances:

Heartbeat Metrics Fetch Interval (seconds)	This interval tells FusionReactor how long to wait between each fetch of a monitored instance's metric data (its heartbeat). The heartbeat contains the data that is shown on the Enterprise Dashboard. Note: increasing this value will lower the refresh rate of the metric data shown on the Enterprise Dashboard. Decreasing this value may cause heartbeat failures. It is NOT recommended to set this value below 5 seconds.
Heartbeat Timeout (seconds)	This timeout governs how long FusionReactor will wait for a heartbeat from a monitored server instance. If no heartbeat is received within this time, FusionReactor will count it as a heartbeat failure.
Heartbeat Failure Threshold	This setting tells FusionReactor how many times the heartbeat may fail on a monitored server instance before considering it to be DOWN. FusionReactor will then report the server instance as [CONN] on the Enterprise Dashboard and will run the self-healing script for the instance (if configured). If you are receiving monitoring emails and log messages for instances known to be available, try increasing this value. A value of zero will cause the first missing heartbeat to signal the server instance is DOWN.

Request Count:

Graph Y Scale	This number will be used as the maximum value on the request count graphs within the Enterprise Dashboard. NOTE: This value is only used for display purposes. Entering a number here will NOT limit the number of requests a server will accept!
---------------	---

Memory:

Memory Status (percent)	These two numbers (Warn and Critical) are used as threshold values to work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the used memory percentage exceeds the Warn value then it will appear in orange. If it exceeds the Critical value then it will appear in red.
-------------------------	---

Average Request Time:

Average Request Time (ms)	These two numbers (Warn and Critical) are used as threshold values to work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the average request time (over the past 60 seconds) exceeds the Warn value then it will appear in orange. If it exceeds the Critical value then it will appear in red.
Graph Y Scale (ms)	This number will be used as the maximum value on the average request time graphs and bars within the Enterprise Dashboard. NOTE: This value is only used for display purposes. Entering a number here will NOT limit the amount of time requests can run on a server!

JDBC:

JDBC Status	These two numbers (Warn and Critical) are used as threshold values to
-------------	---

(ms)	work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the average DB time exceeds the Warn value then it will appear in orange. If it exceeds the Critical value then it will appear in red.
Graph Y Scale (ms)	This number will be used as the maximum value on the average DB time graphs and bars within the Enterprise Dashboard. NOTE: This value is only used for display purposes. Entering a number here will NOT limit the amount of time database requests can run on a server!

Crash Protection:

CP Status	These two numbers (Warn and Critical) are used as threshold values to work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the number of triggered crash protections (for a given type: request quantity, request timeout or low memory) exceeds the Warn value then it will appear in orange. If it exceeds the Critical value then it will appear in red.
-----------	--

Slow Requests:

Slow Page Status	These two numbers (Warn and Critical) are used as threshold values to work out if a server should be put into a warning or critical mode when it is displayed in the Enterprise Dashboard. If the number of requests running longer than the slow page threshold within the last 60 seconds exceeds the Warn value then it will appear in orange. If it exceeds the Critical value then it will appear in red. You can change the slow page threshold from the Metrics Setting page.
------------------	--

Using Monitored Server Scripts (Enterprise Version Only)

FusionReactor Enterprise Scripting enhances FusionReactor's Enterprise Monitoring to trigger configurable self-healing scripts when a server's responsiveness status changes.

This feature can be used to perform any task which can be launched from a platform script.

Examples might include:

- Interacting with SNMP or enterprise monitoring/reporting systems
- Performing automated restarts of affected instances
- Sending custom email messages
- Writing custom log messages

Since FusionReactor Enterprise Scripting is able to launch any platform executable (shell script, executable binary etc.), it may be tailored to virtually any environment.

NOTE: To use this functionality **AT LEAST** two different instances of FusionReactor Enterprise are **REQUIRED**. It is important to understand this limitation of scripting. A script should not be configured to run against the same server instance that FusionReactor is running in. If the server instance fails then the FusionReactor running in this server instance may not be able to run the script, and the may not be able to restart itself for example.

What Are Enterprise Scripts?

FusionReactor (Enterprise Edition) can trigger a script when a server becomes unresponsive, and when it becomes responsive again. This mechanism might allow you to pro-actively attempt to restart a failed server or instance, integrate FusionReactor into an existing monitoring environment or provide custom logging or reporting. Any program which is runnable on your FusionReactor monitoring system may be used as an Enterprise Script target.

On Windows systems, anything that is runnable as a binary, or can be run from the command prompt, for instance:

- Binary programs
- Batch files

On Unix and Unix-based systems, this includes everything you can start from a command line, for instance:

- Binary programs
- Shell scripts (including Bash, Ruby, Python and Perl)
- Java programs (when launched from an appropriate shell script)

How Do I Configure an Enterprise Script?

Enterprise scripts are configured by editing the Script of the server's property page, accessible in the Enterprise -> Manage Servers page, then clicking on the Modify icon of the required server. You should take care to ensure the full path and filename to the script are correct. Also note that it is recommended to only configure one script for a server. If the same server has multiple scripts configured (for example you have entered the server into the dashboard in multiple groups) or you monitor the same server from different machines, then more than one script may be launched at the same time if the server has problems. The scripts may interfere with each other especially if they are trying to restart the same instance.

When Does FusionReactor Run Enterprise Scripts?

Enterprise Scripts are run whenever the Enterprise Monitor detects that a monitored instance has changed state:

- an instance which was previously available is no longer providing Enterprise data
- an instance which was previously unavailable has begun to provide Enterprise data

Scripts are run only if an instance changes state while it is being observed. Additionally, scripts are only run if:

- The Enterprise Server Alerting system is running
 - This is configured in Enterprise Settings -> Server Shutdown/Start Up Alerts
 - At least "On Shutdown" must be selected. If you wish to run scripts when an instance becomes available, "On Shutdown and Start Up" must be selected.
 - If you do not wish to additionally receive email for these events, disable notification in FusionReactor -> Settings -> Email Server -> Notification.

How Does FusionReactor Run These Scripts?

Launch Mechanism

FusionReactor runs these scripts by spawning them using Java system commands. The scripts will be run in the context of the user under which your J2EE (ColdFusion) server runs. This user

must have at least **read + execute** access to these scripts. Any files or other executables called by the script must also be accessible by this user.

The script will be run with the current working directory (CWD) of the J2EE application server. Because of the variety of platforms available, this may be unpredictable. Any scripts you write should therefore not use the current directory notation (usually a single dot) to address files. If you plan to access files within the script, their paths should be specified completely.

Script Arguments

FusionReactor supplies several command-line arguments to the script. These arguments may be used by the script to perform logging or restart operations. Note that these arguments are supplied by FusionReactor and you do not need to enter them in the script field on the Managed Servers page. The supplied arguments are (in order):

Script Parameter	Description
UP or DOWN	reflecting the instance status
instance name	as registered in the Manage Servers screen
IP Address	as returned from a DNS lookup of the machine name part of the URL used to monitor this instance
Process ID	If available, the process ID of the J2EE application server on the remote machine. If the FusionReactor native library is not available, or FusionReactor could not read this value, this field will be -1
Last Seen Time	The time, measured in milliseconds from midnight on January 1st 1970, which the server was last successfully polled for Enterprise data. If the server has not been observed as running during this session, this field will be -1

Logging

When FusionReactor fires a script, an appropriate message is written to the Crash Protection log, located in FusionReactor/instance/<instance_name>/log/crashprotection-0.log. This log is shared with other Crash Protection messages, and not all fields are used by Enterprise scripting.

For the exceptional cases **SCRIPTREADFAILED** and **SCRIPTEXCEPTION**, FusionReactor will log the message associated with the exception to the FusionReactor log.

Operational Impacts of Scripting

There are a few points which should be borne in mind when configuring scripting.

System Restarts and Self-Monitoring

If FusionReactor is configured to monitor itself, i.e. is monitoring the same instance in which it is configured, scripting should not be used for operations.

In these circumstances, we recommend transitioning your environment to a High Availability monitoring solution. This entails installing a new J2EE server (Tomcat, for instance), and installing FusionReactor into that. This container will be used purely as a FusionReactor host, and will be used to monitor other containers. It may be necessary to create scripts which perform system reboots. Again, we recommend a careful evaluation of the impacts of this type of script before implementation. A script which restarts a system should not attempt to restart the system on which the monitoring solution runs.

Manual Restarts

If a script is configured for a given instance, it will be fired when that instance becomes unavailable. FusionReactor does not differentiate between overloaded (or failing) instances, and instances which have been deliberately stopped. Therefore, if you stop an instance manually, through Windows' Services panel for instance, FusionReactor will fire the configured script. As an operational matter, the affected instances should be "**offlined**" from FusionReactor before being shut down. This can be done within the Enterprise Dashboard, by clicking the (+/-) button on the top right of server icon or from within Manage Servers by selecting the Modify icon for the affected server, then changing its Status to Offline. FusionReactor will not monitor these systems. When maintenance is complete, the servers should be "**onlined**" again by reversing the process.

Using the Example Scripts

We have provided several restart scripts to get you started. This section will help you understand how to install and configure them.

Installation

The example scripts are provided in /FusionReactor/etc/cp/, thereafter the structure is split into scripts which will run on Unix platforms, and those which will run on Windows platforms. You are free to run these scripts from this location, but we would recommend you copy these templates before editing them. You will then always have a pristine copy available for new scripts.

Worked Example: Controlling Windows ColdFusion 8 from Windows

In order to get you started, we've provided you with a worked example. In our example scenario, we will use a FusionReactor Enterprise Edition instance on a ColdFusion 7 instance to monitor a ColdFusion 8 installation, also on Windows, also running FusionReactor Enterprise Edition. Enterprise Dashboard

The first stage in preparing the environments is to ensure that both systems are running smoothly, and the monitor is able to poll the target system for enterprise data. Add the remote system to the monitor and check that Enterprise Dashboard is retrieving information from it.

Script Preparation

For this example, we'll be using the restart-Coldfusion8-OnWindows.bat script from the FusionReactor/etc/cp/windows folder. For our example, we copy the example script to a temporary folder, from where we can work on it:

```
copy restart-Coldfusion8-OnWindows.bat c:\tmp
```

In order to customize the script, we open it in an editor. All provided example scripts are commented extensively.

There are a couple of variables we must customize in the script:

- We set the LOGFILE (line 43) to c:\tmp\script.log
- We change the USER and PWD (lines 53 and 54 respectively) to reflect the Windows user with permissions to restart ColdFusion.

Adding the Script to Manage Servers

The final step in the configuration is to add the script to the monitored server's configuration. We edit the server's Enterprise Dashboard configuration by clicking on Manage Servers, then clicking the edit icon of the monitored server. We enter the script location in the Script field.

Testing the Script

The script can be tested by simply using the Windows Service control panel to stop the monitored ColdFusion 8 service. Observing the script log file c:\tmp\script.log file shows the output of the script. The ColdFusion 8 service can be observed restarting in the control panel.

Conclusion

We've shown you how to configure Enterprise Scripting to restart a ColdFusion 8 server. The scope for what scripts can do is immense, since there are no restrictions on what they may call. It would be a simple task, for example, to integrate FusionReactor into an SNMP monitoring solution, write custom log messages or send SMS text messages.

Local Monitoring: Monitoring Instances on the same computer as the Enterprise Monitor

Running all FusionReactor instances on the same computer makes it easy to develop scripts for use with Crash Protection.

Simple Watchdog

Two instances run on the same computer. One of the instances acts as an Enterprise Monitor (Watchdog) for the other, monitored server. The watchdog server has added the monitored server to the list of managed servers in the Enterprise Dashboard of FusionReactor. A script for the monitored server is configured that will be executed by the watchdog server if the monitored server is unavailable or becomes available again. The script is used to restart the monitored server automatically after it became unavailable.

However, if the watchdog server itself becomes unavailable, the server is not longer monitored and can not be restarted automatically any more. To get around this the monitored server could also monitor the watchdog server which is described in the next section.

Cross Monitoring

One instance/server is created for the task of Enterprise Monitor to monitor the other Operational servers. Every operational server is added to the Enterprise Monitor's list of managed servers in the Enterprise Dashboard of FusionReactor. For every server a script is configured that will be executed when the server becomes unavailable (or available again). The script is used to restart

the operational server automatically if it becomes unavailable. One of the operational servers is also configured to monitor the Enterprise Monitor, and a script is configured to restart the Enterprise Monitor should it become unavailable. In this way the Enterprise Monitor watches over all of the operational servers, and one of the operational Servers watches over the Enterprise Monitor. Alternatively a second Enterprise monitor could be configured with the task of watching over the first Enterprise monitor if you do not want operational servers to perform monitoring tasks.

Distributed Monitoring

Running, monitoring and restarting FusionReactor instances in a distributed environment requires remote connections between the different machines. If a monitored server becomes unavailable this is monitored by a different machine on the network. This machine then calls a script which has to connect to the remote machine and restart the remote server/instance. Depending on the operating system the participating machines use, this can be done with (SSH) or some similar technology.

Simple Watchdog

Two servers run on different machines that have a network connection. One of the servers acts as the Enterprise Monitor (watchdog) for the other, monitored server. The watchdog server has the monitored server entered in its list of managed servers in the Enterprise Dashboard of FusionReactor. A script for the monitored server is configured that will be executed by the watchdog server if the monitored server becomes unavailable (or available again). The script is used to login to the remote computer and restart the server/instance automatically after it becomes unavailable. This approach can have the same drawbacks as mentioned in the Local Server Simple Watchdog section before.

Cross Monitoring

All servers run on different machines that can reach each other over the network. One instance/server is created for the task of Enterprise Monitor to monitor the other Operational servers. Every operational server is added to the Enterprise Monitor's list of managed servers in the Enterprise Dashboard of FusionReactor. For every server a script is configured that will be executed when the server becomes unavailable (or available again). The script is used to login to the remote computer and restart the operational server automatically if it becomes unavailable. One of the operational servers is also configured to monitor the Enterprise Monitor, and a script is configured to login to the Enterprise Monitor computer and restart the Enterprise Monitor instance should it become unavailable. In this way the Enterprise Monitor watches over all of the operational servers, and one of the operational Servers watches over the Enterprise Monitor. Alternatively a second Enterprise monitor could be configured with the task of watching over the first Enterprise monitor if you do not want operational servers to perform monitoring tasks.

Summary

FusionReactor Crash Protection Scripts are powerful functionality that can lead to an increase of server availability in local as well as distributed server environments. **You must however be aware of that a script that is executed by a Crash Protection rule is executed by the server instance that is monitoring the problem server and not by the problem server itself. If this is not the same computer, remote scripting has to be used to react to the situation in an appropriate way. Also note, that you should not use scripting in a single instance (self-monitoring) scenario, because the script may not reliably start.**

Metrics

System Metrics

If Metrics are enabled (see Metrics Settings) then the Request Metrics page will display various statistics about this server. These metrics are split into several sections:

Graphs	Here you can see 6 graphs: Request Load, Request Time, JDBC Load, JDBC Request Time, Memory Demand, and CPU. Click on any graph to maximize it and click on a maximized graph to return to the standard view. You can also flip one or all of the graphs between the minute and hour view by selecting one of the options from the little clock menu at the top right corner of each graph.
Overview	This section shows general statistics about this server.
Recent	This section shows statistics about actions which have occurred on this server within the last 60 seconds. The Recent Slow Pages link will show you all requests from the current request history which ran longer than the slow page threshold (defined on the Metrics Settings page).
Crash Protection (Page Aborts)	This section shows how many pages have been aborted due to each of the three types of Crash Protection. Clicking on the various Crash Protection types will show you all requests from the current request history which were aborted for that reason.
Return Code	This section lists all HTTP return codes which have been returned by this server. Each return code is also a link which will show you all requests from the current request history which have that return code.

NOTE: Many of the values shown on this page are counts based on all pages which have been run on this server. If you see that the number of requests which returned "200 OK" is in the thousands, the number of requests actually listed will depend upon the Request History size (which is defined in Request Settings.) It is quite possible that, if enough requests have been processed since the last "500 Internal Server Error", then clicking on that link may not show any requests at all.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.

Longest Requests

The Longest Requests page shows a list of the longest running requests on this server. To change the size of the Longest Request list, change the History Size value on the Metrics Settings page. The Longest Requests table has the following columns:

Finished	The time at which the request was completed.
Ret.Code/IP	The final HTTP return code of the request and the IP address of the caller. If the request was killed or if it triggered Crash Protection then you will see a note here letting you know what happened.

ID	The request ID and the name of the thread which is servicing the request.
URL/Parameters	The URL path and parameters passed to the request. You can turn on or off the hostname in the URL by modifying the URL Format value on the Request Settings page.
Time (ms)	The number of milliseconds that the request took to complete. If available then you will also see the amount of actual CPU time that this request required.
Memory (KB)	A snapshot of the state of system memory at the time the request was started. The following values are shown: Used memory as a percentage of the total available to the JVM, the actual used memory value in KB, and the margin of free memory available before the JVM must extend its own heap (if it has not already reached its limit).

Sorting Requests

You can change the order in which requests are displayed on this page by clicking on any of the column titles. Clicking on a title for a second time will reverse the order on that column.

Managing Requests

Each request listed on this page has two icons next to it. They are:

Request Detail	If you click this button then you will be taken to the Request Details page which gives you access to all information associated with a request. eg. Cookies, JDBC, Headers and User Trace Markers.
Add to CP Restrictions	If you click this button then you will be taken to the Crash Protection Restrictions page which will be filled out for this request.

NOTE: Over extended periods of time you may find that this list gets filled with very long running requests. In order to see some of the medium running requests again you can empty the list by clicking on the "Reset Longest" button.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.

Slow Requests

The Slow Requests page shows a list of currently running and completed requests which have taken longer to run than the current Slow Request Threshold (which is defined on the Metrics Settings page.) To change the size of the Slow Request list, change the History Size value on the Metrics Settings page. The Slow Requests table has the following columns:

Started <i>or</i> Finished	Currently running requests show the time at which the request was accepted. Finished requests show the time at which the request completed.
IP <i>or</i> Ret.Code/IP	The final HTTP return code of the request is shown for completed requests. Both current and completed requests show the IP address of the caller. If a completed request was killed or if it triggered Crash

	Protection then you will see a note here letting you know what happened.
ID	The request ID and the name of the thread which is/was servicing the request.
URL/Parameters	The URL path and parameters passed to the request. You can turn on or off the hostname in the URL by modifying the URL Format value on the Request Settings page.
Time (ms)	The number of milliseconds that the request has taken/took to complete. If available then you will also see the amount of actual CPU time that this request used.
Memory (KB)	A snapshot of the state of system memory at the time the request was started. The following values are shown:Used memory as a percentage of the total available to the JVM, the actual used memory value in KB, and the margin of free memory available before the JVM must extend its own heap (if it has not already reached its limit).

Sorting Requests

You can change the order in which requests are displayed on this page by clicking on any of the column titles. Clicking on a title for a second time will reverse the order on that column.

Managing Requests

Each currently running request listed on this page has four icons next to it. They are:

Stack Trace	Clicking this button will take you to the Stack Trace page, allowing you to see the current Java stack trace of the request.
Request Detail	Clicking this button will take you to the Request Details page, allowing you to access all information associated with a request; e.g cookies, JDBC, headers and User Trace Markers.
Kill Request	Clicking this button allows you to manually kill any tracked request. You will be asked if you are sure. Click on OK to kill the request. WARNING: Although FusionReactor first attempts to signal the request to let it exit safely, if this times out or you skip the timeout, then the request will be killed in a very abrupt manner. Please be aware that in some very limited cases, this can cause the JVM to become unstable.
Add to CP Restrictions	If you click this button then you will be taken to the Crash Protection Restrictions page which will be filled out for this request.

Each completed request listed on this page has two icons next to it. They are:

Request Detail	If you click this button then you will be taken to the Request Details page which gives you access to all information associated with a request. eg. Cookies, JDBC, Headers and User Trace Markers.
Add to CP Restrictions	If you click this button then you will be taken to the Crash Protection Restrictions page which will be filled out for this request.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.

Longest JDBC

The Longest JDBC page shows a list of the longest running JDBC requests on this server. To change the size of the Longest Request list, change the History Size value on the Metrics Settings page. The Longest JDBC table has the following columns:

SQL	This is the SQL which was actually ran followed by a JDBC stack trace (the list of CFCs and tags which were called to get to this query) and finally you will see the URL which was called from the browser. If the stack trace is longer than one line then it will be displayed collapsed. Click on the line or the arrow to the left of the line to expand the stack trace.
Data Source	If you are using a wrapped data source and have named it (Refer to the JDBC Driver Wrapper User Guide) then that name will appear here.
Start Time	This is the time at which this JDBC query began.
Total Time (ms)	This is the amount of time it took CF to complete the query.
DB Time (ms)	This is the amount of time the database actually spent processing the query.
Row Count	This is the number of rows which was returned by the query.

Managing Requests

Each JDBC request listed on this page has an icon next to it:

Request Detail	If you click this button then you will be taken to the Request Details page for the page which ran this JDBC request. This will give you access to all information associated with that request. eg. Cookies, JDBC, Headers and User Trace Markers.
----------------	---

NOTE: Over extended periods of time you may find that this list gets filled with very long running JDBC requests. In order to see some of the medium running requests again you can empty the list by clicking on the "Reset Longest" button.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.

Metrics Settings

The Metrics Settings page lets you modify various elements relating to Metrics. The page has the following items:

Metrics:

Enable Metrics	Enable Metrics in order to access the Request Metrics, Longest Requests and Slow Requests pages. You can reset all metrics by disabling and re-enabling Metrics.
----------------	--

Slow and Longest Running Requests:

History Size	This is the number of requests which will be shown on the Longest Requests and Slow Requests pages.
Slow Request Threshold (s)	If a request takes longer than this number of seconds then it is considered a slow running request. It will be displayed on the Slow Requests page and the slow page count on the Request Metrics page will be increased. NOTE: <i>This value does NOT effect that used in Crash Protection to decide Runtime Timeout.</i>

Slow and Longest Running Requests:

History Size	This is the number of JDBC requests which will be shown on the Longest JDBC page.
--------------	---

Resources

Request Activity Graph

The Request Activity Graph is a graph of the request activity samples currently stored in memory by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

The graphs shows two independent values overlaid (transparent); the number of requests completed per second since the previous sample on the graph (Blue) and the number of requests running at the time that the sample is taken (Gray). This graph gives you insight into the server Request activity; if the number of active requests stays high and constant then your server is under load. It could be that the requests are taking too long to execute for the load being placed upon the server.

Placing your mouse pointer on a data point will show a tooltip with details about that sample.

Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.

Note: After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

[Importing and Graphing Data in Excel](#)

Request Time Graph

The Request Average Execution Time Graph is a graph of the request performance samples currently stored in memory by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

The graphs shows two independent values overlaid (transparent); the average execution time of the completed requests since the previous sample on the graph (Blue) and the average execution time of active requests at the point that the sample is taken (Gray). This graph gives you insight into the server Request performance; if the average execution time of the active requests is continuing to increase it could be that some requests are hanging.

Placing your mouse pointer on a data point will show a tooltip with details about that sample.

Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.

Note: After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

Importing and Graphing Data in Excel

JDBC Activity Graph

The JDBC Request Activity Graph is a graph of all the JDBC request activity samples currently stored by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

The graphs shows two independent values overlaid (transparent); the number of JDBC requests completed per second since the previous sample on the graph in blue, and the number of requests running at the time that the sample is taken. This graphs give you insight into the JDBC Request activity; if the number of active requests stays high and constant then your server is under load. It could be that the requests are taking too long to execute for the load being placed upon the server.

Placing your mouse pointer on a data point will show a tooltip with details about that sample.

Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.

Note: After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

Importing and Graphing Data in Excel

JDBC Time Graph

The JDBC Request Average Execution Time Graph is a graph of all the JDBC request performance samples currently stored by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

The graphs shows two independent values overlaid (transparent); the average execution time of the completed JDBC requests since the previous sample on the graph (Blue) and the average execution time of active JDBC requests at the point that the sample is taken (Gray). This graph gives you insight into the JDBC Request performance; if the average execution time of the active JDBC requests is continuing to increase it could be that some JDBC requests are hanging.

Placing your mouse pointer on a data point will show a tooltip with details about that sample.

Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.

Note: After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

[Importing and Graphing Data in Excel](#)

Memory Graph

The Memory Graph will show you a graph of the memory samples currently stored in memory by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

From this graph you can see the amount of unallocated memory, the amount of allocated memory and the amount of memory actually being used when the sample was taken.

Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.

Why would I want to see a memory graph?

Because of its visual nature, the memory graph can very quickly give you a good idea how the system is performing. Not only that but, because it can cover larger periods of time, it can warn you of trends which could become a problem in the future. This kind of trend is much more

difficult to spot if you are only checking the system periodically and are only looking at the health of the system right at this point in time.

Note: After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

Importing and Graphing Data in Excel

CPU Graph

The CPU Graph will show you a graph of the CPU samples currently stored in memory by FusionReactor. The time span of these samples will depend upon your history size and interval time (see Resource Settings). You can zoom into the graph and scroll through it in order to get a better view of this data. By default the Resource Settings store the last 12 hours of data (sampled every 5 seconds).

Refreshing the Page

Because this graph shows a large amount of information, there is no automatic refresh on this page, but you can click the refresh button in the page header to get the latest set of data.

Note: After a server restart the graphs are reset but the resource logs will contain entries from before the server restart. The amount of data stored in the resource log files is also typically far greater than the number of requests FusionReactor stores in memory to show on the graph screens. If the graphs no longer contain the data from the period you wish to examine, then you should analyze the Resource Log files.

Also See

Importing and Graphing Data in Excel

List All Threads

The List All Threads page shows all threads running within the JVM. The threads table has the following columns:

Group	This is how Threads are organised within the JVM.
Name	This is the name of a Thread.

Managing Threads

Each thread listed on this page has two icons next to it. They are:

Stack	If you click this button then you will be taken to the Stack Trace page which
-------	---

Trace	lets you see exactly what is currently going on inside a thread.
Kill Thread	Clicking on this link will prompt you if you are sure. OK the dialog to kill the thread.

WARNING: All JVM threads are listed in this page, including those pertaining to JVM system and J2EE container functions. Killing a thread is a last resort to solving problems and we must caution you that, while FusionReactor does allow you to stop any thread, it is highly inadvisable to stop any threads in the "system" thread group, or any threads integral to the running of your J2EE container.

Refreshing the Page

At the top of the page you will see the standard page header. At any time you can click the refresh button to update the main content. Alternatively you can select a time period from the drop-down box to update the content automatically. Selecting "None" from the drop-down box will stop the automatic refresh.

Resource Settings

The Resource Settings page lets you modify various elements relating to resource sampling and logging. The page has the following items:

Resource Sampling (CPU/Memory):

Initial Delay (ms)	This value defines how long after startup FusionReactor will wait before starting to periodically gather resource information. This is done so that the server can recover from the extra overhead of startup and your graphs will represent normal operating behaviour.
Sample Interval (ms)	Change this value to alter the amount of time between successive resource samples.
History Size	This value defines the maximum number of resource samples stored in memory. Once the history size reaches its maximum, old values will be removed to make room for new ones.

Resource Logging (CPU/Memory/Request/JDBC):

Resource Logging	If Resource Logging is enabled then resource information will be written to the Resource Log.
View Size (KB)	This value indicates how much of the log file should be displayed when you view the Resource Log from inside FusionReactor.
File Size (KB)	Specifies the size of each file in the Resource Log rotating file set.
File Count	Specifies the number of files in the Resource Log rotating file set.

Resource Log

The Resource Log page is a simple display page which shows the contents of the Resource Log file. The Resource Logs (a set of rotating log files accessible over the Resource Settings page) are simple space-separated files so they can be easily imported into database or spreadsheet applications. By default these files are located in the

/FusionReactor/instance/[your_instance_name]/log directory.

The log file is formatted as follows:

Field	Number	Description
Date (Formatted)	1 [A]	The date on which this resource probe was triggered.
Time (Formatted)	2 [B]	The time at which this resource probe was triggered.
Date/Time (ms)	3 [C]	This is the millisecond time at which at which this resource probe was triggered.
Version	4 [D]	Version number of this log entry.
Server Startup Time (ms)	5 [E]	Startup time of the application server in milliseconds (ms) since epoch.
MEM	6 [F]	<i>Fixed Separator</i>
Used Memory (percentage)	7 [G]	The amount of used memory (expressed as a percentage.).
Max Memory (KB)	8 [H]	The total amount of physical memory available to this instance.
Used Memory (KB)	9 [I]	The amount of used memory.
Total Memory (KB)	10 [J]	The amount of allocated memory.
Free Memory (KB)	11 [K]	The amount of free memory (within the allocated block).
CPU	12 [L]]	<i>Fixed Separator</i>
CPU Time (percentage)	13 [M]	This is the amount of CPU time used by CF since the last CPU value was measured. It is expressed as a percentage of the total CPU time (across all CPU cores) that was available during that period.
Total CPU (percentage)	14 [N]	This is the amount of CPU time used as a whole (all running applications) by the server since the last CPU value was measured. It is expressed as a percentage of the total CPU time (across all CPU cores) that was available during that period.
REQ	15 [O]	<i>Fixed Separator</i>
Total Request Count	16 [P]	The total number of requests that have been called since the last restart.
Total Finished Request Count	17 [Q]	The total number of requests that have been completed since the last restart.
Request Count	18 [R]	The number of requests that were running at the time of this log message.
Queued Request	19 [S]	The number of requests that were queued at the time of this log message.

Count		
Finished Request Count	20 [T]	The number of requests that were completed since the last log message.
Requests per second	21 [U]	The number of requests per second that where being served at the time of this log message.
Average Request time	22 [V]	The average request execution time at the time of this log message.
JDBC	<i>23 [W]</i>	<i>Fixed Separator</i>
Total JDBC Count	24 [X]	The total number of JDBC requests that have been called since the last restart.
Total Finished JDBC Count	25 [Y]	The total number of JDBC requests that have been completed since the last restart.
JDBC Count	26 [Z]	The number of JDBC requests that were running at the time of this log message.
JDBC Finished	27 [AA]	The number of JDBC requests that were completed since the last log message.
JDBC Requests per second	28 [AB]	The number of JDBC requests per second that where being served at the time of this log message.
Average JDBC time	29 [AC]	The average JDBC request execution time at the time of this log message.

Deleting the CPU/Memory Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the CPU/Memory Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Also See

Importing and Graphing Data in Excel

Crash Protection

Crash Protection Restrictions

Crash Protection restrictions are designed to let you specify which pages you want to be monitored by the Crash Protection system, or alternatively, which pages you do NOT wish to be monitored. All remaining pages to be protected. To select which mode the Restrictions page works in, change the Restrictions option on the Crash Protection Settings page.

The Crash Protection Restrictions page has two main areas: At the top of the page there is a form for you to enter new restrictions and below this is a list of all existing restrictions for this feature. To enter a restriction you will need to fill the following fields:

Request	This is the page you want to restrict. Firstly select if you are creating an "Exact Match" or "Regular Expression". Next enter the location of your page. By default the request will be an exact match and should look something like <code>"/batchStats.cfm"</code> or <code>"/batch/dailyStats.cfm"</code> (beginning with a <code>'/'</code> and without parameters). A good use of Regular Expressions here would be to restrict everything within a given folder. For examples of Regular Expressions, please see Regular Expressions and Restrictions.
Hostname	By default, FusionReactor will ignore the hostname and simply match based on the web path of a given request, but you can also opt to include it. If you enable Hostname checking then the request field should also include a hostname and look something like <code>"productionMachine/batchStats.cfm"</code> . You should NOT include HTTP or HTTPS at the start of this string.
Parameters	By default, FusionReactor will ignore any parameters passed to a page, however, as with the Hostname, you can also opt to include them in your restriction. With parameter checking enabled, the request field should include the necessary parameters. For example <code>"/batchStats.cfm?period=YEAR"</code> . Once again, Regular Expressions can be useful here when specifying multiple possible parameter values, or even specifying parameters where the order isn't fixed.
Exclude From/Protection Type	The meaning of this field changed depending upon what mode the Crash Protection restrictions page is working in (Crash Protection Settings page). You can either select whether the specified request will be excluded from Crash Protection, or if it should be included.
Statistics	This field is only visible when the Crash Protection restrictions page is in exclude mode (Crash Protection Settings page). If you are excluding a request from Crash Protection because you know it will run for a long time, you can also exclude it from the general server statistics.

Once you submit the form the page will be refreshed and your new restriction will be listed at the bottom of the page.

Managing Existing Restrictions

Below the Add Restriction form you will see a list of all existing restrictions. Each restriction has two icons next to it:

Modify Restriction	Clicking on this icon will take you to the Modify Restriction page. From this page you can change any of the information which is listed above.
--------------------	---

Delete Restriction	Clicking on this icon will prompt you if you are sure you want to delete this restriction. If you OK the dialog then the restriction will be removed from the Existing Restrictions list and the page will be refreshed.
--------------------	--

Crash Protection Settings

Unattended Monitoring via Crash Protection is a way to protect your server from various types of load. You can protect from excessive numbers of requests all being run at the same time. You can protect from individual requests running for long periods of time and you can protect from requests running when the server is low on memory. FusionReactor can notify you when any of these situations occur.

The Crash Protection Settings page lets you modify various elements relating to the different forms of Crash Protection.

How do I set up Unattended Monitoring with Crash Protection?

From the FusionReactor Administrator, click "CP Settings" from the table of contents (Crash Protection section) and you will see the Crash Protection Settings screen. From here you can activate any of the crash protection triggers and set up what action should be taken if they trigger:

Crash Protection:

Timeout Protection (s)	By entering a number (in seconds) in this field you will enable Timeout Crash Protection, from that point on, any request which takes longer to execute than this time will trigger Crash Protection. To turn this feature off again, simply remove any value from the field. You can specify what should occur when Timeout Crash Protection is triggered: <ul style="list-style-type: none"> • Abort and Notify - The request will be aborted and a notification email will be sent • Notify - The request is left running and a warning email is sent.
Memory Protection (%)	By entering a number (percentage) in this field you will enable Memory Crash Protection. From that point on, if the amount of free memory on the server is less than this number when a request arrives, Crash Protection will be triggered. To turn this feature off again, simply remove any value from the field. You can specify what should occur when Memory Crash Protection is triggered: <ul style="list-style-type: none"> • Abort and Notify - The request will be aborted and a notification email will be sent • Queue and Notify - The request will be added to a queue in order to be executed once the server has recovered and a notification email will be sent. You can specify the queue timeout in the Survival Strategy section of this page. • Notify - The request is left running and a warning email is sent.
Request Protection	By entering a number in this field you will enable Request Quantity Crash Protection. From that point on, if the amount of currently running requests on the server is greater than this number when a request arrives, Crash Protection will be triggered. To turn this feature off again, simply remove any value from the field. You can specify what should occur when Request Quantity Crash Protection is triggered: <ul style="list-style-type: none"> • Abort and Notify - The request will be aborted and a notification

	<p>email will be sent</p> <ul style="list-style-type: none"> • Queue and Notify - The request will be added to a queue in order to be executed once the server has recovered and a notification email will be sent. You can specify the queue timeout in the Survival Strategy section of this page. • Notify - The request is left running and a warning email is sent.
--	--

NOTE: In order to receive notification email you must set up your Email Server on the FusionReactor Settings page. You can also change the behaviour of the Request Abort and Request Queue in the Survival Strategy part of this page.

Crash Protection Restrictions:

Restrictions	From here you can enable or disable the Crash Protection Restrictions feature. When restrictions are enabled you can define which pages will be run through the Crash Protection system and which will be left to run normally. You can defined individual restrictions on the CP Restrictions page.
Behavior	<p>This value defines how Crash Protection Restrictions behave.</p> <ul style="list-style-type: none"> • Protect Matching Requests - No requests will be subject to protection (they will never be aborted or timed out) unless they are matched by a Restriction rule as defined on the CP Restrictions page. • Exclude Ignore Requests - All requests will be subject to protection (they may be aborted or timed out) unless they are matched by a Restriction rule, in which case they will not be touched by Crash Protection under any circumstances as defined on the CP Restrictions page.

What is a Survival Strategy and Self-healing?

The survival strategy is how FusionReactor will attempt to self-heal the situation if crash protection is triggered. It can be configured to kill problem requests and queue new requests until the situation is stabilized. It will also alert you so that you have time to resolve the problem if you need to take action.

Survival Strategy:

Abort Strategy	<p>When a request is aborted due to Crash Protection there are two possible behaviours:</p> <ul style="list-style-type: none"> • Display Message - Enter your message in the text box provided and that message will be presented to the user who made the aborted request. • Redirect to URL - Enter a URL in the text box provided and the user who made the aborted request will instead be redirected to this URL. <p>NOTE: Redirects may not work when a page has already sent data to the browser.</p>
Add Parameters	<p>After an abort, if you "Redirect to URL" then this option defines whether FusionReactor will add extra parameters to your URL so you can see exactly what happened. If you enable parameters the your URL will include the following:</p> <ul style="list-style-type: none"> • detection_method - This is the type of Crash Protection which

	<p>fired. It will be one of: "freemem", "requests", or "timeout"</p> <ul style="list-style-type: none"> • threshold_value - This is the threshold value associated with the type of Crash Protection, as defined earlier on this page. • actual_value - This is the actual which caused the Crash Protection to fire. • source_url - Contains the URL sent on the original request e.g. http://127.0.0.1/test.cfm (since 3.0.1) • source_url_params - Contains the URL parameters sent on the original request (since 3.0.1) • source_method - This is the HTTP method used on the original request e.g. GET, POST (since 3.0.1) • start_time - This is the start time of the request from Epoch e.g. 1209376581900 (since 3.0.1)
Queue Timeout / (Request Abort) (seconds)	If a request is queued due to Crash Protection, this value represents the longest amount of time that the request will be held before FusionReactor abandones it.

Crash Protection Logging:

CP Logging	If Crash Protection Logging is enabled then Crash Protection information will be written to the Crash Protection Log.
View Size (KB)	This value indicates how much of the log file should be displayed when you view the Crash Protection Log from inside FusionReactor.
File Size (KB)	Specifies the size of each file in the Crash Protection Log rotating file set.
File Count	Specifies the number of files in the Crash Protection Log rotating file set.

How do I set up Notifications?

From the FusionReactor Administrator, click "Settings" from the table of contents (FusionReactor section) and you will see the FusionReactor Settings screen. The Email section of this screen deals with the notification.

Why would I want to enable Crash Protection?

There are many cases in which crash protection can be an invaluable tool. These include:

Easing Support

Crash Protection automates the manual process of checking each of your servers and verifying that they are operating within satisfactory parameters. Instead of having to periodically go through each server you can have warnings sent directly to your inbox.

Request Bursts

If you have a server which receives bursts of requests all at once then FusionReactor can use queuing spread that load out over time and so protect your server from becoming overwhelmed.

Preventing Crashes

If you periodically have very resource hungry jobs running then FusionReactor can stop new requests being run which could potentially cause the server to become unstable.

During Development

If you need to write a page which cannot use more than a fixed amount of the system resources then crash protection is a good way of tracking where and when you have problems. As well as emails, crash triggers can also be viewed in a logfile.

Crash Protection Log

The Crash Protection Log page is a simple display page which shows the contents of the Crash Protection Log file. The Crash Protection Logs are a set of rotating files which you can configure from the Crash Protection Settings page. By default these files are located in the `/FusionReactor/instance/[your_instance_name]/log` directory.

The log file is formatted as follows:

Field	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time at which this log entry was written.
Date/Time (ms)	3 [C]	The number of milliseconds since the epoch.
Version	4 [D]	Version number of this log entry.
Server Startup Time (ms)	5 [E]	Startup time of the application server in milliseconds (ms) since epoch.
CP Type	6 [F]	The type of Crash Protection which caused this log entry to be written.
Action	7 [G]	The action which was carried out on this request by Crash Protection.
Memory Percentage	8 [H]	The percentage of free memory at the time of this log entry.
Memory Threshold	9 [I]	The threshold percentage of free memory set up in Crash Protection settings.
Request Count	10 [J]	The number of currently running requests at the time of this log entry.
Request Threshold	11 [K]	The threshold number of currently running requests set up in Crash Protection settings.
Execution Time	12 [L]	The current execution time of this request at the time of this log entry.
Execution Threshold	13 [M]	The threshold execution time set up in Crash Protection settings.
Server ID	14 [N]	The ID of the server for Server Up / Down notifications
Script	15 [O]	The script that you set up in the Crash Protection settings (including passed parameters if available).
Date (Formatted)	16 [P]	The date on which this request was started or finished (depending upon the Request Status.)

Time (Formatted)	17 [Q]	The time at which this request was started or finished (depending upon the Request Status.)
Date/Time (ms)	18 [R]	This is the millisecond time at which this request was started or finished (depending upon the Request Status.)
Version	19 [S]	Version number of this log entry.
Server Startup Time (ms)	20 [T]	Startup time of the application server in milliseconds (ms) since epoch.
Request ID	21 [U]	The FusionReactor Request ID.
Request Status	22 [V]	The current state of the request. This will tell you if a request has started, finished, was killed, queued, etc...
CP Reason	23 [W]	The reason given by Crash Protection for the current action.
Thread ID	24 [X]	The name of the thread responsible for responding to this request.
Client IP Address	25 [Y]	The IP address of the machine making the request.
Request Method	26 [Z]	This will usually be "GET" or "POST"
Request URL	27 [AA]	This is the requested URL. The format of this column will change depending on URL Format key which can be found on the Request Settings page.
Execution Time (ms)	28 [AB]	The amount of milliseconds it took to complete the request. (For incomplete requests, this column will be 0.)
Used Memory (percentage)	29 [AC]	The amount of memory (expressed as a percentage) which was used when this request started or finished.
Max Memory (KB)	30 [AD]	The total amount of physical memory available to this instance.
Used Memory (KB)	31 [AE]	The amount of memory which was used when this request started or finished.
Total Memory (KB)	32 [AF]	The amount of memory which was allocated by the instance when this request started or finished.
Free Memory (KB)	33 [AG]	The amount of free memory (within the allocated block) when this request started or finished.
Query String	34 [AH]	If the URL has any parameters then they will appear here.
Return Status Code	35 [AI]	This is a HTTP return code such as 200 (OK,) 404 (Not found,) or 500 (Internal Server Error.) For "Started:" rows this column will be 200.
CPU Time (ms)	36 [AJ]	This is the amount of actual CPU time which this request required. (For incomplete requests, this column will be 0.)
AMF Request	37 [AK]	If AMF decoding is enabled (Request Settings page) and there is AMF to decode, then the method names will appear here.
JSESSIONID	38 [AL]	The J2EE Session Id for this request.
CFID	39 [AM]	The CF Id for this request.
CFTOKEN	40 [AN]	The CFTOKEN for this request.
JDBC Query	41 [AO]	The number of JDBC queries run by this request at the time of

Count		this log.
JDBC Total Time (ms)	42 [AP]	The total amount of time spent running JDBC queries at the time of this log.
JDBC Total Execution Time (ms)	43 [AQ]	The total amount of time spent by the database running JDBC queries at the time of this log.
JDBC Total Rows	44 [AR]	The total number of rows returned by JDBC queries at the time of this log.
Bytes Sent	45 [AS]	The amount of data which was sent back to the client.
Time to First Byte (ms)	46 [AT]	The number of milliseconds it took to deliver the first bit of data.
Time to Last Byte (ms)	47 [AU]	The number of milliseconds it took to deliver the complete content.
Time to Stream Open (ms)	48 [AV]	The number of milliseconds before the data stream was opened.
Time to Stream Close (ms)	49 [AW]	The number of milliseconds until the completed data stream was closed.

Deleting the Crash Protection Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the Crash Protection Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Also See

Importing and Graphing Data in Excel

JDBC

JDBC Settings

The JDBC Settings page lets you modify various elements relating to JDBC. The page has the following items:

SQL Statement Recording (Request Detail Page):

Record SQL	Specifies how SQL is stored along with each request. You can either store queries in slowest first order (Slowest:default) or in the order that they are run (In Order). Alternatively you can just get a summary of run queries (Summary Only).
Record up to (queries)	<p>This limits the number of queries stored for each request. If you are storing them in order then you will see the first queries which are run. The number of SQL statements recorded on the request affects the amount of memory that FusionReactor could use. It is recommended that you increase this value in stages and check the memory usage after each increase once the server has run at least the number of request that are kept in the Request History. Each query recorded only uses a very small amount of memory but this may become significant if you store hundreds or thousands of queries per request when multiplied by the number of requests maintained in your request history setting.</p> <p>When queries are recorded in Slowest Order and the number of queries to record has been reached, a slower query will be correctly inserted into the list of recorded queries according to the time it has taken to run, and the last entry in the list will removed and added to the Summary Information.</p>
Only queries slower than (ms)	This threshold value lets you record only SQL statements that run slower than the amount of time you specify here. If the SQL statement runs faster than the time you specify it's data will still be added to the Summary information.

Additional Information:

Query Location	When this feature is enabled and a JDBC request is executed, FusionReactor will store a Java Stack Trace which can be used to find out the location of each query that ran. You change the amount of information which gets stored in this stack trace by going to the JDBC Stack Trace Filter page. The resulting stack traces are visible from the Request Detail page and within the JDBC log files.
----------------	---

JDBC Logging (Log Files):

JDBC Logging	If JDBC Logging is enabled then JDBC information will be written to the JDBC Log.
View Size (KB)	This value indicates how much of the log file should be displayed when you view the JDBC Log from inside FusionReactor.
File Size (KB)	Specifies the size of each file in the JDBC Log rotating file set.
File Count	Specifies the number of files in the JDBC Log rotating file set.
Only queries	Only log queries that run longer than this time will be written to the log files.

slower than: (ms)	(Optional)
----------------------	------------

NOTE: These settings pertain only to DSNs wrapped with the FusionReactor JDBC Driver Wrapper. For more information, please see the JDBC Driver Wrapper: User Guide and JDBC Tutorial PDF files (from the www.fusion-reactor.com website.)

JDBC Stack Trace Filter

The JDBC Stack Trace Filter page lets you set up various filters which are used to make the Java Stack Trace more manageable. The complete Java Stack Trace contains far more information than you are likely to need, so this page has been added to let you tailor it to your own preferences. By default, FusionReactor will show you any Stack Trace rows which contain references to cfm, cfc or jsp pages.

To add a new filter, just type a string into the field provided and click the "Add Filter Pattern" button. Any lines in the stack trace which contain that string will be stored along with the JDBC request.

Managing Existing Filters

Below the Add Filter Pattern form you will see a list of all existing filter patterns. Each pattern has two icons next to it:

Modify Pattern	Clicking on this icon will take you to the Modify JDBC Java Stack Trace Filter page. From this page you can change the filter pattern.
Delete Pattern	Clicking on this icon will prompt you if you are sure you want to delete this filter pattern. If you OK the dialog then the filter pattern will be removed from the Existing JDBC Java Stack Trace Filters list and the page will be refreshed.

JDBC Log

The JDBC Log page is a simple display page which shows the contents of the JDBC Log file. The JDBC Logs (which are a set of rotating files which you can configure from the JDBC Settings page) are simple space separated files, so they can be easily imported into database or spreadsheet applications. By default these files are located in the `/FusionReactor/instance/[your_instance_name]/log` directory.

The log file is formatted as follows:

Filed	Number	Description
Date (Formatted)	1 [A]	The date on which this log entry was written.
Time (Formatted)	2 [B]	The time on which this log entry was written.
Date/Time (Milliseconds)	3 [C]	This is the millisecond time at which this request was started or finished (depending upon the Request Status.)
Version	4 [D]	Version number of this log entry.
Server	5 [E]	Startup time of the application server in milliseconds (ms) since

Startup Time (ms)		epoch.
Request ID	6 [F]	The FusionReactor Request ID.
Thread ID	7 [G]	The name of the thread responsible for responding to this request.
Client IP Address	8 [H]	The IP address of the person making the request.
Request Method	9 [I]	Specifies the HTTP method used during the request. This will usually be "GET" or "POST"
Request URL	10 [J]	This is the requested URL. The format of this column will change depending on URL Format setting which can be found on the Request Settings page.
Log Type	11 [K]	One of: "METRIC", "NOTIFICATION" or "REMINDER"
DB Start Time	12 [L]	The time (in milliseconds) at which the JDBC query began.
DB End Time	13 [M]	The time (in milliseconds) at which the DB finished with the query.
Total End Time	14 [N]	The time (in milliseconds) at which the query was completely finished (after results were transferred back etc.).
DB Time	15 [O]	The amount of time (in milliseconds) this JDBC query spent in the DB.
Total Time	16 [P]	The total amount of time (in milliseconds) it took to complete this query.
Row Count	17 [Q]	The number of rows returned by this query.
Prepared Statement	18 [R]	This is a simple TRUE/FALSE value to tell you if this statement had already been prepared prior to its execution.
Row Limited	19 [S]	This is a simple TRUE/FALSE value to tell you if this query was row limited.
DataSource	20 [T]	If you have assigned a name to your data source wrapper then it will appear here. To find out how to name your data sources, please refer to page 14 of the FusionReactor JDBC Driver Wrapper: User Guide (pdf).
Statement	21 [U]	The actual statement which was run.
Stack Trace	22 [V]	This is the stack trace which was stored when the query was run. You can turn on or off this feature from the JDBC Settings page and alter the stack trace information from the JDBC Stack Trace Filter page.
Query String	23 [W]	This is the query string from the request running the JDBC statement.
Message	23 [X]	For Notifications and Reminders, this will contain the notification or reminder message.

NOTE: Depending on the type of log message, not all columns will have values.

Deleting the JDBC Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the JDBC Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Also See

Importing and Graphing Data in Excel

Instances

Instance Manager (Enterprise Version Only)

The Instance Manager page is used to keep track of the different servers running on a machine and which ones have FusionReactor installed on them. The page is split into two main sections. At the top of the page is a form letting you add new servers and at the bottom is a list of all known servers along with the instances on each server.

Adding a Server

To add a new server to the Instance Manager, firstly select the type of server you want to add from the drop-down box and then enter the file path to the root of that server. Clicking the "Add Server" button will cause FusionReactor to attempt to find the server. If it can be found then the server will appear at the bottom of the page along with all instances currently defined within it.

Scanning the System for New Servers

Instead of specifying a single server directly, you can also get FusionReactor to automatically scan the file system for servers. To do this, enter the starting directory for your search and how many folders deep FusionReactor should search, then click the "Scan System" button.

NOTE: Please be aware that asking FusionReactor to search more than 3 or 4 levels deep on a complex file system can take several minutes to complete.

Managing Servers and Instances

Below the server adding forms you will see a list of all servers currently set up. Each server has an icon next to it:

Remove Server	If no instances within a server are currently running FusionReactor then you can remove it from the Instance Manager completely. NOTE: This option just removes it from FusionReactor. It does NOT remove the server from the machine.
---------------	--

Each server also displays a list of instances. Each instance has three icons next to it:

Install FusionReactor	This button will let you install FusionReactor onto any instance.
Remove FusionReactor	If FusionReactor has been added to an instance then this button will let you remove it again.
Add to Dashboard	If FusionReactor has been added to an instance then clicking this button will automatically add it to the Enterprise Dashboard so you can easily monitor it.

Compression

Compression Settings

The Compression Settings page lets you turn on or off the compression features as well as letting you define how the MIME Type Restrictions behave.

Compression Support:

GZIP Requests	When this option is enabled, FusionReactor will attempt to decompress HTTP requests which are sent to the server in a compressed format.
GZIP Responses	When this option is enabled, FusionReactor will attempt to compress outgoing HTTP responses unless otherwise instructed by the MIME Type Restrictions or the URL Exclusions. NOTE: <i>FusionReactor will only compress responses when the client request indicates that the client can support compression by sending a Accept-Encoding: gzip header</i>

MIME Type Restrictions:

MIME Type Restrictions	This value defines how MIME Type Restrictions behave. <ul style="list-style-type: none">• Ignore requests that match the rules - By default, all requests will be compressed when compression is enabled unless the page's document mime type matches one of the rules defined on the MIME Type Restrictions page (then it will NOT be compressed).• Compress requests that match the rules - no requests will be compressed unless the page's document mime type matches one of the rules defined on the MIME Type Restrictions page (then it will be compressed).
------------------------	--

MIME Type Restrictions

MIME Type Restrictions are designed to let you specify which types of file you want to be compressed, or alternatively, which types you do NOT wish to be compressed (leaving all other file types to be compressed). To select which mode the Restrictions page works in, change the Restrictions option on the Compression Settings page.

The MIME Type Restrictions page has two main areas: At the top of the page there is a form for you to enter new restrictions and below this is a list of all existing restrictions for this feature. To enter a restriction you will need to fill the following fields:

MIME Type	This is the document type that you want to restrict. Firstly select if you are creating an "Exact Match" or "Regular Expression". Next enter your MIME Type. For examples of Regular Expressions, please see Regular Expressions and Content Filters.
-----------	---

Once you submit the form, the page will be refreshed and your new restriction will be listed at the bottom of the page.

Managing Existing Restrictions

Below the Add Restriction form you will see a list of all existing restrictions. Each restriction has two icons next to it:

Modify Restriction	Clicking on this icon will take you to the Modify Restriction page. From this page you can change any of the information which is listed above.
Delete Restriction	Clicking on this icon will prompt you if you are sure you want to delete this restriction. If you OK the dialog then the restriction will be removed from the Existing Restrictions list and the page will be refreshed.

Exclude URLs

The Exclude URLs list lets you define a set of pages which will never be compressed. The Exclude URLs page has two main areas: At the top of the page there is a form for you to enter new exclusions and below this is a list of all existing exclusions for this feature. To enter an exclusions you will need to fill the following fields:

Request	This is the page you want to exclude. Firstly select if you are creating an "Exact Match" or "Regular Expression". Next enter the location of your page. By default the request will be an exact match and should look something like "/batchStats.cfm" or "/batch/dailyStats.cfm" (beginning with a '/' and without parameters.) A good use of Regular Expressions here would be to restrict everything within a given folder. For examples of Regular Expressions, please see Regular Expressions and Restrictions.
Hostname	By default, FusionReactor will ignore the hostname and simply match based on the web path of a given request, but you can also opt to include it. If you enable Hostname checking then the request field should also include a hostname and look something like "productionMachine/batchStats.cfm". You should NOT include HTTP or HTTPS at the start of this string.
Parameters	By default, FusionReactor will ignore any parameters passed to a page, however, as with the hostname, you can also opt to include them in your restriction. With parameter checking enabled, the request field should include the necessary parameters. For example "/batchStats.cfm?period=YEAR". Once again, Regular Expressions can be useful here when specifying multiple possible parameter values, or even specifying parameters where the order isn't fixed.

Once you submit the form the page will be refreshed and your new exclusion will be listed at the bottom of the page.

Managing Existing Exclusions

Below the Add Exclusion form you will see a list of all existing exclusions. Each exclusion has two icons next to it:

Modify Exclusion	Clicking on this icon will take you to the Modify Exclusion page. From this page you can change any of the information which is listed above.
Delete Exclusion	Clicking on this icon will prompt you if you are sure you want to delete this exclusion . If you OK the dialog then the exclusion will be removed from the Existing Exclusion list and the page will be refreshed.

Content Filters

Search and Replace

The Search and Replace feature is arguably one of the most powerful features in FusionReactor. It allows you to make temporary changes to your pages without having to modify source code. Fix broken links, spelling mistakes or html problems. Search and Replace Filters enable you to change the content sent to the web browser by replacing or removing text on the fly. The Search and Replace page has two main areas: At the top of the page there is a form for you to enter new filters and below this is a list of all existing filters for this feature. To enter a Search and Replace filter you will need to fill the following fields:

Find What	Firstly you should select if you want to search for an Exact String, or if you want to use Regular Expressions to search for a pattern. For examples of Regular expressions, see Regular Expressions and Content Filters. Once you have selected the search type, you should enter the your search string in the text field provided.
Replace With	This field contains the text you want to replace your search string with. If you just want to search for a string and remove it from your page then simply leave this field blank.
Replace	This option lets you decide if you want to replace all instances of your search string, or just the first one found on the page.

Managing Existing Search and Replace Filters

Below the Add Group form you will see a list of all existing filters. Each group has two icons next to it:

Modify Group	Click this icon to change a group name or to change the servers which are currently a part of this group.
Delete Group	Clicking on this icon will prompt you if you are sure. OK the dialog box to delete the group.

Content Filter Restrictions

Content Filter restrictions are designed to let you specify which pages you want run through the content filter system, or alternatively, which pages you do NOT wish to be filtered (leaving all remaining pages to run through content filtering.) To select which mode the Restrictions page works in, change the Restrictions option on the Content Filter Settings page.

The Content Filter Restrictions page has two main areas: At the top of the page there is a form for you to enter new restrictions and below this is a list of all existing restrictions for this feature. To enter a restriction you will need to fill the following fields:

Request	This is the page you want to restrict. Firstly select if you are creating an "Exact Match" or "Regular Expression". Next enter the location of your page. By default the request will be an exact match and should look something like <code>"/batchStats.cfm"</code> or <code>"/batch/dailyStats.cfm"</code> (beginning with a <code>'/'</code> and without
---------	--

	parameters.) A good use of Regular Expressions here would be to filter everything within a given folder. For examples of Regular Expressions, please see Regular Expressions and Restrictions.
Hostname	By default, FusionReactor will ignore the hostname and simply match based on the web path of a given request, but you can also opt to include it. If you enable Hostname checking then the request field should also include a hostname and look something like "productionMachine/batchStats.cfm". You should NOT include HTTP or HTTPS at the start of this string. With this option you could dynamically change the name of a website based on the URL used to access it.
Parameters	By default, FusionReactor will ignore any parameters passed to a page, however, as with the Hostname, you can also opt to include them in your restriction. With parameter checking enabled, the request field should include the necessary parameters. For example "/batchStats.cfm?period=YEAR". Once again, Regular Expressions can be useful here when specifying multiple possible parameter values, or even specifying parameters where the order isn't fixed.

Once you submit the form the page will be refreshed and your new restriction will be listed at the bottom of the page.

Managing Content Filter Restrictions

Below the Add Restriction form you will see a list of all existing restrictions. Each restriction has two icons next to it:

Modify Restriction	Clicking on this icon will take you to the Modify Restriction page. From this page you can change any of the information which is listed above.
Delete Restriction	Clicking on this icon will prompt you if you are sure you want to delete this restriction. If you OK the dialog then the restriction will be removed from the Existing Restrictions list and the page will be refreshed.

Filter Settings

The Content Filter Settings page lets you turn on or off the Search and Replace feature as well as letting you define how the Filter Restrictions behave.

Content Filters:

Search and Replace	When this option is Enabled, FusionReactor will run responses through the Search and Replace feature unless otherwise instructed by the Filter Restrictions.
--------------------	--

Content Filter Restrictions:

Filter Restrictions	<p>This value defines how Filter Restrictions behave.</p> <ul style="list-style-type: none"> • Filter requests that match the rules - By default, no requests will be run through the Search and Replace filters. If a request matches one of the rules defined on the Filter Restrictions page then it will be filtered. • Ignore requests that match the rules - By default, all requests
---------------------	---

	<p>will be run through the Search and Replace filters. If a request matches one of the rules defined on the Filter Restrictions page then it will not be filtered.</p>
--	--

FusionReactor

Enabling/Disable

This page lets you completely enable or disable FusionReactor. Simply click on the "Disable FusionReactor" or "Enable FusionReactor" button to change the state of the system. The current state of the system can be seen from the status icon which is displayed at the bottom of the table of contents.

FusionReactor Settings

The Request Settings page lets you modify various elements relating to the capture, storage and display of Request information. The page has the following items:

Email Server:

From Address	This will be the email address which notification emails will be sent from.
To Address	This is the email address which will receive FusionReactor notifications.
Mail Server	Enter your email server here. This can be entered either as a machine name or an IP address. NOTE: If your SMTP server does not use port 25 (default), you can specify :portnumber after the server address/IP address (for example, mail.company.com:587). NOTE: If your mail server requires authentication, you can specify a user name and password in the format [user]:[password]@[mail server]
CP Email Interval (mins)	This email interval specifies the minimum number of minutes between Crash Protection notification emails. If you are experiencing regular difficulties with a server then use this option to avoid filling your email inbox!
Notification	This lets you completely enable or disabled email notification.

FusionReactor Internal Web Server:

Web Server	Enable this option to allow access to FusionReactor over its internal web server. When enabled, browsing the server using the machine name and port specified below will let you connect directly to FusionReactor.
Web Server Address	This is the machine name or IP address which you will use to access the internal web server. NOTE: You can leave this field blank to listen on all addresses.
Web Server Port	This is the port number which you will use to access the internal web server.

Access To FusionReactor From The External Web Server:

Enable Access	Enable this option to allow access to FusionReactor through the external web server (typically port 80)
---------------	---

FusionReactor Web Root:

Web Root	This will be the prefix for FusionReactor requests. If the web root was set to "/fusionreactor/" and you were accessing the internal web server on a local machine over port 8088 then your final URL will look like this:
----------	--

	http://127.0.0.1:8088/fusionreactor/
--	---

FusionReactor Log File:

Logging	If FusionReactor Logging is enabled then general information will be written to the FusionReactor Log.
View Size (KB)	This value indicates how much of the log file should be displayed when you view the FusionReactor Log from inside FusionReactor.
File Size (KB)	Specifies the size of each file in the FusionReactor Log rotating file set.
File Count	Specifies the number of files in the FusionReactor Log rotating file set.

FusionReactor Restrictions:

Restrictions	<p>This value defines how FusionReactor Restrictions behave.</p> <ul style="list-style-type: none"> • Monitor requests that match the rules - By default, no requests will appear in FusionReactor. If a request matches one of the rules defined on the FusionReactor Restrictions page then it will appear. • Ignore requests that match the rules - By default, all requests will appear in FusionReactor. If a request matches one of the rules defined on the FusionReactor Restrictions page then it will be invisible.
--------------	---

Stack Trace : (Configuration Item only appears when using a Java 1.5 or higher JVM)

Stack Trace Method	<p>Lets you select the method used by FusionReactor to generate and capture stack traces. Older versions of Java (before Java version 1.5) could not easily generate and capture stack traces of the running JVM but FusionReactor adds this functionality even on to such older versions. It is recommended to use the Java's Stack Trace method on Java version 1.5 or higher. On older versions of Java this configuration item does not appear because FusionReactor's default stack trace method is the only method that can be used to generate stack traces.</p> <ul style="list-style-type: none"> • Use Java's Stack Trace - Use Java's built in implementation of generating stack traces • FusionReactor Default - Use FusionReactor's
--------------------	--

FusionReactor Restrictions

FusionReactor restrictions are set up so that FusionReactor can ignore certain requests, or alternatively, you can specify a set of pages which you want to monitor with FusionReactor while all others are ignored. To select which mode the Restrictions page works in, change the Restrictions option on the FusionReactor Settings page.

The FusionReactor Restrictions page has two main areas: At the top of the page there is a form for you to enter new restrictions and below this is a list of all existing restrictions for the application. To enter a restriction you will need to fill the following fields:

Request	This is the page you want to restrict. Firstly select if you are creating an "Exact Match" or "Regular Expression". Next enter the location of your page. By default the request will be an exact match and should look something like <code>"/batchStats.cfm"</code> or <code>"/batch/dailyStats.cfm"</code> (beginning with a <code>/</code> and without parameters.) A good use of Regular Expressions here would be to restrict everything within a given folder. For examples of Regular Expressions, please see Regular Expressions and Restrictions.
Hostname	By default, FusionReactor will ignore the hostname and simply match based on the web path of a given request, but you can also opt to include it. If you

	enable Hostname checking then the request field should also include a hostname and look something like "productionMachine/batchStats.cfm". You should NOT include HTTP or HTTPS at the start of this string.
Parameters	By default, FusionReactor will ignore any parameters passed to a page, however, as with the Hostname, you can also opt to include them in your restriction. With parameter checking enabled, the request field should include the necessary parameters. For example "/batchStats.cfm?period=YEAR". Once again, Regular Expressions can be useful here when specifying multiple possible parameter values, or even specifying parameters where the order isn't fixed.

Once you submit the form the page will be refreshed and your new restriction will be listed at the bottom of the page.

Managing Existing Restrictions

Below the Add Restriction form you will see a list of all existing restrictions. Each restriction has two icons next to it:

Modify Restriction	Clicking on this icon will take you to the Modify Restriction page. From this page you can change any of the information which is listed above.
Delete Restriction	Clicking on this icon will prompt you if you are sure you want to delete this restriction. If you OK the dialog then the restriction will be removed from the Existing Restrictions list and the page will be refreshed.

Change Password

From this page you can change the password for the various users within the system. To change a password, select the user type and enter a new password (plus verification.)

If you want to disable a user type then click on one of the "Disable [user type]" buttons.

To enable a user type which has been disabled, just set a new password for that user type.

FusionReactor Log

The FusionReactor Log page is a simple display page which shows the contents of the FusionReactor Log file. The FusionReactor Logs are a set of rotating files which you can configure from the FusionReactor Settings page. By default these files are located in the /FusionReactor/instance/[your_instance_name]/log directory.

Deleting the FusionReactor Log

Clicking on the "Delete Log File" button will display an "Are you sure?" dialog box. If you OK this dialog then the FusionReactor Log will be emptied and the screen will be refreshed.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

About

Clicking on the about page will show you your current Licence information along with a form to let you upload a new licence file. To upload a license you will first need a license file (license.lic). Click the "Browse..." button and locate the license file. Once entered, click the "Upload New License File" button to update FusionReactor. To complete your license install, click on the link which says "Click on this link to activate the current license" and FusionReactor will attempt to activate your license automatically. If you do not have a direct internet connection then just follow the instructions given to you on the page.

Keeping FusionReactor Up To Date

At the bottom of the About page you will see a link named "check for FusionReactor updates". When you click on this link you will be taken to the FusionReactor website. If there are any updates available for your version of FusionReactor then you will be given links to download them.

Help

Clicking on this link will take you to the help system which you are reading right now!

Logout

Click the Logout link to exit the FusionReactor Administrator and be returned to the login page.

Request Detail

Stack Tracing

Stack traces are split into two sections. Firstly comes a set of information associated with this thread, or request. Please note that you will see more information here when Stack Tracing a request rather than a thread.

After these information fields you will see the Stack Trace itself. Note that the format of Stack Traces may alter depending on your operating system or Java version. Generally you will see a list of class names, method calls and line numbers. The most current list item will be listed first.

Killing a Thread or Request

If you click on the "Kill" button then you will see a prompt asking if you are sure. OK this dialog box to kill the request or thread.

WARNING: The "Kill" button does not give the requests time to complete in their own time. Using this feature can cause the JVM to become unstable. Also, while FusionReactor does allow you to stop any thread, it is highly inadvisable to stop any threads in the 'system' thread group, or any threads integral to the running of your J2EE container.

Refreshing the Page

Clicking on the "Refresh" button will update the screen to display the latest log entries.

Request Detail Page

You can get to the Request Detail page from any page which lists requests: Running Requests, Request History, Longest Requests, Slow Requests, and pages linked from Request Metrics.

At the top of the Request Details page you will see a copy of the request line from the page you came from. Underneath that there are six tabs.

The Tabs

Main	The Main tab shows general information about a request such as script name, execution time, memory levels and flags.
Headers	The Headers tab shows the HTTP headers for the request and response objects.
Cookies	The Cookies tab takes the Cookie string from the Headers tab and renders it in a more readable format along with other cookie information such as maximum age and whether it is secure.
JDBC	The JDBC tab will show any queries which have been run by your page. If you also have User Trace Markers on this page then you can click the "Include Markers" link to create a merged view of both markers and queries. At the bottom of the JDBC tab you will find the total and average times for queries running on this page.
Markers	The Markers tab shows all User Trace Markers which were called during this

	request along with the amount of time which elapsed between them.
AMF	If Flash Remoting Support is enabled (Request Settings) then the AMF tab displays requests which came in AMF format, used in Flash Remoting and in applications such as Flex.

Refreshing the Page

You can refresh the page at any time by pressing the refresh button which is located at the top right corner of the Request Detail page.

Analyzing Data

Importing and Graphing Data in Excel

Most FusionReactor Log files can be easily imported into Excel (and other tools that can import delimited files) in order to analyze them or create graphs..

There are 4 major log files that you can import:

- request-*.log
- resource-*.log
- jdbc-*.log
- crashprotection-*.log

How Do I import a FusionReactor Log File

1. Make a copy of the log file you want to import. (It's better to work on copies than the live log files otherwise you may stop them from being written to)
2. Start Excel and from the [File] menu select [Open...] and then the log file. You may have to select Files of type: All files (*.*) to see the log files in the file picker
3. The Text Import Wizard dialog will appear
4. For the Original data type choose : Delimited and then press [Next >] at the bottom of the dialog
5. Select Space as the Delimiter and then press [Finish] at the bottom of the dialog
6. The Spreadsheet should now be filled with data in separate columns
7. Click on top of one or more columns that you wish to graph (to highlight the entire column) and then click on the small graph icon (Chart Wizard) in the toolbar (or from the menu [Insert]->[Chart]). Note you can find out what data the Excel columns contain by using the field number [letter] in the descriptions of the log files above.
8. Choose the graph style (e.g. line graph, scatter graph or bar graph) and then click finish.

Note: you can find out what data is stored in the columns using the descriptions of the log files above or by looking into the request-headers.txt, resource-headers.txt, jdbc-headers.txt or crashprotection-headers.txt files stored in the folder with the log files.

Users

Types Of User Account

FusionReactor supports three different types (or levels) of user account:

- Administrators
- Managers
- Observers

Note: FusionReactor does not have individual users with user names and passwords, but three fixed accounts each of which has its own password.

Administrators

Administrator is the most powerful type of user account in FusionReactor. They are able to access all of the functions available in FusionReactor including (Note: some points are only available on Enterprise Versions):

- Creating and Maintaining New Instances of FusionReactor on a Computer
- Configuring the FusionReactor instances visible to the Enterprise Dashboard
- Configure Crash Protection and all other settings for an instance of FusionReactor
- Managing All User Accounts
- Uploading and Activation of License Keys
- Examine All Request and Metric Data
- Kill Running Requests

Who Should use the Administrator Account?

The Administrator should be used by users who need to create and maintain the instances of FusionReactor on computer, setup and manage the instances that are visible on any Enterprise Dashboard and maintain the user accounts. For example, in hosting environments this account might be used by the hosting company to configure FusionReactor for customers on their application server instances.

Managers

Managers are "Instance Administrators". They Manage or Administer a single instance of FusionReactor. They are able to (Note: some points are only available on Enterprise Versions) :

- Configuring the FusionReactor instances visible on this instance's Enterprise Dashboard
- Add instances visible on this instance's Enterprise Dashboard (if they have the Administrator account password for the instance they wish to add)
- Online and Offline Servers in the Enterprise Dashboard
- Configure Crash Protection and all other settings for this instance of FusionReactor
- Managing Their User Account and the Observer Account
- Examine All Request and Metric Data for this instance
- Kill Running Requests

Who Should use the Manager Account?

The Manager account should be used by users who need to manage or administer an application server. For example, in hosting environments this account might be used by the customer's application administrator.

Observers

Observers are only able to view the following information on a specific instance (Note: some points are only available on Enterprise Versions):

- View the Running Request
- View the Request History
- View the Request Log
- View the Enterprise Dashboard
- View the System Metrics
- View the Longest Requests, Slow Request and Longest JDBC Request Reports
- View the Resource Reports and Graphs
- View the Resource Log, Crash Protection Log and JDBC Log
- View the Help and About Information

Who Should use the Observer Account?

The Observer account should be used by users who need to examine the performance and metrics from an application server. For example, in hosting environments this account might be used by the customer's application developers to examine how well the application performs in production.

Enterprise Dashboard

When a user clicks on a link on the Enterprise Dashboard for a monitored server they are transferred to that server instance of FusionReactor with the same level of User account that they are logged in as. For example, this means that a Manager on an instance that is monitoring another instance of FusionReactor, will be granted Manager level access to the other instance of FusionReactor.

Managing The User Accounts

FusionReactor has three user accounts (Administrator, Manager and Observer).

Enabling a User Account

The user accounts are enabled by entering a password for the account (Note: The Administrator account is always enabled and must have a password assigned) and then clicking on the [Save Password] button. Note that when you enter a password and click on the [Save Password] button the account becomes active immediately.

Disabling a User Account

The Manager and Observer accounts can be disabled using the [Disabled Manager] and [Disable Observer] buttons the Change Password screen. An account will be disabled immediately after clicking on the Disable button for the account. To enable an account again you must enter a password for the account.

Enabling Accounts (Setting Passwords) During Installation

During the installation process you must enter a password for the Administrator account, but you can also enter a password for the Manager and Observer accounts. If you do not enter a password for an account during the installation process, the account will be disabled. You can enable the account at any moment by entering a password for the account on the Change Password screen. Note that when you enter password and click on the [Save Password] button, the account becomes active immediately.

Regular Expressions

Regular Expressions and Restrictions

Regular Expression Restrictions for FusionReactor use the Java 1.4 RegEx language. For more technical information on Regular Expressions please refer to the online Java documentation

Quick Overview of RegEx Special Characters

.	matches any character
\d	matches any digit
\w	matches any word character [a-zA-Z_0-9]
\s	matches any whitespace character
[]	contents define a character class
[abc]	matches a, b or c
[^abc]	matches anything except a, b or c
	alternation operator ('or')
X Y	matches X or Y
X?	matches X once, or not at all
X*	matches X zero or more times (matching as much as possible)
X*?	matches X zero or more times (matching as little as possible)
X+	matches X one or more times
()	grouping characters
^	start of line
\$	end of line

Example Restrictions

/myfolder/myfile.cfm	This will match any URL which contains "/myfolder/myfile.cfm"
/myfolder/(.*)	This will match anything within "myfolder"
/myfile\.cfm	This will match any files named myfile.cfm which are inside a directory.
\.((cfm) (jsp))	This will match any .cfm or .jsp file.
/myfile\.cfm/?(.*)mode=add	This expects parameter checking to be enabled. It will match any file named myfile.cfm which has the string "mode=add" as part of it's parameters. (Note: This will also match mymode=add.)
/batch[0-9].cfm	This will match files names batch0.cfm, batch1.cfm ... batch9.cfm

Regular Expressions and Content Filters

Regular Expression Restrictions for FusionReactor use the Java 1.4 RegEx language. For more technical information on Regular Expressions please refer to the online Java documentation.

NOTE: Since the regular expression buffer operates continuously, it has no concept of line breaks; the '\$' (end of line) and '^' (start of line) match characters are not available. Also, pattern grouping and the backref operator '\n' are not available.

Quick Overview of RegEx Special Characters

.	matches any character
\d	matches any digit
\w	matches any word character [a-zA-Z_0-9]
\s	matches any whitespace character
[]	contents define a character class
[abc]	matches a, b or c
[^abc]	matches anything except a, b or c
	alternation operator ('or')
X Y	matches X or Y
X?	matches X once, or not at all
X*	matches X zero or more times (matching as much as possible)
X*?	matches X zero or more times (matching as little as possible)
X+	matches X one or more times
()	grouping characters

Example Search and Replace Filters

FusionReactor	This will match the exact string "FusionReactor" anywhere within a page.
[Ff]usion[Re]actor	This matches FusionReactor without needing to capitalize the F and the R
[Ff][Uu][Ss][Ll][Oo][Nn]	This matches the word "Fusion" in any case.
	This will match any link to a specific machine (which you could replace with a different link to a "Not currently available" page or something.
<meta http-equiv="refresh"(. *?)">	This will match meta refresh tags which you could then remove.

Installation

Installation Guide

The FusionReactor Setup allows you to easily install FusionReactor on Adobe ColdFusion MX, Adobe ColdFusion MX 7, Adobe ColdFusion 8, Adobe LiveCycle Data Services, Apache Tomcat 4.1 to 6, JBoss 3.2.x to 5.0.x, Jetty or Railo Railix and many other application servers running on Windows, Linux or Solaris SPARC platforms, and is the preferred way to install FusionReactor.

The FusionReactor installation guide shows you in detail how to use the FusionReactor Setup to install FusionReactor on Windows and Linux. Installation on other platforms is done similarly. The FusionReactor installation guide is available online : [FusionReactor Installation Guide \(pdf\)](#)

License File - Uploading and Activation

To upload a license file (license.lic), go to the About page and you will see a form with a file upload button. Click the "Browse..." button and locate the license file. Once entered, click the "Upload New License File" button to update FusionReactor.

After submitting the page you need to login again. You will see your new license details on the right hand of the screen.

License Activation

To complete your license install you must click on the link which says "Click on this link to activate the current license" and FusionReactor will attempt to activate your license automatically. If you do not have a direct internet connection then just follow the instructions given to you on the page.

Activation Period

Once you have uploaded a license key you have 10 days to activate the license. You are recommended to activate your license immediately after upload. If you do not activate your license within 10 days you will be unable to use FusionReactor and the Enterprise Dashboard will show an [ACT] warning light on the server until you activate the license. To activate the license after the 10 days activation period has expired you must login as the Administrator and then follow the license activation instructions on the screen or you can upload a different license file, but the activation period will not be reset and you will be required to activate the license immediately to access to FusionReactor.

Multi-Key License Files

A single license file may contain multiple license keys, one for each license you purchased. The same license key file should be uploaded and activated on each server you have purchased a license for. Each time you activate the license file the license activation count will increased, until the total number of licenses that the key file contains has been activated. If FusionReactor reports

that you have no Activations remaining on this Multi-Key file and you believe this to be an error, please contact FusionReactor support.

default-web.xml File Modifications

The following extract of the default-web.xml file will enable you to see the changes made by the FusionReactor Installer in the web descriptor.

```
<?xml version="1.0" encoding="UTF-8"?>
<web-app>
  <!--FusionReactor filter - created Thu Jan 10 17:50:57 CET 2008-->
  <filter>
    <filter-name>FusionReactor</filter-name>
    <filter-
class>com.intergral.fusionreactor.filter.FusionReactorFilter</filter-
class>
    <init-param>
      <param-name>config</param-name>
      <param-
value>C:/FusionReactor/instance/coldfusion.cfm8.int000b/conf/reactor.c
onf</param-value>
    </init-param>
  </filter>
  <filter-mapping>
    <filter-name>FusionReactor</filter-name>
    <url-pattern>/*</url-pattern>
  </filter-mapping>
  .
  .
  .
</web-app>
```

Examples of where the default-web.xml file is typically located :

ColdFusion 6 (Windows)	/CFusionMX/runtime/servers/coldfusion/SERVER-INF/default-web.xml
ColdFusion 7 (Windows)	/CFusionMX7/runtime/servers/coldfusion/SERVER-INF/default-web.xml
ColdFusion 8 (Windows)	/ColdFusion8/runtime/servers/coldfusion/SERVER-INF/default-web.xml

NOTE: If FusionReactor stops working you should check that the web descriptor contains correct settings similar to the ones above.

NOTE: On some J2EE servers you may have to look in the web.xml and not the default-web.xml file. See the Installation Guide for further details..