Home

Introducing FusionAnalytics

Out of the box, FusionAnalytics is a comprehensive analytics tool covering a wide spectrum of analytical requirements. It can provide you with everything from high-level monthly trend reporting, all the way down to monitoring the status of individual database requests on a millisecond to millisecond basis. It is designed to be an invaluable tool for everyone, from the manager trying to estimate future hosting costs and requirements, right down to the developer trying to optimize the memory footprint of a troublesome application page.

At it's core though, FusionAnalytics is actually a powerful platform capable of running any number of Analytics Apps (each of which could be tailored to a specific purpose and can draw data from any external application or server capable of generating log files). At launch time we currently only support the FusionAnalytics for FusionReactor App, but we will be making other Analytics Apps available over time.

Additionally, you can also customize the standard App or even register for a free Developer License and build your own FusionAnaytics App from scratch. Information on customizing or building new Apps will be coming soon.

If you are new to FusionAnalytics, or don't have any plans to create or modify FusionAnalytics Apps then the User Guide is a good place to begin (once you have everything installed of course). Even if you do plan on building your own Apps in the future, the standard App will give you a good idea of what is possible with the platform.

For more information about FusionAnalytics please see: http://www.fusion-analytics.com

Welcome

This is the documentation hub for FusionAnalytics. You can use the links on the left to navigate through the collection.

Feedback

We welcome feedback on all our products and publications. To discover the various ways of contacting us, please use our contact page. We will address your feedback as quickly as possible.

License and Copyright Information

FusionAnalytics End User License Agreement for FusionAnalytics

Third Party Software Notices and/or Additional Terms and Conditions for FusionAnalytics are made a part of and incorporated by reference into such product's End User License Agreement

Copyright statement for all documentation contained in this collection.

Back to the top

Quick Start Guide

FusionAnalytics Quick Start Guide

Overview

Welcome to the Quick Start Guide for FusionAnalytics. Follow our 10 step tutorial to get FusionAnalytics up and running as soon as possible.

Part 1: Installation

Step 1: Database Configuration
Follow our 3 simple tutorials to get your database ready for FusionAnalytics:

- Configure Network Settings of MSSQL Server
- Create a New Database Login
- Create a New Database

Step 2: Installing FusionAnalytics Server
This guide can assist you with your first installation of FusionAnalytics.

Step 3: Installing the FusionAnalytics AIR Client
Make sure you have the FusionAnalytics AIR Client ready, so that you can view your data.

Part 2: Configuration

Before you start...

- Have you installed both components of FusionAnalytics Server?
- Have you installed the FusionAnalytics Client on machines that you want to monitor data from?
- Do you have FusionReactor installed?
- Is your database available and ready for data to be imported to it?
- Do you have the password for the FusionAnalytics ***'admin'*** account ready?
- Is JavaScript enabled in your web browser?

Step 4: Accessing FusionAnalytics
How to get to the different parts of FusionAnalytics.

Step 5: Setting Up the AIR Client
Set up your AIR Client so that you can view your data once it has imported.

Step 6: Configuring FusionReactor for FusionAnalytics
Making sure your log files get sent to the right place.

Step 7: Setting Up Users
How to create and edit different accounts.

Step 8: Setting Up the Email Server
Receiving important notifications and reports from FusionAnalytics.

Step 9: Setting Up Application Filters
How to make the most of your application.

Step 10: Setting Up TAP and Daily Status Reports
Configuring the reporting feature of FusionAnalytics.

Next Steps

Database Configuration

- Configure Network Settings
- Create a New Database Login
- Create a New Database

Accessing FusionAnalytics

Setting Up the AIR Client

Configuring FusionReactor for FusionAnalytics

Setting Up Users

Setting Up the Email Server

Setting Up Application Filters

Setting Up TAP and Daily Status Reports

Database Configuration

Database Configuration

Overview

This section will help you prepare your database ready for use with FusionAnalytics.

Visit Microsoft's SQL Server website for more information on the different versions and details on how to buy. You can also use the **free Express Edition**, see below for more information.

> ⚠ Please note that version 1.0 of FusionAnalytics currently **only** supports Microsoft SQL Server 2005, 2008 and 2008 R2.

Microsoft SQL Server Express

The Express Edition is a **free** version of Microsoft's SQL Server. It still uses the same core database engine, but on a smaller scale.

There are no limitations on the number of databases or users that can be set up, but there are constraints on the size of each individual database, and the overall performance of the system.

Constraint	Maximum Value
Database size	4GB for Express 2005 and 2008. 10GB for Express 2008 R2
Memory used	1 GB
CPUs used	1

⚠ *The maximum database size applies per database, excluding log files.*

> ℹ **Download for free:** Microsoft SQL Server 2008 R2 RTM - Express with Management Tools.

Database size

The following tables should give you a rough idea of how big your database will be. This information is for Microsoft SQL Server in simple recovery mode (if you follow the Quick Start Guide you will configure your database in this way).

- High resolution data is all information stored in the database from every sample taken by FusionAnalytics.
- Low resolution data is data that has been quantized to reduce the database size.

See Scaling and Quantization for more information on data quantization or Application Configuration (FADC) to find out how to control the length of time this information is stored.

Data estimates when keeping 100% of data (high and low res)				
Average requests per day	Estimated size per day	Estimated size per week	Estimated size per month	Estimated size per year
1,000	15 MB	100 MB	500 MB	5 GB
10,000	70 MB	500 MB	2 GB	25 GB
100,000	0.5 GB	5 GB	20 GB	250 GB
250,000	2 GB	15 GB	60 GB	750 GB
500,000	4 GB	30 GB	120 GB	1.5 TB
1,000,000	8 GB	60 GB	250 GB	3 TB
5,000,000	40 GB	300 GB	1.2 TB	15 TB
10,000,000	80 GB	600 GB	2.5 TB	30 TB
Of the above information, 5% comes from FusionReactor log archives and 95% from the FusionAnalytics database.				

Data estimates when keeping high res data for 1 month and low res data for 1 year				
Average requests per day	Estimated size after 1 day	Estimated size after 1 week	Estimated size after 1 month	Estimated size after 1 year
1,000	15 MB	100 MB	500 MB	3.5 GB
10,000	70 MB	500 MB	2 GB	5.5 GB
100,000	0.5 GB	5 GB	20 GB	25 GB
250,000	2 GB	15 GB	60 GB	65 GB
500,000	4 GB	30 GB	120 GB	125 GB
1,000,000	8 GB	60 GB	250 GB	255 GB
5,000,000	40 GB	300 GB	1.2 TB	1.2 TB
10,000,000	80 GB	600 GB	2.5 TB	2.5 TB

Data estimates when keeping high res data for 1 week and low res data for 1 year				
Average requests per day	Estimated size after 1 day	Estimated size after 1 week	Estimated size after 1 month	Estimated size after 1 year
1,000	15 MB	100 MB	350 MB	3.5 GB
10,000	70 MB	500 MB	750 MB	3.6 GB
100,000	0.5 GB	5 GB	5.5 GB	8.5 GB
250,000	2 GB	15 GB	15.5 GB	20 GB
500,000	4 GB	30 GB	30.5 GB	35 GB
1,000,000	8 GB	60 GB	60.5 GB	65 GB
5,000,000	40 GB	300 GB	300 GB	305 GB
10,000,000	80 GB	600 GB	600 GB	605 GB

⚠ The above figures are only an estimate based on test data. Using longer URLs, more parameters, more database requests and other details will change the results. The average URL used in the above data was 60 characters long from http:// to the file extension with an average of 1.3 JDBC requests per standard request.

Next Steps

Prepare and Configure your database for FusionAnalytics:

Configure Network Settings

Create a New Database Login

Create a New Database

Once your database is ready to use, visit the installation guide to get started.

Configure Network Settings

Configure Network Settings

FusionAnalytics uses JDBC to connect to the Microsoft® SQL Server database. To allow FusionAnalytics to communicate to the database the network settings of the database server must be configured appropriately. To configure the network connections of the database server start **SQL Server Configuration Manager**

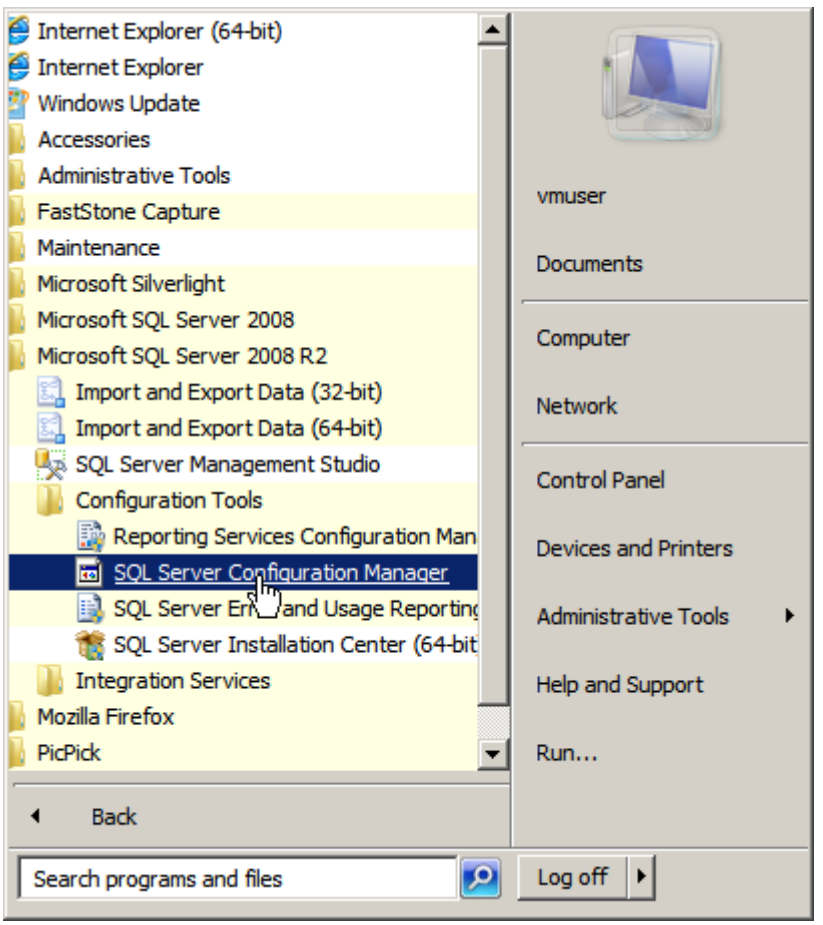

Figure 1: Starting SQL Server Configuration Manager

then select `Protocols for SQLExpress` and double click the line with the `TCP/IP` protocol shown below

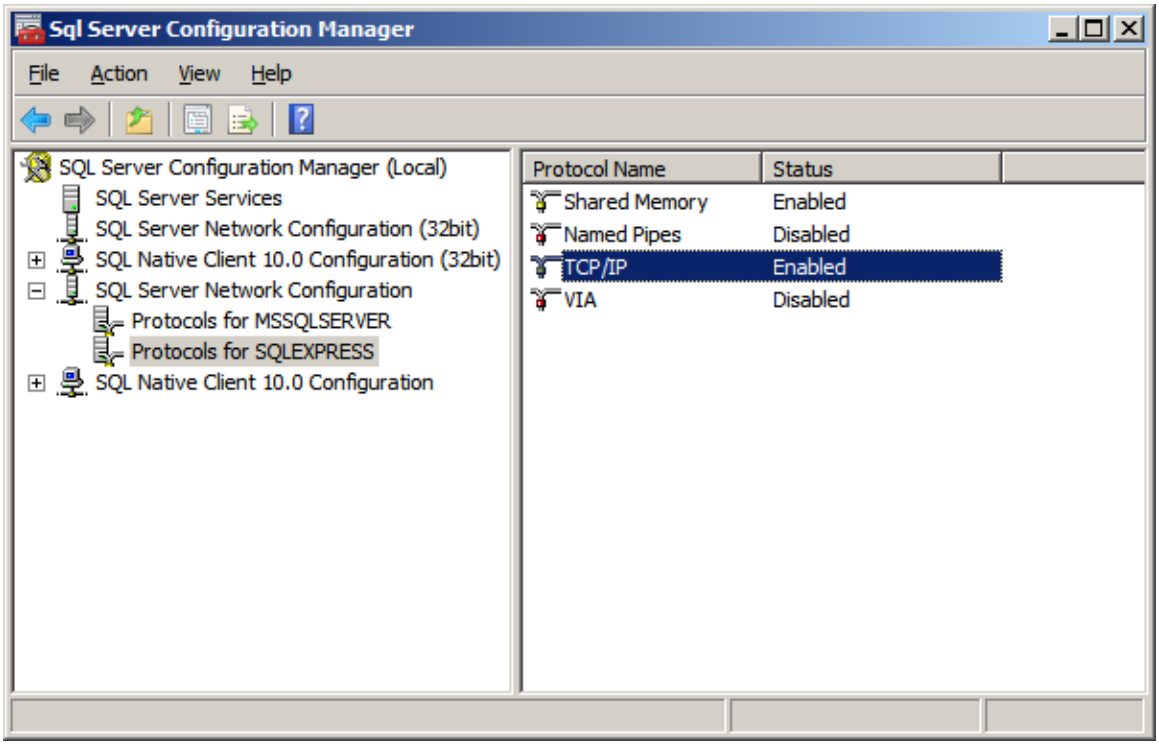

Figure 2: Configuring Protocols

In the **TCP/IP Properties** dialog set `Enabled` to `Yes`

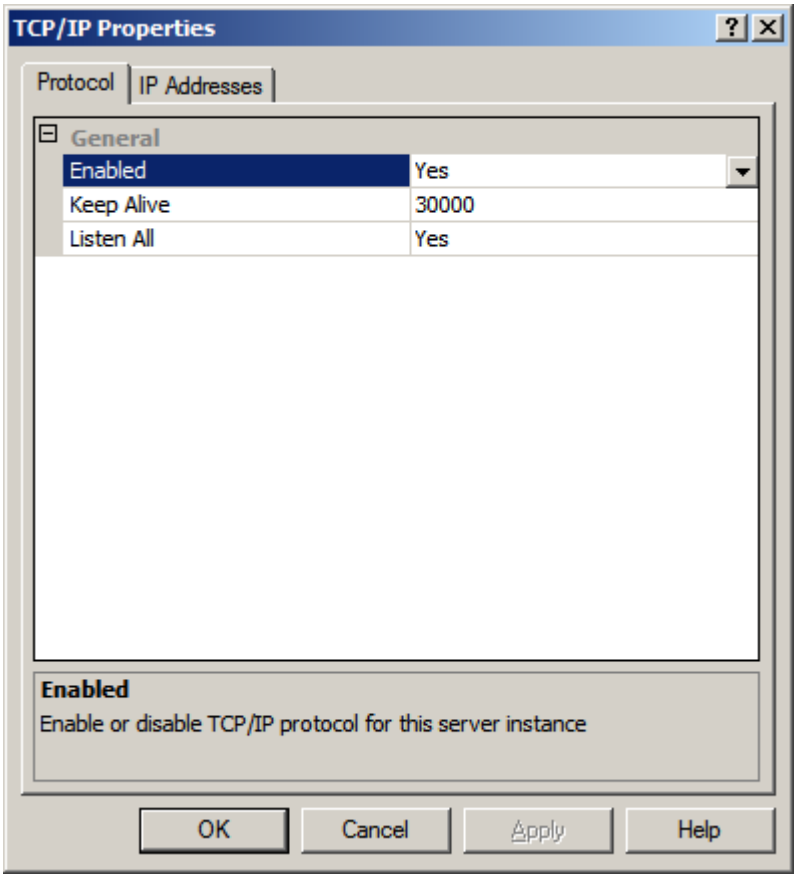

Figure 3: Enabling TCP/IP

then click on the `IP Addresses` tab and at the bottom of the property list enter `1433` as value for `TCP Port`.

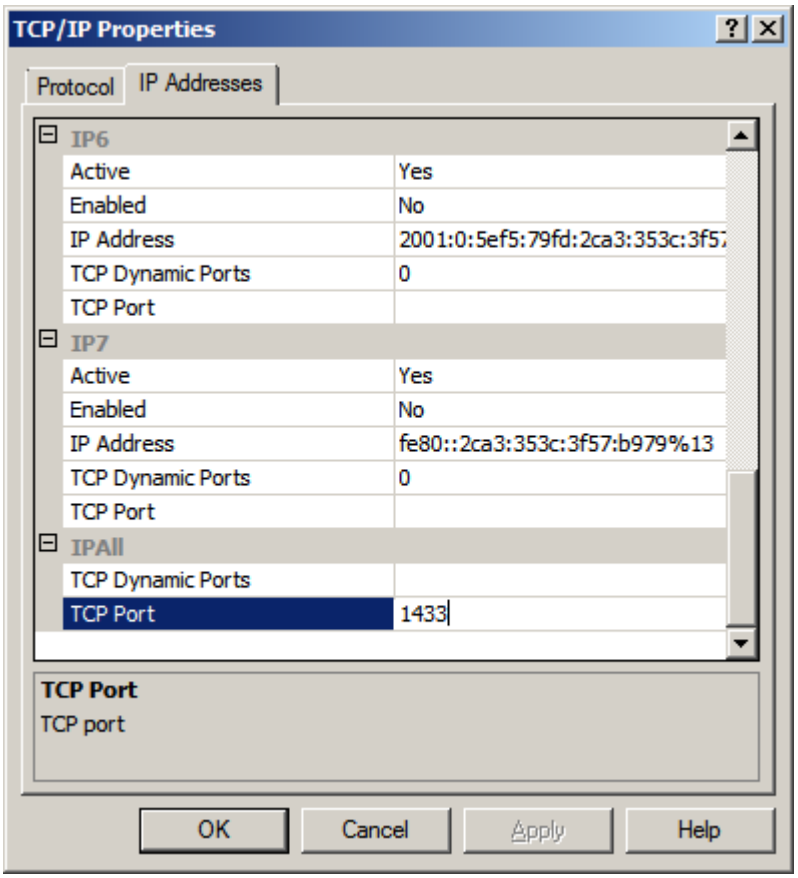

Figure 4: Configuring TCP Port

To apply the configuration changes the SQL Server must be restarted. This can be done from the **SQL Server Configuration Manager** shown below.

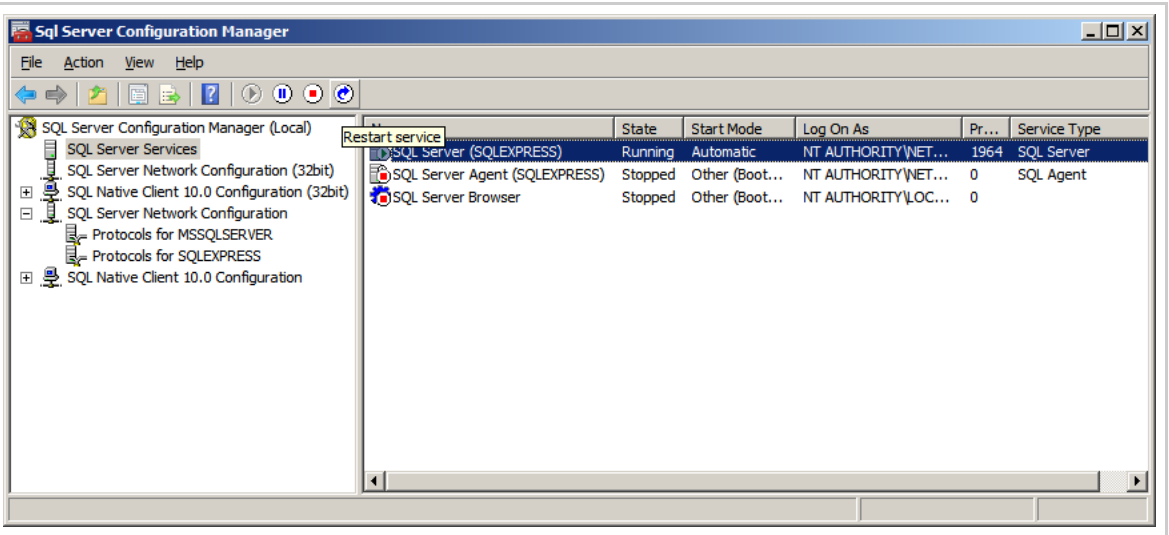

Figure 5: Restarting SQL Server

Back to the top

More Database Configuration Tutorials:

Configure Network Settings

Create a New Database Login

Create a New Database

Next Steps

Create a New Database Login
Next, you will need to create a new database login for FusionAnalytics to use.

Create a New Database Login

Create a New Database Login

1. Open the SQL Server Management Studio
2. In the **Object Explorer** panel (located on the left side), expand the server where you want to create a new login

Enabling SQL Server Authentication

1. Right-click on the server and go to **Properties**
2. In the **Server Properties** pop-up that appears, go to **Security**
3. Make sure that **SQL Server and Windows Authentication mode** is selected and click **OK**

Back to the top

Create a New User

1. Right-click the **Security** folder and go to **New Login...**
2. In the **Login - New** dialogue that appears, enter a new login name
   This tutorial will use ***fadb*** as the new login name
3. Select **SQL Server authentication**
4. Enter an appropriate password in the **Password** and **Confirm Password** fields
5. Uncheck **Enforce password expiration**
6. Make sure **Enforce password policy** remains checked
7. Click **OK** to create the user

Check your new user can log in
In the **Object Explorer** panel, click **Connect Database Engine**. Select the correct **Server Name** and enter the new login and password details. Login should be successful but you should not be able to expand any of the **Database** folders.

Back to the top

More Database Configuration Tutorials:

Configure Network Settings

Create a New Database Login

Create a New Database

Next Steps

Create a New Database
Next, you will need to create a new database for FusionAnalytics to use.

Create a New Database

Create a New Database

1. Open the SQL Server Management Studio
2. In the **Object Explorer** panel (located on the left side), expand the server where you want to create a new database
3. Right-click the **Databases** folder and go to **New Database...**

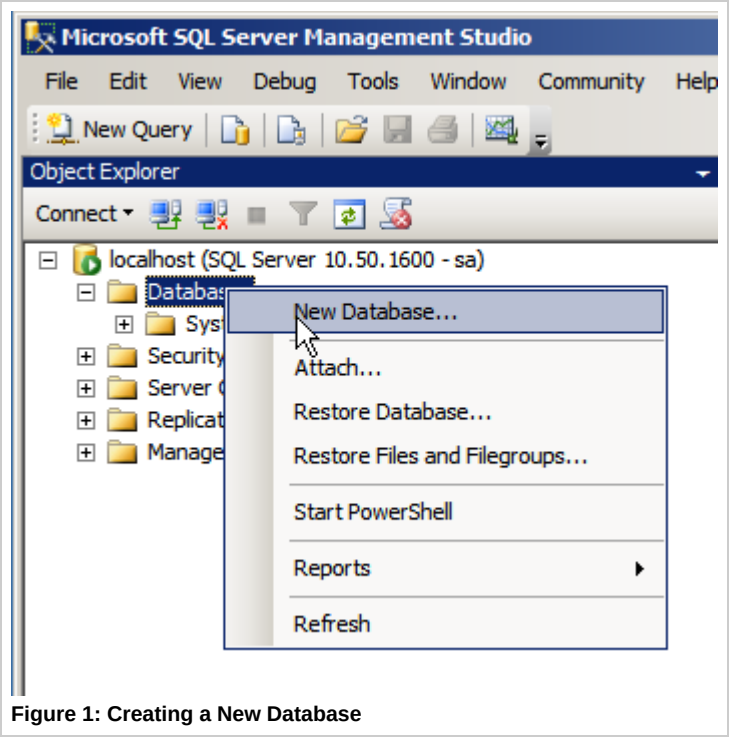

Figure 1: Creating a New Database

4. In the **New Database** pop-up that appears, enter a new database name and owner

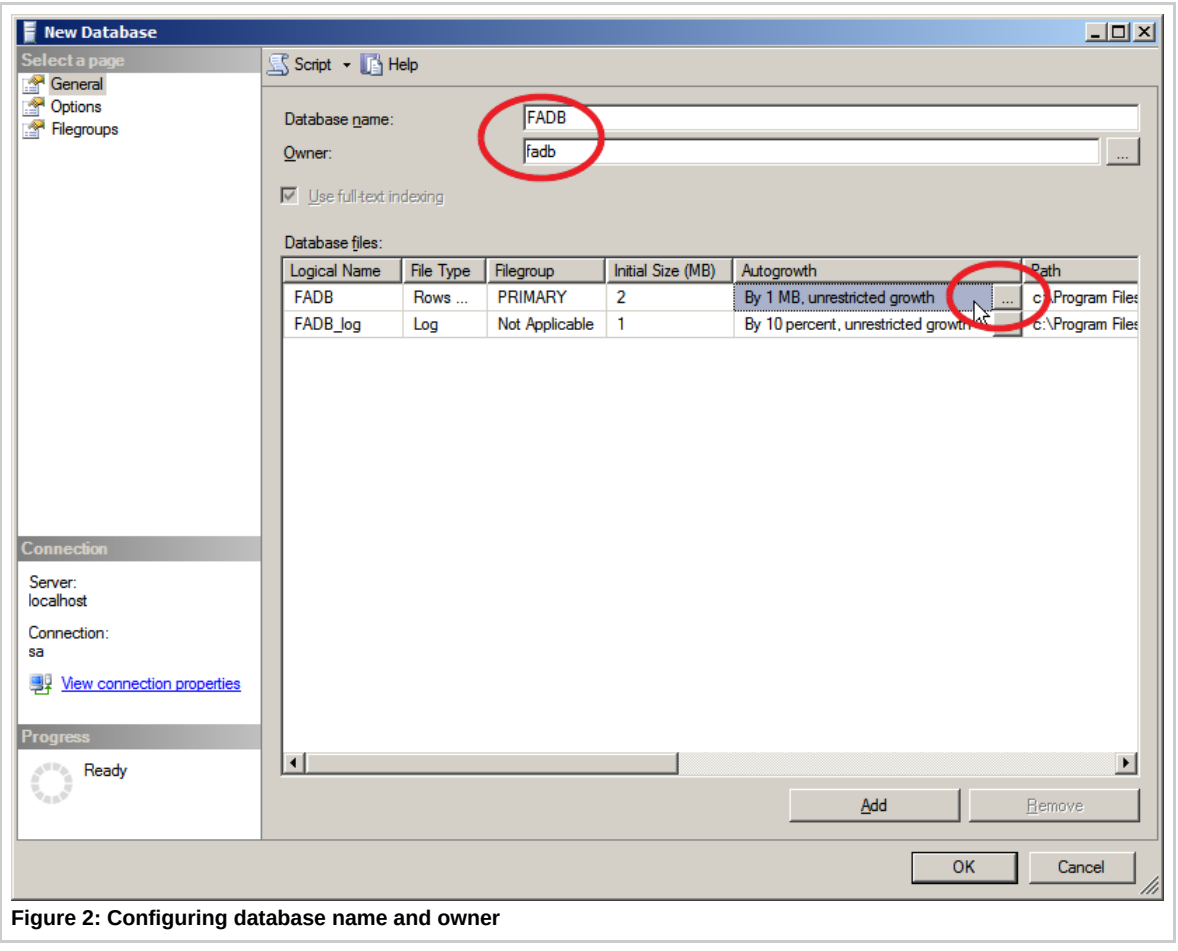

Figure 2: Configuring database name and owner

> This tutorial will use fadb as the new database name

5. Next click on the little button with the three dots underneath the Autogrowth column in the first line. In the Change Autogrowth for FADB dialog change the setting for Autogrowth from its default value In Megabytes to In Percent as shown below.

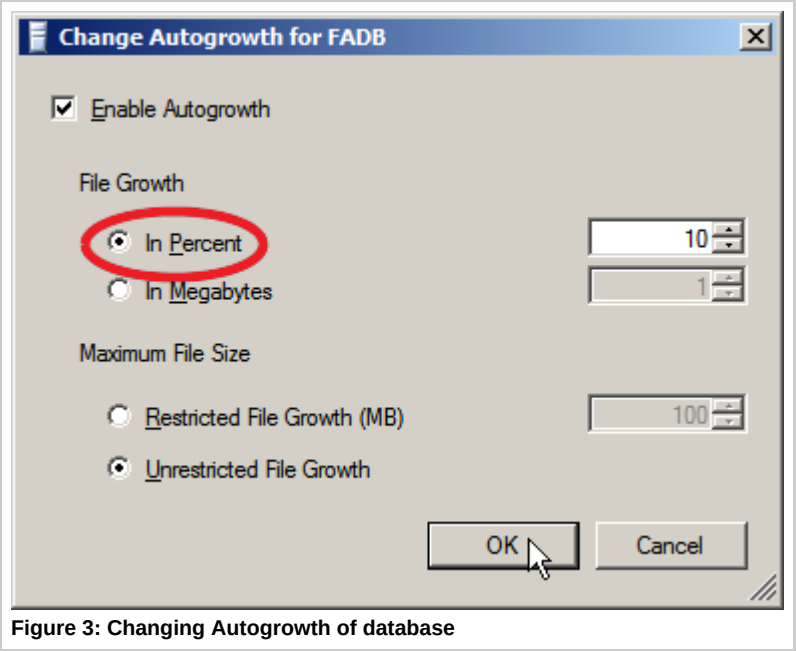

Figure 3: Changing Autogrowth of database

6. Back in the New Database dialog change the Recovery Model of the database from Full to Simple mode

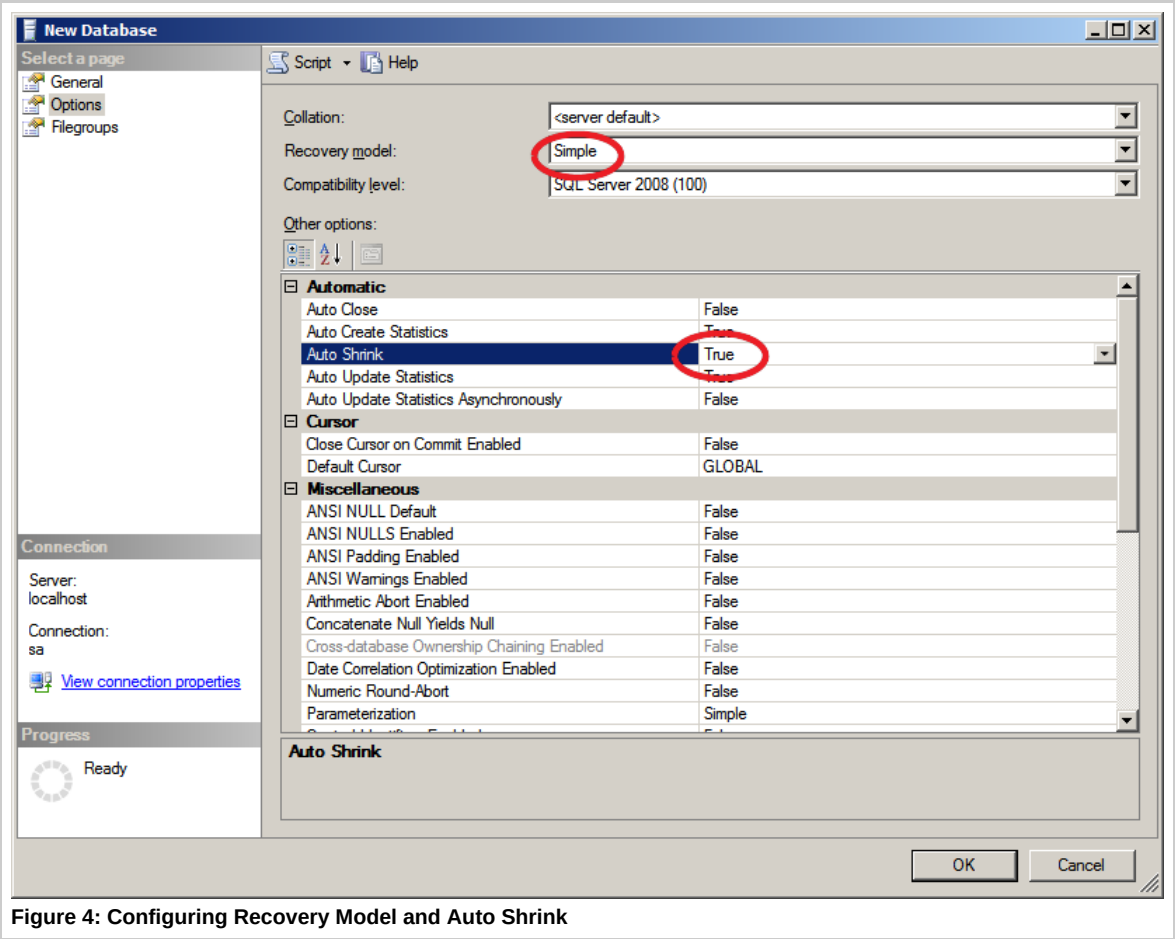

Figure 4: Configuring Recovery Model and Auto Shrink

and enable the database to automatically shrink its size by enabling the Auto Shrink property.

> To prevent the database log files from taking up too much disk space, the **Recovery Model** is set to **Simple**.

7. Click **OK** to create the new database

> **Check your new database has been created**
> The new database should appear when you expand the **Databases** folder.

Server Memory
If you do not limit the amount of memory the database can use it will be consumed very quickly.

8. In Server Management Studio right click on your server name and click properties.

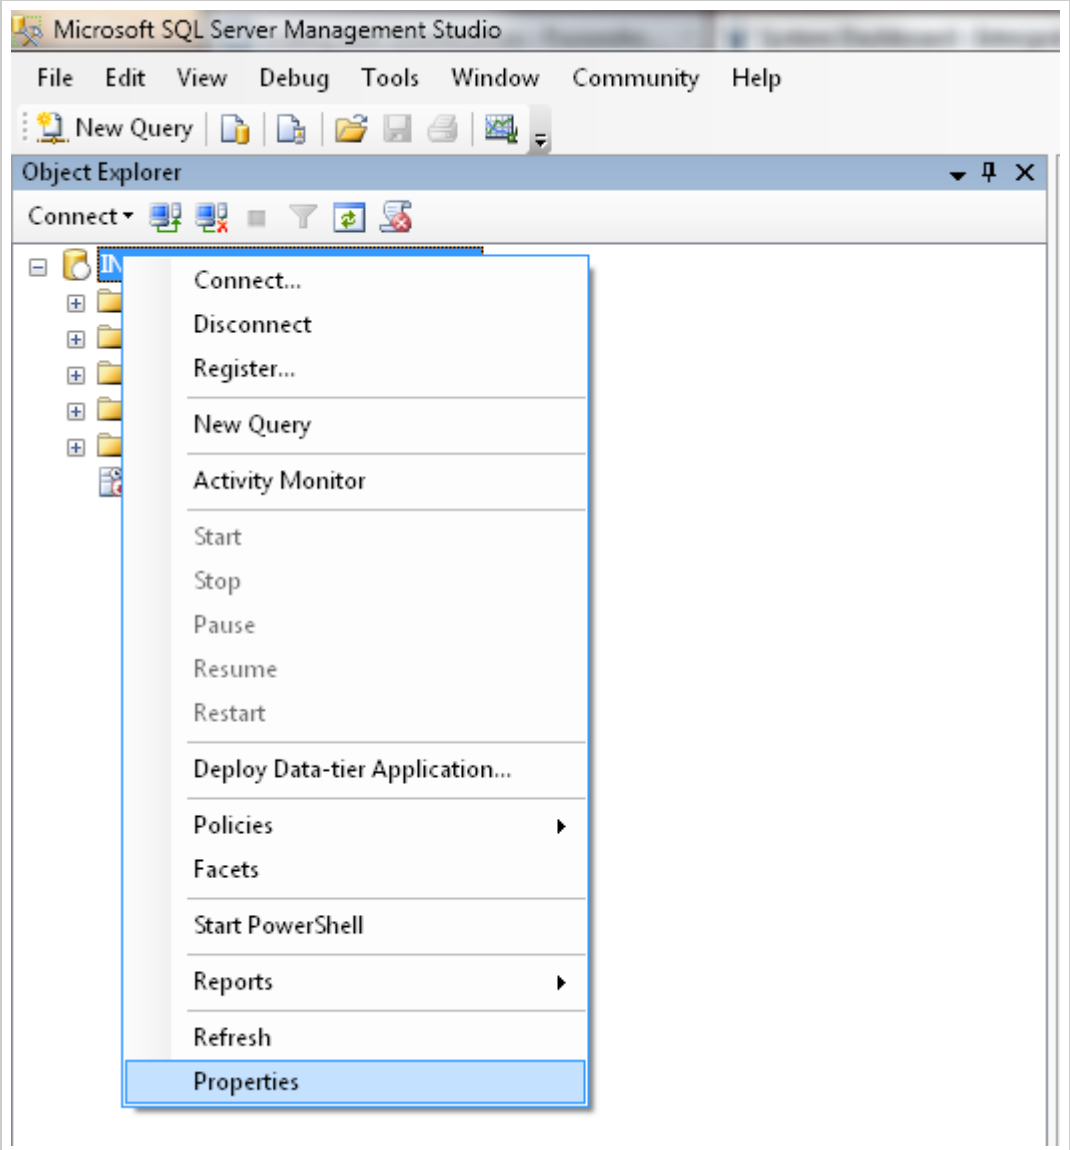

Figure 5: Server properties

9. In the properties dialog click Memory. Set the Maximum Server Memory here.

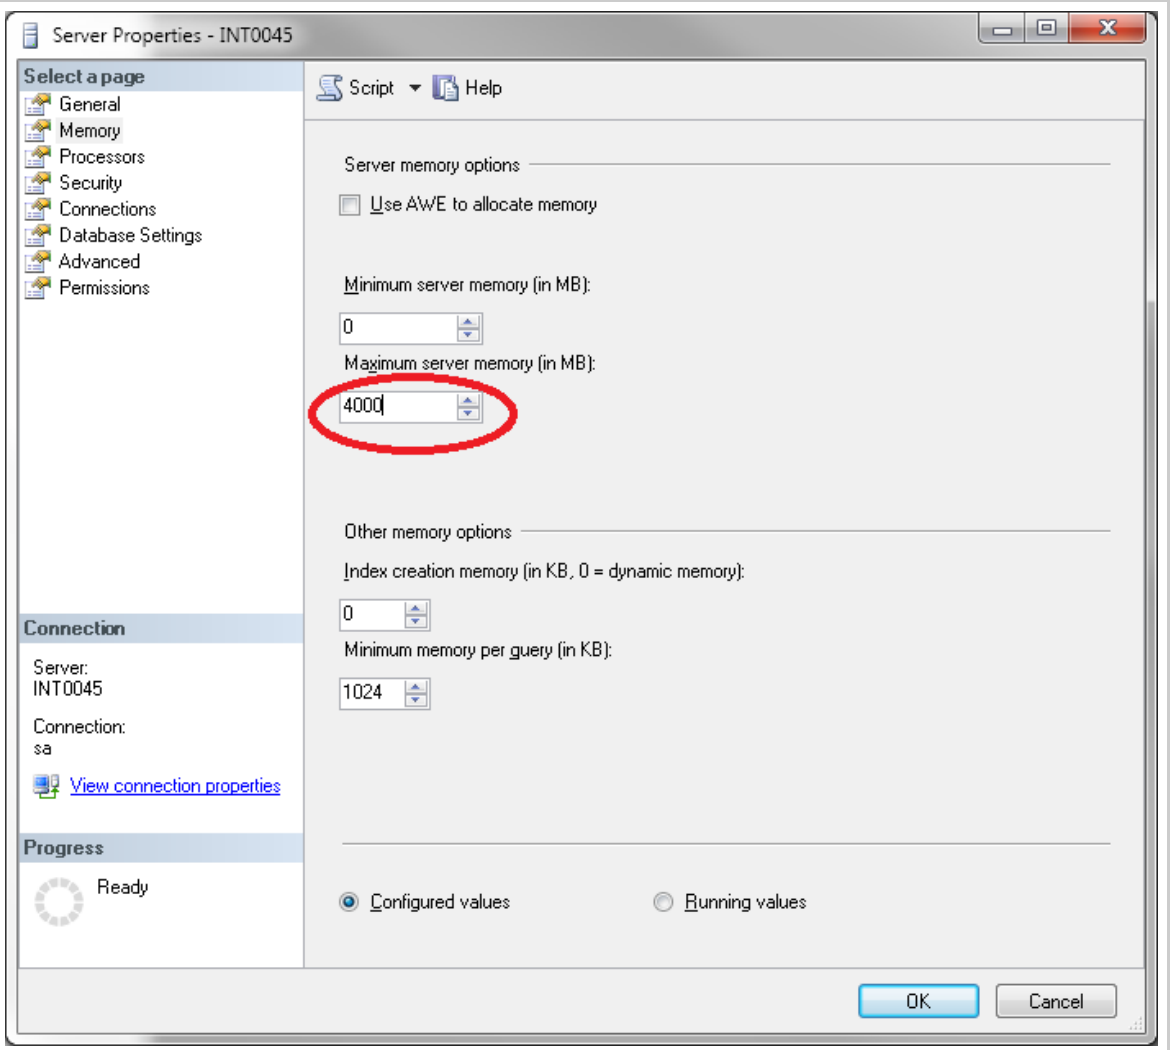

Figure 6: Allocating memory

We recommend allocating **between 1 third and 1 half** of available memory to the database

Back to the top

More Database Configuration Tutorials:

Configure Network Settings

Create a New Database Login

Create a New Database

Next Steps

Installing FusionAnalytics Server
Visit the installation guide when you are ready to get started with FusionAnalytics.

Accessing FusionAnalytics

Accessing FusionAnalytics

There are three different parts of FusionAnalytics that you can access using your **'admin'** account. These are DataCollector, DataServices and the Client. For more information on what each component does, please see the FusionAnalytics User's Guide.

FusionAnalytics Server

The following URLs can be used for direct access to each component. The FusionAnalytics Dashboard provides links to the DataCollector and DataServices components.

Component	URL
Dashboard	http://localhost:8400
FusionAnalytics DataCollector (FADC)	http://localhost:8400/fadc
FusionAnalytics DataServices (FADS)	http://localhost:8400/fads

> ⚠ You can change the *'localhost'* part of the URL to the machine name that you want to access (if it is not the local machine).
>
> If you changed the HTTP port number from the default (8400), then you will need to use that instead.

Back to the top

FusionAnalytics Client

To access FusionAnalytics, you can use either the AIR client or a web client.

Component	URL
FusionAnalytics Client	http://localhost:8400/fads/myFADSApplicationName ℹ You will need to replace **'myFADSApplicationName'** with your actual FADS application name. If you followed the installation tutorial this will be ***coldfusion9***.

To access the AIR Client:

You can double-click the FusionAnalyticsAir icon on your desktop.

The next step of this tutorial will help you with Setting Up the AIR Client.

To access the Web Client:

You can login to FADS and click the Launch button on the application you want to view (Figure 1).

You can also enter the above URL into a web browser's address bar.

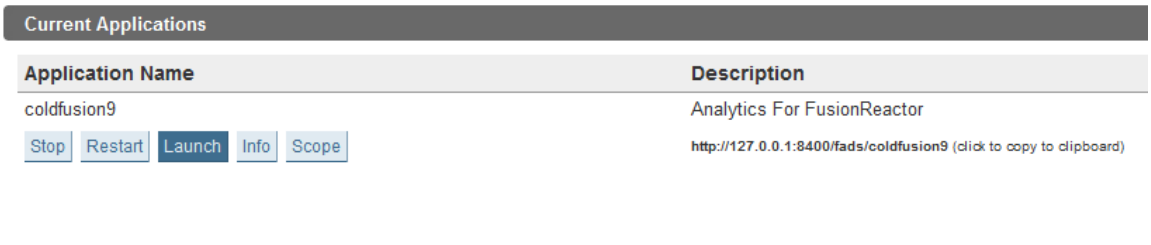

Figure 1: Launch Application

Logging In

You can login to FADC, FADS and the Client using:

Username:	admin
Password:	*The password you defined in the **Administrator Password** screen during installation.*

We will set up more users later in the configuration tutorial.

Back to the top

More Configuration Tutorials:

Accessing FusionAnalytics

Setting Up the AIR Client

Configuring FusionReactor for FusionAnalytics

Setting Up Users

Setting Up the Email Server

Setting Up Application Filters

Setting Up TAP and Daily Status Reports

Next Steps

Setting Up the AIR Client
Set up your AIR Client so that you can view your data once it has imported.

Setting Up the AIR Client

Setting Up the AIR Client

AIR Client Configuration

The FusionAnalytics AIR Client uses URLs to identify and connect to the server.

They follow the same pattern: **http://127.0.0.1:8400/fads/myFADSApplicationName**

This is broken down as follows:

URL Element	Description
http://	**HTTP is the default.** See the tutorial on Setting Up FusionAnalytics with HTTPS for secure connections.
127.0.0.1	**The machine where FusionAnalytics Server is installed.** This example uses the local machine, so the address is 127.0.0.1.
8400	**The machine port.** If you changed the default port during the installation, use that instead; if not, use the default (8400).
/fads	**The fixed path to DataServices.**
/myFADSApplicationName	**The name of the DataServices application you want to connect to.** If you used the suggested name from the installation guide, this will be ***coldfusion9.***

Enter this into the URL box at the top and click **Add Application**. A shortcut will be created (Figure 1), which you can double-click to go to the login page for that application.

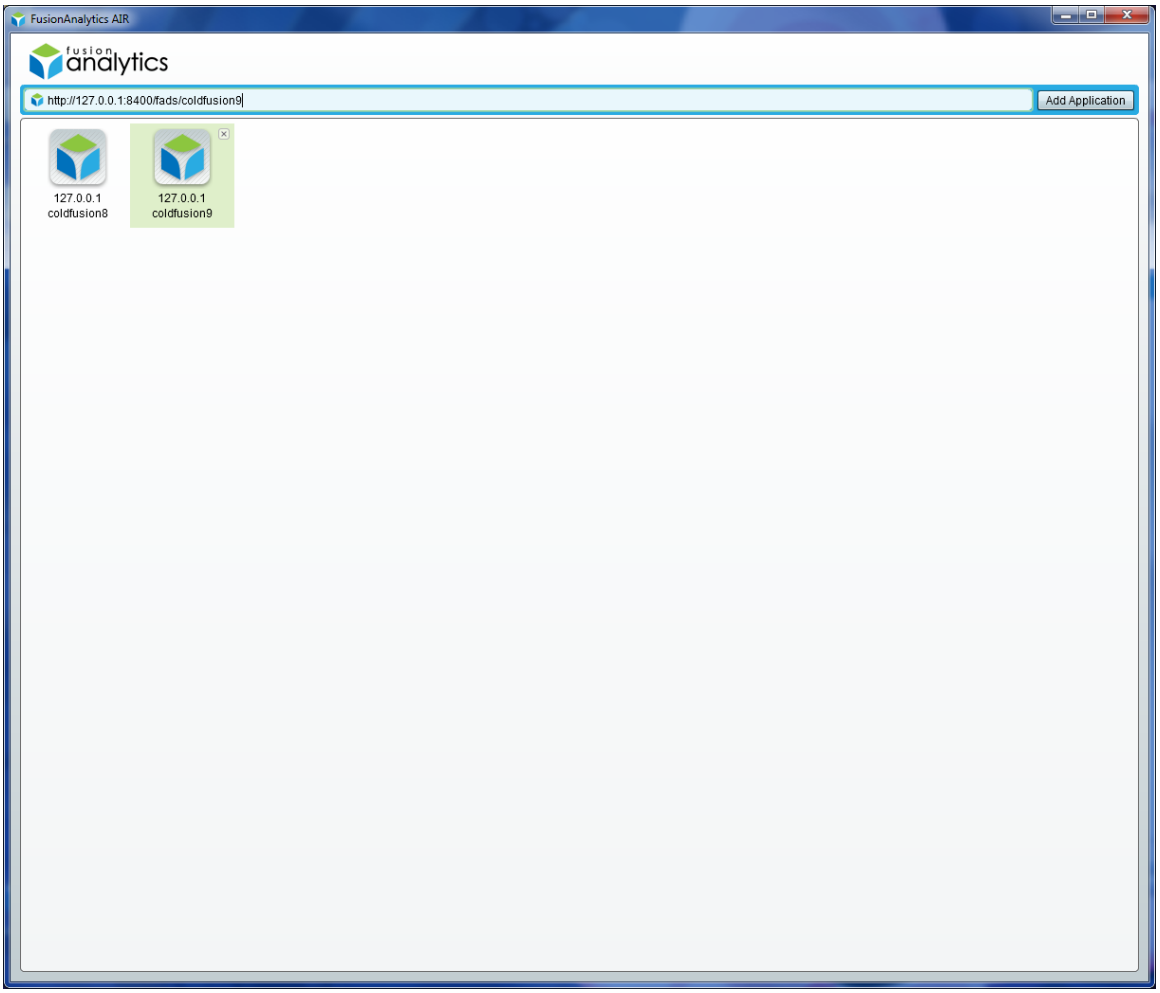

Figure 1: AIR Client Shortcuts

At the login page, you should then enter your username and password (Figure 2).

To begin with, you can use the ***'admin'*** username and password that you set up during the installation of FusionAnalytics Server.

We will set up more users later in the Setting Up Users tutorial.

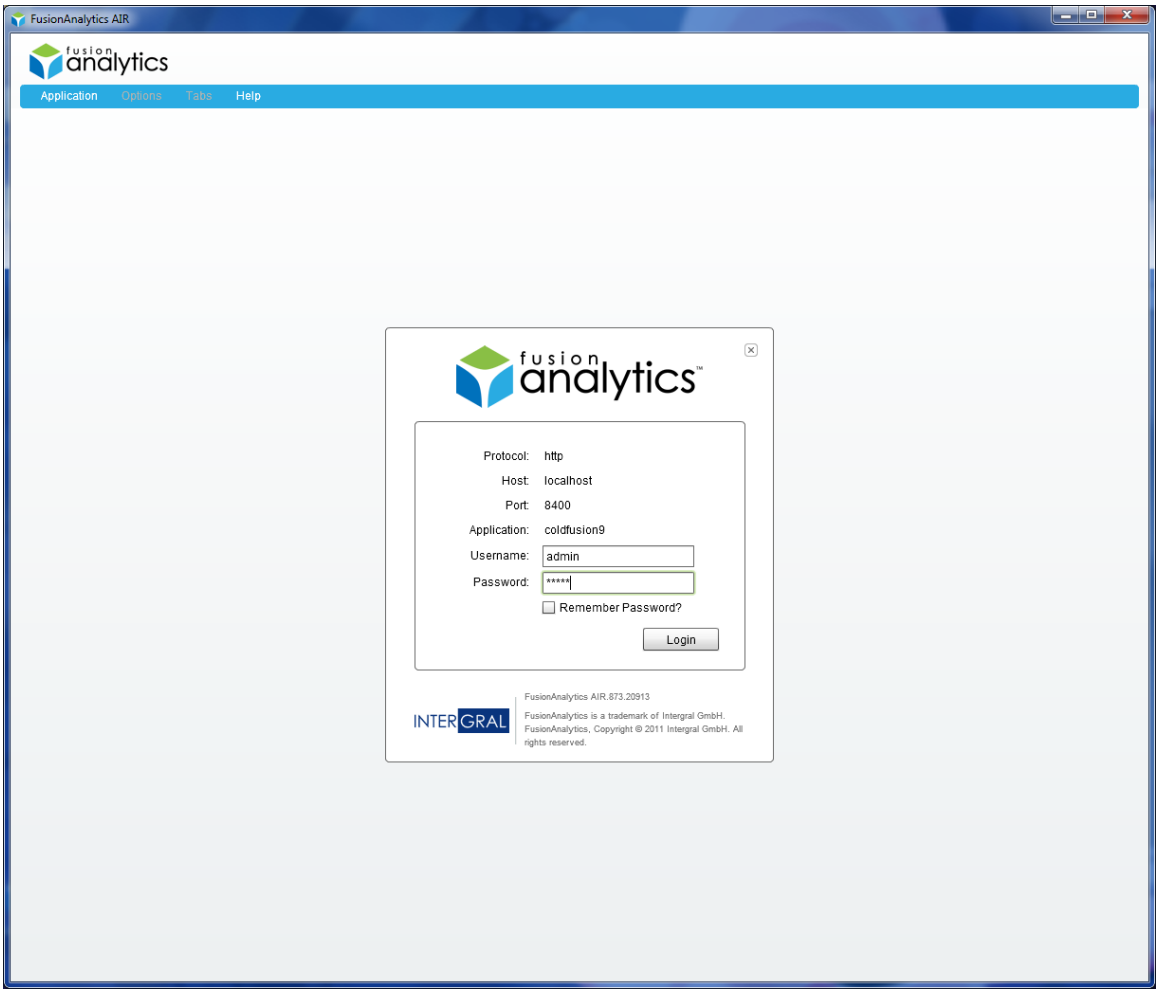

Figure 2: Login Screen

Back to the top

An Alternative to the AIR Client

If you do not have an AIR client available, you can also access the client in your web browser.

The same URL can be entered into the address bar **or** you can log into FusionAnalytics DataServices and click **Launch** on the application that you want to view (Figure 3).

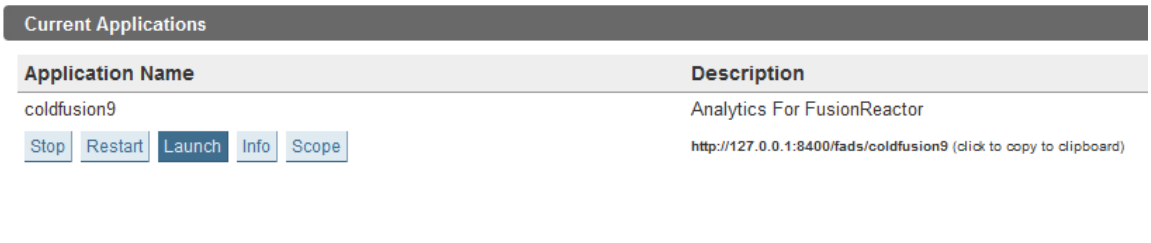

Figure 3: FADS Launch Button

Back to the top

On This Page:

More Configuration Tutorials:

Accessing FusionAnalytics

Setting Up the AIR Client

Configuring FusionReactor for FusionAnalytics

Setting Up Users

Setting Up the Email Server

Setting Up Application Filters

Setting Up TAP and Daily Status Reports

Next Steps

Configuring FusionReactor for FusionAnalytics
Set up FusionReactor so that data from your server is sent to FusionAnalytics for you to analyze.

Configuring FusionReactor for FusionAnalytics

Configuring FusionReactor for FusionAnalytics

To begin, you will need to login to the FusionReactor Administration Manager (FRAM) and go to **Instances Instance Manager**.

Find the instance that you want to use to send data to FusionAnalytics and click the *FusionAnalytics Connector Settings* button (Figure 1).

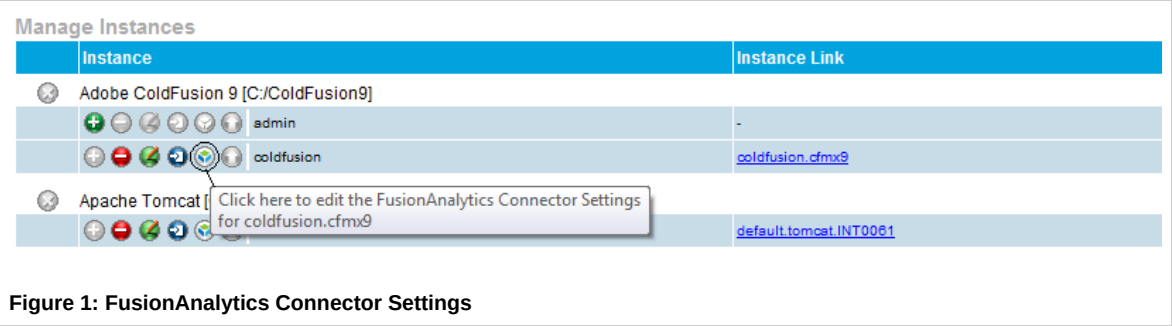

Figure 1: FusionAnalytics Connector Settings

This will display the **FusionAnalytics Connector Manager** pop-up (Figure 2). Click **Next** to begin.

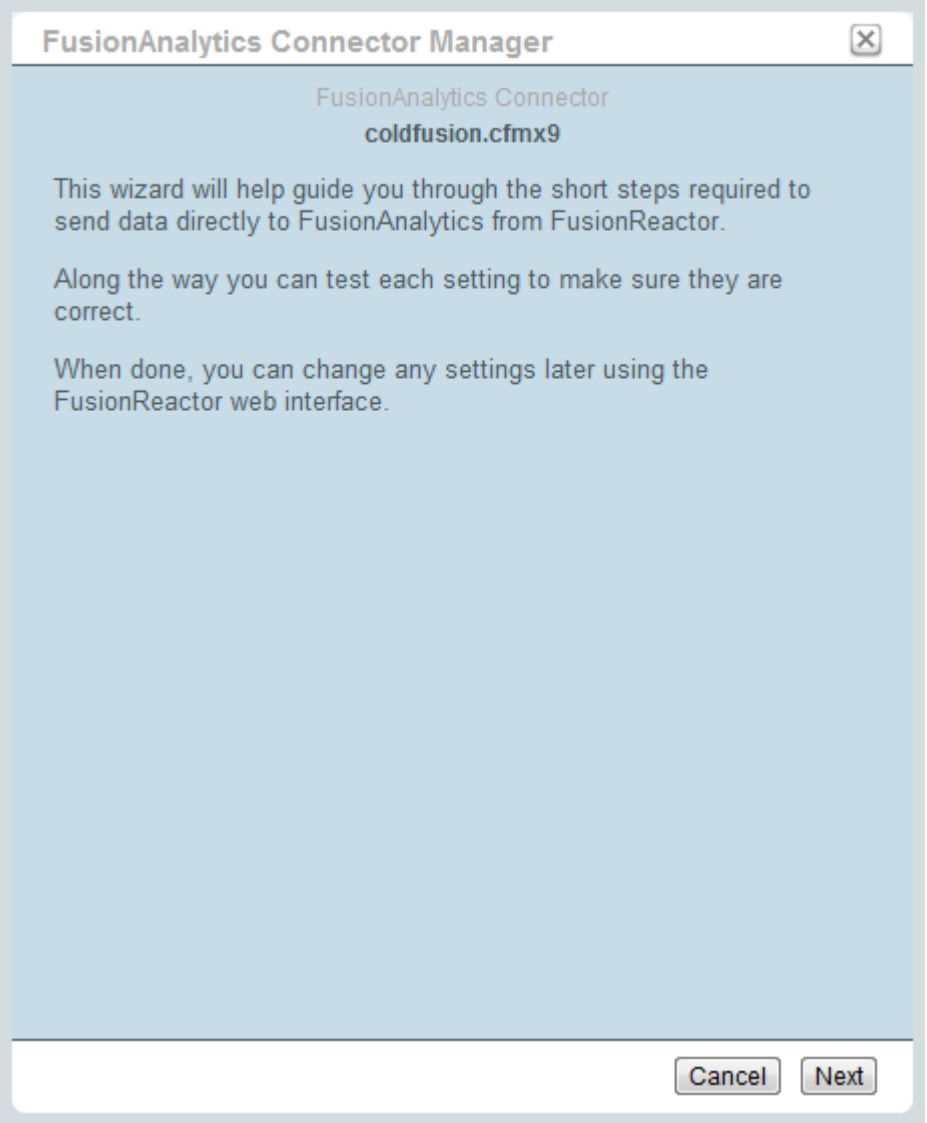

Figure 2: FusionAnalytics Connector Manager

The first step is to configure the **Log Settings** (Figure 3). This sets how often log files will be transferred to FusionAnalytics.

You can change these settings later in the instance under **Logging Log Settings**.

> *This tutorial will rotate the logs to FusionAnalytics once an hour, as shown in Figure 3.*

You can also choose to lock these settings, so that other users cannot change them. See Locked in the FusionReactor User's Guide for more information.

When you have finished, click **Save Log Settings**.

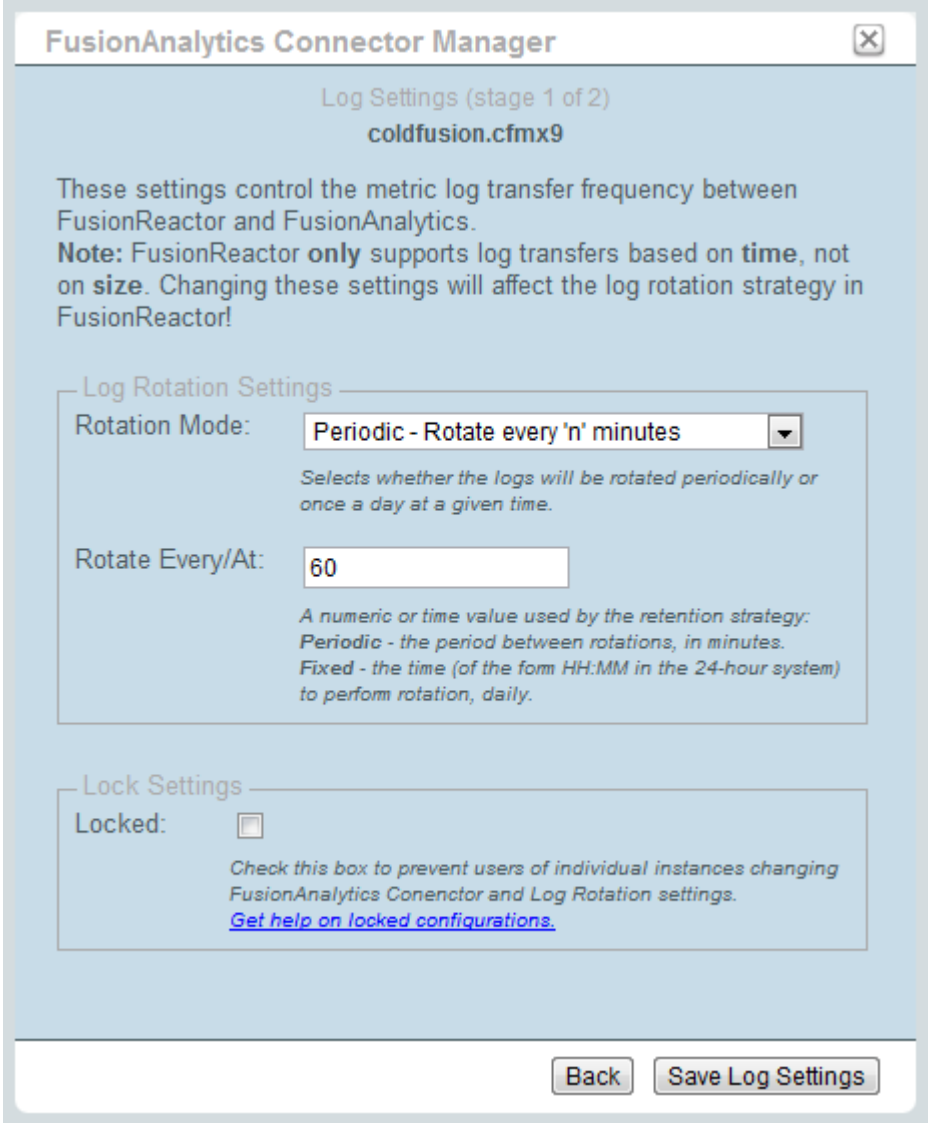

Figure 3: Log Settings

The next page takes you to the **FusionAnalytics Settings** (Figure 4).

> *For this tutorial we will keep the Connector Mode and FRAM Service as the default options. We will only add a FusionAnalytics Target.*

Under **FusionAnalytics Targets**, click the  icon to add a new target. This specifies where the log files from FusionReactor will be sent.

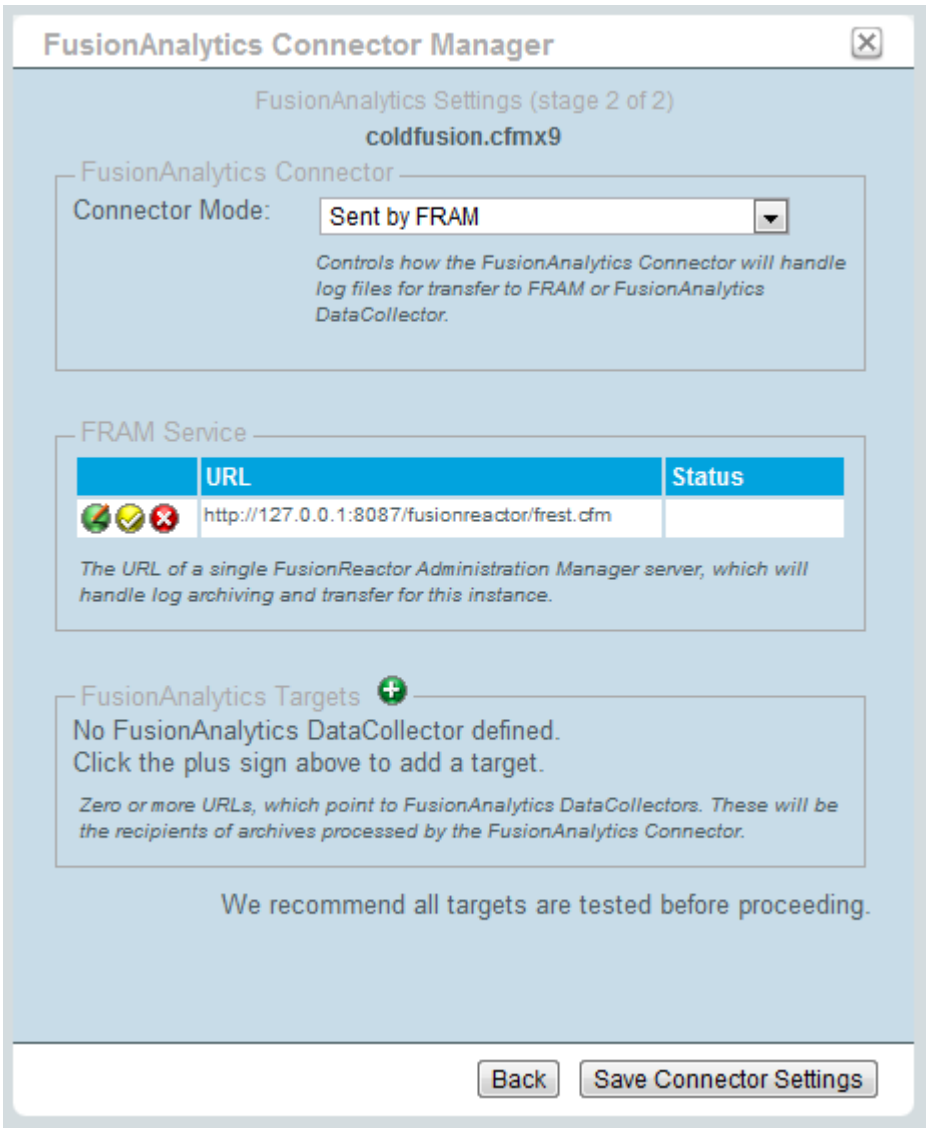

Figure 4: FusionAnalytics Settings

In the screen that appears (Figure 5), fill in the following:

Field	Description
Host	**The machine where FusionAnalytics is installed.** This example uses the local machine, so the IP address is 127.0.0.1 or localhost.
Port	**The machine port.** If you changed the default port during the installation, use that instead; if not, use the default (8400).
Application Name	**The name of the DataCollector Application that you want to collect the data.** If you followed the example in the installation guide, this will be ***coldfusion9***.
DataCollector	**The name of the actual DataCollector that you want to receive the data.** Each DataCollector application can have multiple DataCollectors that receive data. The default setting is one DataCollector named ***faDataCollector-1***.
Username	**The username for a user who has access to the DataCollector application.** This is only used if the DataCollector requires authentication. If you are using the default setting, this is not needed.
Password	**The password for the above username.**
HTTPS	**Check this box if your DataCollector uses HTTPS.** For more information on this see Setting Up FusionAnalytics with HTTPS.
Status	Once you have entered in the details, click **Test Target** to check that you are connecting to a valid DataCollector.

If you are following the tutorial, you should fill in the form as shown in Figure 5.

If you want to learn more about DataCollectors, including authentication and multiple DataCollectors within an application, see Applications (FADC) in the FusionAnalytics User's Guide.

Click **Save Target Settings** once you have checked that the status is **OK**.

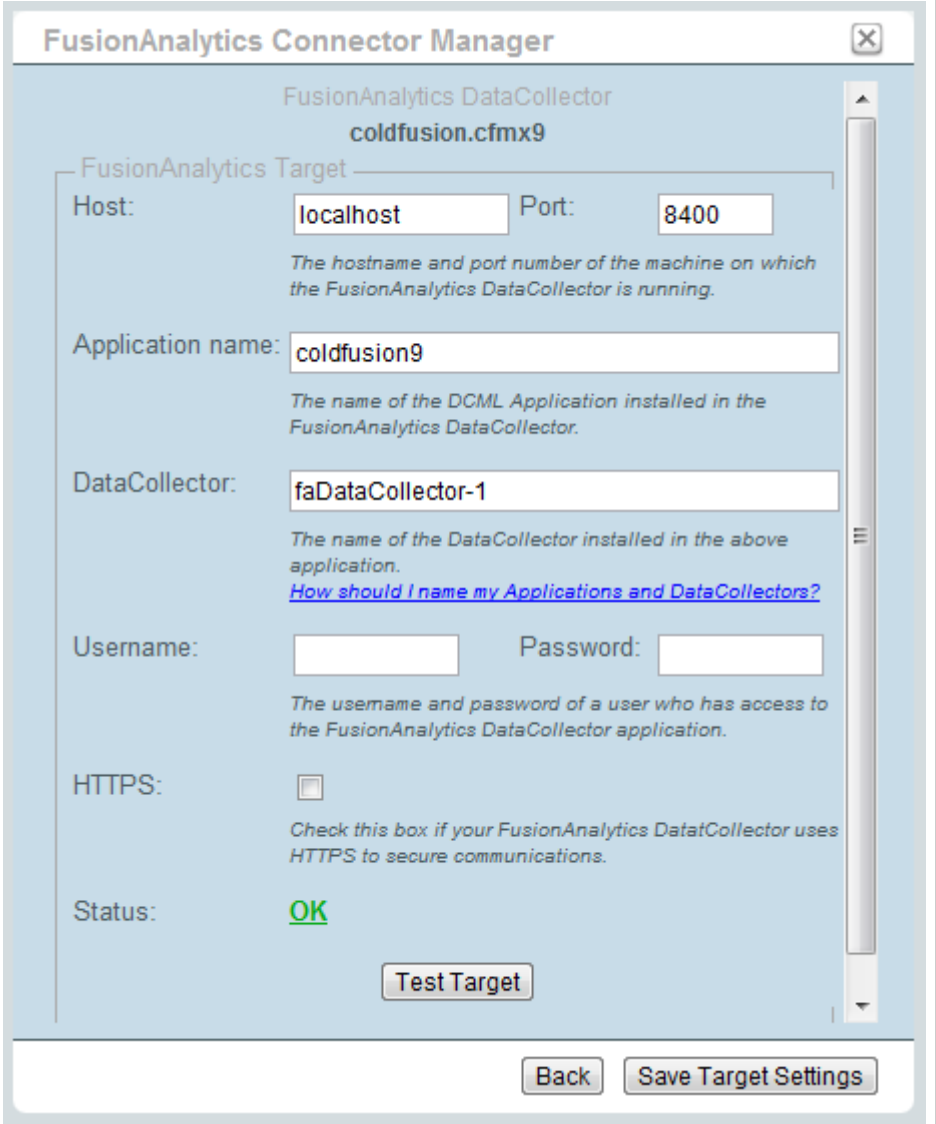

Figure 5: FusionAnalytics Target configuration

You will now return to the **FusionAnalytics Settings** page (Figure 6).

You can check that both the FRAM Service and FusionAnalytics Targets are correct by clicking the yellow **Check Target** button.

If both targets are OK, click **Save Connector Settings.**

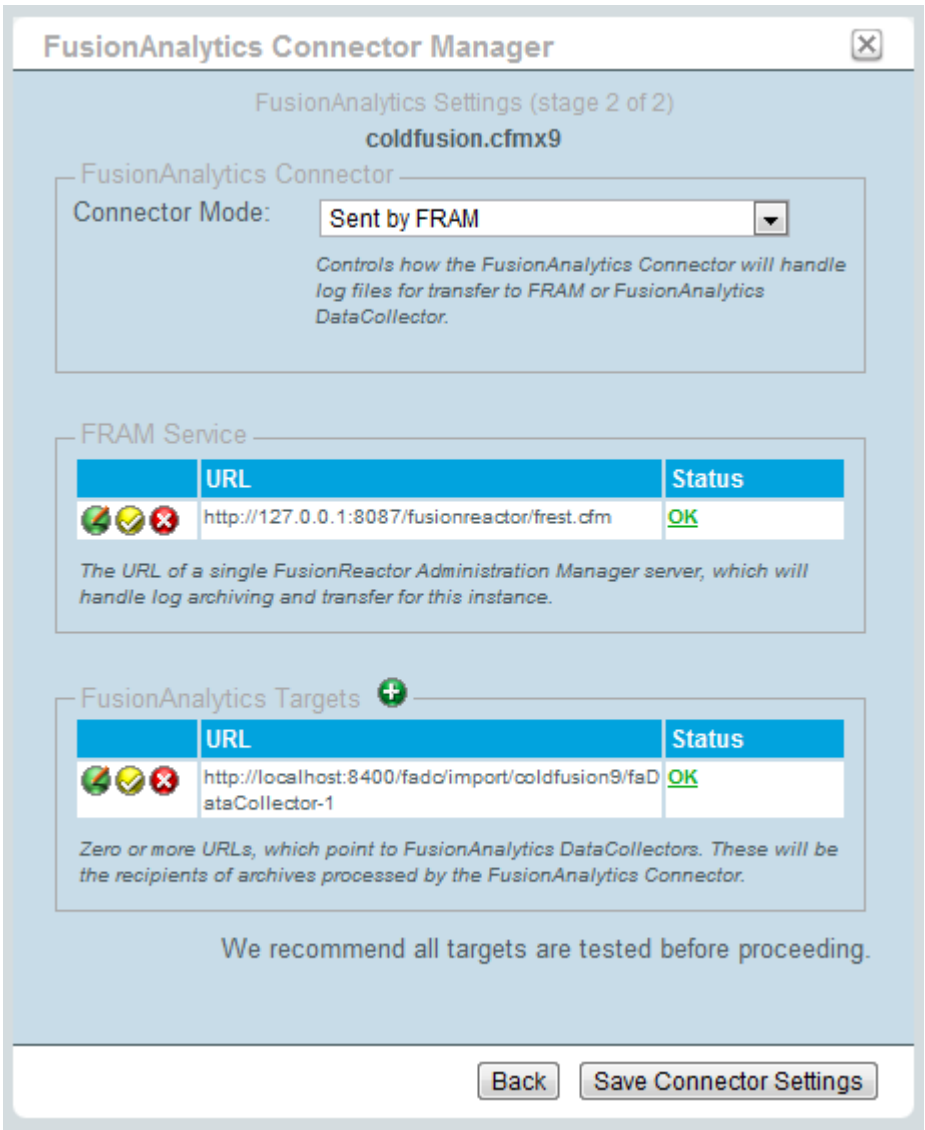

Figure 6: Testing the FusionAnalytics Settings

You have now reached the end of the FusionAnalytics Connector Manager wizard (Figure 7), click **Finish** to return to the Instance Manager.

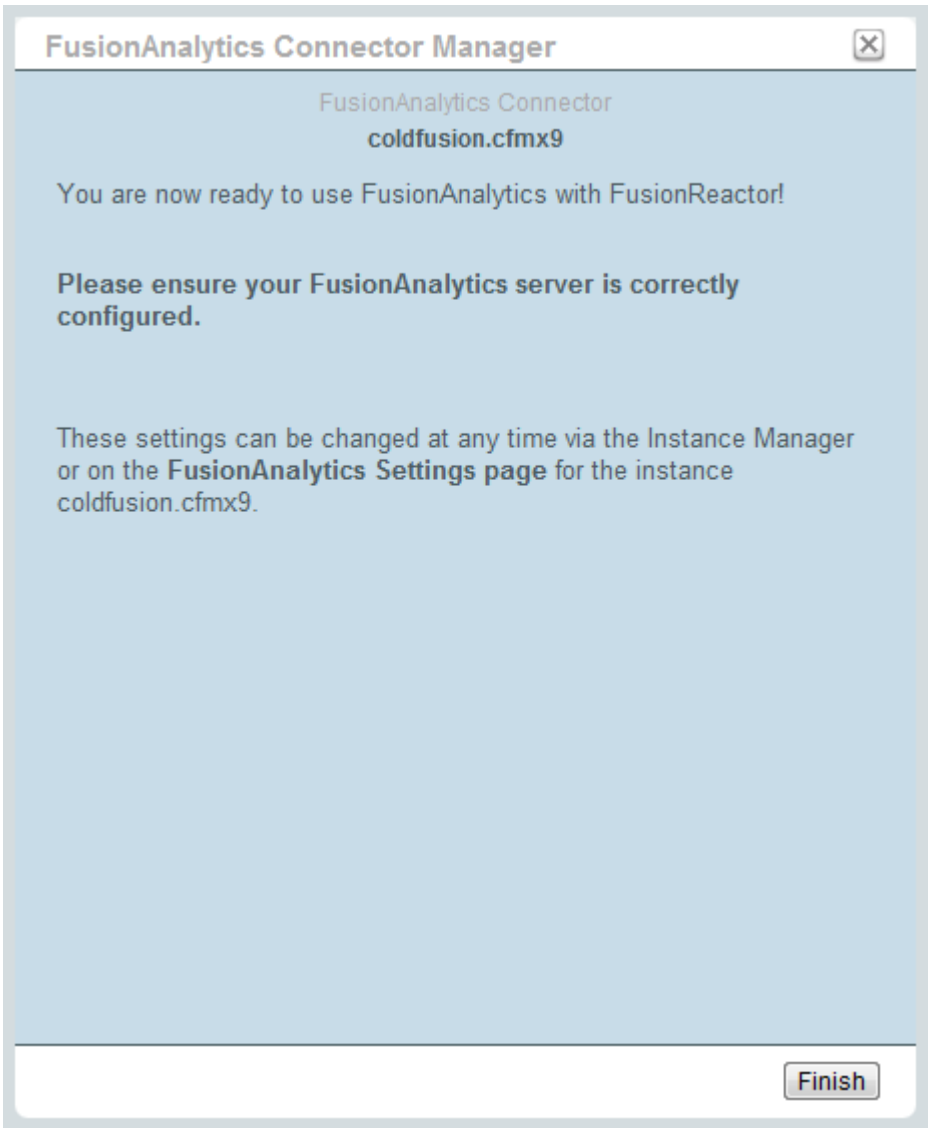

Figure 7: Final screen of the FusionAnalytics Connector Manager

Seeing your data in FusionAnalytics
When you connect FusionReactor to FusionAnalytics, the data will not instantly appear. You will need wait for a transfer to occur.

If you want to force a transfer, in FusionReactor go to **Analytics Connector Status**. Click **Transfer** to send the logs to FusionAnalytics. You will need to wait for each log to be processed by FADC before the data will be available in the Client.

Back to the top

More Configuration Tutorials:

Accessing FusionAnalytics

Setting Up the AIR Client

Configuring FusionReactor for FusionAnalytics

Setting Up Users

Setting Up the Email Server

Setting Up Application Filters

Setting Up TAP and Daily Status Reports

For more information:

FusionAnalytics Documentation

Using FusionAnalytics with FusionReactor

FusionReactor 4 Documentation

Edit FusionAnalytics Connector

Logging

Analytics

Next Steps

Setting Up Users
Add more users to and allow them to view specific applications.

Setting Up Users

Setting Up Users

You can configure more user accounts for FusionAnalytics under **Users Users Summary**.

To add a new user click the **Create A New User** button at the top of the screen (Figure 1).

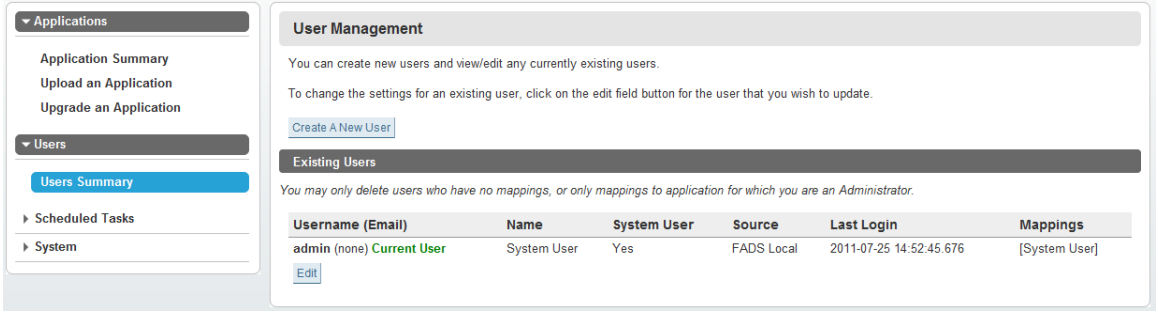

Figure 1: Users Summary

Adding a New User

You must enter a new **Username** and **Password** (Figure 2). **First Name, Last Name** and **Email** fields are optional.

The password strength meter can help you choose a secure password and check that the two fields match.

Click **Add New User** when you have finished entering their details.

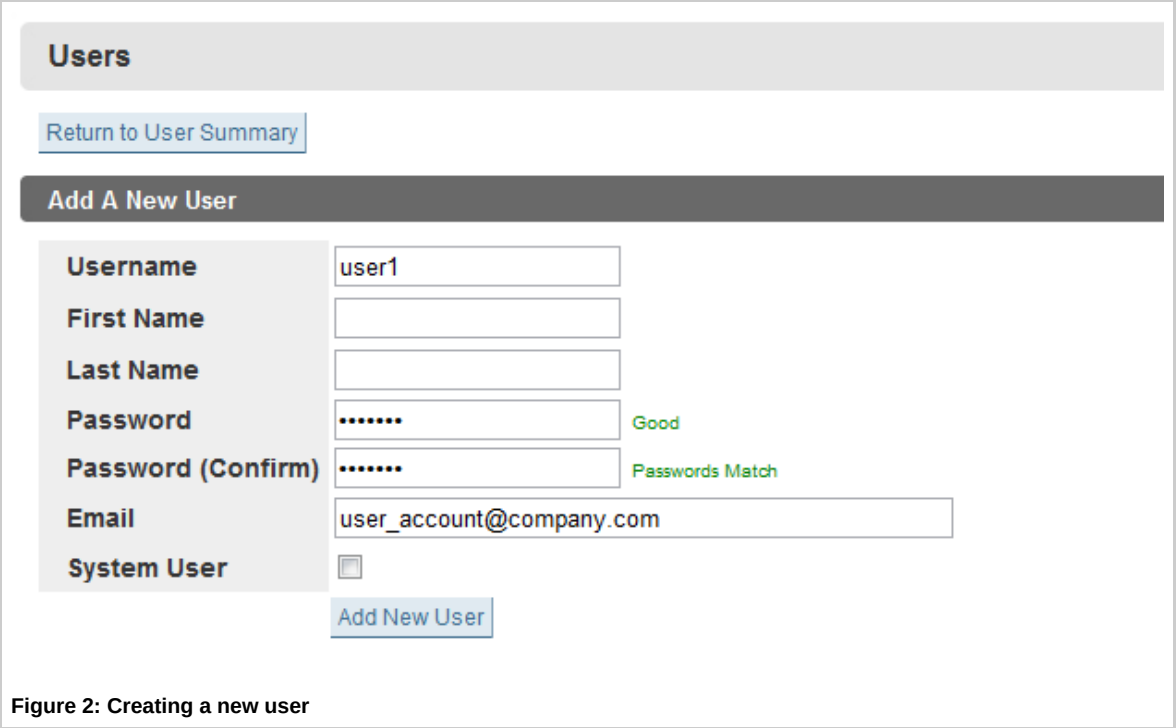

Figure 2: Creating a new user

Back to the top

Adding Mappings to a User

Once you have added the user, you can add any application mappings they might need (Figure 3).

There will be a list of applications under **Available Application Role Mappings.** Choose a role type from the drop down list and click **Add Role.**

The application will now appear under a list of **Current Application Role Mappings.**

For information on the differences between *Administrator* and *User* accounts, see User Mappings in the FusionAnalytics User's Guide.

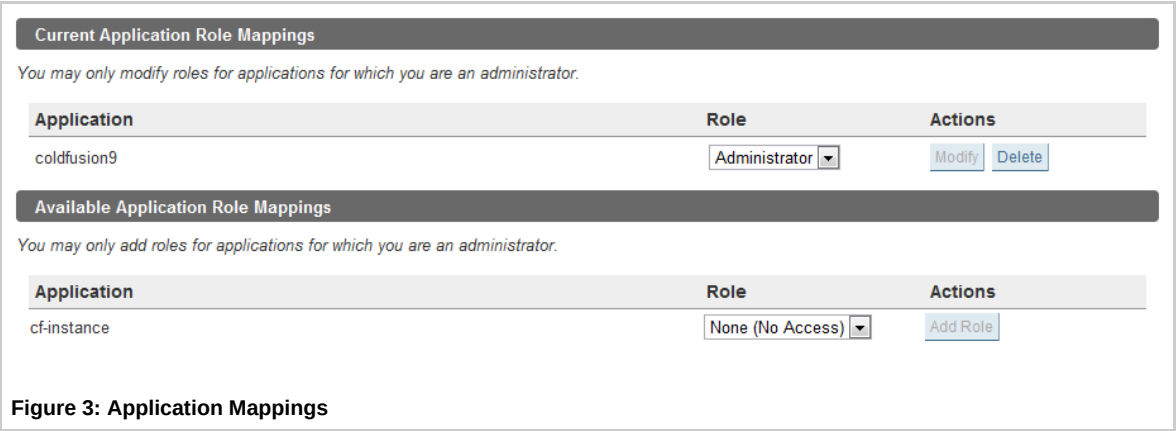

Figure 3: Application Mappings

Logging in as your new user
For an Administrator: Log out of FADS and log in as the new user. You will only see applications you are mapped to. You can also try logging into the client.

For a User: Try logging into the client using the new user account.

Back to the top

More Configuration Tutorials:

Accessing FusionAnalytics

Setting Up the AIR Client

Configuring FusionReactor for FusionAnalytics

Setting Up Users

Setting Up the Email Server

Setting Up Application Filters

Setting Up TAP and Daily Status Reports

Next Steps

Setting Up the Email Server
_Configure the email server so that you receive important notifications and reports from FusionAnalytics.

Setting Up the Email Server

Setting Up the Email Server

To receive notifications and reports from FusionAnalytics, you need to set up your email server.

> This must be completed for **both** Fusion Analytics DataCollector (FADC) and FusionAnalytics DataServices (FADS).

In FADC this can be found under **System Email Settings** (Figure 1).

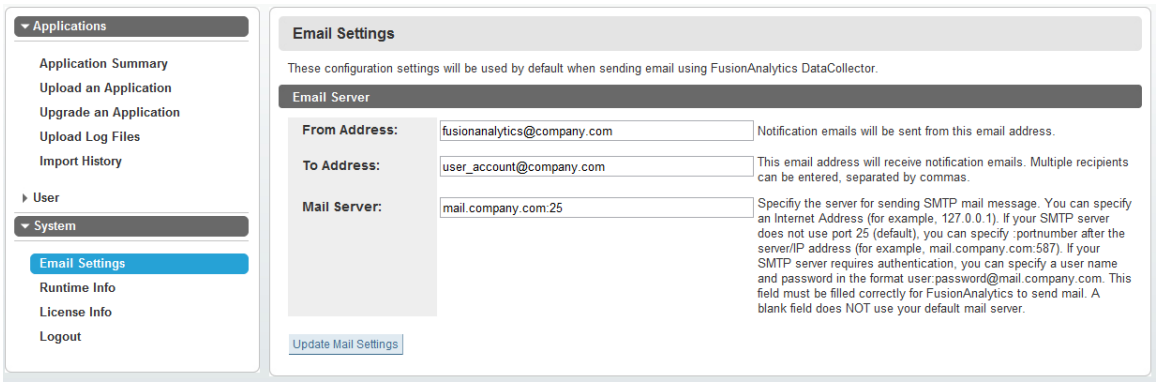

Figure 1: FADC Email Settings

In FADS, it is under **System Server Settings** (Figure 2).

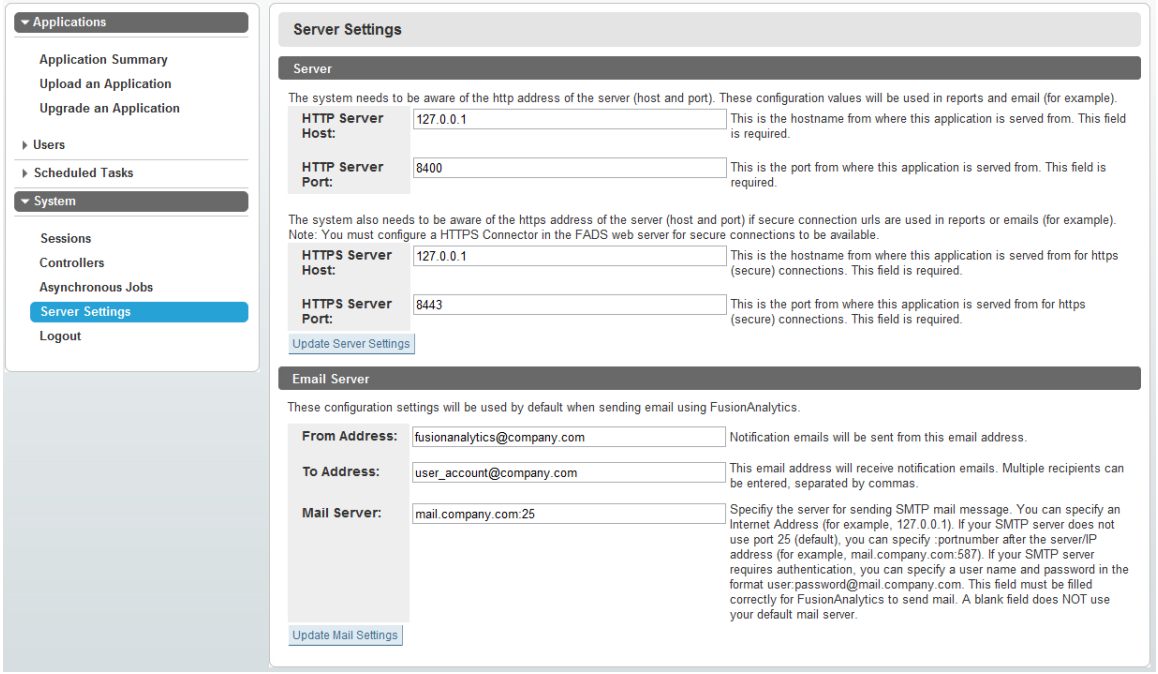

Figure 2: FADS Email Settings

The **From Address** should be the email address that you want any notifications sent from.

The **To Address** is the email address(es) that all notifications are sent to; you can enter multiple addresses, as long as they are separated by a comma.

> ⚠ These addresses do not apply to the reporting feature of FusionAnalytics (which allows you to receive TAP and Daily Status reports on the status of your server by email). You can set different email addresses for these.

The **Mail Server** field is needed for **all** reports and notifications. This will need to be set to the name or IP address of the server that sends SMTP mail messages; e.g. **mail.company.com:25.**

By default the port number is set to 25, but you can specify a port number after the address if needed; e.g. **127.0.0.1:587**.

If you need to provide authentication, you will need to put this in front of the address: **username:password@mail.company.com:25**.

Back to the top

More Configuration Tutorials:

Accessing FusionAnalytics

Setting Up the AIR Client

Configuring FusionReactor for FusionAnalytics

Setting Up Users

Setting Up the Email Server

Setting Up Application Filters

Setting Up TAP and Daily Status Reports

What Next?

Setting Up Application Filters
Configure the application filters to view specific types of request.

Setting Up Application Filters

Setting Up Application Filters

To get the most out of FusionAnalytics, you should set up the Application Filters.

These allow you to filter specific requests within the application, based on their URL or user-defined parameters.

> ⚠ You must be logged in as an Administrator or System User to have access to the Application Filters.

In the Client, go to **Administrator Manage Applications** (Figure 1). Enter a name for your application and click **Add Application Filter.**

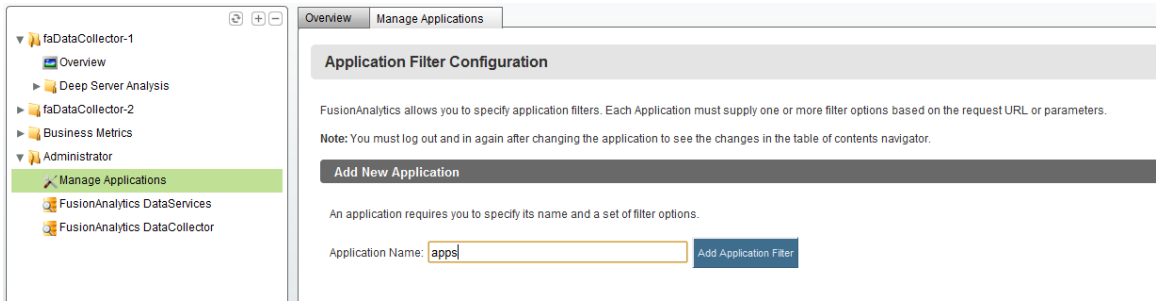

Figure 1: Application filters

The new application filter will appear in a list of **Current Application Filters**. It will be 'opened' and ready to configure (Figure 2).

You can open and close the filters by clicking the  button.

You can add any filters that you need by defining the URL path of the request, or setting specific request parameters.

To add a **Request URL** filter, select a path option and enter an argument. For **Request Params** you will need to enter both a filter and argument. Click  to add the filter.

> *In the example below (Figure 2), the **apps** filter **would** show you the following requests:*
>
> - http://localhost:8500/apps/testpage.cfm
> - http://localhost:8500/apps/newfolder/testpage2.cfm
>
> *It **would not** include these requests:*
>
> - http://localhost:8500/apps/ignorethispage.cfm
> - http://localhost:8500/apps2/newpage.cfm

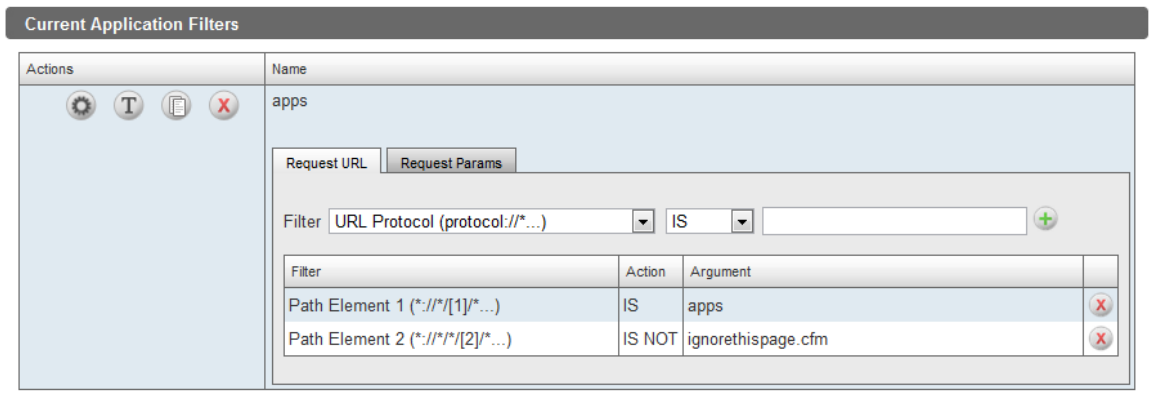

Figure 2: Newly added application in the Current Application Filters

To view the result of your new application filters, click the **Refresh Tree** button on the menu to the left (the table of contents). Now expand **Deep Server Analysis Applications**. You will see the filters that you have defined (Figure 3).

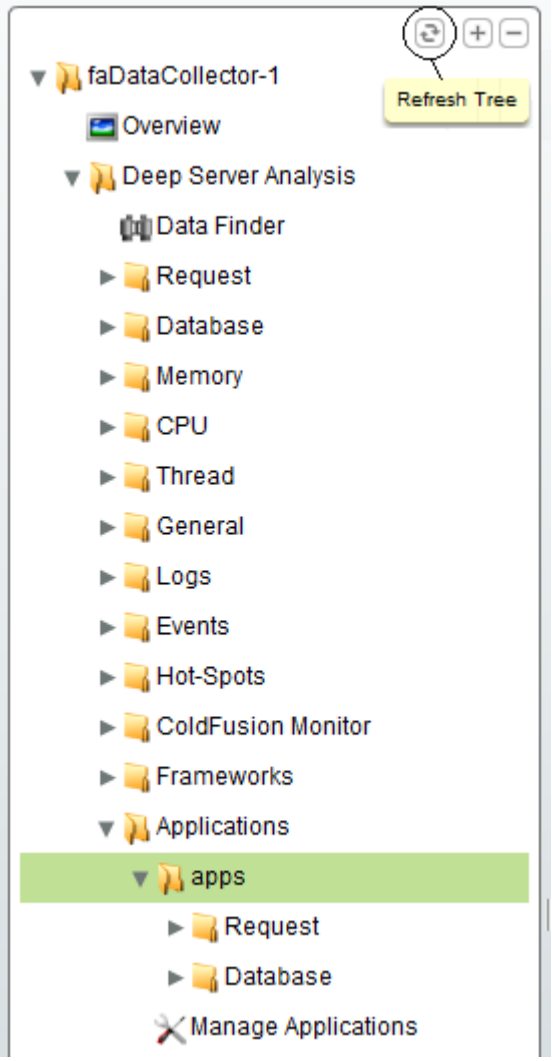

Figure 3: Filters in the TOC

Within each filter you can view a selection of charts and graphs, such as Slowest Requests, Most Popular Requests and Slowest DB Requests, specifically for the requests you have filtered.

To add multiple filters, just enter a new, unique name in the **Add New Application** section and click **Add Application Filter**.

For a more detailed look at the application filters, including deleting, modifying and copying, see Manage Applications in the FusionAnalytics User's Guide.

Back to the top

More Configuration Tutorials:

Accessing FusionAnalytics

Setting Up the AIR Client

Configuring FusionReactor for FusionAnalytics

Setting Up Users

Setting Up the Email Server

Setting Up Application Filters

Setting Up TAP and Daily Status Reports

Next Steps

Setting Up TAP and Daily Status Reports
The final stage in this configuration tutorial is setting up daily reports that can give you information on how your servers are performing.

Setting Up TAP and Daily Status Reports

Setting Up TAP and Daily Status Reports

The reporting feature of FusionAnalytics allows you to receive daily reports on the status of your server.

There are 2 types of report:

TAP Reports
These provide a daily application score and trend analysis, based on your system's Traffic, Availability and Performance (TAP).

Daily Status Reports
Aimed at more technical users, these reports give a complete breakdown of the performance for the whole application and individual Application Filters.

Viewing Application Tasks

In FADS, under the **Scheduled Tasks Application Tasks Summary** you can find the two types of reports that can run for each application (Figure 1).

These are already configured for you with each new application.

The default setting is for these to run every day at 4.15am (TAP Report) and 4.30am (Daily Status Report).

Figure 1: Pre-scheduled Application Tasks for the *coldfusion9* application

When the reports have run, they can be viewed in the Client, under **Business Metrics Reports**.

Receiving Reports as Emails

> ⚠ You **must** have completed the Setting Up the Email Server part of the Quick Start Guide tutorial.

A major feature of FusionAnalytics is that it is able to email you a daily report, so that you can easily keep track and analyze its performance.

Click **Edit** to enter the configuration screen for a report (Figure 2 shows editing a Daily Status Report).

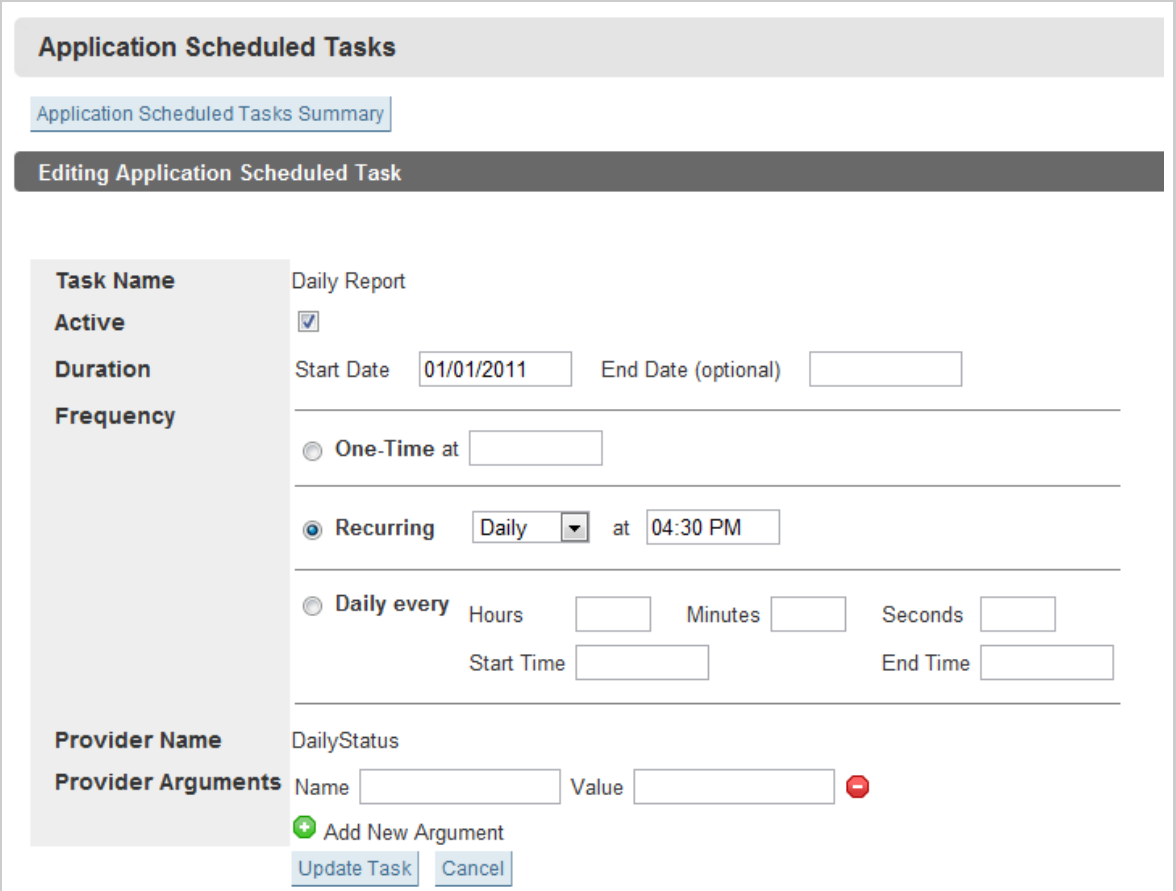

Figure 2: Editing a Daily Report

Under provider arguments, enter ***'users'*** as the name and a valid username as the value (Figure 3).

Make sure the users have email addresses set. This can be changed under **Users Users Summary.**

Multiple user names can be entered if they are separated by a comma (but no space).

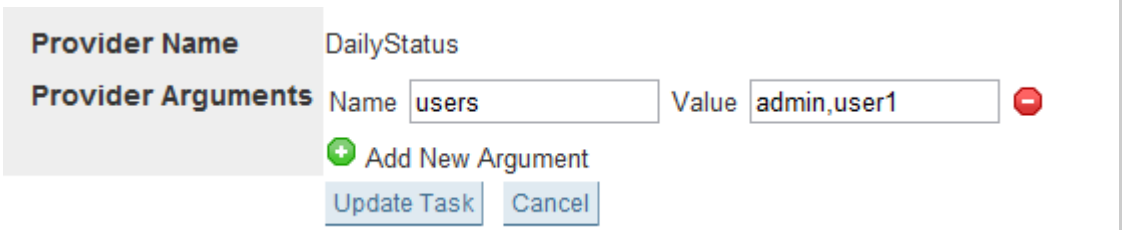

Figure 3: Entering provider arguments

Click **Update Task**. The next time the reports run, they will be emailed to the selected users.

> You will need to repeat this for the TAP Report, so that both reports get emailed!
>
> Both these reports gather data from the previous day, so you will need to wait **24 hours** before these reports begin to show any data.

Back to the top

More Configuration Tutorials:

Accessing FusionAnalytics

Setting Up the AIR Client

Configuring FusionReactor for FusionAnalytics

Setting Up Users

Setting Up the Email Server

Setting Up Application Filters

Setting Up TAP and Daily Status Reports

Next Steps

FusionAnalytics User's Guide
You have now finished configuring FusionAnalytics! Visit the User's Guide for more information on specific parts of FusionAnalytics.

FusionAnalytics Reporting
Visit this section for more information on these reports, including more configuration options and understanding the report data.

Installation Guides

Installation Guides

Overview

This section provides information about the general system requirements of FusionAnalytics and all of the installation instructions required to get FusionAnalytics running.

Next Steps

System Requirements

- Why We Recommend Installing FusionAnalytics on Separate Machines

Installing FusionAnalytics Server

- Prerequisites for Installation
- Installation
- Installing your license

Installing the FusionAnalytics AIR Client

Installing FusionReactor

Installing FusionReactor Extensions for ColdFusion

Uninstalling FusionAnalytics

System Requirements

System Requirements

Overview

FusionAnalytics Server

Operating Systems	 Microsoft Windows Server 2008 Microsoft Windows Server 2003 Microsoft Windows 7 Microsoft Windows XP Professional Microsoft Windows Vista Microsoft Windows 2003 Web Edition, Standard Edition, Enterprise Edition
Memory	4 GB (MIN)
Hard Disk Space	25 GB (MIN) - Using ***Microsoft SQL Server 2008 R2 Express***
Architecture	64-BIT / 32-BIT

 Memory and hard Disk Space requirements are **completely** dependent on the amount of metric data you with to store and have available in FusionAnalytics

 Please note that version 1.0.0 of FusionAnalytics will **only** support Microsoft SQL Server 2008 R2, Microsoft SQL Server 2008 and Microsoft SQL Server 2005

FusionAnalytics Client

Adobe AIR

- Supported from version 2.7 onwards
- Visit the Adobe website for more information

 **Download for free:** Latest version of Adobe AIR

Next Steps

Why We Recommend Installing FusionAnalytics on Separate Machines

Why We Recommend Installing FusionAnalytics on Separate Machines

Why is it recommended to install FusionAnalytics on separate machines?

FusionAnalytics comes in three main parts:

- DataCollector
- DataServices
- Client.

Both DataCollector and DataServices are part of the FusionAnalytics Server.

We recommend installing the Server and Client on separate machines as this will let you monitor your applications from your desktop. The Client can also be installed on multiple machines, allowing different users to access their own applications from their own desktops.

Where should each setup be installed?

The FusionAnalytics Server should be installed at a location that can be accessed by anyone using a client.

The FusionAnalytics Client should be installed at a location from where you would like to view your data.

Installing FusionAnalytics Server

Installing FusionAnalytics Server

Overview.

This guide will help you with your first installation of FusionAnalytics.

Prerequisites for Installation
Make sure you are ready to begin the installation.

Installation
Once your system is ready, this section will guide you through the install process.

Installing your license
You can visit the tutorial on license files for help uploading and activating your licenses.

Quick Start Guide
Go back to the Quick Start Guide to view your next steps.

Take me back to the **Installation Guides**.

Next Steps

Prerequisites for Installation

Installation

Installing your license

Prerequisites for Installation

Prerequisites for Installing FusionAnalytics Server

1. You have checked the minimum system requirements for the system where you want to install FusionAnalytics (see System Requirements).
2. You have prepared your database (see Database Configuration).
3. You have the database login details available, with appropriate privileges on the target database.
4. You have the appropriate FusionAnalytics installer.
5. JavaScript is enabled in your default browser.

Next Steps

Installation
Once you have completed the above steps, you can start the installation.

Installation

Installing FusionAnalytics Server

Executables

Executable	System Architecture
FusionAnalytics_windows_1_0_0.exe	32-bit
FusionAnalytics_windows-x64_1_0_0.exe	64-bit

Back to the top

Setup

To start the installation, double-click the installer.

> ⚠ You will need to have Administrator privileges on the system you are installing on.

When the installer starts you will see the Welcome screen (Figure 1). You can click the **Next** and **Back** buttons to navigate through the installer.

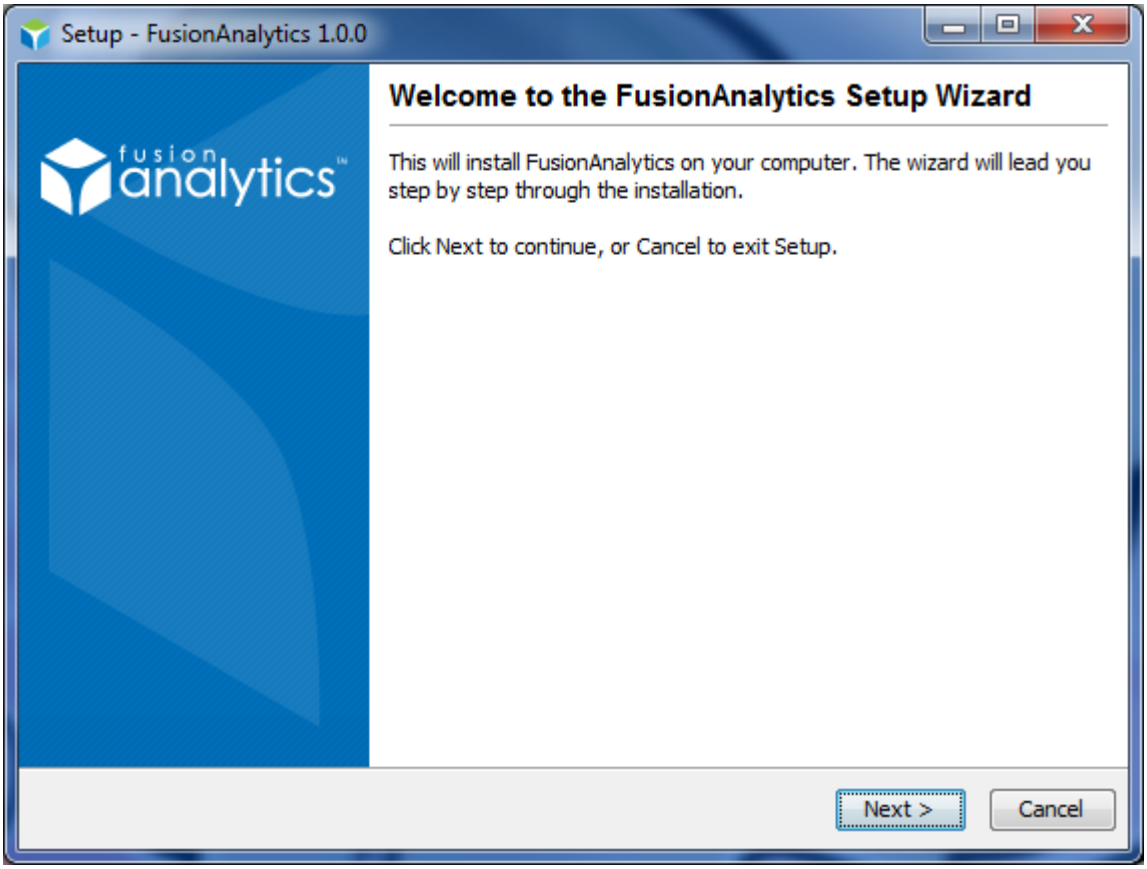

Figure 1: Welcome Screen

The next screen will show you the **License Agreement** (Figure 2). If you agree to its terms and conditions, you need to select *'I accept the agreement'* before continuing.

⚠ You will not be able to install FusionAnalytics if you do not accept the license agreement.

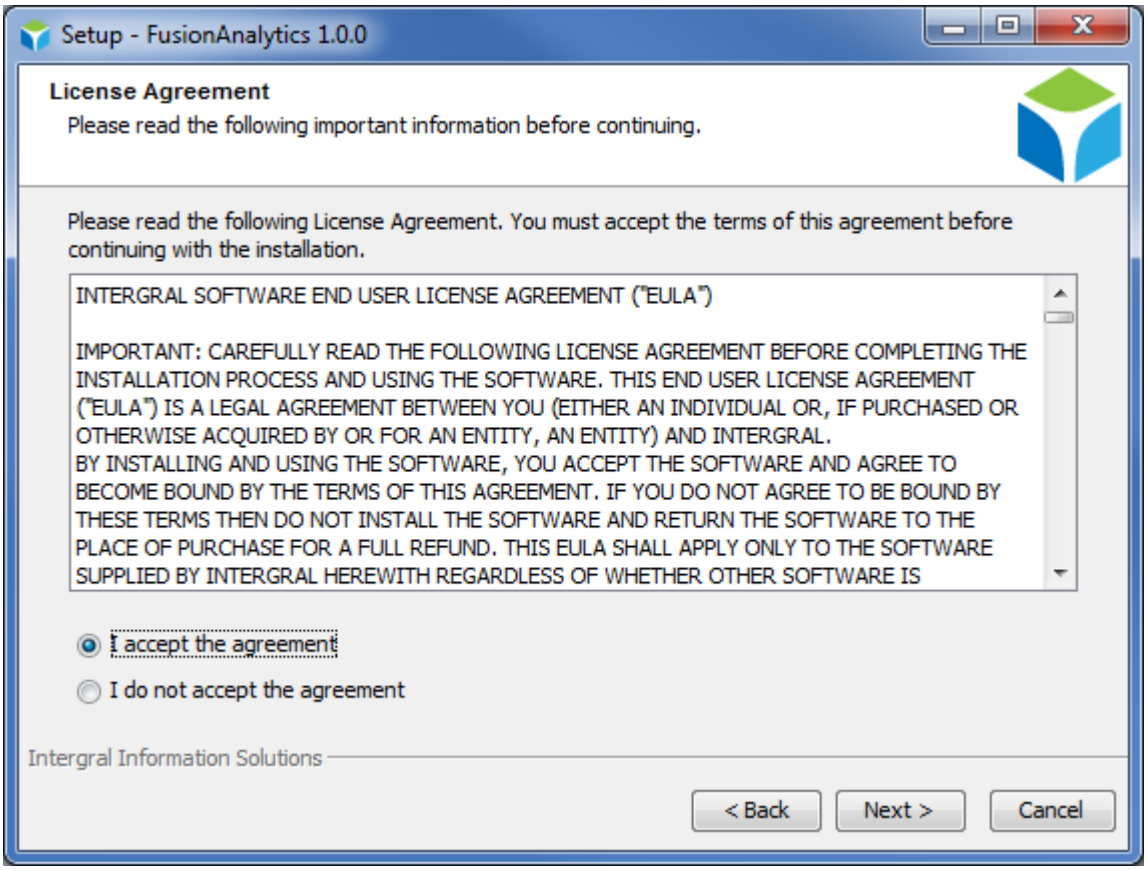

Figure 2: License Agreement

Next you will see the **Select Components** screen. The default (and **recommended**) option is to install the whole FusionAnalytics package (Figure 3).

A description of what each component does can be seen under its name (or by clicking the green question mark).

If you would only like to install one of the components on the current machine, make sure only the component you want to install is selected.

⚠ If you only install one component, you will skip the Application Name and Database Configuration screens. After install, you will need to upload an application and connect it to a database manually.

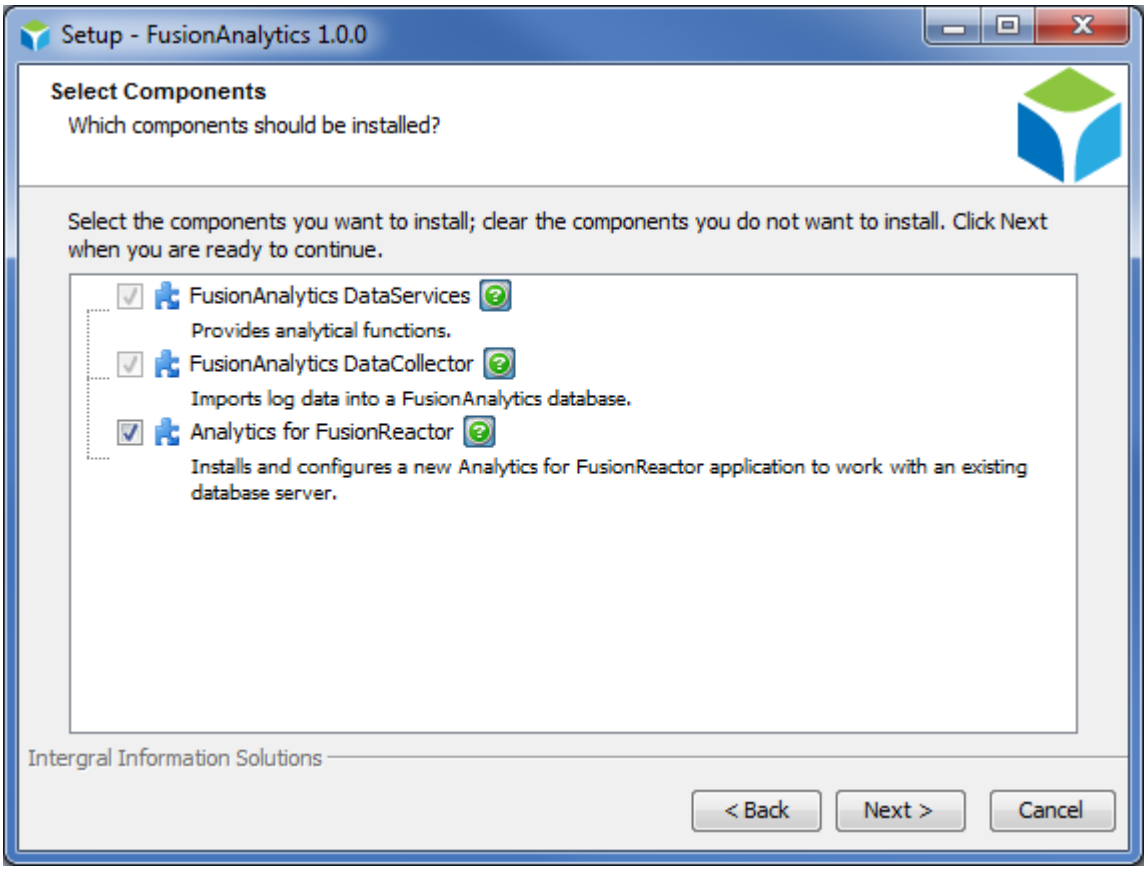

Figure 3: Select Components

The next screen allows you to select the **Destination Directory**, which is where FusionAnalytics will be installed (Figure 4).

You can change this location by clicking the **Browse...** button.

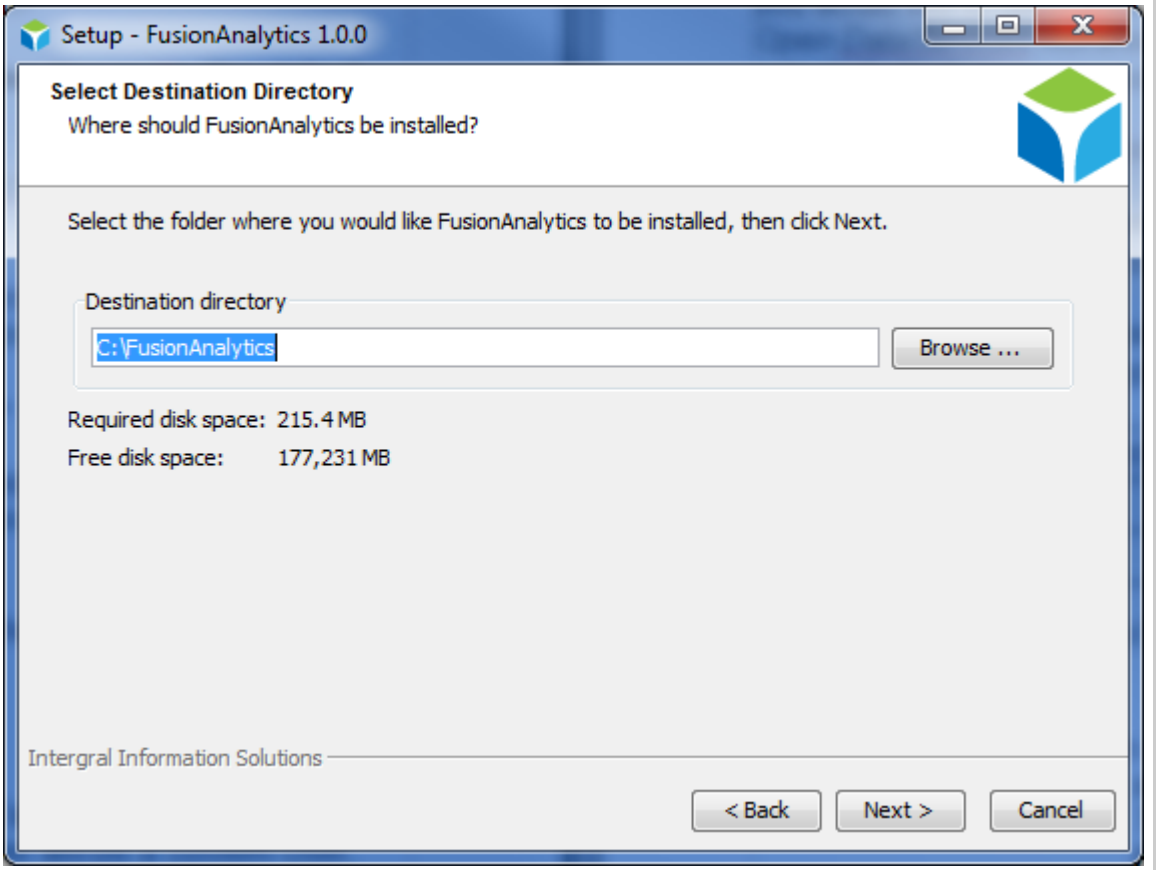

Figure 4: Destination Directory

The following screen allows you to choose a **Start Menu Folder** (Figure 5).

If you do not want to create a folder in the start menu, uncheck the box marked *Create a Start Menu folder*.

You can also choose if you want to create shortcuts for all users or only the current user, by checking or unchecking the box marked *Create shortcuts for all users*.

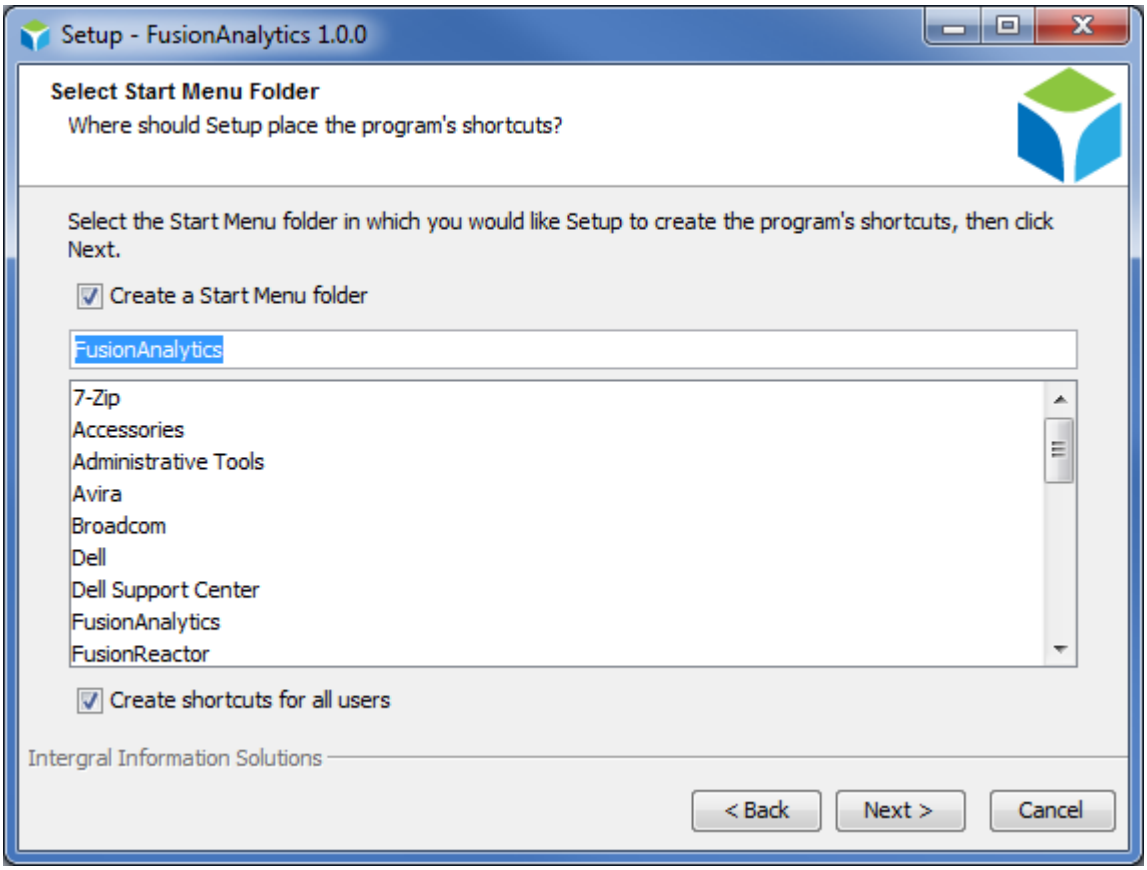

Figure 5: Start Menu Folder

On the next screen you can define the **TCP/IP Ports** to be used by FusionAnalytics (Figure 6). If you want different ports to be used, check the box marked *Adjust Settings* and adjust the port numbers.

The setup will check if the ports are free before continuing. If they are not available, they will need to be changed.

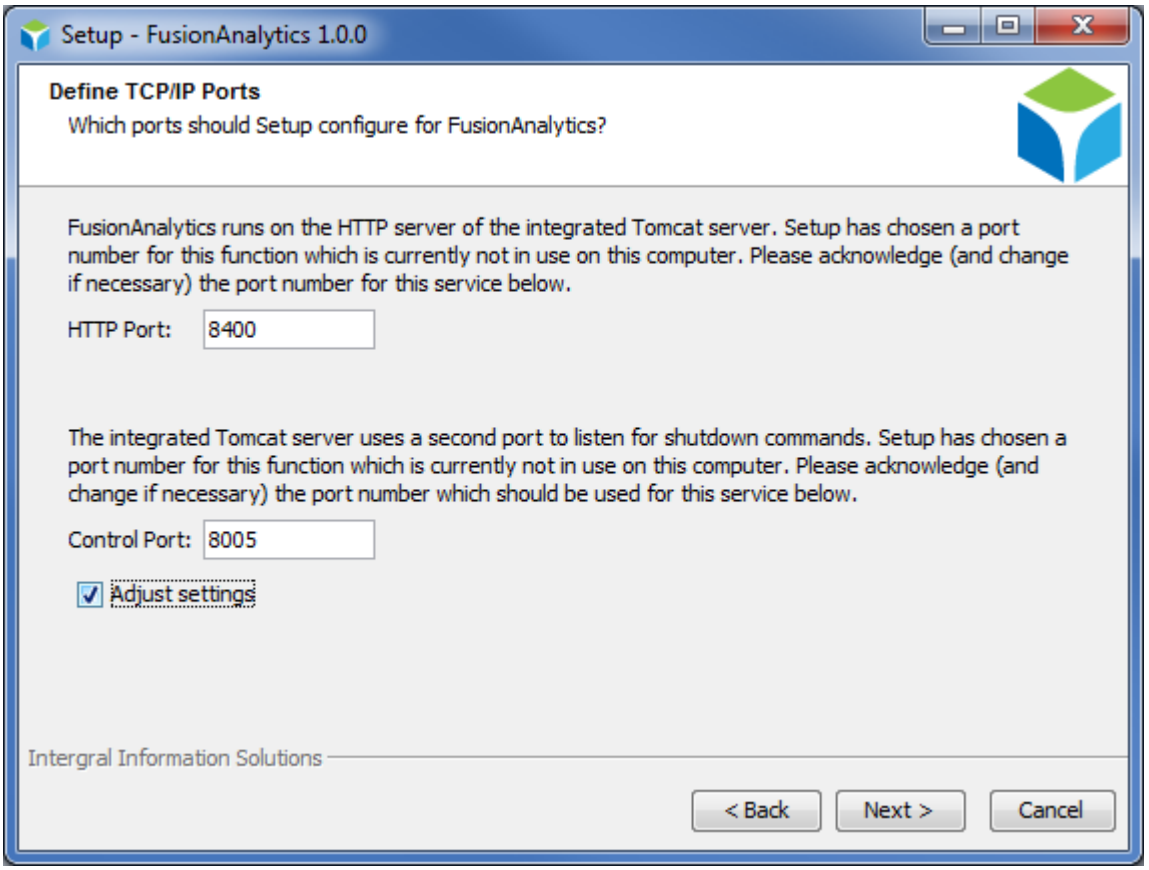

Figure 6: TCP/IP Ports

The default user that will be created for FusionAnalytics is called **'admin'**; for this you will need to set the **Administrator Password** (Figure 7). The **'admin'** user will be the System User.

You need enter the password in the Password and Confirmation fields to make sure they match, before clicking **Next**.

> ⚠ Make sure you choose a secure password, as this **'admin'** default system user account will have access to everything within FusionAnalytics.

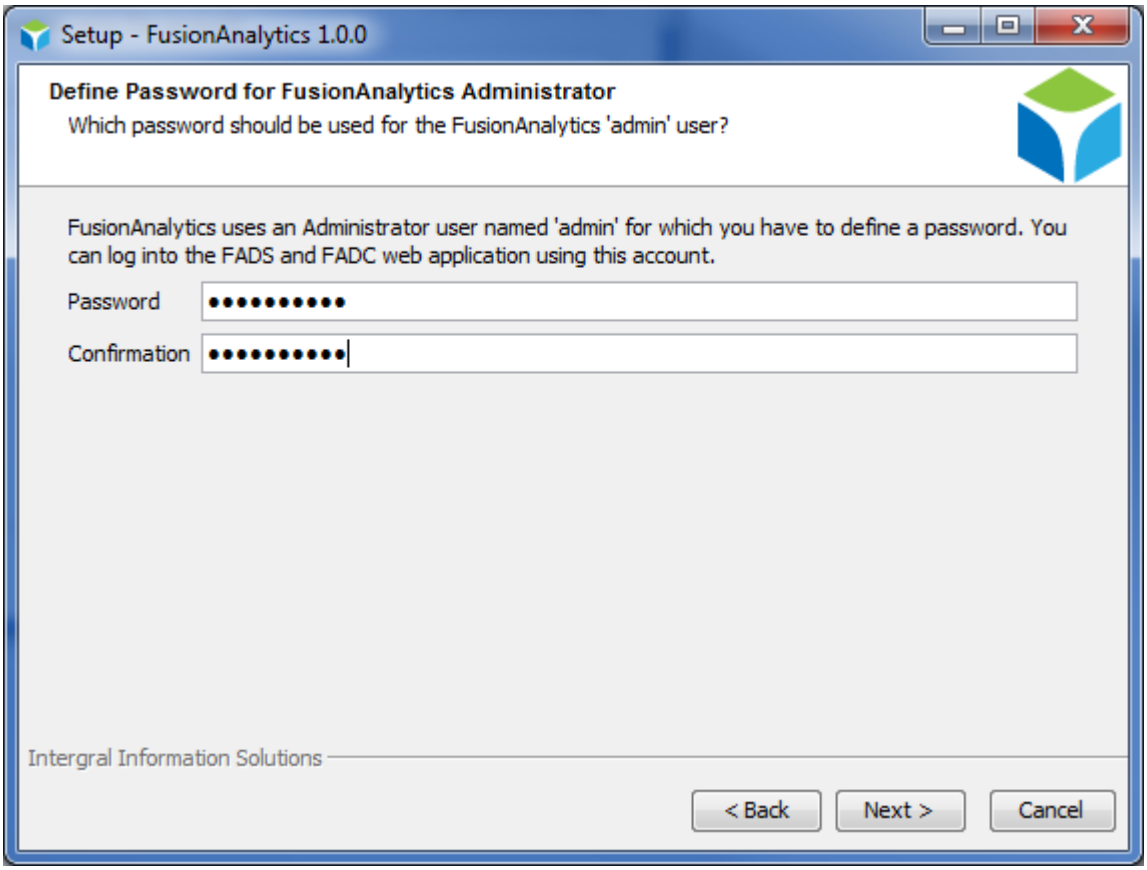

Figure 7: Administrator Password

Next, you will be asked to enter an **Application Name** (Figure 8). This will create one application for FusionAnalytics DataCollector and one application for FusionAnalytics DataServices.

This needs to be a unique instance name, since later you can upload further applications, which may point to different databases and log files.

> *For this guide we will use **coldfusion9**, as this is what we will be monitoring.*

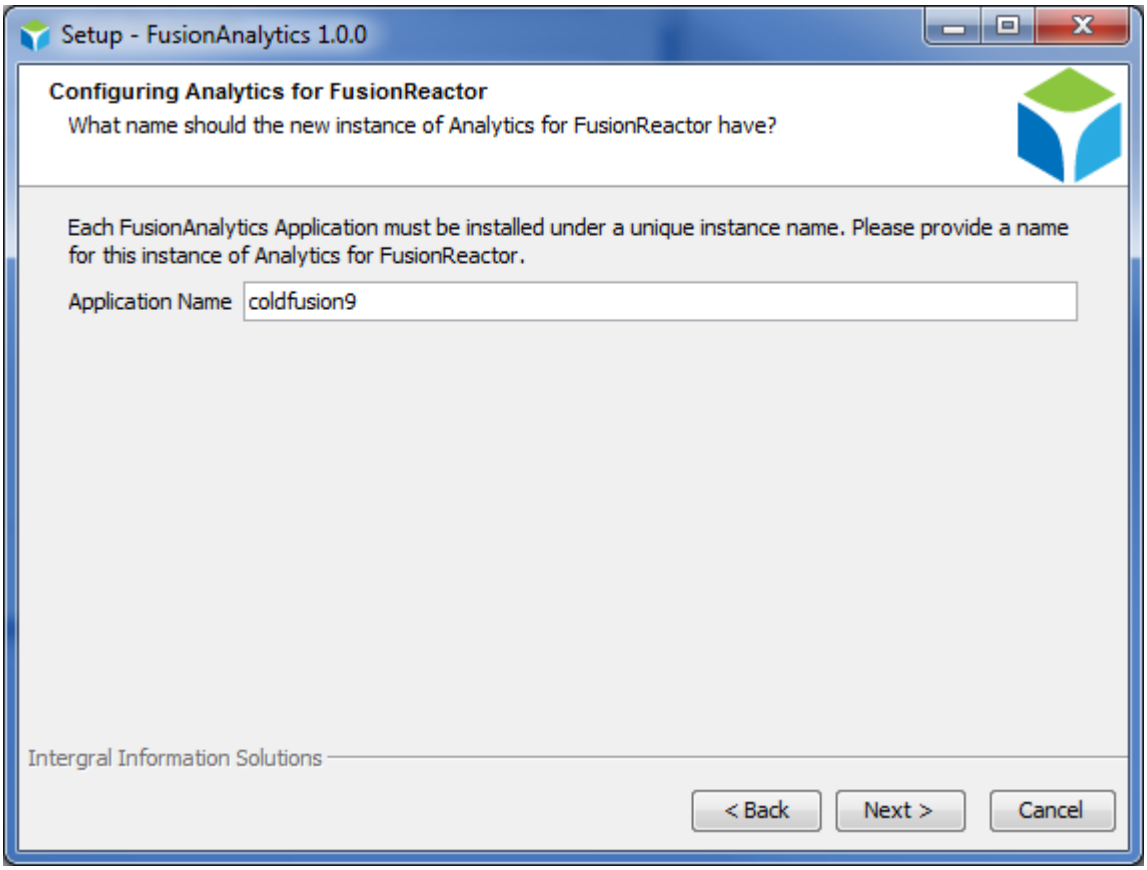

Figure 8: Application Name

The next screen (Figure 9) allows you to enter the **Database Configuration** for the database that you have already prepared (see Database Configuration).

You will need to enter the database name, machine name or IP address, and the username and password which has the appropriate privileges on that database.

Before continuing, click **Test Connection** to make sure a connection can be correctly established.

> *If you have followed the Database Configuration tutorial, you should use the **fadb** database we set up, with **fadb** as the user.*

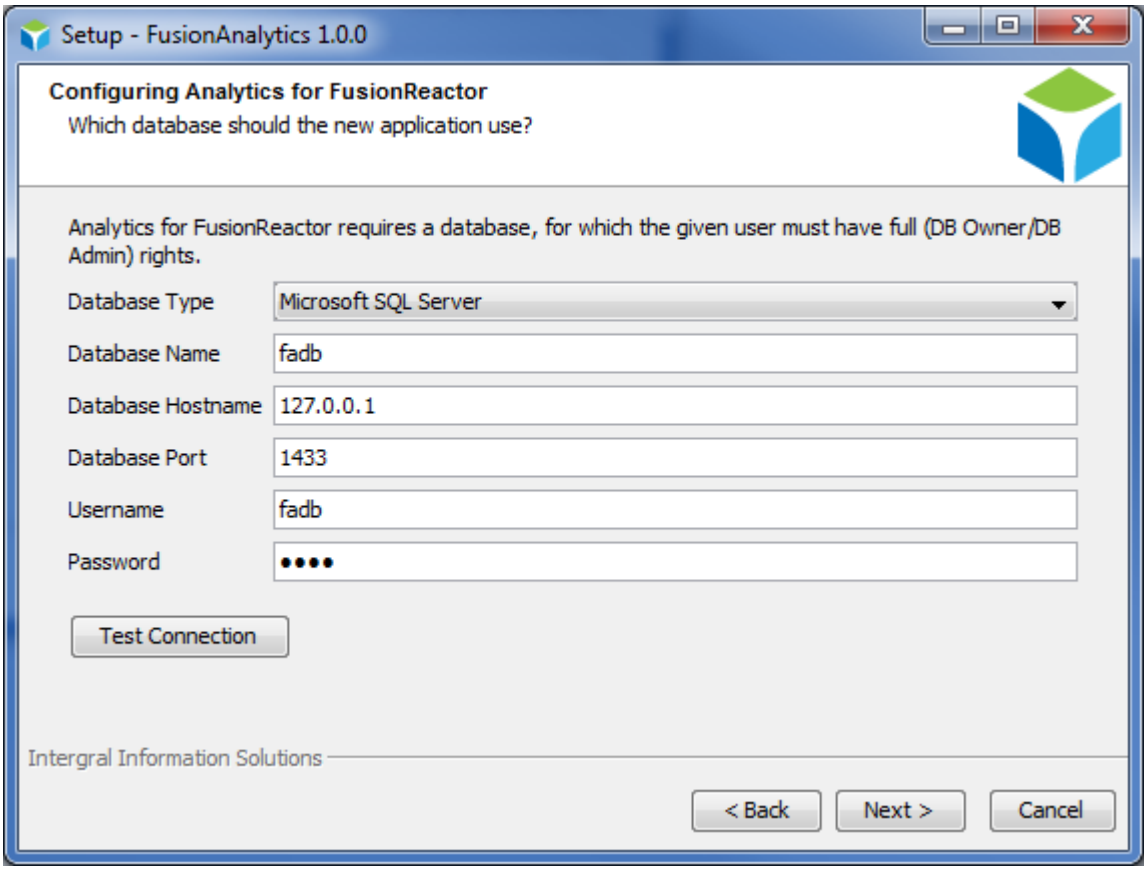

Figure 9: Database Configuration

In the **Additional Tasks** screen (Figure 10) you can select which parts of FusionAnalytics you would like to open after the install finishes.

It is recommended that you select both *Open DataServices Administrator* and *Open DataCollector Administrator*.

After you have selected which ones you want to open, click **Next**.

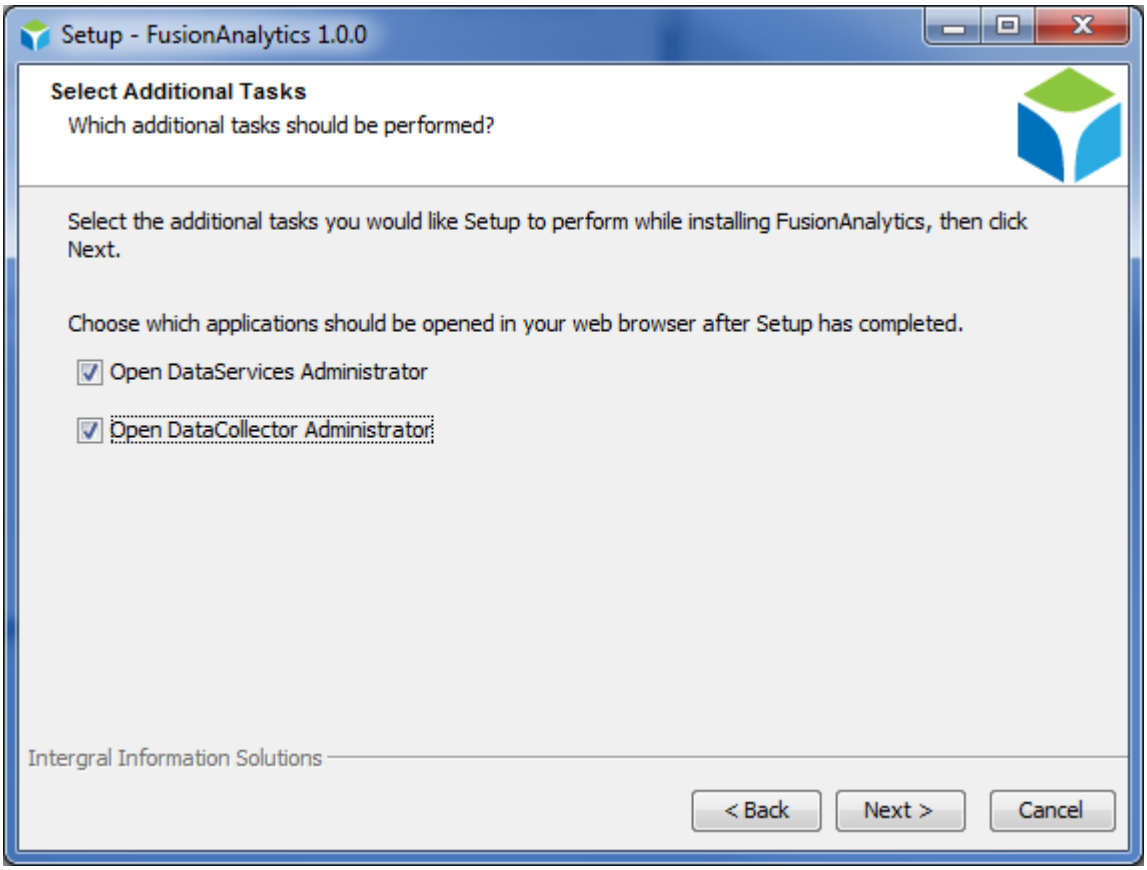

Figure 10: Additional Tasks

The installer will now install the files to the selected directory and the installation wizard will complete. Click **Finish** to close the installer.

> ⚠ If you chose to open either of the Administrator windows at the end of the installation, it can take 30 seconds or more (depending on the performance of your server) for the service to start and the pages to be displayed.

> ℹ You can log into both components using the ***'admin'*** user we configured in the installation.

Back to the top

On this page:

Next Steps

Installing your license
After logging in for the first time, you will need to upload and activate your licenses.

Installing the FusionAnalytics AIR Client
Installing the AIR client will allow you to start analyzing the data from your server.

Quick Start Guide
Go back to the Quick Start Guide to review your next steps.

Installing your license

Installing your license

Note: If you have installed FusionAnalytics for the first time you will be granted a Free Trial Period (FTP) license which allows you to use the system for 10 days.

Uploading a License

> If you are using multiple product families you will need to upload a license for each.

You can upload your license in either FADS or FADC from the navigation menu > System > **License Info** (Figure 1).

Figure 1: License Info

This will take you to the License Info screen (Figure 2). Click the **Install License** button.

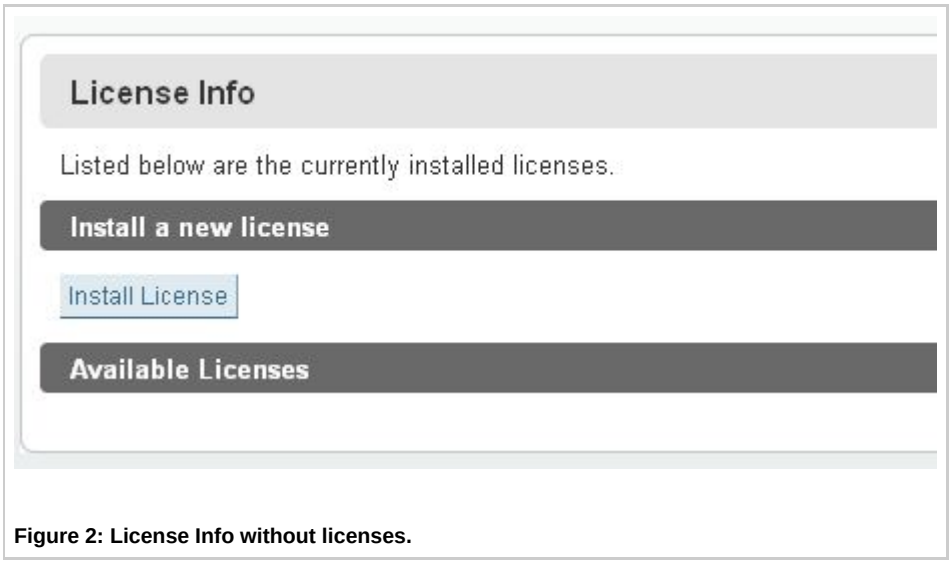

Figure 2: License Info without licenses.

You will be directed to the License Upload screen (Figure 3). Click the browse button and navigate to the license file you wish to upload. Then click **upload**.

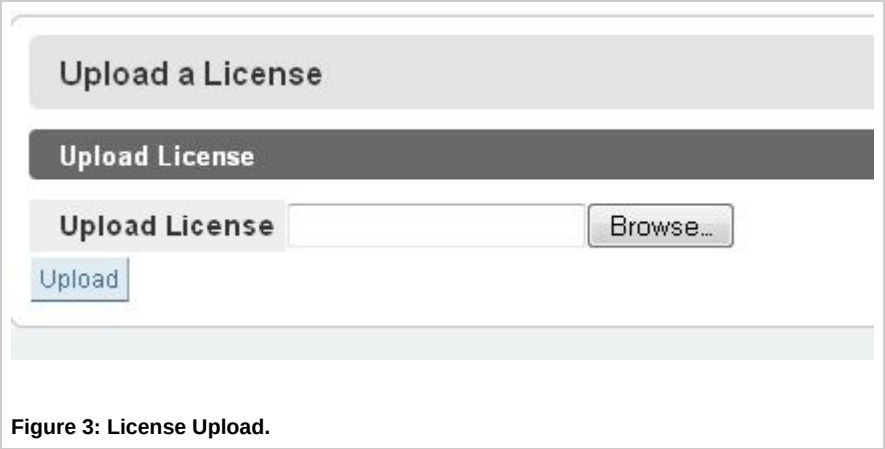

Figure 3: License Upload.

If your upload was successful you should see something like the following:

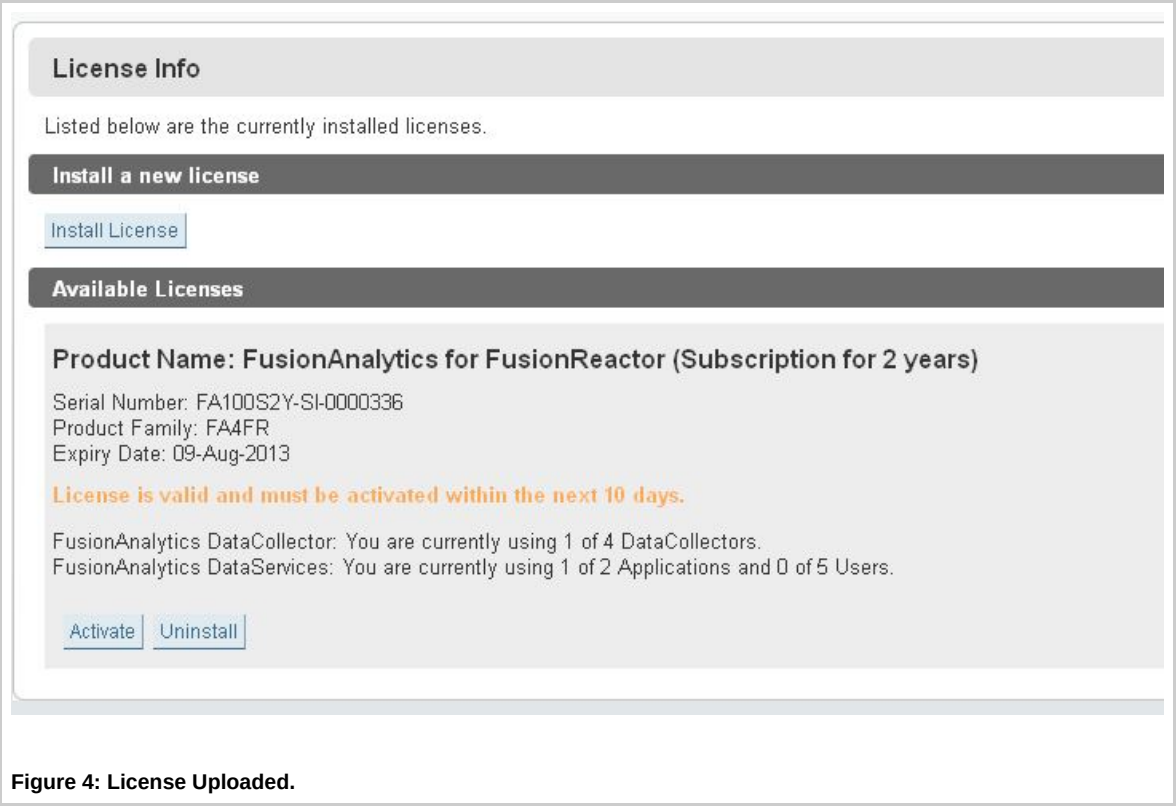

Figure 4: License Uploaded.

Your license will now be fully functional. You **must** activate your license within the number of days shown or you will no longer be able to use analytics for that product family until you activate.

> ⚠ If you encountered an error uploading your license, please check you are using a valid license file.

Back to the top

Activating a License

Once your license is uploaded it needs to be activated. To activate your license simply click the Activate button and click "OK" on the prompt.

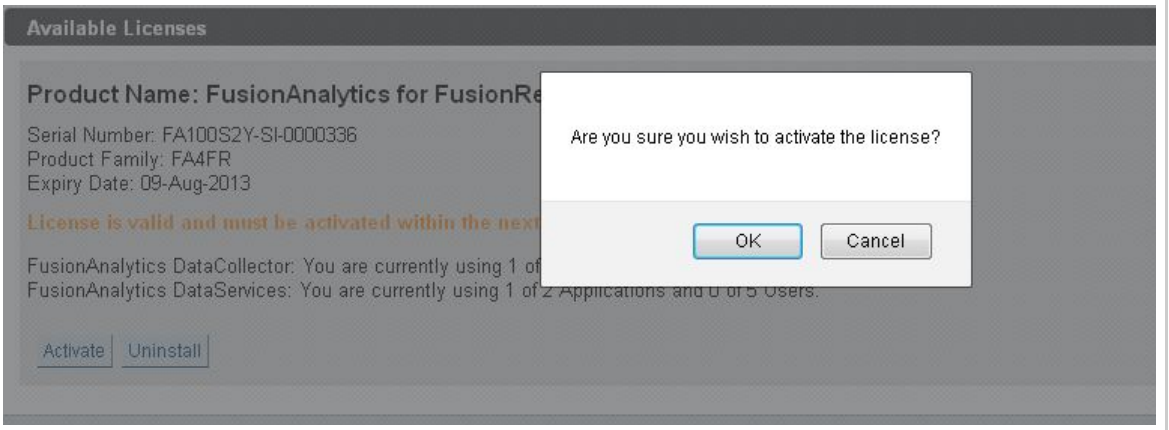

Figure 5: Activate confirmation.

There are multiple things that can happen at this stage:

- Your license activates straight away. See License Activated
- You are taken to the manual activation page because you have no internet connection. See Manual Activation
- Your license has no activations left. See License Rejected

Back to the top

Manual Activation

If the server running Analytics is not connected to the internet you will be taken to the manual activation page. Follow the steps on screen to manually activate your license.

Figure 7: Manual activation

Back to the top

License activated straight away

If your license activated instantly, you should see a screen similar to the following:

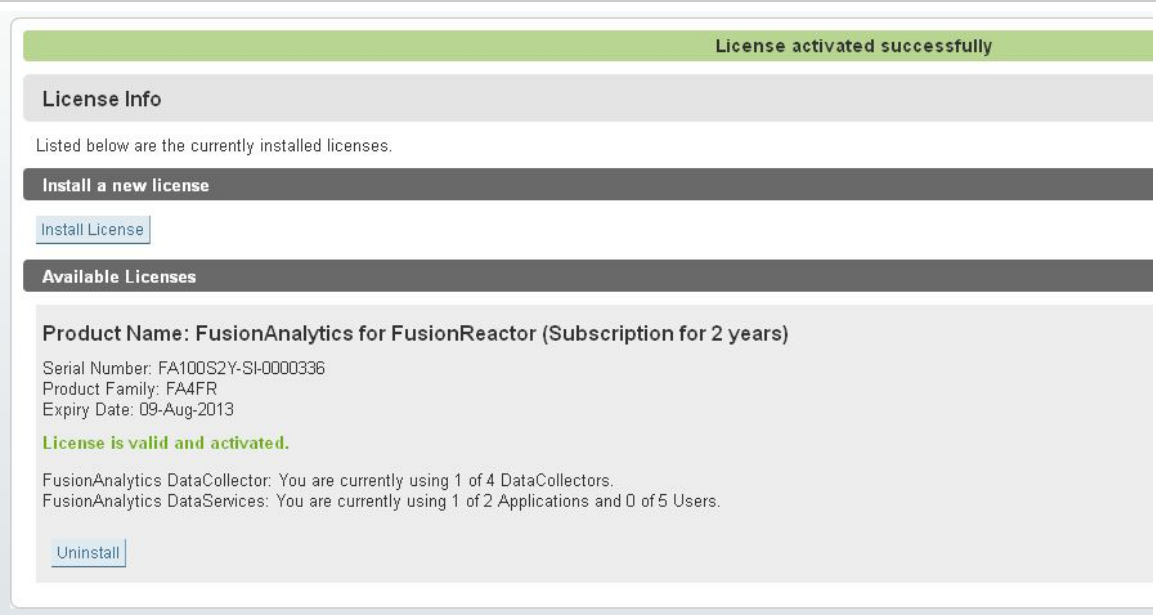

Figure 6: License activated.

Your license is now fully activated.

Back to the top

Activation rejected due to no activations remaining

If you have run out of activations, you will have to uninstall this license on another machine that is using it before you are able to activate it again.

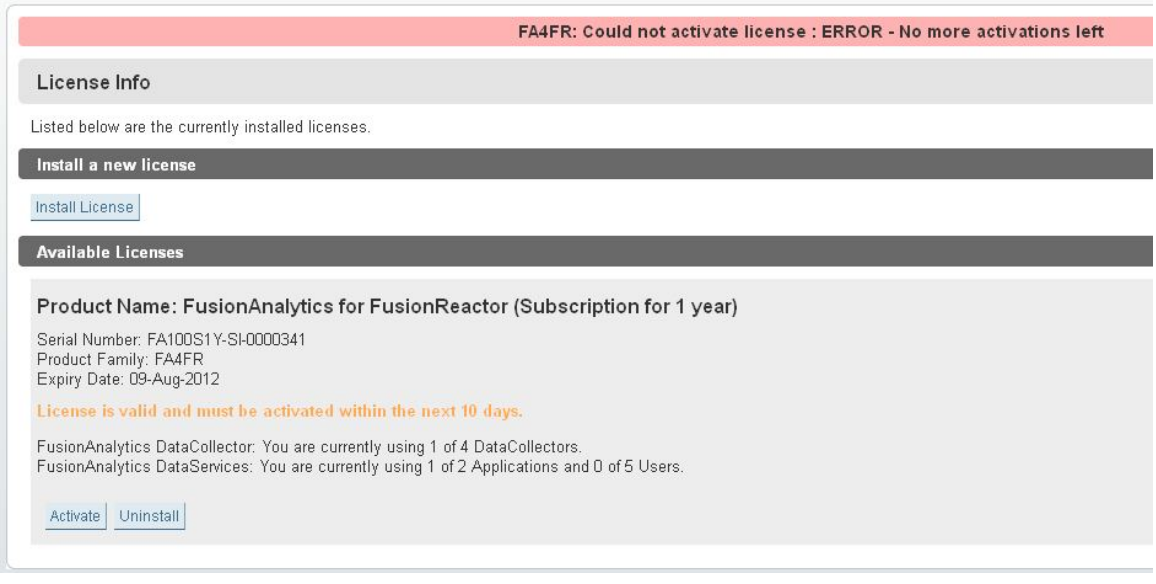

Figure 8: No activations remaining

Back to the top

On This Page:

Free Trial Period (FTP) License
If you have installed FusionAnalytics for the first time you will be granted a Free Trial Period (FTP) license which allows you to use the system for 10 days.

Next Steps

Installing the FusionAnalytics AIR Client
If you have not already done so, you will need to install the AIR client to start analyzing the data.

Installing the FusionAnalytics AIR Client

Installing the FusionAnalytics AIR Client

The AIR Client should be installed on any machine that you want to use to monitor your server. This **does not** have to be on the same machine where you have installed the FusionAnalytics Server. You can use the AIR Client to monitor the data from multiple applications from your desktop.

Before continuing, you will need to make sure that the machine you want to install the FusionAnalytics AIR Client on has Adobe AIR already installed on it. See System Requirements for more information.

Download for free: Latest version of Adobe AIR

About Adobe AIR

At a basic level, you can consider Adobe AIR to be a bit like a new web browser specifically designed to run Adobe Flash content. The FusionAnalytics client is made available as Flash content so it can be accessed in normal web browsers (by clicking the Launch button on the FADS Application Summary page) or by running it in Adobe AIR. Because Adobe AIR doesn't have to worry about all of the other tasks that web browsers have to deal with, you often find that you get better performance from applications which use Adobe AIR. Additionally, Adobe AIR can support more functionality because it is running directly out of your operating system rather than as a plugin within another piece of software. For these reasons we generally recommend that you use the FusionAnalytics AIR client.

Unlike Flash content viewed from within a browser, when you download and run an AIR application such as the FusionAnalytics AIR client, it is installed as if it were a desktop application. You can have a shortcut on your computer desktop and run it directly from there.

Back to the top

Executable

Executable	System Architecture
fusionanalytics-air-1.0.1.air	32-bit and 64-bit

Back to the top

Setup

Once AIR is installed, double click the FusionAnalytics AIR Client installer (.air file).

When the installer starts, you should click **Install** (Figure 1).

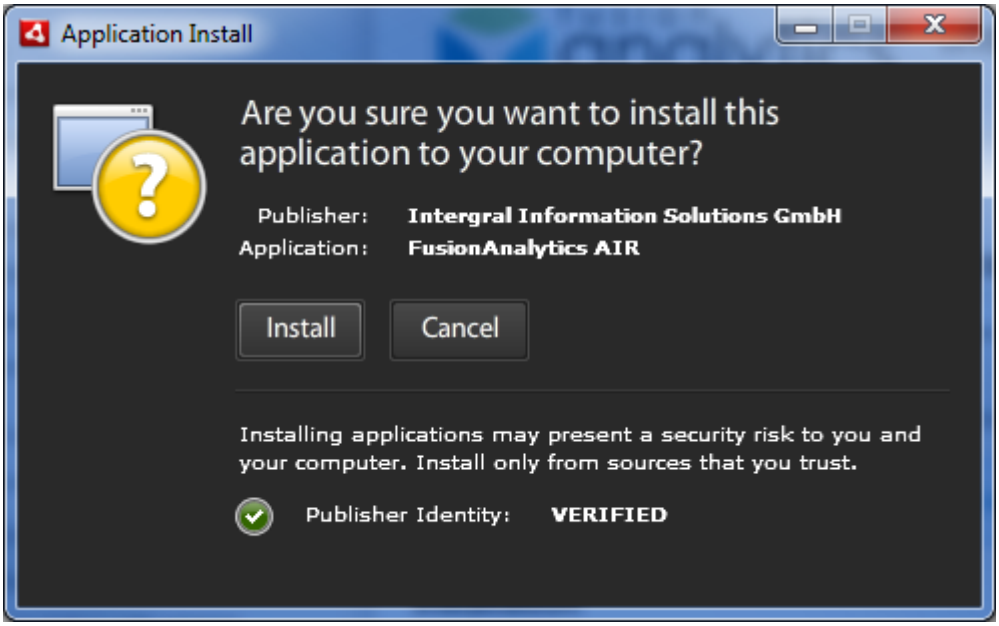

Figure 1: AIR Client Installer

The next screen (Figure 2) allows you to configure installation options.

It is recommended that you keep *Start application after installation* checked, as we will start configuring the client after it is installed.

If needed, you can change the location of the installation directory by clicking the folder icon.

Click **Continue** to carry on with the installation.

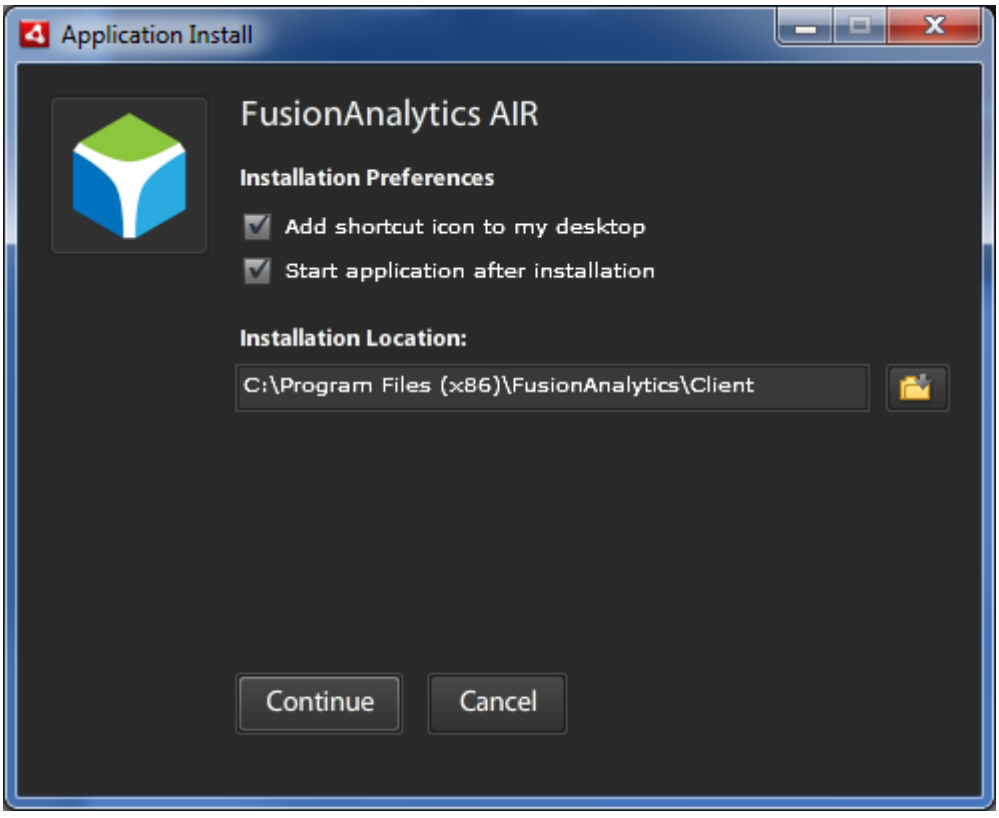

Figure 2: Installation options

When the installation has completed, the AIR Client can be opened by double-clicking the icon on the desktop named: **FusionAnalyticsAir**.

Back to the top

Uninstalling the Adobe AIR Client

To uninstall the FusionAnalytics AIR Client you need to follow the standard steps for uninstalling an application on your computer. For example, on Windows you would open the "Add or Remove Programs" dialog from the Control Panel. Scroll down to "FusionAnalytics AIR", highlight it and then click the Remove button. The applicaion will be automatically uninstalled from your system.

Back to the top

Next Steps

Accessing FusionAnalytics
Once you have FusionAnalytics installed, you are ready to start the configuration tutorial.

Quick Start Guide
Go back to the Quick Start Guide to review your next steps.

Installing FusionReactor

Installing FusionReactor

To get the most out of FusionAnalytics, it should be used with FusionReactor 4. Using them together will produce highly detailed log files for your server and allow you to analyze the results.

Visit www.fusion-reactor.com for more information and details on how to buy. There is a 10-day trial period, enabling you to check out all the functionality of the software before purchase.

Please see the FusionReactor documentation for help on installing FR. Once installed, FR will start producing logs suitable for importing into FusionAnalytics.

These can be shipped automatically to FusionAnalytics periodically using the **FusionAnalytics Connector**. See the section in this manual, with links to individual topics, or use **FRAM**, the **FusionReactor Administration Manager**, to set up FR for FusionAnalytics transfer, using handy popups.

Installing FusionReactor Extensions for ColdFusion

Installing FusionReactor Extensions for ColdFusion

FusionReactor ColdFusion Extensions is a plugin for FusionReactor, which can enhance the information that you receive in FusionAnalytics.

It is designed specifically for ColdFusion servers, allowing extra log files to be generated giving you a much higher level of detail when analyzing requests in FusionAnalytics.

Download for Free: FusionReactor Plugins

For a detailed installation tutorial, please see the FusionReactor documentation: FusionReactor Extensions for ColdFusion (FREC).

Uninstalling FusionAnalytics

Uninstalling FusionAnalytics

If you need to uninstall FusionAnalytics for any reason, you can either use the FusionAnalytics Uninstaller in the Start Menu *(for FusionAnalytics Server only)* or the Uninstall a Program option in the Windows Control Panel (Figure 1).

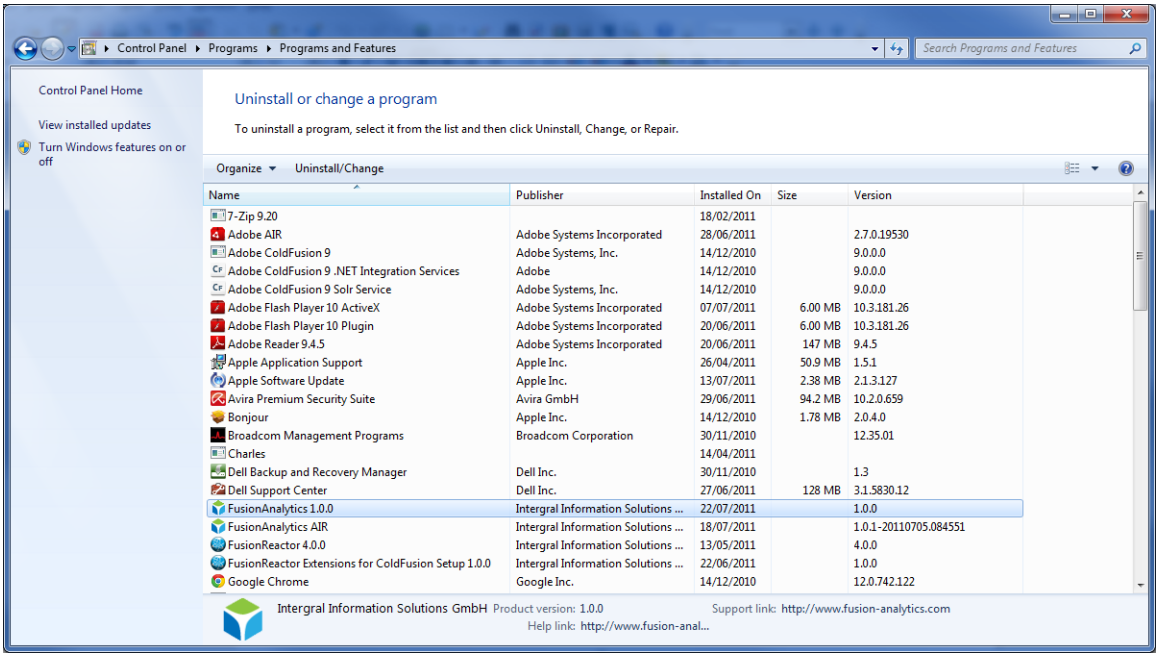

Figure 1: Uninstall a Program

The FusionAnalytics Server and FusionAnalytics AIR Client must be removed separately.

The database used by FusionAnalytics must also be removed separately, either by deleting it from the SQL Server Management Studio or uninstalling Microsoft SQL Server.

If you remove the Server, you will no longer have access to those applications in the Client.

If you remove the AIR Client, the applications will still exist and can be accessed in another client.

FusionAnalytics Server

Once the uninstaller has been started, you will see the FusionAnalytics Uninstall wizard (Figure 2).

If you are sure you want to remove FusionAnalytics DataCollector and FusionAnalytics DataServices, click **Next**.

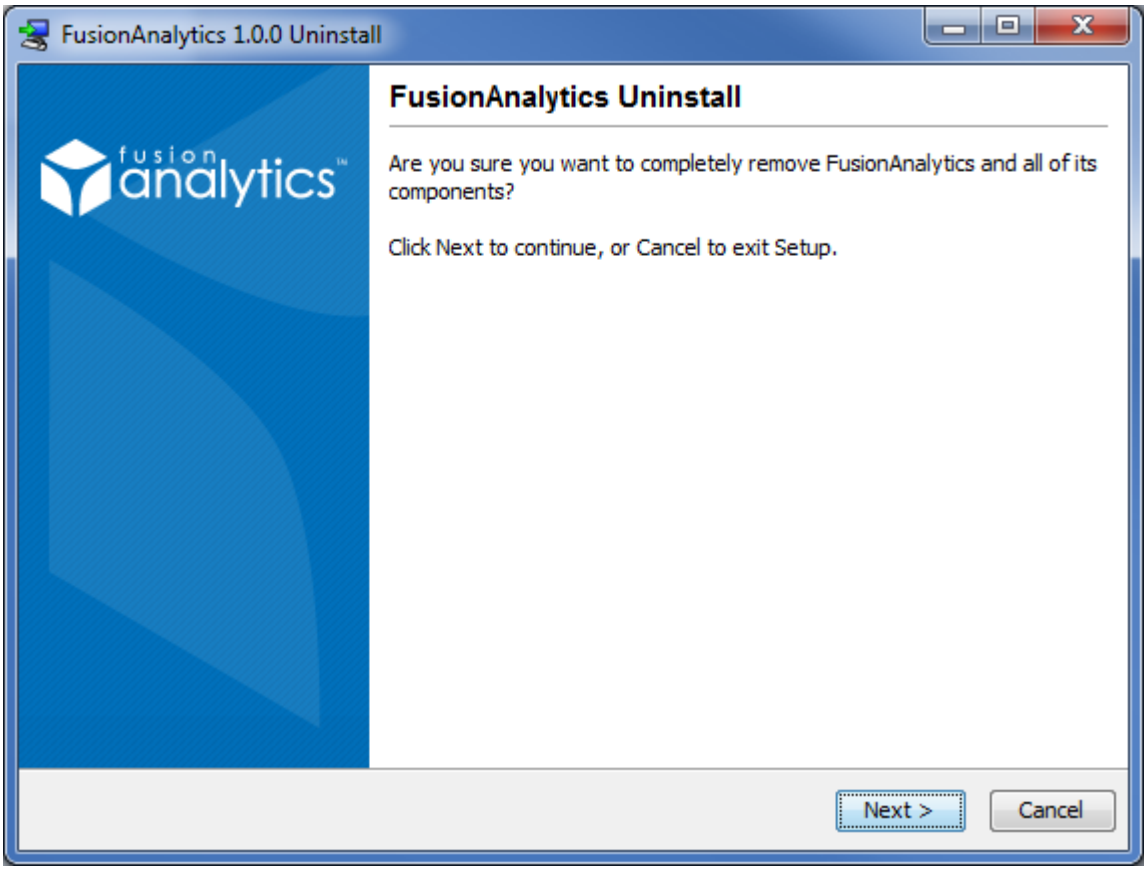

Figure 2: Uninstaller Wizard

After the uninstaller has finished, FusionAnalytics Server will have been removed from your computer and you will no longer have access to its applications in the client.

Click **Finish** to close the uninstaller wizard.

Back to the top

FusionAnalytics AIR Client

You can only uninstall the client from the Windows Control Panel (see above).

Once it has been uninstalled the applications will still exist and can be accessed from other clients.

Back to the top

On this page:

FusionAnalytics User's Guide

FusionAnalytics User's Guide

Overview

Introducing FusionAnalytics - After FusionAnalytics has been installed the User's Guide will help you getting started and up to speed with FusionAnalytics. This guide will discuss how FusionAnalytics is organized and how it can be used including detailed documentation on both

the Client and Server.

Next Steps

Getting Started

- How FusionAnalytics Is Organized
- Expandable and Extensible Platform
- Scaling and Quantization
- Exploring Data

FusionAnalytics Client

- Getting Started with the FA Client
- The FusionAnalytics Web Client
- The FusionAnalytics AIR Client
- Using the Table of Contents (TOC)
- Date Navigation
- Time Zones
- Filtering Data
- Data Visualizations
    - DataGrid Visualization
    - Chart Base Visualization
    - Time Block Chart Visualization
    - Placemark Chart Visualization
    - Sixway Chart Visualization
- Application Menus
- Using Tabs
- Using State Stores
- Context Menus
- Business Metrics
    - Reports
- Administrator
    - Manage Applications
    - FusionAnalytics DataServices
    - FusionAnalytics DataCollector

FusionAnalytics Server

- Server Structure
    - Directory Structure
    - FADC Applications
    - FADS Applications
- Server Configuration
    - FusionAnalytics Log Files
    - Windows Service Configuration
    - FusionAnalytics DataCollector Settings
    - Setting Up FusionAnalytics with HTTPS
- FusionAnalytics DataCollector (FADC)
    - Applications (FADC)
        - Application Summary (FADC)
            - Application Configuration (FADC)
            - Application Detail (FADC)
        - Upload an Application (FADC)
        - Upgrade an Application (FADC)
        - Upload Log or ZIP Files (FADC)
        - Import History (FADC)
    - User (FADC)
        - Change Password (FADC)
    - System (FADC)
        - Email Settings (FADC)
        - Runtime Info (FADC)
        - License Info (FADC)
        - Logout (FADC)
    - DCML (FADC)
- FusionAnalytics DataServices (FADS)
    - Applications (FADS)
        - Application Summary (FADS)
            - Application Configuration (FADS)
            - Application Detail (FADS)
            - Application Scope (FADS)
        - Upload an Application (FADS)
        - Upgrade an Application (FADS)
    - Users (FADS)
        - Users Summary (FADS)
            - User Mappings
    - Scheduled Tasks (FADS)

Getting Started

Getting Started

Overview

This document will help you to explore and evaluate FusionAnalytics, a tool to visualize log data.

We'll introduce the architecture of the system, show you how to use the client and server packages, explain the terminology, and touch briefly on the extension points.

Let's get started by explaining what FusionAnalytics is, and what it can help you do.

Next Steps

How FusionAnalytics Is Organized

Expandable and Extensible Platform

Scaling and Quantization

Exploring Data

How FusionAnalytics Is Organized

How FusionAnalytics Is Organized

FusionAnalytics is organized into three main areas, listed below.

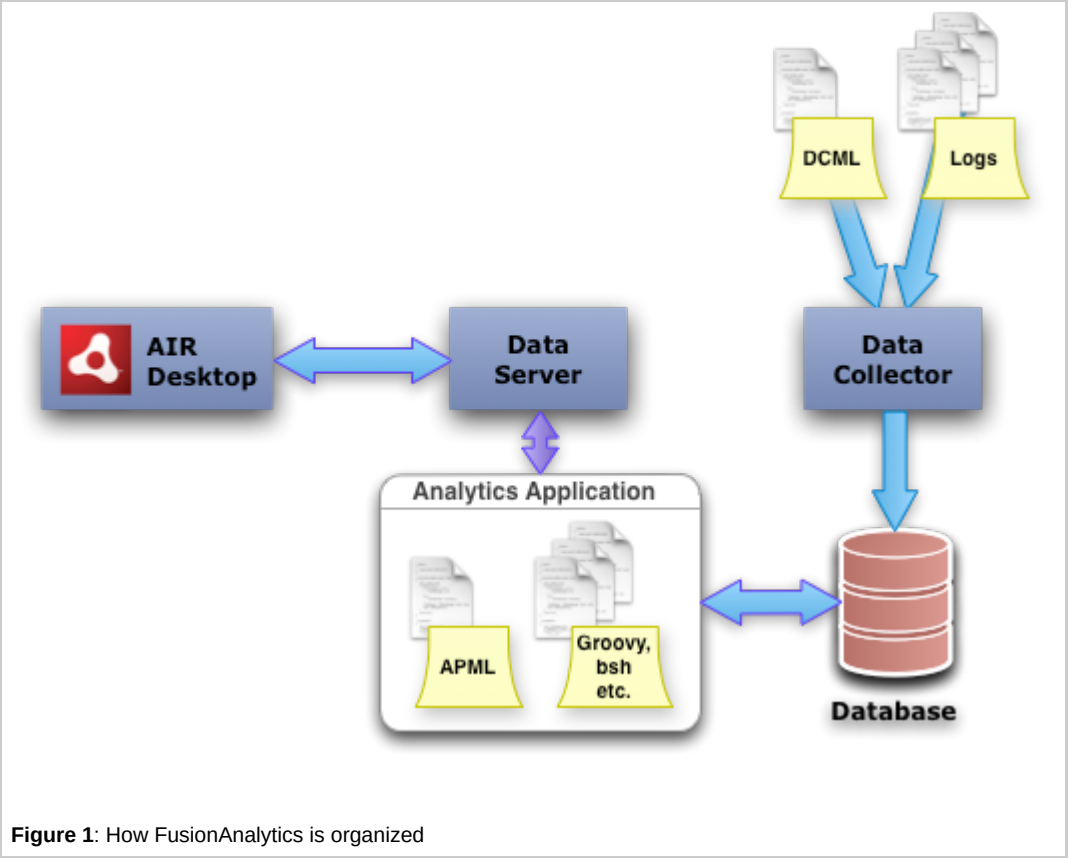

Figure 1: How FusionAnalytics is organized

FusionAnalytics DataCollector

The **DataCollector** is responsible for gathering information, transforming it, and writing it to a database. The DataCollector provides fine-grained control over what raw information is loaded, and how it is transformed. It's also able to aggregate data to speed up analytics queries.

Learn more...

Back to the top

FusionAnalytics DataServices

DataServices hosts **Analytics Applications**, supplying services like caching, persistent storage, database connections and reporting. DataServices is responsible for communicating with the **AIR Desktop**, answering queries for analytical information, and managing authentication and sessions.

Learn more...

Back to the top

FusionAnalytics AIR Desktop

The **AIR Desktop** is the user interface to the FusionAnlytics system. This is the component which allows you to navigate and control an Analytics Application, and displays the generated analytical information.

Learn more...

Back to the top

FusionAnalytics Applications

FusionAnalytics applications are applications hosted by the DataServer and/or DataCollector. They can for example load FusionReactor data into a database via the DataCollector to provide insight and analytics on that data via the DataServices. FusionReactor is able to store many requests online for an immediate diagnostic view about the health of a server, but is not able to provide information about long-term data – this is to ensure that FusionReactor is small and fast and doesn't provide much overhead when installed.

FusionReactor *can*, however, generate extensive log data. This data can be mined by the FusionAnalytics DataCollector and stored in a database. Together with the advanced functionality provided by FusionAnalytics, this data can then be used to provide valuable insight into long-term trends.

The high-resolution data is naturally still available - you can zoom down from months to milliseconds using our intuitive timeline control.

Learn more...

Back to the top

Expandable and Extensible Platform

Expandable and Extensible Platform

FusionAnalytics is a platform for visualizing **any** data. The DataCollector and DataServices don't place any restrictions on the type of data in use.

In order to specialize the platform for a particular system, we devised the concept of FusionAnalytics **Applications**. These run within the FusionAnalytics platform, and adapt it to a specific type of data.

There must be two applications for every type of data you want to visualize:

- A **DCML application** (DataCollector Markup Language) which tells the **DataCollector** how to load and interpret log data, and into which database tables it should be loaded, and
- An **APML application** (Analytics Provider Markup Language) which tells **DataServices** which views are available in the client, and how to generate data from the database to serve those views.

Both application packs, DCML and APML, are based on an XML syntax. Both can be extended using scripting languages.

Scaling and Quantization

Scaling and Quantization

Most views in FusionAnalytics are navigated using our **TimeLine** control. This innovative component allows you to zoom in and out, as well as scrolling forwards and backwards in time. The control supports a huge range---you can zoom in to a millisecond or out to a decade.

Naturally, the data must also support this.

The Analytics for FusionReactor application makes this possible by generating sets of data at different resolutions. We call this **quantized** data since it has been combined and aligned to one of several preset values: one minute, one hour, one day and 30 days.

Most of the views we provide automatically switch to the most relevant data set as you zoom in and out, always providing the quantization level which fits best, and never overloading the visualizations with too many data points. A dual-level (Client and DataServices) cache further speeds up data delivery, using a dynamic hinting process to pre-fill what we think you'll want to look at next.

If you're only interested at long timescales, the high-resolution (millisecond-scale) data can be purged after it's been quantized, saving space in the database.

Now let's have a quick look over the architecture of the system, and then get FusionAnalytics installed and importing data.

High Resolution Data

High resolution data is stored differently depending on where the data has come from:

- For requests and JDBC, all information about every request is stored in the database.
- For resources (cpu, memory etc) FusionReactor takes samples at a set time interval (default 5 seconds). The current status and details about the resource is stored at each sample.

High resolution data is very detailed information about your server. All of the information taken from every sample/request is stored in the database allowing you to see accurate and detailed statistics. While this is very useful, it consumes a large amount of disk space and so we

recommend only storing high resolution data for a short period of time for extremely accurate deep server analysis.

Low Resolution Data

Low resolution data has been quantized to greatly reduce the amount of disk space that is required. As explained above, data is quantized and aligned to one minute, one hour, one day and 30 days which can then be used when analyzing data in the FusionAnalytics perspectives. This information is useful if you want to analyze your server over long periods of time.

How the quantizer works

When a sample is taken by FusionReactor or a requests are sent to the server, the information gathered is stored in the logs. This is very precise and there is a lot of data gathered over a short period of time. At some point after the logs are transferred from FusionReactor to FusionAnalytics, the quantizer runs, gathering data for the appropriate alignments listed above and take averages creating compressed copies of the high resolution data. These averages are stored as quantized data.

Take a look at the image below. This shows an example of high resolution resource data taken over 1 hour. When the quantizer runs it takes an average of these samples and stores it as the quantized value for that time frame.

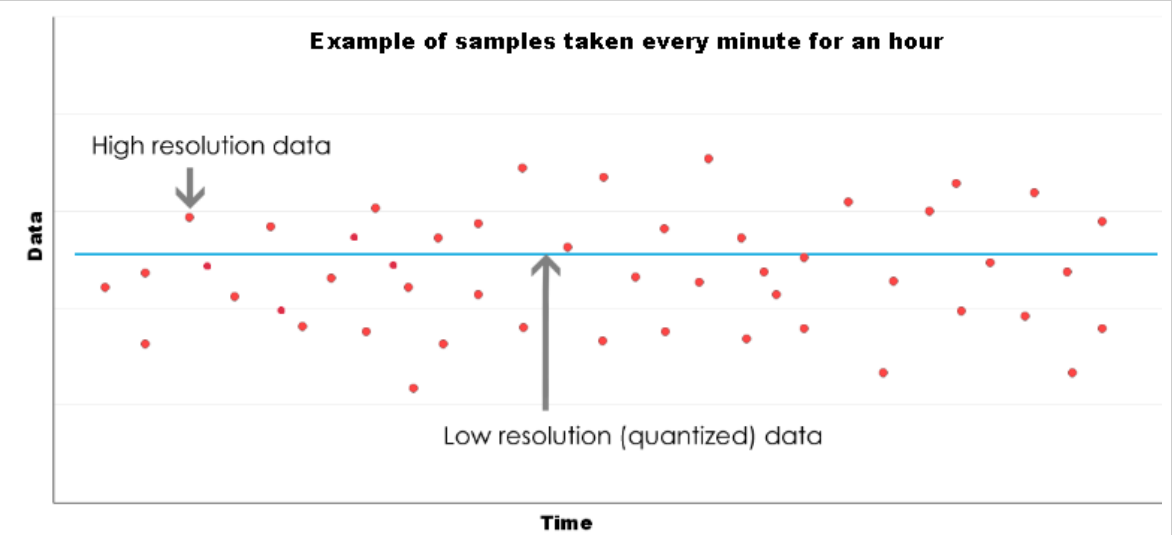

Figure 1: How the quantizer works

Data Truncation

If you choose to keep high resolution for one week, then any information over a week old will be deleted. This means you will always the most recent high resolution data for the previous 7 days. Similarly if you choose to keep low resolution data for one year you would always have low resolution stored for the previous 356 days.

In FusionAnalytics you can choose how long you wish to store High Resolution and Low Resolution data for each FusionAnalytics for FusionReactor application. Please see the Application Configuration (FADC) for details on how to do this.

Exploring Data

Exploring Data

FusionAnalytics offers different tools and techniques that can be used to help Explore Data. Some of these tools include:

Filtering Data

FusionAnalytics gives you a lot of data so it is useful to be able to focus in on just the information you want. The type of filters available to you will differ depending upon the type of perspective. You will see a tab at the bottom of the perspective for each type of filter available (if any)... For more information please see Filtering Data.

Back to the top

Date Navigation

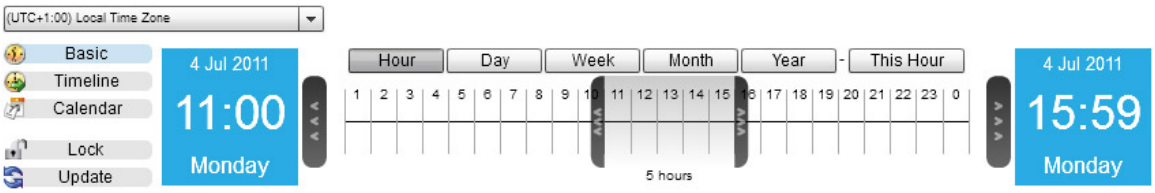

Most of the perspectives in FusionAnalytics are time based so, if you want to explore your data further, you will want to change the span of time that you are currently viewing. At the bottom of time-based perspectives you will see a tab named "TimeLine". By default the TimeLine will be in Basic Mode... For more information please see Date Navigation.

Back to the top

Time Zones

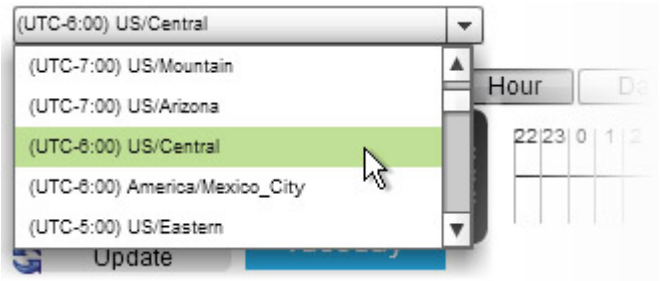

FusionAnalytics allows you to view graph and table data in the Time Zone of your choice. By default, this will be set to your local Time Zone. (ie, the Time Zone in which the FusionAnalytics Client is running.) If you want to change your viewing Time Zone then you can select a new Time Zone from the drop down box in the top left of the Date Navigation area... For more information please see Time Zones.

Back to the top

Context Menus

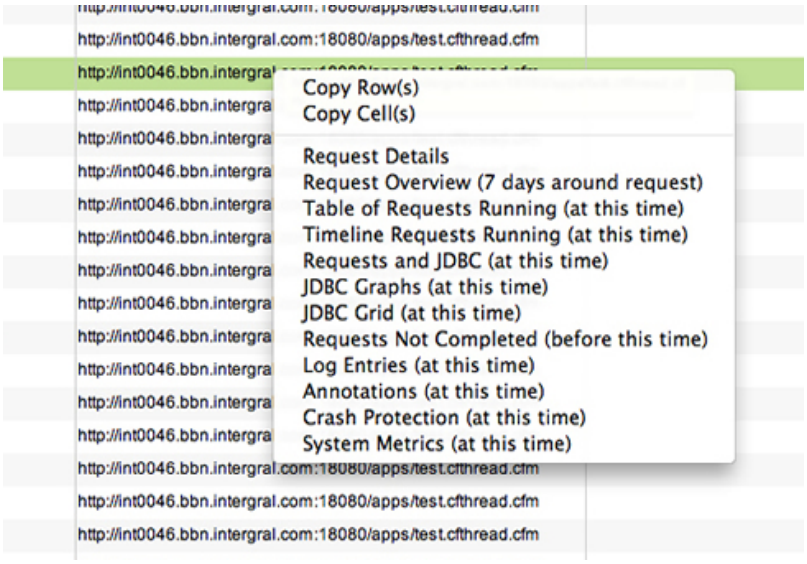

Context Menus are used throughout the Client to help navigation and link data within applications. Context Menus are attached to and can be used with most Data Visualizations... For more information please see Context Menus.

Back to the top

Using Tabs

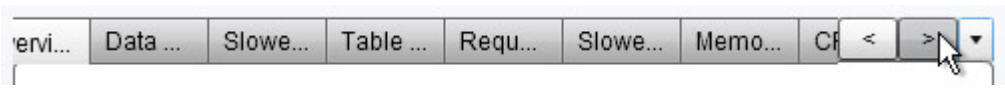

Whenever you click on an item from the Table of Contents (TOC) a new perspective will be opened in the content area. This perspective will be opened in a new tab. Tabs appear along the top of the content area and can be interacted with in various ways. The simplest way to switch between tabs is simply to click on the tab you want to view. The content of that tab will then be made visible.... For more information please see Using Tabs.

Back to the top

Using State Stores

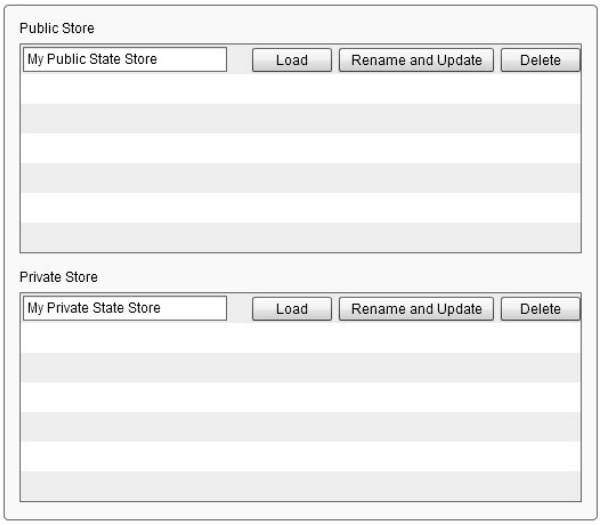

The State Store feature is designed to allow you to store the current state of a perspective so that you can return to it at a later point in time. The State Store will save the current date range as well as any currently defined filters.... For more information please see Using State Stores.

Back to the top

Data Visualizations

FusionAnalytics provides different Data Visualizations that are used to render data. For example:

- The DataGrid Visualization: used to render data in a "list" based format.
- The Chart Base Visualization: used to display / represent data within a chart / graph making data interpretation more "visible"
- The Time Block Chart Visualization: used to render a large amount of data points providing a more detailed visualization (picture) of data at any one time

For more information on the above Visualizations (and more) please see Data Visualizations.

Back to the top

FusionAnalytics Client

FusionAnalytics Client

Overview

This section will discuss the FusionAnalytics Client in detail including information on how to set the clients (both Web and AIR) up and how to add, launch and use applications within the clients. This section will further discuss the different components and visualizations included in the client and how they can be used (including the best ways) to use the application (tools and components) to explore and analyze data.

Next Steps

Getting Started with the FA Client

The FusionAnalytics Web Client

The FusionAnalytics AIR Client

Using the Table of Contents (TOC)

Date Navigation

Time Zones

Filtering Data

Data Visualizations

- DataGrid Visualization
- Chart Base Visualization
- Time Block Chart Visualization
- Placemark Chart Visualization
- Sixway Chart Visualization

Application Menus

Using Tabs

Using State Stores

Context Menus

Business Metrics

- Reports

Administrator

- Manage Applications
- FusionAnalytics DataServices
- FusionAnalytics DataCollector

Getting Started with the FA Client

Getting Started with the FA Client

Login

After launching the The FusionAnalytics Web Client the first screen you will be taken to will be the login screen. Enter your log in credentials (Username and Password) here and click Login to enter the client.

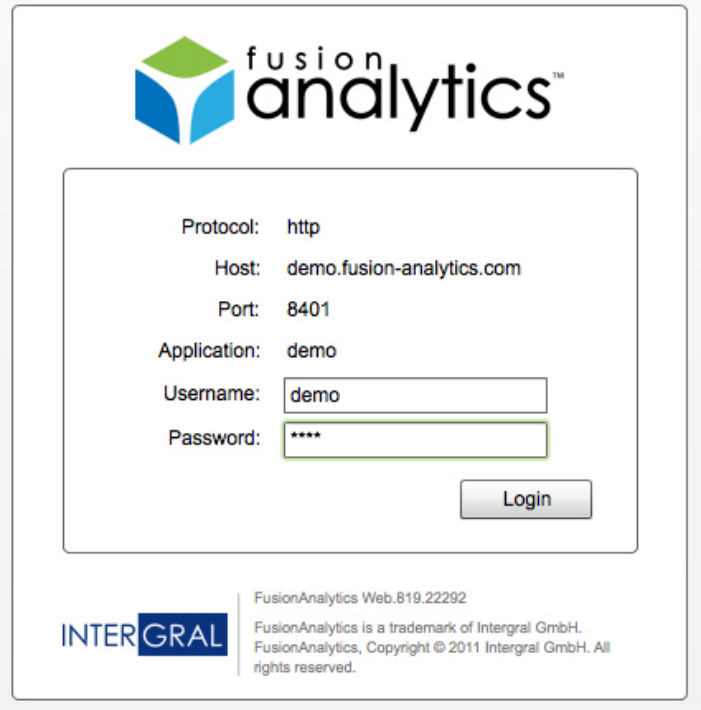

Figure 1: Client login panel

Back to the top

Default Perspective

In FusionAnalytics, perspectives are essentially different views that can be displayed on the screen. Once you have entered the client, the first perspective you will see is the "Overview Perspective".

The Overview Perspective is the default Perspective and (like most Perspectives in the application) is date/time-based. The perspective contains two visualizations, FusionReactor System Metrics and Server Summary, providing an instant insight into server utilization including :

- Breakdown server utilization by hostname (per project/customer/language)
- Total Response Size
- Average Execution Time - per site
- Total Execution Time

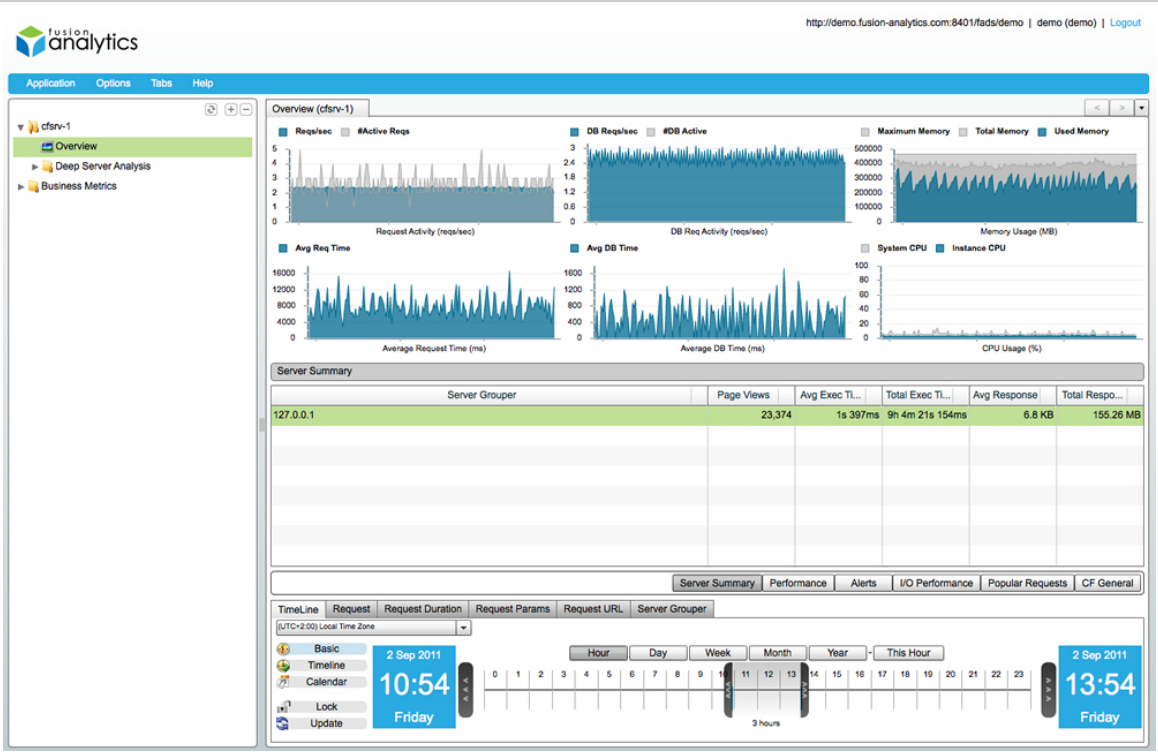

Figure 2: Default Perspective

You will notice that the application layout is split into two columns:

- Table Of Contents (Left side) - This navigation menu allows you to navigate around FusionAnalytics and access different perspectives and features.
- Perspective Canvas (Right side) - This is the panel where perspectives are displayed. This is where your server information is visible and the timeline control can be adjusted.

Date Navigation

At the bottom of the Overview Perspective you will see a tab named "Timeline" containing the Date Navigation component :

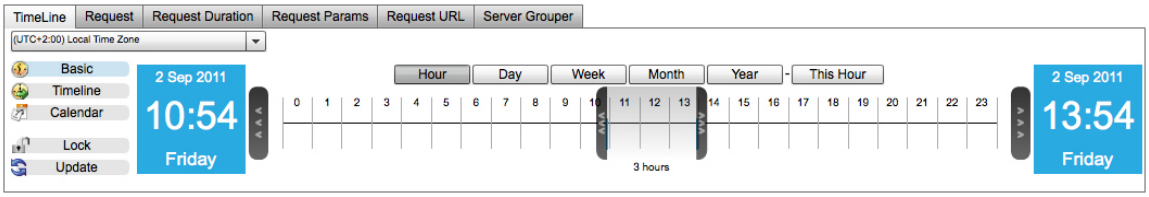

Figure 3: Timeline Control

You can use the Date Navigation component to explore data further allowing you to change the span of time that you are currently viewing. By default the Date Navigation component will be in **Basic Mode**. In this mode you can select a unit of time (Hours, Days, Weeks, Months or Years) and then select a range of those units to be the span of time you are interested in seeing analytics data for. Clicking on one of the "unit" buttons will switch to that unit of time. You can also change the mode (from Basic to **Timeline Mode** or **Calendar Mode**) if preferred by switching between the three mode toggles on the left of the control.

Back to the top

Locking and Updating

⚠ You may have noticed that even though you have been changing the time span, the data within the current perspective has NOT yet been updated. At any point you can click the "Update" button and your perspective will be refreshed showing the currently selected span of time. Even if you have not changed the current time span, you can click the "Update" button to get the most recent data from FusionAnalytics.

There are many situations in which you would want graphs to automatically be updated as you scroll through time. In this case you can lock the TimeLine. Click the "Lock" button to toggle this feature on and off. Whilst locked, your perspective will be updated whenever you change

the currently selected timespan.

Figure 4: Lock and Update Buttons

> **Note**: If you have the Timeline "Locked" whilst browsing through data then FusionAnalytics will constantly be requesting new data to show. In this case you may find that the performance is slower than having the Timeline "Unlocked". In general, we recommend that you have the Timeline "Unlocked" and use the "Update" button to refresh data once you have navigated to the point in time you are interested in.

Back to the top

Reports

Using the Table Of Contents (left column) navigate to "*Business Metrics > Reports*" to open the Reports perspective. The reports Perspective includes a DataGrid containing all of the reports that are relevant to the Application.

> **Note**: FusionAnalytics will ONLY show reports which have been generated on the dates which you have selected in the Date Navigation component.

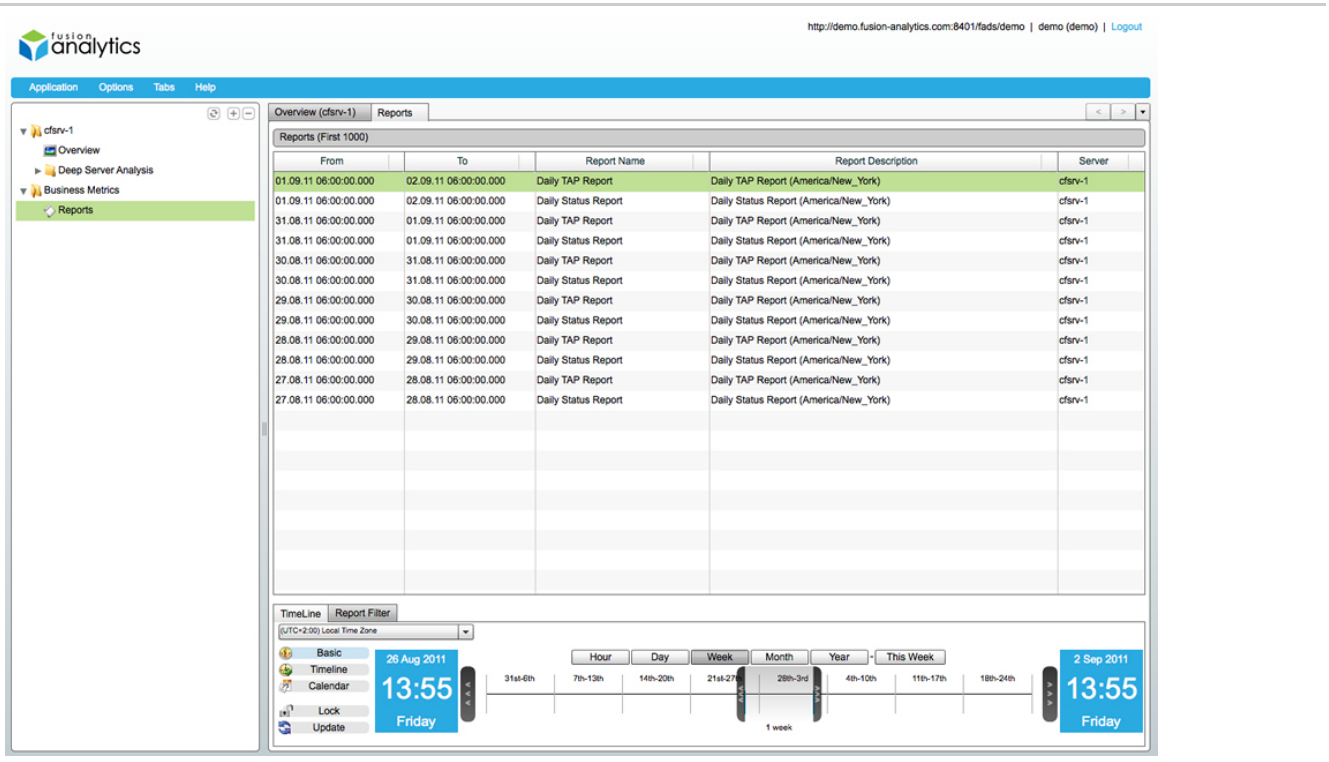

Figure 5: Reports Perspctive

To open a report, either double click on the report you would like to open or right click on the report and select "Open Report".

Note: The first time the report is opened Authentication may be required to view the reports (within the Web Client). If so simply enter the credentials for the application and click OK:

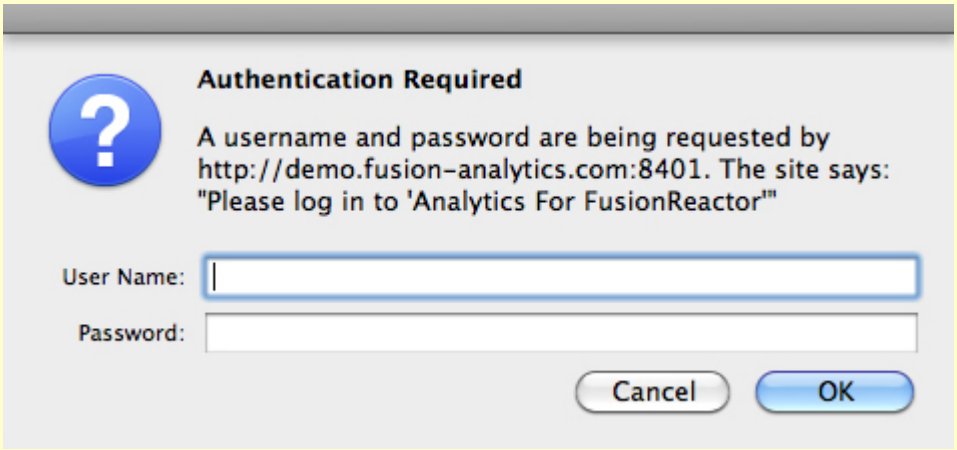

Once authenticated, the selected report (in this case TAP) will then open in a new tab:

Tap Report

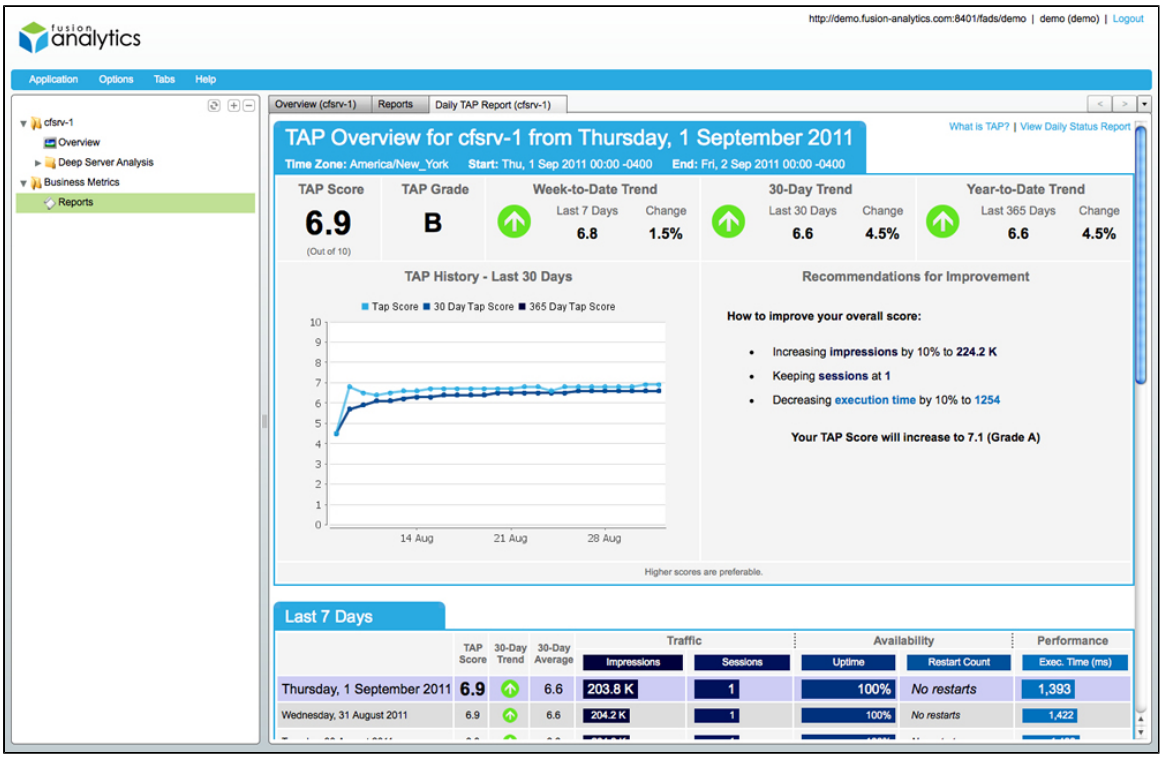

Figure 6: TAP Report

TAP Report: The objective of the Traffic, Availability and Performance (TAP) index is to give application & program managers, business owners and solution stake holders, a less-technical orientated metric on their servers. FusionAnalytics makes suggestions on how to improve your TAP score, as well as measuring and comparing TAP performance over time.

To open another report simply repeat the same process by going back the the Reports Perspective (within the Tab Navigator) and selecting and opening a new report:

Daily Status Report

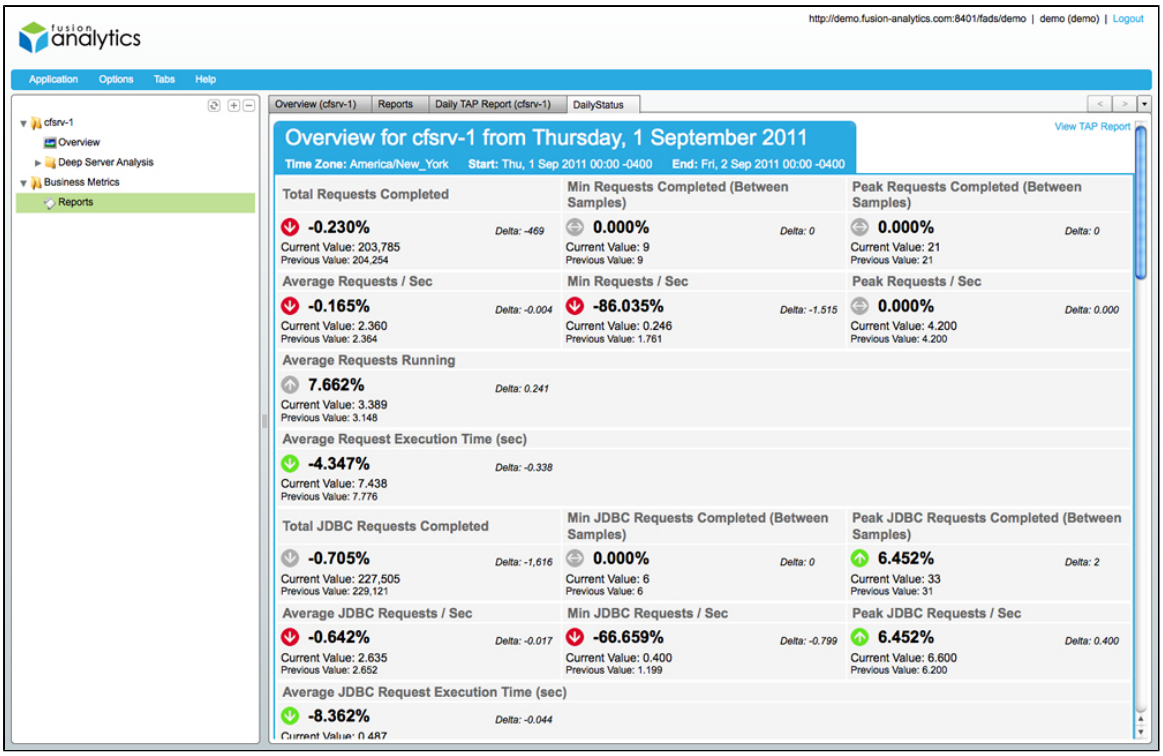

Figure 7: Daily Status Report

Daily Status Report: The Daily Status Reports are intended for use by technical developers, project leads as well as project managers and system administrators. The Status Report contains a complete breakdown of your application / server performance as well as specific slow running requests which need improvement.

Back to the top

Deep Server Analysis

At the heart of FusionAnalytics is the wealth of information gathered and transferred from FusionReactor. This instrumentation includes requests, SQL statements, response times, CPU and memory information, as well as URL parameters, heap memory usage and JMX statistics. FusionAnalytics takes all of this information and builds a series of combined visual perspectives, allowing you to visualize all transactions, SQL queries, data transfer, server events etc. taking place over time, so you can build a complete historical view of what happened on your server. This data is essential to support investigative analysis, continuous server analysis and application improvement. With over 80 different analytic views and reports, FusionAnalytics gives you unlimited insight into exactly what's happening on your server.

Note: All deep server related perspectives are listed under the "Deep Server Analysis" menu item within the Table of Contents.

For Example: Using the TOC Navigate to "*Deep Server Analysis* > *Request* > *Slowest Request*" to open the Slowest Request Perspective:

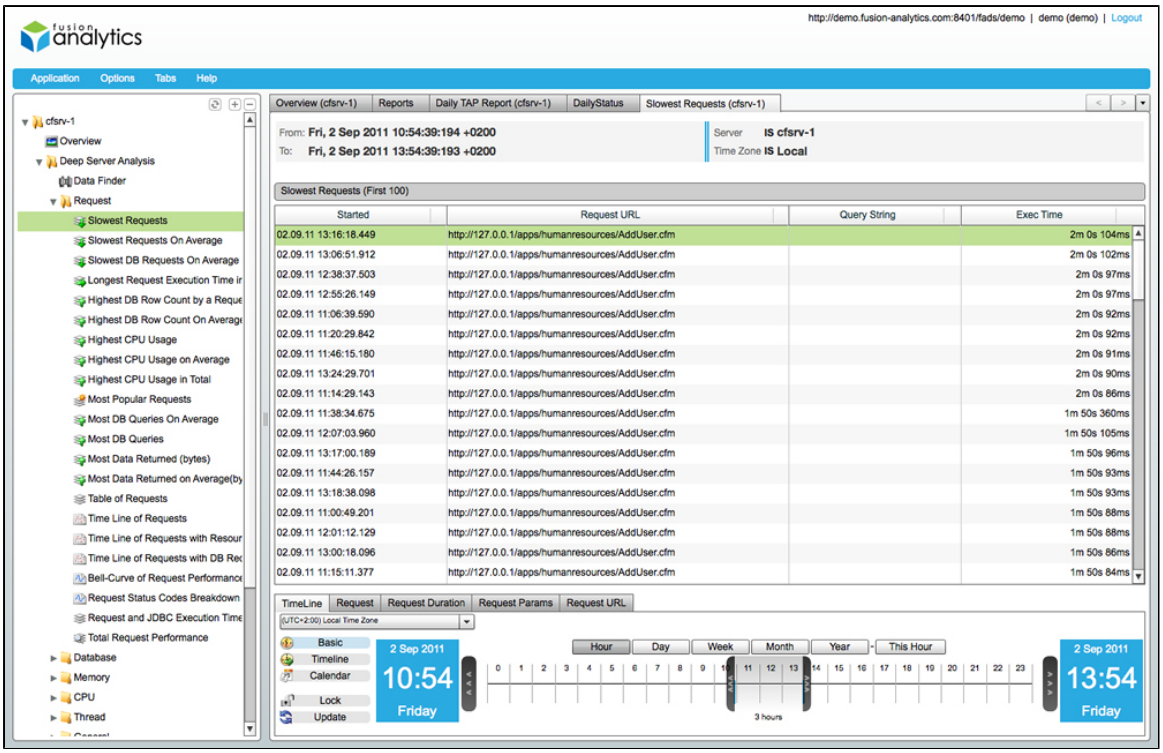

Figure 8: Slowest Requests Perspective

Within the Slowest Requests Perspective, the top 100 slowest requests (within the the specified time frame) are listed allowing detailed analysis to be performed to help the identification of key areas that can be optimized. You can see more in depth information for each request by right clicking on a request and selecting the "Request Details" menu item.

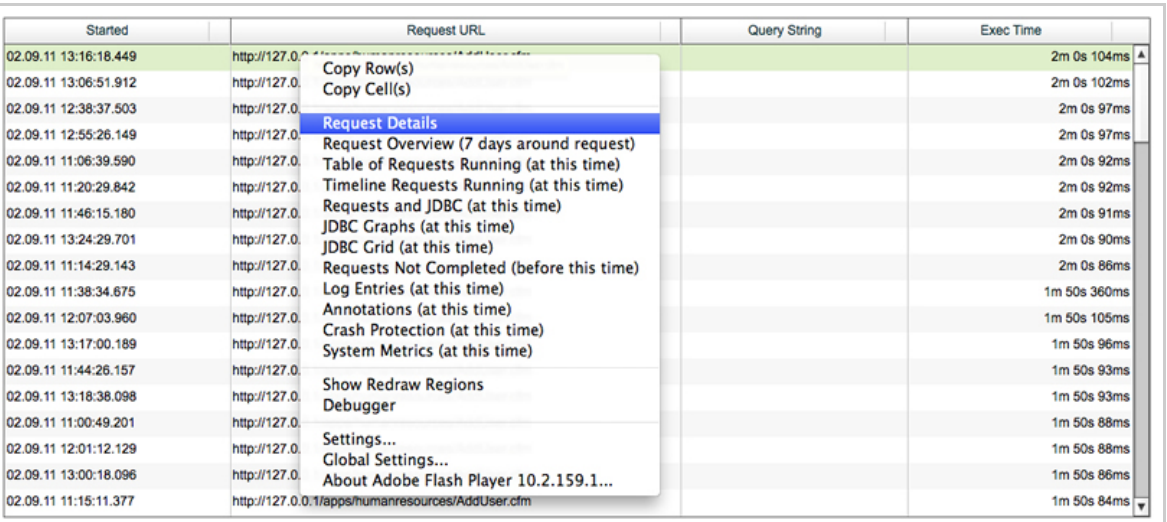

Figure 9: Slowest Requests Context Menu

This will open the Request Details perspective for the selected request in a new tab:

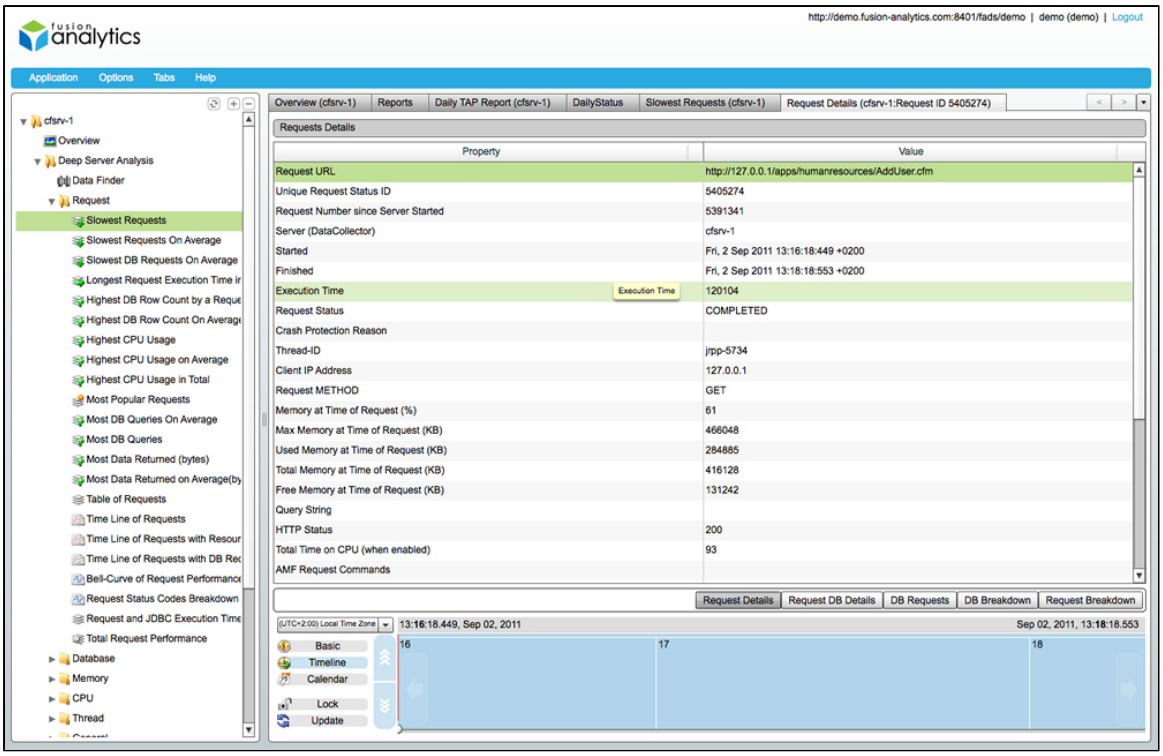

Figure 10: Request Details Perspective

The Request Details perspective shows all the details of the actual (selected) request, including *URL*, *Start/Finish time*, *Status code*, *execution time*, *thread-ID*, *request method*, *memory metrics* (at time of execution) etc.

The perspective offers further useful information related to the request including *Request DB Details*, *DB requests*, *DB Breakdown* and *Request Breakdown* which can be accessed using the toggle buttons that are rendered below the Request Details DataGrid. For example, to see where the database time is being spent on a query by query basis relevant to the the selected request click the "DB Requests" toggle button just above the timeline.

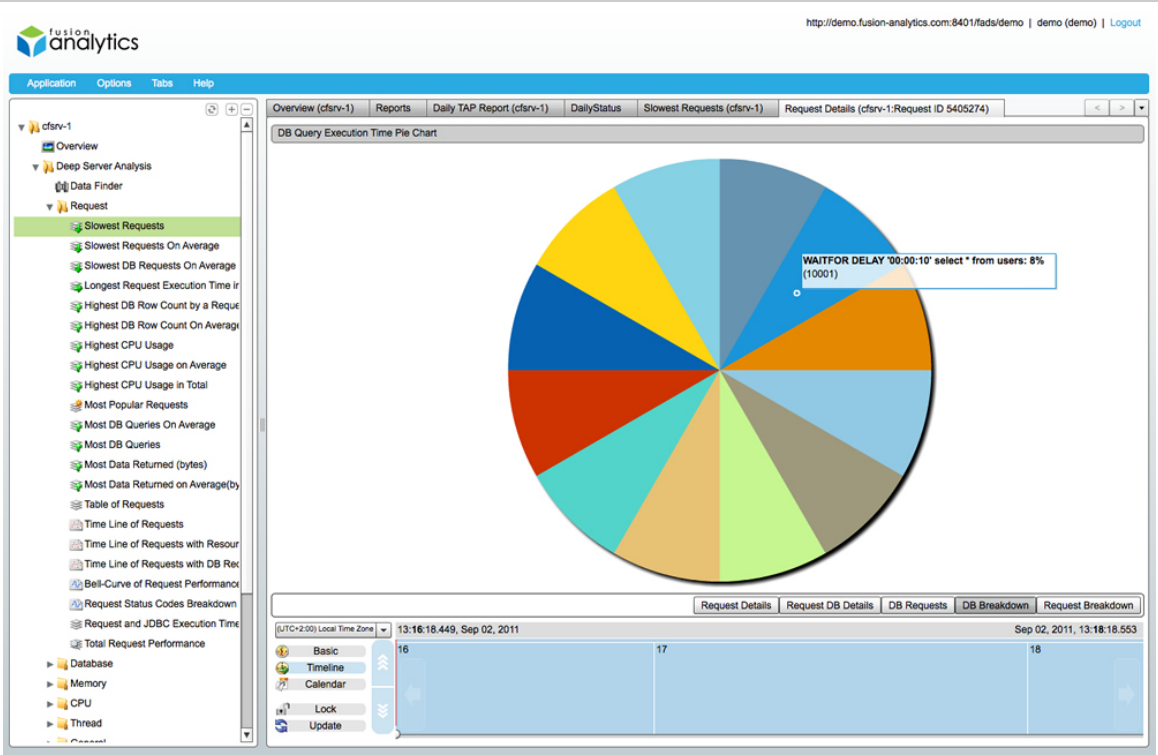

Figure 11: DB Query Execution Time

With the database breakdown tab, mousing over a segment of the pie chart will reveal the actual SQL statement being called.

To compare processes CPU usage and process memory usage, again using the Table of Contents navigate to Deep Server Analysis > CPU > CPU / Memory Usage to open the CPU and Memory usage Perspective:

Figure 12: CPU/Memory Usage Perspective

Remember: If you want to change the time span at any time you can do so at any time using the Date Navigation Component

Back to the top

On This Page:

TOC and Content Layout
The application canvas is split into two columns. The TOC is rendered in the left column and the Perspectives (content) are loaded in the right column:

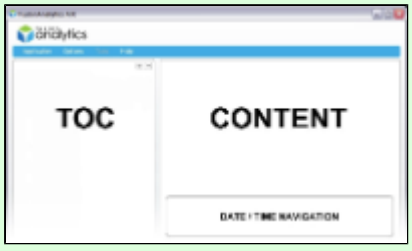

The FusionAnalytics Web Client

The FusionAnalytics Web Client

You can access the Web Client from FusionAnalytics DataServices (FADS) or by pasting the application URL directly into a browser (if known).

If you open FADS then, by default you will be taken to the application overview page. For each application defined within FusionAnalytics (listed under "Applications" > "Application Summary") you will see a Launch button.

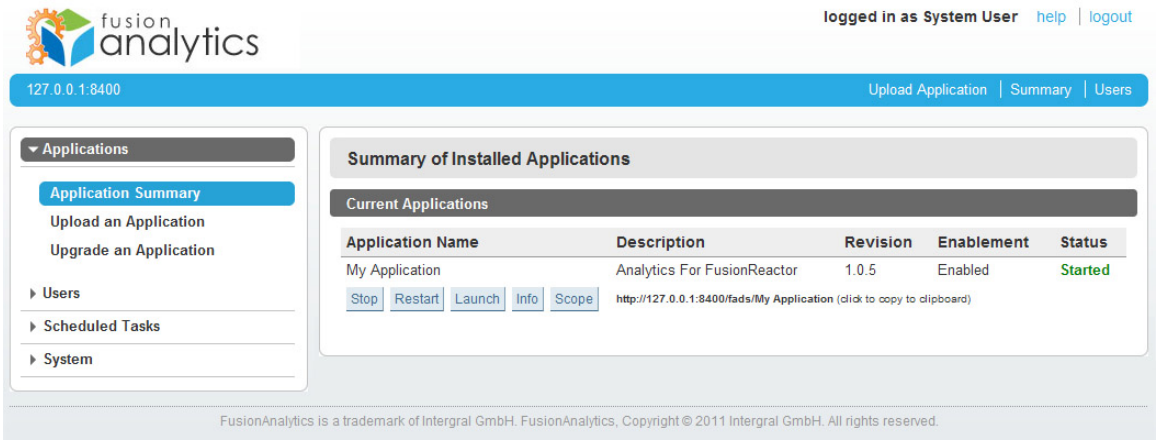

Figure 1: FusionAnalytics DataServices (FADS)

If you click on this Launch button then the Web Client will be opened in a new browser window.

You will also see the application URL which can be copied and used to access the Web Client directly by pasting the URL into a web browser or alternatively into the AIR Client.

> ⚠ This is the same (first) screen that you will see if you have entered the application URL directly into a browser. An Example application URL is:
>
> - http://serverName:8400/fads/applicationName

After launching an application the first screen you will be taken to will be the Login screen. Enter your log in credentials (Username and Password) here to access the rest of the client:

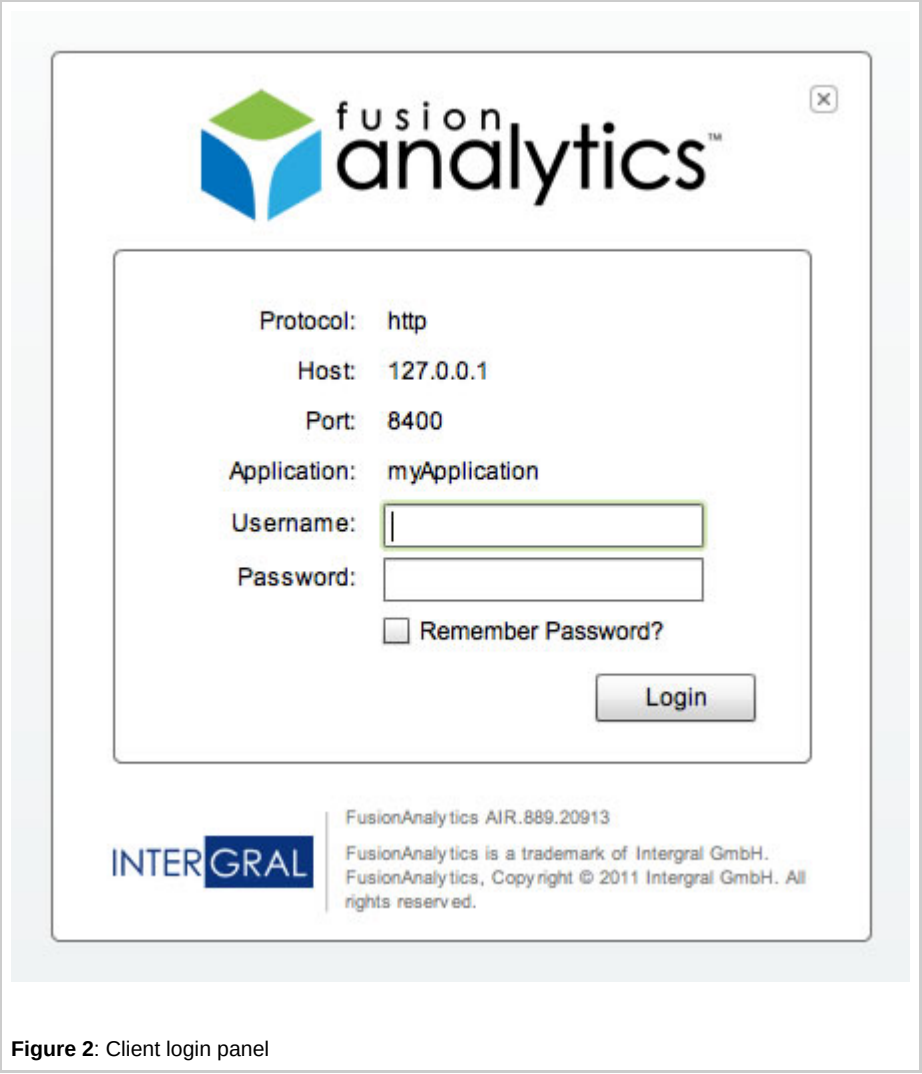

Figure 2: Client login panel

Back to the top

Launching FADS Applications
For more information on FADS Applications and launching the Applications using the Web Client / browser See Applications (FADS).

Next Steps

The FusionAnalytics AIR Client
Learn how to access multiple different FusionAnalytics Applications using the FusionAnalytics AIR Client

FusionAnalytics DataServices (FADS)
Learn about FADS and how it can be used in relation to System (Settings, Sessions, Controllers, Jobs ...), Applications (Summary, Configuration, Detail, Scope ...), Users and Scheduled Tasks (both System and User tasks).

The FusionAnalytics AIR Client

The FusionAnalytics AIR Client

Unlike the Web Client which you launch directly from FusionAnalytics DataServices (FADS), the FusionAnalytics AIR Client allows you to access multiple different FusionAnalytics Applications (even running on different servers) from a central location.

Adding Applications

When you first launch the AIR Client, you will be taken to the Server summary screen. This will initially be empty so you will need to add a link to an existing FusionAnalytics Application. To get this URL you need to go to the FusionAnalytics DataServices screen. By default you will see the Application Summary page.

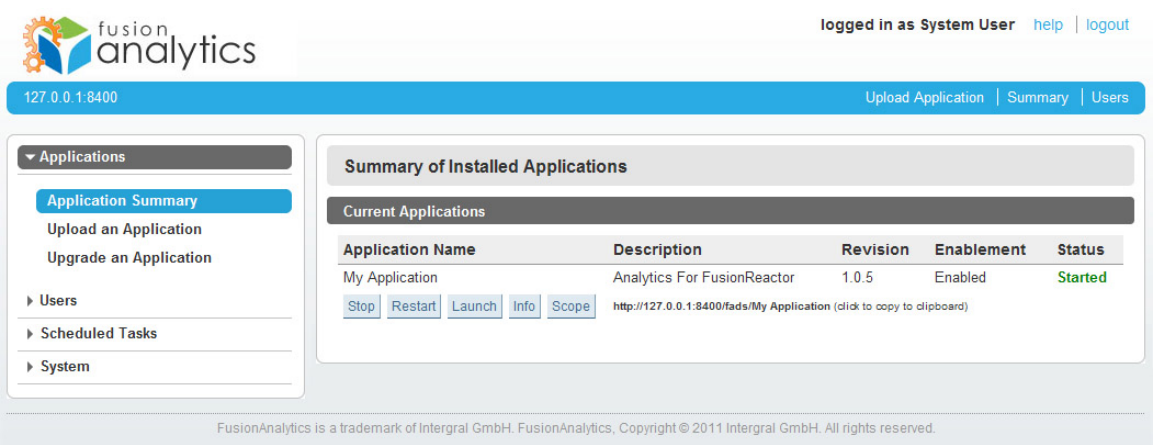

Figure 1: FUsionAnalytics DataServices (FADS)

Locate the application that you want to add to your AIR Client and you will see a URL displayed. If you click the URL then it will be copied to your clipboard.

Figure 2:click to copy to clipboard

Once you have copied the Application URL, return to the AIR client. At the top of the Server Overview screen you will see an input box with a button to the right saying "Add Application". Paste the Application URL into the text box and click the button.

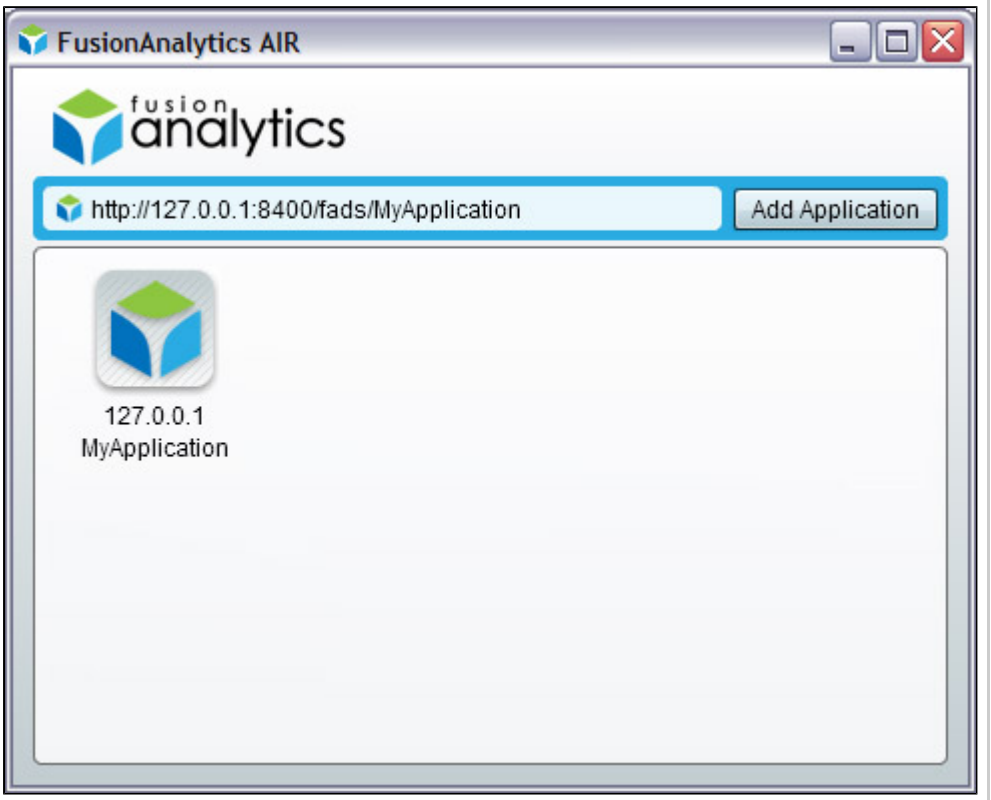

Figure 3: Paste application URL into the FusionAnalytics AIR client

A new icon will appear in the main screen area representing the application.

Back to the top

Editing / Deleting Applications

If you want to edit or delete this application link then you can right click on this icon.

Figure 4: Editing / Deleting Applications

Note: when you mouse over / highlight an application icon, you will see a small "delete" button in the top right of the application icon. You can also delete / remove an application from the AIR client by clicking this button *(this will have the same effect as clicking the "delete" menu item from the icons context menu*).

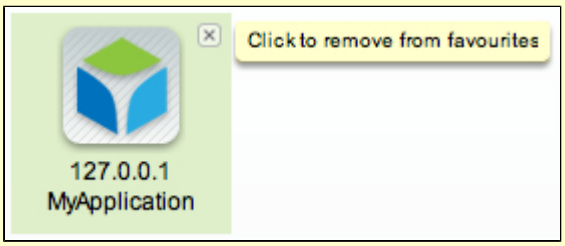

Selecting Delete will remove the icon from your AIR Client. (Note: This will have no effect on the actual application.) If you select Edit then the Edit Dialog will appear. Here you can change the details of this application link.

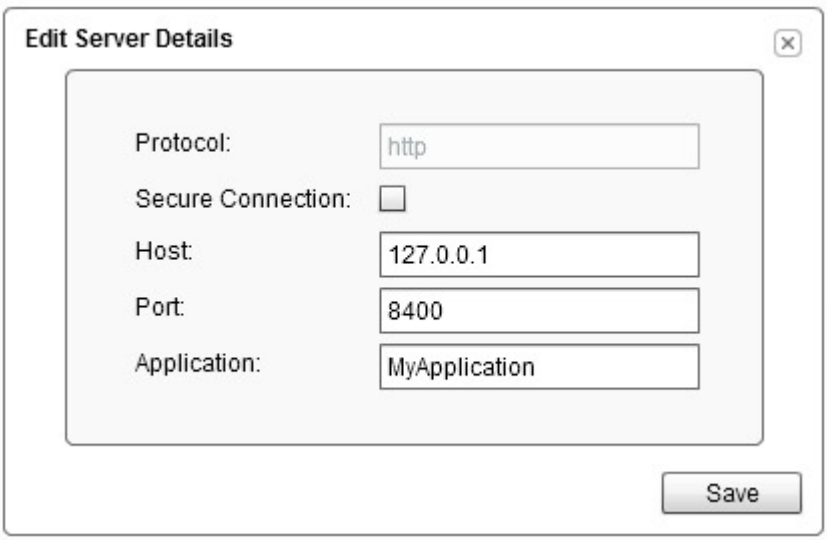

Figure 5:Edit Server Details panel

Back to the top

Logging into Applications

To log into a FusionAnalytics application from the AIR client, simply double-click on the requested server icon and you will be taken to the Login Page for that application. Enter your Username and Password here to access the rest of the client.

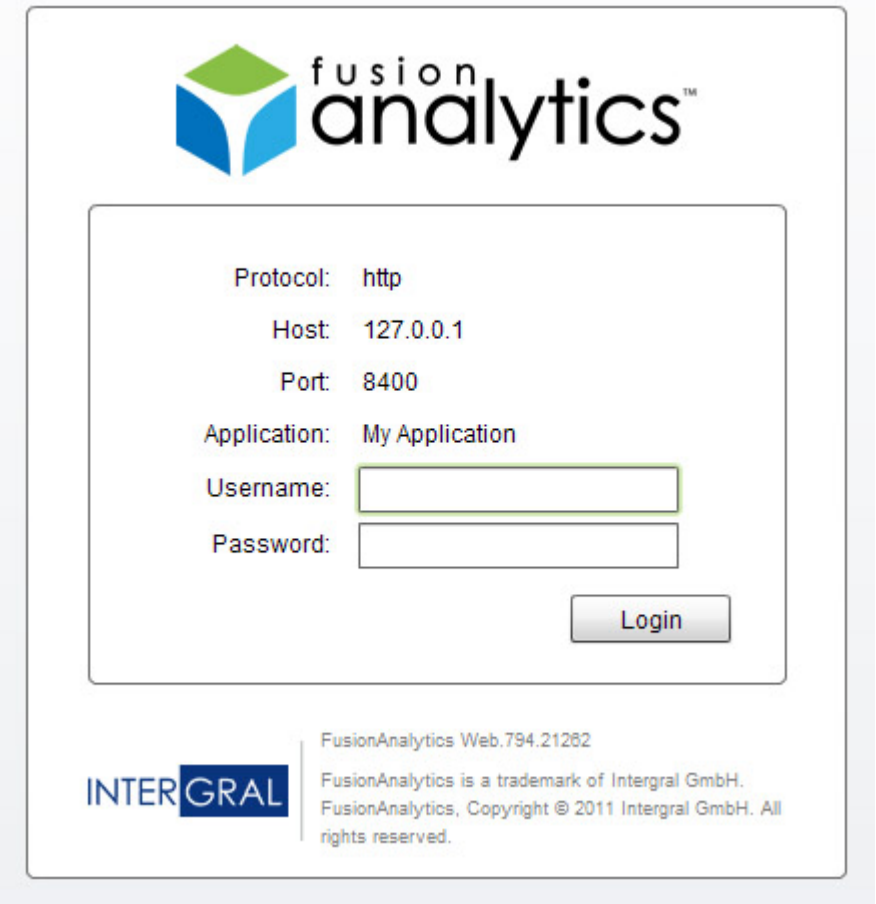

Figure 6: client login panel

Remember Password
To remember the login credentials for this application click the "Remember Password" check box (the username remembered by default for each application) and then the next time you log into th is sever the password filled will be pre-filled so you can simply (and quickly) click the "Login" button to log into the application.

Unlike the Web Client, once you have logged into the AIR client you will have an extra "Logout" option in the "Application" menu. Selecting Logout from this menu will return you to the Server Summary screen.

Back to the top

On This Page:

Next Steps

The FusionAnalytics Web Client
Learn how to log into and use the FA Web Client

FusionAnalytics DataServices (FADS)
Learn about FADS and how it can be used in relation to System (Settings, Sessions, Controllers, Jobs ...), Applications (Summary, Configuration, Detail, Scope ...), Users and Scheduled Tasks (both System and User tasks).

Using the Table of Contents (TOC)

Using the TOC

When you log into an application through the FusionAnalytics Client (either the Web Client client or the AIR Client) you will be taken to a screen showing the TOC (Table of Contents) and Content Area. By default the Content Area will be showing the Overview perspective:

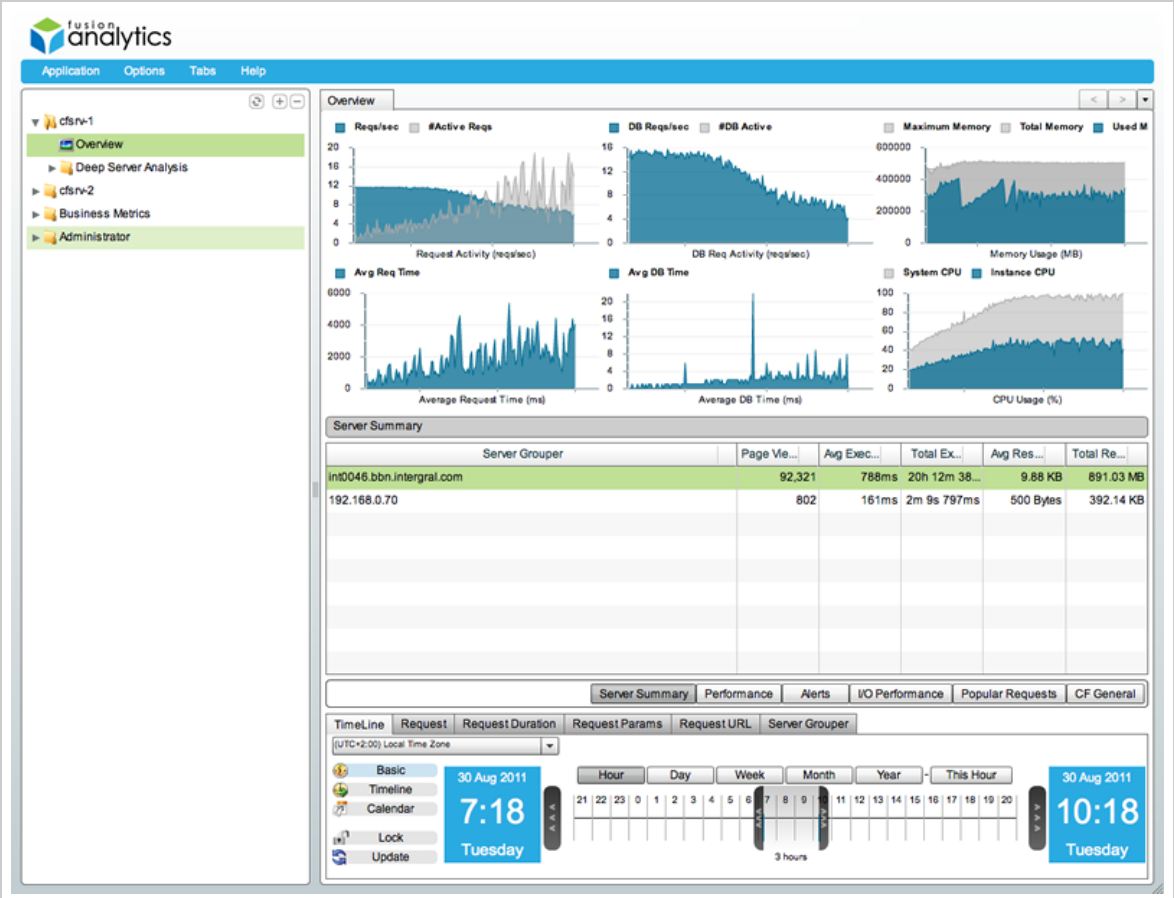

Figure 1: Default Perspective showing the TOC

If you wish to view any of the other perspectives within FusionAnalytics then you can launch them from the TOC. The TOC is a simple tree navigation component containing perspectives. Launch a new perspective by left-clicking on it. You can also open folders by left-clicking them (expanding the menu items).

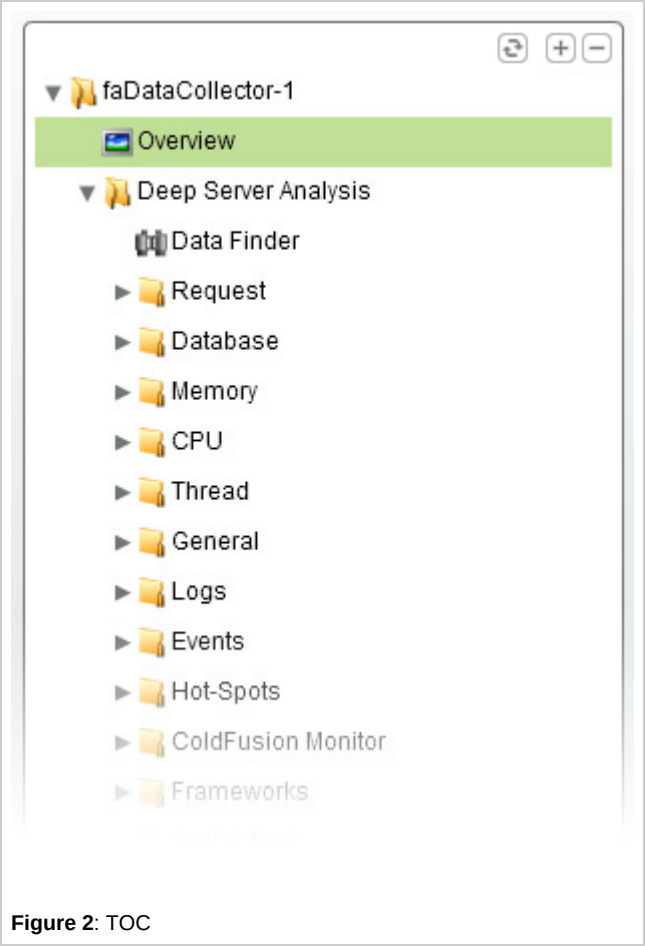

Figure 2: TOC

You will notice three buttons rendered horizontally across the top of the TOC - The TOC can be expanded and collapsed using the and buttons and can also be refreshed (refreshing the TOC's data provider if required / the provider changes) using the button:

	TOC Refresh	*Refresh the TOC data provider*
	Expand TOC	*Expand all TOC menu items*
	Collapse TOC	*Collapse all TOC menu items*

Back to the top

TOC and Content Layout
The application canvas is split into two columns. The TOC is rendered in the left column and the Perspectives (content) are loaded in the right column:

Note: for multiple Perspective handling, the application Tab Navigator is used.

Next Steps

The FusionAnalytics AIR Client
Learn how to access multiple different FusionAnalytics Applications using the FusionAnalytics AIR Client

The FusionAnalytics Web Client
Learn how to log into and use the FA Web Client

FusionAnalytics DataServices (FADS)
Learn about FADS and how it can be used in relation to System (Settings, Sessions, Controllers, Jobs ...), Applications (Summary, Configuration, Detail, Scope ...), Users and Scheduled Tasks (both System and User tasks).

Using Tabs
Learn how to use the Application Tab Navigator

Date Navigation

Date Navigation

Most of the perspectives in FusionAnalytics are time based so, if you want to explore your data further, you will want to change the span of time that you are currently viewing. At the bottom of time-based perspectives you will see a tab named "TimeLine". By default the TimeLine will be in Basic Mode.

Basic Mode

The default TimeLine mode is Basic mode. In this mode you can select a unit of time (Hours, Days, Weeks, Months or Years) and then select a number of those units to be your timespan. Clicking on one of the "unit" buttons will switch to that unit of time.

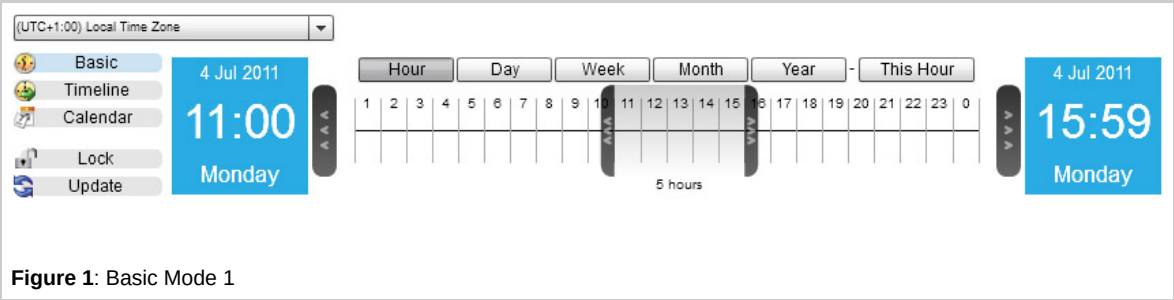

Figure 1: Basic Mode 1

Note: When you switch between different modes, your current span will be retained. If this means that the start or end of your span doesn't match up with the snap points for the new unit of time, then new temporary snap points will be added.

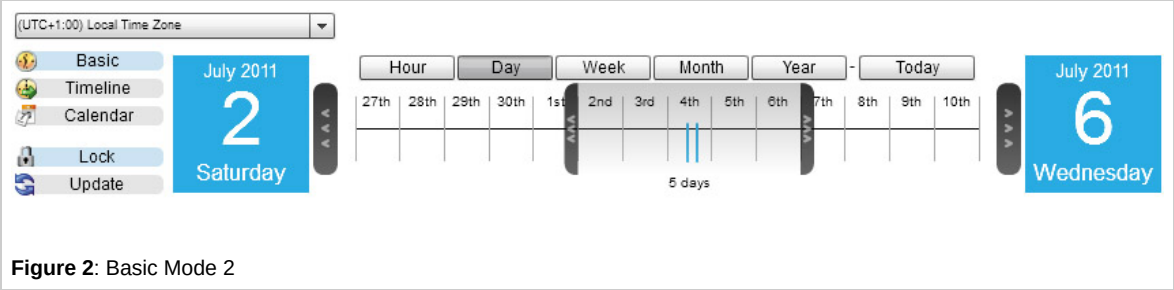

Figure 2: Basic Mode 2

As you switch between modes, the blue calendars on the left and right of the TimeLine will change in order to better represent time spans within each unit. The left-hand calendar displays the start point of the current span and the right-hand calendar displays the end point of the current span.

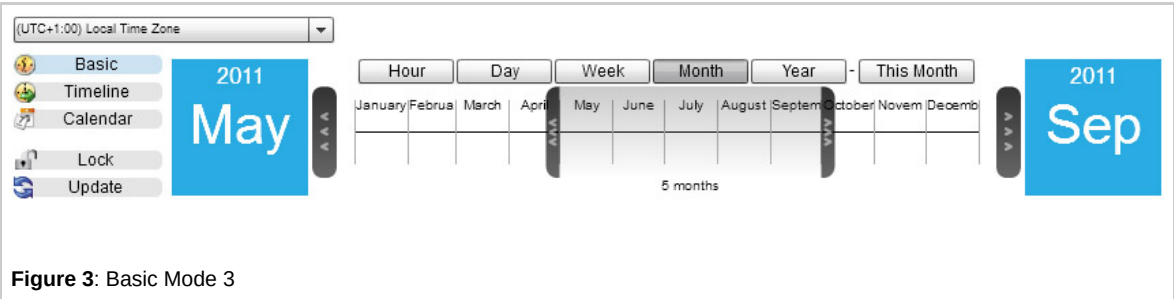

Figure 3: Basic Mode 3

Clicking on the Arrow buttons between the calendars and the central time span will push the current window in that direction. Click and drag the central span handles in order to change the start or end points of the current time span. You can also click and drag from the center of the span in order to move both the start and end points at the same time.

At any point you can return to "now" by clicking the button to the right of the unit buttons. This button changes label depending upon which time unit you currently have selection. eg. "This Hour", "Today", "This Week"...

Back to the top

Timeline Mode

The Timeline view represents time as one long piece of ticker-tape. Click the "Timeline" button on the far left in order to switch to this mode. The window of time at which you are currently looking is the timespan currently selected. There are various ways in which you can interact with this time span. Firstly, you can click and drag the timeline to the left or right. Secondly, you can click the left and right buttons at the edges of the Timeline in order to move the current time span forward or back.

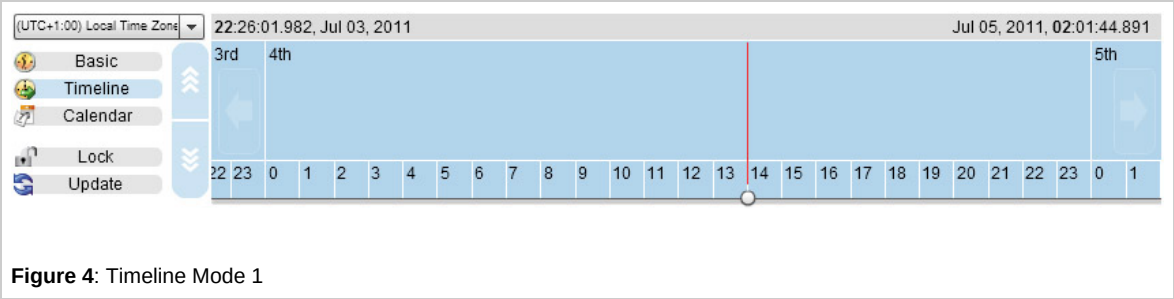

Figure 4: Timeline Mode 1

If you want to increase of decrease the size of the currently visible time span then you can click the up and down arrows at the left side of the Timeline. Clicking these will step up or down between the different units available. You can also use the mouse wheel to zoom the Timeline in and out. Unlike the Basic mode, the Timeline mode has many degrees of time unit, starting from millisecond, tens of milliseconds or hundreds or milliseconds, going right up to hundreds of years. As you scroll in and out, you will notice that the left and right buttons move up and down with the rest of the Timeline labels. A scroll button in the "minutes" areas of the Timeline will move you forward or back one minute, whilst the same button in the "years" area will move you by a full year. The buttons are context sensitive.

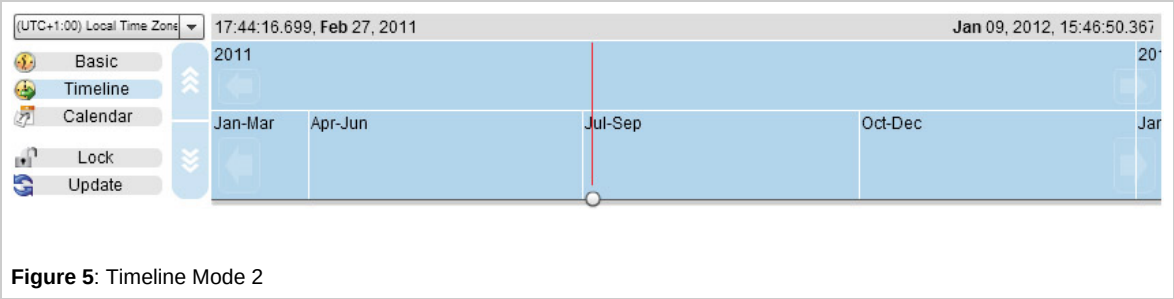

Figure 5: Timeline Mode 2

You can also double-click on any unit of time within the timeline. If you clicked on a unit from the lower half of the Timeline then the current displayed span of time will be changed so that it perfectly matches up with the unit that you double clicked on. If you clicked on a unit from the upper half of the Timeline then the currently display span of time will be zoomed out to perfectly match with the next highest unit above the one you double clicked on.

If you also like to use keyboard commands then the left and right cursor keys can be used to step forward and back by the smaller of the two visible time units. The up and down cursor keys can be used to zoom the Timeline in and out in the same way that the up and down buttons on the right do.

> You will notice in the screen shots that there is a red marker line sitting over the Timeline. In this case you will also see marker lines on all supported graphs within your currently displayed Perspective. Moving the marker line using the handle at the bottom of the Timeline area will also move the marker on all of these graphs. This is just make help you when you are comparing data accross multiple graphs. You can double click the handle at the bottom of the timeline to turn the cursor on and off.

Back to the top

Calendar Mode

The final mode for time span selection is the Calendar mode. Click the "Calendar" button on the far left to switch to this mode. In Calendar mode you can click within the calendar components or the click on the numeric stepper components in order to change the start and end points of the currently selected time span. By default the start and end points will be linked (which is represented by the central check box labelled "Link" having a tick within it.) When the start and end points are linked, you will be unable to change the amount of time within the current time span, so if you had an hour-long time span and increased the start point by a minutes. The end point would also be increased by a minute.

If you click the "Reset" button for the start or end times then the Hour, Min, Sec and MS fields will all be reduced to zero. This can be helpful if you want to isolate a span of time encompassing a number of full days. (Note that this will still try and retain the current span of time though so clicking one and then the other whilst the dates are Linked will not leave you with both times at midnight, unless they both started out with exactly 24 hours separating them.)

If you unlink the start and end points (by clicking the central check box) then you can change the start and end points independently. (Although obviously, the start point must always occur before the end point.)

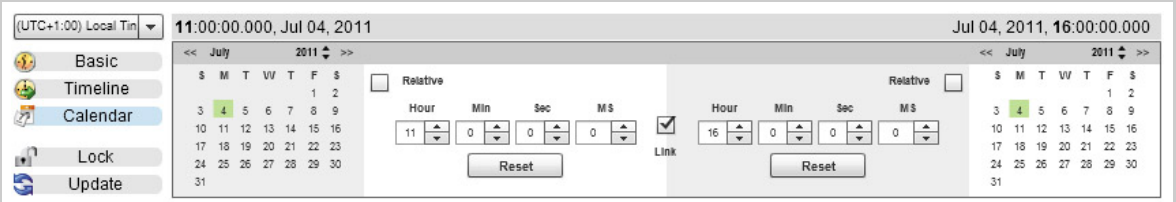

Figure 6: Calendar Mode 1

If you have the start and end points linked then you can change either the start point or the end point to "Relative" mode. In the image below, the start point is dealt with as before, but the end point is defined by picking a unit of time and then specifying how many of those units separate the start and end points. You can click and hold any of the numberic steppers in order to make them cycle through numbers quicker.

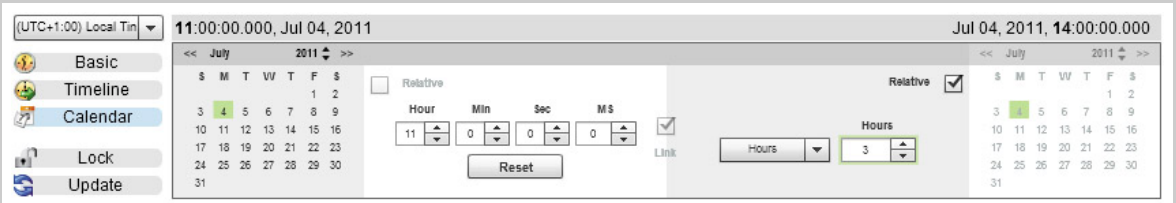

Figure 7: Calendar Mode 2

The Calendar components on the left and right-hand side of the display area can be changed by simply clicking of a desired Date. The left and right arrows at the top of each calendar allow you to change the current Month and to the right of the Year there are up and down stepper buttons as well.

Back to the top

Locking and Updating

If you have been trying the different date navigation modes outlined above then you will have noticed that, even though you have been changing the currently selected time span, the data shown within your current perspective will not have been updated to reflect your changes. At any point you can click the "Update" button and your perspective will be refreshed in order to show your currently selected span of time. Even if you have not changed the current time span, you can click the "Update" button to get the most recent data from FusionAnalytics.

There are many situations in which you would want graphs to automatically be updated as you scroll through time. In this case you can lock the TimeLine. Click the "Lock" button to toggle this feature on and off. Whilst locked, your perspective will be updated whenever you change the currently selected timespan.

Figure 8: Lock and Update Buttons

Both the "Lock" and "Update" features work with all three of the Timeline display modes.

> ⚠ If you have the Timeline "Locked" whilst browsing through data then then you will be constantly asking FusionAnalytics for new data to show. In this case you may find that the performance is slower than having the Timeline "Unlocked".

Back to the top

Timezone Control

You will notice in the top left of the Date Navigation component (used for all modes) is a Time zone Control (drop down):

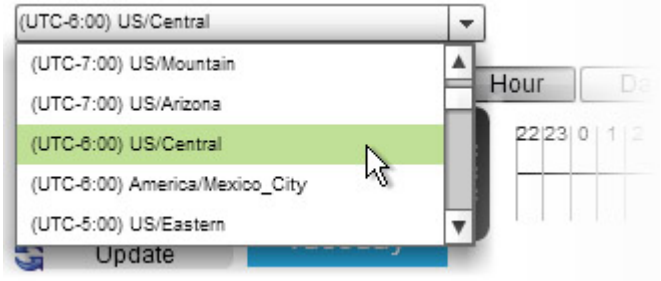

|

Figure 9: Timezone Control |

The Time zone Control allows you to view Perspective data in the time zone of your choice. By default, this will be set to your local time. For more information about Time zones and the Time zone control See: Time Zones.

Back to the top

On This Page:

Time Zones
See Time Zones to learn about using and controlling Time Zones within the Client

Next Steps

Time Zones
Learn about Time Zones within the Client and how to use the Time Zone control

Time Zones

Time Zones

FusionAnalytics allows you to view graph and table data in the Time Zone of your choice. By default, this will be set to your local Time Zone. (ie, the Time Zone in which the FusionAnalytics Client is running.) If you want to change your viewing Time Zone then you can select a new Time Zone from the drop down box in the top left of the Date Navigation area:

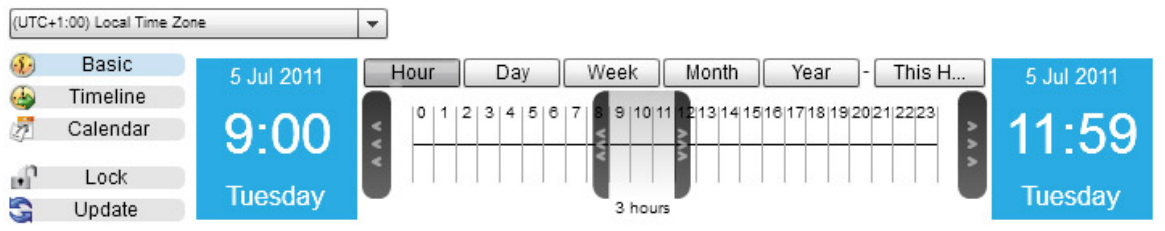

Figure 1: Timeline Control - Basic Mode

Many perspectives have an overview at the top and you can also see the current Time Zone setting in this location:

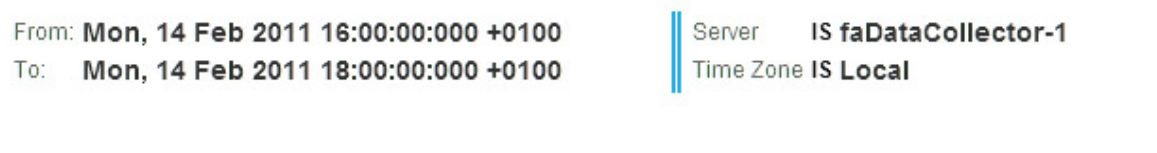

Figure 2: Perspective Filters

If you change the current Time Zone then all the times and dates displayed in the client will be updated to take this new Time Zone into consideration.

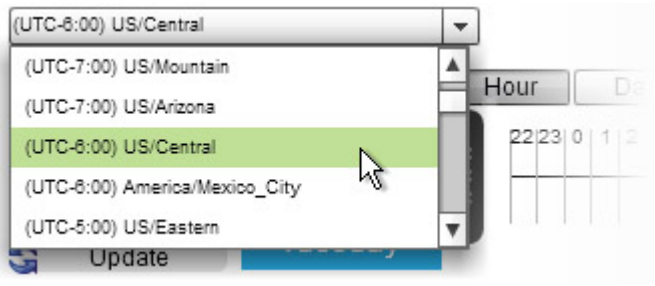

Figure 3: Timezone Control

Dates displayed in the overview section, in data grids and on graphs will include the current Time Zone Offset so that you know what Time Zone you are currently using.

Figure 4: Perspective Filters

Note: Many Time Zones have an additional hour offset over summer. This is referred to as Daylight Savings Time (DST). If you are looking at a Time Zone which is UTC - 6:00, but is currently using Daylight Savings Time, then the actual offset will be UTC - 5:00. In this situation the offset will be displayed as -0500 but the string "(DST)" will be added to the date in order to identify this change, as in the following image which represents a span of time that crosses the Daylight Savings Time boundary.

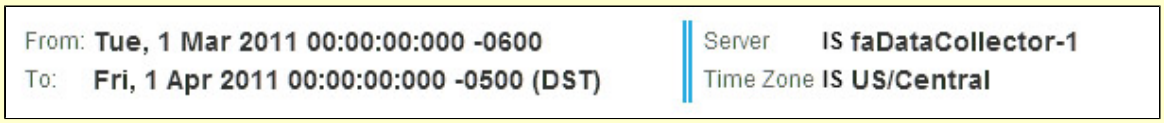

Back to the top

Using Time Zones
The client offers over 50 different Time Zones that can be used to view your data, defaulting to your local (current) Time Zone and additionally supporting UTC Time.

Next Steps

Date Navigation
Learn how to use the Time Line Component for advance Data Navigation.

Filtering Data

Filtering Data

FusionAnalytics gives you a lot of data so it is useful to be able to focus in on just the information you want. The type of filters available to you will differ depending upon the type of perspective. You will see a tab at the bottom of the perspective for each type of filter available (if any).

Figure 1: Filters Tab

Timeline

The TimeLine filter can be used to select a specific period of time. For more information on this filter, see the Date Navigation section. The TimeLine filter also contains the TimeZone Control allowing you to view graph and table data in the TimeZone of your choice.

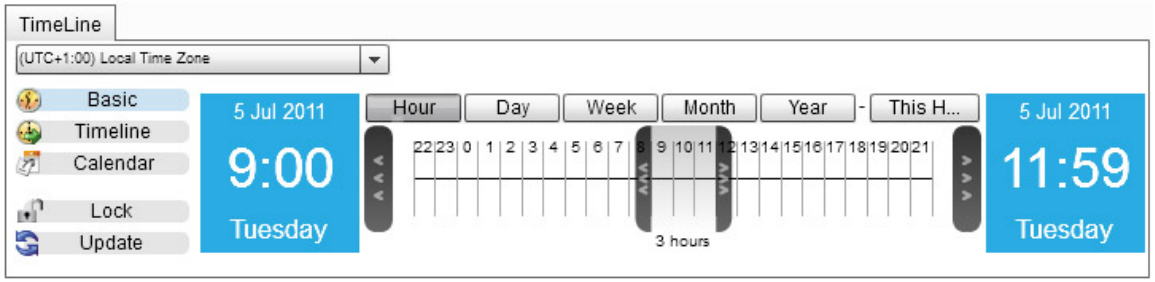

Figure 2: Timeline Tab

Back to the top

Request

This filter type allows you to limit the view to only include requests with specific limitations. These are your options:

Request ID	IS, ISNOT, >, <, >=, <=
JSession ID	IS, ISNOT, LIKE, NOTLIKE
Method	IS, ISNOT, LIKE, NOTLIKE
CFID	IS, ISNOT, LIKE, NOTLIKE
CFTOKEN	IS, ISNOT, LIKE, NOTLIKE
Client IP Address	IS, ISNOT, LIKE, NOTLIKE
SERVER	IS, ISNOT
THREAD ID	IS, ISNOT, LIKE, NOTLIKE
CP Reason Code	IS, ISNOT, LIKE, NOTLIKE
AMF Request	IS, ISNOT, LIKE, NOTLIKE
Status Code	IS, ISNOT, >, <, >=, <=
Request Execution Time	IS, ISNOT, >, <, >=, <=
Memory Percent	IS, ISNOT, >, <, >=, <=
Number of Queries	IS, ISNOT, >, <, >=, <=
Total Time in Queries	IS, ISNOT, >, <, >=, <=
Execution Time in Queries	IS, ISNOT, >, <, >=, <=
Total DB Rows Retrieved	IS, ISNOT, >, <, >=, <=
Response Size	IS, ISNOT, >, <, >=, <=
Time To First Byte	IS, ISNOT, >, <, >=, <=
Time To Last Byte	IS, ISNOT, >, <, >=, <=

Once you have entered your filter type, comparison type and value, you can click the "+" button to add the new filter to the list of filters below. Clicking the "Apply Filter" button will update the perspective to include your defined filters.

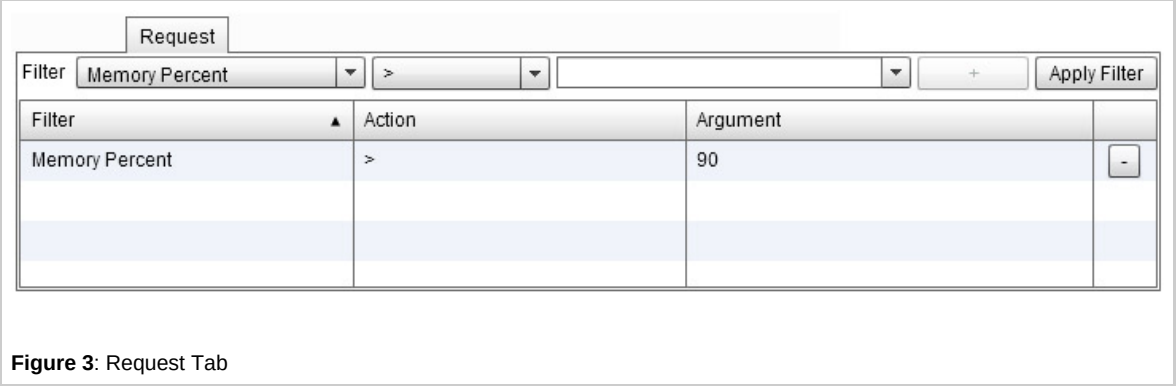

Figure 3: Request Tab

> **Note**: The Request filter is based on a MultiComboFilter component. For more information about using this component see the MultiComboFilter Overview section at the bottom of this page.

Back to the top

Request Duration

This filter allows you to use a slider to define a minimum and maximum request duration to be displayed. The minimum and maximum values are both in milliseconds.

Figure 4: Request Duration Tab

> You can change / adjust the minimum and maximum values by wither dragging the slider thumbs, or to be more precise by using the numeric steppers below the slider control by clicking the "up" and "down" stepper buttons or simply typing values into the control.

Back to the top

Request Params

This filter allows you to select values for a given request parameter. You have the option of picking "fuseaction", "event", or you can enter your own Request parameter to filter. For each filter you can either choose "IS" or "ISNOT".

Once you have entered your filter type, comparison type and value, you can click the "+" button to add the new filter to the list of filters below. Clicking the "Apply Filter" button will update the perspective to include your defined filters.

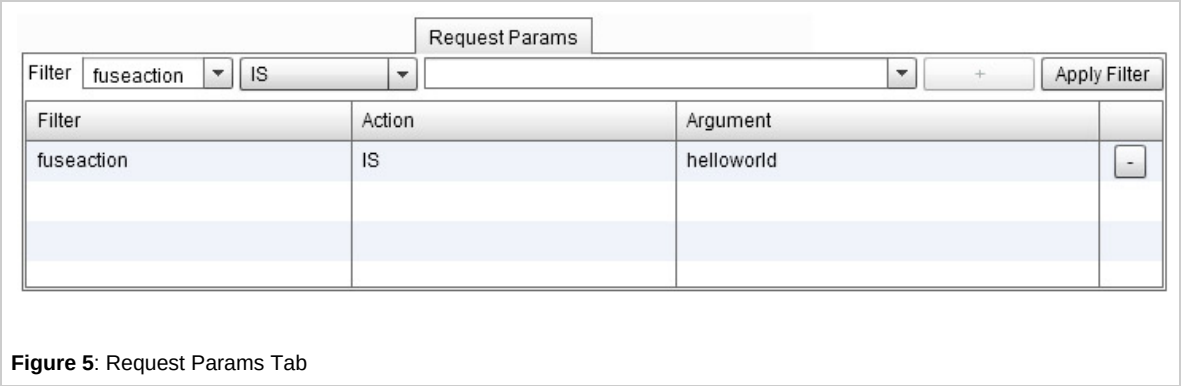

Figure 5: Request Params Tab

> **Note**: The Request Params filter is based on a MultiComboFilter component. For more information about using this component see the MultiComboFilter Overview section at the bottom of this page.

Back to the top

Request URL

This filter lets you isolate some section of the request URL and filter just that part of the URL. These are the options:

URL Protocol (protocol://*...)
URL Username (://username:...)
URL Password (://:password@*...)
URL Hostname (*://hostname/...)
URL Port (://:port/*...)
Path Element 1 (:///1/...)
Path Element 2 (:////2/...)
Path Element 3 (://///3/...)
Path Element 4 (:///////4/...)
Path Element 5 (://////////5/...)
Path Element 6 (:////////6/...)
Path Element 7 (://////////7/...)
Path Element 8 (:///////////8/...)

Once you have entered your filter type, comparison type and value, you can click the "+" button to add the new filter to the list of filters below. Clicking the "Apply Filter" button will update the perspective to include your defined filters.

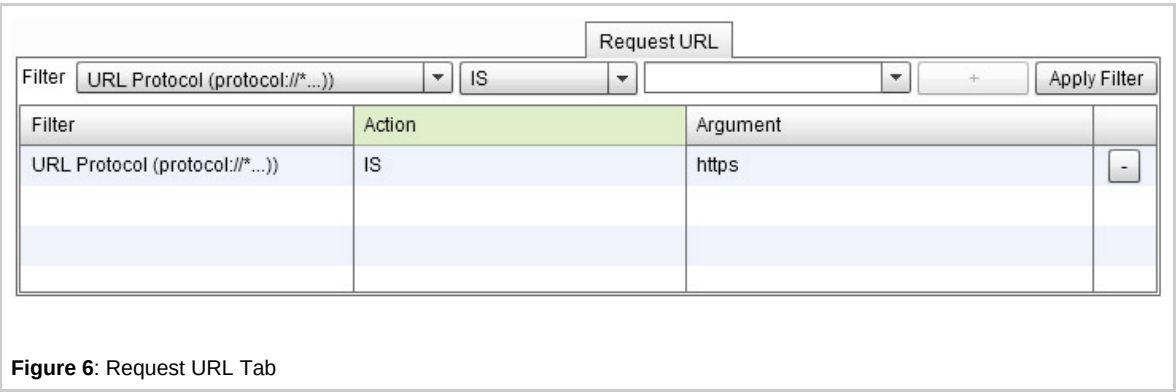

Figure 6: Request URL Tab

> **Note**: The Request URL filter is based on a MultiComboFilter component. For more information about using this component see the MultiComboFilter Overview section at the bottom of this page.

Back to the top

MultiComboFilter Overview

As discussed within this section the are some filters () that are based on the MultiComboFilter component. The MultiComboFilter is a component that accepts different values (DataProviders) that can be used to filter different data types. The component allows the user to add filter values and compare these filters to arguments (also user specified) using different action types (E.g. IS, ISNOT, LIKE, NOTLIKE, >, <, >=, <= etc ...).

As an Example use of the MultiComboFilter component lets take the "Request" filter:

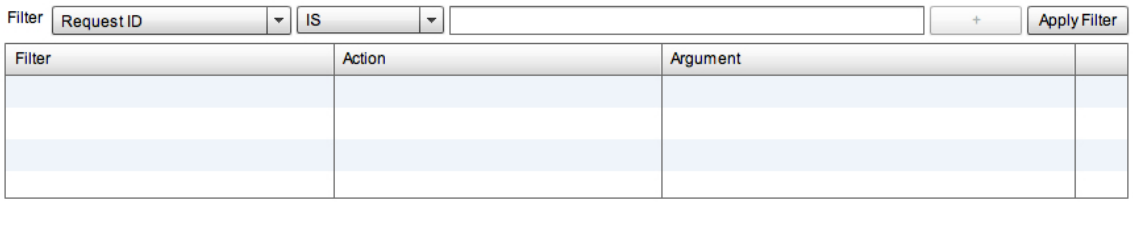

Figure 7: MultiComboFilter 1

The MultiComboFilter component consists of a horizontal filter control bar containing the three filter inputs (rendered as either a combo box / drop down list or a text input), an add "+" button and an "Apply Filter" button (all laid out horizontally across the top of the component) and a Grid placed below to display the added filters. The three filter inputs include:

1. The Filter itself (the first input)
2. The filter Action (the second input)
3. The filter Argument (the third / last input)

All three filter inputs are linked together / are dynamic with Filter input controlling both the DataProviders and Item Renderers of the Action and Argument inputs, therefore the available Actions and the Argument type (if it can be any user defined input / string or a distinct list of items) depends on the filter that has been selected. E.g. The "Client IP Address" Filter contains the Actions "IS, ISNOT, LIKE, NOTLIKE" and can be any user defined input so the Argument input is rendered as a "Text Input" :

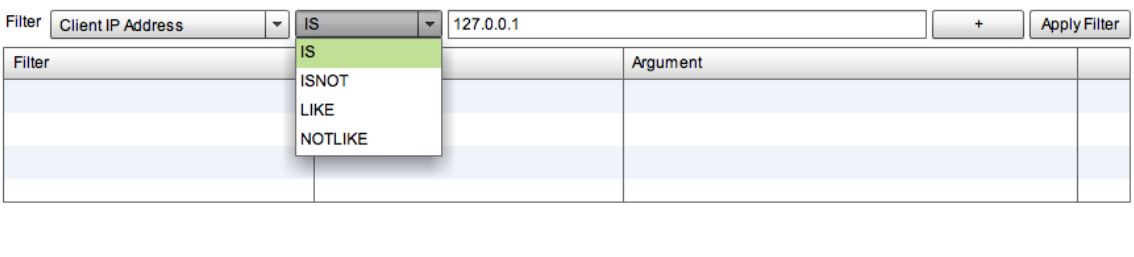

Figure 8: MultiComboFilter 2

Whereas the "Status Code" Filter contains the Actions "IS, ISNOT, >, <, >=, <=" however its Arguments are a distinct list of values so the Argument input is rendered as a Combo Box / Drop Down list:

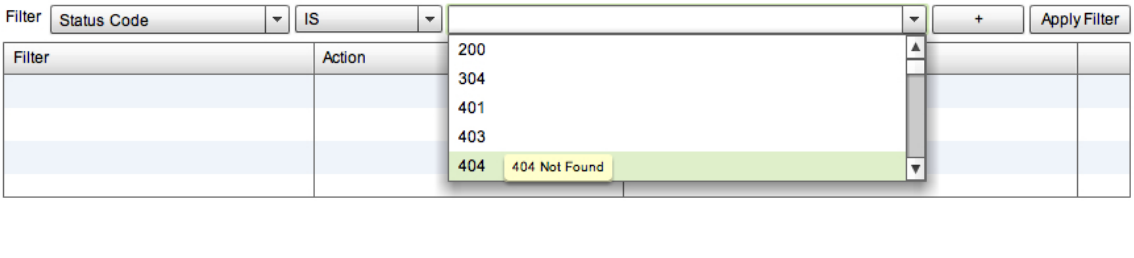

Figure 9: MultiComboFilter 3

Once you have selected a filter and added an Action and Argument, you can add the filter to the filters grid by clicking the add "+" button and / or apply the filter by clicking the "Apply Filter" button.

The Filter inputs / values contain validation rules that applied to the Arguments for each Filter type (E.g. filters types can be of type Number not allowing character values etc...) :

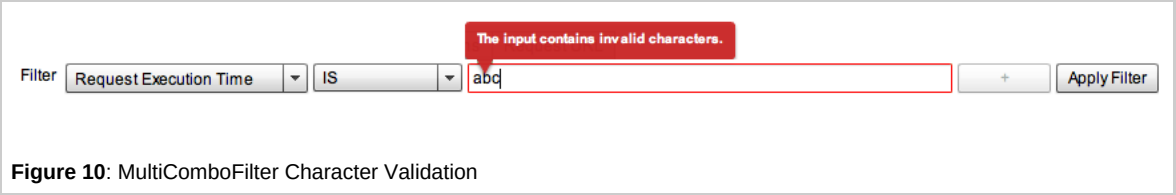

Figure 10: MultiComboFilter Character Validation

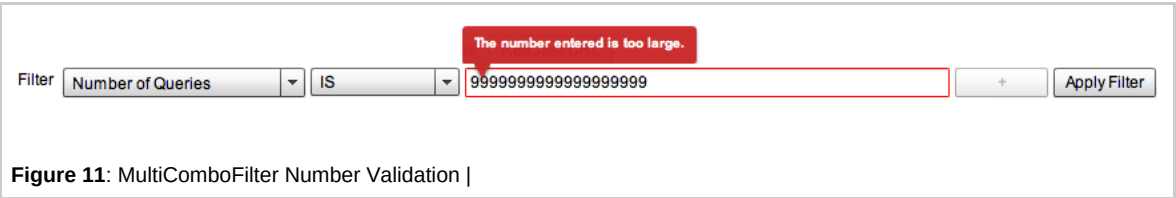

Figure 11: MultiComboFilter Number Validation |

Note: The add "+" button will be disabled while the content within the Argument input is either empty or invalid (in relation to the selected Filter) to prevent any invalid filter queries against the DataBase :

Once valid content has been entered into the Argument input the add "+" button is enabled and the Filter can be added to the Filters Grid

When the "Apply Filter" button is clicked only filters that have been added to the Filters Grid are applied. If the Add button is enabled (if you are adding a new filter) and the "Apply Filter" button is clicked, the new filter is both added to the Filters Grid and applied with all other filters in the Grid. If the Add button is disabled (there is no new filter to add) all the filters currently within the FIlters DG are applied

Adding Filters
When adding single filters the "Apply Filter" button is very useful and faster to use that Adding the Filters to the Filters Grid and then Applying the Filters using a second step, however, if adding multiple filters, from a performance perspective it is better to add the filters individually and then Apply them all together as a second step using the "Apply Filter" button (as when the "Apply Filter" button is used it runs a query and updates the results each time, therefore when adding ten Filters, ten Queries will be performed being a lot more expensive than just one as when using the "Add" button to add the filters).

Once Filters have been added they are displayed in the Filters Grid where they can bu further edited or deleted:

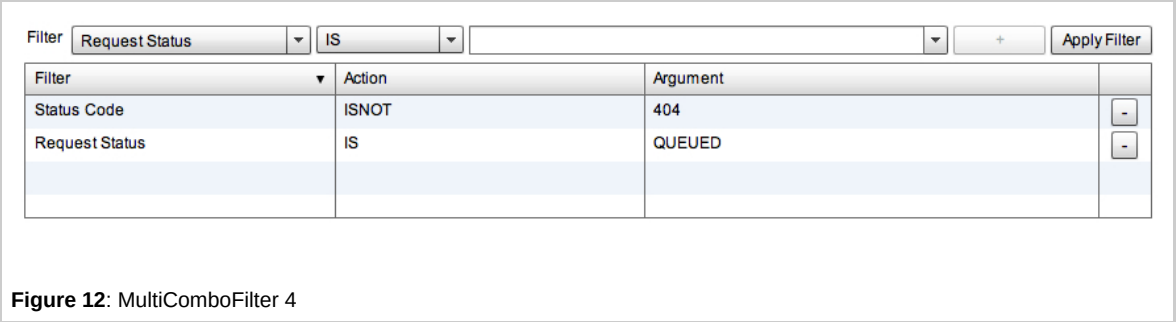

Figure 12: MultiComboFilter 4

The filters within the Filters Grid can be edited in the same way as they were added. To edit a filter simply click the element (Action or Argument) within the row that you want to edit, and the cell will become editable (using its required item renderer) so thqt you can change the value(s). Once you are done editing click the "Apply Filter" button to update the Perspective results using the new filter(s). To remove / delete filters clikc the delete "-" button :

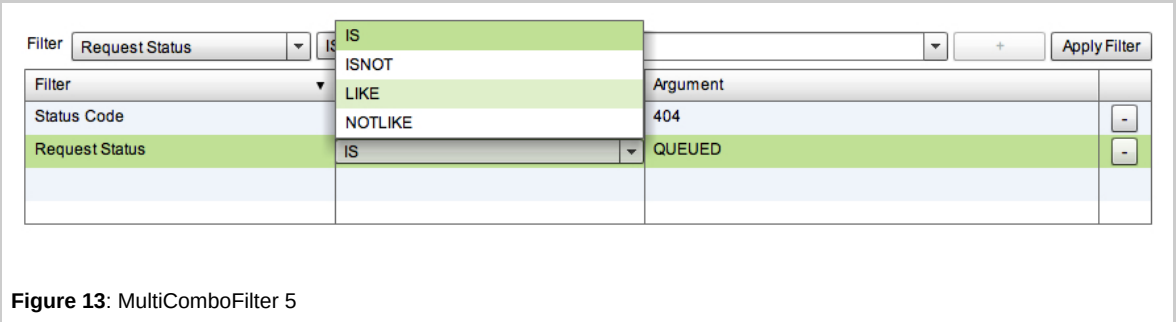

Figure 13: MultiComboFilter 5

Note: The filters within the Filters Grid contain the same validation rules as when adding the Filters. If any arguments within the Filters Gris are not valid the "Apply Filters" button will be diabled until the Argumentshave been changed to a valid value preventing any invalid

Arguments to be executed within a query :

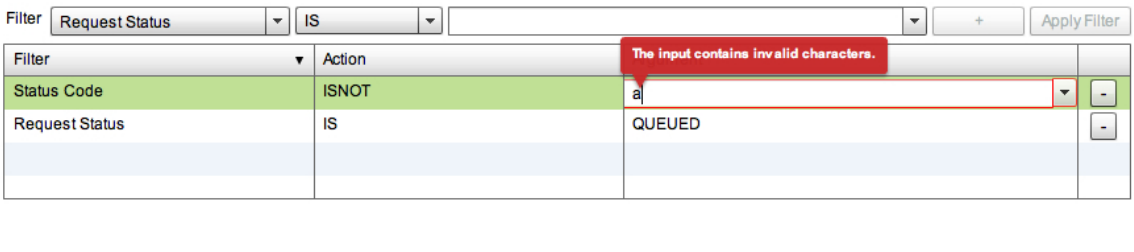

Figure 14: MultiComboFilter 6

Back to the top

On This Page:

Next Steps

Date Navigation
Learn how to use the Time Line Component for advance Data Navigation.

Time Zones
Learn about Time Zones within the Client and how to use the Time Zone control

Data Visualizations

Data Visualizations

Overview

This section will discuss the different Data Visualizations that are provided and used within the different application Perspectives and how they can be used to help visualize and analyze data in the most efficient manner(s)

Next Steps

DataGrid Visualization

Chart Base Visualization

Time Block Chart Visualization

Placemark Chart Visualization

Sixway Chart Visualization

DataGrid Visualization

DataGrid Visualization

The DataGrid Visualization is used to render data in a "list" based format having a similar layout to a (vertical) table with column headers at the top and rows of data (records) listed below. The DataGrid Visualization offers user controlled re-sizable and sortable columns, text

wrapping (for larger strings), cell based Tool Tips, single and multiple row selection (including selection copy functionality), Context Menus and data paging to allow easy and usable scrolling through large data sets.

The following image shows a typical DataGrid Visualization:

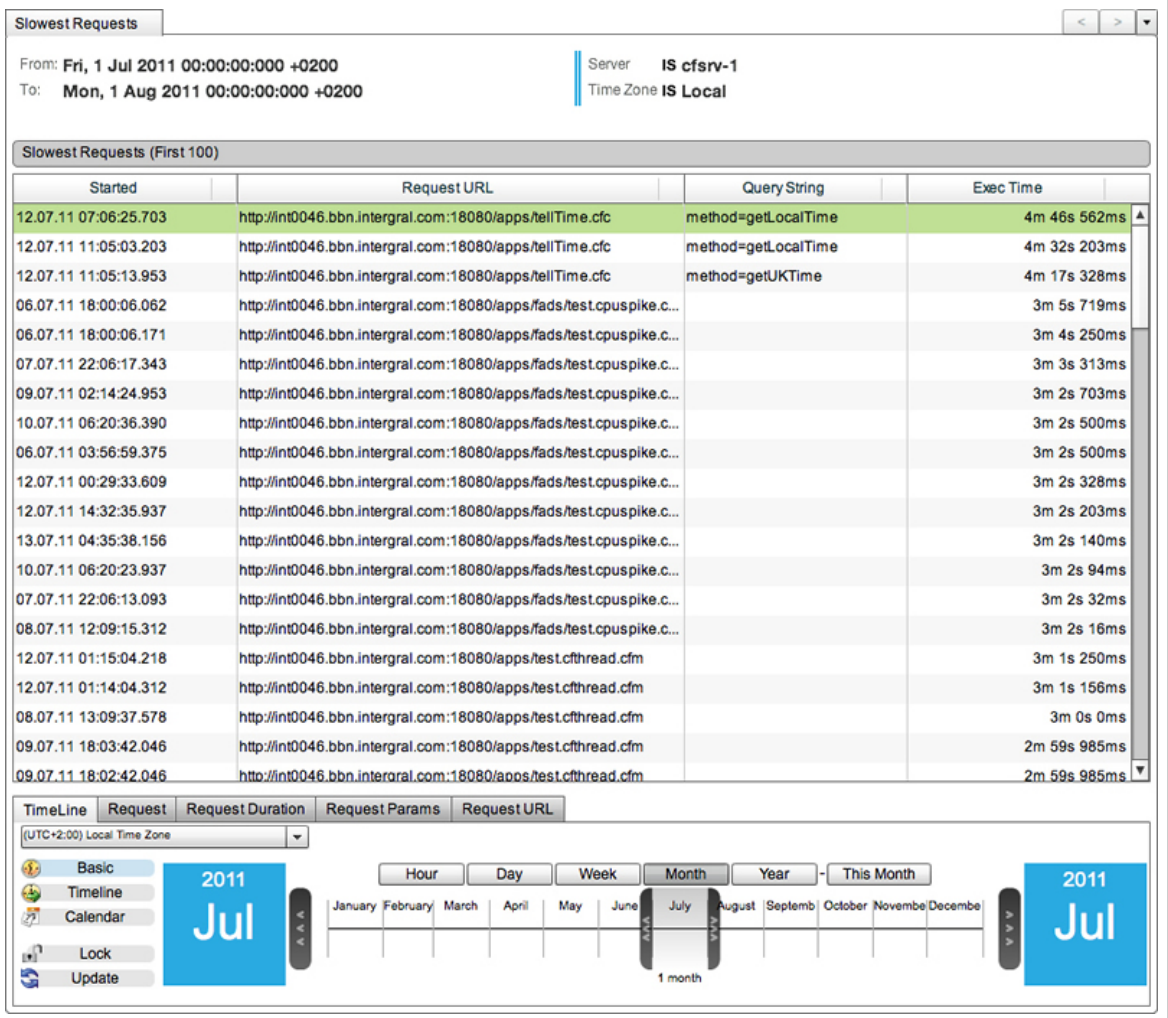

Figure 6: DataGrid Visualization

Back to the top

Data Sorting

The DataGrid Visualization allows data to be sorted by "multiple" columns simply by clicking the column headers:

When the DataGrid is initialized it displays the data in the order that is is delivered from its DataProvider (DP). Once the DataGrid has been initialized (loaded) it allows the user to sort the data by single or multiple columns. Ordering data by a column is very straight forward, simply click a column header to order by that column, and click it again to re-order by the same column but in the reverse order etc. to add a second, third, fourth (multiple) ... order filter / sort repeat this process by clicking on a additional column headers to apply multiple column sorts. The direction in which the sort has been applied (ascending / descending) is displayed using a small arrow within the column header(s) and the order in which the sort has been applied to the columns is represented by a number (again displayed in the column header):

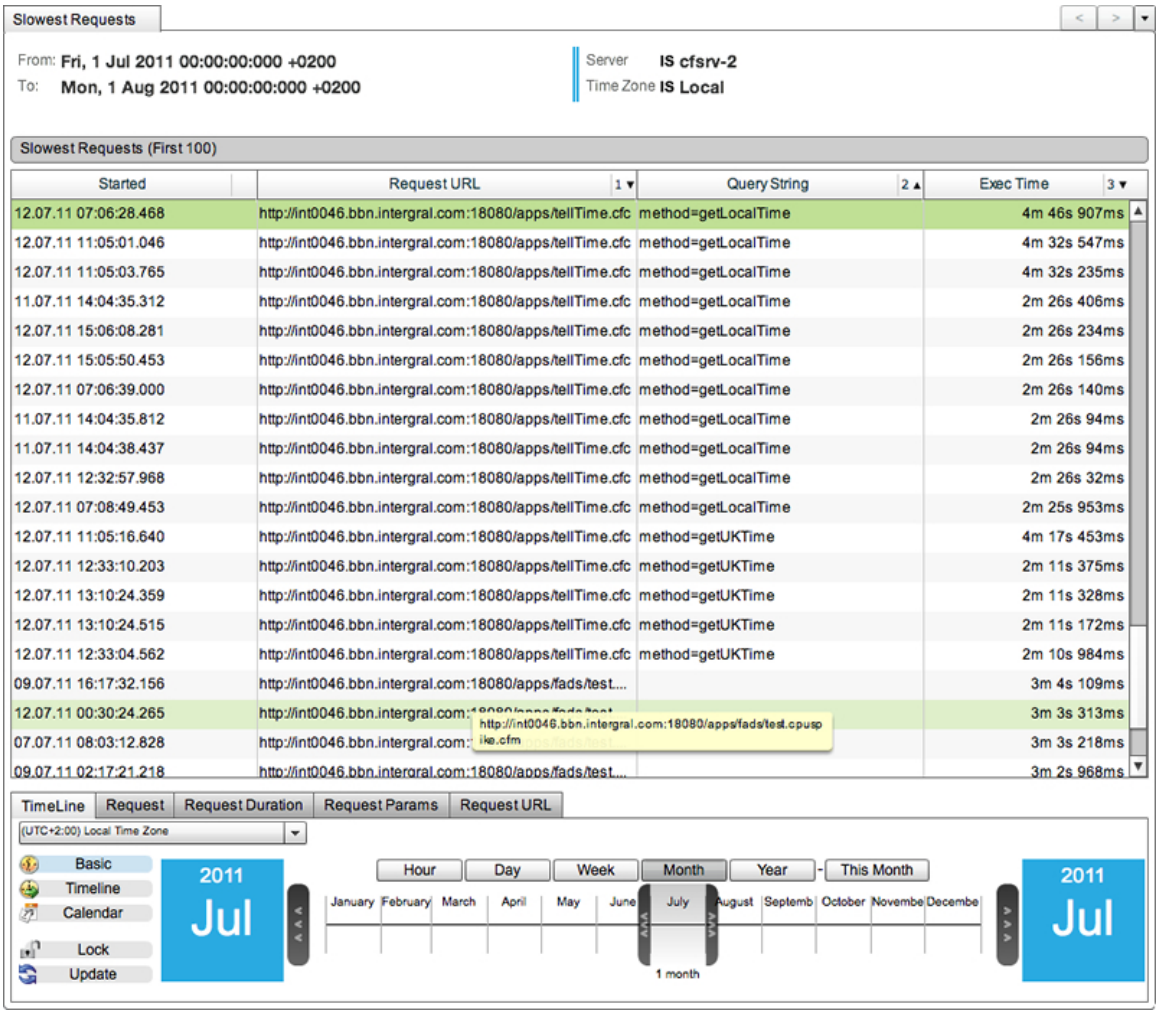

Figure 2: DataGrid Visualization - Data Sorting

Back to the top

Column Resizing

Columns can be re-sized by hovering over the column border (within the column header) until the cursor changes from a pointer to the "re-size" cursor: ⇹. When the re-size cursor is visible, click the column edge and drag it left or right to re-size it.

Columns resizing is useful when column content is too large / long to be rendered using the default column widths. In these cases the columns can be re-sized (as discussed), however this can also (sometimes) affect other column content preventing this content (all DataGrid content) to be fully displayed at the same time. The DataGrid Visualization further supports text truncating to help with this issue. Text truncating truncates strings that are too long to be rendered within a column. The string is cut to fit into a column and a trialing "..." is appended to the end of the string allowing the user to see that the string has been truncated.

Back to the top

Grid Tool Tips

The DataGrid Visualization additionally supports Tool Tips which are displayed when on a mouse over event for each cell within the DataGrid. The Tool Tip displays the (full) text content of the cell that the mouse is currently over (useful for seeing truncated content).

Back to the top

Grid Context Menu

Within most application Perspectives using DataGrid Visualizations, Context Menu support has also been implemented. To use the Context Menus, simply right click on a row within the DataGrid and a Context Menu will be displayed. Each Context Menu has a DataProvider that is relevant to the current Perspective / Visualization and the Context Menu renders each DataProvider item as a menu item within the Context

Menu itself which the user can click / open.

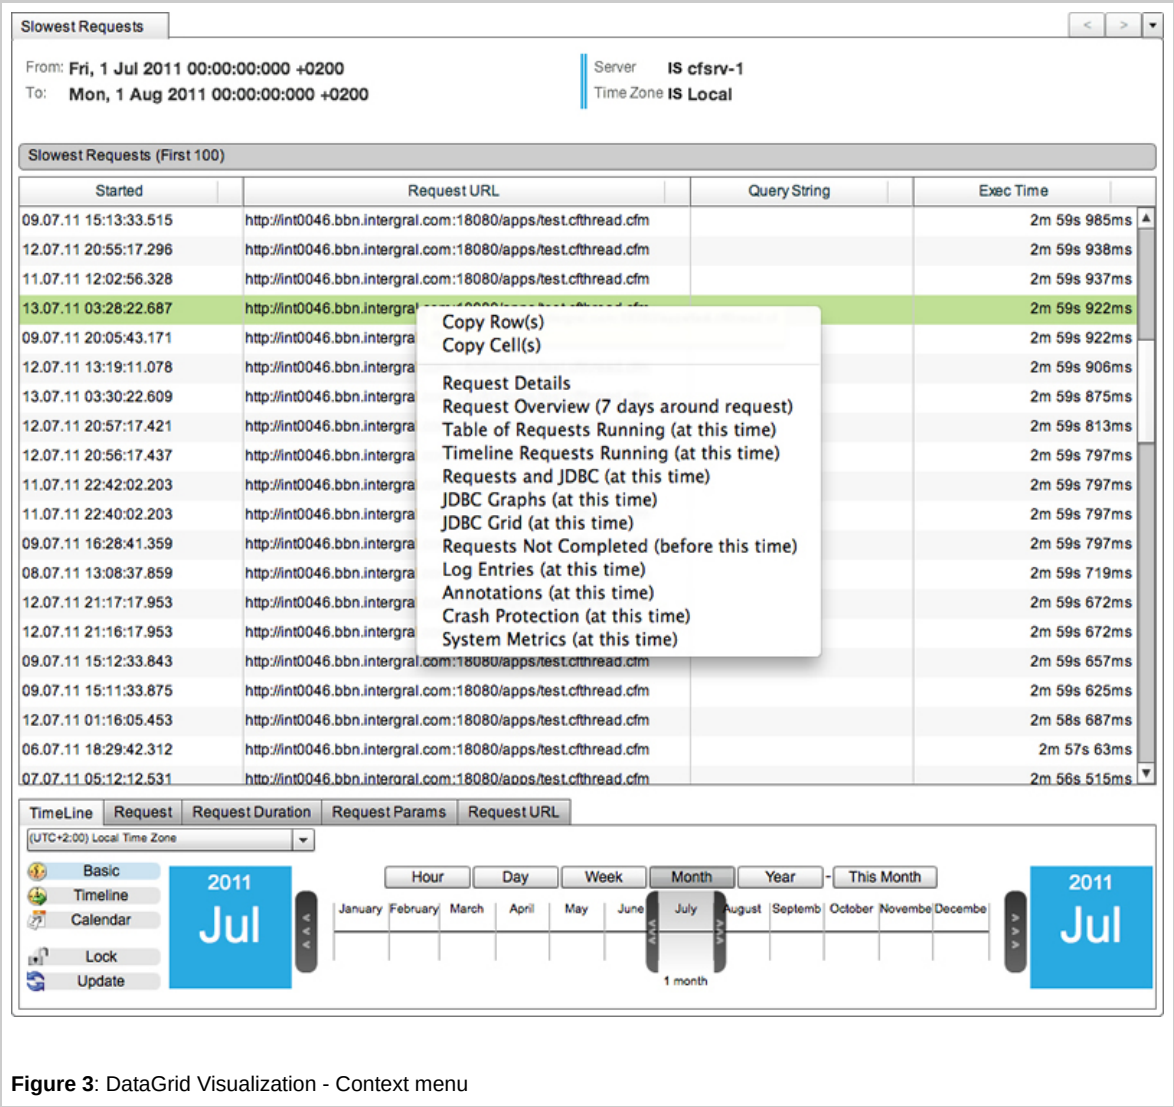

Figure 3: DataGrid Visualization - Context menu

As you can see, the Context Menu Items are generally "links" to other Perspectives relevant to the current Perspective (date / time) within the application.

Back to the top

Copying Grid Content

You will also notice that the Context Menu within the DataGrid Visualization contains two additional menu items:

- Copy Row(s)
- Copy Cell(s)

The DataGrid Visualization supports both (multiple) row and cell section. To select multiple rows, select a row, hold the "Shift" key and select a second row or select a row and drag the mouse (up or down) across the rows you wish to select - The selected rows are then highlighted in green.

When the "Copy Row(s)" menu item is selected, the currently selected rows from "all" columns within the DataGrid will be copied to the computers clipboard (in the order that they are currently rendered within the DataGrid). E.g.:

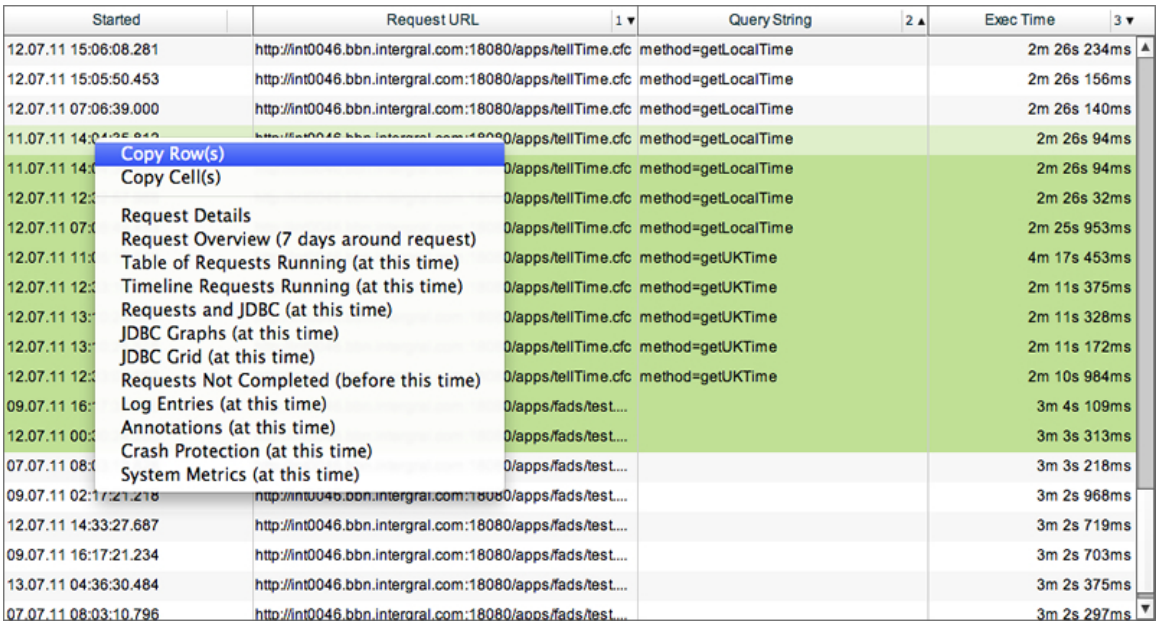

Figure 4: Copy Rows

The above ("Copy Row(s)") menu item click will copy the selected rows to the clipboard where they can be further pasted resulting in:

```
Started Request URL Query String Exec Time
11.07.11 14:04:35.812 http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc
method=getLocalTime 2m 26s 94ms
11.07.11 14:04:38.437 http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc
method=getLocalTime 2m 26s 94ms
12.07.11 12:32:57.968 http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc
method=getLocalTime 2m 26s 32ms
12.07.11 07:08:49.453 http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc
method=getLocalTime 2m 25s 953ms
12.07.11 11:05:16.640 http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc  method=getUKTime
4m 17s 453ms
12.07.11 12:33:10.203 http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc  method=getUKTime
2m 11s 375ms
12.07.11 13:10:24.359 http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc  method=getUKTime
2m 11s 328ms
12.07.11 13:10:24.515 http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc  method=getUKTime
2m 11s 172ms
12.07.11 12:33:04.562 http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc  method=getUKTime
2m 10s 984ms
09.07.11 16:17:32.156 http://int0046.bbn.intergral.com:18080/apps/fads/test.cpuspike.cfm    3m 4s
109ms
12.07.11 00:30:24.265 http://int0046.bbn.intergral.com:18080/apps/fads/test.cpuspike.cfm    3m 3s
313ms
```

When the "Copy Cell(s)" menu item is selected, the currently selected rows from only the current (mouse over) column will be copied to the computers clipboard (in the order that they are currently rendered within the DataGrid). E.g.:

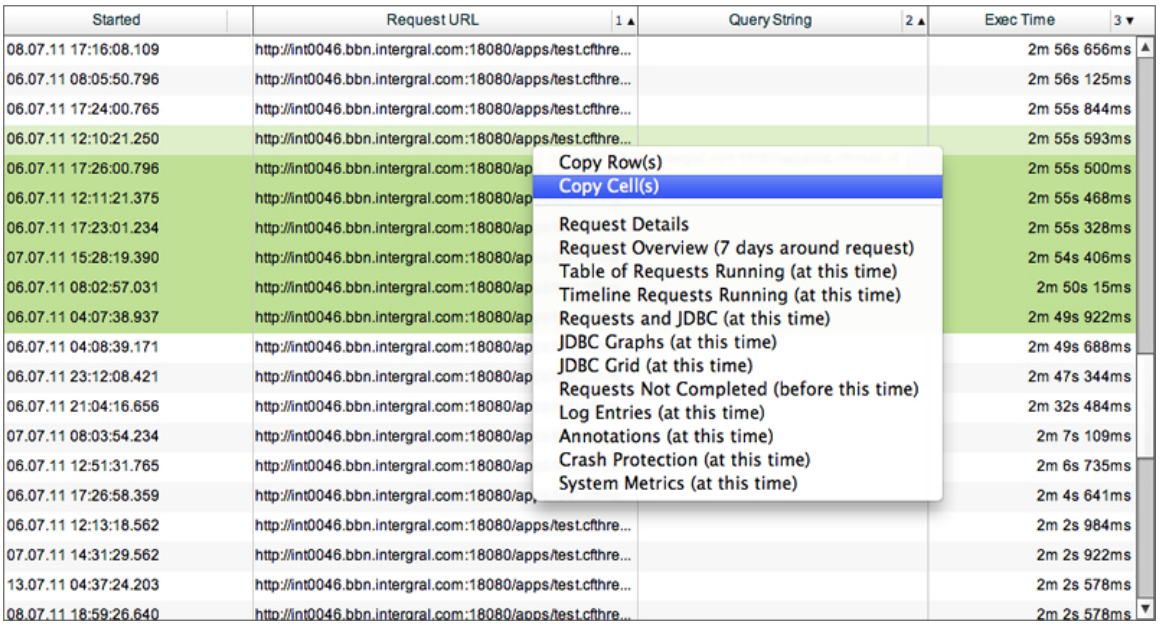

Figure 5: Copy Cells

The above ("Copy Cell(s)") menu item click will copy the selected cells to the clipboard where they can be further pasted resulting in:

```
http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm
http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm
http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm
http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm
http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm
http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm
http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm
```

Back to the top

Copying Grid Data to your Clipboard
See the Copying Grid Content section on this page to learn how to copy multiple sets of both Row and Cell data from the DataGrid Visualization to your Clipboard.

Next Steps

Date Navigation
An explanation of what Date Navigation is and how it can be used

Using Context Menus
An explanation of what the Context Menus are and how they works / are used

---

DataGrid Visualization - *Description of the DataGrid Visualization*
Chart Base Visualization - *Description of the Chart Base Visualization*
Time Block Chart Visualization -*Description of the Time Block Chart Visualization*
Placemark Chart Visualization -*Description of the Placemark Chart Visualization*
Sixway Chart Visualization -*Description of the Sixway Chart Visualization*

Take me back to Data Visualizations.

Chart Base Visualization

Chart Base Visualization

The Chart Base Visualization supports (and is the base for) all chart series types within the Application. The chart types that the application supports include:

- Area Series (Area Chart)
- Line Series (Line Chart)
- Column Series (Column Chart)
- Bar Series (Bar Chart)

The Chart Base Visualization is used to display / represent data within a chart / graph making data interpretation more "visible" and in some cases easier to analyze and understand. Like most charts, the Chart Base Visualization uses an "X" and "Y" axis to map data, and in turn represent data points from the Charts DataProvider in a two-dimensional manner rendered as a series (Line, Area, Column, Bar etc.).

Most Chart Based Visualizations within the application are "time-bound" Visualizations therefore mapping time-based data (time-stamps) to the chart axis allowing interaction with the Perspectives Timeline Control (if attached).

> **Note**: Each "time-bound" Chart Base Visualization uses a Chart Cursor that is controlled from the Timeline Control. This is the red line that is vertically rendered over the chart canvas and Timeline control. The cursor is controlled by dragging the handle at the bottom of the Timeline control which in turn updates and moves the cursor within the Chart Base Visualization. This is to help when comparing data across multiple charts within a Perspective.

The following image shows a typical Chart Base Visualization:

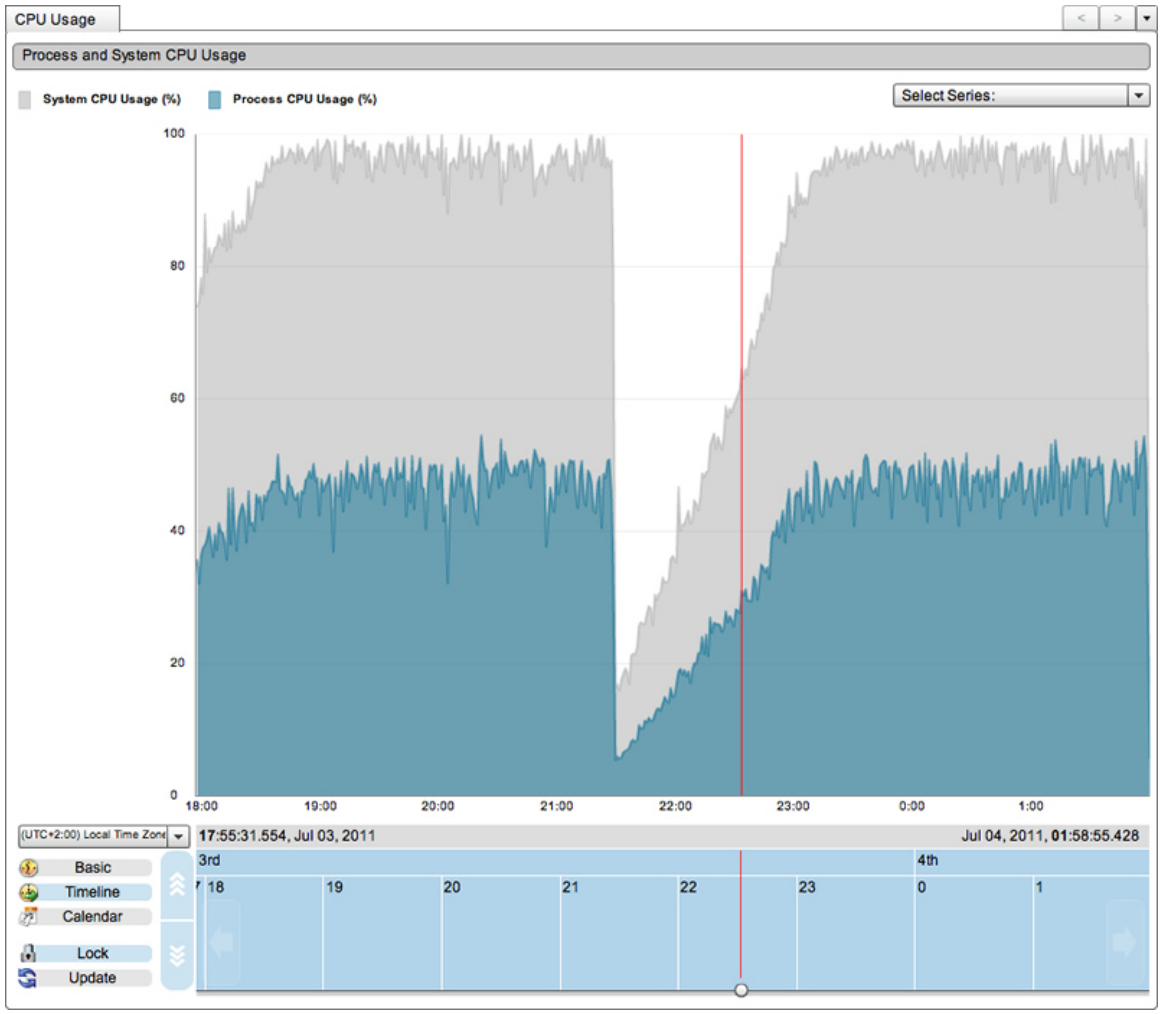

Figure 1: Chart Base Visualization

Back to the top

Chart Tool Tips

Each axis has labels attached that relate to the data being rendered within the chart at any time. Each data point within each chart series contains a Tool Tip that displays information for each data point in relation to the Chart axis and Series DataProvider.
It is easy to "roughly" see the value / time of a specific data point just by looking at its position in relation to the chart axis, however to know the exact information of a specific point, Tool Tips are required and prove to be very helpful for accurate analysis between data points:

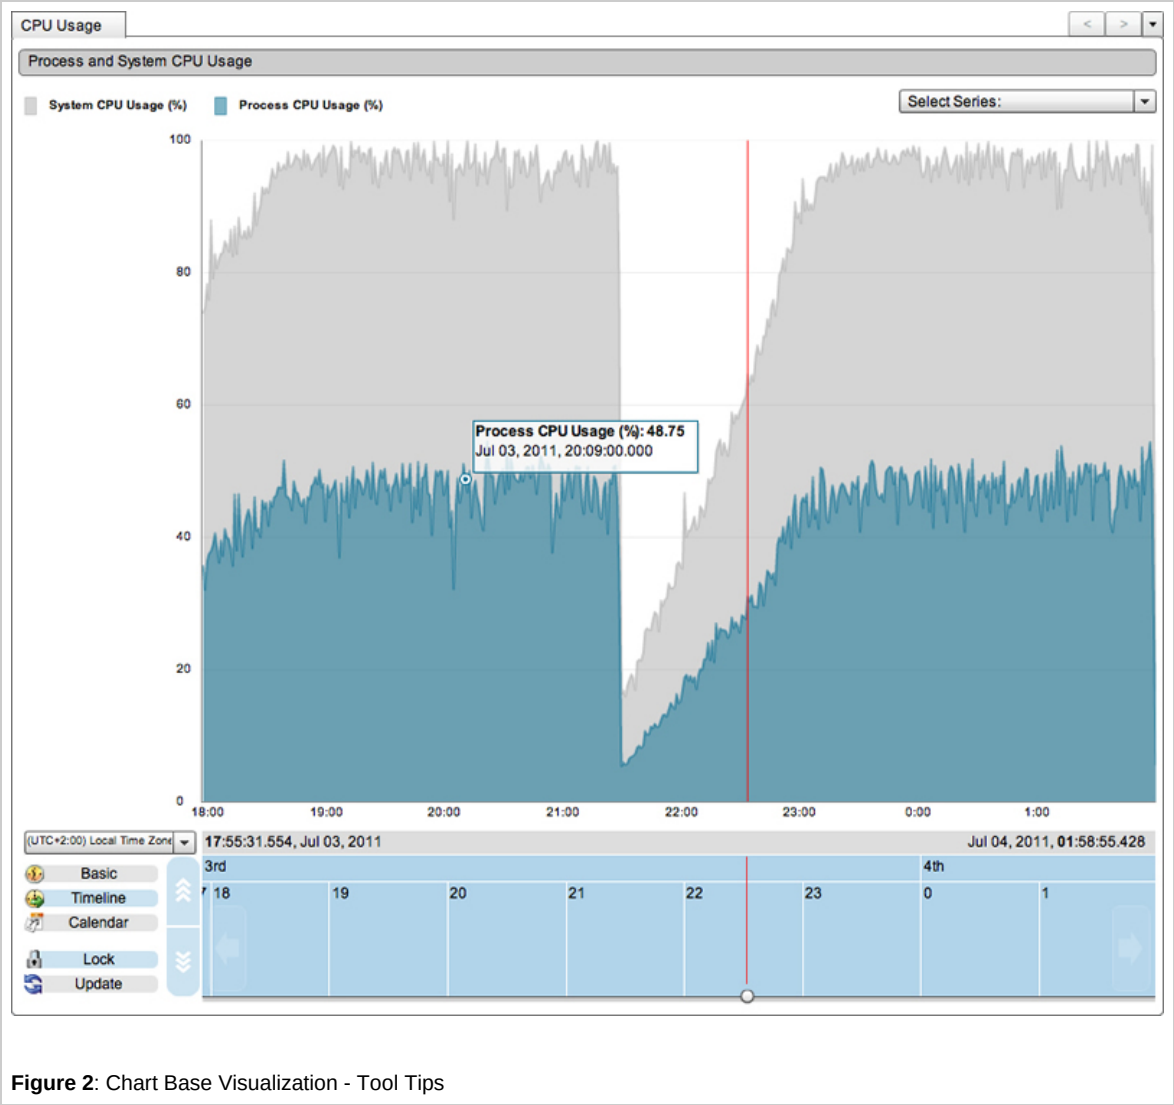

Figure 2: Chart Base Visualization - Tool Tips

Back to the top

Chart Legend / Series Select

The Charts include a (color-coded) legend in that is rendered to the top left of the chart. The legend represents each of the (currently) selected chart series, displaying the series name and measurement. This legend is linked to the Series Select Component that is rendered to the top right of the chart (if included). The Series Select Component controls what series are displayed at any one time. For example if a Chart Base Visualization contains many series, or if a series is rendered behind another (preventing it from being visible) the chart can often be difficult to interpret / analyze. In such cases the Series Select Component can be used to turn on and off (show / hide) specific series making analysis both easier and more accurate:

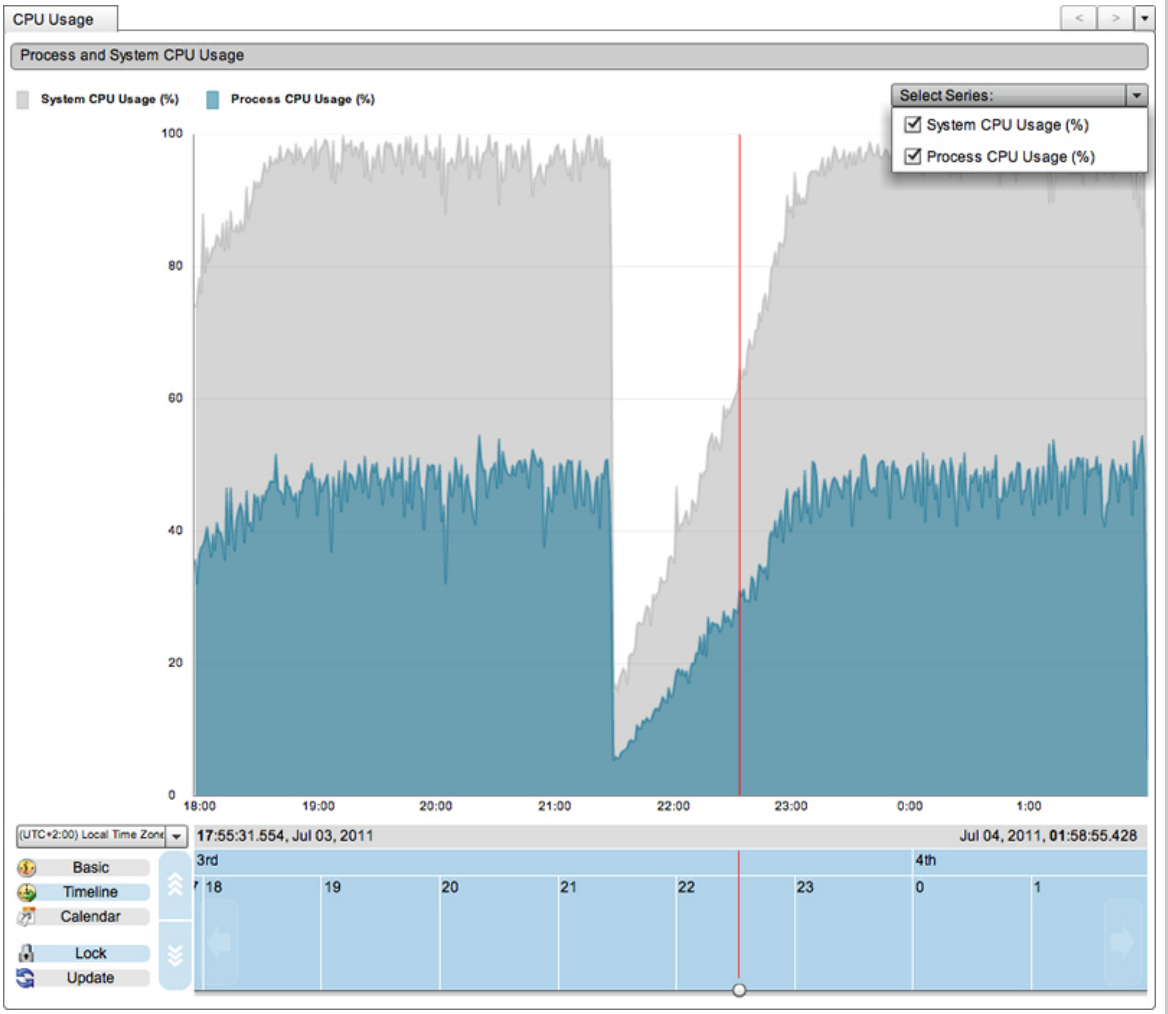

Figure 3: Chart Base Visualization - Series Select

To turn off / hide a series, simply open the Series Select Component combo box, find the series that you want to remove, and "un-check" it. This will update and remove the series from both the chart and the chart legend. To re-add the series follow the same process but select / "check" the series and in turn the chart will be updated and the series re-added and rendered again:

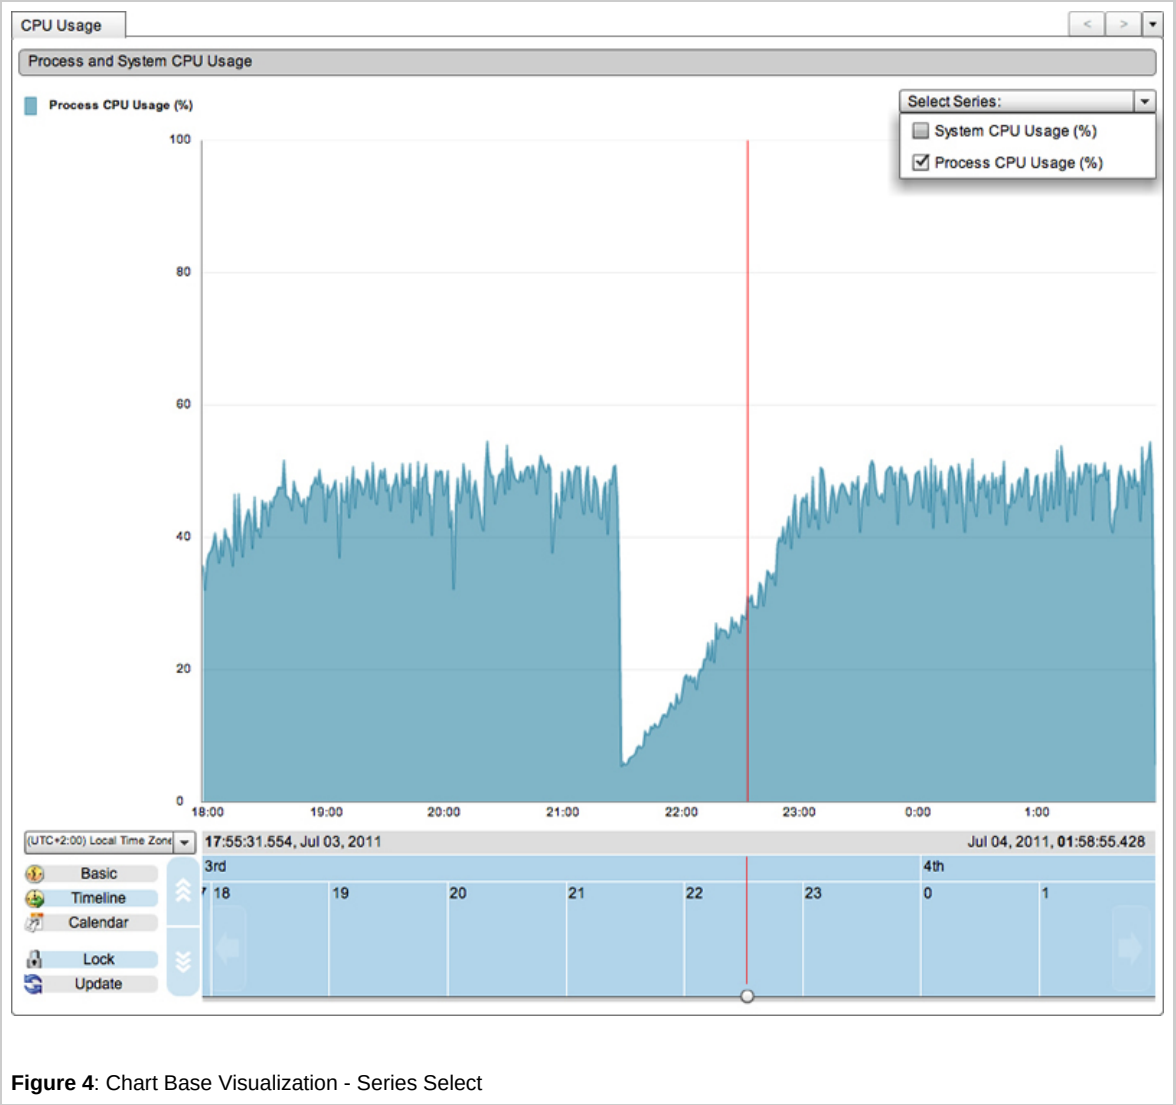

Figure 4: Chart Base Visualization - Series Select

Back to the top

Chart Context Menus

The Chart Base Visualizations support Context Menus. To use the Context Menus, simply right click a data point within a chart series and a Context Menu will be displayed. Each Context Menu has a DataProvider that is relevant to the Perspective / Visualization, and the Context Menu renders the each DataProvider item as a menu item that the user can click / open. As you can see, the Context Menu Items are generally "links" to other Perspectives relevant to the current Perspective (date / time) within the application:

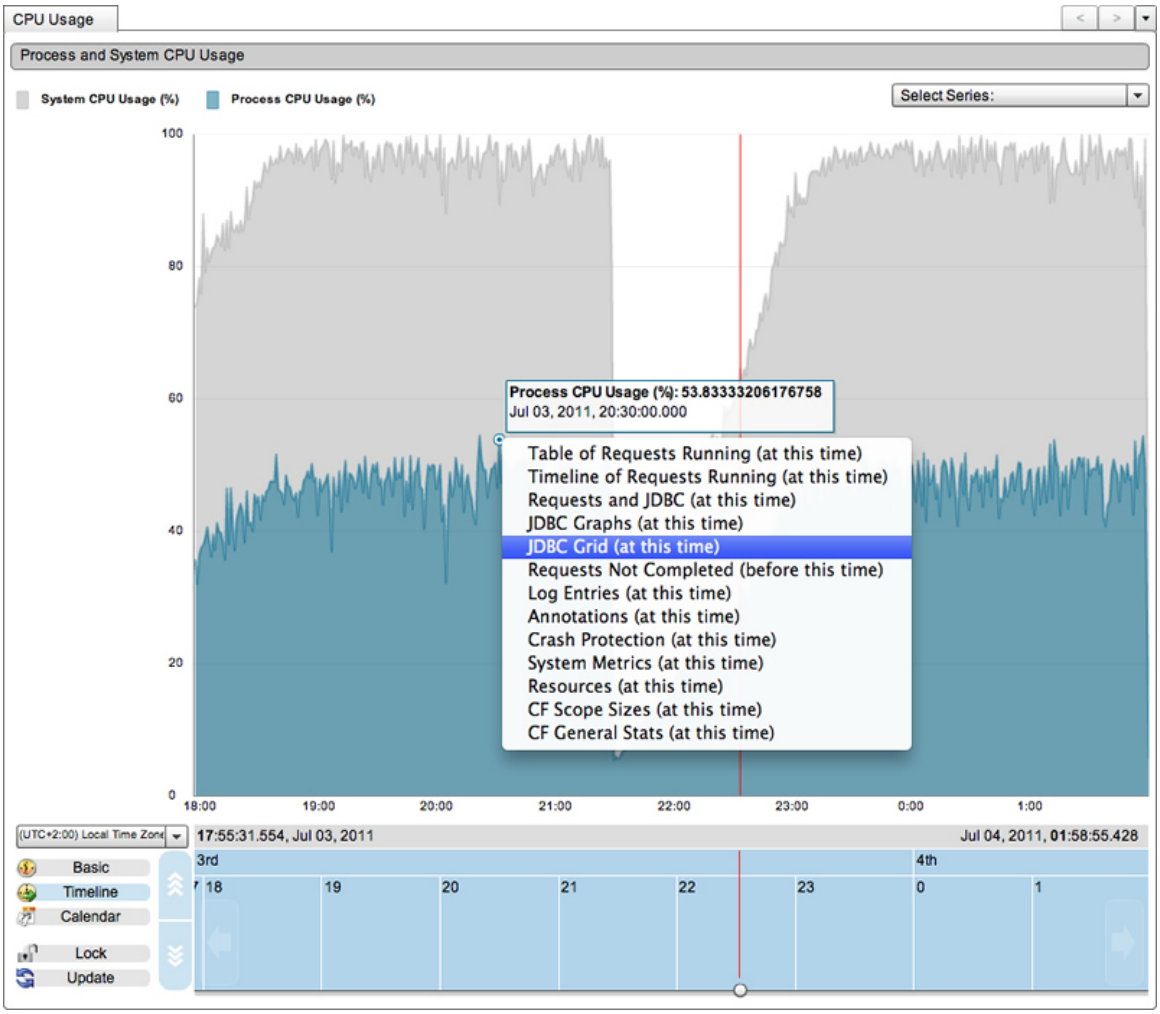

Figure 5: Chart Base Visualization - Context Menu |

Back to the top

Next Steps

Date Navigation
An explanation of what Date Navigation is and how it can be used

Using Context Menus
An explanation of what the Context Menus are and how they works / are used

---

DataGrid Visualization - *Description of the DataGrid Visualization*
Chart Base Visualization - *Description of the Chart Base Visualization*
Time Block Chart Visualization -*Description of the Time Block Chart Visualization*
Placemark Chart Visualization -*Description of the Placemark Chart Visualization*
Sixway Chart Visualization -*Description of the Sixway Chart Visualization*

Take me back to Data Visualizations.

Time Block Chart Visualization

Time Block Chart Visualization

The Time Block Chart Visualization is a very powerful chart allowing the rendering of many data points (a lot more than the standard charts) providing a larger and more detailed visualization (picture) of data at any one time further allowing more powerful and accurate data analysis and interpretation.

The Time Block Chart Visualization is used to display / represent data within as a "block" based representation providing different yet in depth and more detailed / advanced image of data, in turn making data both interpretation and analysis more accurate and problem idetification easier to both find and understand.

The Time Block Chart Visualization is "time-bound" mapping time-based data (time-stamps) to the chart axis allowing interaction with a Perspectives Timeline Control (if attached). The Time Block Chart Visualization uses both vertical and horizontal axis to map data. The vertical axis contains category data and the horizontal axis is mapped to time stamps, in turn representing data points from the Charts DataProvider rendered as Time Blocks in a two-dimensional manner.

> **Note**: As the time Block Chart is time-bound it includes a Chart Cursor which is controlled from the Timeline Control. This is the red line that is vertically rendered over the chart canvas and Timeline control. The cursor is controlled by dragging the handle at the bottom of the Timeline control which in turn updates and moves the cursor within the Time Block Chart Visualization. This is to help when comparing data across multiple charts within a Perspective.

The following image shows a typical Time Block Chart Visualization:

Figure 1: Time Block Chart Visualization

Back to the top

Chart Tool Tips

Each individual Time Block rendered within the Time Block Chart has Tool Tip assigned containing information about that individual Time Block.

It is easy to "roughly" see the time (duration) and category of each block just by looking at the Time Blocks positions in relation to the chart axis, however to know the exact category / time (duration) of a specific Time Block (especially in charts rendering lots of Time Blocks) and additional information related to the Time Blocks, Tool Tips are required and prove to be very helpful for accurate analysis between time Blocks:

Figure 2: Time Block Chart Visualization - Tool Tips

Back to the top

Copying Tool Tip Content

The Tool Tips used within the Time Block Chart Visualization are custom Tool Tips that allow more information than the standard Tool Tips along with Tool Tip formatting, scrolling (allowing for useable Tool Tip containing large amounts of data) and additionally the ability to copy the Tool Tip content. The ability to copy the content is both especially helpful and important for the Time Block Chart as it allows values to be copied and easily used when Filtering Data, again making data filtering and analysis much more efficient.

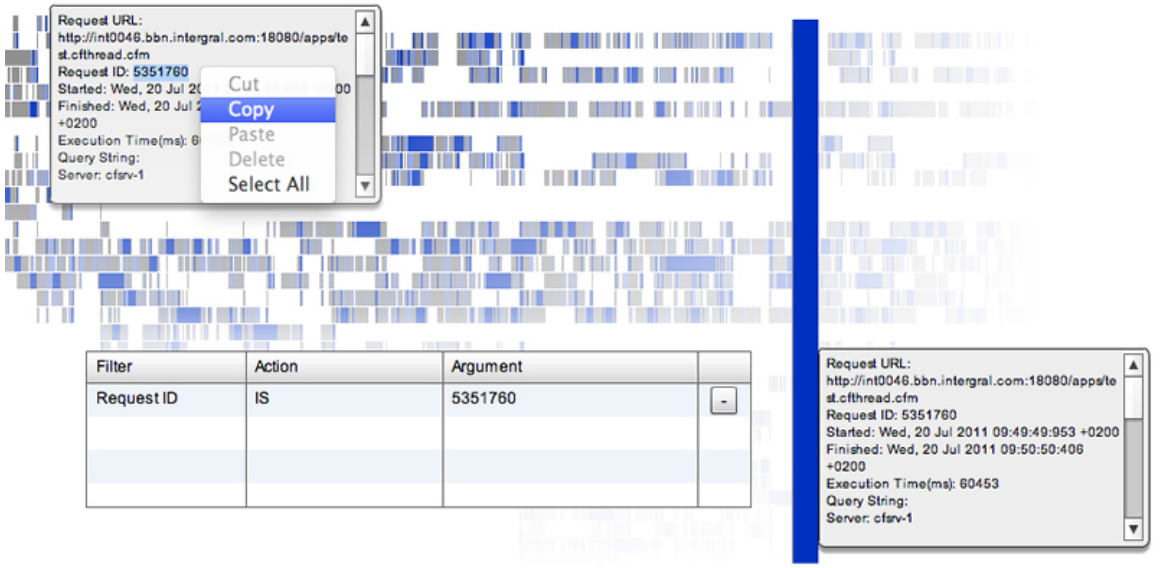

Figure 3: Time Block Chart Visualization - Copy Tool Tip

Back to the top

Chart Context Menus

The Time Block Chart Visualization supports Context Menus. To use the Context Menus, simply right click on a Time Block rendered within the chart and a Context Menu will be displayed. Each Context Menu has a DataProvider that is relevant to the Perspective and selected Time Block and the Context Menu renders each DataProvider item as a menu item within the Context Menu that the user can click / open. As you can see (in the screen shot), the Context Menu Items are generally "links" to other Perspectives relevant to the current Perspective (date / time) within the application.

Figure 4: Time Block Chart Visualization - Context menu

Back to the top

Chart Scrolling

To allow fast and understandable vertical rendering of large quantities of data, the Time Block Chart Visualization supports vertical scrolling allowing many data categories (rows) to be rendered within any one chart at one time. Chart scrolling is activated when the amount of rows (size) rendered within the chart is larger than the available chart real estate (height). When scrolling is activated, the chart can be navigated both vertically (using the scroll bar) and horizontally (using the Timeline Control).

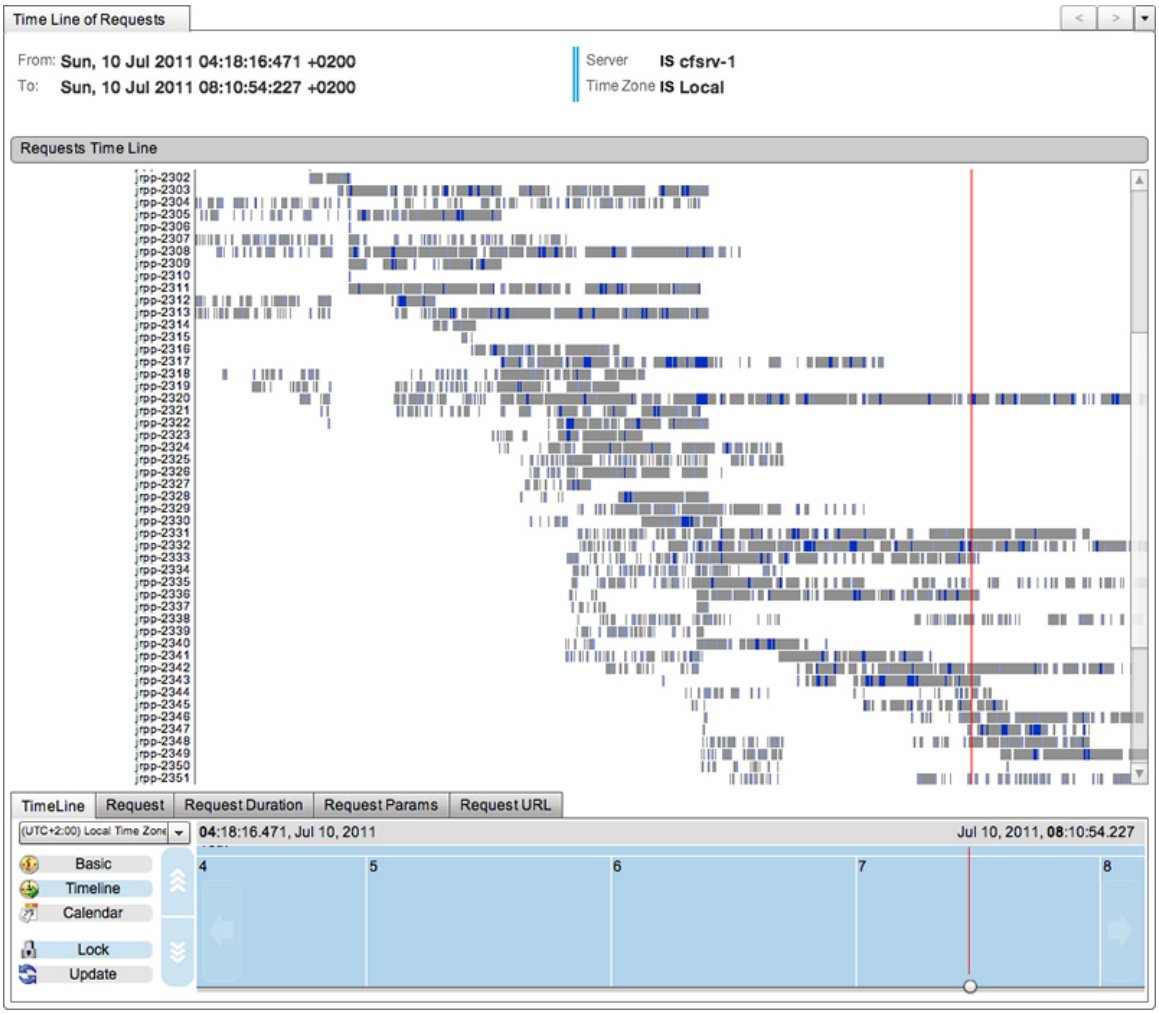

Figure 5: Time Block Chart Visualization - Chart Scrolling

Back to the top

On This Page:

Data Visualizations:

Copying Tool Tip Content
See the Copying Tool Tip Content section on this page to learn how to copy tool tip content. The ability to copy the content very helpful as it allows values to be easily used when Filtering Data.

Also See

Date Navigation
An explanation of what Date Navigation is and how it can be used

Using Context Menus
An explanation of what the Context Menus are and how they works / are used

Filtering Data
An explanation of the data Filters, how they work and how they can be used

---

DataGrid Visualization - *Description of the DataGrid Visualization*
Chart Base Visualization - *Description of the Chart Base Visualization*
Time Block Chart Visualization -*Description of the Time Block Chart Visualization*
Placemark Chart Visualization -*Description of the Placemark Chart Visualization*
Sixway Chart Visualization -*Description of the Sixway Chart Visualization*

Take me back to Data Visualizations.

Placemark Chart Visualization

Placemark Chart Visualization

The Placemark Visualization allows many points to be displayed / represented in an easy, understandable (visual) and fast (high performance) manner.

The chart is "time-based" containing one horizontal axis which is linked to and controlled by the Perspective Timeline Control (if in place) allowing the user to efficiently navigate between time periods to see / find different Placemarks.

> **Note**: Each "time-bound" Placemark Chart Visualization uses a Chart Cursor that is controlled from the Timeline Control. This is the red line that is vertically rendered over the chart canvas and Timeline control. The cursor is controlled by dragging the handle at the bottom of the Timeline control which in turn updates and moves the cursor within the Chart Base Visualization. This is to help when comparing data across multiple charts within a Perspective.

The following image shows a typical Placemark Chart Visualization:

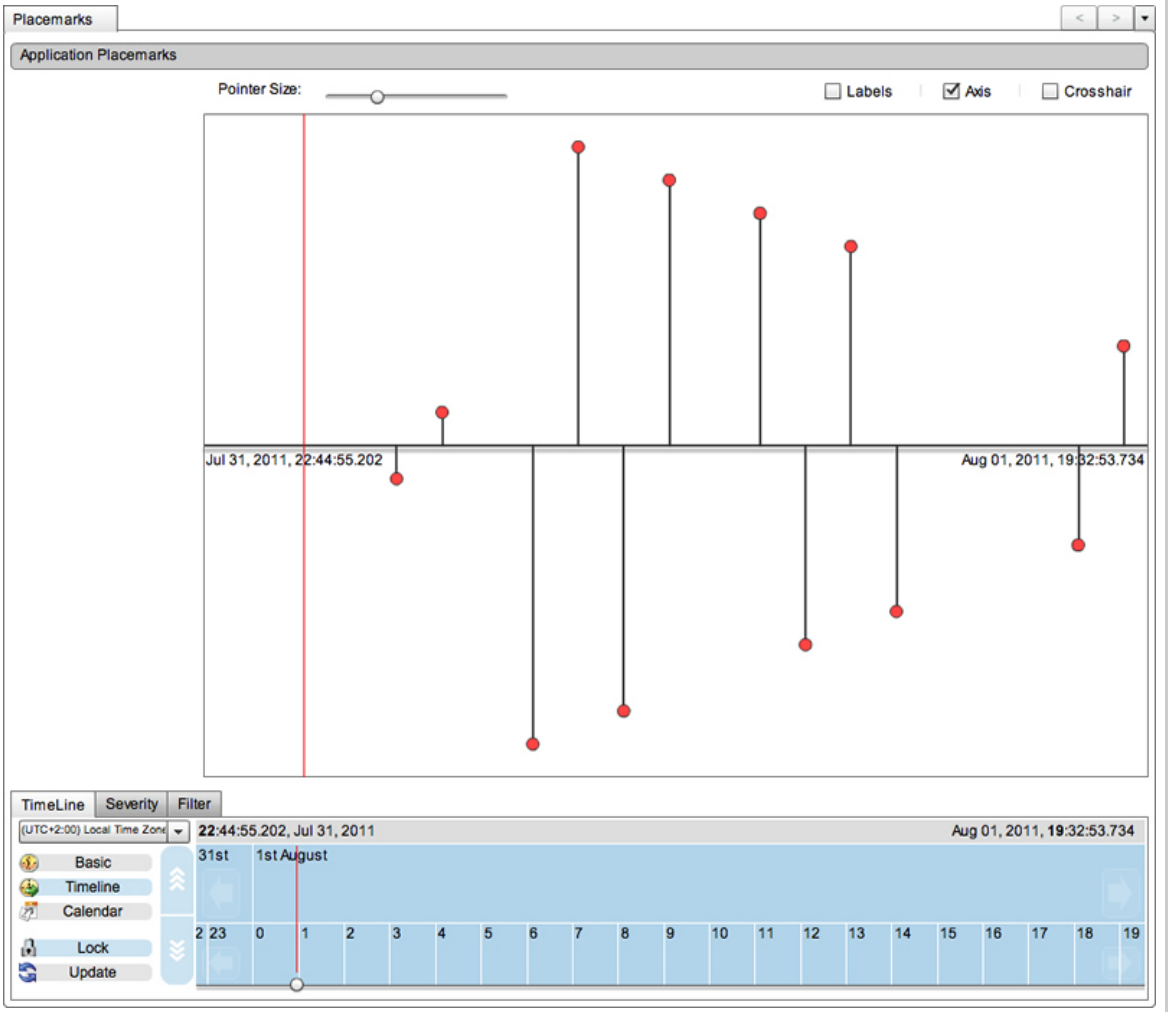

Figure 1: Placemark Chart Visualization

Back to the top

Placemark Grouping

The placemark Chart Visualization supports "grouping" allowing placemearks to be rendered both efficiently and in a manner that is easy to read / evaluate, e.g. If the canvas contains too many placemarks to render within a selected time period, instead of rendering all placemarks (which would make the chart confusing) the placemarks within a similar time period and of a similar type (E.g. Info, Error, Warning etc) are grouped together therefore rendering less placemarks on the screen and making analysis easier. If an item contains grouped placemarks the items label (or Tool Tip if labels are turned off) will display the type and number of the grouped items:

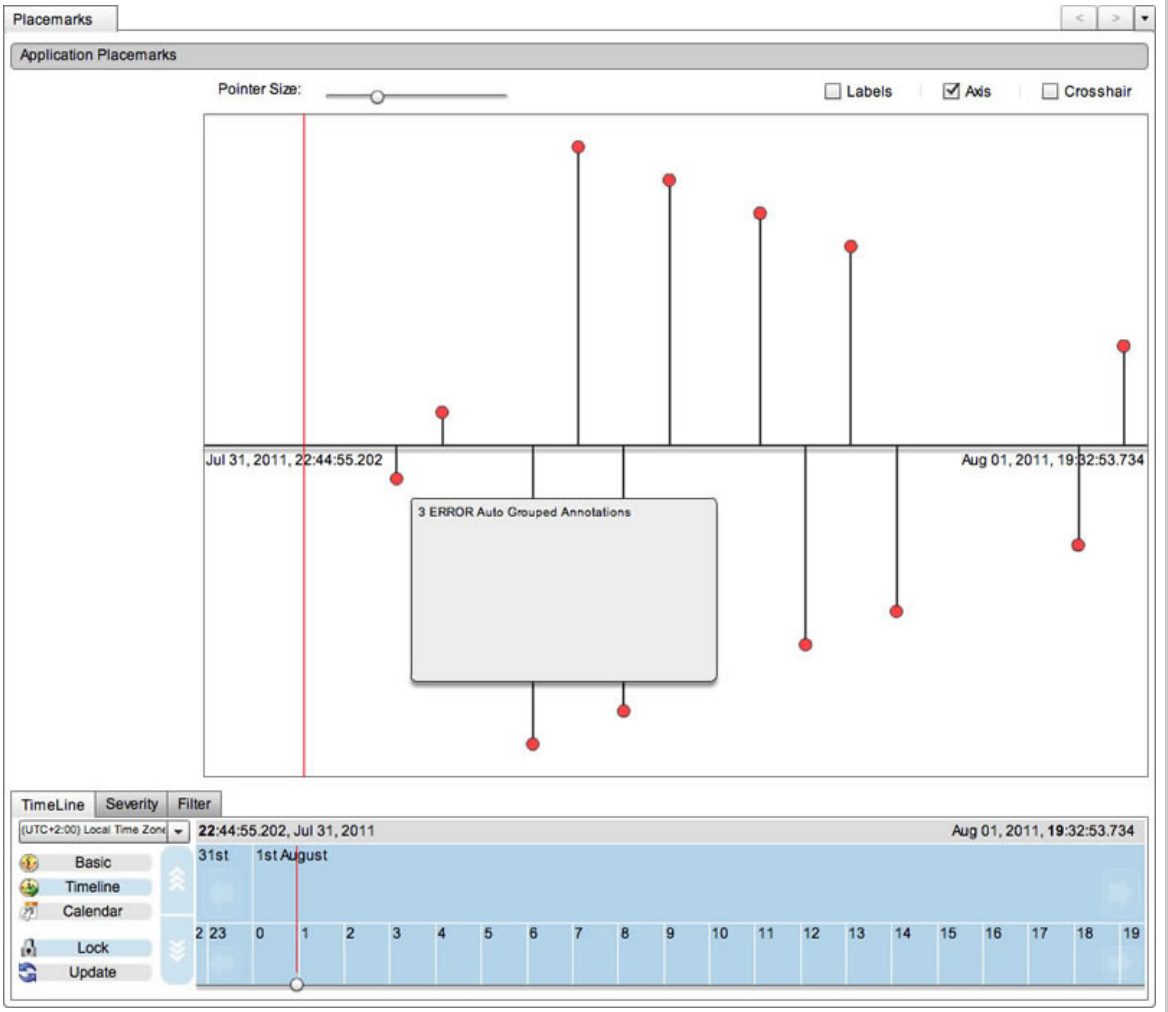

Figure 2: Placemark Chart Visualization - Grouping

Each grouped item within the chart can be zoomed into using the Timeline Control until the chart is at a level that it can render the grouped items individually. At this point a switch will be made and the grouped items will be un-grouped and rendered individually

> ⚠ The chart supports multi-levels of grouping, e.g. items can be both grouped and sub-grouped (grouped items containing additional grouped items). In such cases multiple levels of switches (zooming) will be required to see all of the individual placemarks.

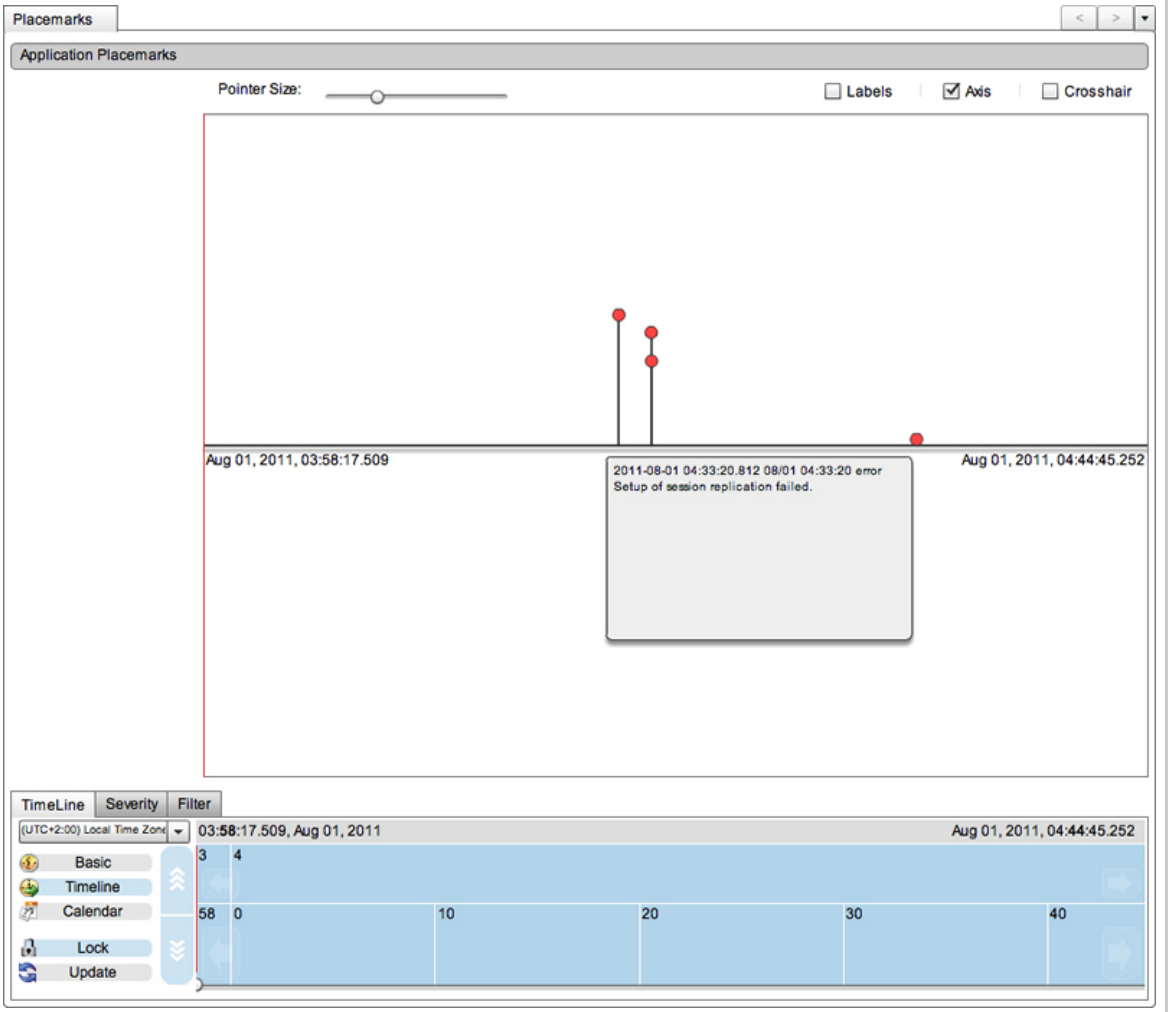

Figure 3: Placemark Chart Visualization - Grouping

Once placemark items are rendered at there lowest (individual) level the labels (or Tool Tips) of the placemarks will be relevant / distinct to each item.

Back to the top

Adding Placemarks

The Placemanrk Visualization allows users to add custom placemarks to the chart for specific points in time. These placemarks are stored into the application database and retrieved / rendered whenever the chart is loaded within the "saved" placemark time frame.
To add a placemark to the chart, simply right click on the charts canvas at the time (horizontal) and position (vertical) where you want to save the placemark, and select the "Add Placemark" menu item from the displayed context menu:

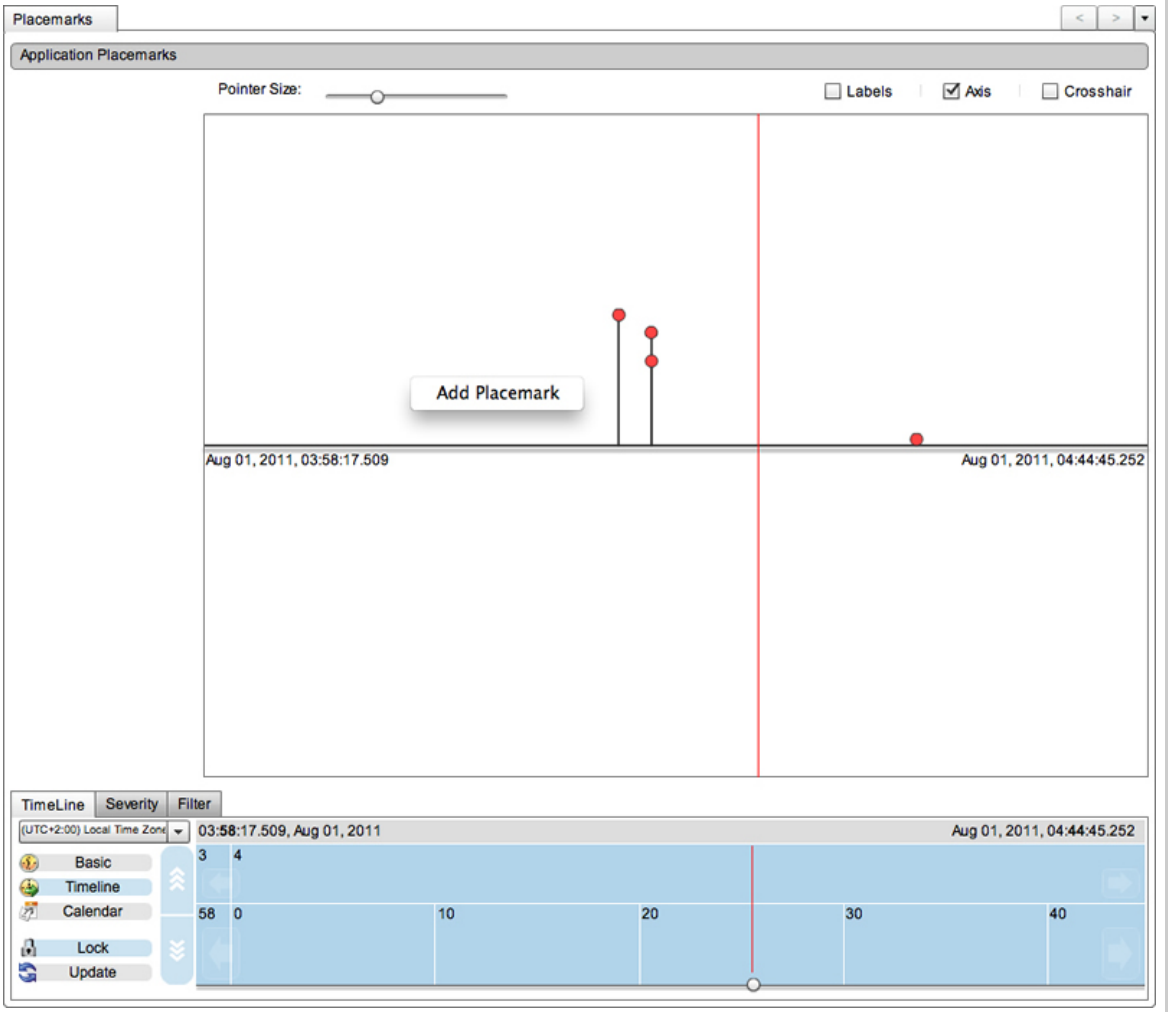

Figure 4: Placemark Chart Visualization - Add Placemark

This will show the "Add Placemark" dialogue where the user can add details for the "new" placemark:

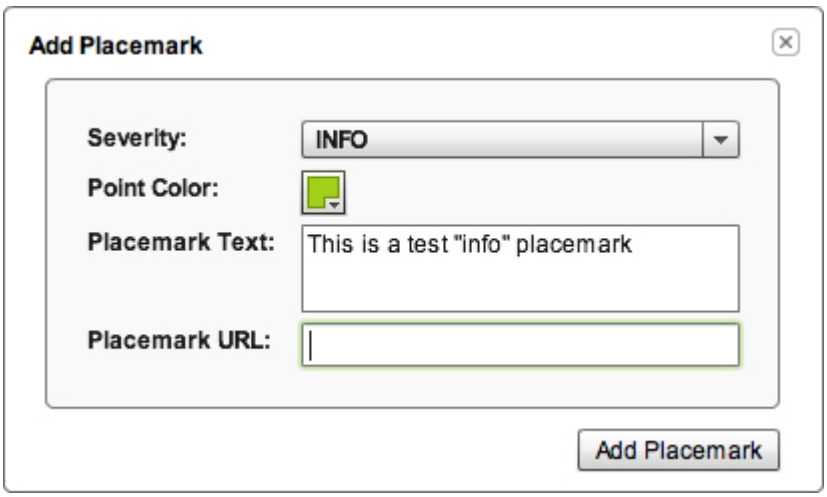

Figure 5: Add Placemark Panel

When adding a new placemark you can enter the following information:

Item	Description
Severity	Choose a severity for the Placemark. The options available include: **INFO** (Information), **WARN** (Warning) and **ERROR** (Error)

Point Color	Choose a color for the placemart from the color pallet provided in the dialogue. Different colors can be chosen allowing to help categorize / distinguish different placemarks and placemark types
Placemark Text	Enter some text / description for the tool tip. This is the text that will be rendered within the placemark label (or tooltip)
Placemark URL	Optional: Enter a URL for the placemark. When the placemark is double clicked the URL will be opened within a new tab within the application

When you have finished adding the details for the new placemark, click the "Add placemark" button to remove the dialogue and add the placemark. You will notice that the placemark has been added in the exact location (that was clicked) on the canvas, the fill color of the placemark is the color that was selected and the placemark label (Tool Tip) is the text that was entered. The placemarks have double click actions enabled, so if a URL was added to the the placemark, you can double click the item and the URL will be loaded in a new tab window within the application

> ⚠ If no URL was added a double click event will have no effect.

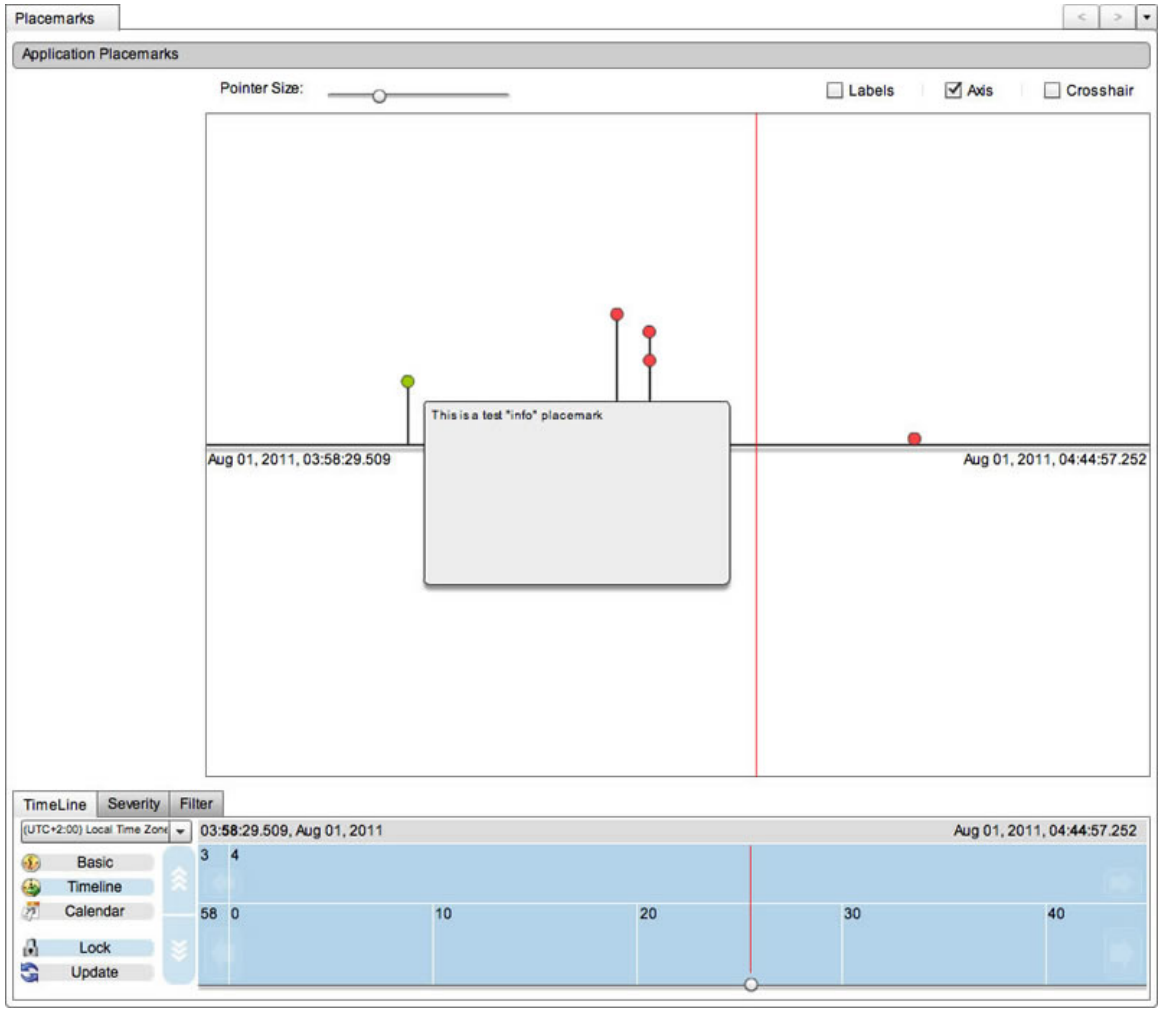

Figure 6: Placemark Chart Visualization - Added Placemark

Back to the top

Editing and Deleting Placemarks

To edit or delete a custom placemark, simply right click on the relevant placemark and select the "Edit Placemark" menu item from the context menu:

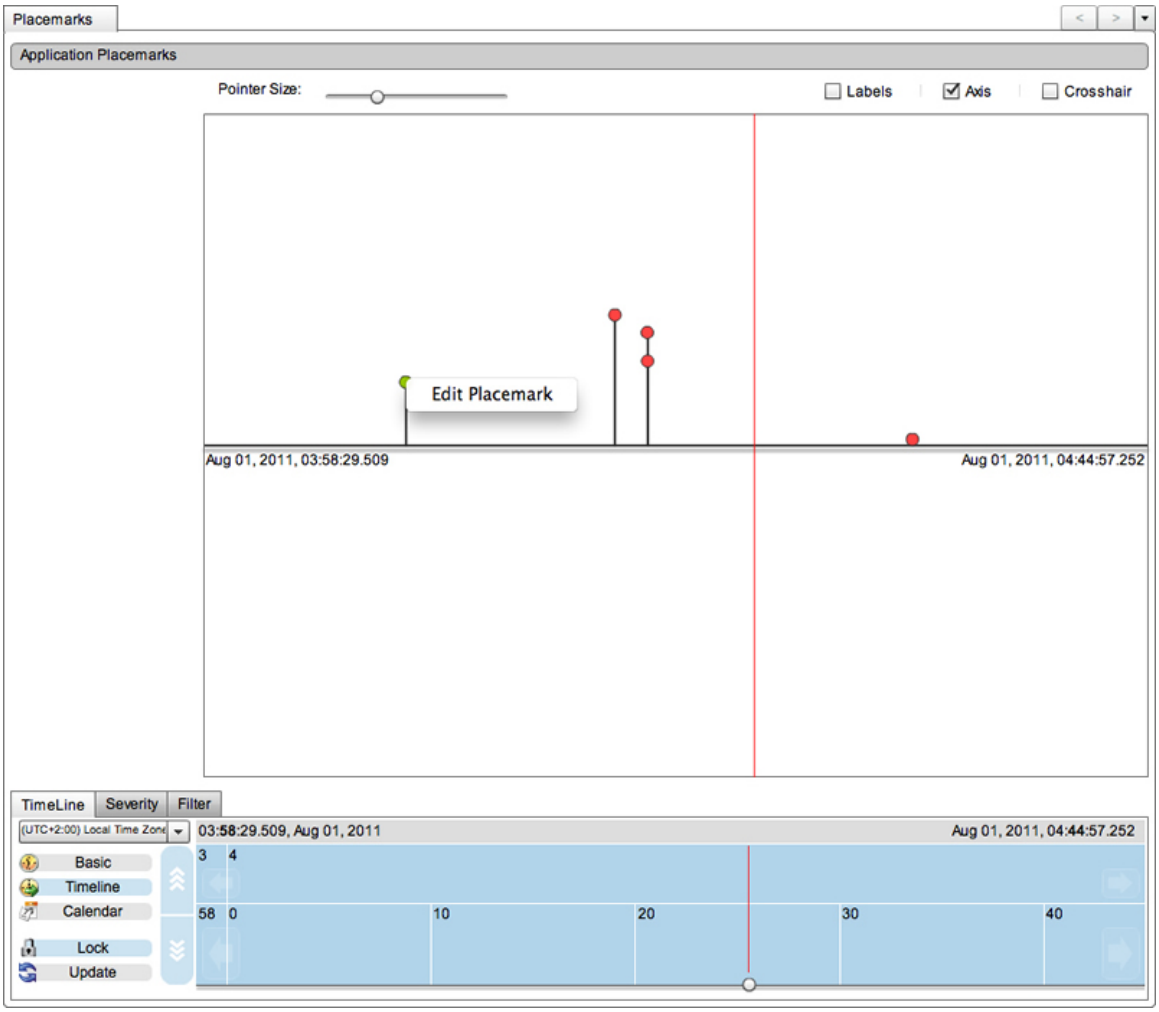

Figure 7: Placemark Chart Visualization - Edit Placemark

This will display the "Edit Placemark" dialogue:

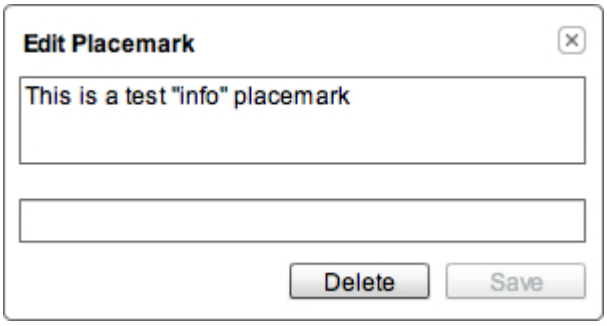

Figure 8: Edit Placemark Panel

The "Edit Placemark" dialogue will display the information relevant to the selected placemark item. Within this dialogue you can edit / change the placemark (description) and URL. click the "Save" button to save / modify the item and return to the placemark canvas. Alternatively you can also choose to delete the placemark by clicking this "Delete" button which will delete the item from both the the database and the canvas.

Back to the top

Placemark Control Bar

Above the placemark canvas is the placemark control car containing a pointer slider and label, axis and crosshair check boxes.

Back to the top

Pointer Size Slider

The Pointer Slider controls the size of the placemarks that are rendered on the canvas. The ability to change the placemark size is useful as it allows the user to change the placemark size depending on the real estate available and the amount of placemarks being rendered at any time for easier chart readability.

Back to the top

Show Labels Check

The Labels Check controls weather or not the placemark labels are displayed. If there are a lot of placemarks being rendered at anyone time, including grouped placemarks, it can affect the charts performance therefore the labels can be turned off while scrolling / navigating through data, and then turned on again once the required time frame has been found.

> ⚠ Even when labels has been turned off, the placemark text is still available within the individual placemark tool tip.

Back to the top

Show Chart Axis Check

The Axis Check controls whether or not the axis is displayed. Again this is useful in regards to rendering. If there are a lot of placemarks rendered at anyone time and / or a placemark is rendered in top of / over the axis then it can sometimes be hard to read. In such cases the chart axis can be switched off.

> ⚠ When the chart axis is turned off chart performance can also be improved, although only really noticeable when a lot of placemarks are being rendered.

Back to the top

Show Crosshair Check

The Corosshair Check controls the visibilily of the crosshair cursor. The crosshair cursor follows the mouse pointer drawing both a horizontal line across the mouse pointer (the width of the canvas) and a vertical line through the pointer (the height of the canvas) with the two lines meeting / crossing in the middle. The crosshair is useful if placemark positions are important e.g. for identifying alignment patterns between them.

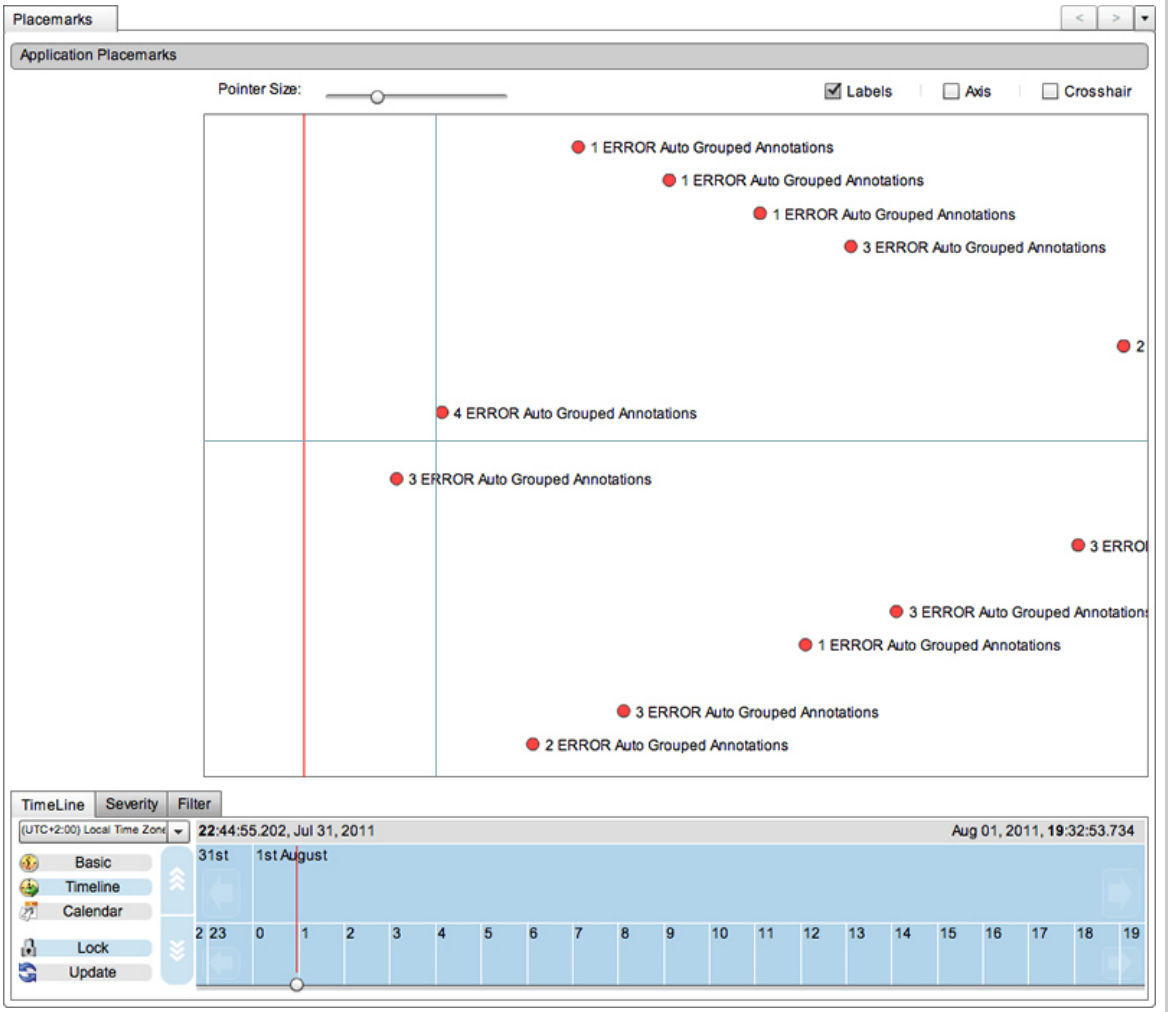

Figure 9: Placemark Chart Visualization - Crosshair

Back to the top

On This Page:

Data Visualizations:

DataGrid Visualization

Chart Base Visualization

Time Block Chart Visualization

Placemark Chart Visualization

Sixway Chart Visualization

Next Steps

Date Navigation
An explanation of what Date Navigation is and how it can be used

---

DataGrid Visualization - *Description of the DataGrid Visualization*
Chart Base Visualization - *Description of the Chart Base Visualization*
Time Block Chart Visualization -*Description of the Time Block Chart Visualization*
Placemark Chart Visualization -*Description of the Placemark Chart Visualization*
Sixway Chart Visualization -*Description of the Sixway Chart Visualization*

Take me back to Data Visualizations.

Sixway Chart Visualization

Sixway Chart Visualization

The SixWay Chart Visualization contains 6 Area Charts which are (by default) rendered evenly in a 3x2 grid structure (when the component is initialised). The contained charts include:

- Request Activity Chart
- Database Request Activity Chart
- Memory Usage Chart
- Average Request Time Chart
- Average DB Time Chart
- CPU Usage Chart

The following image shows the SixWay Chart Visualization in its initialized state:

Figure 1: Sixway Chart Visualization

Back to the top

Chart Transitions

The SixWay Chart Visualization is transitional and each individual chart within the Visualization can be maximized for better visibility.

To maximize one of the charts simply click the chart that you want to maximize and the selected chart will be maximized to the size of the allocated real estate and positioned to the right. The remaining five charts will be minimized and positioned vequally and laid out vertically to the left of the selected chart (see screen shot below).To maximize another chart simply click one of the minimized charts on the left, and the selected chart will then be maximized into the space on the right replacing the current maximized chart which in turn will be minimized down back to the left:

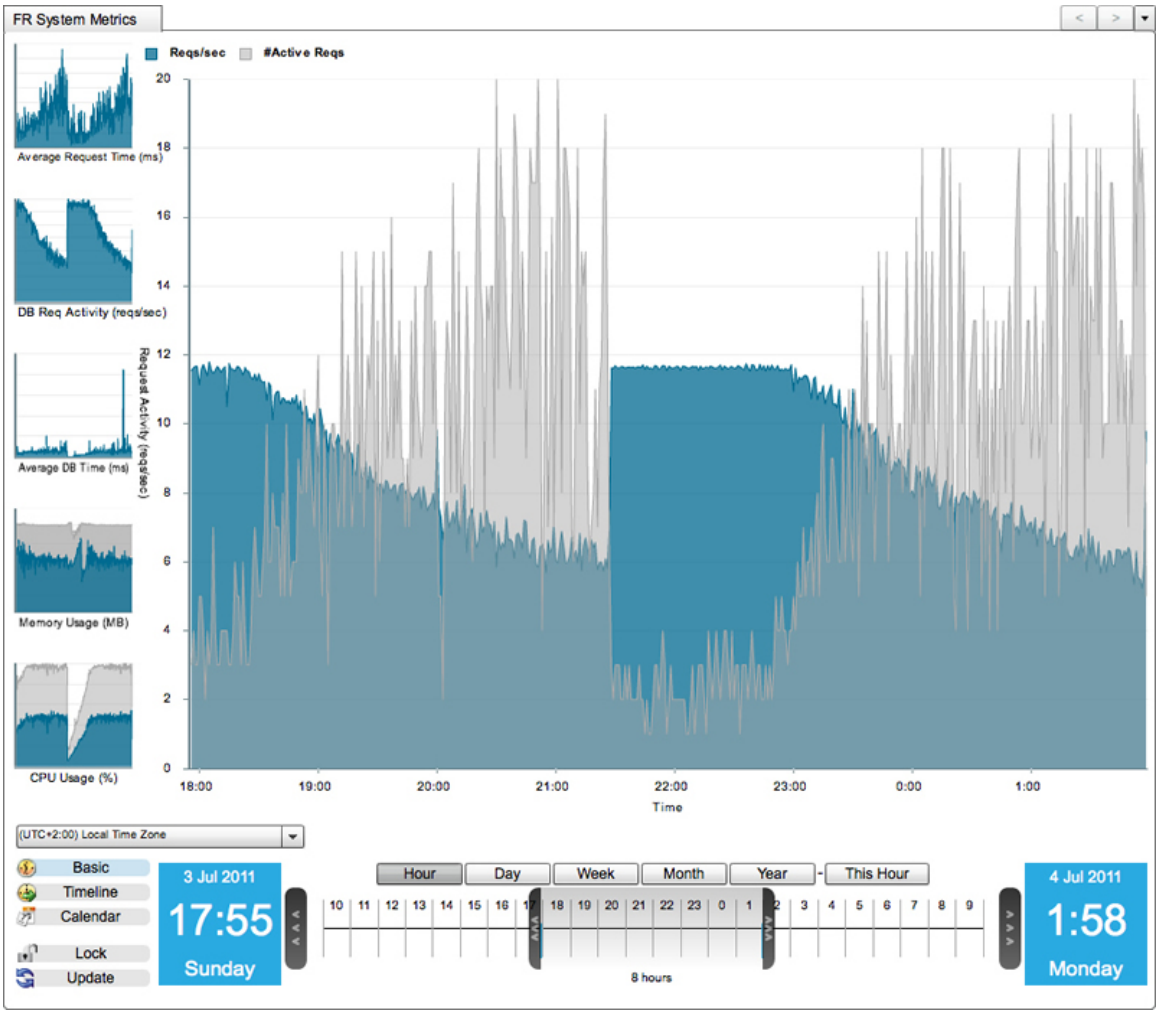

Figure 2: Sixway Chart Visualization - Chart Transitions

To return to the default layout (with all charts of equal space and size) simply click the current maximized chart to return.

Back to the top

Chart Context Menus

The SixWay Chart Visualization supports Context Menus. To use the Context Menus, simply right click on a data point within one of the six charts chart series and a Context Menu will be displayed. Each series within each of the 6 charts has its own Context Menu DataProvider relevant to the Perspective / Visualization. The Context Menu renders the each DataProvider item as a menu item within the Context Menu that the user can click / open. As you can see, the Context Menu Items are generally "links" to other Perspectives relevant to the current Perspective (Date / time) within the application.

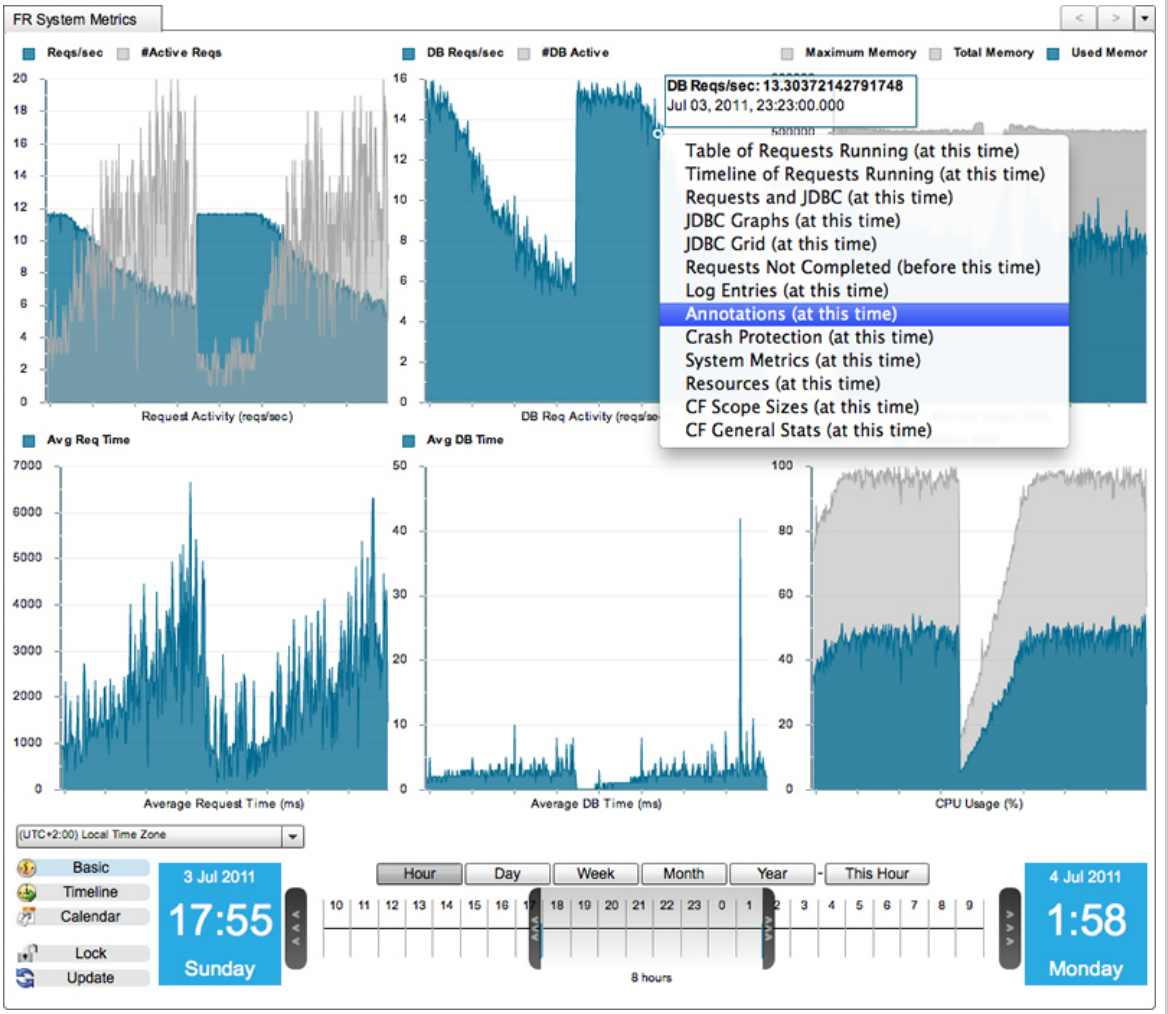

Figure 3: Sixway Chart Visualization - Context Menu

For more information about the SixWay Chart Visualization and its content / usage please see the Overview and FusionReactor System Metrics Perspectives pages along with the Chart Base (the individual charts within the SixWay Chart Visualization are based upon a "light weight" version of the Chart Base).

Data Visualizations:

DataGrid Visualization

Chart Base Visualization

Time Block Chart Visualization

Placemark Chart Visualization

Sixway Chart Visualization

Next Steps

Date Navigation
An explanation of what Date Navigation is it can be used

Using Context Menus
An explanation of what the Context Menus are and how they works / are used

DataGrid Visualization - *Description of the DataGrid Visualization*
Chart Base Visualization - *Description of the Chart Base Visualization*
Time Block Chart Visualization -*Description of the Time Block Chart Visualization*
Placemark Chart Visualization -*Description of the Placemark Chart Visualization*
Sixway Chart Visualization -*Description of the Sixway Chart Visualization*

Take me back to Data Visualizations.

Application Menus

Application Menus

The application menus can be found at the top of the FusionAnalytics client. Click on a menu to open it:

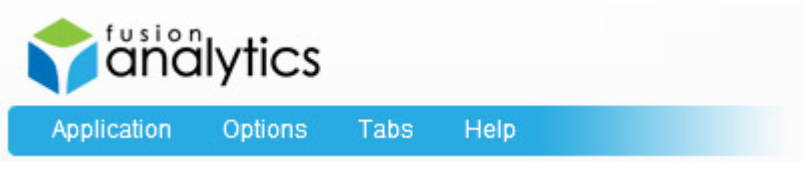

Figure 1: Application Menus

Application

The Application menu contains three elements:

About	Show the FusionAnalytics About dialog box.
Logout	In the AIR client, Logout will return you to the server browser screen. In the Web client, Logout will return you to the login screen for the current server.
Exit	This is only available in the AIR client and will completely close the application.

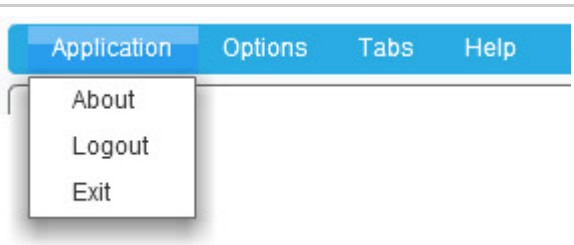

Figure 2: Application Menus > Application

Back to the top

Options

The Options menu has two available options:

Reset Cache	Much like an internet browser, the FusionAnalytics client caches some information in order to improve performance. Just like an internet browser, this option allows you to throw away all of this cached information to make sure you are seeing the most recent data from the server.
State Store	Selecting this option brings up the State Store dialog.

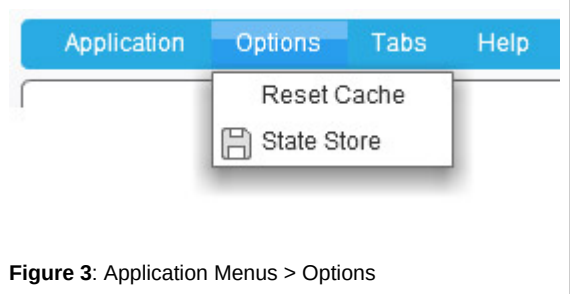

Figure 3: Application Menus > Options

Back to the top

Tabs

The Tabs menu is only available if you have tabs open. If you close all tabs then this menu will be greyed out.

Print Tab	Only available in the AIR client. If you select this menu item then a print dialog will be displayed in order to print the contents of the currently selected tab.
Expert Tab to PDF	Only available in the AIR client. If you select this menu item then an export dialog will be displayed in order to create a PDF file from the contents of the currently selected tab.
Close Tab	Close the currently selected tab.
Close Others	Close all tabs but leave the currently selected tab open.
Close All	Close all tabs.

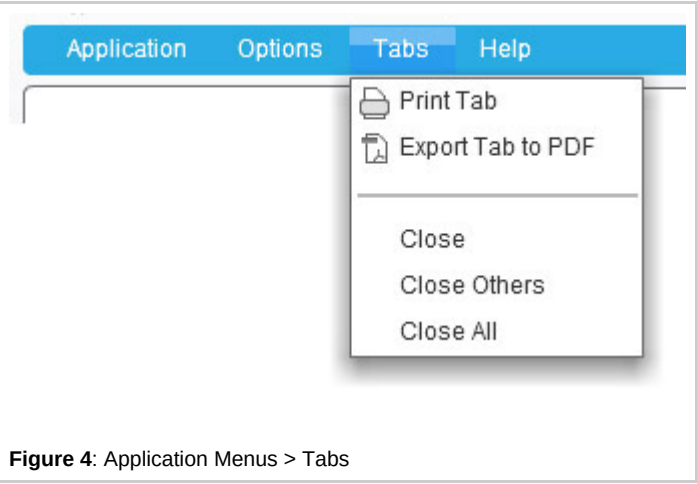

Figure 4: Application Menus > Tabs

Back to the top

Help

The Help menu has two menu items:

About	Show the FusionAnalytics About dialog box.
FusionAnalytics Help	Open the FusionAnalytics online documentation in a web browser window.

Figure 5: Application Menus > Help

Back to the top

On This Page:

Next Steps

Using State Stores
Learn about saving States within your Application using the State Store

Using Tabs
Learn how to use the Application Tab Navigator

The FusionAnalytics AIR Client
Learn how to access multiple different FusionAnalytics Applications using the FusionAnalytics AIR Client

The FusionAnalytics Web Client
Learn how to log into and use the FA Web Client

Using Tabs

Using Tabs

Whenever you click on an item from the Table of Contents (TOC) a new perspective will be opened in the content area. This perspective will be opened in a new tab. Tabs appear along the top of the content area and can be interacted with in various ways. The simplest way to switch between tabs is simply to click on the tab you want to view. The content of that tab will then be made visible.

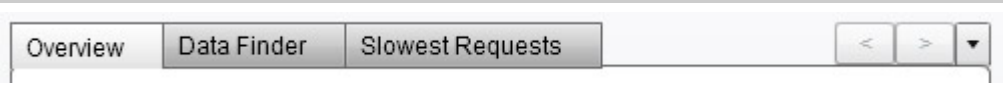

Figure 1: Tab Navigator

If you look at the far right of the tab bar then you will see a button with a down arrow on it. Click this to display a drop down list of currently opened tabs. The currently selected tab will have a marker next to it. You can click on an item within that list in order to open that tab.

Figure 2: Currently Selected Tab

If you have opened more tabs than can be displayed across the top of the content area then you will be able to scroll through them by using the left and right buttons which appear next to the drop down list.

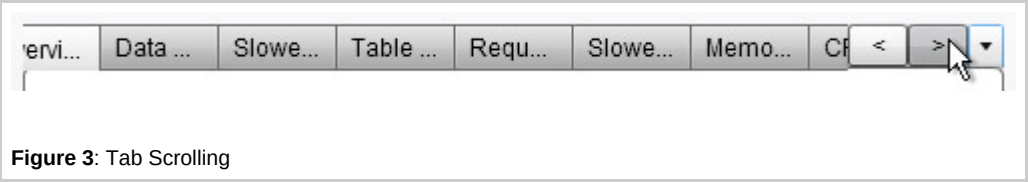

Figure 3: Tab Scrolling

If you want to close one of the tabs then move your mouse over it. When the mouse moves over a tab a close button will appear. Click on that button and the tab will be removed.

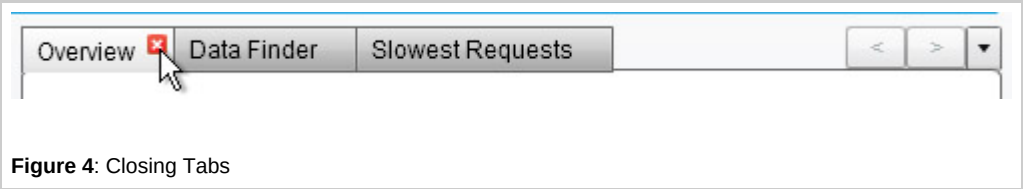

Figure 4: Closing Tabs

If you right-click on a tab then you will see a context menu. From here you can either close a tab, close all tabs or close all but the tab you right-clicked on.

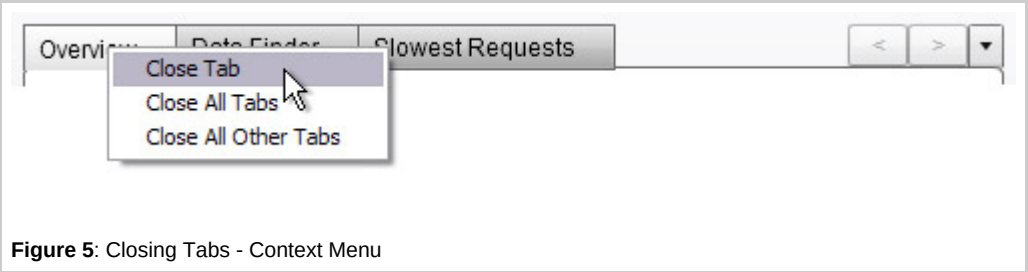

Figure 5: Closing Tabs - Context Menu

The tab closing menu items can also be accessed from the top level Application Menus. If you select "Close Others" from the application menu then all tabs will be closed except the currently selected item.

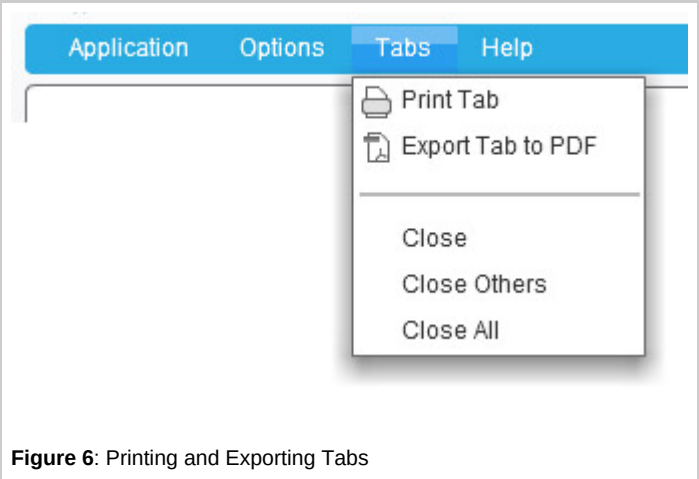

Figure 6: Printing and Exporting Tabs

If you are using the AIR client then you will also have Print Tab and Export Tab to PDF options available to you. These items are not available from within the Web Client. For more information on the Print and Export functionality, see the Application Menus section.

Back to the top

TOC and Content Layout
The application canvas is split into two columns. The TOC is rendered in the left column and the Perspectives (content) are loaded in the right column:

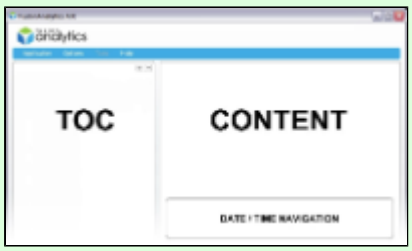

Next Steps

Using the Table of Contents (TOC)
Learn about the client layout and Using the Table of Contents (TOC)

Application Menus
Learn how to use the Application Menus to close all tabs

The FusionAnalytics AIR Client
Learn how to access multiple different FusionAnalytics Applications using the FusionAnalytics AIR Client

The FusionAnalytics Web Client
Learn how to log into and use the FA Web Client

Using State Stores

Using State Stores

The State Store feature is designed to allow you to store the current state of a perspective so that you can return to it at a later point in time. The State Store will save the current date range as well as any currently defined filters.

You can open the State Store dialog box using the Application Menus - Options > State Store :

Figure 1: State Store Panel

Back to the top

Creating a New State Store

In order to create a new State Store you have to enter a name, choose whether your State Store will be visible to all users who view this perspective, or just visible to the current user, and then click the Save State button.

Figure 2: Creating a new State Store

Once you have saved your state store it will appear in one of the boxes below. Public State Stores are listed first, followed by Private State Stores below.

Back to the top

Loading, Updating and Deleting States

If you close the State Store dialog box and the perspective, then re-open the perspective and open the State Store dialog box then you will see your saves State Store listed in the appropriate section. To retrieve your saved state, simply click on the "Load" button associated with your State Store. You can also remove the State Store with the "Delete" button. The "Rename and Update" button will overwrite all of the currently saved State Store settings with those currently being used in the Perspective. It will also update the name of the State Store.

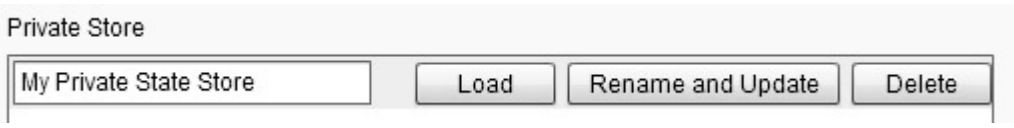

Figure 3: Loading, Updating and Deleting States

To close the State Store dialog box, click the "X" button in the top right hand corner of the dialog.
Back to the top

On This Page:

Opening the State Store dialog box
To access and open the State Store dialog box you the Application Menus - Options > State Store:

Next Steps

Application Menus
Learn how to use the Application Menus to close all tabs

Date Navigation
Learn how to use the Time Line Component for advance Data Navigation.

Filtering Data
Learn how to filter date using the provided Application Filters

Context Menus

Context Menus

Context Menus are used throughout the Client to help navigation and link data within applications. Context Menus are attached to and can be used with most Data Visualizations:

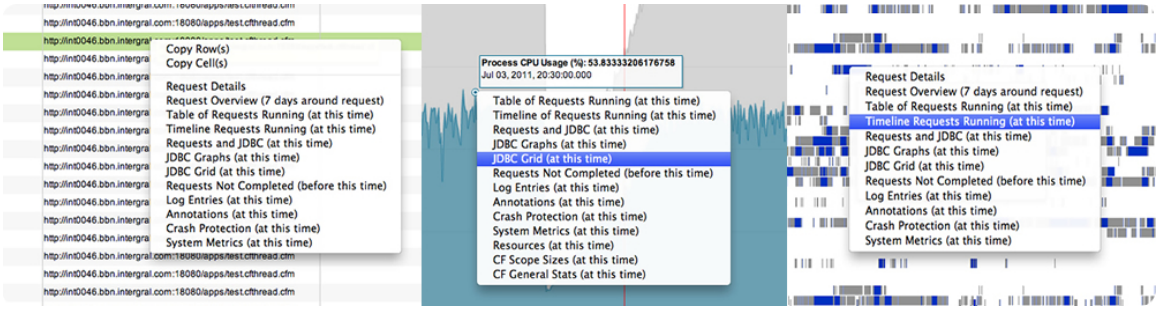

Figure 1: Application Context Menus

See the Data Visualization pages for more information about using Context Menus within each Visualization :

- DataGrid Visualization
- Chart Base Visualization
- Time Block Chart Visualization
- Sixway Chart Visualization

Context menus are created and displayed when the user "right" clicks on a data point within a visualization, e.g. on a DataGrid row or a Chart Series Item. The Context Menus contain menu items that are relevant to the current Perspective and the clicked data point. The menu items are (generally) links to other perspectives within the application relating to the selected data point and time frame:

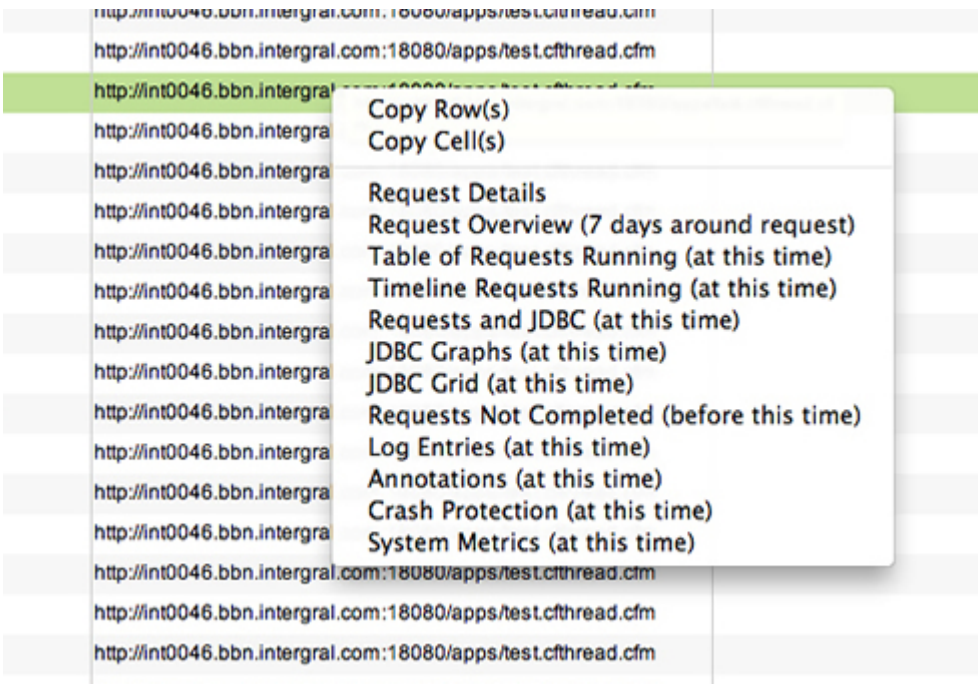

Figure 2: Context Menu rendered in a DataGrid Visualization

You will notice that the Context Menus within the DataGrid Visualization contain two additional menu items at the top of the menu: "Copy Row(s)" and "Copy Cell(s)" allowing you to cony content form the DataGrid.

For more information about copying DataGrid content see the Copying Grid Content section of the DataGrid Visualization page.

The remaining menu items are links to other Perspectives within the application which when clicked open the Perspective (selected from the menu item) in a new Tab within the Perspective Canvas. The selected Perspective will be loaded at a time frame relevant to the menu item selected from the menu, e.g. :

(at this time)
(before this time)
(7 days around this time)

Back to the top

Next Steps

Copying Grid Content
Learn about using the Context Menus to copy Data within the DataGrids

---

DataGrid Visualization - *Description of the DataGrid Visualization*
Chart Base Visualization - *Description of the Chart Base Visualization*
Time Block Chart Visualization -*Description of the Time Block Chart Visualization*
Sixway Chart Visualization -*Description of the Sixway Chart Visualization*

Business Metrics

Business Metrics

Overview

This section contains information about FusionAnalytics Reporting (including TAP and Daily Status reports). The section will discuss how the client can be used to navigate to and open the available reports from within the "Reports" Perspective.

Next Steps

Reports

Reports

Reports

Opening Reports

To open Application Report, using the TOC navigate to "Business Metrics > Reports". This will open the "Reports" Perspective within a new Application Tab Navigator. The "Reports" Perspective contains a DataGrid Visualization displaying all Reports relevant to the current application.

You will see two types of reports displayed within the Grid, TAP Reports and Daily Status Reports. To open a Report either right click on a report (to pop a Context Menu) and select the "Open Report" menu item from the Context Menu or simply double click on the Report to open the report in a new Perspective:

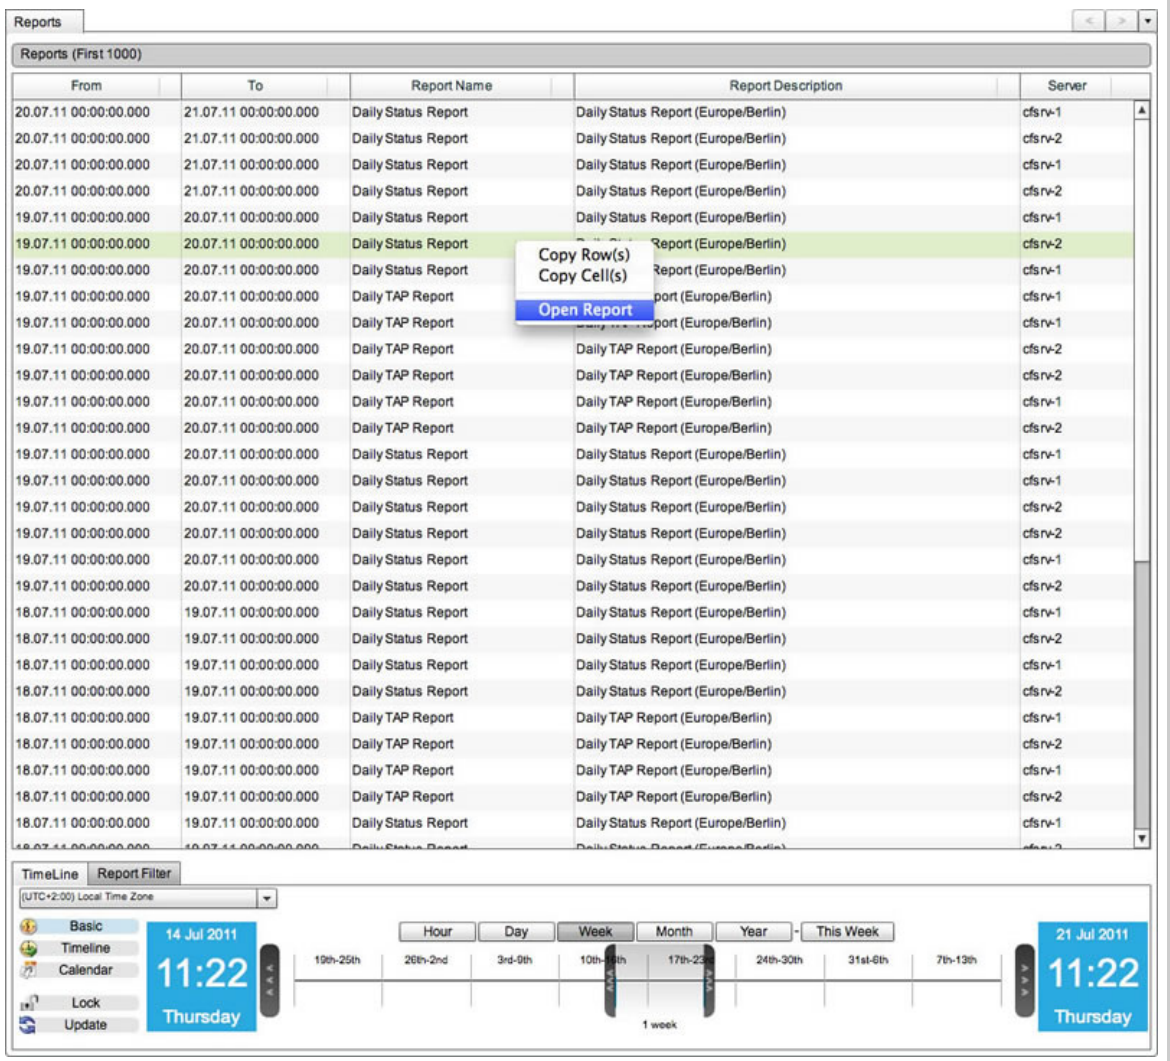

Figure 1: Reports Perspective - Opening Reports

Back to the top

TAP Reports

The TAP Reports give you a less technical view of how your server is performing. Every day you will get a TAP Score and Grade for how well your server is doing. It will also compare this to your previous scores. This score gives you a standardized comparison of applications regardless of function or hardware capability.

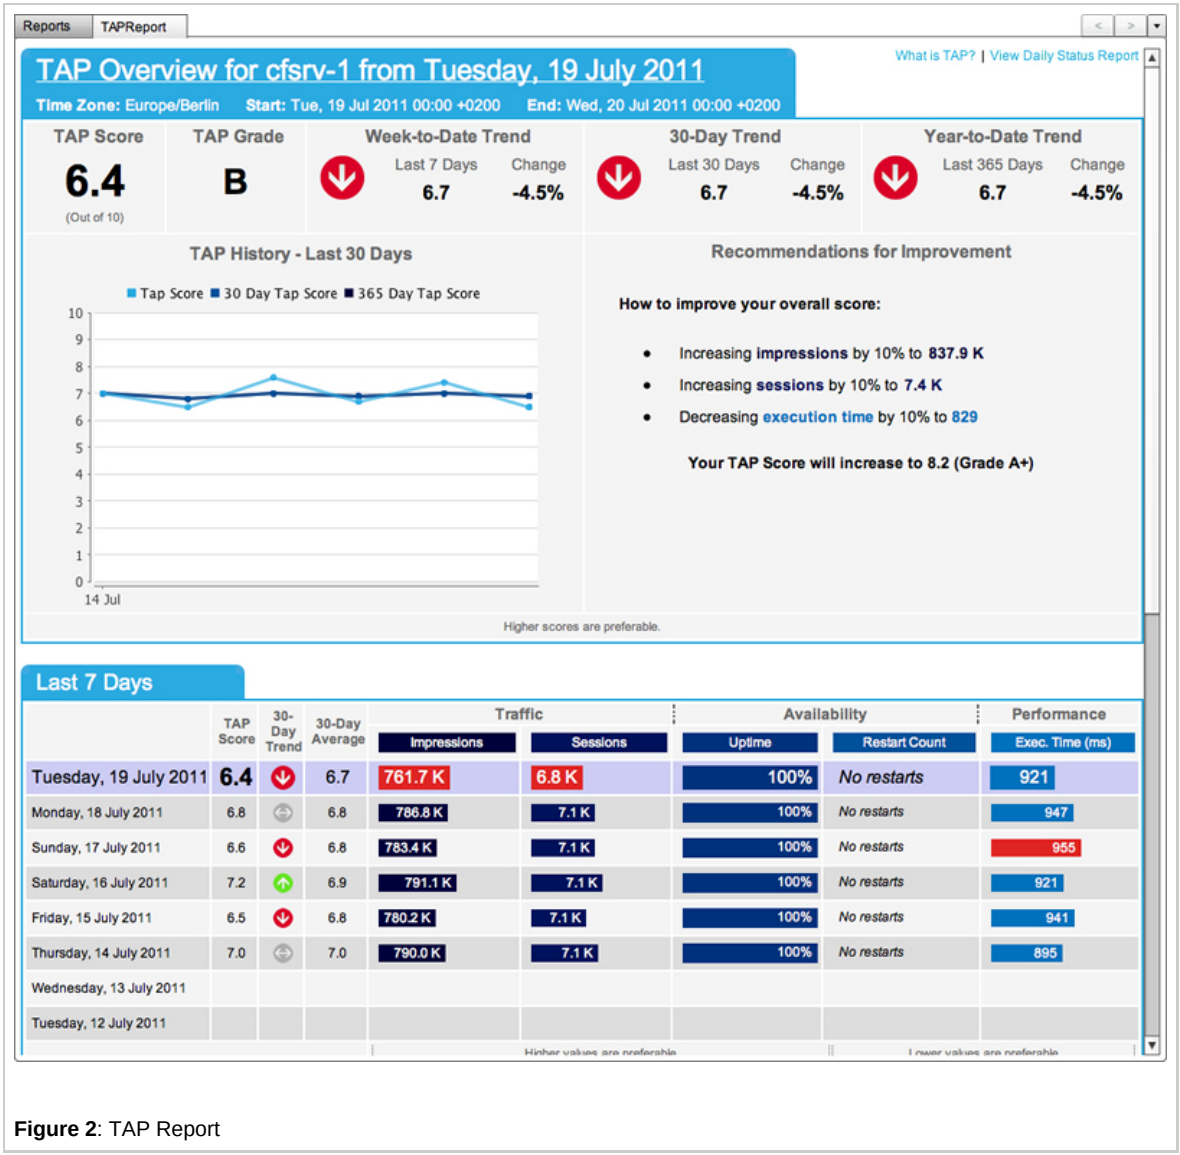

Figure 2: TAP Report

Learn more about TAP Reports

Back to the top

Daily Status Reports

The Daily Status Reports provide detailed information about the requests, database activity, resource usage and even a breakdown of the individual applications running. Each day you can see how your server has done compared to the previous day. This allows you to easily analyze which areas need improvement, helping you to improve the performance of your server.

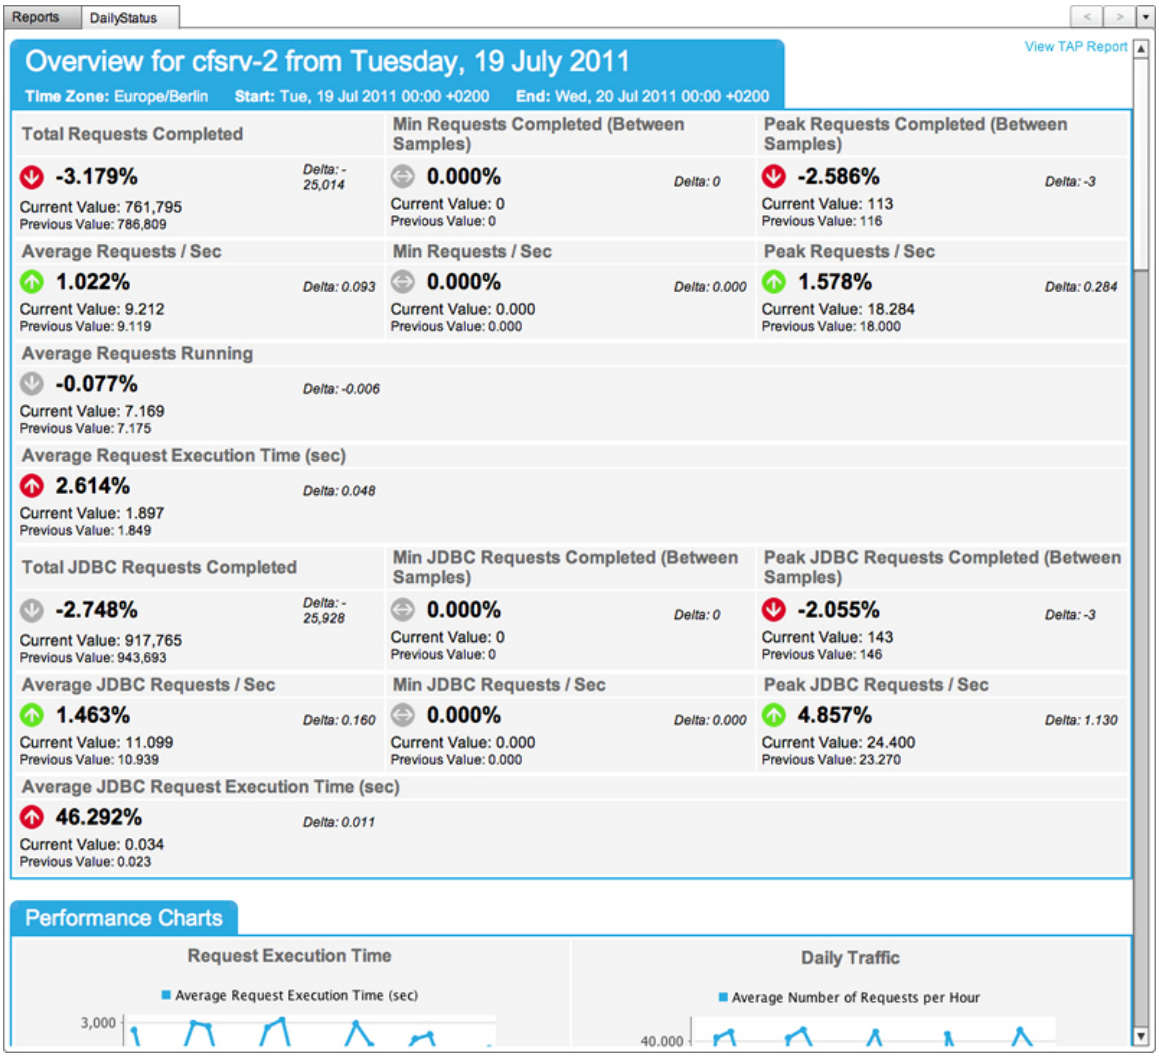

Figure 3: Daily Status Report

Learn more about Daily Status Reports

Back to the top

On This Page:

Next Steps

TAP Reports
Learn all about the TAP Reports

Daily Status Reports
Learn all about the Daily Status Reports

Using the Table of Contents (TOC)
Learn about the client layout and Using the Table of Contents (TOC)

Using Tabs
Learn how to use the Application Tab Navigator

DataGrid Visualization
Learn how to use and interact with the DataGrid Visualization

Context Menus
An explanation of what the Context Menus are and how they works / are used

Administrator

Administrator

Overview

This section will discuss how the Administrator can be used from within the FusionAnalytics application. The section will describe how you can Manage Applications from within the application and additionally how you can open and use both FusionAnalytics DataServices (FADS) and FusionAnalytics DataCollector (FADC) from within the client.

Next Steps

Manage Applications

FusionAnalytics DataServices

FusionAnalytics DataCollector

Manage Applications

Manage Applications

> ⚠ See the Setting Up Application Filters page for more information

Adding an application filter

To add an application filter, within the "Add New Application" section simply enter the application name (that you want to add the filter to) and click the "Add Application Filter" button:

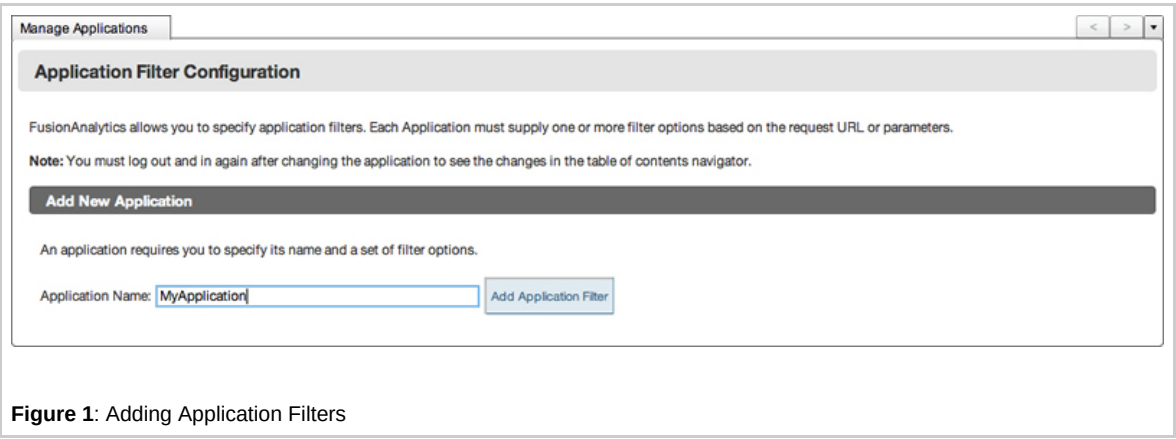

Figure 1: Adding Application Filters

The Application filter will then be added and displayed within the "Current Application Filters" section on the "Application Filter Configuration" page. To expand the Application Filter(s) (if not already expanded) click the "Edit Filters For (Application Name)"  .

Once the Application has been expanded you will see two tabs: "Request URL" and "Request Params" which you can add both filters accordingly. These filters allow you to limit the view to only include requests with specific limitations.

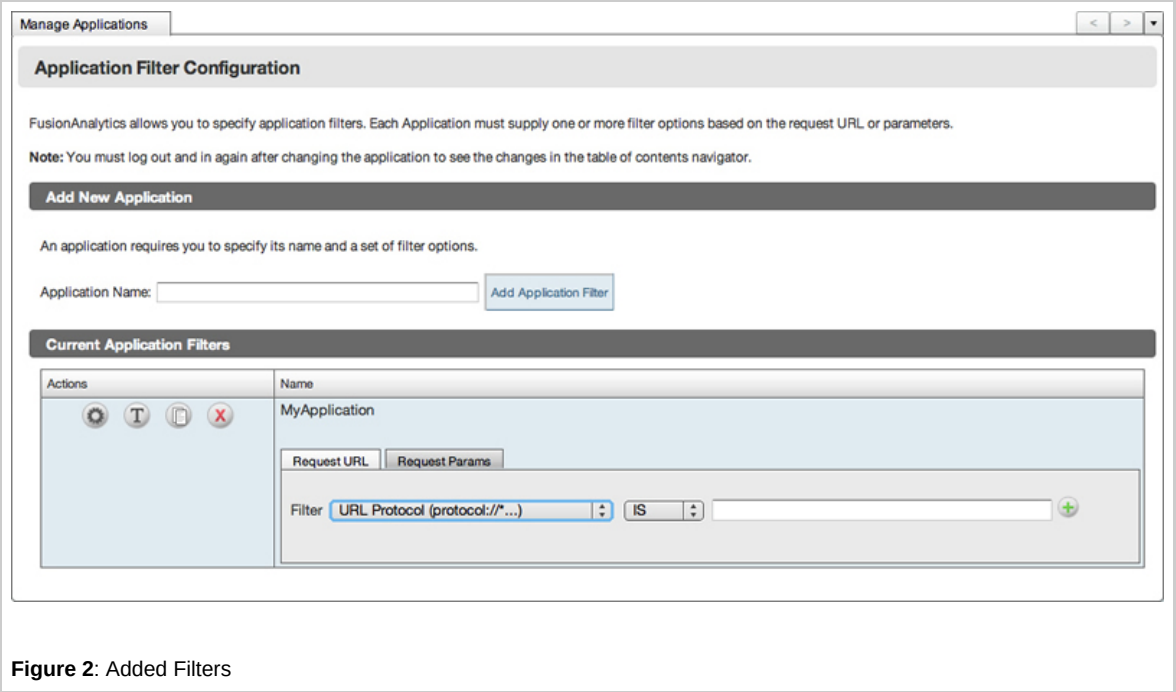

Figure 2: Added Filters

You will notice that the Application filter added will be added to the TOC under the "Applications" node. Note, you will have to click the "refresh" button to update the TOC after an Application filter has been added. For more information about the TOC see ....

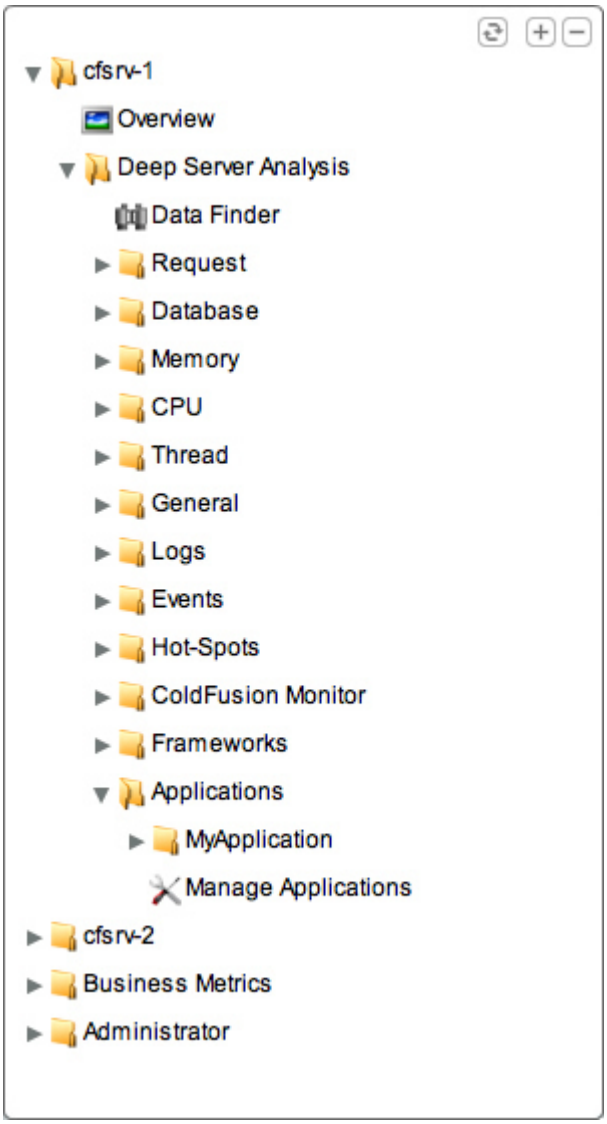

Figure 3: Added Filters listed in TOC

Back to the top

Request URL

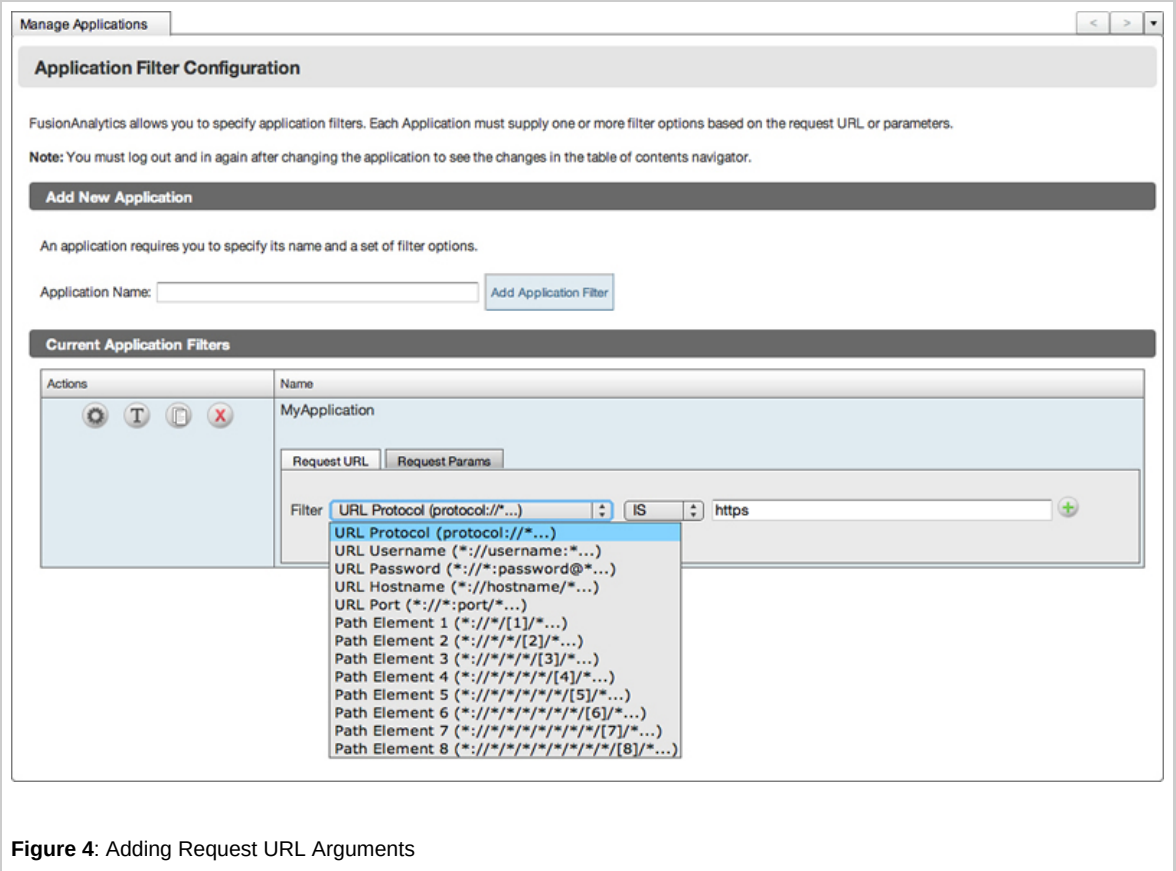

Figure 4: Adding Request URL Arguments

The Request URL filter lets you isolate some section of the request URL and filter just that part of the URL. These are the options:

URL Protocol (protocol://*...)
URL Username (**://username:**...)
URL Password (**://**:password@*...)
URL Hostname (*://hostname/...)
URL Port (**://**:port/*...)
Path Element 1 (**://**[1]/*...)
Path Element 2 (**:////[2]/**...)
Path Element 3 (**:////**[3]/*...)
Path Element 4 (**://////[4]/**...)
Path Element 5 (**://////**[5]/*...)
Path Element 6 (**://////[6]/**...)
Path Element 7 (**://////**[7]/*...)
Path Element 8 (**://////[8]/**...)

Once you have entered your filter type, comparison type and value, you can click the "Add" button to add the new filter to the list of filters below and update the perspective to include your defined filters.

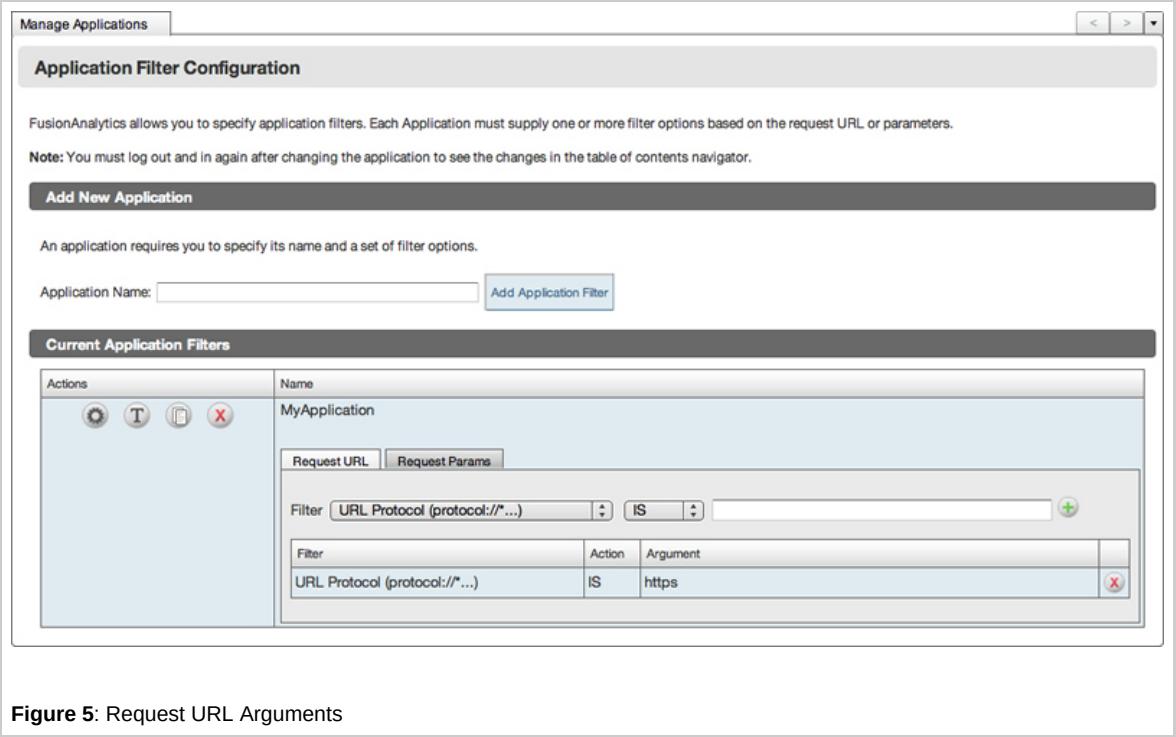

Figure 5: Request URL Arguments

Back to the top

Request Params

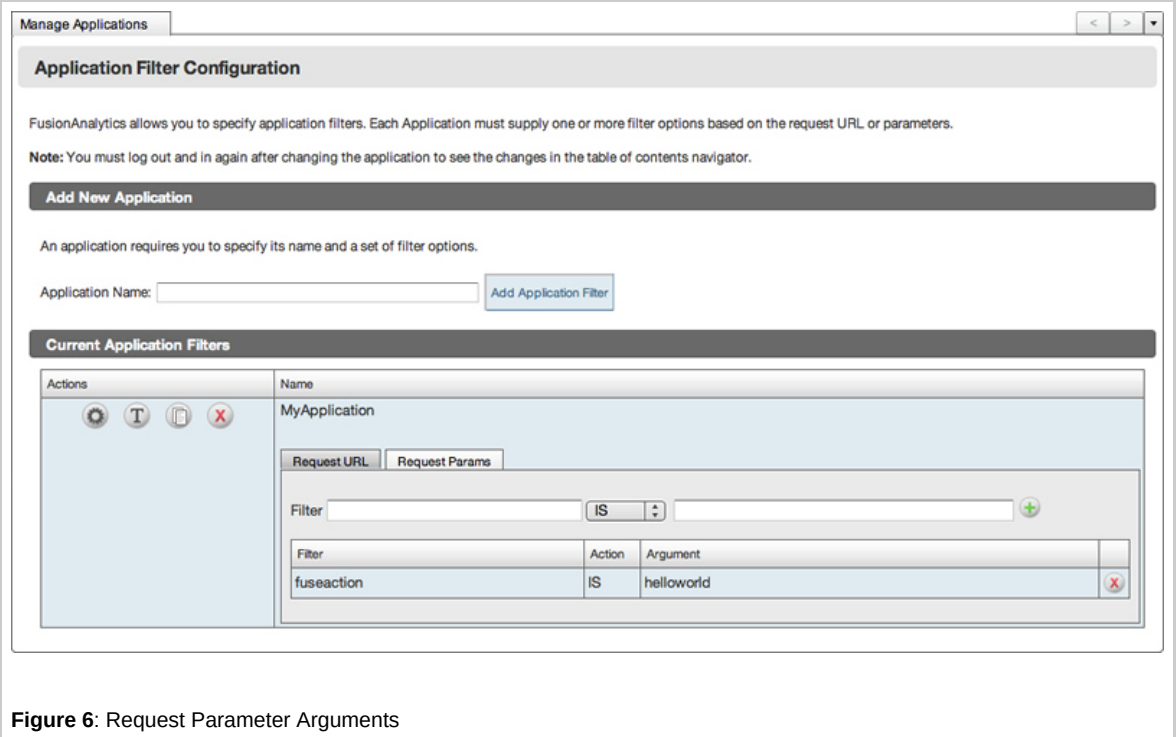

Figure 6: Request Parameter Arguments

The Request Params filter allows you to select values for a given request parameter. You can enter a Request parameter to filter (e.g. fuseaction, event etc...). For each filter you can either choose "IS" or "ISNOT".

Once you have entered your filter type, comparison type and value, you can click the "Add" button to add the new filter to the list of filters below and update the perspective to include your defined filters.

As with the "Request URL Filters", to remove a Filter simply click the "Delete" button within the "Name" column of the Application Filters table in the "Request Params" tab to remove the filter from the filter list and update the perspective removing the deleted filter.

Back to the top

Deleting a filter

To remove / delete a Filter simply click the "Delete" button within the "Name" column of the Application Filters table in the "Request URL" tab.

Note there is a "Delete"  button assigned to each filter. This will remove the filter from the filter list and update the perspective removing the deleted filter.

Back to the top

Renaming a filter

To rename a filter click the "rename" button next to the filter that you want to rename. This will change the state of the filter name at the top of the relevant row within the filters grid into an editable text box allowing you to rename (change / edit) the filter name. **Note**: if you want to cancel the rename process simply click the cancel button (next to the text input box) to return to the filters list. Once you have renamed the filter and are happy with your chosen name click the "modify" button (next to the text input box). A new dialogue will be showed confirming that you want to rename the filter which you can either select "cancel" (return to the rename state and choose another name) or "ok" (to change and save the filler name).

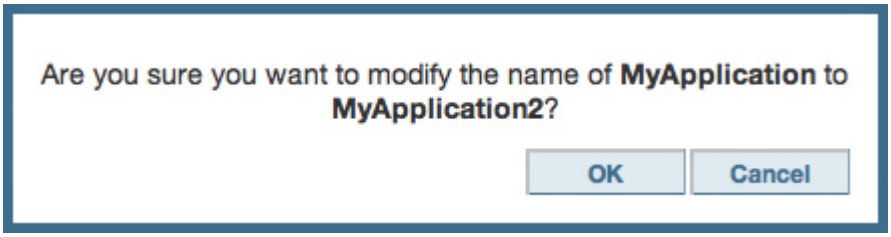

Figure 7: Renaming Filters

Back to the top

Copying a filter

To copy a application filter click the "copy" button for the application that you want to copy and the following popup will appear asking you you "Enter name for copied application" (**Note**: if you want to cancel the rename process simply click the cancel button (next to the text input box) to return to the filters list):

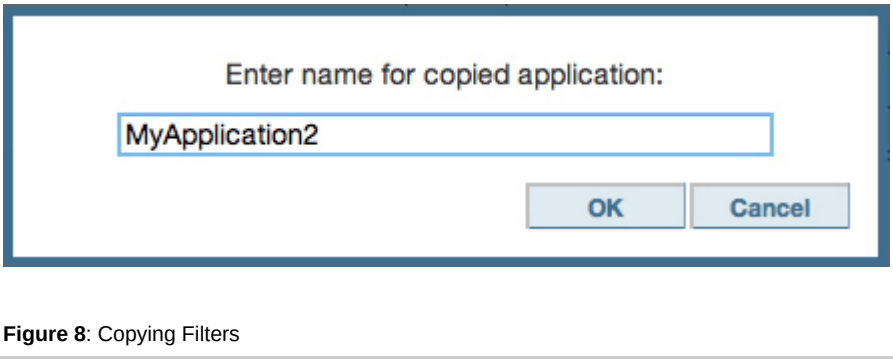

Figure 8: Copying Filters

Enter the name for the copied application and click the "OK" button from within the dialogue. The application will then be copied and added to the current application filter list and will also be added to the TOC under the "Applications" node:

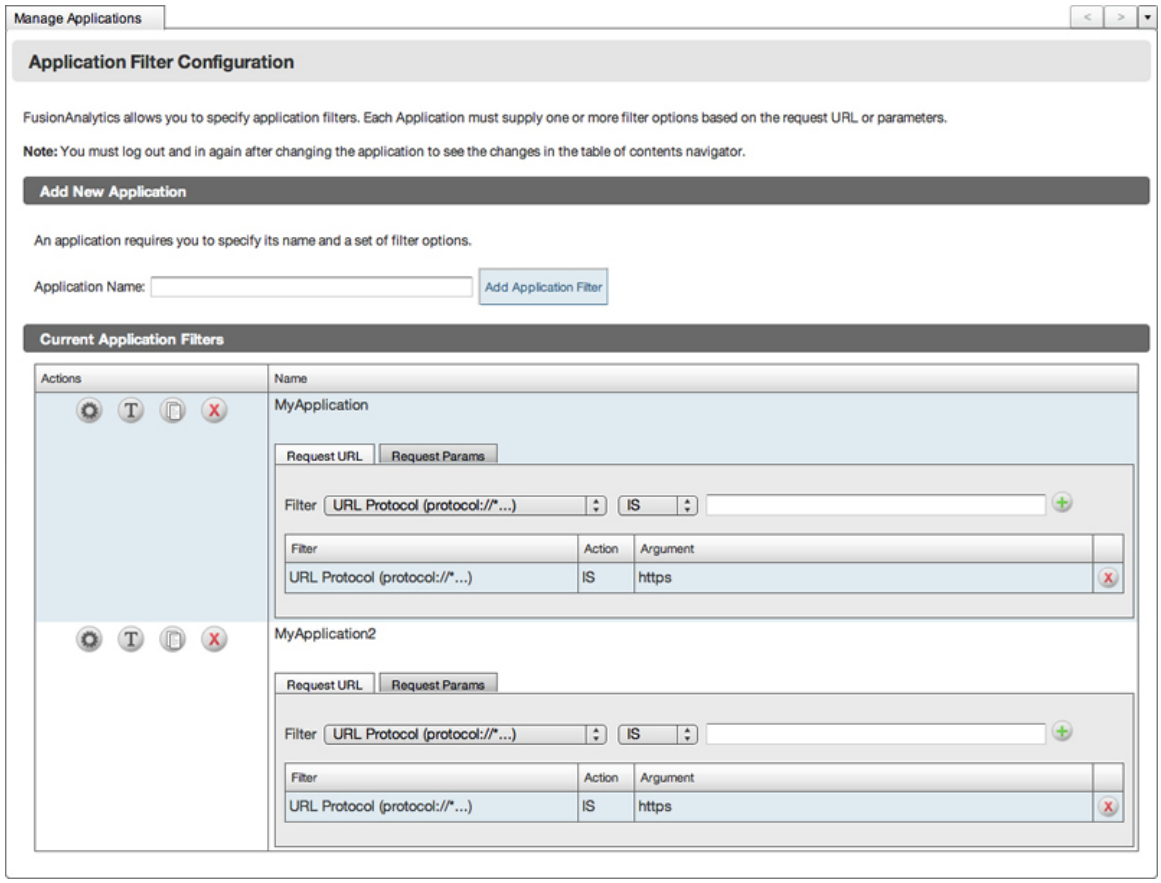

Figure 9: Copied Application Filter

See the Setting Up Application Filters page for more information

Back to the top

On This Page:

Next Steps

Understanding the Daily Status Report
Learn to use and understand the data in the Daily Status Reports.

FusionAnalytics DataServices

FusionAnalytics DataServices using the client

This page allows you to access all of the functionality of FADS exactly like you would in a browser. From the client menu simply navigate to Administrator > FusionAnalytics DataServices.

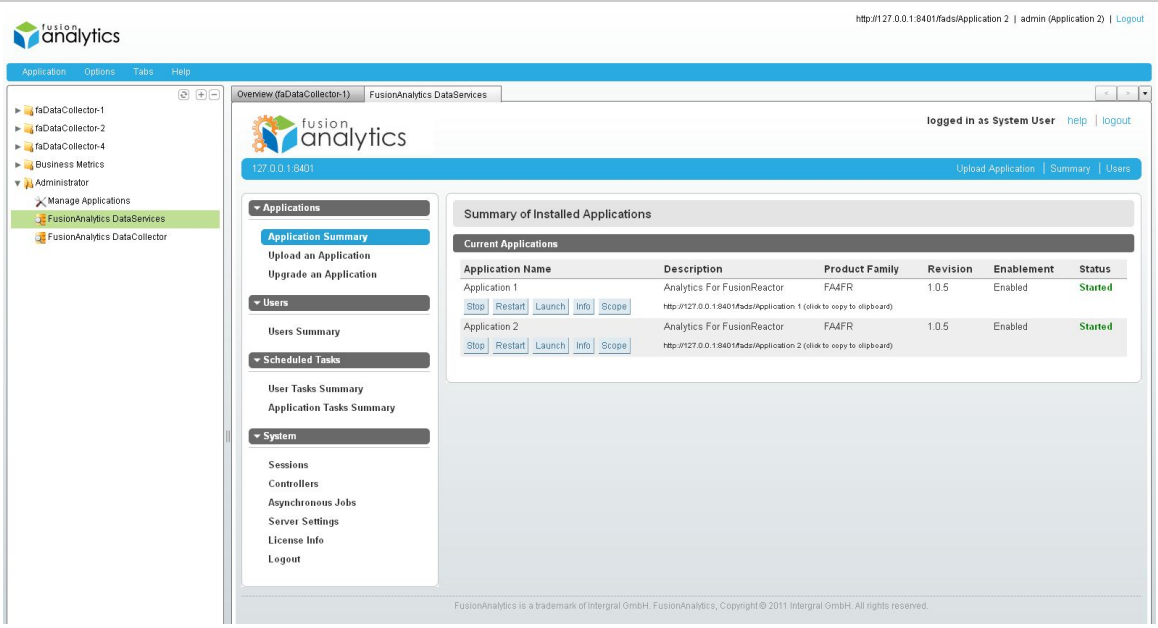

Figure 1: FADS from the client

For information on how to use FADS, see FusionAnalytics DataServices (FADS).

Back to the top

Next Steps

FusionAnalytics Client

FusionAnalytics DataCollector

FusionAnalytics DataCollector using the client

This page allows you to access all of the functionality of FADC exactly like you would in a browser. From the client menu simply navigate to Administrator > FusionAnalytics DataCollector.

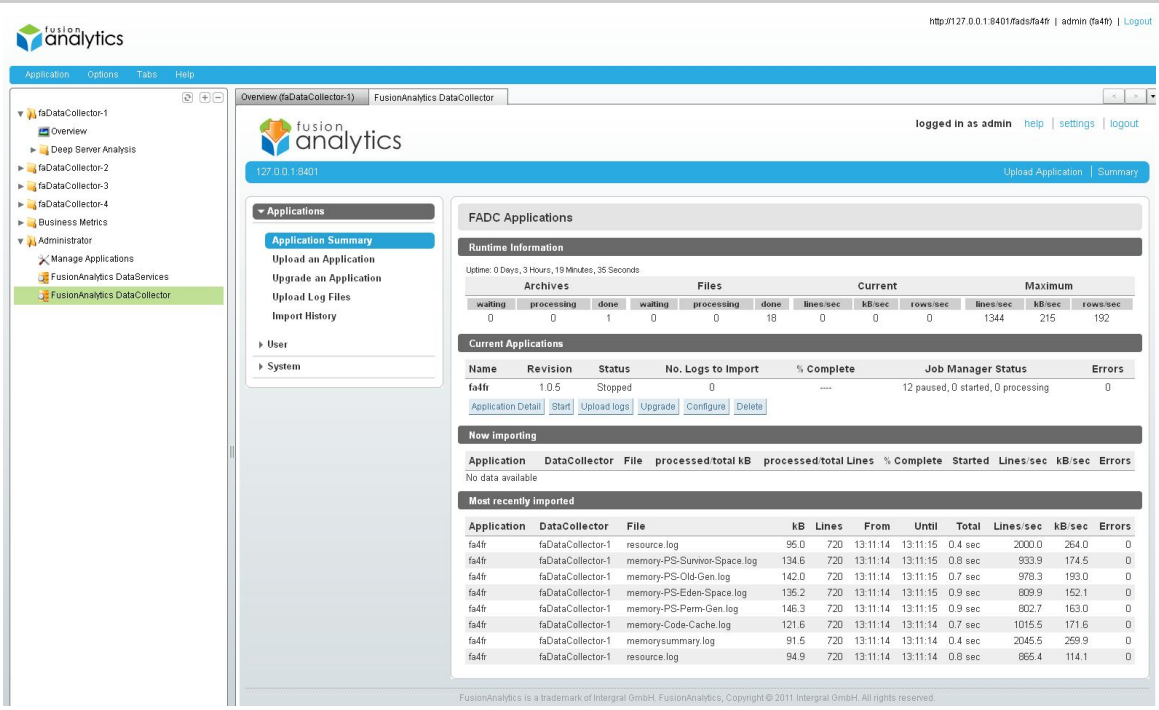

Figure 1: FADC from the client

For information on how to use FADC, see FusionAnalytics DataCollector (FADC)

Back to the top

Next Steps

FusionAnalytics Client

FusionAnalytics Server

FusionAnalytics Server

Overview

In the following we will have a look at the directory structure of FusionAnalytics server and how it is configured. The two web applications FusionAnalytics DataCollector and FusionAnalytics DataServices are described in detail and the license model of FusionAnalytics is explained.

Next Steps

Server Structure

- Directory Structure
- FADC Applications
- FADS Applications

Server Configuration

- FusionAnalytics Log Files
- Windows Service Configuration
- FusionAnalytics DataCollector Settings
- Setting Up FusionAnalytics with HTTPS

FusionAnalytics DataCollector (FADC)

- Applications (FADC)
  - Application Summary (FADC)
    - Application Configuration (FADC)
    - Application Detail (FADC)
  - Upload an Application (FADC)
  - Upgrade an Application (FADC)
  - Upload Log or ZIP Files (FADC)
  - Import History (FADC)
- User (FADC)
  - Change Password (FADC)
- System (FADC)
  - Email Settings (FADC)
  - Runtime Info (FADC)
  - License Info (FADC)
  - Logout (FADC)
- DCML (FADC)

FusionAnalytics DataServices (FADS)

- Applications (FADS)
  - Application Summary (FADS)
    - Application Configuration (FADS)
    - Application Detail (FADS)
    - Application Scope (FADS)
  - Upload an Application (FADS)
  - Upgrade an Application (FADS)
- Users (FADS)
  - Users Summary (FADS)
    - User Mappings
- Scheduled Tasks (FADS)
  - User Tasks Summary
  - Application Tasks Summary
- System (FADS)
  - Sessions
  - Controllers
  - Asynchronous Jobs
  - Server Settings

- License Info (FADS)
- Logout (FADS)

FusionAnalytics Licensing

Server Structure

Server Structure

Overview

This section provides you with an overview about the directory structure of the FusionAnalytics Server. You will learn what the main components of the server are, where the DCML and APML application packages can be found and where the installed applications reside in the directory structure of the server.

Next Steps

Directory Structure

FADC Applications

FADS Applications

Directory Structure

Directory Structure

The FusionAnalytics Setup installs a a Java 6 VM together with a preconfigured Apache Tomcat Server with the FusionAnalytics DataCollector and FusionAnalytics DataServices web applications deployed as exploded web archives (WAR) under the standard `webapps` directory. If you have accepted the default installation path FusionAnalytics is installed into the directory `C:/FusionAnalytics` containing the files as shown in Figure 1.

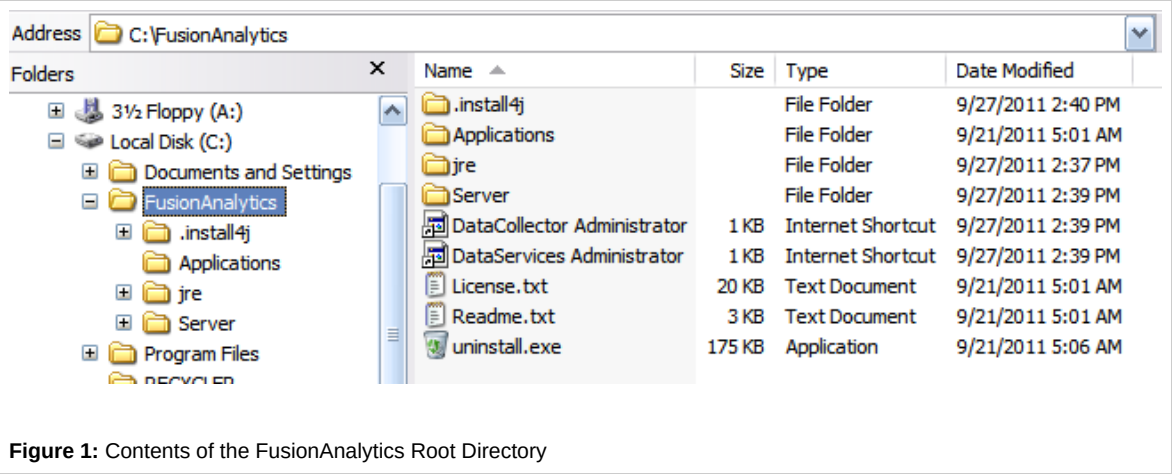

Figure 1: Contents of the FusionAnalytics Root Directory

The directories are used by FusionAnalytics as described in the following table:

Directory	Description
.install4j	A directory created by the install4j setup containing files related to the Setup and also required when uninstalling FusionAnalytics. You should not modify files in this directory.
Applications	Contains the DCML and APML FusionAnalytics for FusionReactor Applications (FA4FR) . If you want to add more FA4FR applications you can choose from these files when uploading the application from FADC resp. FADS
jre	Contains the Java VM used by the FusionAnalytics Server
Server	The root directory of the Tomcat Server coming with FusionAnalytics Setup. In addition to the regular structure of the Tomcat Server FusionAnalytics contains a directory named `data` where FADC specific files and license data is stored.

The Tomcat Server is configured to run as a Windows Service named **FusionAnalytics Server**. The root directory of Tomcat in this environment is `C:/FusionAnalytics/Server`. The two web applications are located in

Web Application	Location
FusionAnalytics DataCollector	`C:/FusionAnalytics/Server/webapps/fadc`
FusionAnalytics DataServices	`C:/FusionAnalytics/Server/webapps/fads`

Back to the top

FADC Applications

FADC Applications

The `data/fadc` directory contains an `Applications` and `Configurations` directory where FADC applications and their configuration are installed to. Figure 1 shows the structure of a FADC application named 'myApplication' after it has been successfully installed.

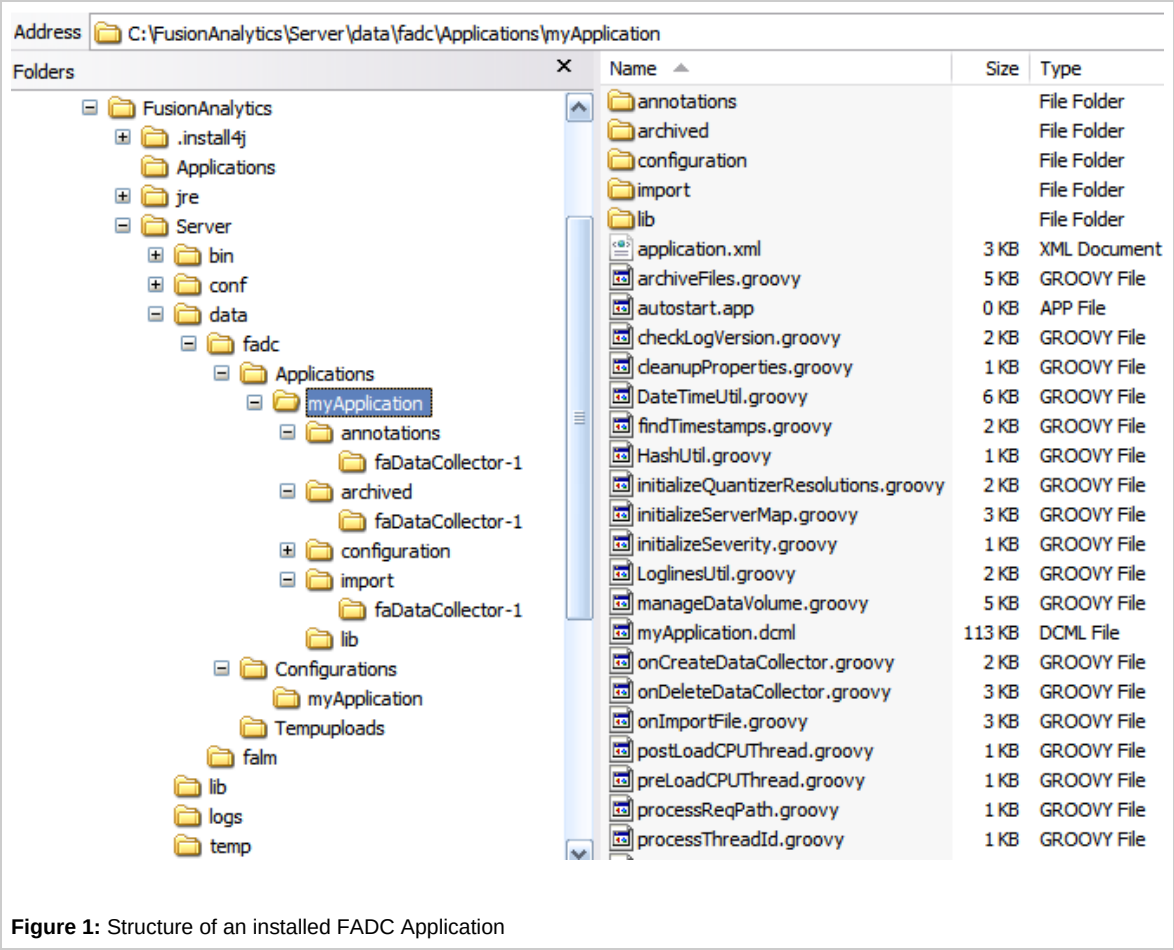

Figure 1: Structure of an installed FADC Application

You can see that the application contains additional directories named `annotations`, `archived`, `configuration` and `import` with a subdirectory for each configured DataCollector (e.g. 'faDataCollector-1').

Directory	Description
annotations	Currently not used
archived	Contains a dedicated directory for each configured DataCollector where a copy of all processed files is kept. Files are kept in 'daily' directories which are managed automatically by FusionAnalytics DataCollector.
configuration	The content of this directory is deployed to the fadc webapp when the application is installed.
import	Contains a dedicated directory for each configured DataCollector where files that should be processed are copied to. You can add files by uploading them via the web interface or manually copying them here.

You can upload single files or a ZIP archive of files to FusionAnalytics DataCollector. Paths inside ZIP archives are fully supported. When you manually copy files into the belonging import directory you can copy single files, ZIP archives or even complete directories.

If you copy large files, archives or directories into the import directory it can take a long time for FusionAnalytics DataCollector to process them. The recommended approach is to start with smaller files until you get a feeling of how much resources are required for a certain amount of input files.

Back to the top

FADS Applications

FADS Applications

FADS applications are deployed to the directory `C:/FusionAnalytics/Server/webapps/fads/WEB-INF/classes/Applications`. Figure 1 shows how the structure of a successfully installed FADS application looks like:

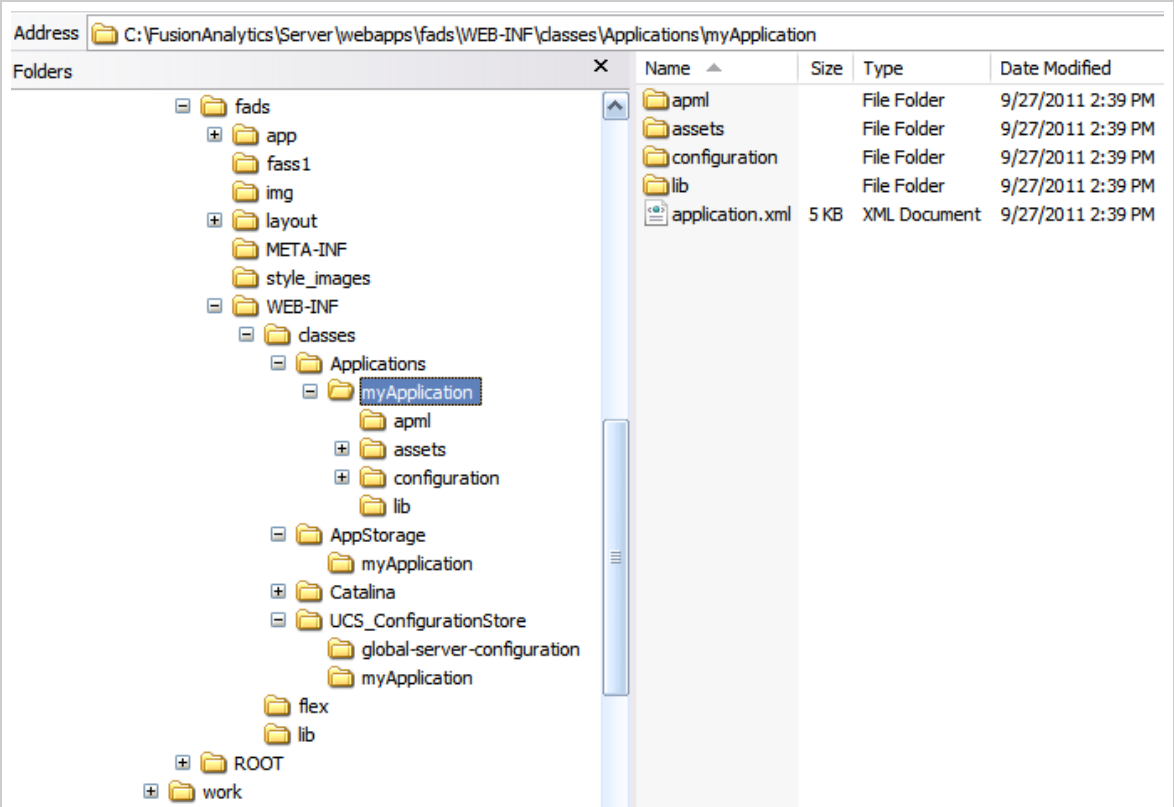

Figure 1: Structure of an installed FADS Application

Back to the top

Server Configuration

Server Configuration

Overview

This section describes the FusionAnalytics Server log files and the default configuration and additional system properties that can be used to change the configuration. As FusionAnalytics runs on a Tomcat Server any configuration step related to Tomcat applies as well.

Next Steps

FusionAnalytics Log Files

Windows Service Configuration

FusionAnalytics DataCollector Settings

Setting Up FusionAnalytics with HTTPS

FusionAnalytics Log Files

FusionAnalytics Log Files

FusionAnalytics does output comprehensive logging information depending on the log level configured. The logging output can be configured by changing the settings in the two belonging log4j configuration files.

For FADC the log4j configuration is defined in file

```
C:/FusionAnalytics/Server/webapps/fadc/WEB-INF/classes/fadc-log4.groovy
```

for FADS it is defined in file

```
C:/FusionAnalytics/Server/webapps/fads/WEB-INF/classes/log4j.xml
```

The log files are located in the directory
`C:/FusionAnalytics/Server/logs`

Log file	Description
FusionAnalytics Server.log	Contains messages written to STDOUT and STDERR
fusionanalytics-datacollector.log	Contains all messages created by FADC
fusionanalytics-datacollector-errors.log	Contains only the error messages created by FADC
fusionanalytics-datacollector-ignored-files.log	Lists the files that have be ignored by FADC
fusionanalytics-datacollector-imported-files.log	Lists the files that have been processed by FADC
fusionanalytics-datacollector-rejected-files.log	Lists the files that have been rejected by FADC because they have been imported before
fusionanalytics-dataservices.log	Contains all messages created by FADS

Keep an eye on file `fusionanalytics-datacollector-errors.log` to see if there were any problems when running FADC

Back to the top

Per default the log level is set to 'INFO' which does produce only little logging data but on the other hand does not slow down the application.

Windows Service Configuration

Windows Service Configuration

FusionAnalytics Setup creates a Windows service named **FusionAnalytics Server**. The service uses a configuration file to define settings for the Java VM and additional settings read by FusionAnalytics DataCollector. Shown below is the default configuration file for the FusionAnalytics service:

```
-Xms128m
-Xmx512m
-XX:MaxPermSize=256m
-Dcom.sun.management.jmxremote.port=3333
-Dcom.sun.management.jmxremote.ssl=false
-Dcom.sun.management.jmxremote.authenticate=false
-Dfusionanalytics.datacollector.importThreads=4
-Dfusionanalytics.datacollector.analyzerThreads=2
-Dfusionanalytics.datacollector.dirObserver.scanInterval=5000
-Dfusionanalytics.lm.store=C:/FusionAnalytics/Server/data/falm/falm.store
-Dgrails.scan.interval=3600000
```

Back to the top

FusionAnalytics DataCollector Settings

FusionAnalytics DataCollector Settings

The FusionAnalytics DataCollector can be configured using the following system properties.

Name	Description	Default Value
fusionanalytics.datacollector.analyzerThreads	The number of file analyzer threads	2
fusionanalytics.datacollector.applicationmanager.threads	The number of threads used to unzip archives	2
fusionanalytics.datacollector.batchSize	The number of statements in a SQL batch before committing it	500
fusionanalytics.datacollector.connectionObserverInterval	The number of seconds the connection observer thread sleeps until trying to restart apps that stopped because they have lost their database connection	300
fusionanalytics.datacollector.connectionpool.size		50
fusionanalytics.datacollector.connectionpool.timeout	The number of seconds before a timeout occurs if a connection cannot be retrieved	30
fusionanalytics.datacollector.dirNameFormat	The format used for the directory names created in the archived and failed directories	yyyy-MM-dd
fusionanalytics.datacollector.dirObserver.scanInterval	The number of seconds between two checks for modified/new/removed files in the application directories	5000
fusionanalytics.datacollector.disableCommitStatement	Disables database commits; useful for performance analysis only	false
fusionanalytics.datacollector.errorsBeforeFailure	not used	0
fusionanalytics.datacollector.fileMonitor.readFileAttempts	The number of attempts to process a drop import file if it is not ready	10
fusionanalytics.datacollector.fileMonitor.timeBetweenAttempts	The number of milliseconds until the next attempt to process a drop import file if it was not ready	6000
fusionanalytics.datacollector.importThreads	The number of import threads	#cores

fusionanalytics.datacollector.janitorInterval	The number of seconds the janitor thread sleeps until trying to remove obsolete files again	300
fusionanalytics.datacollector.jobShutdownTimeout	The number of seconds to wait for a job to finish before it will be killed	15
fusionanalytics.datacollector.lookbackLines		100
fusionanalytics.datacollector.maxLenErrorMessage	The maximum length of an error message to be written to the logs before it will be truncated	1024
fusionanalytics.datacollector.queueInterval	The number of ms to sleep inside the event processing loop before processing the next event	5
fusionanalytics.datacollector.resumeImportsFile	The format used for the directory names created in the archived and failed directories	${java.io.tmpdir}/.fadc-resumeImports.dat
fusionanalytics.datacollector.runtimeinfo.interval	The number of milliseconds between to probes	3000
fusionanalytics.datacollector.runtimeinfo.queueSize	The number of probes used to calculate average data	60
fusionanalytics.datacollector.unzipMaxAttempts	The number of attempts to unpack a zip archive	5
fusionanalytics.datacollector.unzipSecondsUntilNextAttempt	The number of seconds between to attempts to unpack a zip archive	30
fusionanalytics.datacollector.urlencoding		UTF-8
fusionanalytics.datacollector.web.recentImportSize	The number of most recently processed log files to show in the user interface	5

We do not recommend to change the default configuration of FusionAnalytics DataCollector if not really required.

Back to the top

Setting Up FusionAnalytics with HTTPS

Setting Up FusionAnalytics with HTTPS

FusionAnalytics supports secure (HTTPS) connections, however you need to setup / configure the server to use HTTP for this to work.

Note: You should also configure the HTTPS port in FusionAnalytics DataServices (FADS) on the Server Settings hostname/port to tell it that HTTPS is available

Tomcat

You must install a Certificate into Tomcat to enable HTTPS. HTTPS then runs on a different port to the HTTP port (the HTTP port is typically 8400 in FusionAnalytics)

- Modify the `server.xml` file located under <fusionanalytics>/Server/conf

```
<Connector port="8443" protocol="HTTP/1.1" SSLEnabled="true"
maxThreads="150" scheme="https" secure="true"
keystoreFile="C:/keys/keystore.jks" keystorePass="changeit"
clientAuth="false" sslProtocol="TLS" />
```

- This is using port 8443 and assuming the keystore password is "changeit". Make any modifications as necessary

- Restart the server (start it using catalina.bat to make sure no errors are being displayed).
- For more information on how to configure a tomcat server to use SSL see: http://tomcat.apache.org/tomcat-6.0-doc/ssl-howto.html

Back to the top

Creating a Self-Signed Certificate

- You need to have Java installed and access to the keytool (java/bin)
- keytool -genkey -keyalg RSA -alias selfsigned -keystore keystore.jks -storepass changeit -validity 360 -keysize 2048
- Enter "localhost" as the first and last name!
- Use the generated keystore in your tomcat installation.

Back to the top

Configuring FADS

You need to configure the FADS server to tell it that HTTPS is available.

- Open FADS Administrator and log-in as a System User
- Got the System -> Server Settings
- Enter the HTTPS server details into the Server Settings and click on the [Update Server Serttings] button

Back to the top

FusionAnalytics Jobs Configuration

Typically Jobs in FusionAnalytics will generate links to the HTTP configuration but you may want links and images to use HTTPS also. This is controlled by the job itself but so far we are controlling this in the jobs by adding an Argument to the job:

Name	Value
https	true

Back to the top

AIR Client

- Open the AIR client and create an application to localhost with HTTPS as the protocol.
    - You can also edit an existing application to use a secure connection
        - Right click on an existing application and select the "Secure Connection" check box

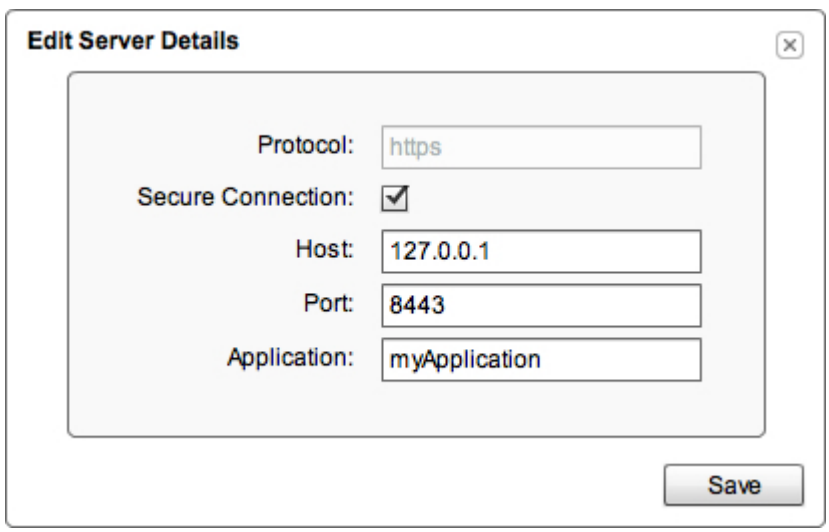

Figure 1: Edit Server Details Panel

- When logging you will be prompted that the certificate is a problem (not secure).
- Choose to view the certificate and install it.
- When the installation wizard pops, chose to install the certificate yourself (not letting the wizard place it in the "right" place).
- Then choose to see the physical locations and select the Trusted Root Certification Authorities and then pick the Registry node.
- The certificate should install correctly and then OK the security dialog.
- This should allow you to log in completely into the application without being prompted again.

- This does not work with the certificate from the FR tickets. Use a keystore generate as above.

Back to the top

Sending data to FusionAnalytics
DataCollectors are also available in HTTPS allowing data to be securely sent from FusionReactor into FusionAnalytics.

FusionAnalytics DataCollector (FADC)

FusionAnalytics DataCollector (FADC)

Overview

FusionAnalytics DataCollector (FADC) provides high speed import of arbitrary log files into a SQL database so that the resulting data can be processed by the FusionAnalytics Server and visualized by the FusionAnalytics Desktop application. FADC uses applications to model different usage scenarios. An application is typically made up of a DCML file and a set of auxiliary Groovy and/or BeanShell scripts, optionally accompanied by further classes or JAR files (e.g. JDBC drivers).

Next Steps

Applications (FADC)

- Application Summary (FADC)
  - Application Configuration (FADC)
  - Application Detail (FADC)
- Upload an Application (FADC)
- Upgrade an Application (FADC)
- Upload Log or ZIP Files (FADC)
- Import History (FADC)

User (FADC)

- Change Password (FADC)

System (FADC)

- Email Settings (FADC)
- Runtime Info (FADC)
- License Info (FADC)
- Logout (FADC)

DCML (FADC)

DCML
A domain specific XML language called DataCollector Markup Language which defines from which files data is imported and how this data is mapped to a relational database structure.

Applications (FADC)

Applications (FADC)

Overview

This section describes the menu items of the Application menu as shown in Figure 1.

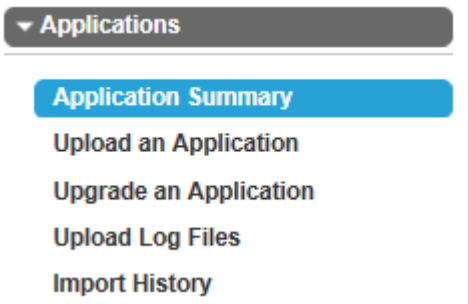

Figure 1: The Applications Menu

Next Steps

Application Summary (FADC)

- Application Configuration (FADC)
- Application Detail (FADC)

Upload an Application (FADC)

Upgrade an Application (FADC)

Upload Log or ZIP Files (FADC)

Import History (FADC)

Application Summary (FADC)

Application Summary (FADC)

Overview

The application summary gives you control over all of your DCML applications as well as information on log importing. The summary is split into the four sections

- Runtime Information
- Current Applications
- Now Importing
- Most recently imported

Figure 1 shows you a screen shot of the Application Summary page.

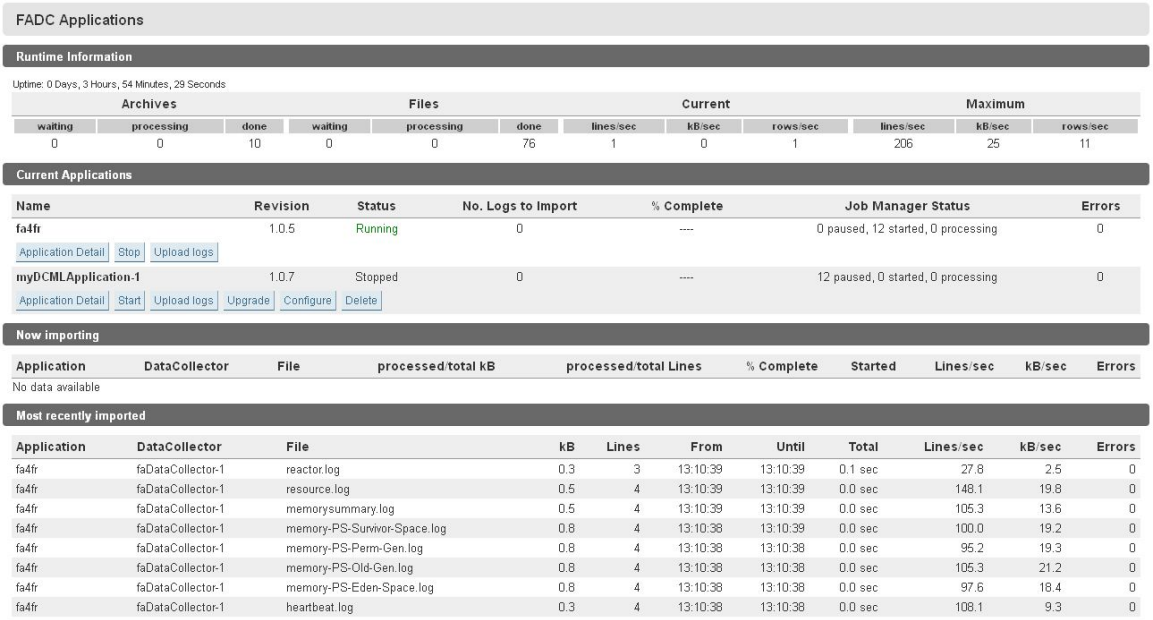

Figure 1: Application Summary

Runtime Information

This section details importing information and statistics such as uptime, how many files have been imported and how many archives are currently been imported.

Current Applications

Similar to FADS, this section allows you to control every aspect of your application. The available buttons are different depending on whether your application is started or stopped.

Options for a started application

Name	Description
Application Detail	Allows you to see information about this application's jobs, status, DataCollectors and database.
Stop	Stops the application.
Upload Logs	Allows you to manually upload logs. See Upload Log or ZIP Files (FADC).

Options for a stopped application

Name	Description
Application Detail	Allows you to see information about this application's jobs, status, DataCollectors and database.
Start	Starts the application
Upload Logs	Allows you to manually upload logs. See Upload Log or ZIP Files (FADC).
Upgrade	Allows you to upgrade your application. See Upgrade an Application (FADC).
Configure	Takes you to the Application Configuration (FADC) page.
Delete	Deletes your application from FADC.

Now Importing

This section details logs as they're being imported.

Most Recently Imported

This section shows what logs were most recently imported, what applications they belonged to and timestamp/transfer information.

Next Steps

Application Configuration (FADC) Application Detail (FADC)

Application Configuration (FADC)

Configuring applications in FADC

The Application Configuration page allows you to configure various properties of an FADC application. Here you can:

- Configure data sources
- Limit data size
- Manage log file retention period/size
- Add, modify and delete DataCollectors
- Record uploaded files to avoid duplicate uploads

Back to the top

FusionAnalytics DB Data Management

Here you can define the amount of time you would like FusionAnalytics to maintain your log/transactional data. FusionAnalytics stores detailed data and quantized data - this simply means that the data has been stored at different resolutions - 1 minute, 1 hour, 1 day, 30 days.

We recommend keeping the lower resolution (quantized) data for a longer period of time (so that you can analyze and compare your application's resources historically. The higher resolution (detailed) data should only be kept for short periods of time, because of the amount of space required to store this data in your database.

This will affect the size of your database.

The times you specify will refer to the dates inside your log files, not when the data was uploaded to FusionAnalytics.

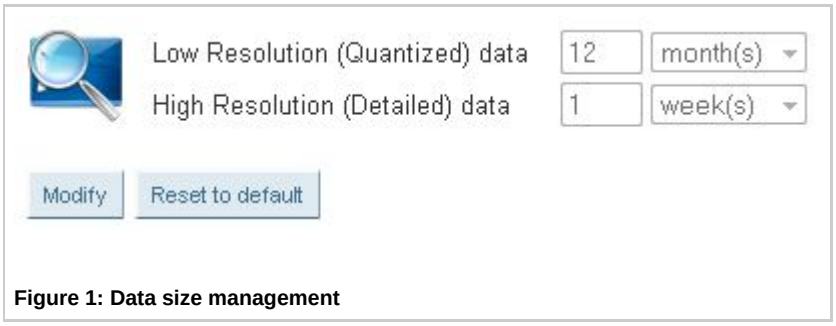

Figure 1: Data size management

Back to the top

FusionAnalytics Processed Log File Management

Here you can define the number of days and the amount of storage which FusionAnalytics should take into account when managing retention of processed log files.

Once the FusionAnalytics DataCollector (FADC) has processed your log files then they will be stored in archive folders - one archive folder (with the naming convention YYYY-MM-DD) is used to store the logs for that day. To ensure that you do not overflow your disk space, you can limit the number of days and the amount of space which can be allocated to storing these log files. Note that specified disk space (GB/MB) takes precedence over the number of days you specify i.e. if you reach your disk space limit, then the oldest log file data will be removed until you are back within acceptable space range. If you are within your space limitation, then we will simply retain the logs for the number of days you specify.

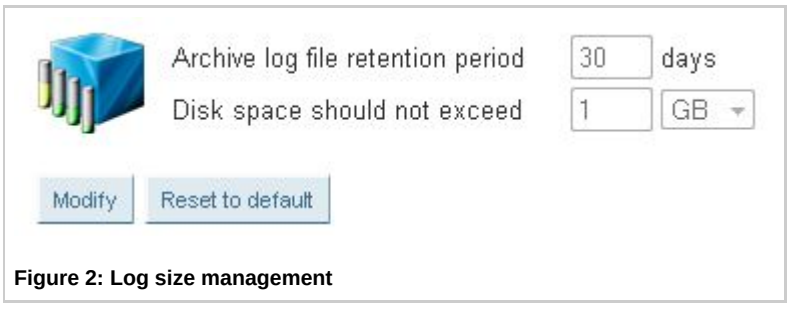

Figure 2: Log size management

Back to the top

FusionAnalytics DataCollectors

FusionAnalytics supports multiple DataCollectors to be used in a single application; there must always be one DataCollector set up for FusionAnalytics to use. DataCollector names can consist of letters, numbers, hyphens and underscores only.

Figure 3: DataCollectors

Back to the top

Miscellaneous Server Settings

If you are using FusionAnalytics to analyze logs generated from a Resin Application Server (e.g. with when using Railo), then you should tick the checkbox in this section.

Back to the top

Record Uploaded Files

FusionAnalytics can be configured to reject files that have been uploaded before. Therefore it records each file uploaded to this application together with the target DataCollector, the time when the upload occurred and a MD5 hashsum of the file in the database. You can limit the number of rows this table is allowed to have before entries - starting with the oldest data - will be deleted.

Back to the top

Data Sources

This is where you can configure data sources for your application. Please fill in the settings that match your database.

Click to enlarge

Field Name	Description	Example
Name	The name of your datasource. Every FA4FR application contains a default data source (frdb) for fusion reactor which cannot be renamed.	DataSource1
Description	Enter a brief description of this data source.	This data source is used to connect FusionAnalytics to FusionReactor
Database Type	The type of database you are using. Currently only Microsoft SQL Server is supported.	Microsoft SQL
Driver Class	The driver class tells FusionAnalytics the name of the JDBC database driver to use.	com.microsoft.sqlserver.jdbc.SQLServerDriver
URL	This specifies the JDBC connection URL used to connect to the database. Customize this field for your own database.	jdbc:sqlserver://127.0.0.1:1433;loginTimeout=5;databaseName=fadb
User	A username that has full access to the database	
Password	The password for the above account	

Pool Size	This is the number of database connections that can be open at any one time.	30
Pool Timeout (Seconds)	The number of seconds that have to pass before the pool times out.	10

Back to the top

Application Detail (FADC)

Application Detail (FADC)

The application detail page gives you control over the DataCollectors and Jobs of your application as well as information on log importing. The summary is split into the sections:

- Application Status
- Data Input
- Data Management
- Database Details

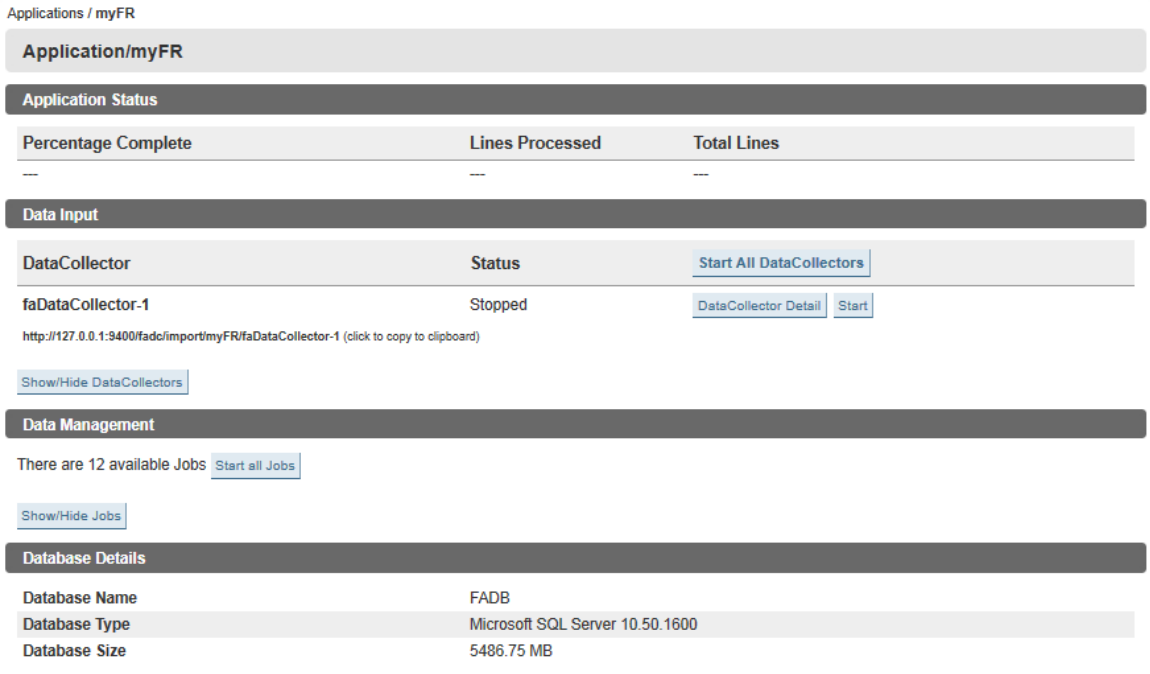

Figure 1: Application Detail

Back to the top

Application Status

This section provides you with information about the running imports of the application. Shown is how much percent of the lines of all files scheduled to be imported have been processed so far.

If new files are added to the application these values are updated as soon as their lines have been counted by FADC.

Data Input

Here you can control the DataCollectors of the application. You can start/stop all DataCollectors with a single click or you can start/stop each DataCollector separately after having clicked on the `Show/Hide DataCollectors` button. In addition to the name of the DataCollector its import URL as used by FusionReactor is shown. A click on the `DataCollector Detail` button opens the DataCollector detail page.

Back to the top

Data Management

This section lists the jobs defined by the application. Click the `Show/Hide Jobs` button to show/hide the jobs of the application. Click the `Start all Jobs` to start all jobs, click `Stop all Jobs` to stop them. Jobs can also be started/stopped separately. If you want to execute a job right now you can click the `Execute Now` button.

A click on the `Job Detail` button opens a new page where you can edit the cron pattern that defines when the jobs executes.

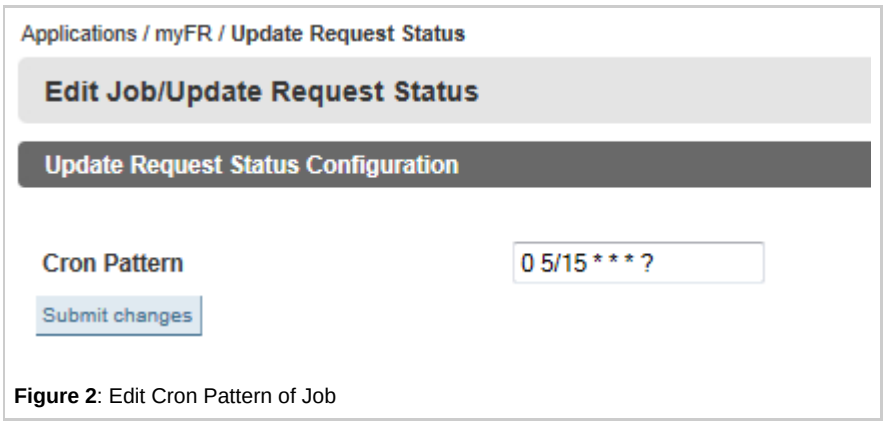

Figure 2: Edit Cron Pattern of Job

Back to the top

On this page:

Starting an application from the Application Summary page results in starting all of its DataCollectors and all of its jobs. From the Application Details page you can start/stop the DataCollectors and jobs of the application independently.

Next Steps

Application Configuration (FADC)

Upload an Application (FADC)

Upgrade an Application (FADC)

Upload an Application (FADC)

Upload an Application (FADC)

Step 1: Navigating to the upload page

In the FADC navigation menu, go to Applications > Upload an Application

Back to the top

Step 2 - Uploading an application

Choose a name for the application and enter it into the name input field. Click the browse button and navigate to the application package you wish to upload. Click the upload button to upload your application.

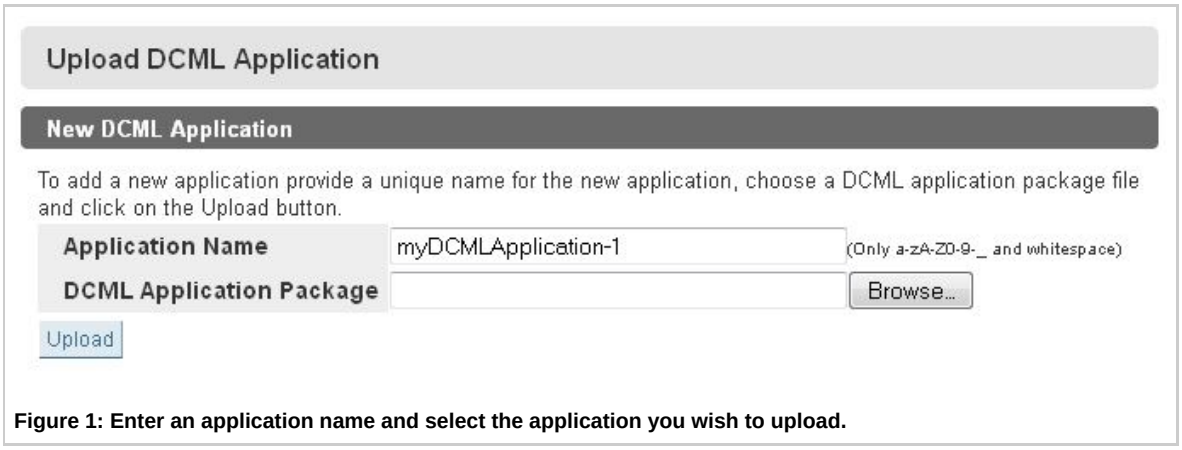

Figure 1: Enter an application name and select the application you wish to upload.

An application name can only use:
- a - z
- A - Z
- 0 - 9
- dash (-)
- underscore (_)
- whitespace

Back to the top

On this page:

Upgrade an Application (FADC)

Upgrading an Application

An application must meet the following prerequisites before upgrading:

- The new application's version number must be higher than the current application's version number.
- Application names can only contain a-z A-Z 0-9 dash(-) underscore(_) whitespace
- The Application must be stopped in order to be upgraded.
- The Application must be able to connect to the database.
- A restart of the sever is required before the upgraded application can be configured again

Step 1: Navigating to the upgrade page

You can upgrade an application in two places:

- In the FADC navigation menu, go to Applications > Upgrade an Application.
- In the FADC navigation menu, go to Applications > Application Summary and press the "Upgrade" button on a stopped application.

Back to the top

Step 2: Selecting the application to upgrade

You will be presented with a drop-down list of all applications that you have permission to upgrade.

1. Select the application you wish to upgrade.
2. Click the browse button and navigate to the package you would like to upgrade to.

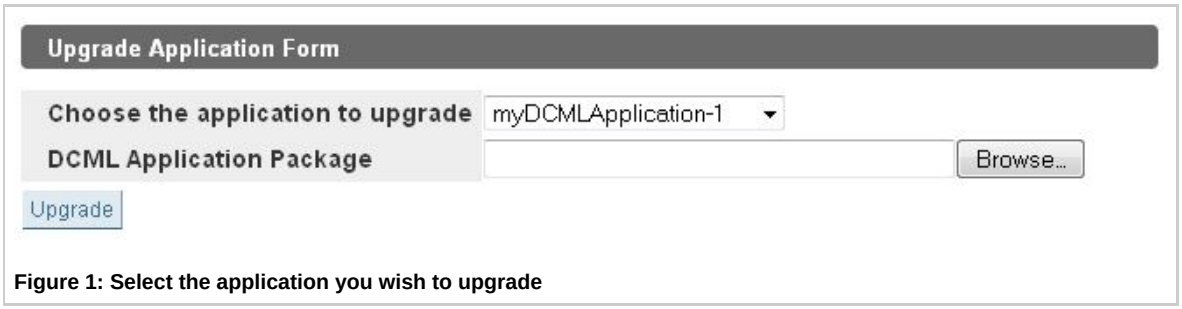

Figure 1: Select the application you wish to upgrade

Finally press "Upgrade" and you will be taken to the Application Configuration (FADC) in order to configure the upgraded application.

Back to the top

Upload Log or ZIP Files (FADC)

Upload Log or ZIP files (FADC)

This page allows you to manually upload a log file or an archive of logs. Select the DataCollector that you want to receive the logs, click the Upload button and navigate to your .log or .zip file. After pressing submit you will see the Application Summary page and your logs will be processed.

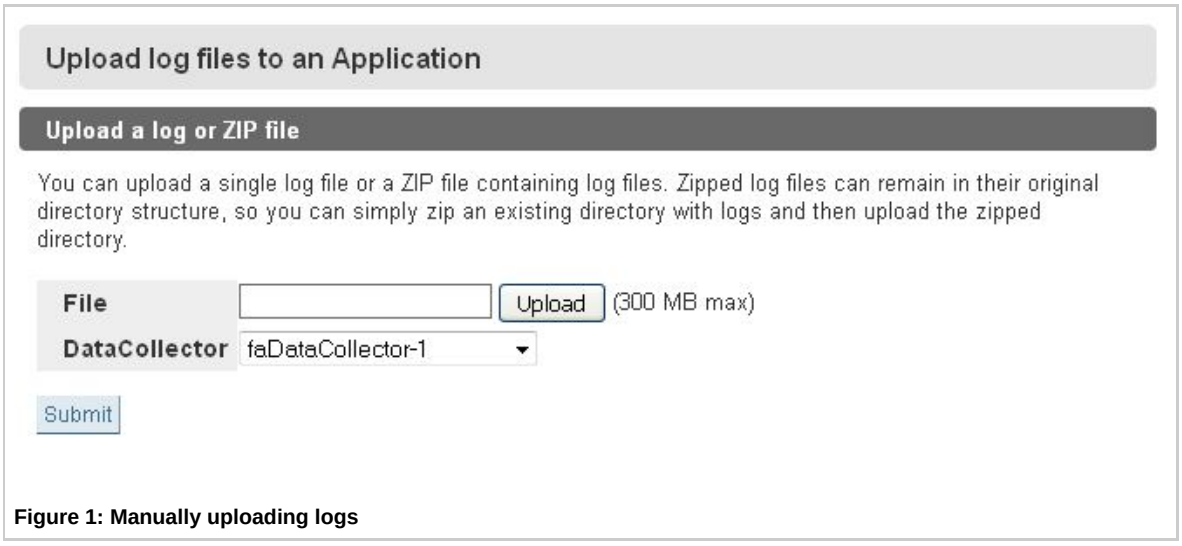

Figure 1: Manually uploading logs

Back to the top

If you have access to the file system of the computer where FADC is installed you can also add files by manually copying them into the belonging import directory of the application. FADC does observe the DataCollector import directories for new files.

Import History (FADC)

Import History (FADC)

The import history page displays up to the last 100 files imported by FADC.

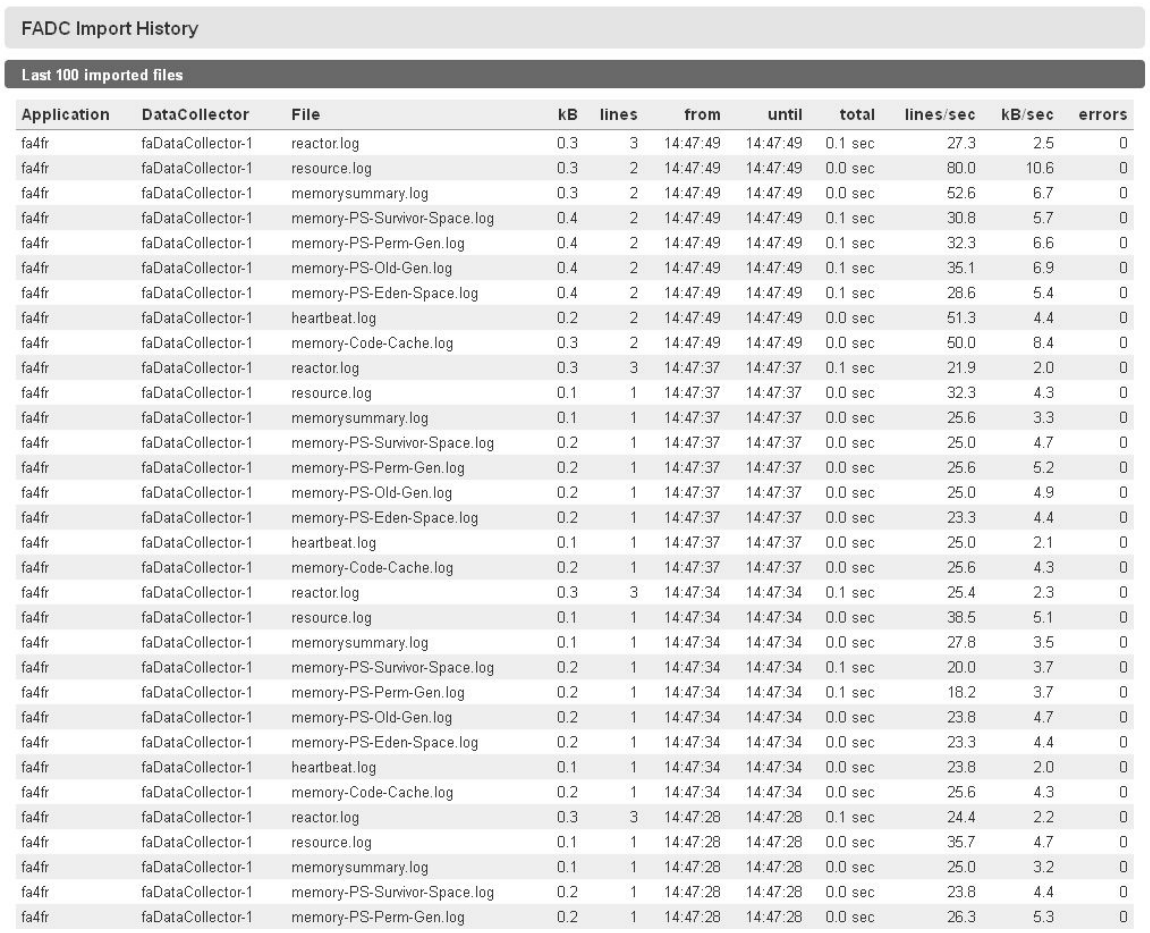

FADC Import History

Last 100 imported files

Application	DataCollector	File	kB	lines	from	until	total	lines/sec	kB/sec	errors
fa4fr	faDataCollector-1	reactor.log	0.3	3	14:47:49	14:47:49	0.1 sec	27.3	2.5	0
fa4fr	faDataCollector-1	resource.log	0.3	2	14:47:49	14:47:49	0.0 sec	80.0	10.6	0
fa4fr	faDataCollector-1	memorysummary.log	0.3	2	14:47:49	14:47:49	0.0 sec	52.6	6.7	0
fa4fr	faDataCollector-1	memory-PS-Survivor-Space.log	0.4	2	14:47:49	14:47:49	0.1 sec	30.8	5.7	0
fa4fr	faDataCollector-1	memory-PS-Perm-Gen.log	0.4	2	14:47:49	14:47:49	0.1 sec	32.3	6.6	0
fa4fr	faDataCollector-1	memory-PS-Old-Gen.log	0.4	2	14:47:49	14:47:49	0.1 sec	35.1	6.9	0
fa4fr	faDataCollector-1	memory-PS-Eden-Space.log	0.4	2	14:47:49	14:47:49	0.1 sec	28.6	5.4	0
fa4fr	faDataCollector-1	heartbeat.log	0.2	2	14:47:49	14:47:49	0.0 sec	51.3	4.4	0
fa4fr	faDataCollector-1	memory-Code-Cache.log	0.3	2	14:47:49	14:47:49	0.0 sec	50.0	8.4	0
fa4fr	faDataCollector-1	reactor.log	0.3	3	14:47:37	14:47:37	0.1 sec	21.9	2.0	0
fa4fr	faDataCollector-1	resource.log	0.1	1	14:47:37	14:47:37	0.0 sec	32.3	4.3	0
fa4fr	faDataCollector-1	memorysummary.log	0.1	1	14:47:37	14:47:37	0.0 sec	25.6	3.3	0
fa4fr	faDataCollector-1	memory-PS-Survivor-Space.log	0.2	1	14:47:37	14:47:37	0.0 sec	25.0	4.7	0
fa4fr	faDataCollector-1	memory-PS-Perm-Gen.log	0.2	1	14:47:37	14:47:37	0.0 sec	25.6	5.2	0
fa4fr	faDataCollector-1	memory-PS-Old-Gen.log	0.2	1	14:47:37	14:47:37	0.0 sec	25.0	4.9	0
fa4fr	faDataCollector-1	memory-PS-Eden-Space.log	0.2	1	14:47:37	14:47:37	0.0 sec	23.3	4.4	0
fa4fr	faDataCollector-1	heartbeat.log	0.1	1	14:47:37	14:47:37	0.0 sec	25.0	2.1	0
fa4fr	faDataCollector-1	memory-Code-Cache.log	0.2	1	14:47:37	14:47:37	0.0 sec	25.6	4.3	0
fa4fr	faDataCollector-1	reactor.log	0.3	3	14:47:34	14:47:34	0.1 sec	25.4	2.3	0
fa4fr	faDataCollector-1	resource.log	0.1	1	14:47:34	14:47:34	0.0 sec	38.5	5.1	0
fa4fr	faDataCollector-1	memorysummary.log	0.1	1	14:47:34	14:47:34	0.0 sec	27.8	3.5	0
fa4fr	faDataCollector-1	memory-PS-Survivor-Space.log	0.2	1	14:47:34	14:47:34	0.1 sec	20.0	3.7	0
fa4fr	faDataCollector-1	memory-PS-Perm-Gen.log	0.2	1	14:47:34	14:47:34	0.1 sec	18.2	3.7	0
fa4fr	faDataCollector-1	memory-PS-Old-Gen.log	0.2	1	14:47:34	14:47:34	0.0 sec	23.8	4.7	0
fa4fr	faDataCollector-1	memory-PS-Eden-Space.log	0.2	1	14:47:34	14:47:34	0.0 sec	23.3	4.4	0
fa4fr	faDataCollector-1	heartbeat.log	0.1	1	14:47:34	14:47:34	0.0 sec	23.8	2.0	0
fa4fr	faDataCollector-1	memory-Code-Cache.log	0.2	1	14:47:34	14:47:34	0.0 sec	25.6	4.3	0
fa4fr	faDataCollector-1	reactor.log	0.3	3	14:47:28	14:47:28	0.1 sec	24.4	2.2	0
fa4fr	faDataCollector-1	resource.log	0.1	1	14:47:28	14:47:28	0.0 sec	35.7	4.7	0
fa4fr	faDataCollector-1	memorysummary.log	0.1	1	14:47:28	14:47:28	0.0 sec	25.0	3.2	0
fa4fr	faDataCollector-1	memory-PS-Survivor-Space.log	0.2	1	14:47:28	14:47:28	0.0 sec	23.8	4.4	0
fa4fr	faDataCollector-1	memory-PS-Perm-Gen.log	0.2	1	14:47:28	14:47:28	0.0 sec	26.3	5.3	0

Figure 1: Import history

Back to the top

User (FADC)

User (FADC)

Overview

The User menu contains the Change Password (FADC) menu item. Choose it if you want to change the password of the 'admin' user.

Next Steps

Change Password (FADC)

Change Password (FADC)

Change Password (FADC)

Here you can change the password used to log into FADC. Simply enter the current password and the password you would like to switch to as well as confirmation of this new password.

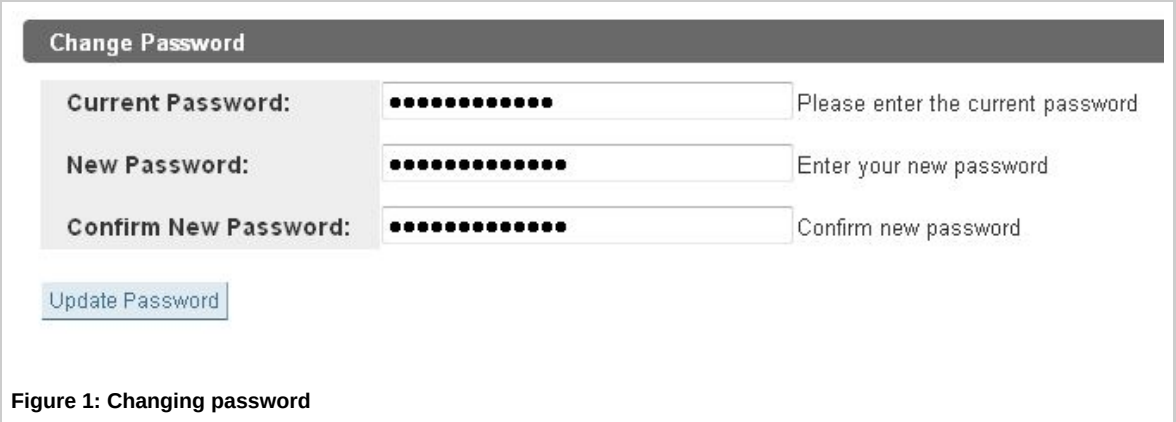

Figure 1: Changing password

Back to the top

System (FADC)

System (FADC)

Overview

From the System menu you can proceed to the Email Settings page, the Runtime Info page and to the License Info page. Finally you can click on the Logout menu to quit your current session.

Next Steps

Email Settings (FADC)

Runtime Info (FADC)

License Info (FADC)

Logout (FADC)

Email Settings (FADC)

Setting Up the Email Server

To receive notifications and reports from FusionAnalytics, you need to set up your email server.

In FADC this can be found under **System Email Settings** (Figure 1).

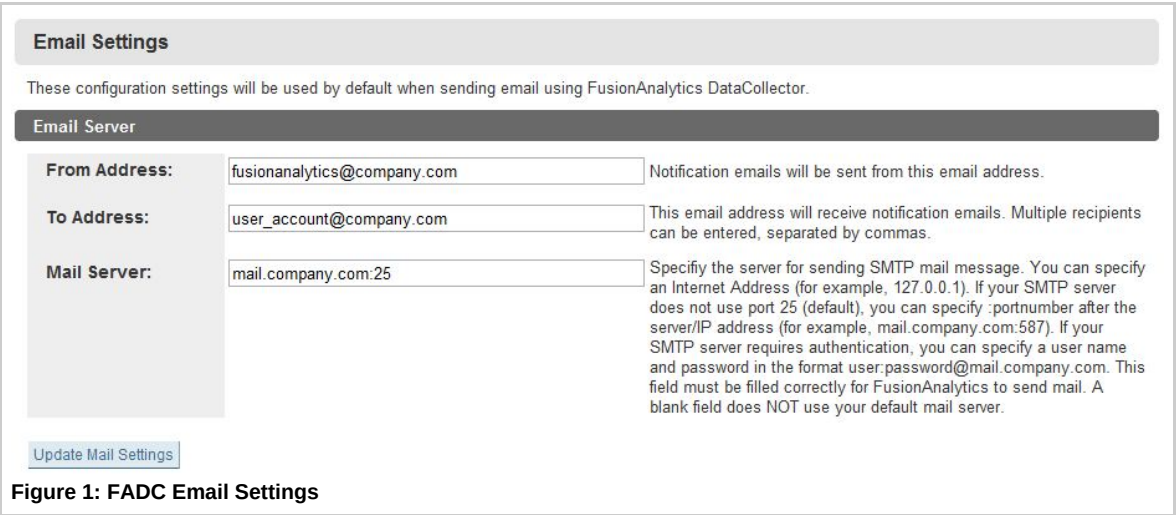

Figure 1: FADC Email Settings

The **From Address** should be the email address that you want any notifications sent from.

The **To Address** is the email address(es) that all notifications are sent to; you can enter multiple addresses, as long as they are separated

by a comma.

> ⚠ These addresses do not apply to the reporting feature of FusionAnalytics (which allows you to receive TAP and Daily Status reports on the status of your server by email). You can set different email addresses for these.

The **Mail Server** field is needed for **all** reports and notifications. This will need to be set to the name or IP address of the server that sends SMTP mail messages; e.g. **mail.company.com:25.**

By default the port number is set to 25, but you can specify a port number after the address if needed; e.g. **127.0.0.1:587**.

If you need to provide authentication, you will need to put this in front of the address: **username:password@mail.company.com:25**.

Back to the top

Runtime Info (FADC)

Runtime Info (FADC)

This page allows you to reset the runtime information displayed on the Application Summary (FADC) page. This will also clear the "Now importing" and "Most recently imported" sections on the Application Summary page.

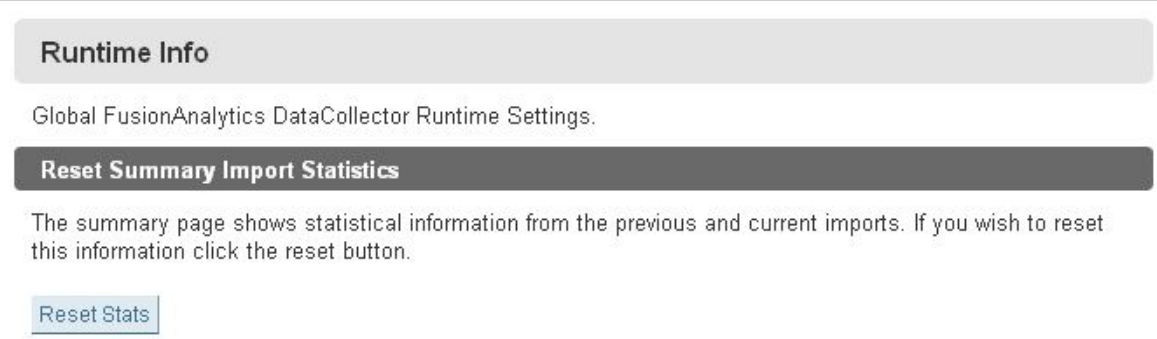

Figure 1: Resetting runtime information

Back to the top

License Info (FADC)

License Info (FADC)

The license info page (System > License Info) displays all of the information related to your FusionAnalytics license. For help installing and activating your license see Installing your license.

The following information is displayed for each license you have installed:

Name	Description	Example
Product Name	The name of your license.	FusionAnalytics for FusionReactor (Subscription for 1 year)
Product Family	The product family your license belongs to.	FA4FR
Expiry Date	The date your license will expire (not the same as activation expiry).	01-Jan-2012
License Status	The current status of your license.	Your license is valid and activated.
FusionAnalytics DataCollector	The number of datacollectors used and the maximum your license allows.	You are currently using 9 of 18 DataCollectors.
FusionAnalytics DataServices	The number of users and applications used and the maximum your license allows.	You are currently using 16 of 20 Applications and 31 of 40 Users.

You will also find buttons to activate or uninstall your license (depending on its current activation state) as well as the option to install a new license.

Back to the top

Logout (FADC)

Logout (FADC)

You can log out of FADC by either:

- clicking the logout link in the top right corner of FADC.
- navigating to System > Logout in the navigation menu.

Logging out will return you to the FusionAnalytics DataServices login screen.

Back to the top

DCML (FADC)

DCML (FADC)

Overview

This section provides you with a basic overview about the DataCollector Markup Language (DCML) which is used to define how a FADC application works.

A DCML file is a XML document which conforms to the FADC XSD schema and is used to

- declare which data should be imported,
- how this data has to be processed before it is imported and
- which additional tasks will be executed to bring the data into a format suitable to be used by FusionAnalytics Server.

A typical DCML configuration contains a root `<FadcConfig>` element with three children elements `<DataInput>`, `<DataOutput>` and `<DataManagement>` which define how incoming files are processed, how the belonging database schema looks like and which auxiliary tasks are used to manage the imported data.

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<FadcConfig version='1.0' xmlns='http://www.fusion-analytics.com/datacollector'>
    <DataOutput dbType='mssql'...>
    <DataInput...>
    <DataManagement...>
</FadcConfig>
```

Listing 1: DCML Structure

DataOutput

The `<DataOutput>` element defines the database type and schema - the tables, indices and statistics that will store the data provided by the `<DataInput>` element. The tables, indices and statistics defined here will be created in the database when you configure the application from the FADC web application.

```
<DataOutput dbType='mssql'>
    <Table name='table_1'...>
    <Index name='idx_table-1' onTable='table_1' fillFactor='90'...>
    <Statistics name='stat_table-1' onTable='table_1'...>
</DataOutput>
```

Listing 2: The DataOutput Element

Tables

A database table consists of a set of columns. In DCML this is defined as shown in Listing 3.

```
<DataOutput dbType='mssql'>
    <Table name='table_1' >
        <Column name='id' type='id' parse='false' />
        <Column name='date'/>
        <Column name='time'/>
        <Column name='dt' type='bigint'/>
        <Column name='message' type='varchar(30)' >
        <Column name='code' type='smallint' required='false'/>
    </Table>
</DataOutput>
```

Listing 3: The Table Element

When importing log files each line of the log file is tokenized into a list of column values matching the column definitions of the table element in the DCML. Per default a line is tokenized using the built-in high speed column parser which is able to parse the most common log line formats. For more control it is also possible to define a regular expression with the `<Table>` elements optional `separatorRegex` attribute which is then used to tokenize the input line into a list of column values.

The type attribute defines the data type of the column in the database. When inserting the data FADC converts it to the correct type automatically.
If a column has a `parse='false'` attribute FADC will not expect to find its value in the log line. If a column has no `type` attribute FADC will parse its value from the log line but ignore it. This is typically used to ignore redundant information in log lines. The `required` attribute is used to define database columns that are allowed to receive null values. Use it to process optional column values.

If you want to import log files into a table the order of the columns of the DCML table definition must match the columns of the lines in the log file.

The order of the `<Column>` elements used in the `<Table>` elements defines how to parse a line of data so that it maps to the columns of the table.

How a log file is processed by FADC

The following example shows how a log file is processed by FADC depending on the table definition of Listing 3.

Our example log file contains two lines which will be tokenized one after each other:

```
2008-11-17 13:36:40.671 1226925400671 "a message and a code" 0
2008-11-17 13:36:40.703 1226925400703 "a message without a code"
```

The first line is tokenized by the internal column parser into a list of strings:

```
"2008-11-17", "13:36:40.671", "1226925400671", "a message and a code", "0"
```

Then the column definitions of Listing 3 are applied to assign the values to their target columns:

```
dt:"1226925400671", message:"a message and a code", code:"0"
```

Column values without a `type` attribute are discarded and an appropriate SQL statement is added to the current batch of insert statements. The second line is then processed in the same way. Since the line does not have a 'code' value the belonging SQL insert statement does not specify a value for the 'CODE' column. This works because the 'CODE' column has a `required='false'` attribute set. Once the maximum size of the batch is reached or no more lines are available the insert operation is committed to the database.

Indices and Statistics

To keep your database efficient DCML allows you to define additional `<Index>` and `<Statistics>` elements under the `<DataOutput>` element. Listing 4 and Listing 5 show example definitions for an index called 'idx_table_1' and statistics named 'stats_table_1'.

```
<Index name='idx_table_1' onTable='table_1' fillFactor='90'>
    <Column name='dt'/>
    <Column name='message'/>
</Index>
```

Listing 4: An Index Definition

```
<Statistics name="stats_table_1' onTable='table_1'>
    <Column name='dt'/>
</Statistics>
```

Listing 5: A Statistics Definition

DataInput

The `<DataInput>` element defines the sources from which data is imported. It can contain one or more `<DataCollector>` elements which logically group a set of data belonging together (e.g. the log files of a FusionReactor instance and the ColdFusion server log files).

```
<DataCollector name='faDataCollector-1' timezone='Europe/Berlin'>
    <Events.../>
    <OnCreate program="onCreateDataCollector.groovy"/>
    <OnDelete program="onDeleteDataCollector.groovy"/>
    <OnImport program="onImportFile.groovy"/>
</DataCollector>
```

Listing 6: A DataCollector Definition

Data is imported if the regular expression defined by the `matchFile` attribute of a `<Match>` element as shown in Listing 7

```
<Events>
    <Match name='sample' matchFile='(?i)sample.*[.]log'>
        <TargetTable table='table_1'/>
        <Line regex='.*'/>
    </Match>
</Events>
```

Listing 7: A Match Definition

matches the name of a file found in an import directory. The `<TargetTable>` element defines with its `table` attribute into which database table data read from the file will be inserted. The `<Line>` element defines with its `regex` attribute which parts of each line should be taken - here the regular expression `'.*'` specifies to use the complete line.

Each application in FADC has an import directory containing a directory for each of its DataCollectors. These directories are monitored by FADC so that any file found there is detected and processed by FADC.

Below the `<Events>` element you can optionally use `<OnCreate>`, `<OnDelete>` and `<OnImport>` elements to execute a user defined script triggered by the belonging event.

DataManagement

The `<DataManagement>` element defines auxiliary jobs written in Groovy or BeanShell which are executed depending on their `Trigger` element. The job defined in Listing 8 will result in FADC executing the Groovy script 'archiveFiles.groovy' at 1:00am every night. The script is expected to be in the DCML applications root directory.

```
<DataManagement>
    <Job program='archiveFiles.groovy' name='Archived Logfiles Management'>
        <Trigger name='archiveFiles' cronPattern='0 0 1 * * ?'/>
    </Job>
</DataManagement>
```

Listing 8: A Job Definition

Cron Patterns

FADC uses the Quartz Scheduler to manage its jobs. Read the tutorial available on their site to learn more about using cron patterns.

Back to the top

On This Page:

DataCollector Markup Language

A domain specific XML language used to define from which files data is imported and how this data is mapped to a relational database structure.

The FADC XSD documentation is available here

FusionAnalytics DataServices (FADS)

FusionAnalytics DataServices (FADS)

Overview

This section provides information on using and configuring FusionAnalytics DataServices.

Next Steps

Applications (FADS)

- Application Summary (FADS)
  - Application Configuration (FADS)
  - Application Detail (FADS)
  - Application Scope (FADS)
- Upload an Application (FADS)
- Upgrade an Application (FADS)

Users (FADS)

- Users Summary (FADS)
  - User Mappings

Scheduled Tasks (FADS)

- User Tasks Summary
- Application Tasks Summary

System (FADS)

- Sessions
- Controllers
- Asynchronous Jobs
- Server Settings
- License Info (FADS)
- Logout (FADS)

Applications (FADS)

Applications (FADS)

Overview

This section details how to use applications within FusionAnalytics DataServices.

Next Steps

Application Summary (FADS)

- Application Configuration (FADS)
- Application Detail (FADS)
- Application Scope (FADS)

Upload an Application (FADS)

Upgrade an Application (FADS)

Application Summary (FADS)

Application Summary (FADS)

Overview

This section goes into detail about how to navigate and use the Application Summary page in FusionAnalytics DataServices.

The Application Summary page gives you control over all of your FADS applications and is the first page that is displayed in FADS. The options available for an application vary depending on whether it is started or stopped.

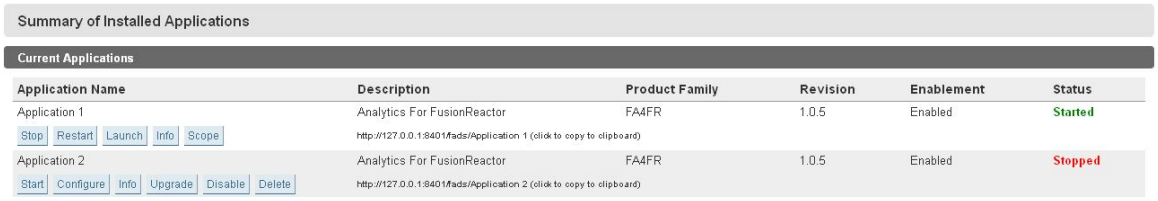

Figure 1: Application Summary

> **Please Note:** You must be an administrator of at least 1 application to see the Application Summary, and only applications for which you are an administrator or user will be displayed.

Options for a started application

> When an application is started, it can be accessed via the Adobe AIR Client or the Web Client by users with sufficient permissions. A started application cannot be configured.

Name	Description
Stop	Stops the application. This will remove all scheduled jobs from the scheduler, preventing them from running until the application is started again.
Restart	An application restart causes the currently running application to be stopped and started automatically with one button click.
Launch	Launches the FusionAnalytics DataServices Web Client
Info	Displays the Application Detail page. See Application Detail (FADS)
Scope	The Application Scope will display the application's essential elements and their details. See Application Scope (FADS)

Options for a stopped application

Name	Description
Start	Starts the application. This will add all scheduled jobs to the scheduler, ensuring they run at the next allotted time.
Configure	Displays the application configuration page, used to configure database connections and other application properties. See Application Configuration (FADS)
Info	Displays the Application Detail page. See Application Detail (FADS)
Upgrade	Allows the application to be upgraded to a later version.

Enable/Disable	When an application is enabled, it will be started when the server restarts regardless of its previous stopped/started state and its jobs will be added to the scheduler. A disabled application will not be started when the server restarts and its jobs will not be added to the scheduler.
Delete	Deletes an application, removing all association to it within FADS including scheduled jobs and user mappings. Some elements associated to the application cannot be deleted until the server is restarted.

Next Steps

Application Configuration (FADS) Application Detail (FADS) Application Scope (FADS)

Application Configuration (FADS)

Application Configuration (FADS)

The Application Configuration page allows you to configure various properties of an application.

Name/Value Elements

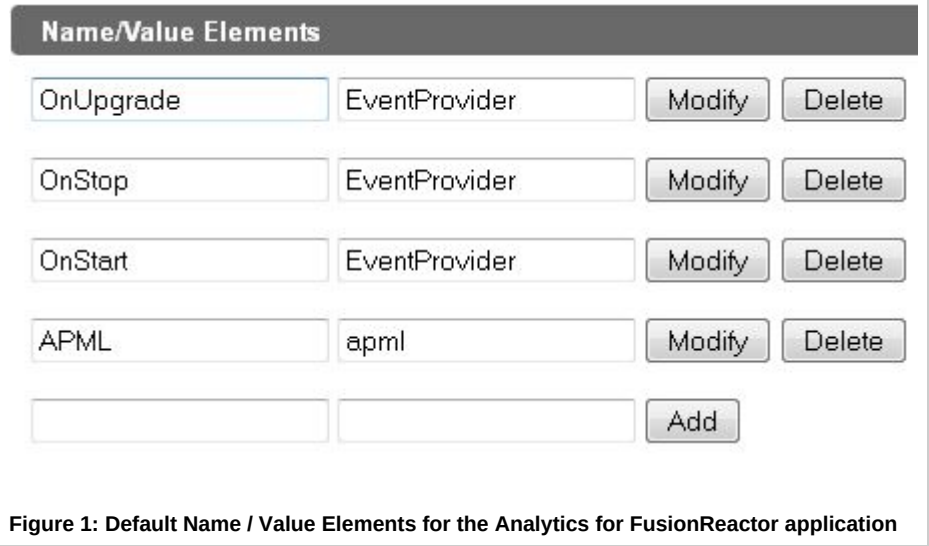

Figure 1: Default Name / Value Elements for the Analytics for FusionReactor application

Here you can add, modify or remove Name/Value Elements for your application.

- *For the default Analytics for FusionReactor application, these have been preconfigured for you and should not be changed.*

Back to the top

Application Scheduled Tasks

Here you can activate or deactivate Scheduled Tasks (FADS) that are run by the application.

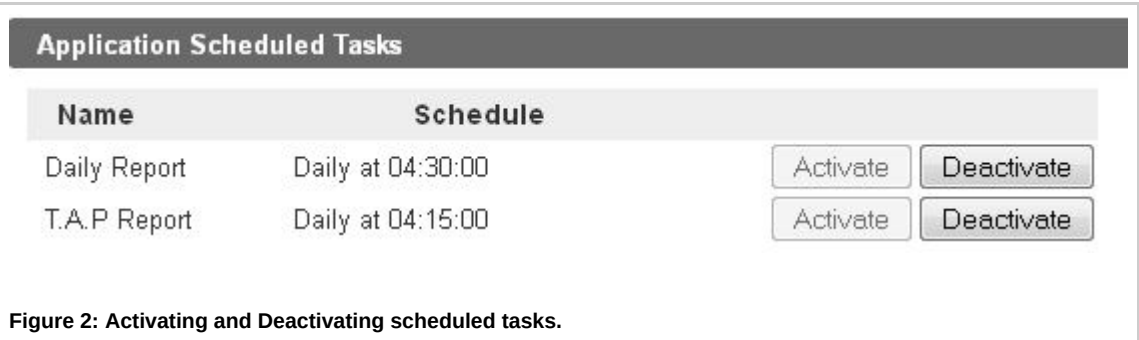

Figure 2: Activating and Deactivating scheduled tasks.

Back to the top

Data Sources

This is where you can configure data sources for your application. Please fill in the settings that match your database.

The "Name" field is simply anything you wish to use as a name for this data source. It is **not** the name of the database. The name of the database must be added to the URL field as shown below.

Every FA4FR application contains a default datasource (frdb) for FusionReactor which cannot be renamed.

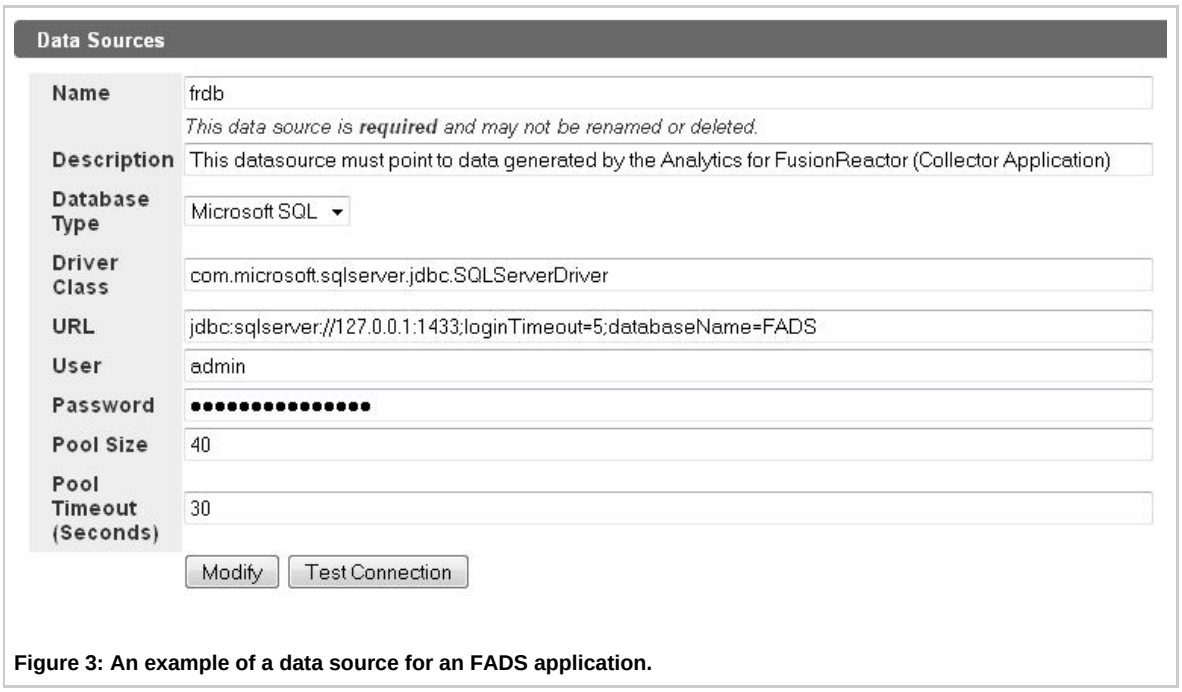

Figure 3: An example of a data source for an FADS application.

Field Name	Description	Example
Name	The name of your datasource. Every FA4FR application contains a default data source(frdb) for fusion reactor which cannot be renamed.	DataSource1
Description	Enter a brief description of this data source.	This data source is used to connect FusionAnalytics to FusionReactor
Database Type	The type of database you are using. Currently only Microsoft SQL Server is supported.	Microsoft SQL
Driver Class	The driver class tells FusionAnalytics the name of the JDBC database driver to use.	com.microsoft.sqlserver.jdbc.SQLServerDriver
URL	This specifies the JDBC connection URL used to connect to the database. Customize this field for your own database.	jdbc:sqlserver://127.0.0.1:1433;loginTimeout=5;databaseName=fadb
User	A username that has full access to the database	
Password	The password for the above account	
Pool Size	This is the number of database connections that can be open at any one time.	40
Pool Timeout (Seconds)	The number of seconds that have to pass before the pool times out.	30

If you upload a new application into FADS and want to use a new database with it, you must upload an FADC application on the same database to automaticaly create all of the required tables and database objects. This will then allow you to launch your FADS application.

Back to the top

Application Detail (FADS)

Application Detail (FADS)

There are three sections to the application detail page:

- Summary
- Manifest Data
- Data Dump

Summary

The summary shows basic information about the application.

Name	Description
Key	The application deployment key.
Name	The application name.
Revision	The application version.
Deployed	The date and time that the application was deployed.
Enablement	Whether the application is enabled or disabled.
Current State	The current state of the application (Stopped, Started, Installed, Upgraded).

Back to the top

Manifest Data

The Manifest Data section displays information related to the installation of the application.

Key	Description
Location	The physical location of the application folder.
Provider	The application provider.
ID	The UUID (Universally unique identifier) of the application.
Support	The support email address.
Configuration	The location of the application configuration.
Analytics	Detailed information about the application. *Hover over any of these for a detailed list.*

Back to the top

Data Dumps

The Data Dump section displays any exception log data that was last output. If the application could not be started because of an error within the application, detailed log data would be shown here.

Back to the top

Application Scope (FADS)

Application Scope (FADS)

The application scope shows essential elements of an application and their details.

Key	Type	Description
id	java.util.UUID	The UUID (Universally unique identifier) of an application
revision	java.lang.String	The current version of the application
application	com.intergral.fads.impl.app.APMLApplication	The application type
datasources	java.util.ArrayList	A list of datasources used by the application
description	java.lang.String	The application description
name	java.lang.String	The application name
deploymentKey	java.lang.String	The application deploymentKey
provider	java.lang.String	The application provider
productFamily	java.lang.String	The product family to which this application belongs
elements	java.lang.String	The Name/Value elements for this application. See Application Configuration (FADS)
supportURL	java.net.URL	The mailto: address of the support inbox for this application.

Back to the top

Upload an Application (FADS)

Uploading and Installing FADS Applications (APML)

Step 1: Navigating to the upload page

In the FADS navigation menu, go to Applications > Upload an Application

Back to the top

Step 2: Uploading an application

Choose a name for the application and enter it into the name input field. Click the browse button and navigate to the application package you wish to upload. Click the upload button to upload your application.

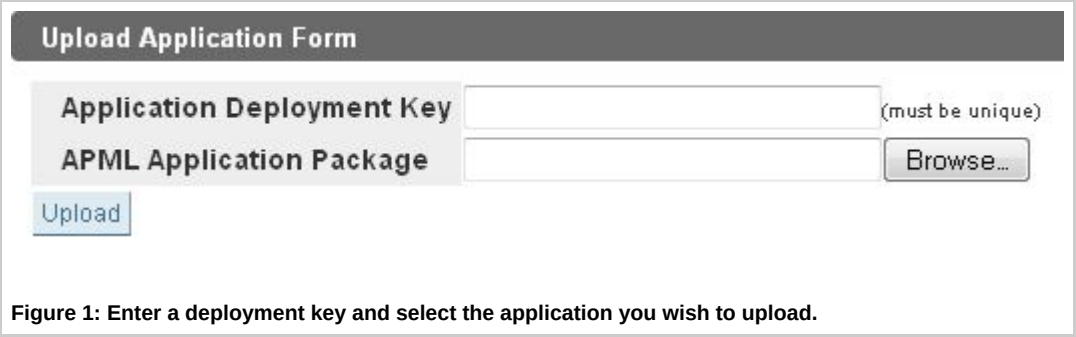

Figure 1: Enter a deployment key and select the application you wish to upload.

If your application uploaded successfully you will be taken to the Application Configuration (FADS) page where you will be able to configure your application.

If your license does not allow any more applications you can either...

- Delete an application belonging to the same product family as the one you're trying to upload
- Upgrade to a license with bigger application capacity

An application cannot be named:

- with a slash character (neither'\' nor '/') in any position
- with a semicolon (';') in any position
- with a colon (':') in any position
- as a single dot character ('.')
- as a single at character ('@')
- the same as an existing application (different capitalization is not considered unique)

Back to the top

On this page:

Upgrade an Application (FADS)

Upgrading an Application

An application must meet the following prerequisites before upgrading:

- The new application's version number must be higher than the current application's version number.
- Deployment keys cannot contain the slash character (neither '/' nor '\'), the semicolon or the colon in any position; nor can they named with a single dot ('.') or at ('@').
- The Application must be stopped in order to be upgraded.
- A restart of the sever is required before the upgraded application can be configured again

Step 1: Navigating to the upgrade page

You can upgrade an application in two places:

- In the FADS navigation menu, go to Applications > Upgrade an Application.
- In the FADS navigation menu, go to Applications > Application Summary and press the "Upgrade" button on a stopped application.

Back to the top

Step 2: Selecting the application to upgrade

You will be presented with a drop-down list of all applications that you have permission to upgrade.

1. Select the application you wish to upgrade.
2. Click the browse button and navigate to the package you would like to upgrade to.

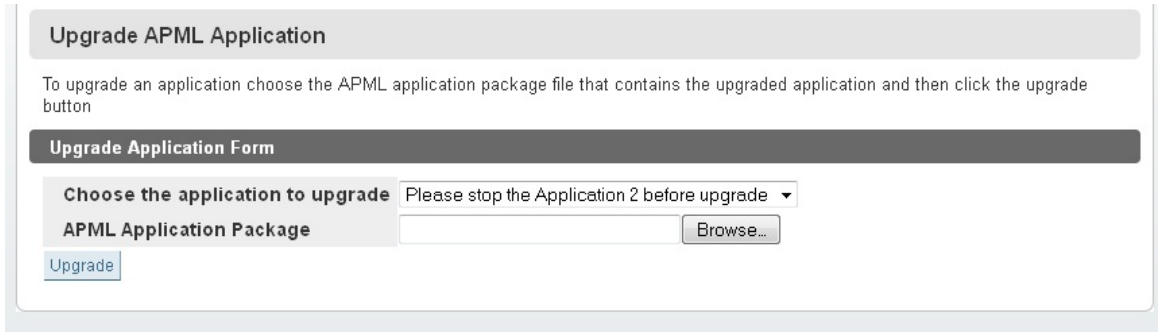

Figure 1: Select the application you wish to upgrade

Back to the top

Step 3: Upgrading the application

Finally press "Upgrade" and you will be taken to the application configuration screen in order to configure the upgraded application. See **Application Configuration (FADS)** for more information on configuring applications.

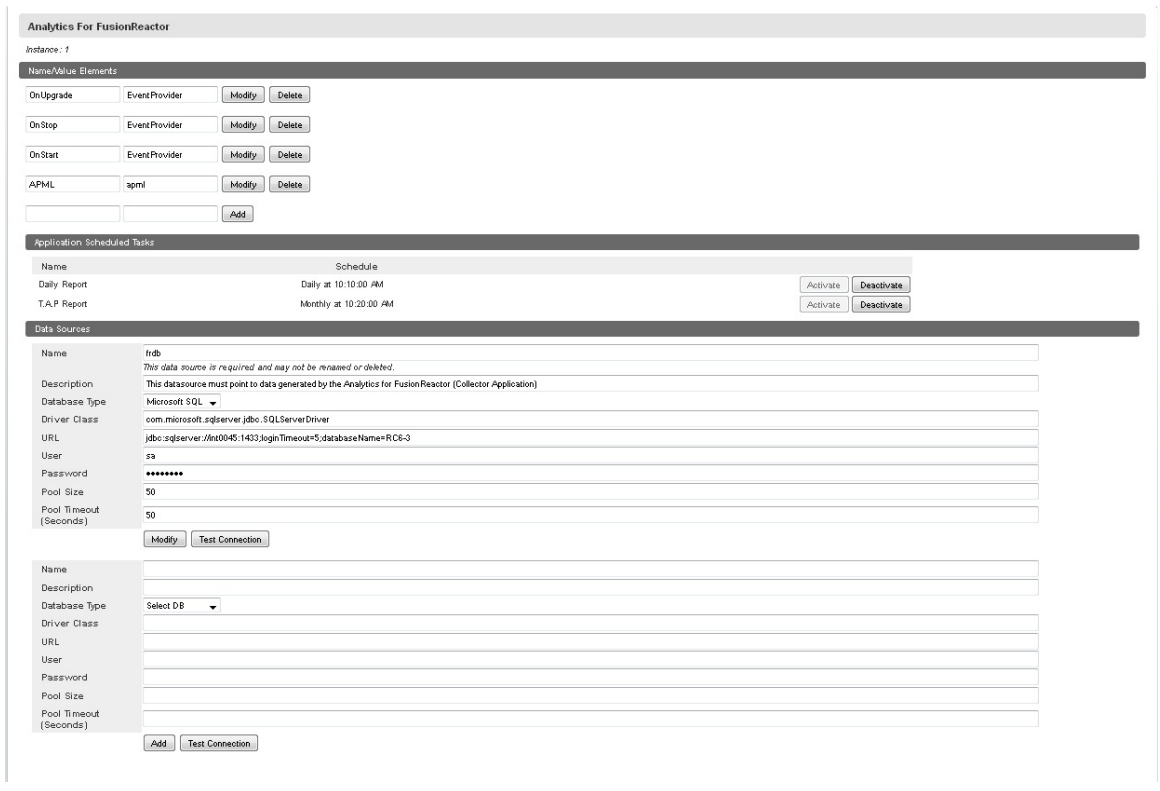

Figure 2: Application Configuration page

Back to the top

Users (FADS)

Users (FADS)

Overview

This section details information about users in FusionAnalytics DataServices.

Next Steps

Users Summary (FADS)

- User Mappings

Users Summary (FADS)

User Summary (FADS)

Overview

This section details creating, updating and mapping users as well as explaining the FusionAnalytics DataServices users interface.

You are able to do the following from the user summary page:

- Create new users
- Delete users
- Edit users
- Control user mappings (See User Mappings).

Creating new Users

To create a new user, navigate to Users > User Summary. You will see the Users Summary screen in the content area. Click on the "Create a New User" button and fill out the displayed form:

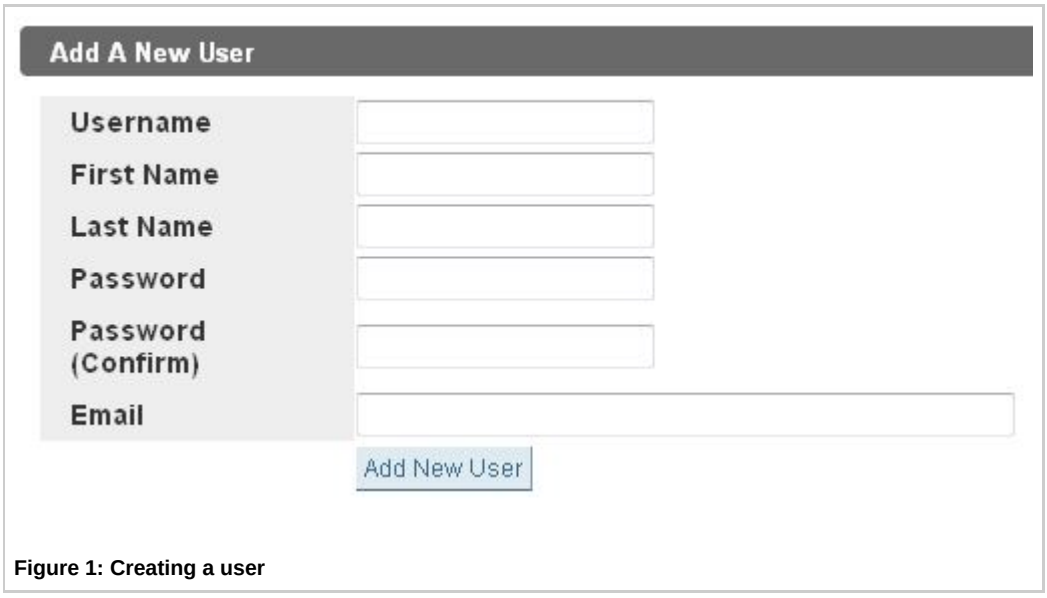

Figure 1: Creating a user

Each user created must have the following fields completed correctly in order to be created:

- Username
- Password

- Password (Confirm)

The following fields are not mandatory for user creation:

- First Name
- Last Name
- Email **(Required if you want this user to receive reports via email)**

> Once a user is created, you cannot change their username. Please ensure you have entered the correct username before submitting.

Deleting users

To delete a user, navigate back to the User Summary page (Users > User Summary). You will see a list of users present on your system:

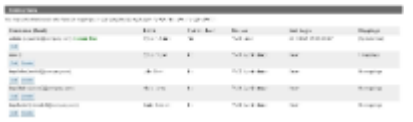

Figure 2: Deleting a user

To delete a user, simply press the Delete button and press OK at the confirmation dialog.

> **You may only delete users who either:**
>
> - have no mappings.
> - are mapped **only** to applications for which you are an administrator.

Editing a user

To edit a user, follow the steps in Deleting Users but click the Edit button instead of Delete. Then fill out the form as described in Creating new users

User Mappings

To learn about user mapping, see User Mappings.

Next Steps

User Mappings

User Mappings

User Mappings

There are two types of user mapping in FADS:

- Administrator
- User

There is also a single System User role.

The Administrator Role

The administrator role grants full access to an application. An application administrator can:

- Log into FADS and FADC.
- Start and stop the application.
- Access all functions made available by the Application Summary. See Application Summary (FADS).
- Add, remove or modify any "user" mappings for that application.
- Launch the web client.
- View the application admin perspectives in the client.
- Edit application scheduled tasks.
- Run application scheduled tasks.
- Deactivate application scheduled tasks.

⚠ A user only requires a single administrator mapping to gain access to FADS and FADC, though only applications they are mapped to will be visible.

Back to the top

The User Role

The user role grants basic access to an application. An application user can:

- Access the application via the AIR client.
- Access the web client via a direct link.

An application user cannot log into FADS or FADC, nor can they view the admin perspectives in the client.

Back to the top

System User role

There is a single System User within FADS. The System user will have complete access to all applications within FADS and FADC, all scheduled tasks for all applications (both application and user) and full access rights in regards to the modification of users. If you open the "Users Summary" page (Users > Users Summary) as the system user you will see that a system user has no mappings displayed with its edit screen. It is automatically mapped to all applications

Back to the top

Mapping a user

To map a user to an application, go to the Users Summary (FADS) and click edit on the user you wish to map. You will be presented with the edit screen and will find the mapping section at the bottom. If you have applications installed and are an administrator to any of these applications you should see the following screen:

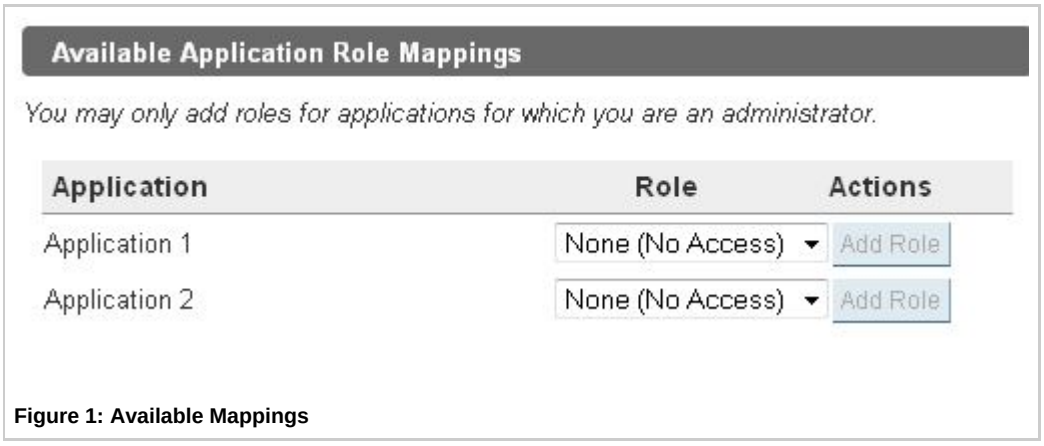

Figure 1: Available Mappings

Simply select the mapping from the drop-down box and press Add Role.

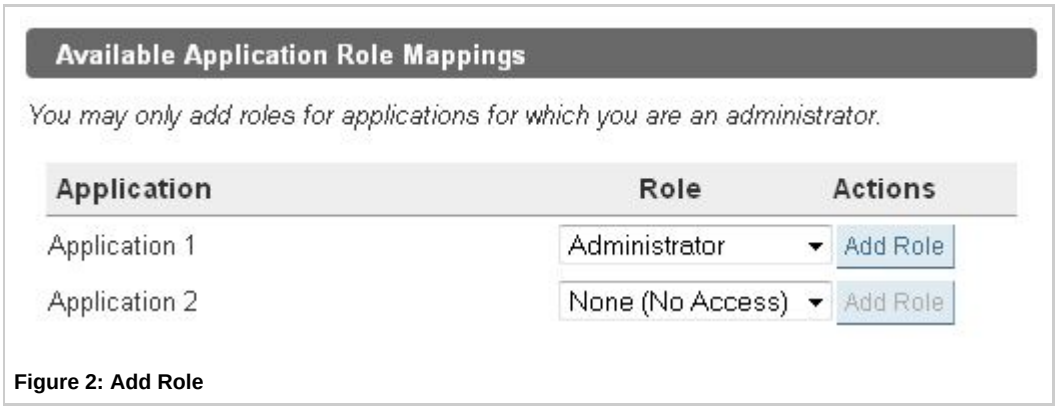

Figure 2: Add Role

Back to the top

Editing a mapping

To edit a user mapping, go to the Users Summary (FADS) and click edit on the user you wish to change. You will be presented with the edit screen where you will find the current mapping section. If you are an administrator to the application you with to edit a mapping for you should see the following screen:

Figure 3: Current Mappings

You can either alter the mapping and press Modify, or delete the mapping completely.

Back to the top

Scheduled Tasks (FADS)

Scheduled Tasks (FADS)

Overview

This section explains how to create and modify both User and Application Scheduled Tasks in FusionAnalytics DataServices.

Next Steps

User Tasks Summary

Application Tasks Summary

User Tasks Summary

FADS User Scheduled Tasks

Steps to Setup:

- Creating a User Job
- Running Jobs

Database Connection

Scheduled tasks that use a database require a database connection to be alive.

Back to the top

Creating a User Job

The following are the types of schedules for jobs

- On-Time: Will run once on the designated date and time and will not be scheduled again unless configured again.
- Recurring Daily: Job will execute every Day at the following time
- Recurring Weekly: Job will execute every Week at the following time
- Recurring Monthly: Job will execute every Month at the following time
- Daily Every: Job will execute within the specific time periods defined every day

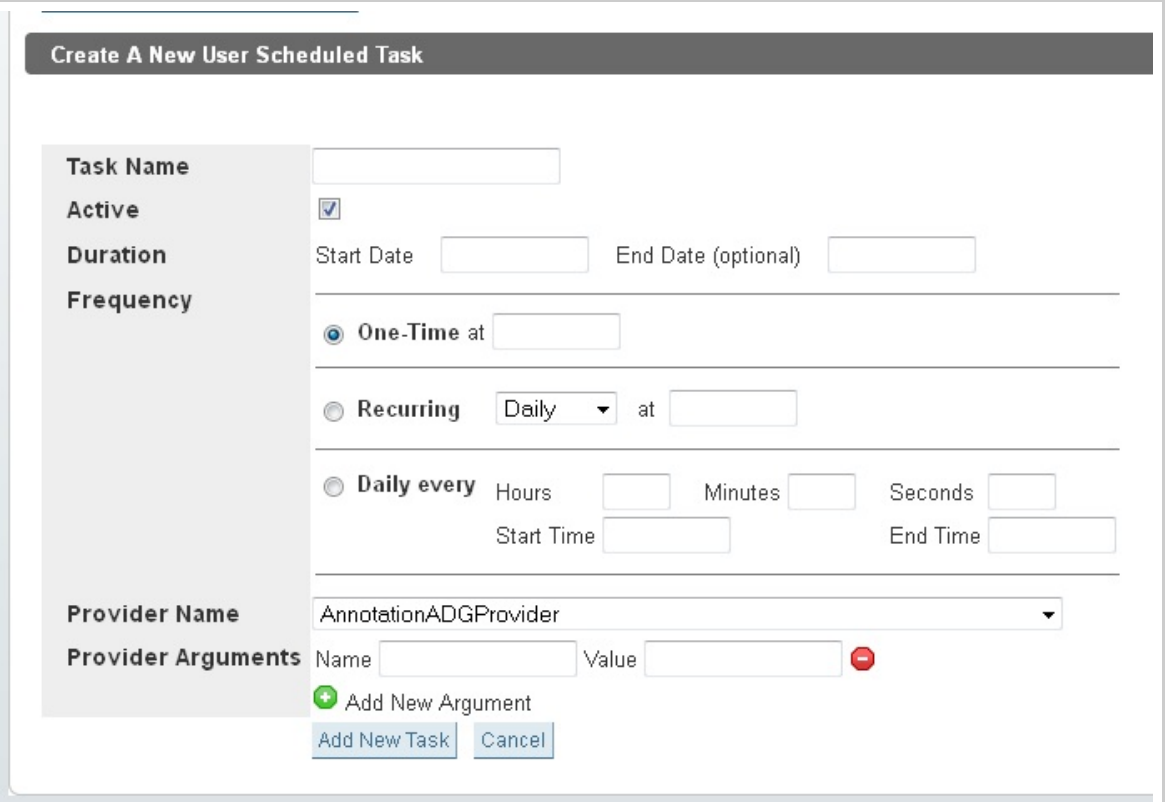

Figure 1: Create a New USer Scheduled Task

The following is a list of validation constraints in place for the entry fields

Name	Date Format
One-Time	hh/mm + (AM/PM)
Recurring	hh/mm + (AM/PM)
Daily Every	hh/mm + (AM/PM)
Start Date	mm/dd/yyyy
End Date (optional)	mm/dd/yyyy

Any invalid input will be identified and appropriate error messages will prompt valid input over the invalid fields.

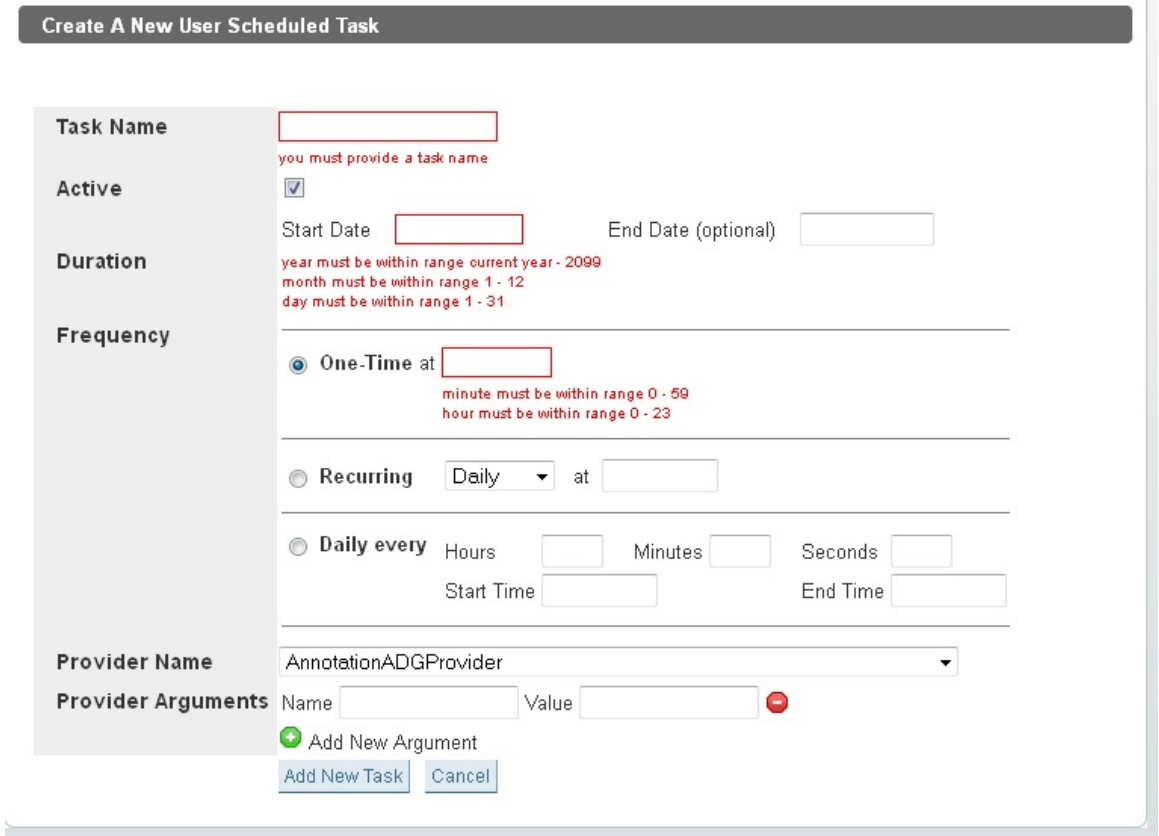

Figure 2: Create a New User Scheduled Task - Validation

Back to the top

Running Jobs

A job can only be executed and remain within the scheduler if the application the job belongs to is started. If the application is stopped then the job is removed from the scheduler and will not be executed even if it's time to run as been reached. If a job's application is started the job will run upon its scheduled date and time. There is an option besides every job, user or application that allows a specific job to be executed at the current time via the 'Run Now' button. This button will be unusable if the job's application is stopped.

Back to the top

Application Tasks Summary

FADS Application Scheduled Tasks

Steps to Setup:

- Editing Application Job
- Running Jobs

Database Connection:

Scheduled tasks that use a database require a database connection to be alive. See Application Configuration (FADS).

Back to the top

Editing Application Job

Application scheduled tasks can only be edited, not created, as they are pre-installed with applications.
To edit an application scheduled job from FADS navigate to Scheduled Tasks > Application Tasks Summary. An admin has the permission to edit every aspect of the application job except the job's name and data provider.

The following are the types of schedules for jobs

- On-Time: Will run once on the designated date and time and will not be scheduled again unless configured again.
- Recurring Daily: Job will execute every Day at the following time
- Recurring Weekly: Job will execute every Week at the following time
- Recurring Monthly: Job will execute every Month at the following time
- Daily Every: Job will execute within the specific time periods defined every day

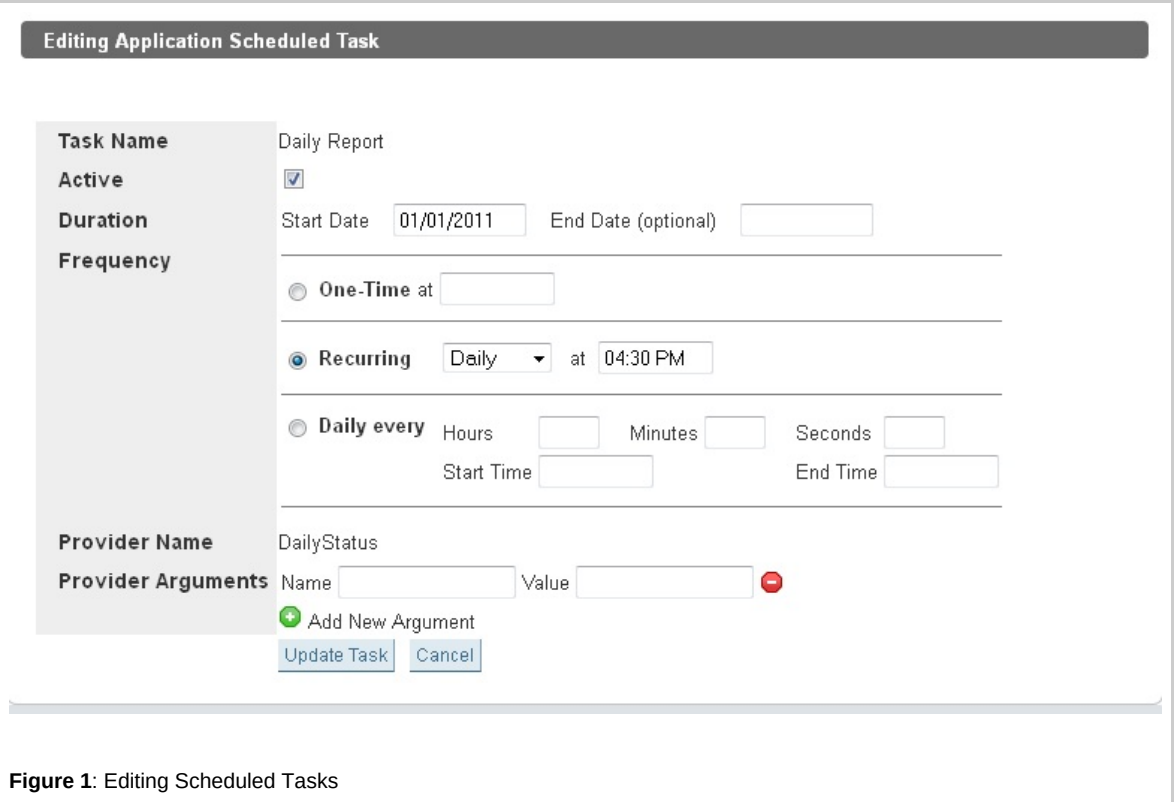

Figure 1: Editing Scheduled Tasks

The following is a list of Validation Constraints in place for the entry fields

Name	Date Format
One-Time	hh/mm + (AM/PM)
Recurring	hh/mm + (AM/PM)
Daily Every	hh/mm + (AM/PM)
Start Date	mm/dd/yyyy
End Date (optional)	mm/dd/yyyy

Back to the top

Running Jobs

A job can only be executed and remain within the scheduler if the application it belongs to is started. If the application is stopped then the job is removed from the scheduler and will not be executed even if it's time to run as been reached. If a job's application is started the job will run upon its scheduled date and time. There is an option besides every job, user or application that allows a specific job to be executed immediately via the 'Run Now' button. This button will be unusable if the job's application is stopped.

Back to the top

Disabling an Application Job

You can disable an application by either:

- visiting its configuration page

Figure 2: Application Scheduled Tasks

- vising the Application Tasks Summary page

Figure 3: Application Tasks Summary Page

Disabling an Application Job will cause it to be removed from the scheduler and prevent it from executing at its designated time. Enabling the job will place it back into the scheduler.

Back to the top

System (FADS)

System (FADS)

Overview

This section details the options available within the System section of FusionAnalytics DataServices.

Next Steps

Sessions

Controllers

Asynchronous Jobs

Server Settings

License Info (FADS)

Logout (FADS)

Sessions

Sessions

A session is considered to be one user using FusionAnalytics through one client. When you log into the FusionAnalytics Client you are starting a session and logging out terminates that session. If you were to log into the AIR client and, at the same time, log into a separate server using the Web client then you will have creates two separate sessions. To see all of the currently active sessions, select "Sessions" from within the "System" section of the Table of Contents.

With no active sessions the content area will look like this:

Figure 1: Session Summary

Example Session: Web Client

Follow these steps in order to generate a session using the Web Client and view it from within FADS:

- Log into FADS as either System user or Admin user
- If one is not already started then start an application you have admin rights to (if administrator)
- Press the "Launch" button to launch the web client

- Log in with valid credentials
- Return to FADS and navigate to the System -> Session page. The session's details will be listed i.e Username of user logged in, Login Time, Session ID and the User's Name

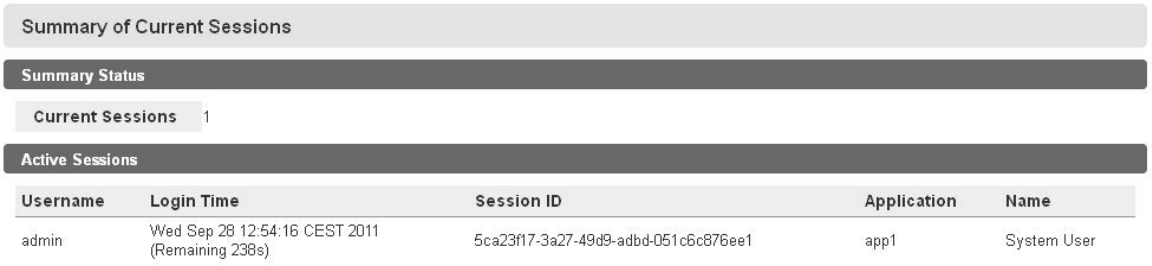

Figure 2: Current Session Summary

Back to the top

Example Session: AIR Client

Follow these steps in order to generate a session using the Web Client and view it from within FADS:

- Log into FADS as either System user or Admin user
- If one is not already started then start an application you have admin rights to (if administrator)
- Click on the URL of the application in order to copy it to your clipboard.

Figure 3: Copy the Application URL from FADS

- Start the AIR client and (if one is not already set up) create a new link to your application by pasting the URL into the top bar and clicking "Add Application"

Figure 4: Paste Application URL into the AIR Client

- Double-click on the server icon to bring up the login screen and log in with valid credentials

Back to the top

Controllers

Controllers

All actions within the Data Server are handled by Controllers. These controllers are a logical breakdown of each responsibility of the Data Server, and are complex, self-contained components. E.g., the Cache Controller which is responsible for intelligently storing and retrieving large amounts of data very quickly. Selecting the “Controllers” menu item from the “System” section in the main Table of Contents will display all the controllers available within DataServices. Some of the more advanced controllers also allow you to perform some action. These actions vary from clearing records, dumping the state to the screen and also stopping/starting particular functionality. Under normal circumstances, you shouldn't need to change anything in the Controllers page; however you may be directed to obtain information from here for debugging, should you make a support request.

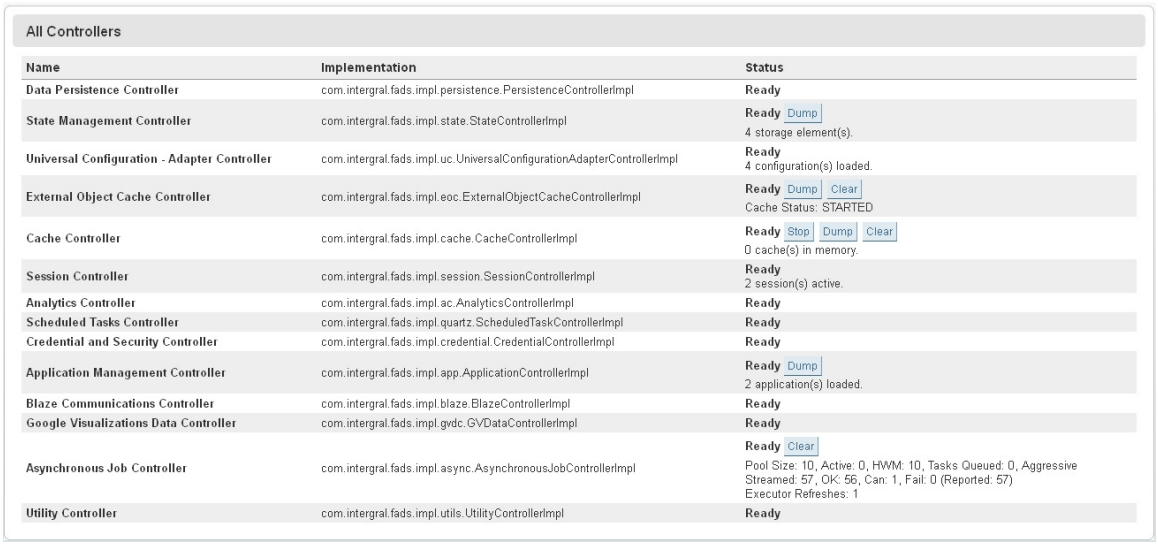

Figure 1: All Controllers

Back to the top

Asynchronous Jobs

Asynchronous Jobs

When the client (either the FusionAnalytics AIR client, or the web client) requests a set of data or an analytical operation from the server, the client creates a job which is streamed into the server.

In order to remain responsive, the client simply waits for the job to complete in the background. We say that these jobs are asynchronous – independent in time.

When DataServices receives the job, it queues it, executing it when it has available resources. With many jobs being queued and executed simultaneously, it is sometimes difficult for the user to comprehend what is happening when a job takes time to execute. By selecting the “Asynchronous Jobs” menu item from the “System” section in the main Table of Contents you will be able to view the current status of the job controller. You are able to see all the queued jobs, currently running jobs and a history of past jobs. It also displays the number of failed jobs.

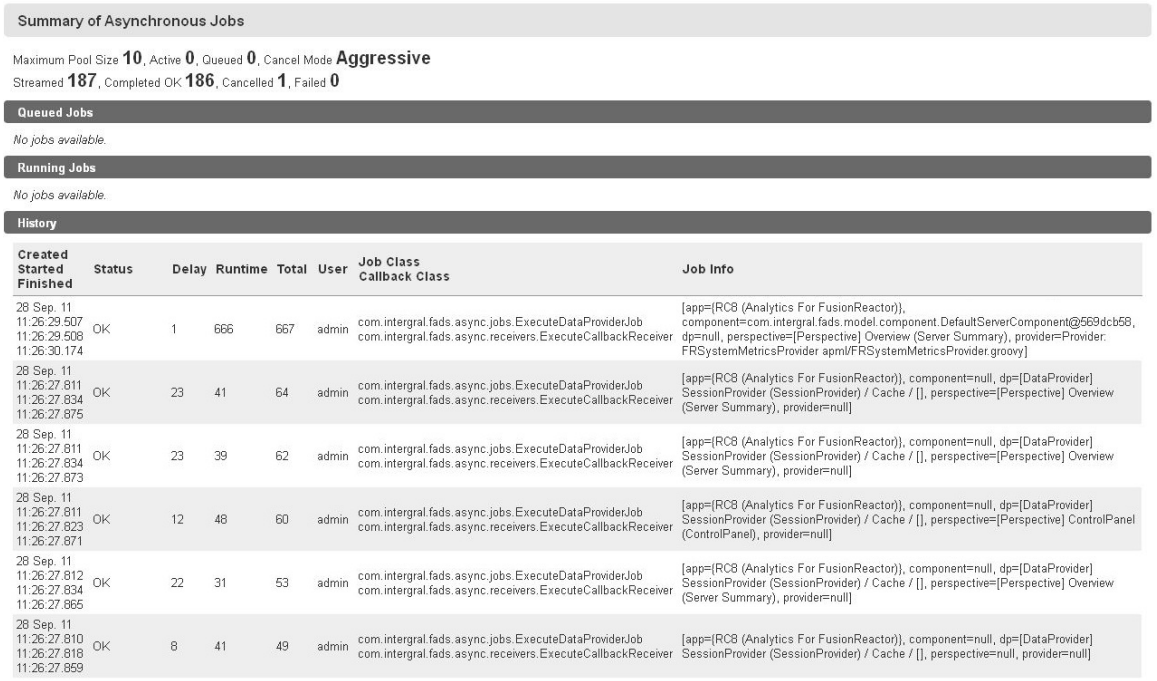

Summary of Asynchronous Jobs

Maximum Pool Size 10, Active 0, Queued 0, Cancel Mode Aggressive
Streamed 187, Completed OK 186, Cancelled 1, Failed 0

Queued Jobs

No jobs available.

Running Jobs

No jobs available.

History

Created Started Finished	Status	Delay	Runtime	Total	User	Job Class Callback Class	Job Info
28 Sep. 11 11:26:29.507 11:26:29.508 11:26:30.174	OK	1	666	667	admin	com.intergral.fads.async.jobs.ExecuteDataProviderJob com.intergral.fads.async.receivers.ExecuteCallbackReceiver	[app=[RC8 (Analytics For FusionReactor)], component=com.intergral.fads.model.component.DefaultServerComponent@569dcb58, dp=null, perspective=[Perspective] Overview (Server Summary), provider=Provider: FRSystemMetricsProvider aprnl/FRSystemMetricsProvider.groovy]
28 Sep. 11 11:26:27.811 11:26:27.834 11:26:27.875	OK	23	41	64	admin	com.intergral.fads.async.jobs.ExecuteDataProviderJob com.intergral.fads.async.receivers.ExecuteCallbackReceiver	[app=[RC8 (Analytics For FusionReactor)], component=null, dp=[DataProvider] SessionProvider (SessionProvider) / Cache / [], perspective=[Perspective] Overview (Server Summary), provider=null]
28 Sep. 11 11:26:27.811 11:26:27.834 11:26:27.873	OK	23	39	62	admin	com.intergral.fads.async.jobs.ExecuteDataProviderJob com.intergral.fads.async.receivers.ExecuteCallbackReceiver	[app=[RC8 (Analytics For FusionReactor)], component=null, dp=[DataProvider] SessionProvider (SessionProvider) / Cache / [], perspective=[Perspective] Overview (Server Summary), provider=null]
28 Sep. 11 11:26:27.811 11:26:27.823 11:26:27.871	OK	12	48	60	admin	com.intergral.fads.async.jobs.ExecuteDataProviderJob com.intergral.fads.async.receivers.ExecuteCallbackReceiver	[app=[RC8 (Analytics For FusionReactor)], component=null, dp=[DataProvider] SessionProvider (SessionProvider) / Cache / [], perspective=[Perspective] ControlPanel (ControlPanel), provider=null]
28 Sep. 11 11:26:27.812 11:26:27.834 11:26:27.865	OK	22	31	53	admin	com.intergral.fads.async.jobs.ExecuteDataProviderJob com.intergral.fads.async.receivers.ExecuteCallbackReceiver	[app=[RC8 (Analytics For FusionReactor)], component=null, dp=[DataProvider] SessionProvider (SessionProvider) / Cache / [], perspective=[Perspective] Overview (Server Summary), provider=null]
28 Sep. 11 11:26:27.810 11:26:27.818 11:26:27.859	OK	8	41	49	admin	com.intergral.fads.async.jobs.ExecuteDataProviderJob com.intergral.fads.async.receivers.ExecuteCallbackReceiver	[app=[RC8 (Analytics For FusionReactor)], component=null, dp=[DataProvider] SessionProvider (SessionProvider) / Cache / [], perspective=null, provider=null]

Figure 1: Summary of Asynchronous Jobs

Back to the top

Server Settings

Server Settings

Navigating to System -> Server Settings you will see the following: Under the heading "Server" you can configure the hostname from where the application is hosted along with the port. The https address of the server (host and port) can also be configured if secure connection urls are used. Under the heading "Email Server" you can configure where notification emails are sent from and to as well as adding a mail server to process these requests. (This is critical configuration for Daily Reports and T.A.P Reports)

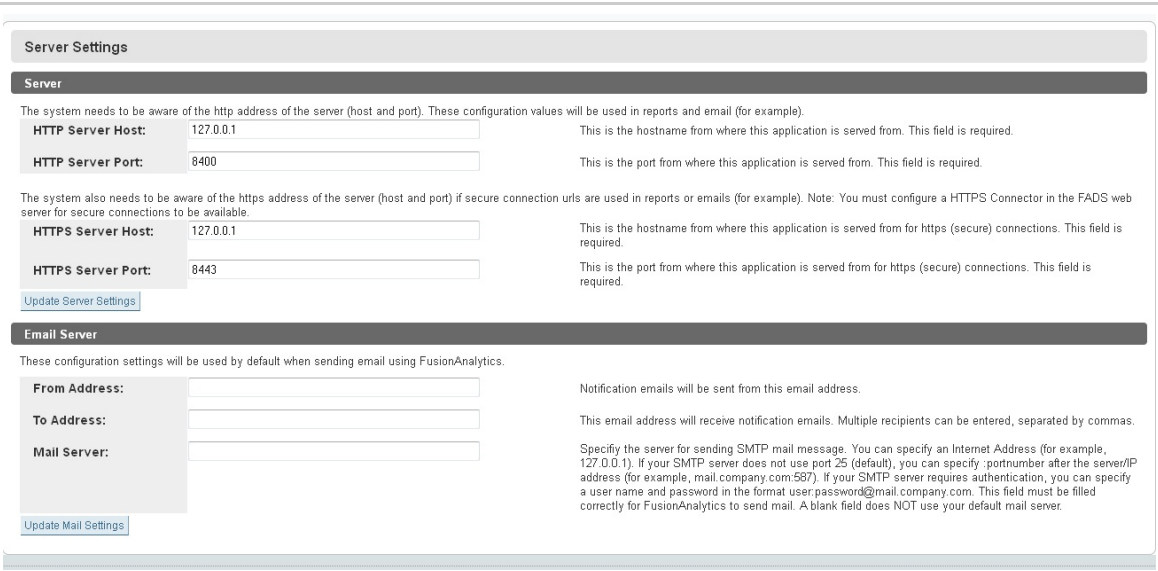

Server Settings

Server

The system needs to be aware of the http address of the server (host and port). These configuration values will be used in reports and email (for example).

HTTP Server Host: 127.0.0.1 — This is the hostname from where this application is served from. This field is required.

HTTP Server Port: 8400 — This is the port from where this application is served from. This field is required.

The system also needs to be aware of the https address of the server (host and port) if secure connection urls are used in reports or emails (for example). Note: You must configure a HTTPS Connector in the FADS web server for secure connections to be available.

HTTPS Server Host: 127.0.0.1 — This is the hostname from where this application is served from for https (secure) connections. This field is required.

HTTPS Server Port: 8443 — This is the port from where this application is served from for https (secure) connections. This field is required.

Update Server Settings

Email Server

These configuration settings will be used by default when sending email using FusionAnalytics.

From Address: — Notification emails will be sent from this email address.

To Address: — This email address will receive notification emails. Multiple recipients can be entered, separated by commas.

Mail Server: — Specifiy the server for sending SMTP mail message. You can specify an Internet Address (for example, 127.0.0.1). If your SMTP server does not use port 25 (default), you can specify :portnumber after the server/IP address (for example, mail.company.com:587). If your SMTP server requires authentication, you can specify a user name and password in the format user:password@mail.company.com. This field must be filled correctly for FusionAnalytics to send mail. A blank field does NOT use your default mail server.

Update Mail Settings

Figure 1: Server Settings*

More information about setting up HTTP can be found here: Setting Up FusionAnalytics with HTTPS.

Back to the top

Setting up HTTPS
To learn more about how to setup a secure (HTTPS) connection please see the Setting Up FusionAnalytics with HTTPS section.

License Info (FADS)

License Info (FADS)

The license info page (System > License Info) displays all of the information related to your FusionAnalytics license. For help installing and activating your license see Installing your license.

The following information is displayed for each license you have installed:

Name	Description	Example
Product Name	The name of your license.	FusionAnalytics for FusionReactor (Subscription for 1 year)
Product Family	The product family your license belongs to.	FA4FR
Expiry Date	The date your license will expire (not the same as activation expiry).	01-Jan-2012
License Status	The current status of your license.	Your license is valid and activated.
FusionAnalytics DataCollector	The number of datacollectors used and the maximum your license allows.	You are currently using 9 of 18 DataCollectors.
FusionAnalytics DataServices	The number of users and applications used and the maximum your license allows.	You are currently using 16 of 20 Applications and 31 of 40 Users.

You will also find buttons to activate or uninstall your license (depending on its current activation state) as well as the option to install a new license.

Back to the top

Logout (FADS)

Logging out of FADS

You can log out of FADS by either:

- clicking the logout link in the top right corner of FADS.
- navigating to System > Logout in the navigation menu.

Logging out will return you to the FusionAnalytics DataServices login screen.

Back to the top

FusionAnalytics Licensing

FusionAnalytics Licensing

How FusionAnalytics is Licensed

FusionAnalytics licenses come in individual packs (S,M,L,XL and XXL). The size of the pack determines how many users you can setup and how many CF instances you can analyze with FusionAnalytics.

Instances in relation to licensing

ColdFusion is installed on a J2EE server instance running within its own Java™ Virtual Machine (JVM). As its name implies, a virtual machine is a self-contained operating environment that behaves as if it were a separate computer, while actually sharing resources with other virtual machines on the same server. You might install another J2EE server instance on another, completely separate JVM, and install

ColdFusion there as well. These two ColdFusion servers would then operate completely isolated from each other---as if they were on separate physical servers. Each individual instance can be monitored (by FusionReactor) and will generate a set of log files corresponding to the CF instance being monitored.

License packs are measured in the number of instances which you would like to analyze using FusionAnalytics.

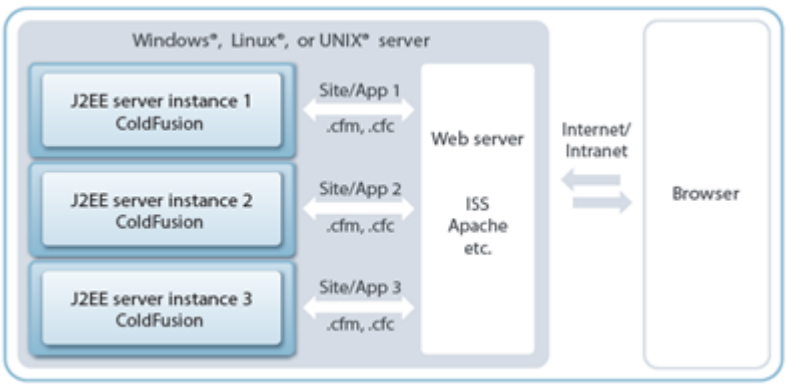

Figure 1: Instances

Take this example – here we have 3 instances of ColdFusion installed on one single server. If you wished to analyze the metric information from all three instances, then you would require:

1 x FusionReactor license (either FusionReactor Standard or Enterprise) – 1 license covers 1 physical server, no matter how many instances you have configured

1 x FusionAnalytics M Pack (containing 3 Data Collectors) – 1 Data Collector for each instance being analyzed

FusionAnalytics Users

You can setup individual users within FusionAnalytics – these may be system administrators, or regular users. Individual users receive their own login to FusionAnalytics and you can select specific users to receive any of the reports available, such as the TAP Reports or Daily Status Reports. Users are associated with applications and users can be given specific roles within FusionAnalytics e.g. Administrator.

FusionAnalytics DataCollectors

FusionAnalytics DataCollector's (FADC) are used to import log files into a database so that the resulting data can be analyzed and visualized by the FusionAnalytics application. You require 1 Data Collector for each instance which you would like to analyze.

FusionAnalytics DataCollectors in relation to licensing

FusionAnalytics DataCollector's (FADC) are used to collect data from a specific 'instance' which you would like to analyze with FusionAnalytics. Each instance is unique and you require 1 data collector for each instance you would like to analyze. So, if you have 3 instances to analyze, you need 3 data collectors etc.

I have more than 20 instances to analyze – what can I do?

If you have more than 20 instances, we can create you a custom license which perfectly matches your needs – please contact sales and let us know how many instances you have and we would be happy to send you a quote.

Back to the top

On this page:

FusionAnalytics for FusionReactor User's Guide

FusionAnalytics for FusionReactor User's Guide

Overview

FusionAnalytics for FusionReactor is the default App for FusionAnalytics at the time of launch. This App is designed to integrate with your existing FusionReactor installations (FusionReactor 4.0 or higher) in order to deliver a wide range of valuable analytics about your servers. This guide covers the various perspectives found in the FusionAnalytics for FusionReactor UI.

Next Steps

Using FusionAnalytics with FusionReactor

Application Perspectives (eg. faDataCollector-1)

FusionAnalytics Reporting

Using FusionAnalytics with FusionReactor

Using FusionAnalytics with FusionReactor

FusionReactor's Analytics Connector can be configured to automatically transfer logs to FusionAnalytics. This gives you a reliable, continuous analysis of FusionReactor's captured data. The **Analytics Connector** is available in all editions of FusionReactor 4.0.0 and above, and is found in the **Analytics** menu item.

We won't go into detail here, because FusionReactor's own documentation explains everything - feel free to follow the links to get more information on a specific subject. Here's the main entry on Analytics within FusionReactor.

- **Connector Status** shows you the current status of FusionReactor's archive (similar to the **Logging -> Log Archive** page) together with the status of the internal components of the connector:
    - The transfer status of each log archive, together with butts for (re)transfer either single log archives or the whole day's worth;
    - The total size of the archive system, so you can keep an eye on disk space;
    - The status of each Connector Worker thread - if it's idle or busy transferring logs to FusionAnalytics, and how much data remains to be sent;
    - The Connector Transfer Queue, showing all the log archives remaining to be sent. These will be processed by worker threads as they become available.
- **FusionAnalytics Settings** allows you to set the mode used (by default, FRAM - the FusionReactor Administration Manager - is used to transfer archives to FusionAnalytics, to alleviate load on the individual FusionReactor instances.
- **FusionAnalytics Targets** allows you to set up the DataCollector target for each FusionReactor instance. Each DataCollector within FusionAnalytics receives data from a single FusionReactor instance, and this is where you configure them.

Many of these settings can also be administered from FRAM, the FusionReactor Administration Manager, in the **Instance Manager**, by clicking the FusionAnalytics icon on a single instance. Popups help you set up and test the FusionAnalytics connector to make sure archives can be transferred.

Back to the top

Application Perspectives (eg. faDataCollector-1)

Application Perspectives (eg. faDataCollector-1)

Overview

This section will discuss the different perspectives that are used within the FusionAnalytics for FusionReactor Application.

At the heart of FusionAnalytics is the wealth of information gathered and transferred from FusionReactor. This instrumentation includes requests, SQL statements, response times, CPU and memory information, as well as URL parameters, Heap memory usage and JMX statistics. FusionAnalytics takes all of this information and builds a series of combined visual perspectives, allowing you to visualize all transactions, SQL queries, data transfer, server events etc. taking place over time, so you can build a complete historical view of what happened on your server. This data is essential to support investigative analysis, continuous server analysis and application improvement. With over 80 different analytic views and reports, FusionAnalytics gives you unlimited insight into exactly what's happening on your server.

The section will describe each Perspective in detail including information regarding what the Perspectives is showing (what data it is rendering) and how it can be used for analysis problem identification

Next Steps

Overview

Deep Server Analysis

- Data Finder
- Request
  - Slowest Requests
  - Slowest Requests on Average
  - Slowest DB Requests On Average
  - Longest Request Execution Time in Total
  - Highest DB Row Count by a Request
  - Highest DB Row Count On Average by a Request
  - Highest CPU Usage
  - Highest CPU Usage on Average
  - Highest CPU Usage in Total
  - Most Popular Requests
  - Most DB Queries On Average
  - Most DB Queries
  - Most Data Returned (bytes)
  - Most Data Returned on Average (bytes)
  - Table of Requests
  - Time Line of Requests
  - Time Line of Requests with Resources
  - Time Line of Requests with DB Requests
  - Bell-Curve of Request Performance
  - Request Status Code Brakedown
  - Request and JDBC Execution Time
  - Total Request Performance
- Database
  - Database Requests Grid
  - Highest DB Request Row Count
  - Slowest Database Requests
  - Database Usage Graphs
  - Requests and JDBC Execution Time
- Memory
  - Memory Usage
  - Memory Spaces
  - Memory Heap Non-Heap Summary
- CPU
  - CPU Usage
  - CPU Memory Usage
- Thread
  - Thread Activity
- General
  - FusionReactor System Metrics
  - Server Summary
  - Resource Usage Data Grid
- Logs
  - Log File Entries
- Events
  - Annotations
  - Placemarks
  - Crash Protection
  - Resources with Annotations
- Hot-Spots
  - Memory Hot-Spots
    - Memory at 100%
    - Memory Highest Usage
    - Memory Events
  - CPU Hot-Spots
    - Process CPU at 100%
    - Process CPU Highest Usage

Overview

Overview

Description

This is the default perspective that you will see when opening the FusionAnalytics Client. The top of the perspective contains a set of graphs showing various server statistics. Under those graphs are a set of other useful statistics.

Usage

Clicking on any of the graphs from the top section will increase the size of that graph.

Additionally, you can change the information shown from the bottom section by selecting the required button from the button bar:

Figure 1: Overview Perspective - Button Bar

Server Summary	A table view of all requests within the current time range, grouped by Server.
Performance	A table view of Requests, ordered by execution time and grouped by Request URL.
Alerts	A view showing application alerts. You can display this information as a graph or in a table.
I/O Performance	A table showing the amount of database activity grouped by Request URL. You can either view the Most Queries on Average, or Most Data Returned on Average.
Popular Requests	A pie chart showing the top 10 most popular requests.
CF General	This is a collection of various CF graphs. A second button bar will give you the option to select from Sessions, CFC Q, Thread Q, Throttle Q, WebService Q, Flash Q, Req Q, Template, Users, Temp Cache or Query Cache

Overview - Perspective View

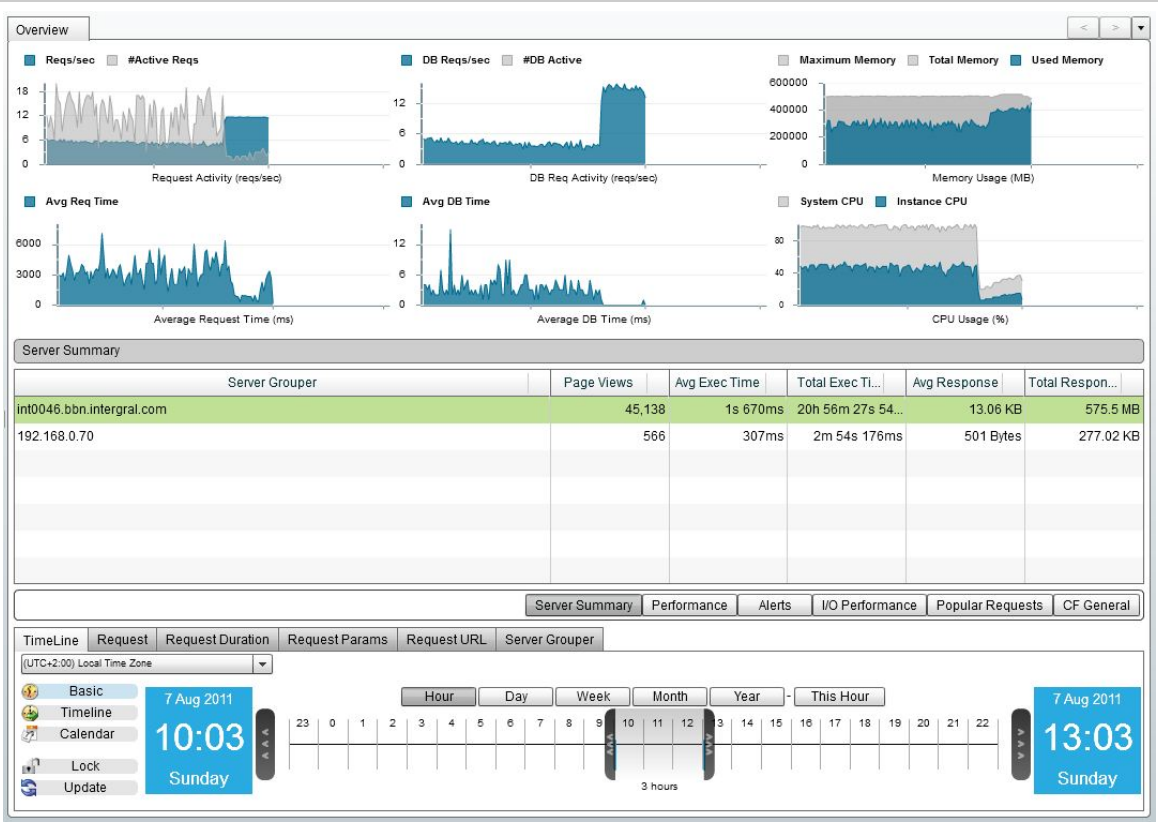

Figure 2: Overview Perspective

Back to the top

Deep Server Analysis

Deep Server Analysis

Overview

Next Steps

Data Finder

Request

- Slowest Requests
- Slowest Requests on Average
- Slowest DB Requests On Average
- Longest Request Execution Time in Total
- Highest DB Row Count by a Request
- Highest DB Row Count On Average by a Request
- Highest CPU Usage
- Highest CPU Usage on Average
- Highest CPU Usage in Total
- Most Popular Requests
- Most DB Queries On Average
- Most DB Queries
- Most Data Returned (bytes)
- Most Data Returned on Average (bytes)
- Table of Requests
- Time Line of Requests
- Time Line of Requests with Resources
- Time Line of Requests with DB Requests
- Bell-Curve of Request Performance
- Request Status Code Brakedown
- Request and JDBC Execution Time
- Total Request Performance

Database

- Database Requests Grid
- Highest DB Request Row Count
- Slowest Database Requests
- Database Usage Graphs
- Requests and JDBC Execution Time

Memory

- Memory Usage
- Memory Spaces
- Memory Heap Non-Heap Summary

CPU

- CPU Usage
- CPU Memory Usage

Thread

- Thread Activity

General

- FusionReactor System Metrics
- Server Summary
- Resource Usage Data Grid

Logs

- Log File Entries

Events

- Annotations
- Placemarks
- Crash Protection
- Resources with Annotations

Hot-Spots

- Memory Hot-Spots
  - Memory at 100%
  - Memory Highest Usage
  - Memory Events
- CPU Hot-Spots
  - Process CPU at 100%
  - Process CPU Highest Usage
  - System CPU at 100%
  - System CPU Highest Usage
- Request Hot-Spots
  - Requests Not Completed
  - Slowest Requests (Hot-Spots)
  - Slowest Requests On Average (Hot-Spots)
  - Slowest DB Requests On Average (Hot-Spots)
- Database Hot-Spots
  - Slowest Database Requests (Hot-Spots)

ColdFusion Monitor

- CF General Statistics
- CF Scope Sizes
- CF Hit Counts
- CF DB Pool Stats

Frameworks

Applications

Data Finder

Data Finder

Description

The Data Finder gives you an overview of all the types of data within FusionAnalytics and, once you have isolated a span of data you can switch the Data Finder into a different mode in order to further investigate that data.

Usage

Use the Date Navigation to move back and forward through the timespan. The Data Finder chart will shows blue bars to indicate when different types of information are available within the specified time span. There are various types of information which can be available, including JDBC, Resource, ColdFusion Monitor, Memory Spaces, Hit Counts, Logs and Request.

To switch the Data Finder to a different type of view, select the required view from the drop down box at the top right of the main graph area.

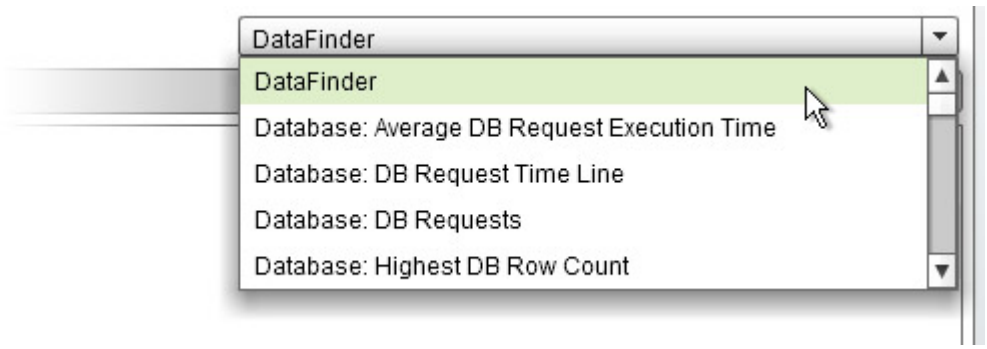

Figure 1: Data Finder Perspective Selector

The various DataFinder views are grouped by data type: Here is a list of the different perspective options you have to select from:

- DataFinder
- Database: Average DB Request Execution Time
- Database: DB Request Time Line
- Database: DB Requests
- Database:Highest DB Row Count
- Database: Most DB Queries by a Request On Average
- Database: Slowest DB Requests On Average by a Request
- Database: Slowest Database Requests
- Events: Annotations Time Line
- Events: Annotations
- Events: Crash Protection Events
- Events: Placemarks
- General: Server Summary
- Log: Log Entries
- Request: Average Request Execution Time
- Request: Highest Average DB Row Count by a Request
- Request: Highest Request DB Row Count
- Request: Most DB Queries by a Request
- Request: Most Data Returned on Average
- Request: Most Popular Requests
- Request: Request Cumulative Execution Time
- Request: Request Performance Bell-Curve
- Request: Request Time Line
- Request: Request Not Completed
- Request: Requests
- Request: Slowest Requests on Average
- Request: Slowest Requests
- Request: Most Data Returned
- Resource: CPU and Memory Usage
- Resource: FusionReactor System Metrics
- Resource: Highest Memory Usage Data Grid
- Resource: Highest Process CPU Usage
- Resource: Highest System CPU Usage
- Resource: Memory at 100% Data Grid
- Resource: Memory Code-Cache Usage
- Resource: Memory Eden-Space Usage
- Resource: Memory Perm-Gen Usage
- Resource: Memory Survivor-Space Usage
- Resource: Memory Tenured-Gen Usage
- Resource: Memory PS-Eden-Space Usage
- Resource: Memory PS-Old-Gen Usage
- Resource: memory PS-Perm-Gen Usage
- Resource: PS-Survivor-Space Usage
- Resource: Process CPU at 100%
- Resource: Process Memory Usage
- Resource: Process and System CPU Usage

- Resource: Resource Usage Data Grid
- Resource: System CPU at 100%
- Resource: Thread Activity Time Line
- CF: CF Active Sessions
- CF: CF Scope Sizes
- CF: CF CFC Queue
- CF: CF CFThread Queue
- CF: CF Concurrent Users
- CF: CF Flash Remoting Queue
- CF: CF Query Cache
- CF: CF Request Queue
- CF: CF Template Cache
- CF: CF Template Queue
- CF: CF Throttle Queue Size
- CF: CF Request Queue
- CF: CF Template Cache
- CF: CF Template Queue
- CF: CF Throttle Queue Size
- CF: CF Web Services Queue

Perspective View

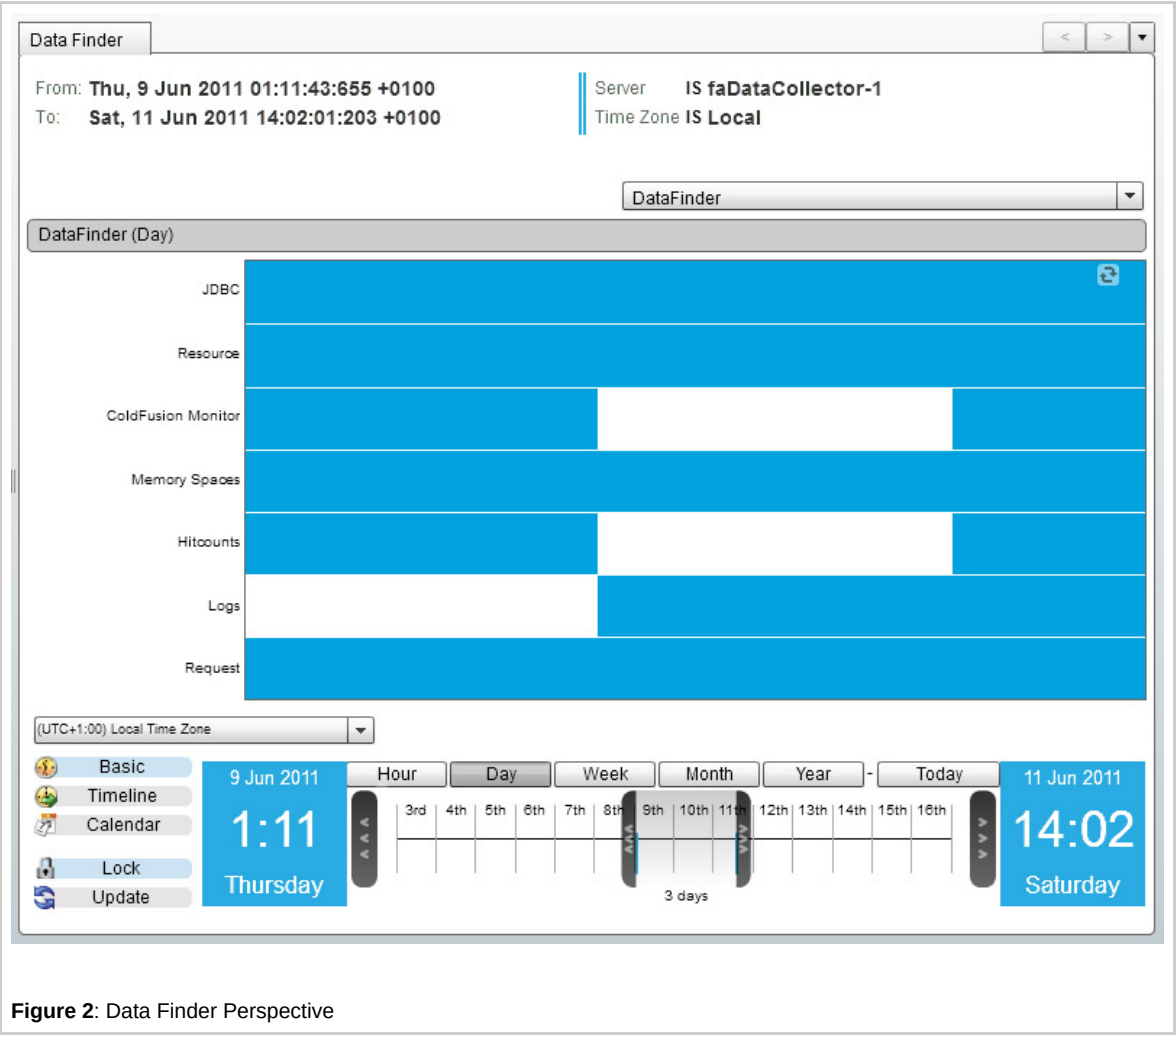

Figure 2: Data Finder Perspective

Back to the top

Request

Request

Overview

Next Steps

Slowest Requests

Slowest Requests

Slowest Requests

Description

This perspective shows a table of the top 100 individual requests by execution time within the specific period of time.

Usage

To change the order in which the requests are shows, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

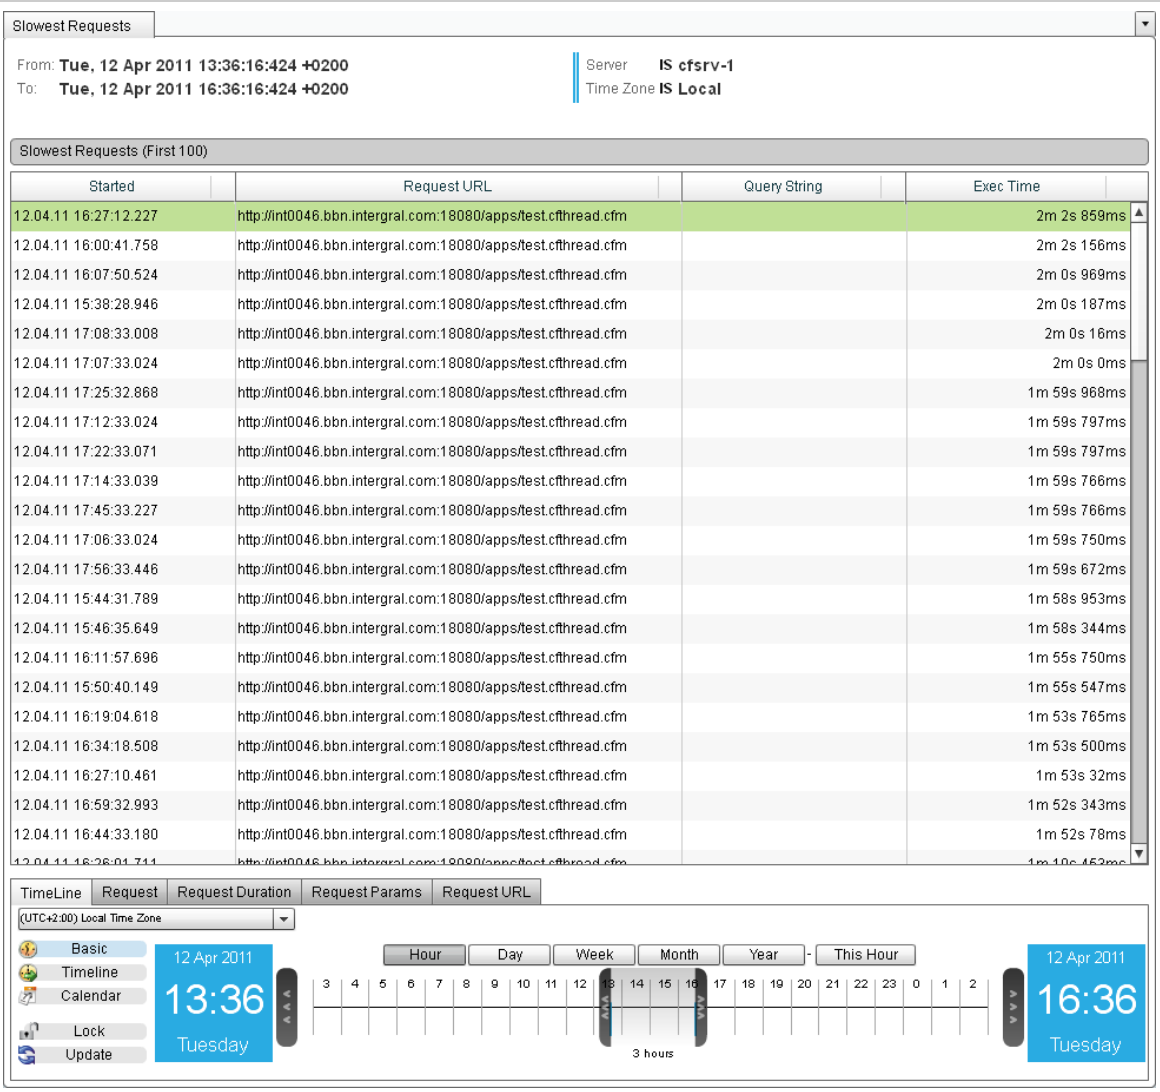

Figure 1: Slowest Requests Perspective

Back to the top

Slowest Requests on Average

Slowest Requests on Average

Description

This perspective shows a table of the top 100 requests by their average execution time within the specific period of time.

Usage

To change the order in which the requests are shows, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

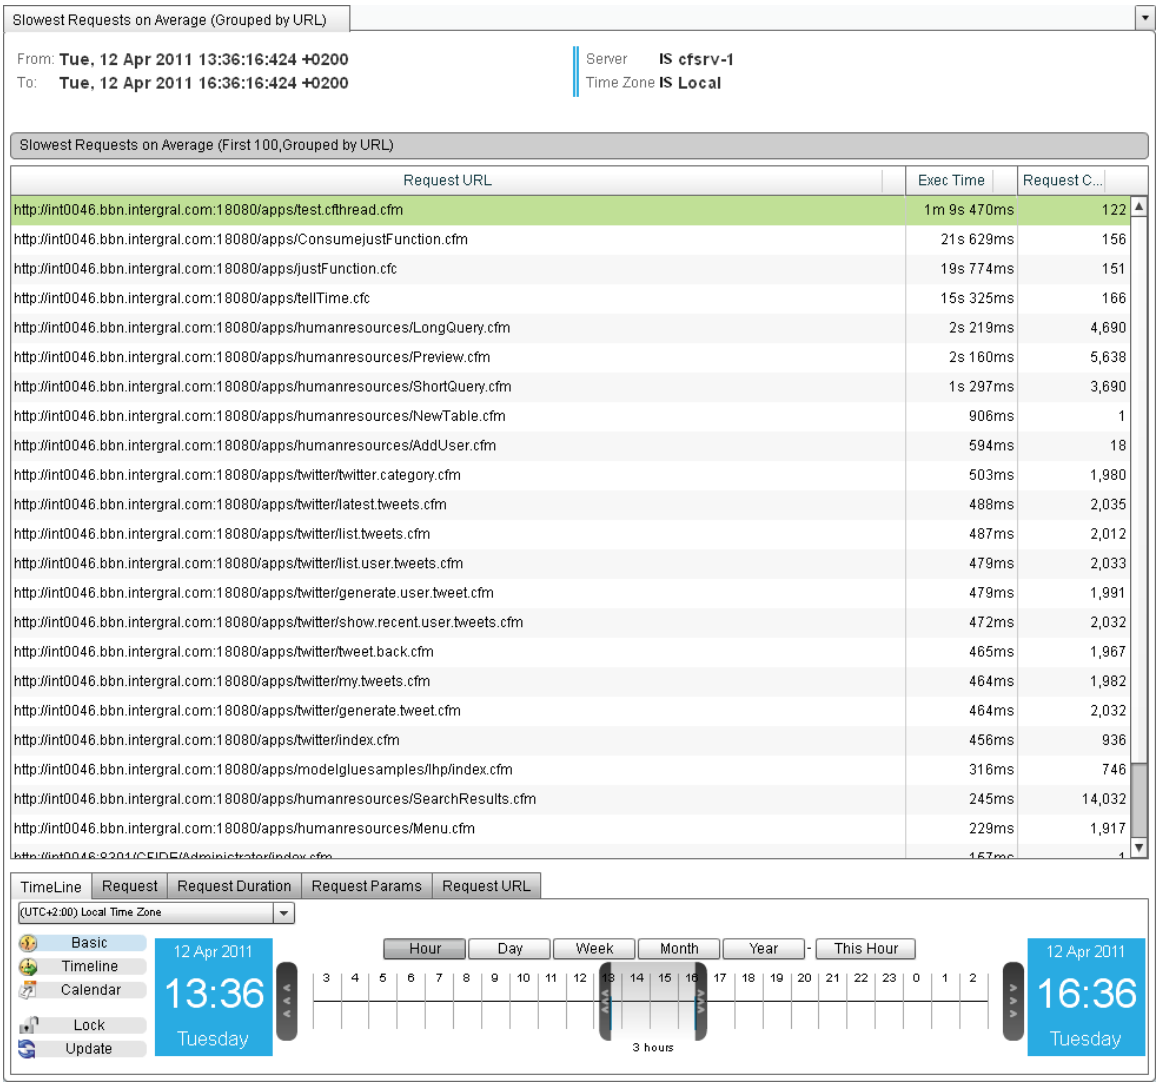

Figure 1: Slowest Requests on Average Perspective

Back to the top

Slowest DB Requests On Average

Slowest DB Requests On Average

Description

This perspective shows a table of the top 100 requests by their average database execution time within the specific period of time.

Usage

To change the order in which the requests are shows, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

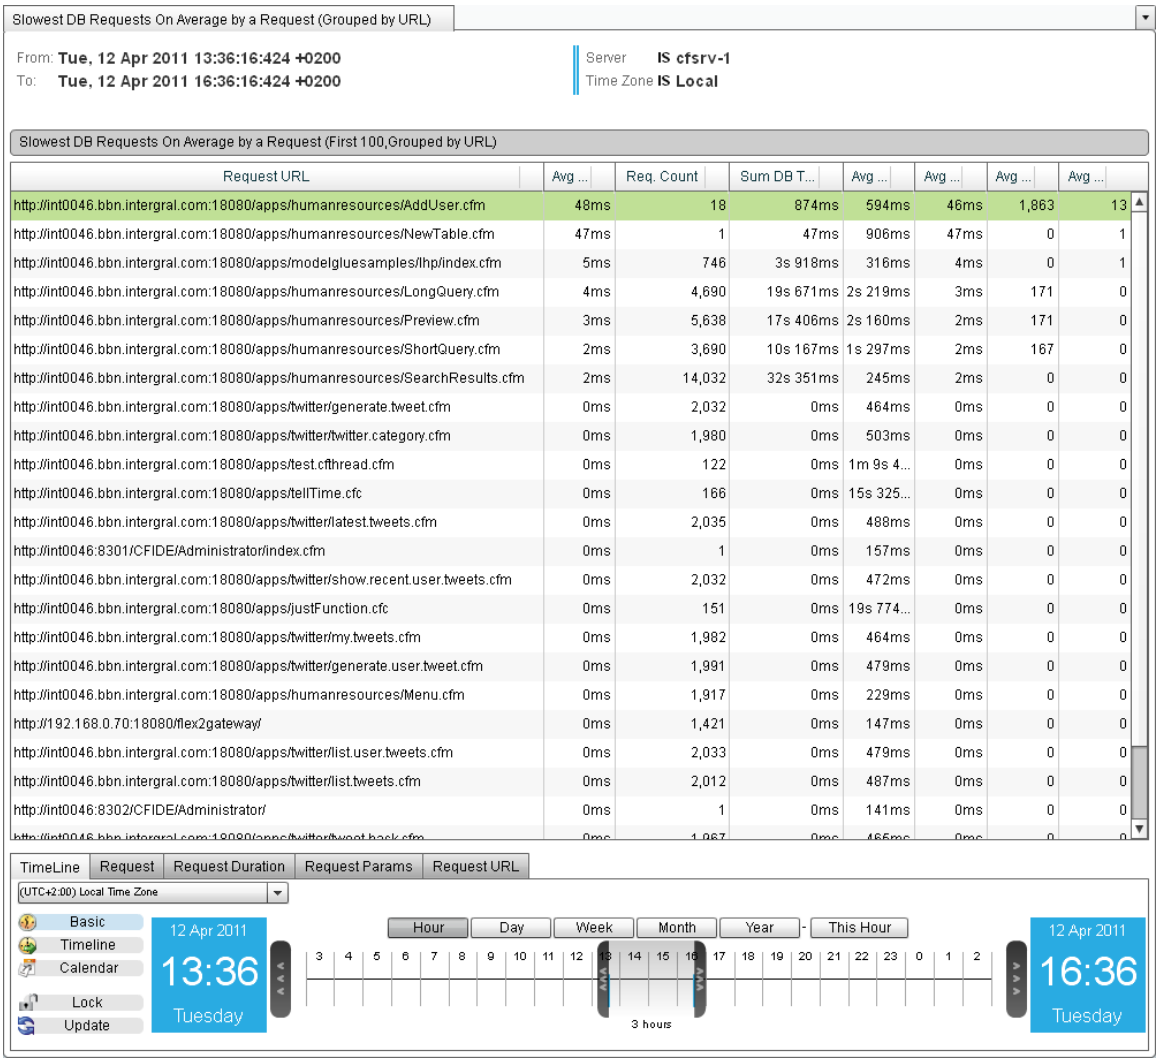

Figure 1: Slowest DB Requests On Average Perspective

Back to the top

Longest Request Execution Time in Total

Longest Request Execution Time in Total

Description

This perspective shows a table of the top 100 requests by their total execution time within the specific period of time.

Usage

To change the order in which the requests are shows, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

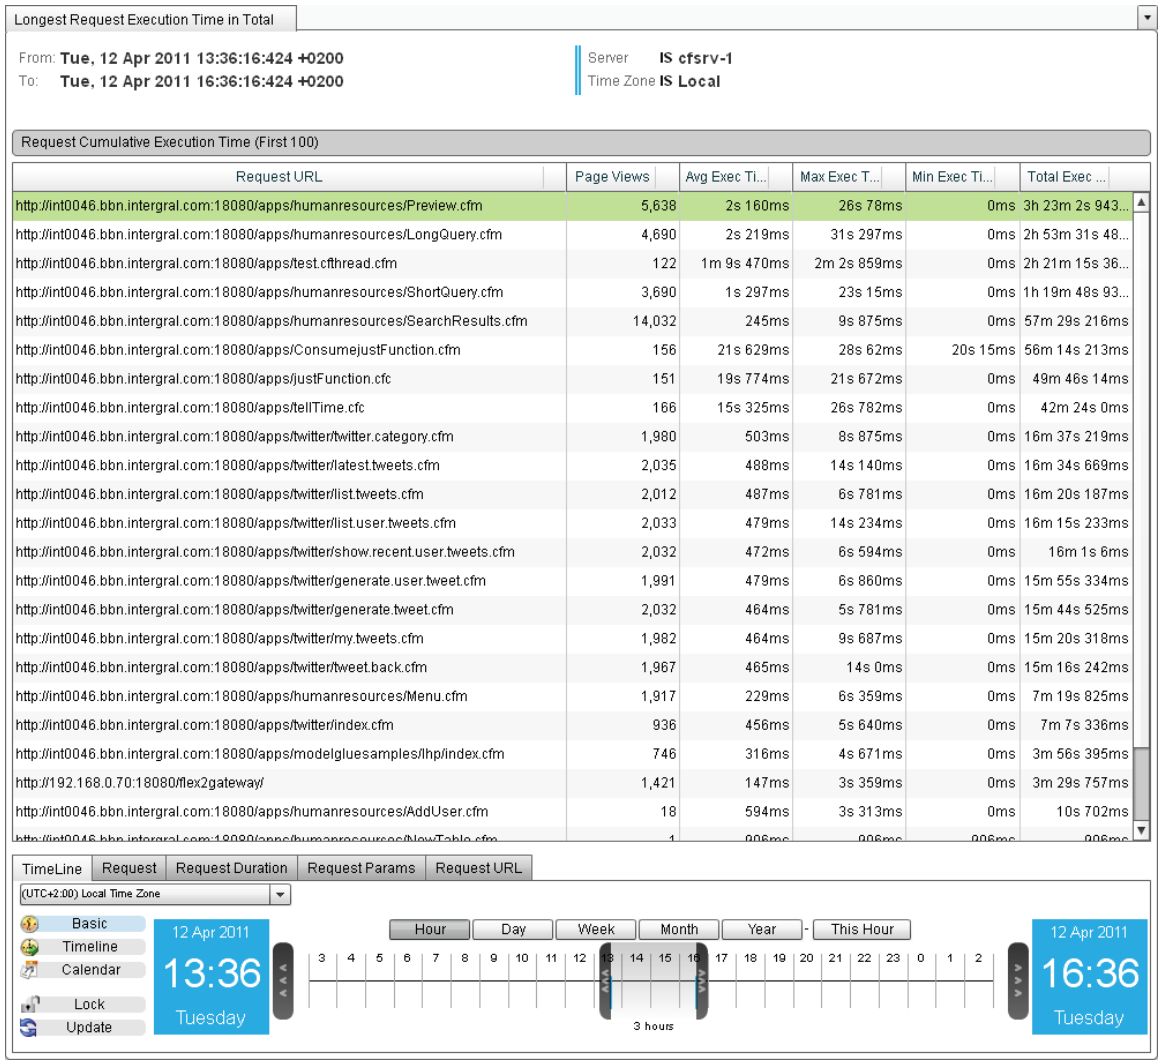

Figure 1: Longest Request Execution Time in Total Perspective

Back to the top

Highest DB Row Count by a Request

Highest DB Row Count by a Request

Description

Highest Total DB Row Count is a perspective to help you identify which request fetched the most rows of data from the database. The report calculates the number of rows fetched from the database by a request. Knowing the number of rows returned from the database can help you to identify many potential issues and give you insight into the resource usage and scale of your applications. Generally fetching rows from the database takes time, resource and can induce network overhead if the database is located on a remote server. Often rows fetched from the database are stored in memory and this can increase the overall memory usage of your application or potentially create spikes in memory usage.

Using this report you should look to fix, tune or optimize the requests that return most rows from the database. This will typically increase application performance, lower resource usage including memory, cpu and network traffic.

Usage

- Discover which requests fetch the most rows of data from the database(s)
- Identify unexpected results - requests with poorly written queries
- Use the information to find bugs in pages and/or queries which are causing too many rows to be returned
- Use the information to optimize code
- Identify changes in application behavior across time

Perspective View

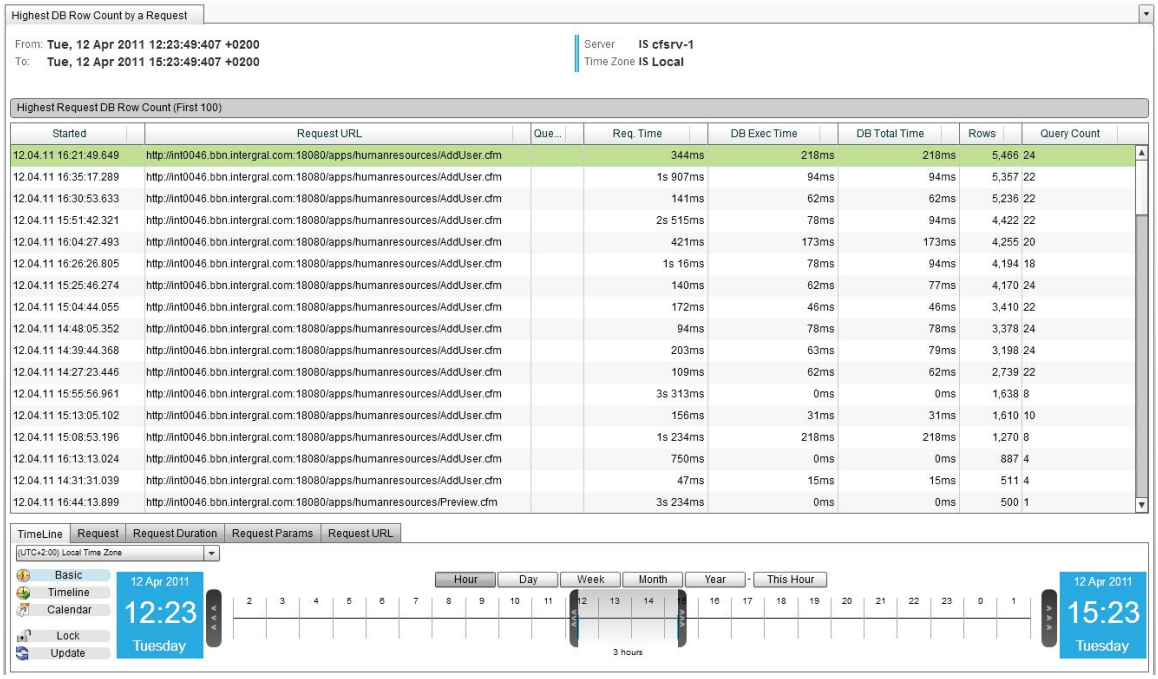

Figure 1: Highest DB Row Count by a Request Perspective

Back to the top

Highest DB Row Count On Average by a Request

Highest DB Row Count On Average by a Request

Description

Highest Total DB Row Count On Average is a report to help you indentify which requests fetch the most rows of data from the database. The report calculates the average number of rows fetched from the database per request run grouped them by their URL. Knowing the number of rows returned from the database can help you to indentify many potential issues and give you insight into the resource usage and scale of your applications. Generally fetching rows from the database takes time, resource and can induce network overhead if the database is located on a remote server. Often rows fetched from the database are stored in memory and this can increase the overall memory usage of your application or potentially create spikes in memory usage.

Using this report you should look to fix, tune or optimize the requests that return most rows from the database. This will typically increase application performance, lower resource usage including memory, cpu and network traffic.

Usage

- Discover which requests fetch the most rows of data from the database(s)
- Indentify unexpected results - requests with poorly written queries
- Use the information to find bugs in pages and/or queries which are causing too many rows to be returned
- Use the information to optimize code
- Indentify changes in application behavior across time

Perspective View

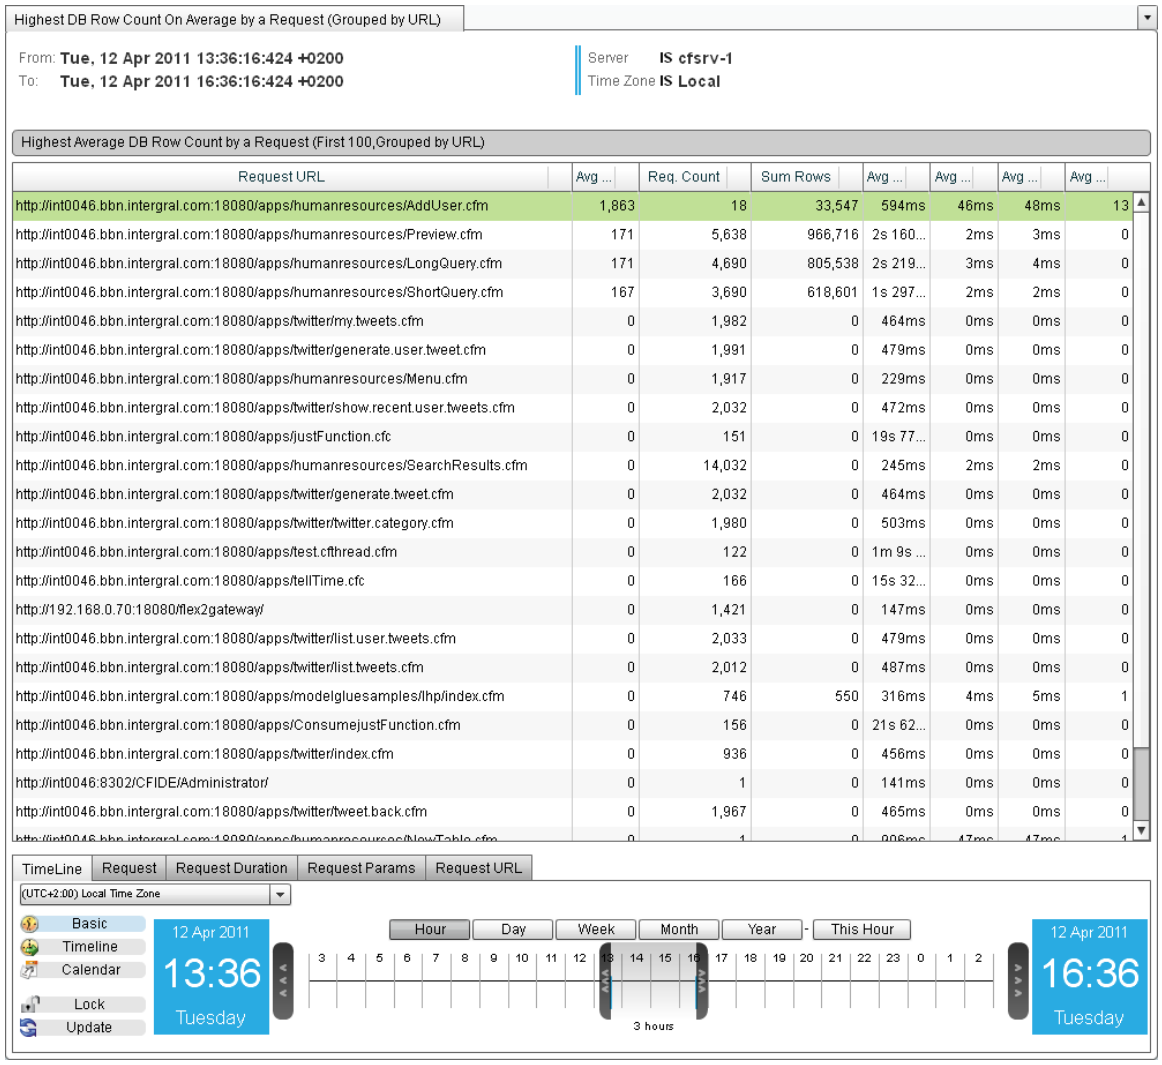

Figure 1: Highest Total DB Row Count On Average Perspective

Back to the top

Highest CPU Usage

Highest CPU Usage

Description

Usage

Back to the top

Highest CPU Usage on Average

Highest CPU Usage on Average

Description

Usage

Back to the top

Highest CPU Usage in Total

Highest CPU Usage in Total

Description

Usage

Back to the top

Most Popular Requests

Most Popular Requests

Description

Usage

Perspective View

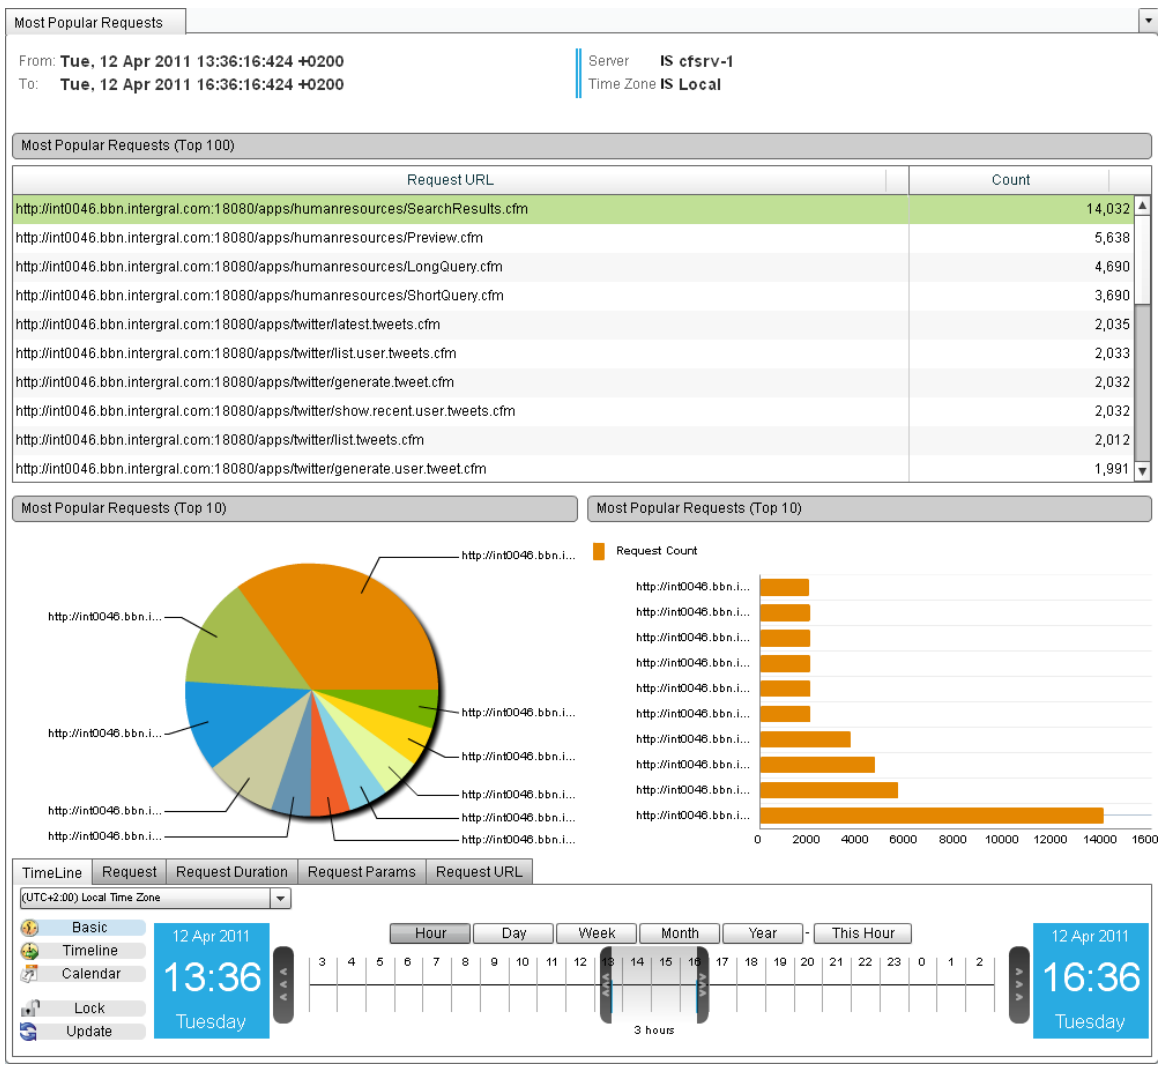

Figure 1: Most Popular Requests Perspective

Back to the top

Most DB Queries On Average

Most DB Queries On Average

Description

This perspective shows a table of the top 100 Most Database Requests on Average (Grouped by URL) within the specific period of time.

Usage

To change the order in which the requests are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

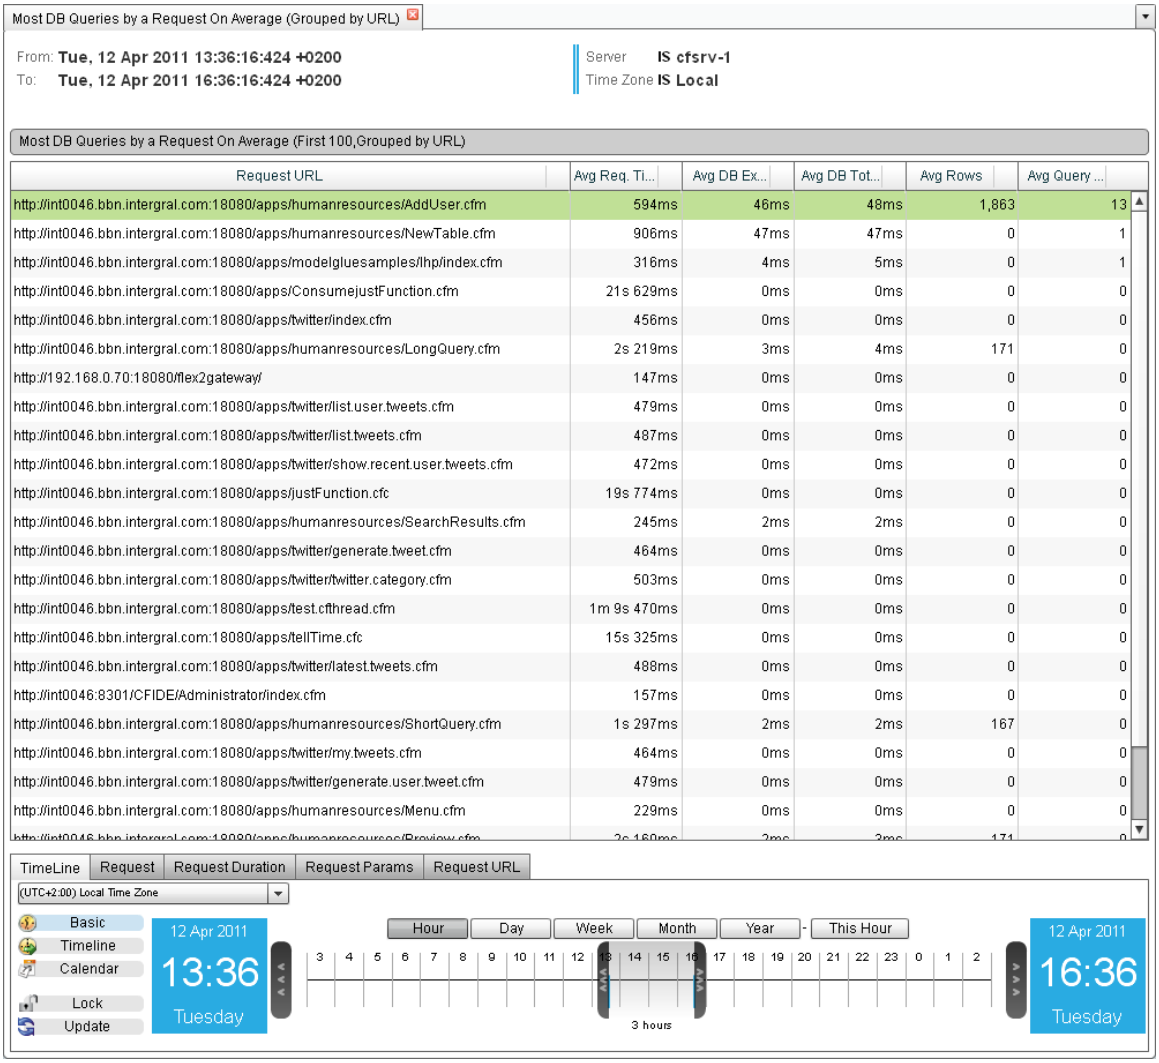

Figure 1: Most DB Queries On Average Perspective

Back to the top

Most DB Queries

Most DB Queries

Description

This perspective shows a table of the top 100 Most Database Queries by a Request within the specific period of time.

Usage

To change the order in which the requests are shows, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

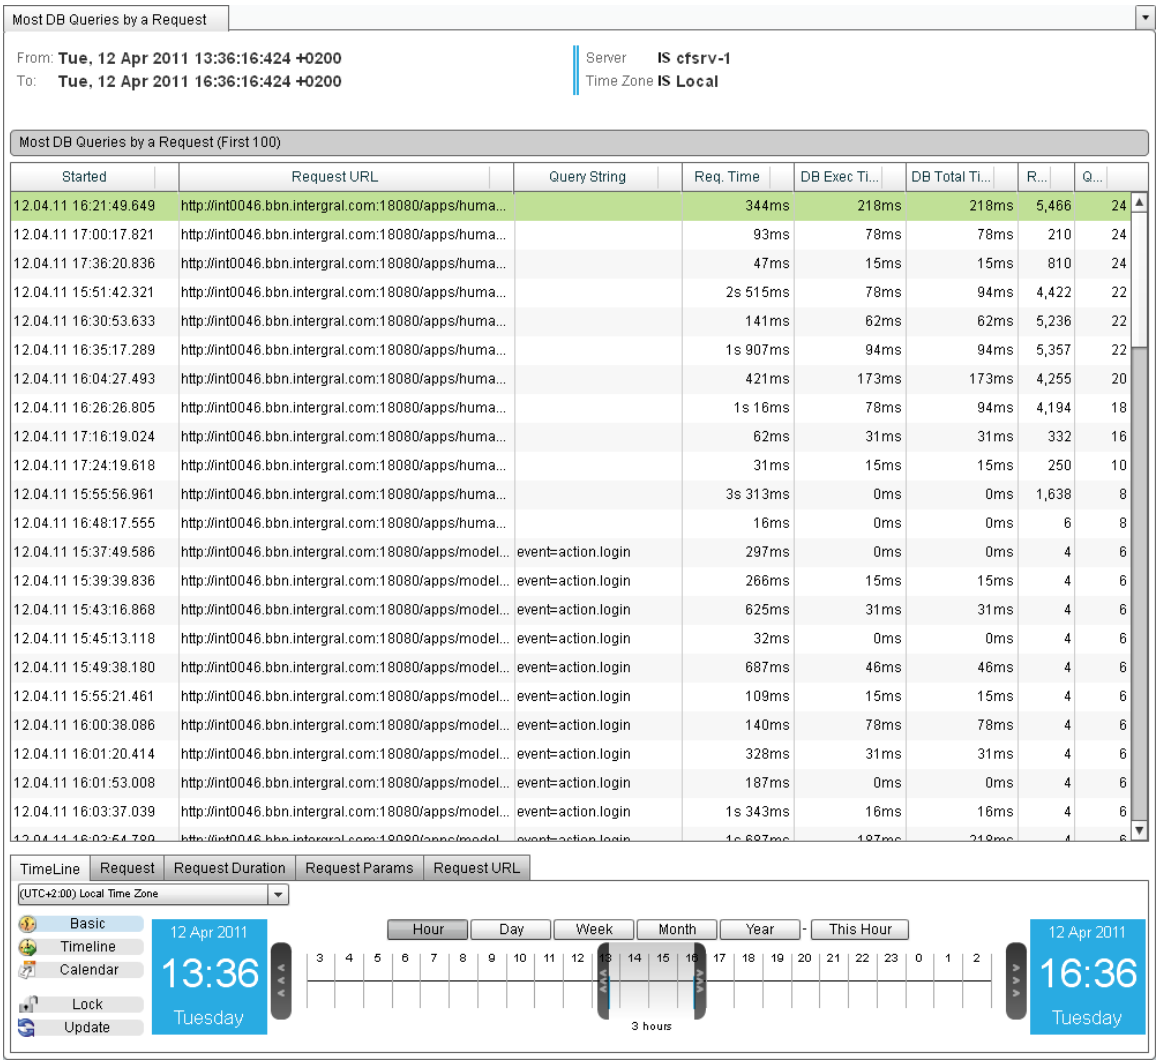

Figure 1: Most DB Queries Perspective

Back to the top

Most Data Returned (bytes)

Most Data Returned (bytes)

Description

Largest Size (bytes) is a report to help you indentify which request returned the most bytes in it's response to the client. Generally the more data returned to the client the more resource is used. Often content (bytes) to be returned are buffered or stored in memory until a request is complete, which can cause large quantities of memory to be used. Many applications servers will flush data to the client each time the buffered response reaches a certain size (configurable in many cases) but often the request processing will be halted until the client consumes the response being sent. This can cause performance degradation in requests and thread starvation. It can also used to a create Denial of Service (DoS) attacks by the client not consuming or only slowly consuming the data, causing requests to backup on the server waiting for a thread to become available.

Using this report you should look to fix, tune or optimize the requests that return the large quantities data. This will typically increase application performance, lower resource usage including memory, cpu and network traffic.

Usage

- Discover which requests return the most bytes to the client
- Identify unexpected results - requests with poorly code
- Use the information to find bugs in requests that cause to much data to be returned
- Use the information to optimize code
- Identify spikes in the application / request behavior (URL parameters that cause issues?)

- Identify a need to compress requests

Perspective View

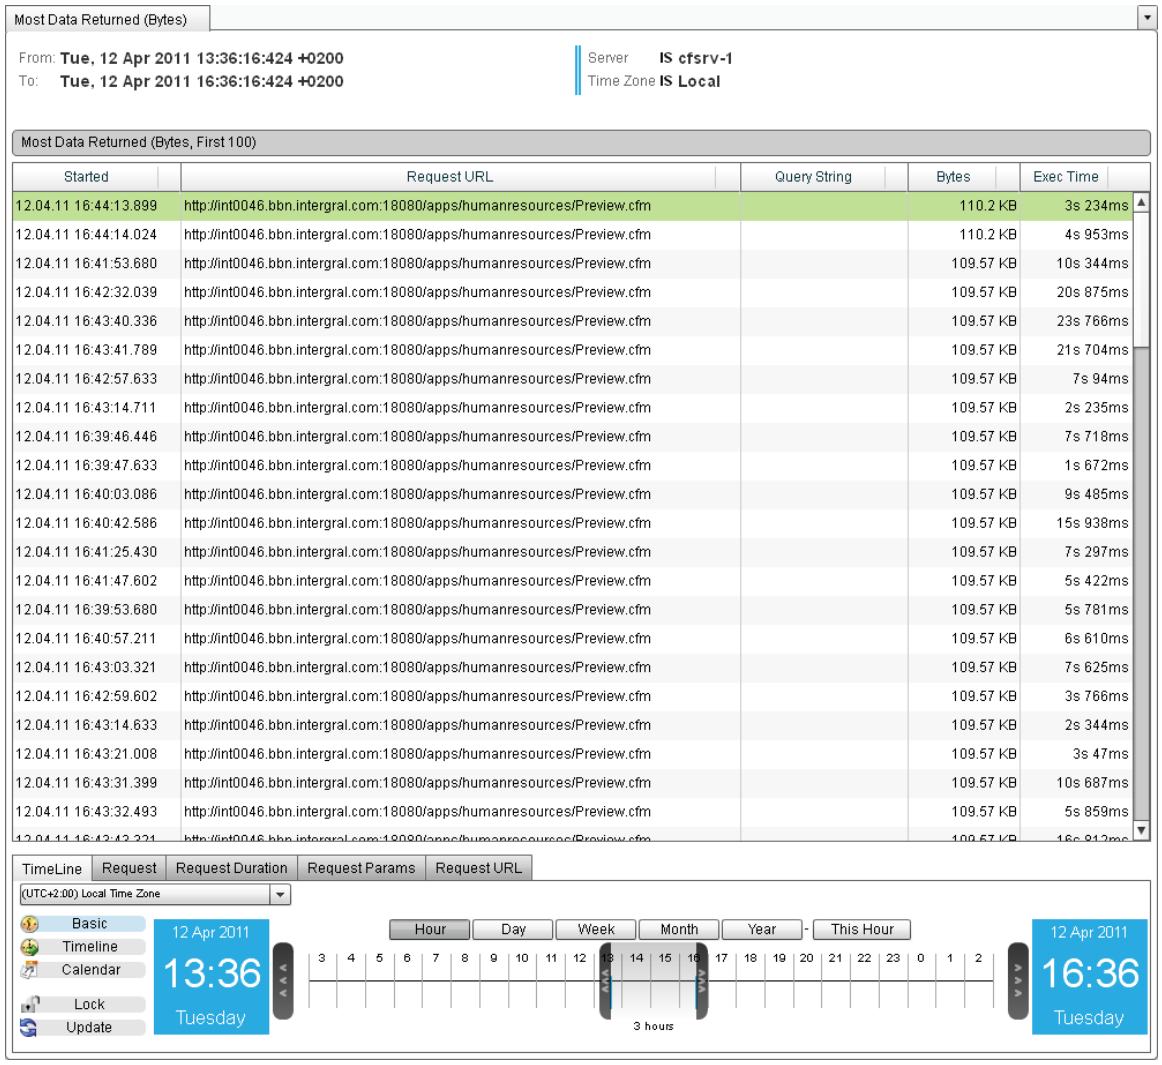

Figure 1: Most Data Returned (bytes) Perspective

Back to the top

Most Data Returned on Average (bytes)

Most Data Returned on Average (bytes)

Description

This perspective shows a table of the top 100 Most Data Returned on Average (Bytes, Grouped by URL) within the specific period of time.

Usage

To change the order in which the results are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

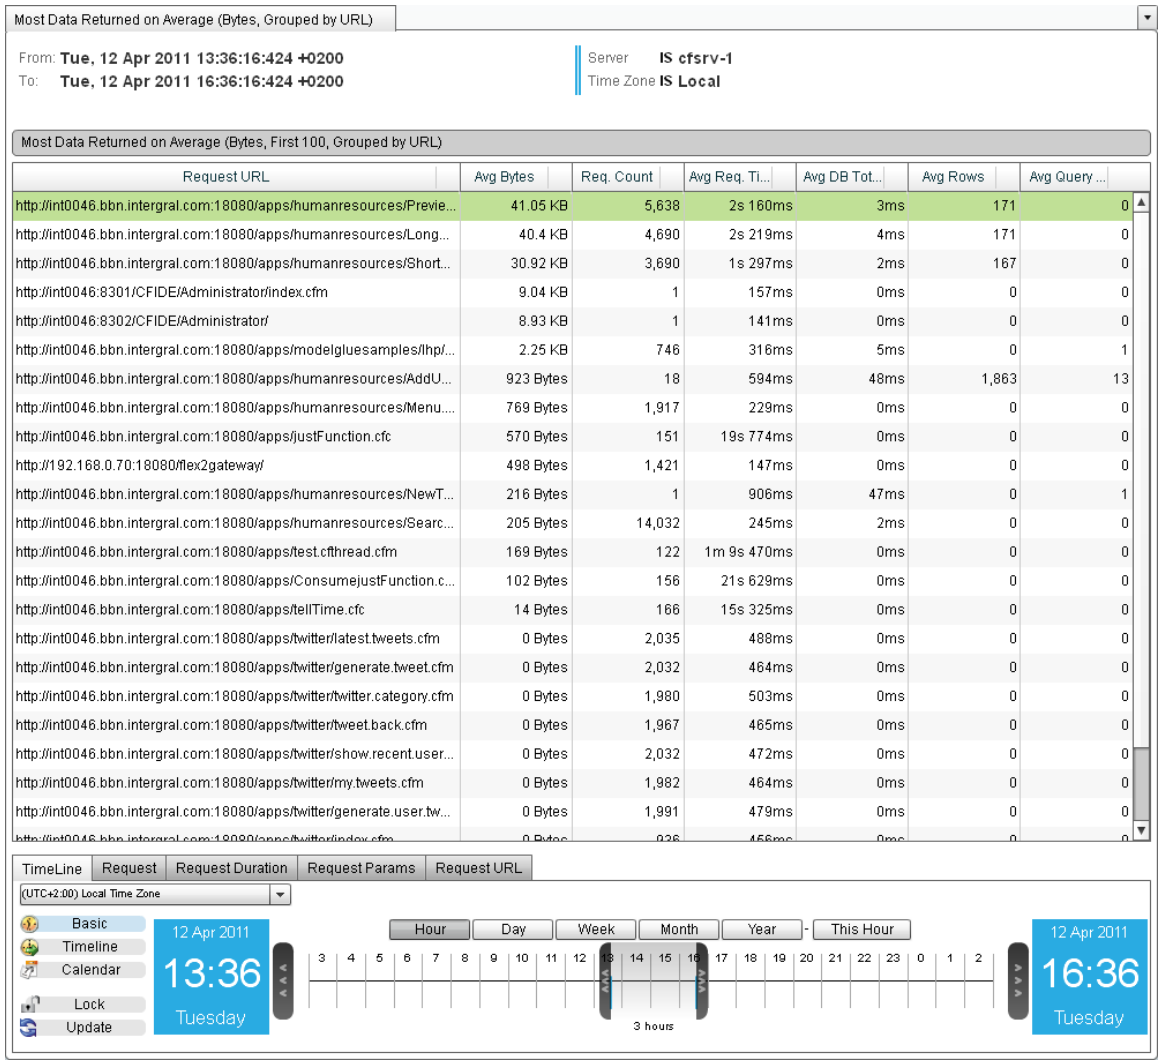

Figure 1: Most Data Returned on Average Perspective

Back to the top

Table of Requests

Table of Requests

Description

This perspective shows a table of the top 1000 Requests within the specific period of time.

Usage

To change the order in which the requests are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

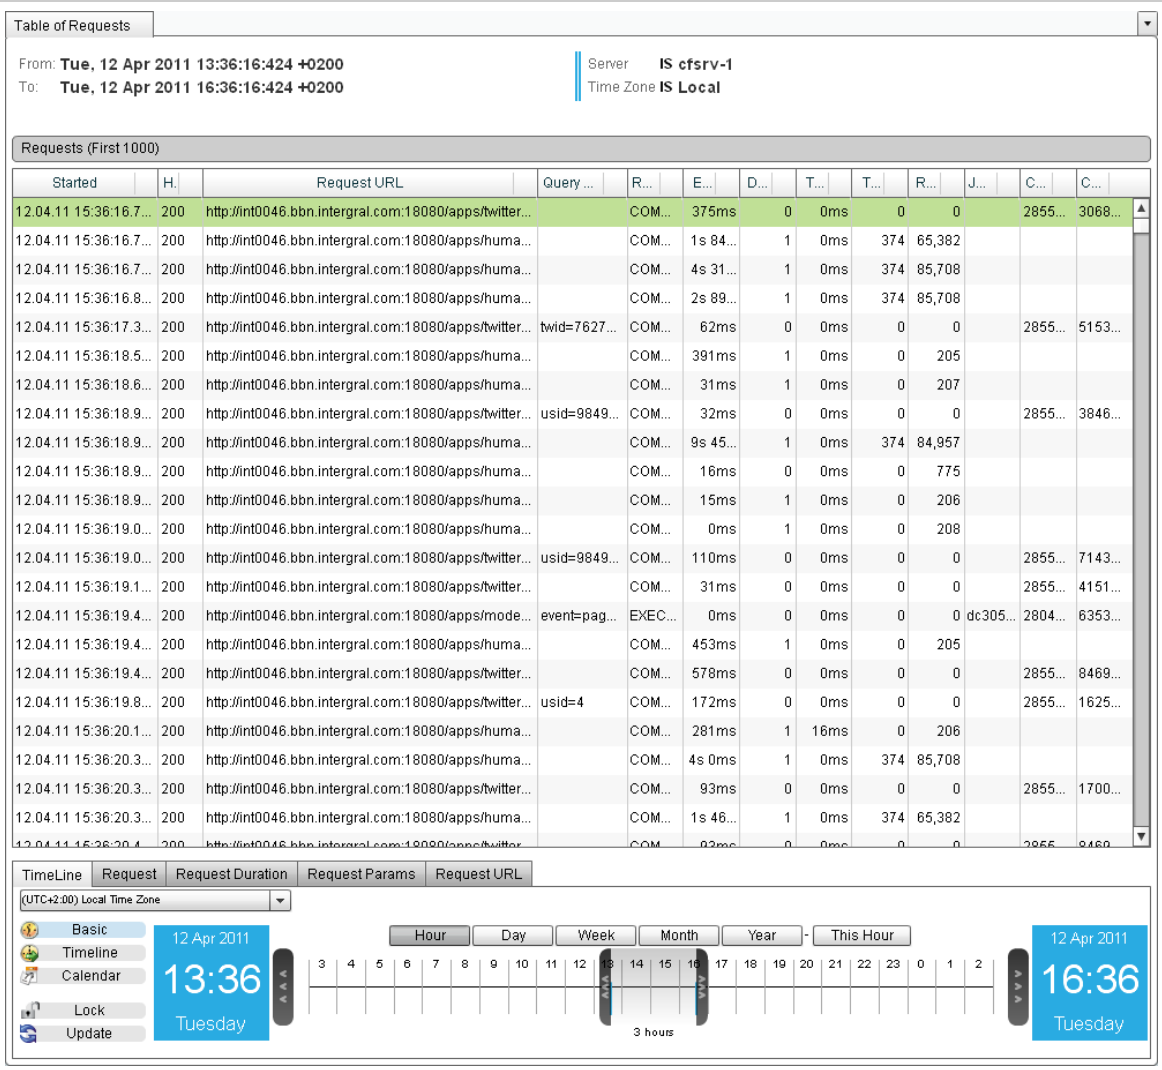

Figure 1: Table of Requests Perspective

Back to the top

Time Line of Requests

Time Line of Requests

Description

This perspective shows a Time Block Chart of Requests within the specific period of time.

Usage

Perspective View

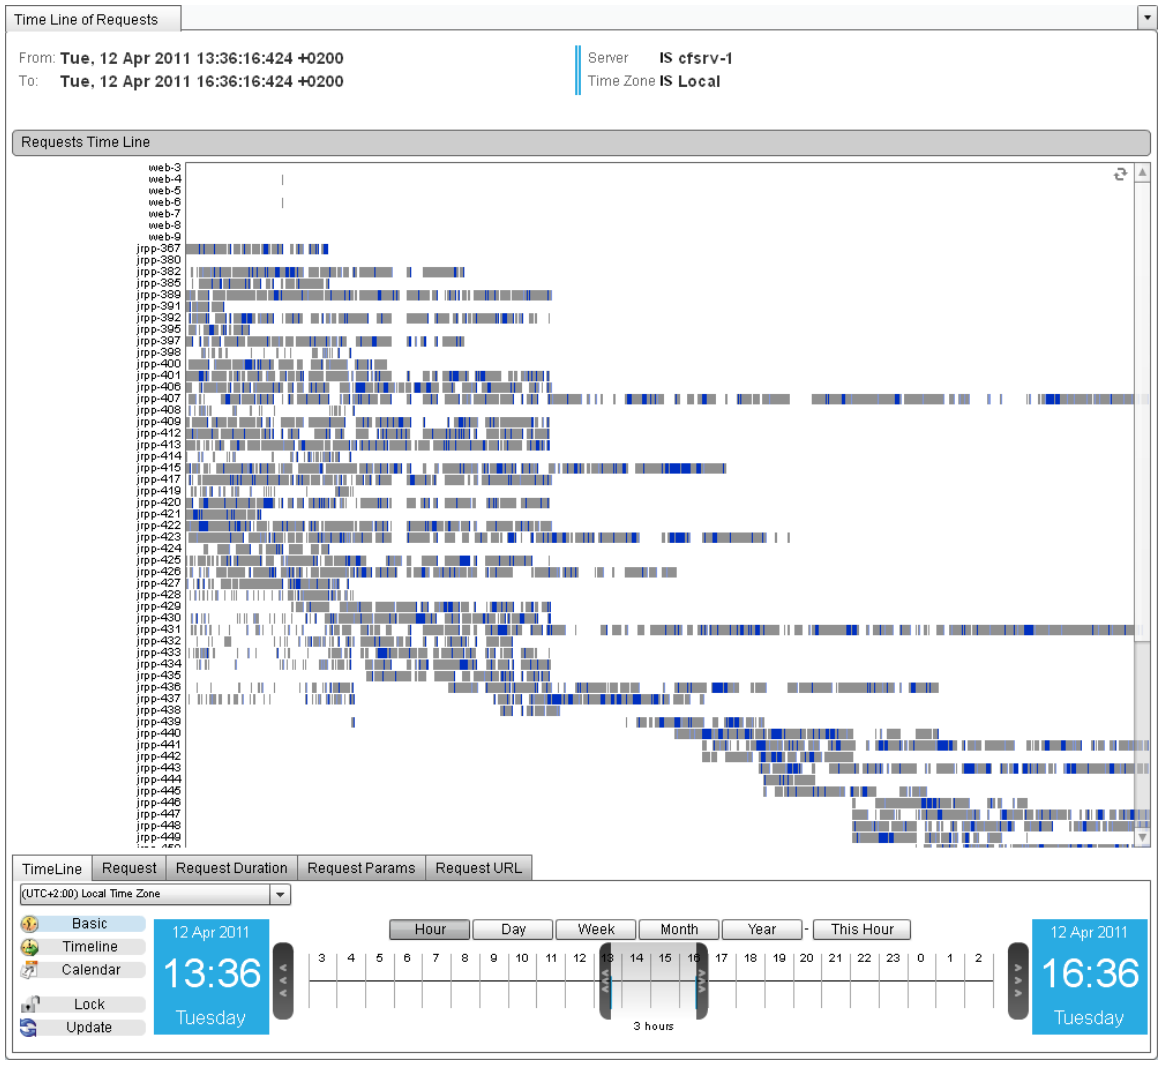

Figure 1: Time Line of Requests Perspective

Back to the top

Time Line of Requests with Resources

Time Line of Requests with Resources

Description

Usage

Perspective View

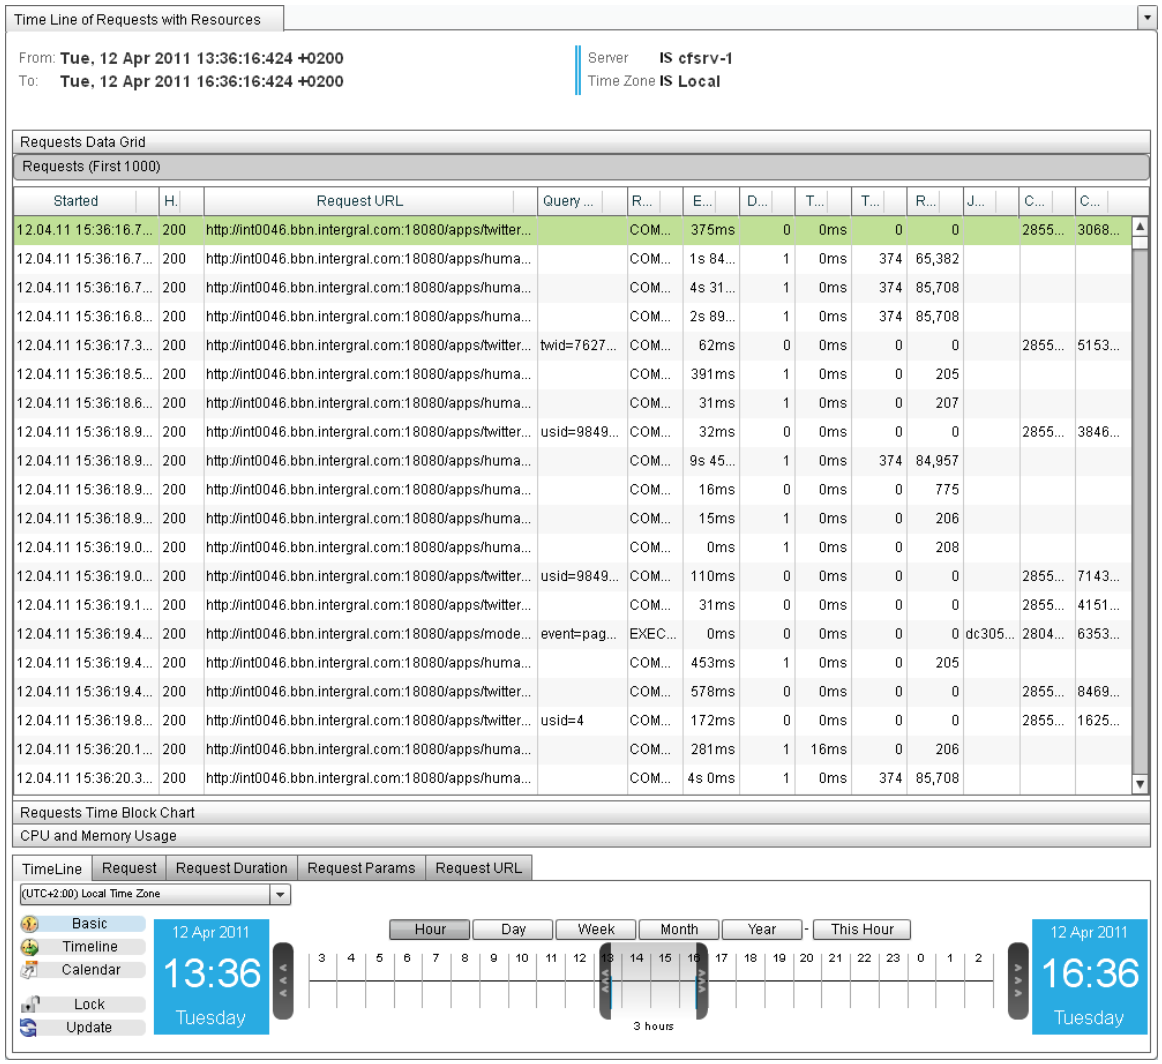

Figure 1: Time Line of Requests with Resources Perspective

Back to the top

Time Line of Requests with DB Requests

Time Line of Requests with DB Requests

Description

Usage

Perspective View

Figure 1: Time Line of Requests with DB Requests Perspective

Back to the top

Bell-Curve of Request Performance

Bell-Curve of Request Performance

Description

This perspective shows a Bell Curve Chart showing Request Performance within the specific period of time.

Usage

Perspective View

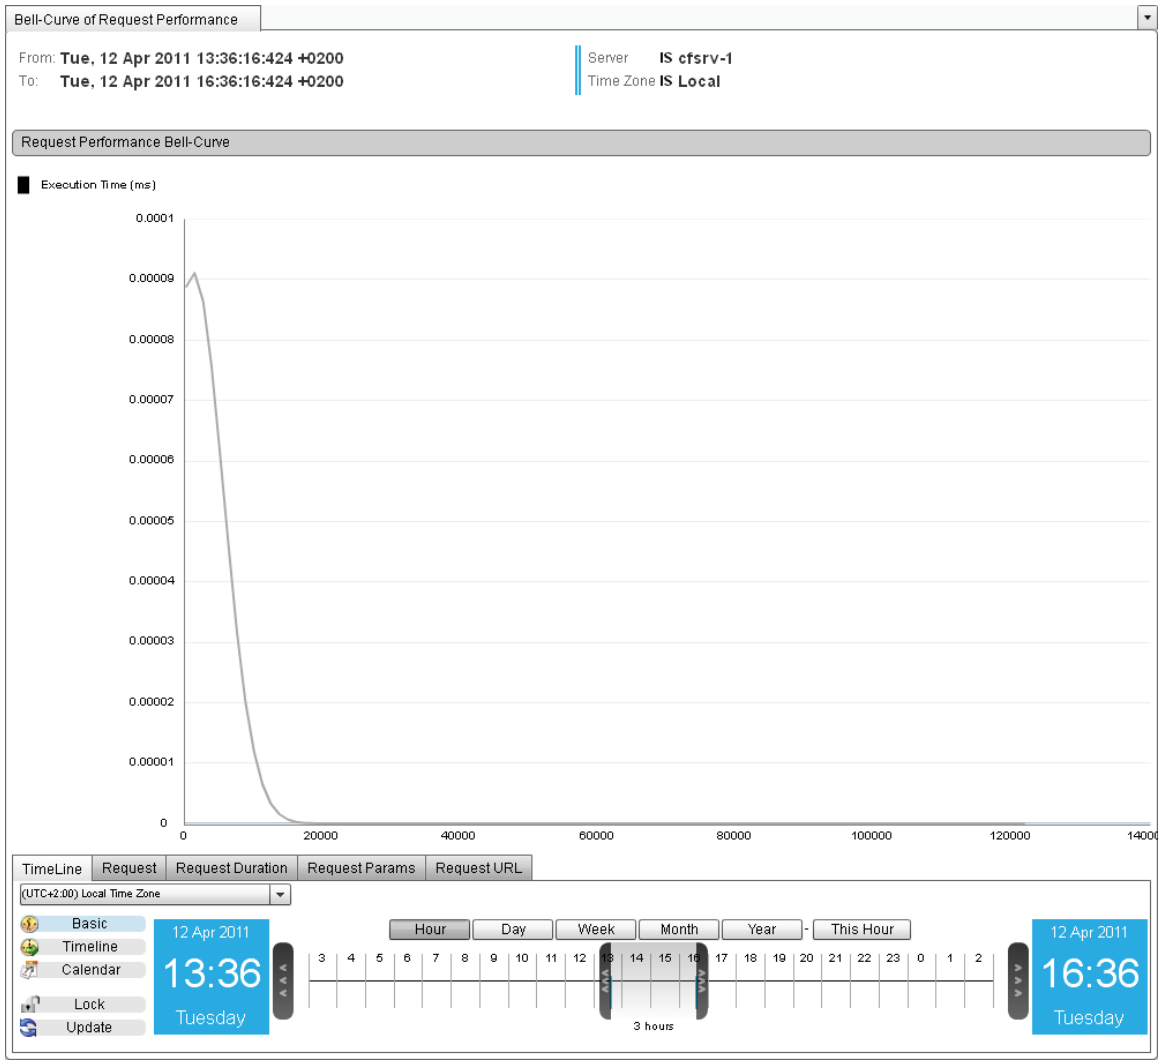

Figure 1: Bell-Curve of Request Performance Perspective

Back to the top

Request Status Code Brakedown

Request Status Code Brakedown

Description

Usage

Perspective View

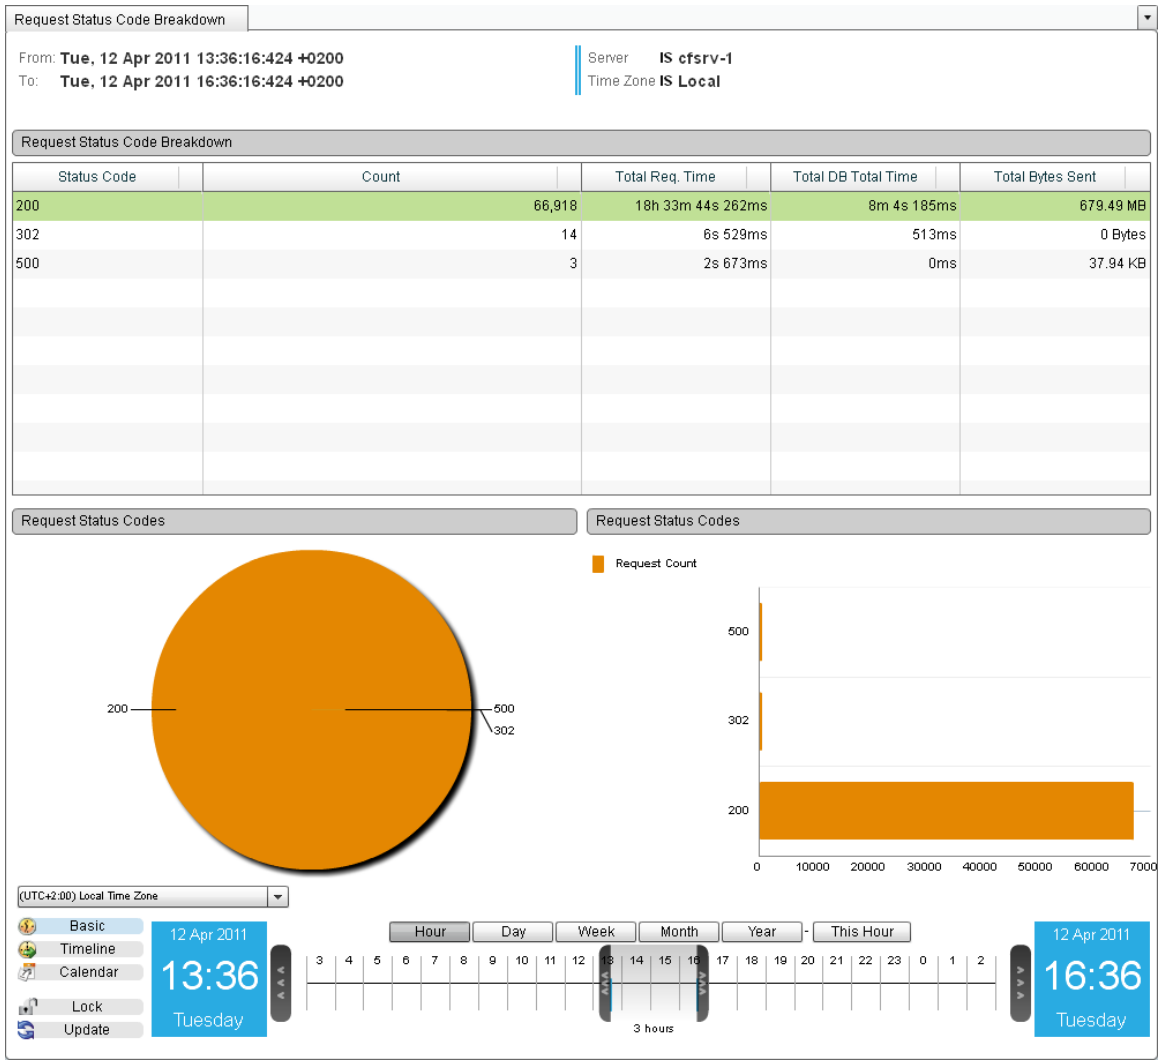

Figure 1: Request Status Code Brakedown Perspective

Back to the top

Request and JDBC Execution Time

Request and JDBC Execution Time

Description

This perspective shows a chart of the Average Request Execution Time within the specific period of time.

Usage

Perspective View

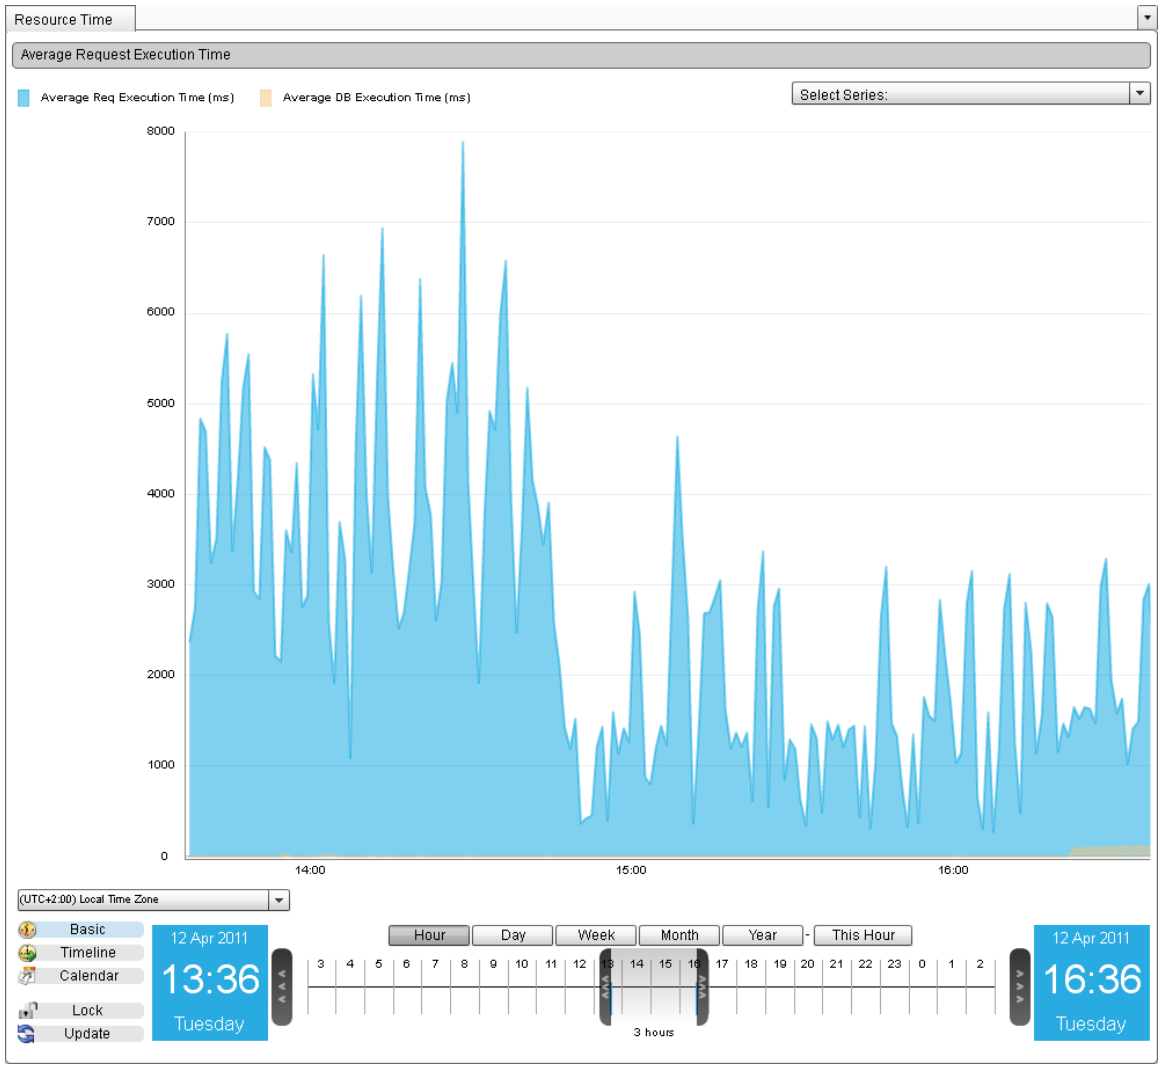

Figure 1: Request and JDBC Execution Time Perspective

Back to the top

Total Request Performance

Total Request Performance

Description

Usage

Perspective View

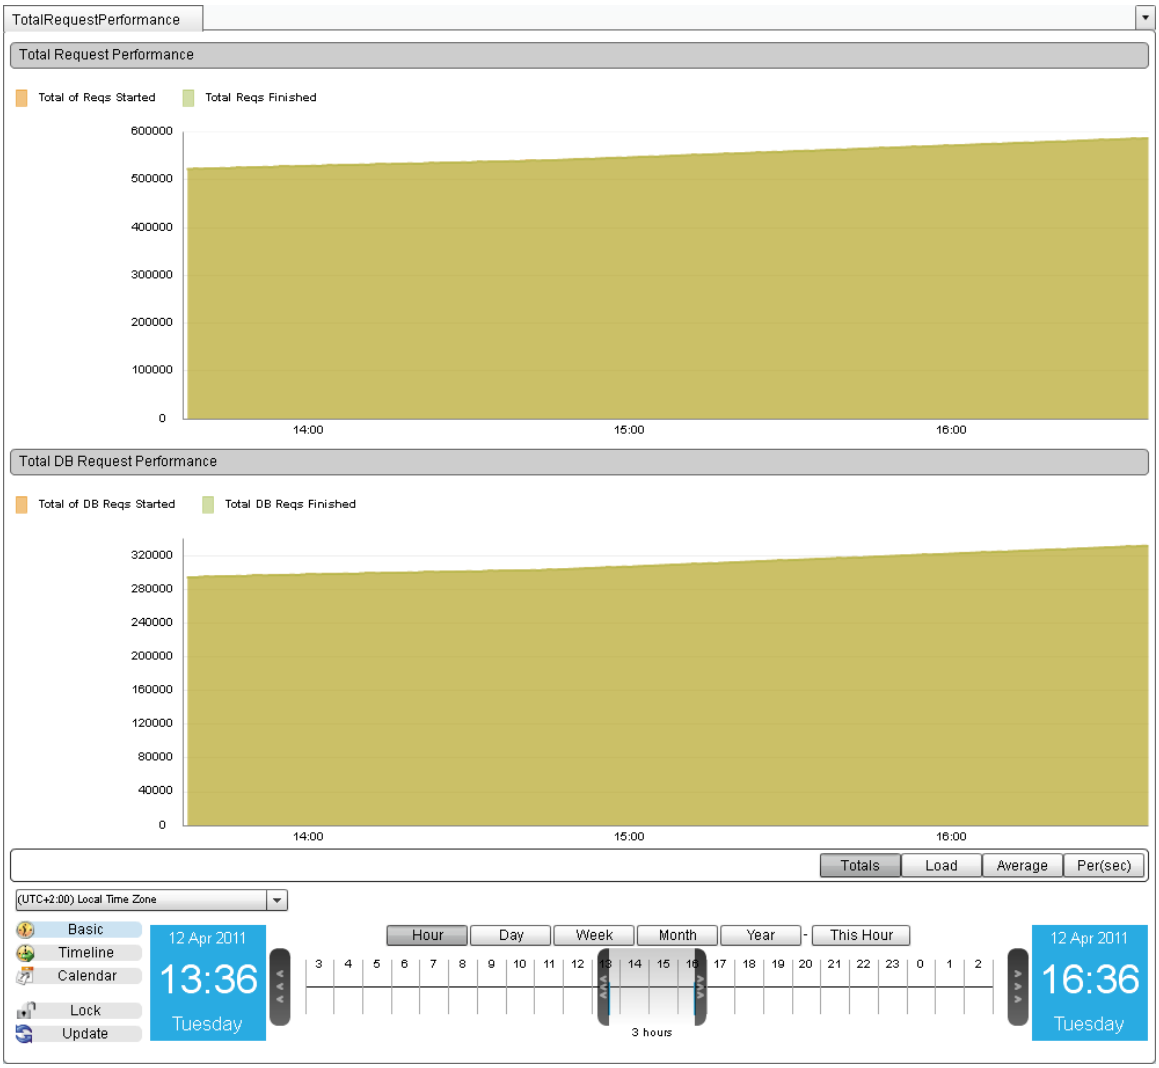

Figure 1: Total Request Performance

Back to the top

Database

Database

Overview

Next Steps

Database Requests Grid

Highest DB Request Row Count

Slowest Database Requests

Database Usage Graphs

Requests and JDBC Execution Time

Database Requests Grid

Database Requests Grid

Description

This perspective shows a table of the top 100 Database Requests within the specific period of time.

Usage

To change the order in which the requests are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

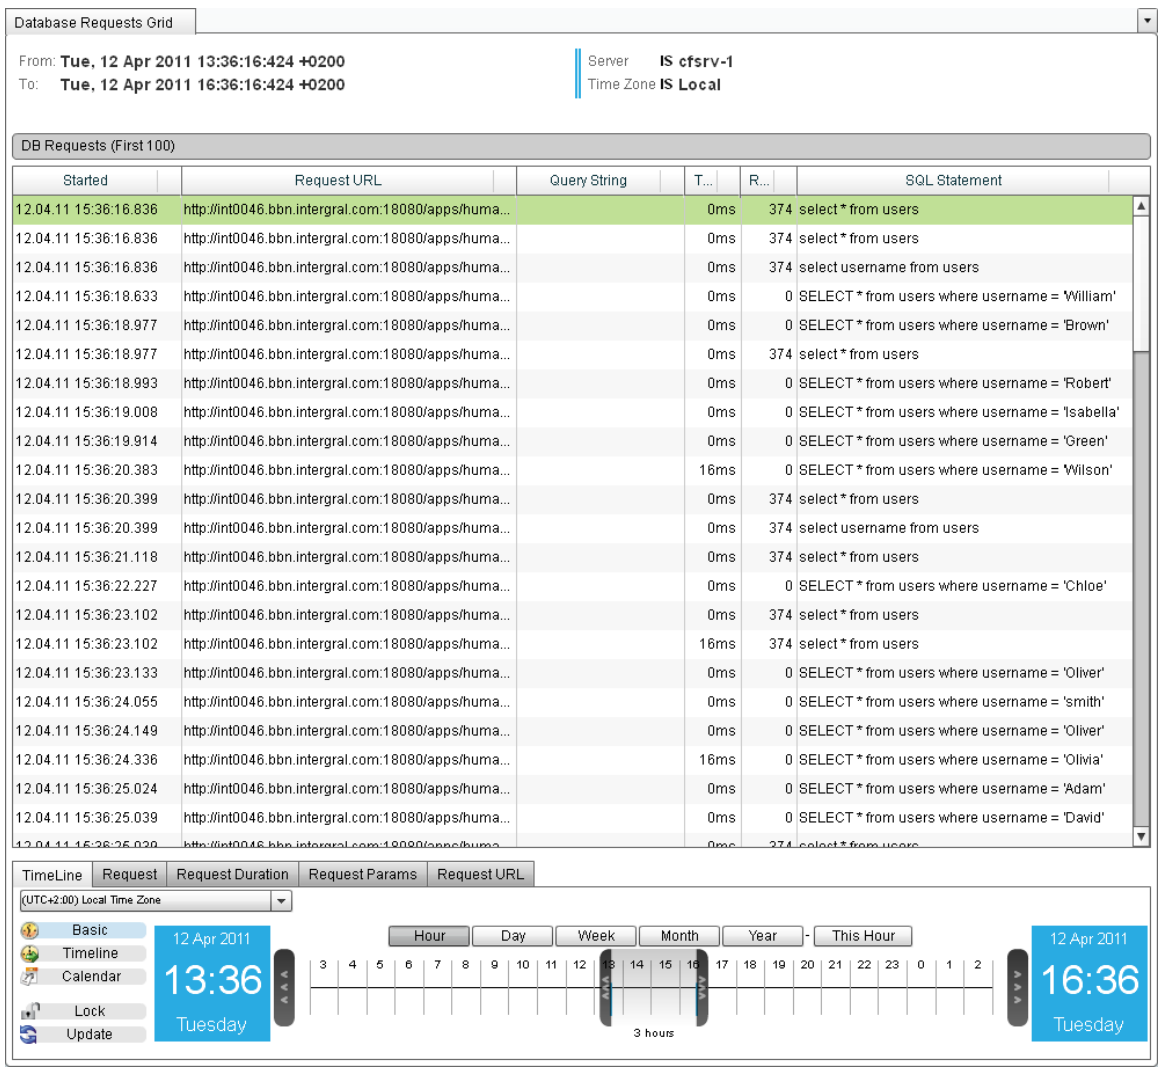

Figure 1: Database Requests Grid Perspective

Back to the top

Highest DB Request Row Count

Highest DB Request Row Count

Description

Highest Total DB Row Count is a report to help you identify which request fetched the most rows of data from the database. The report calculates the number of rows fetched from the database by a request. Knowing the number of rows returned from the database can help you to identify many potential issues and give you insight into the resource usage and scale of your applications. Generally fetching rows from the database takes time, resource and can induce network overhead if the database is located on a remote server. Often rows fetched from the database are stored in memory and this can increase the overall memory usage of your application or potentially create spikes in memory usage.

Using this report you should look to fix, tune or optimize the requests that return most rows from the database. This will typically increase application performance, lower resource usage including memory, CPU and network traffic.

Usage

- Discover which requests fetch the most rows of data from the database(s)
- Identify unexpected results - requests with poorly written queries
- Use the information to find bugs in pages and/or queries which are causing too many rows to be returned
- Use the information to optimize code
- Identify changes in application behavior across time

Perspective View

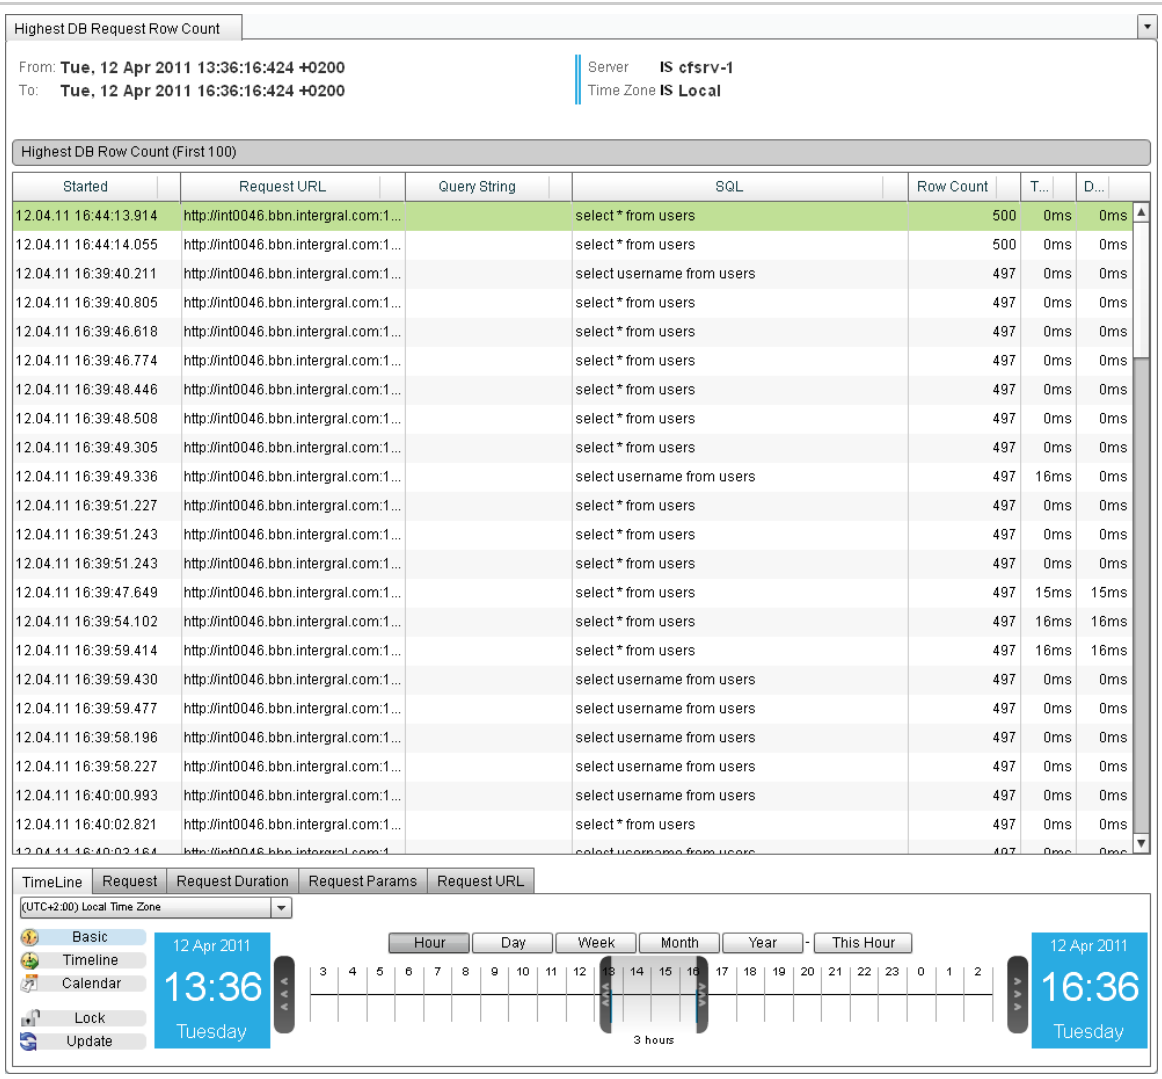

Figure 1: Highest DB Request Row Count Perspective

Back to the top

Slowest Database Requests

Slowest Database Requests

Description

This perspective shows a table of the top 100 Slowest Database Requests within the specific period of time.

Usage

To change the order in which the requests are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

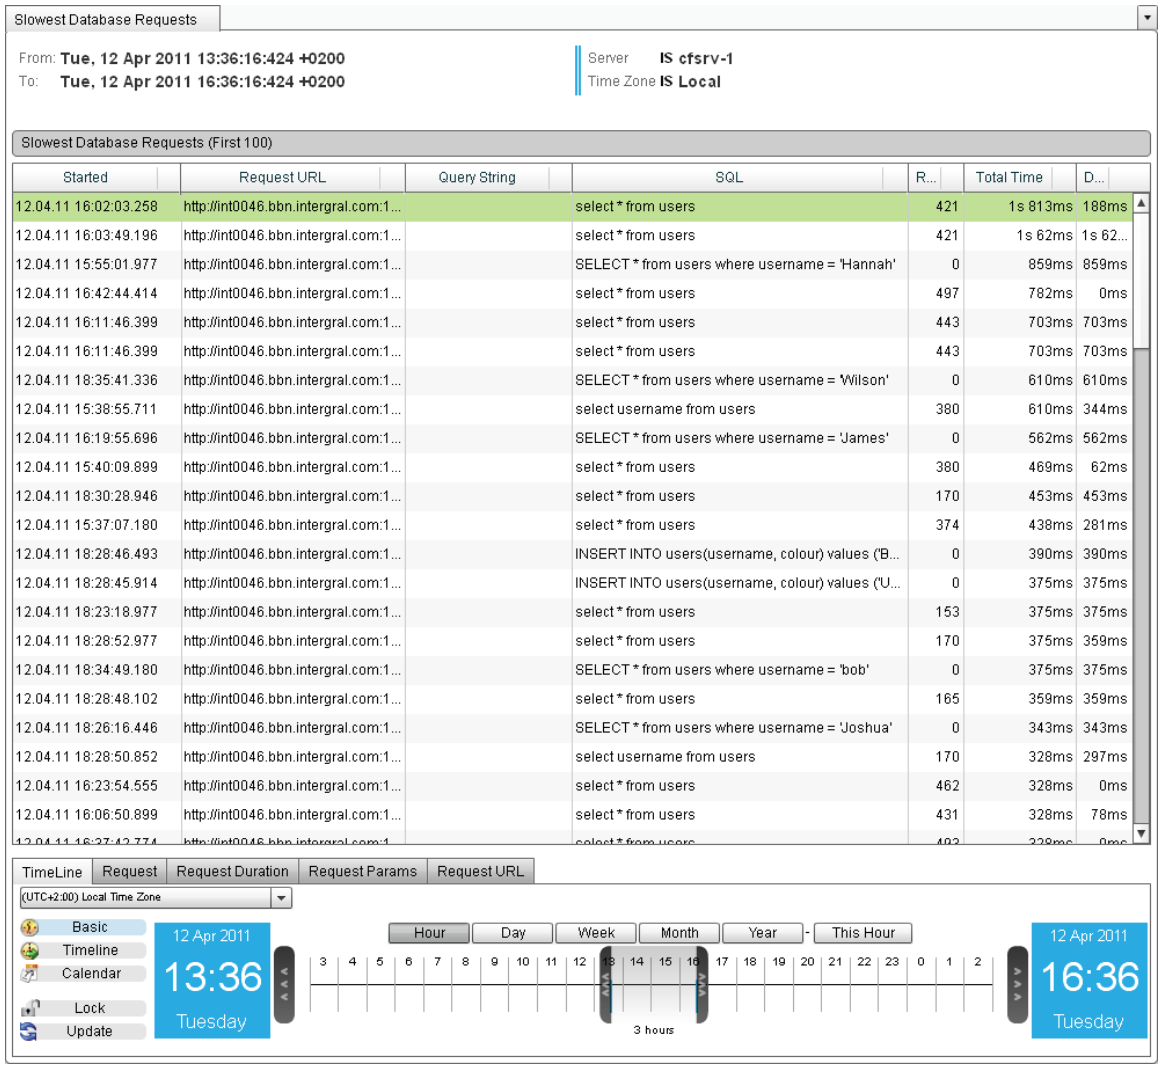

Figure 1: Slowest Database Requests Perspective

Back to the top

Database Usage Graphs

FusionAnalytics DataServices

Description

Usage

Perspective View

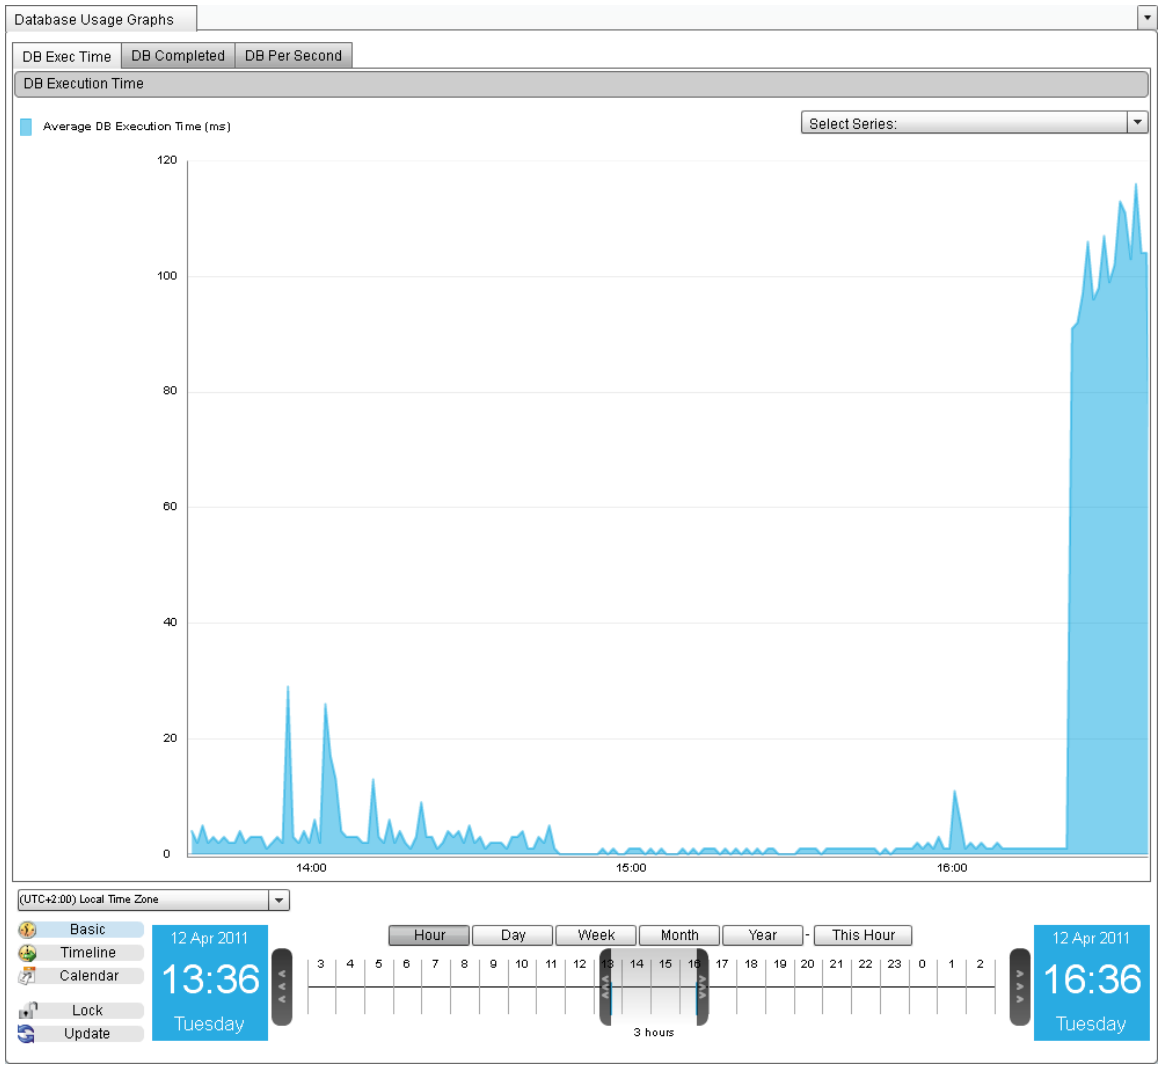

Figure 1: Database Usage Graphs Perspective

Back to the top

Requests and JDBC Execution Time

Requests and JDBC Execution Time

Description

This perspective shows a chart of the Average Request Execution Time within the specific period of time.

Usage

Perspective View

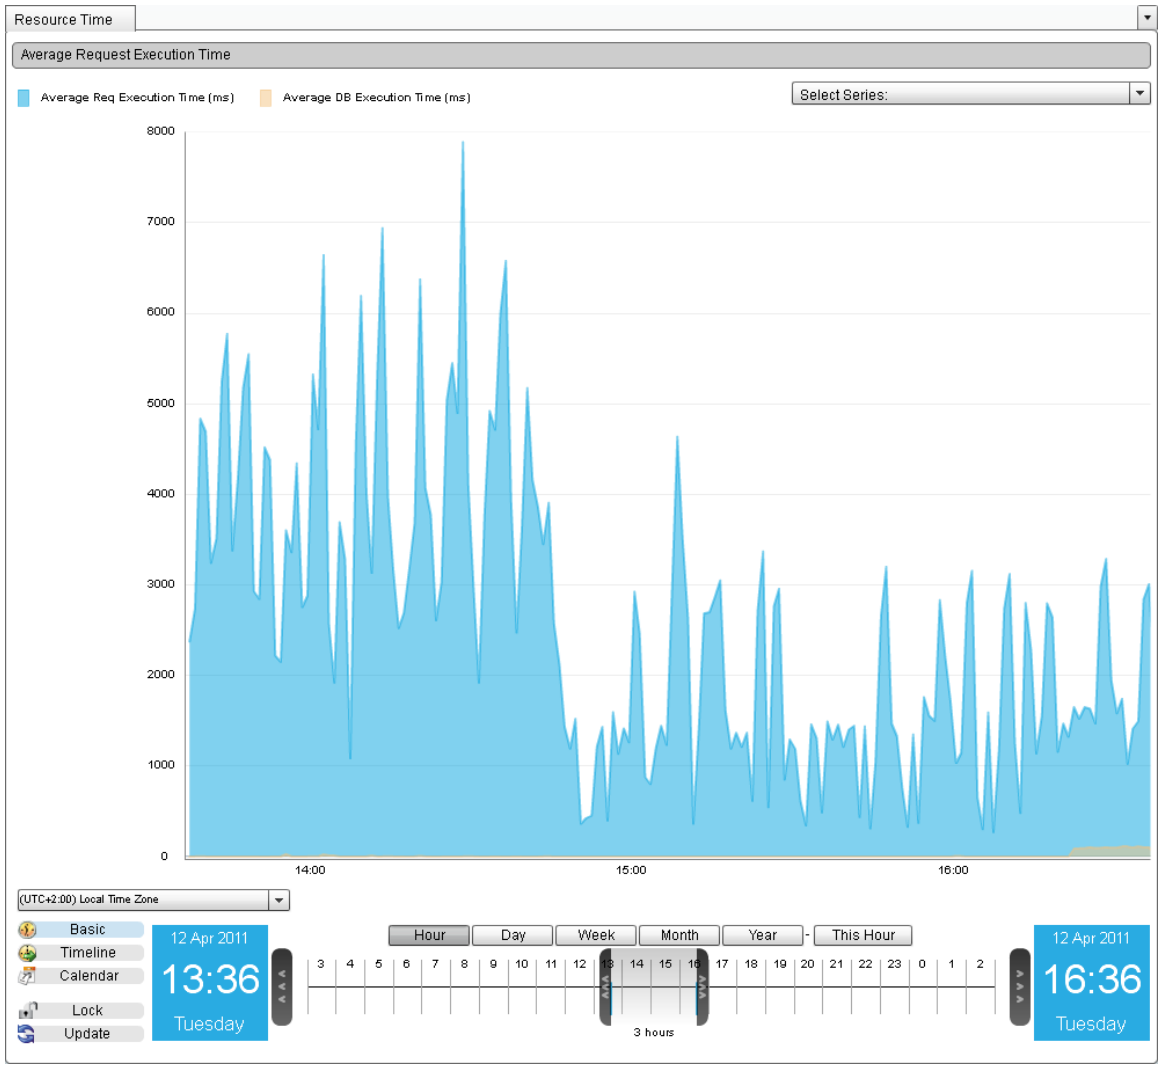

Figure 1: Requests and JDBC Execution Time Perspective

Back to the top

Memory

Memory

Overview

Next Steps

Memory Usage

Memory Spaces

Memory Heap Non-Heap Summary

Memory Usage

Memory Usage

Description

This perspective shows a chart of the Process memory usage within the specific period of time.

Usage

Perspective View

Figure 1: Memory Usage Perspective

Back to the top

Memory Spaces

Memory Spaces

Description

This perspective shows a chart of the Code-Cache memory usage within the specific period of time.

Usage

Perspective View

Figure 1: Memory Spaces Perspective

Back to the top

Memory Heap Non-Heap Summary

Memory Heap Non-Heap Summary

Description

This perspective shows a chart of the Memory Heap/Non-Heap Summary (Usage) within the specific period of time.

Usage

Perspective View

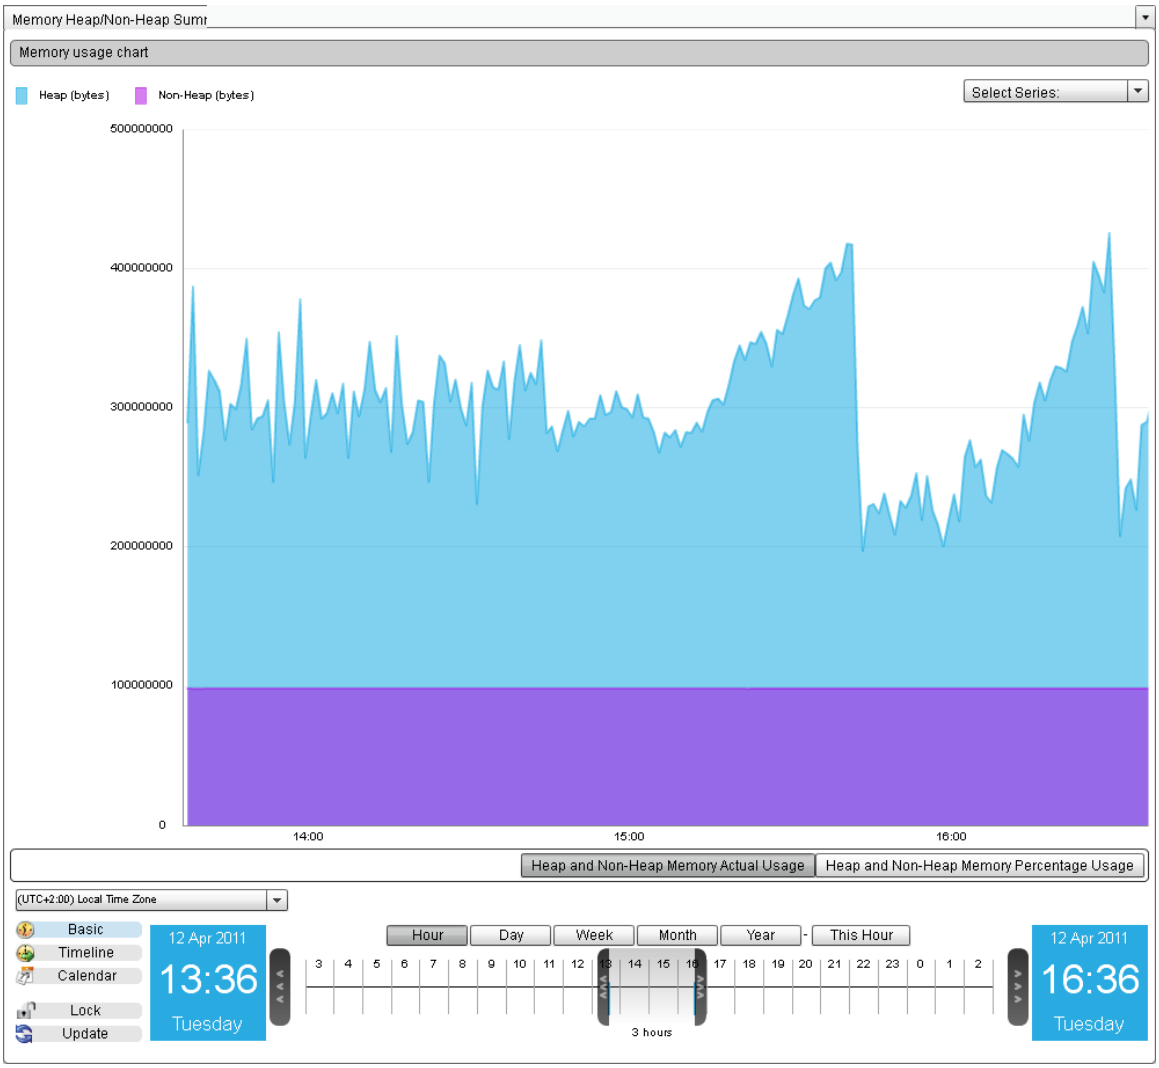

Figure 1: Memory Heap Non-Heap Summary Perspective

Back to the top

CPU

CPU

Overview

Next Steps

CPU Usage

CPU Memory Usage

CPU Usage

CPU Usage

Description

This perspective shows a chart of the Process and System CPU Usage within the specific period of time.

Usage

Perspective View

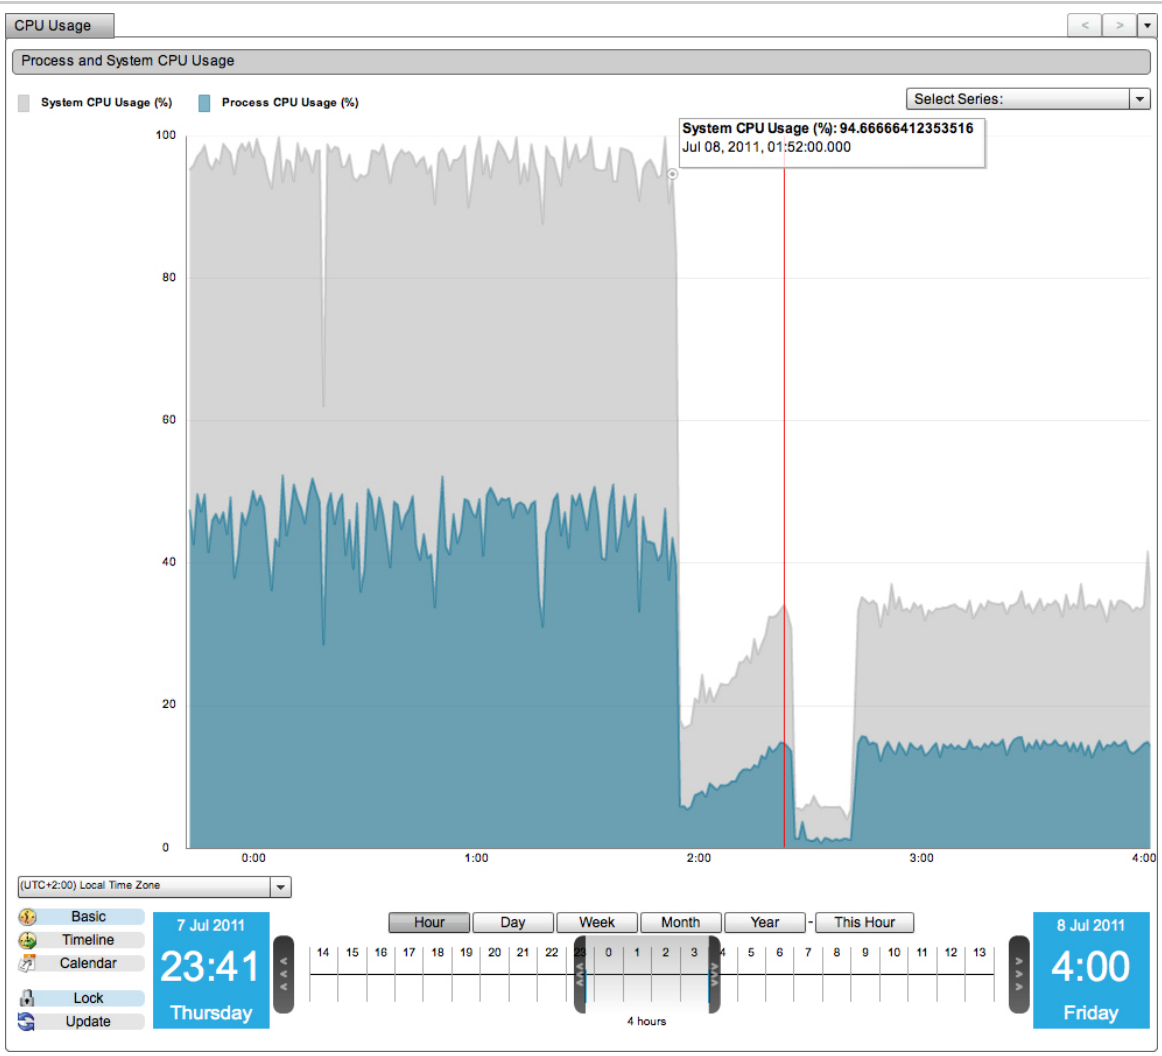

Figure 1: CPU Usage Perspective

Back to the top

CPU Memory Usage

CPU Memory Usage

Description

This perspective shows a chart of the CPU and Memory Usage within the specific period of time.

Usage

Perspective View

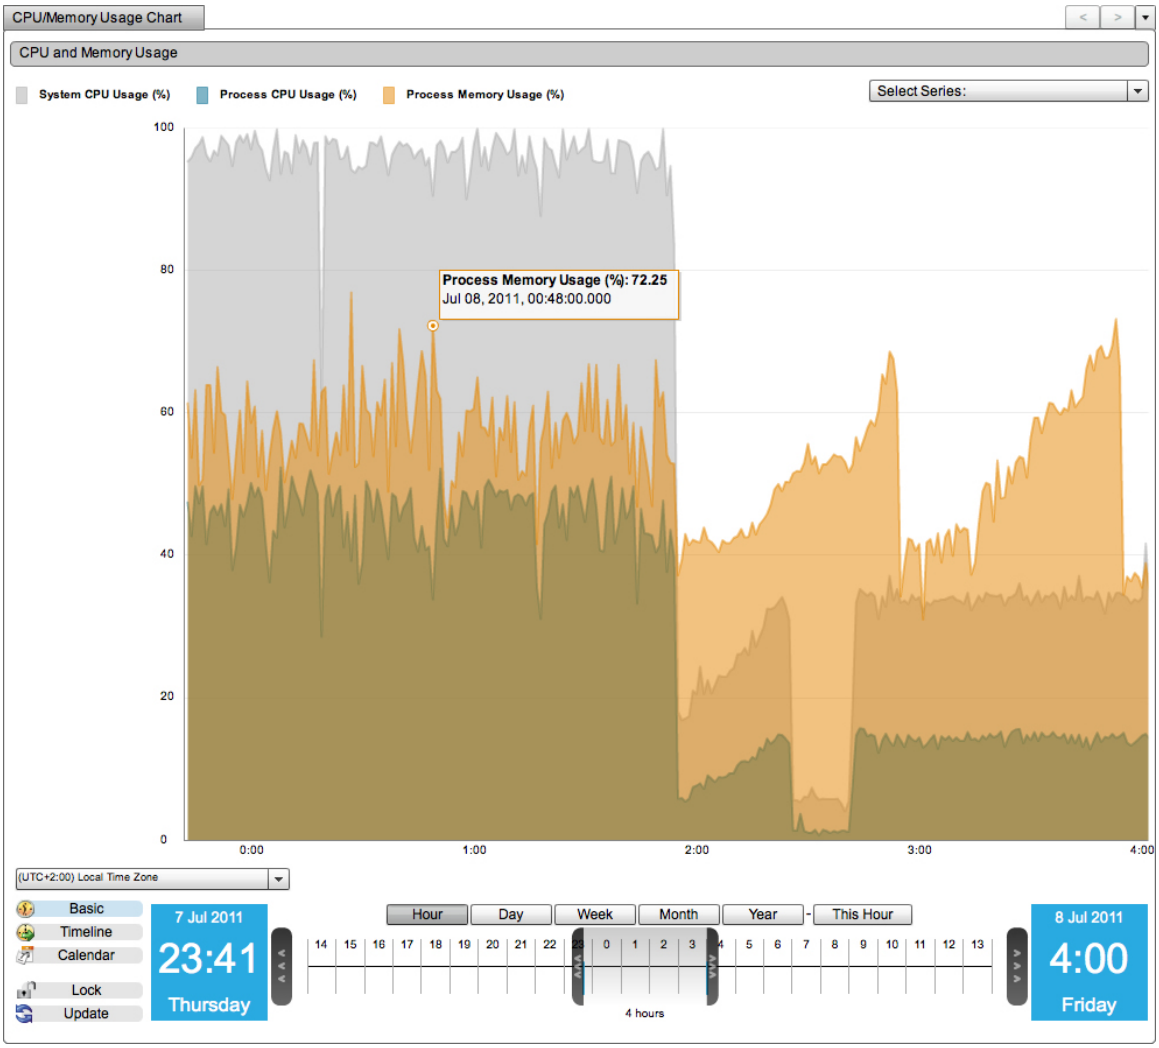

Figure 1: CPU Memory Usage Perspective

Back to the top

Thread

Thread

Overview

Next Steps

Thread Activity

Thread Activity

Thread Activity

Description

Usage

Perspective View

Back to the top

General

General

Overview

Next Steps

FusionReactor System Metrics

Server Summary

Resource Usage Data Grid

FusionReactor System Metrics

FusionReactor System Metrics

Description

This perspective shows the FusionReactor System Metrics Six-Way Charts within the specific period of time.

Usage

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.You can click on the garbage can on the Memory Demand graph to request the system garbage collector.

Perspective View

Figure 1: FusionReactor System Metrics Perspective

Back to the top

Server Summary

Server Summary

Description

This perspective shows a Server Summary table within the specific period of time.

Usage

To change the order in which the results are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

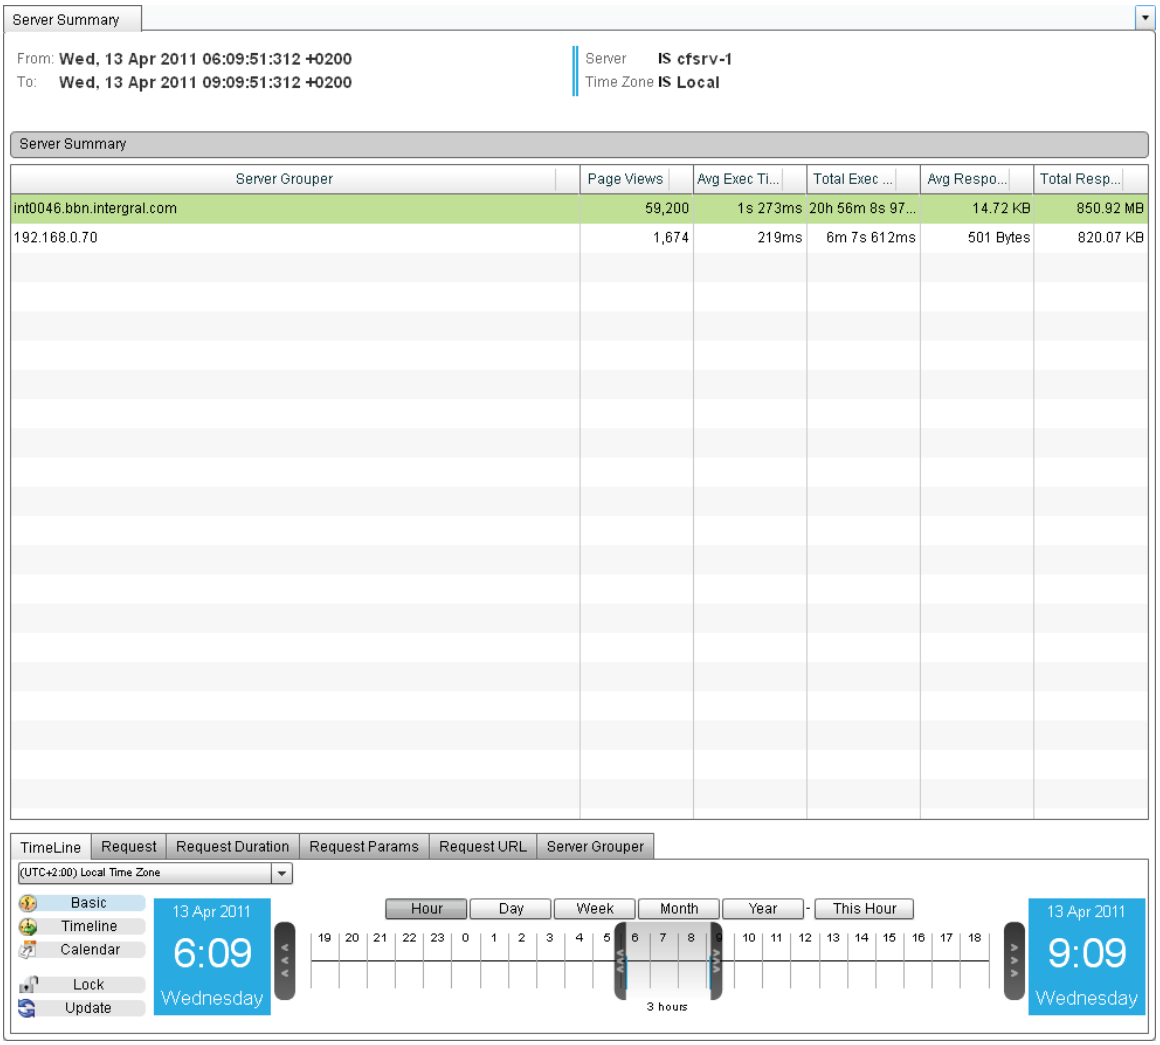

Figure 1: Server Summary Perspective

Back to the top

Resource Usage Data Grid

Resource Usage Data Grid

Description

This perspective shows a table of the top 100 Resource Usage Data within the specific period of time.

Usage

To change the order in which the results are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

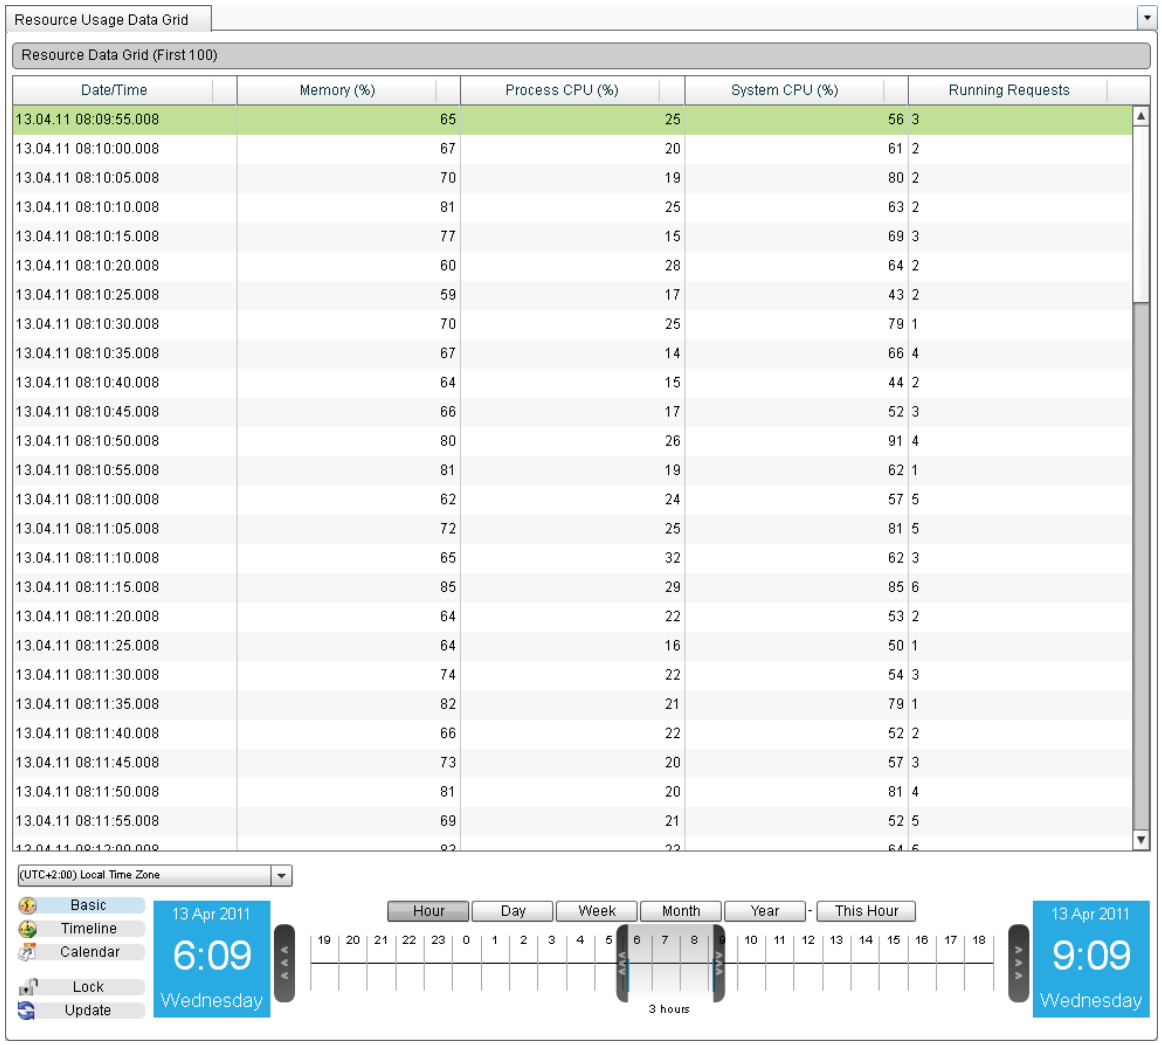

Figure 1: Resource Usage Data Grid Perspective

Back to the top

Logs

Logs

Overview

Next Steps

Log File Entries

Log File Entries

Log File Entries

Description

This perspective shows a table of the top 100 Log File Entries within the specific period of time.

Usage

To change the order in which the log entries are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

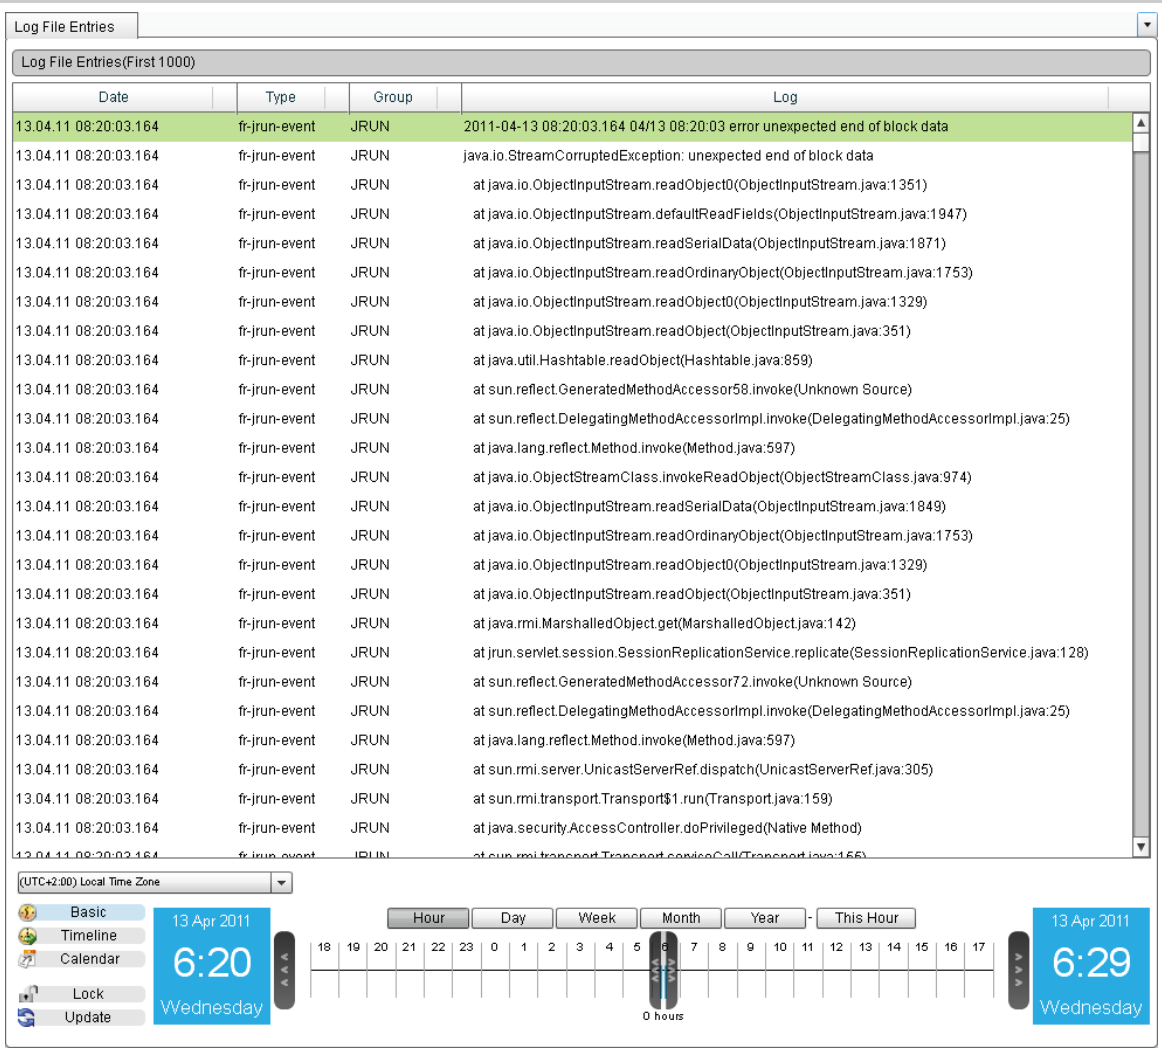

Figure 1: Log File Entries Perspective

Back to the top

Events

Events

Overview

Next Steps

Annotations

Placemarks

Crash Protection

Resources with Annotations

Annotations

Annotations

Description

Usage

Perspective View

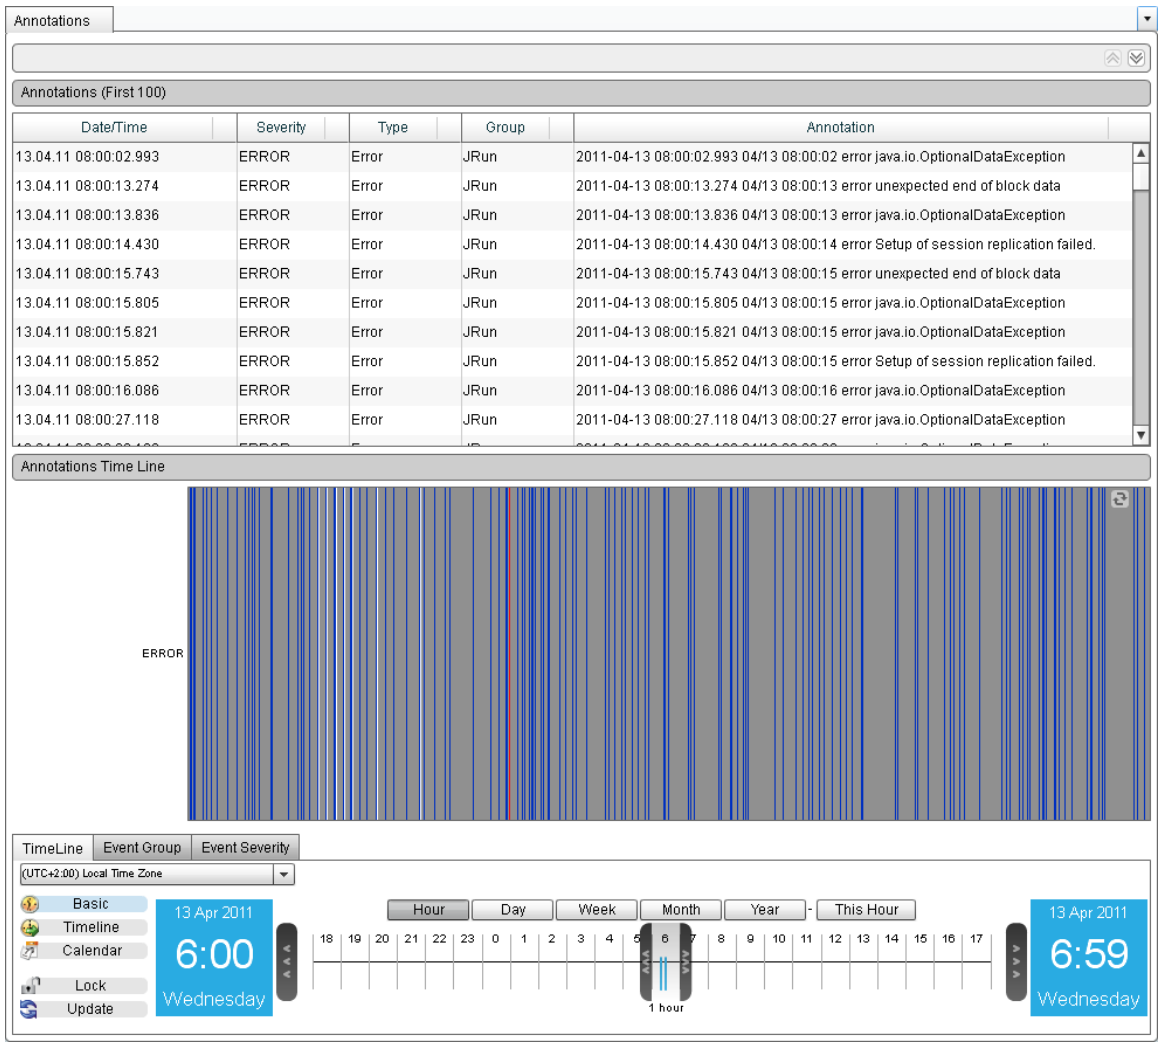

Figure 1: Annotations Perspective

Back to the top

Placemarks

Placemarks

Description

This perspective shows a chart containing the Application Placemarks within the specific period of time.

Usage

Perspective View

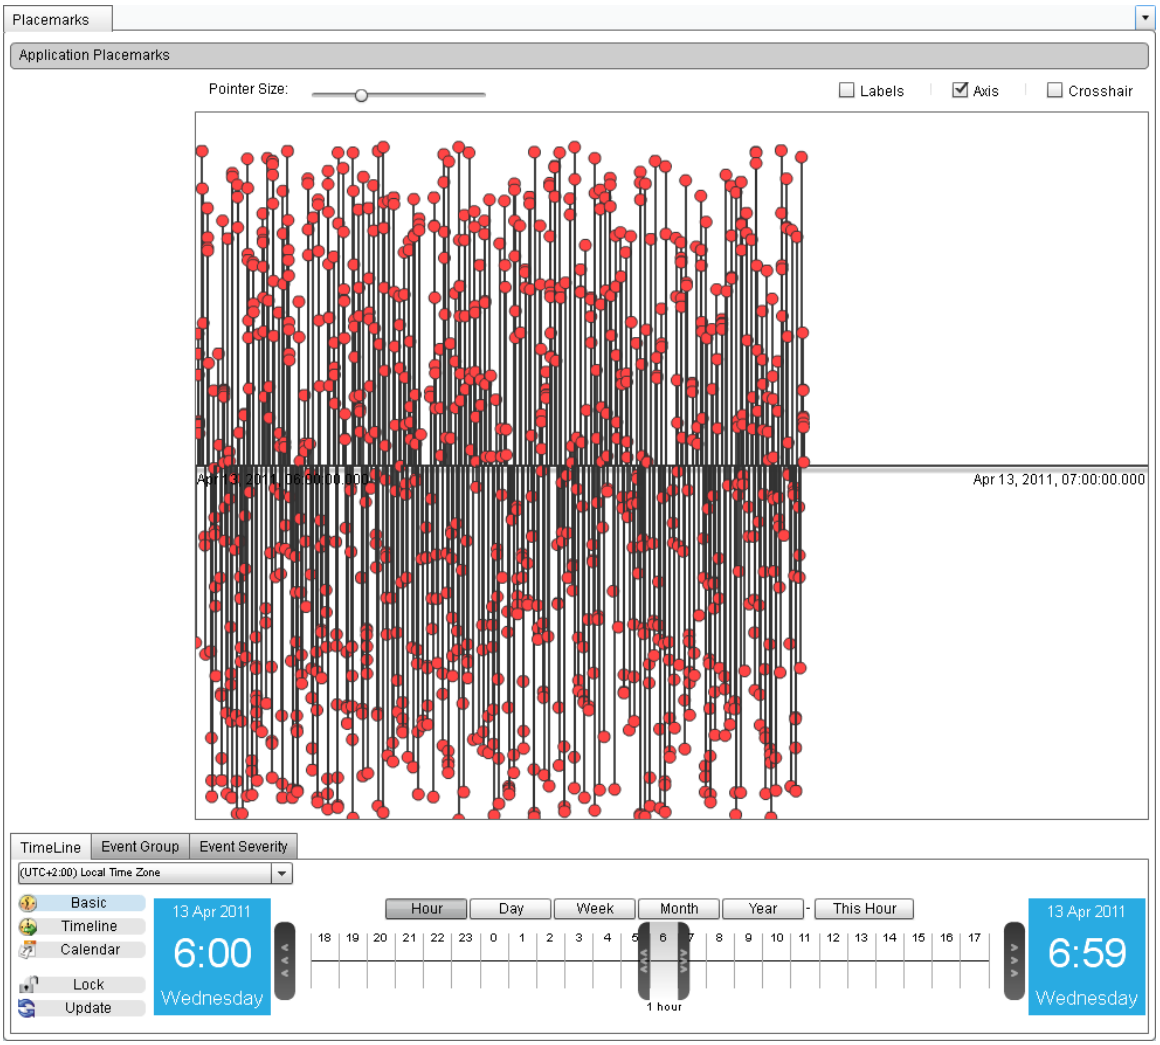

Figure 1: Placemarks Perspective

Back to the top

Crash Protection

Crash Protection

Description

Usage

Perspective View

Back to the top

Resources with Annotations

Resources with Annotations

Description

Usage

Perspective View

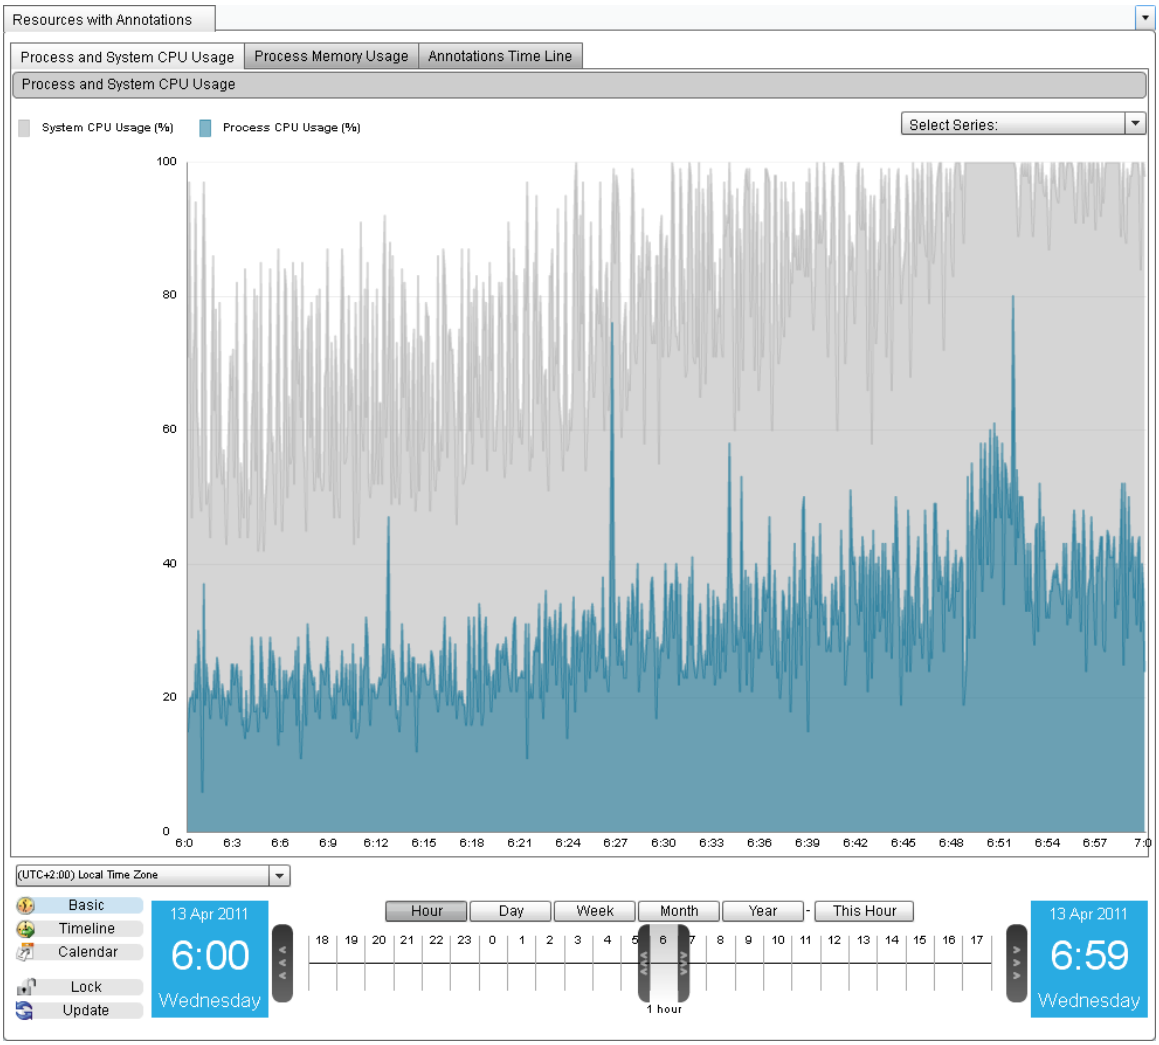

Figure 1: Resources with Annotations Perspective

Back to the top

Hot-Spots

Hot-Spots

Overview

Next Steps

Memory Hot-Spots

- Memory at 100%
- Memory Highest Usage
- Memory Events

CPU Hot-Spots

- Process CPU at 100%
- Process CPU Highest Usage
- System CPU at 100%
- System CPU Highest Usage

Request Hot-Spots

- Requests Not Completed
- Slowest Requests (Hot-Spots)
- Slowest Requests On Average (Hot-Spots)
- Slowest DB Requests On Average (Hot-Spots)

Database Hot-Spots

- Slowest Database Requests (Hot-Spots)

Memory Hot-Spots

Memory Hot-Spots

Overview

Next Steps

Memory at 100% Memory Highest Usage Memory Events

Memory at 100%

Memory at 100%

Description Usage Perspective View

Back to the top

Memory Highest Usage

Memory Highest Usage

Description

This perspective shows a table of the top 100 Highest Memory Usage within the specific period of time.

Usage

To change the order in which the results are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

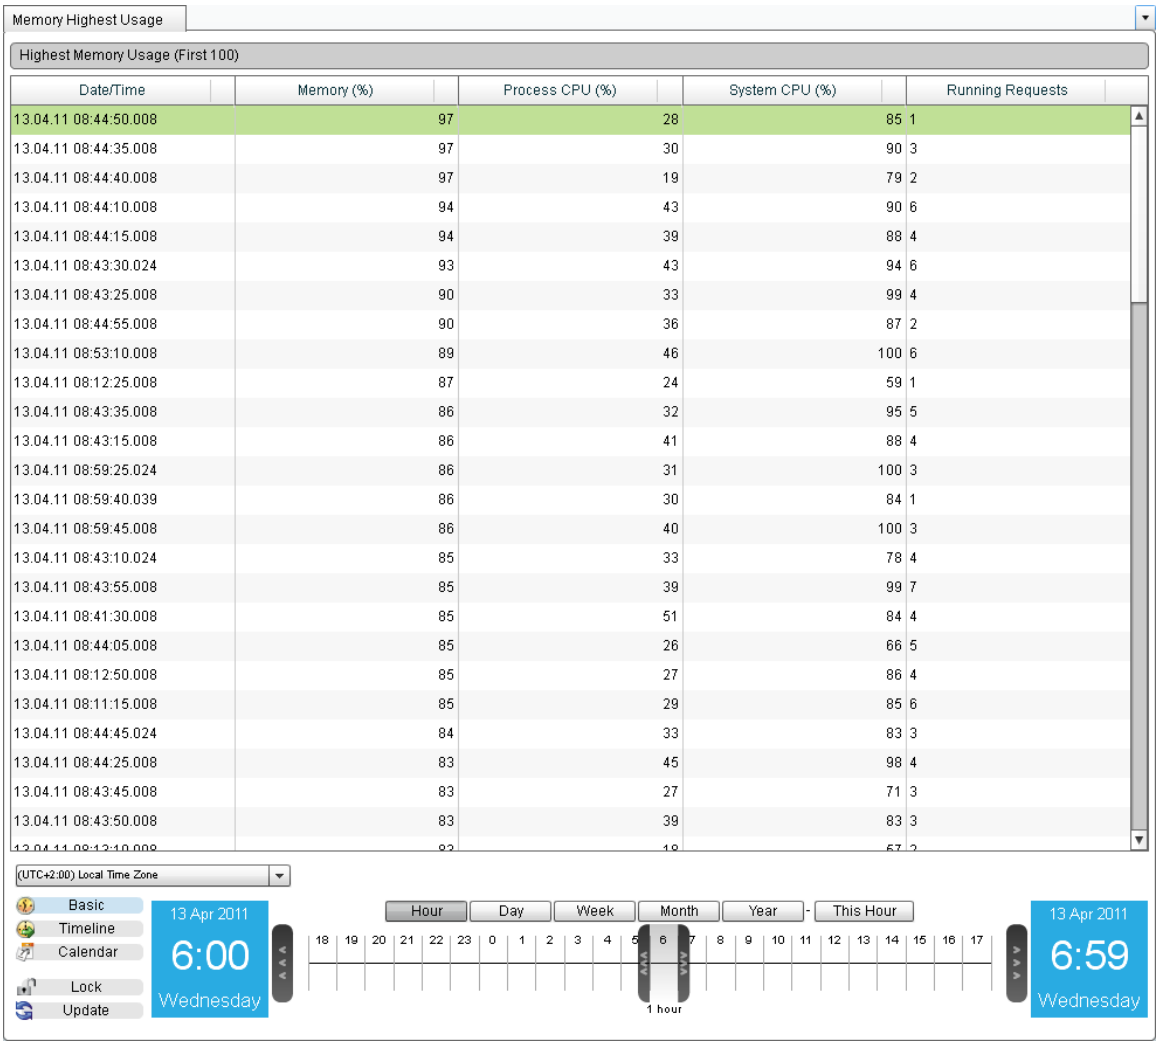

Figure 1: Memory Highest Usage Perspective

Back to the top

Memory Events

Memory Events

Description Usage Perspective View

Back to the top

CPU Hot-Spots

CPU Hot-Spots

Overview

Next Steps

Process CPU at 100% Process CPU Highest Usage System CPU at 100% System CPU Highest Usage

Process CPU at 100%

Process CPU at 100%

Description Usage Perspective View

Back to the top

Process CPU Highest Usage

Process CPU Highest Usage

Description

This perspective shows a table of the top 100 Highest Process CPU Usage within the specific period of time.

Usage

To change the order in which the results are shows, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

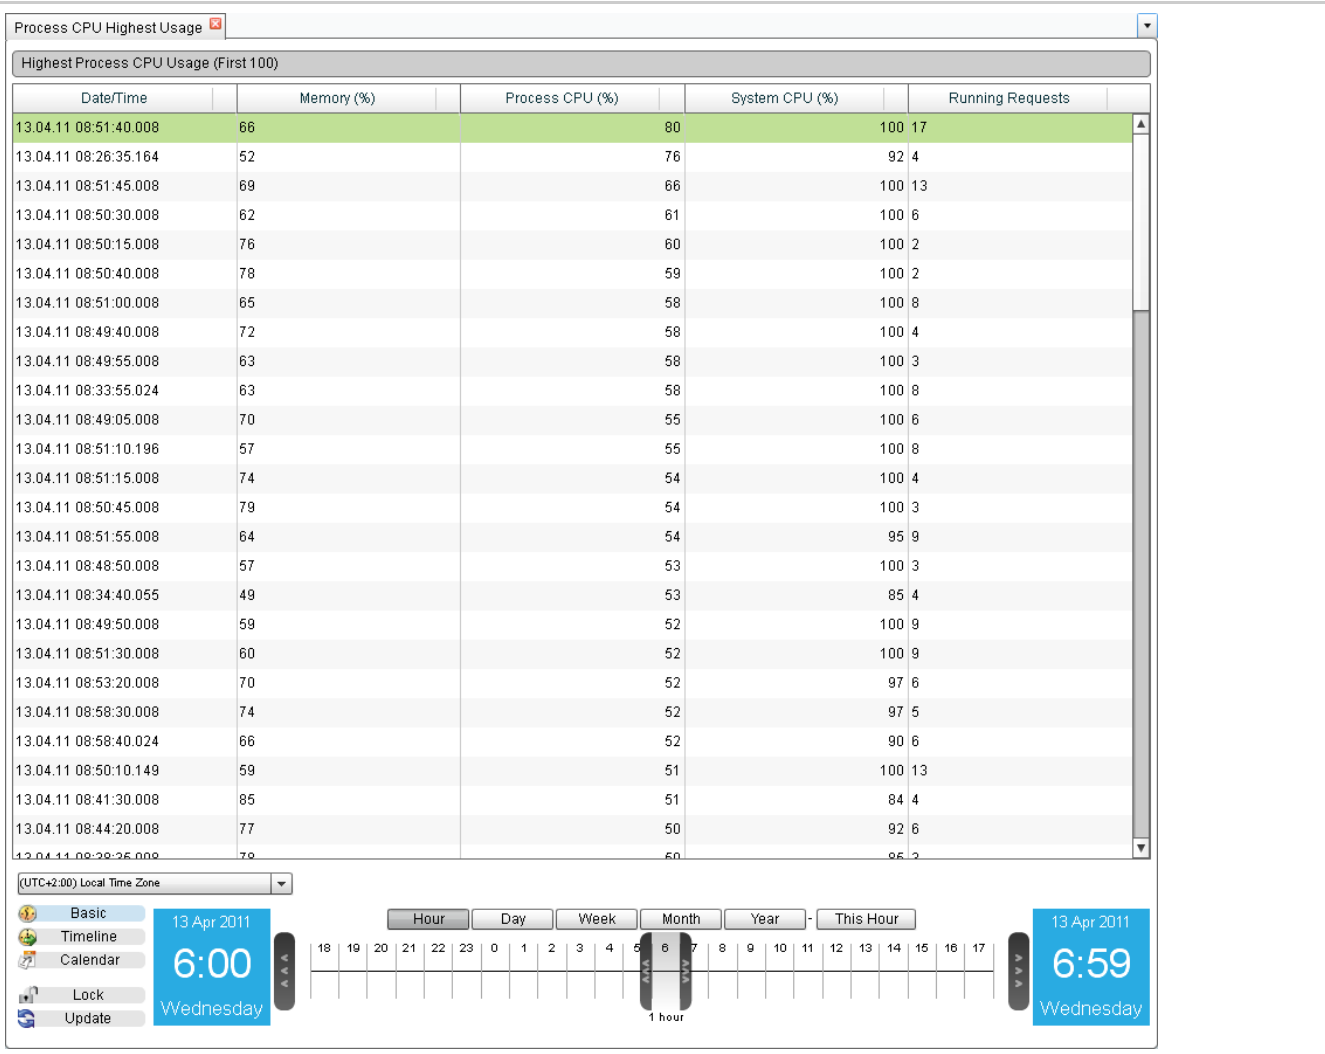

Figure 1: Process CPU Highest Usage Perspective

Back to the top

System CPU at 100%

System CPU at 100%

Description

This perspective shows a table of the top 100 System CPU at 100% within the specific period of time.

Usage

To change the order in which the results are shows, click on one of the table column headings. Use the Date Navigation in order to change

the period of time over which this perspective covers.

Perspective View

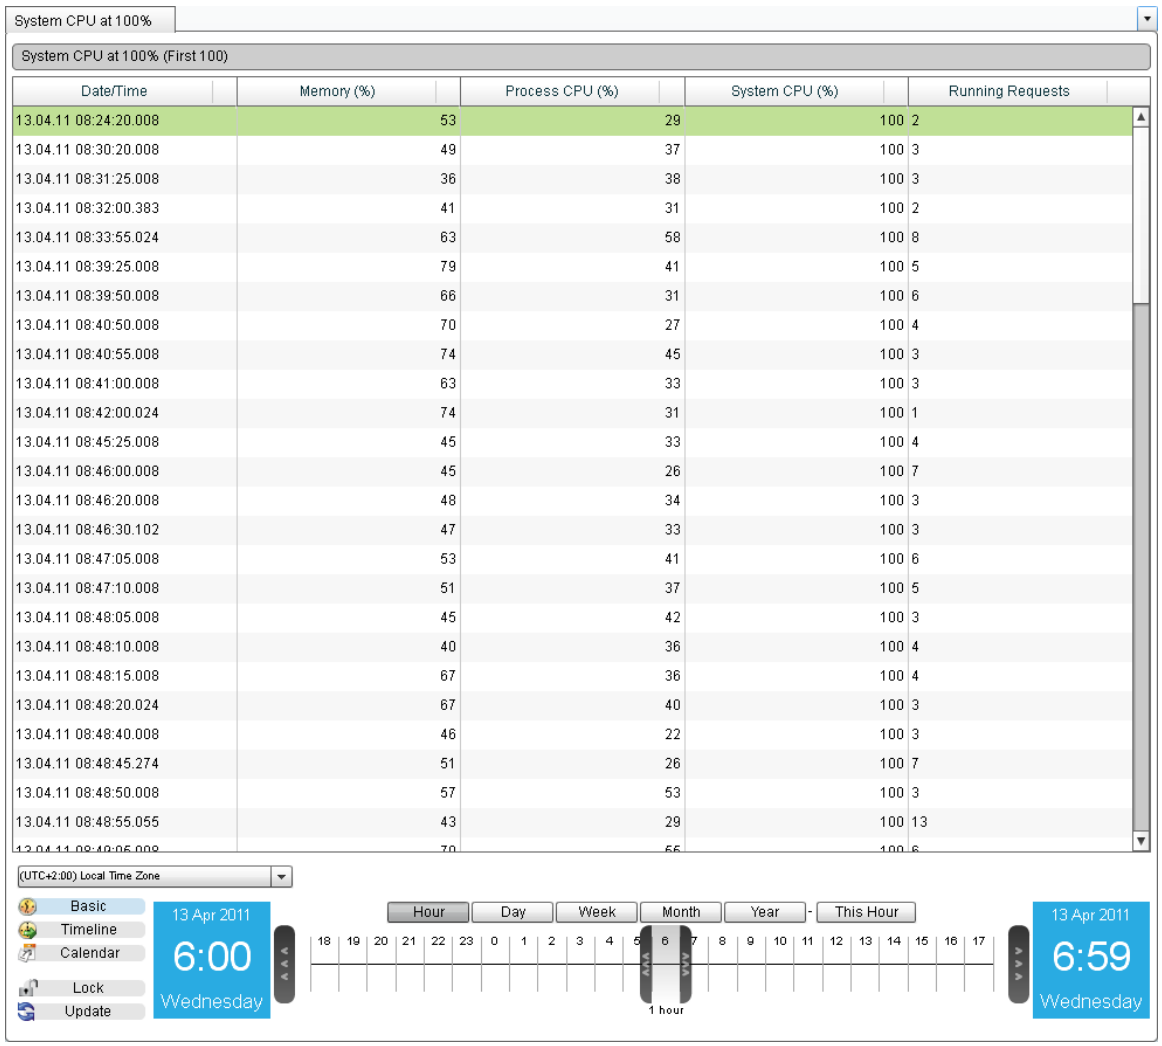

Figure 1: System CPU at 100% Perspective

Back to the top

System CPU Highest Usage

System CPU Highest Usage

Description

This perspective shows a table of the top 100 Process and System CPU Highest Usage within the specific period of time.

Usage

To change the order in which the results are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

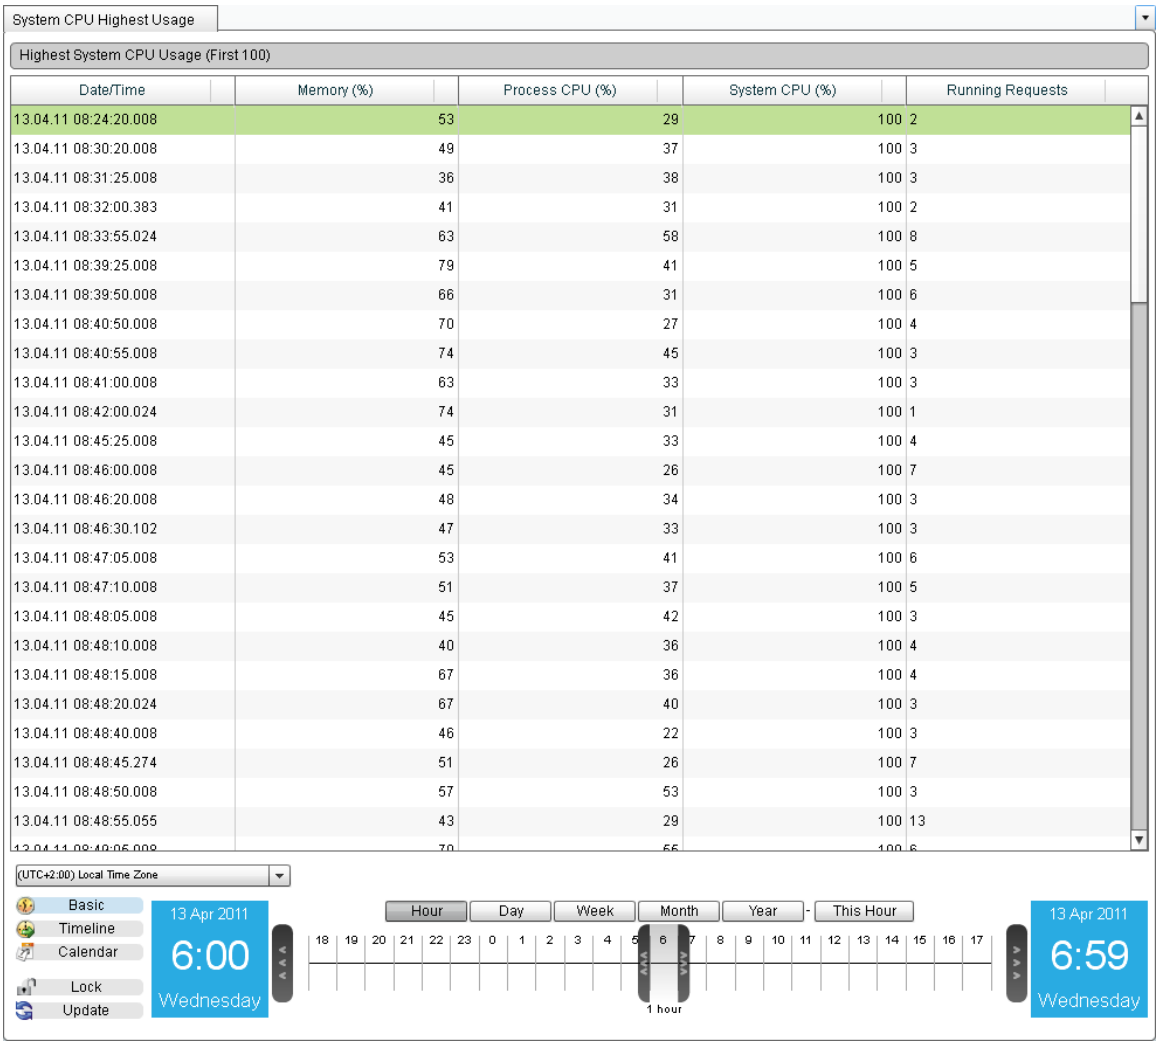

Figure 1: System CPU Highest Usage

Back to the top

Request Hot-Spots

Request Hot-Spots

Overview

Next Steps

Requests Not Completed Slowest Requests (Hot-Spots) Slowest Requests On Average (Hot-Spots) Slowest DB Requests On Average (Hot-Spots)

Requests Not Completed

Requests Not Completed

Description Usage Perspective View

Back to the top

Slowest Requests (Hot-Spots)

Slowest Requests (Hot-Spots)

Description

This perspective shows a table of the top 100 Slowest Requests within the specific period of time.

Usage

To change the order in which the requests are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

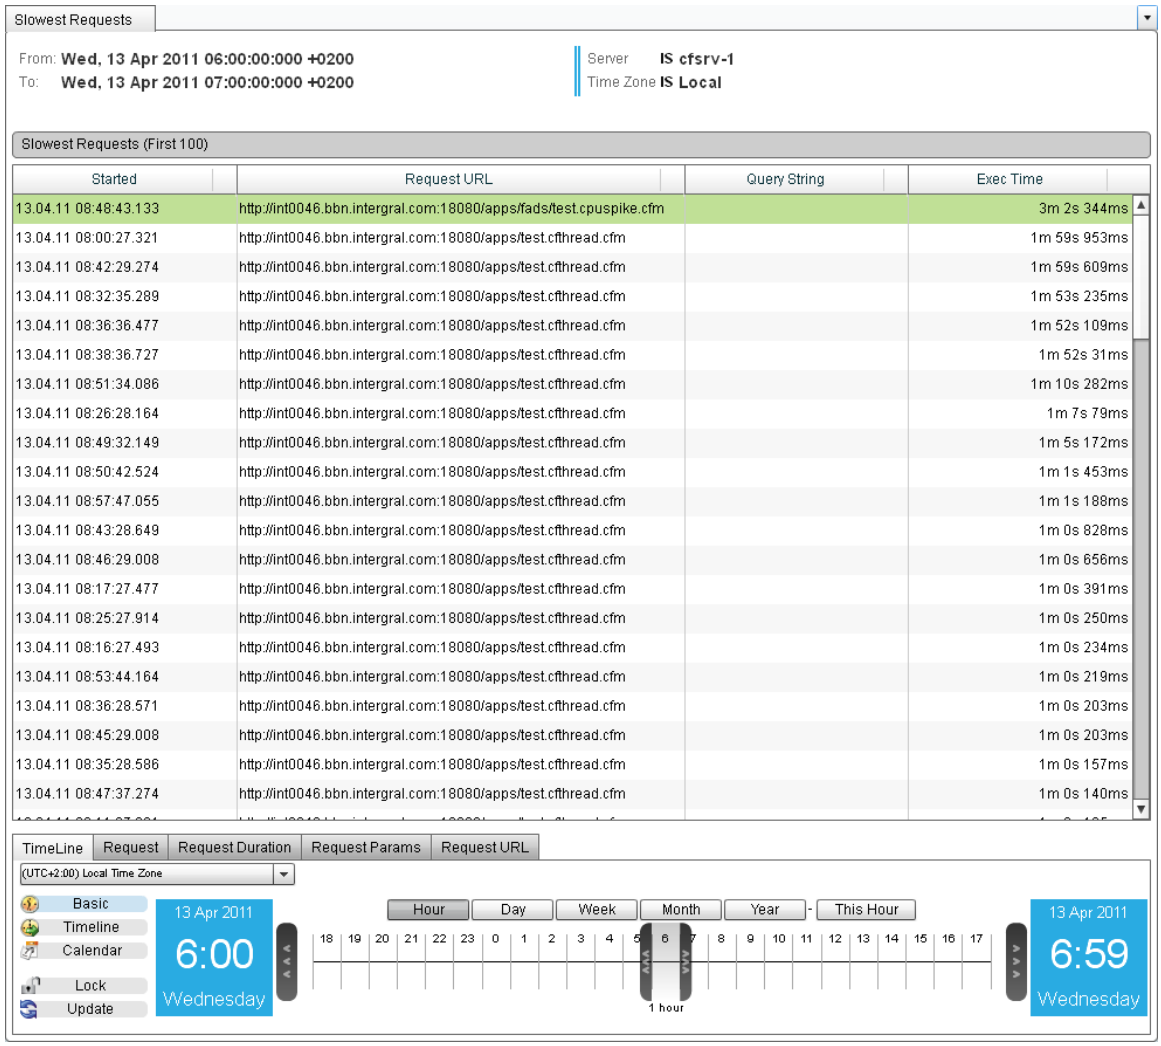

Figure 1: Slowest Requests (Hot-Spot) Perspective

Back to the top

Slowest Requests On Average (Hot-Spots)

Slowest Requests On Average (Hot-Spots)

Description

This perspective shows a table of the top 100 Slowest Requests on Average (Grouped by URL) within the specific period of time.

Usage

To change the order in which the requests are shown, click on one of the table column headings. Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View

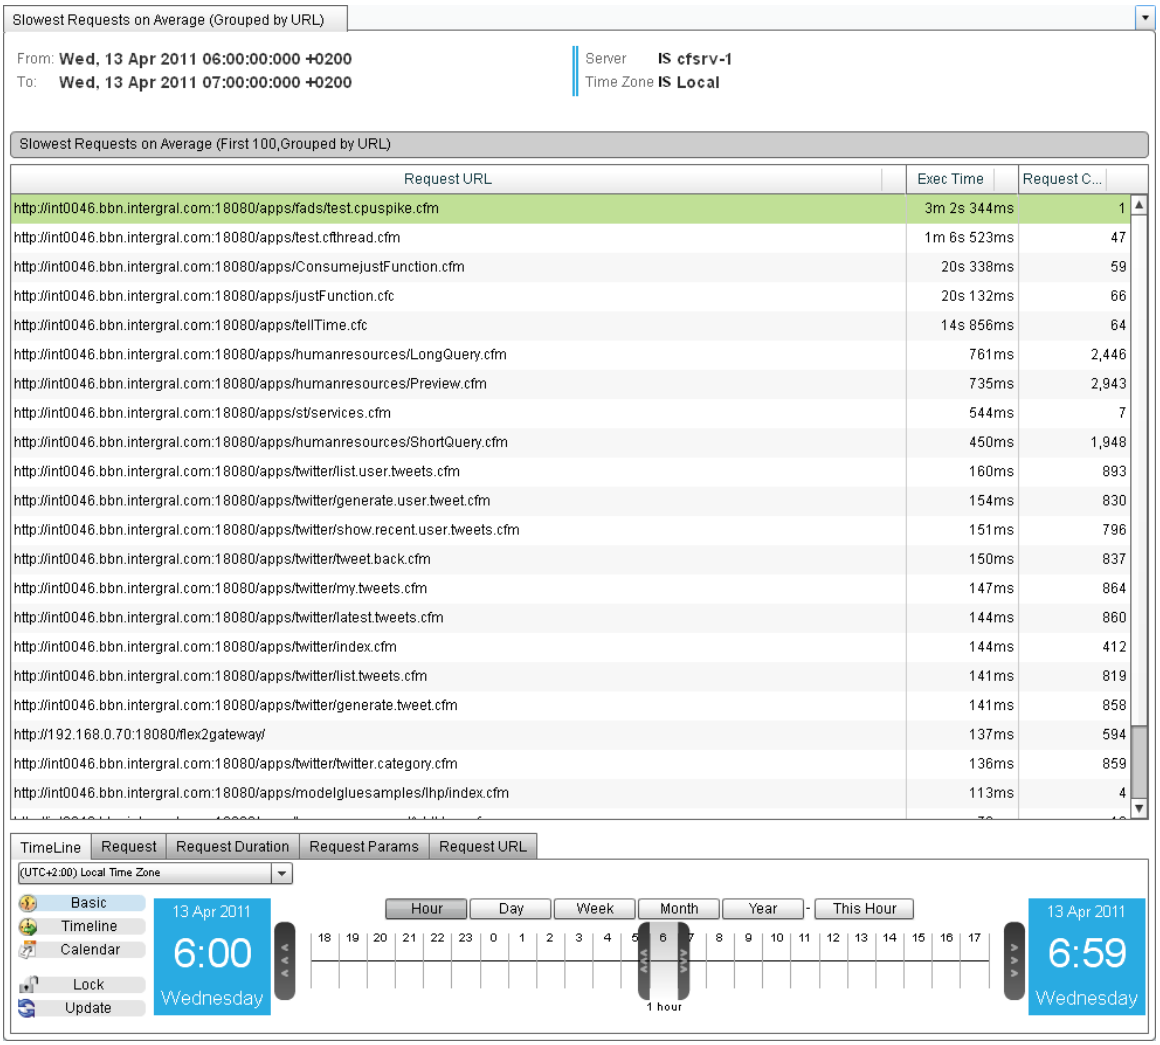

Figure 1: Slowest Requests On Average (Hot-Spot) Perspective

Back to the top

Slowest DB Requests On Average (Hot-Spots)

Slowest DB Requests On Average (Hot-Spots)

Description

This perspective shows a table of the top 100 Slowest Database Requests on Average by a Request (Grouped by URL) within the specific period of time.

Usage

To change the order in which the requests are shown, click on one of the table column headings. Use the FA10:Date Navigation] in order to change the period of time over which this perspective covers.

Perspective View

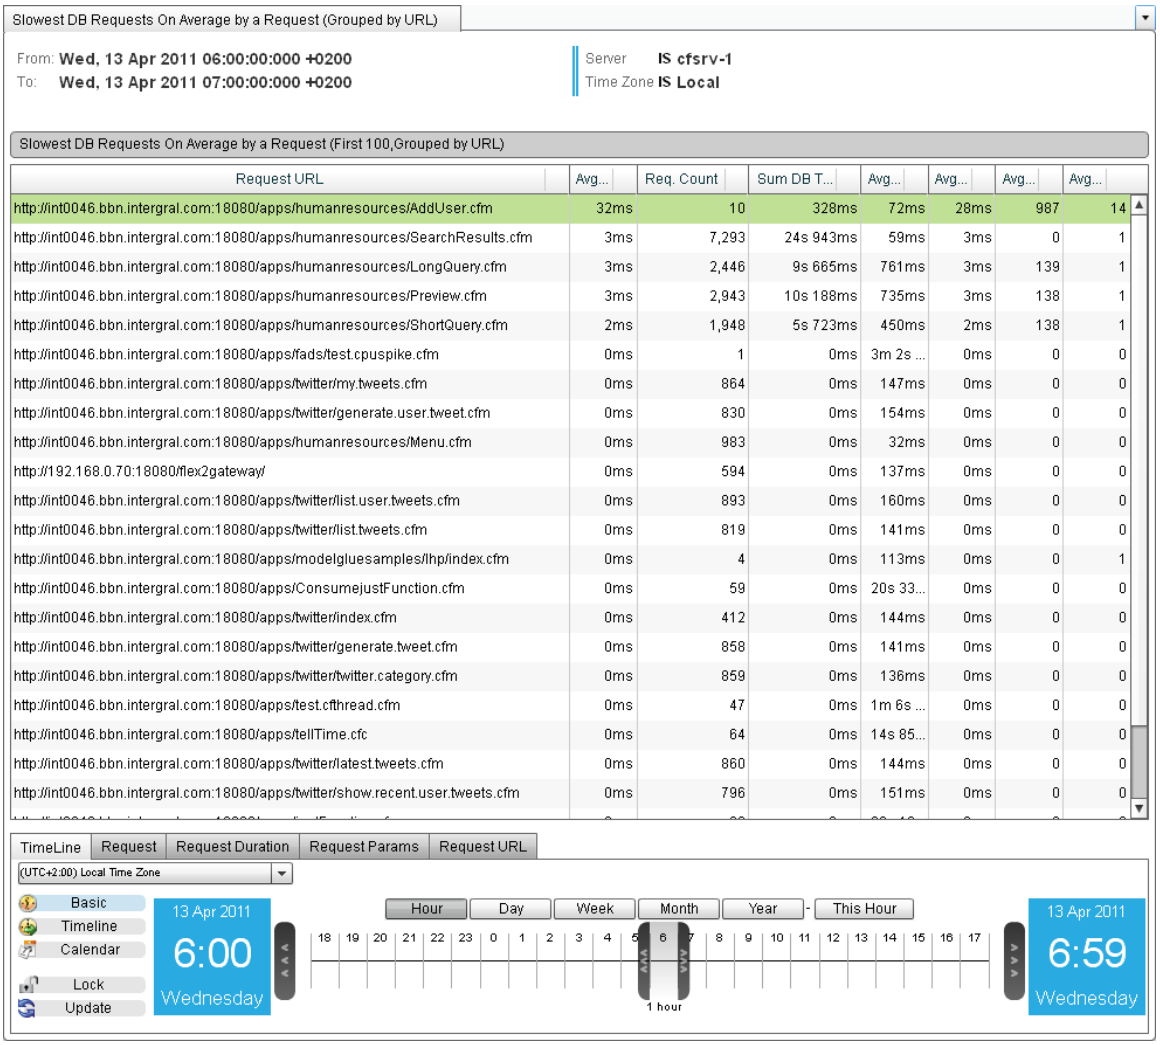

Figure 1: Slowest DB Requests On Average (Hot-Spot) Perspective

Back to the top

Database Hot-Spots

Database Hot-Spots

Overview

Next Steps

Slowest Database Requests (Hot-Spots)

Slowest Database Requests (Hot-Spots)

Slowest Database Requests (Hot-Spots)

Description

This perspective shows a table of the top 100 Slowest Database Requests within the specific period of time.

Usage

To change the order in which the requests are shown, click on one of the table column headings. Use the FA10:Date Navigation] in order to change the period of time over which this perspective covers.

Perspective View

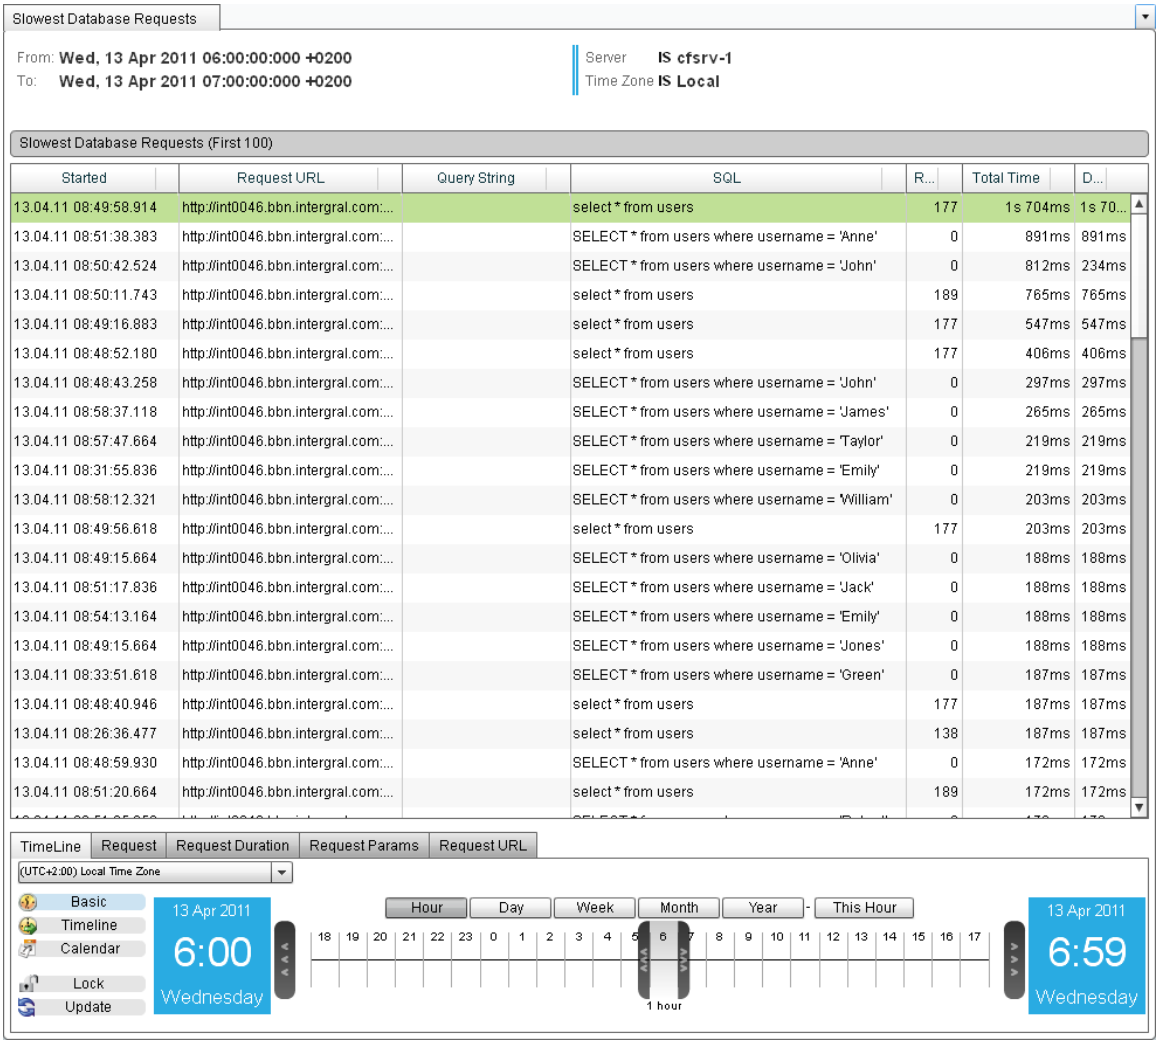

Figure 1: Slowest Database Requests (Hot-Spot) Perspective

Back to the top

ColdFusion Monitor

ColdFusion Monitor

Overview

Next Steps

CF General Statistics

CF Scope Sizes

CF Hit Counts

CF DB Pool Stats

CF General Statistics

CF General Statistics

Description

This perspective shows a chart of the CF General Statistics within the specific period of time.

Usage

Perspective View

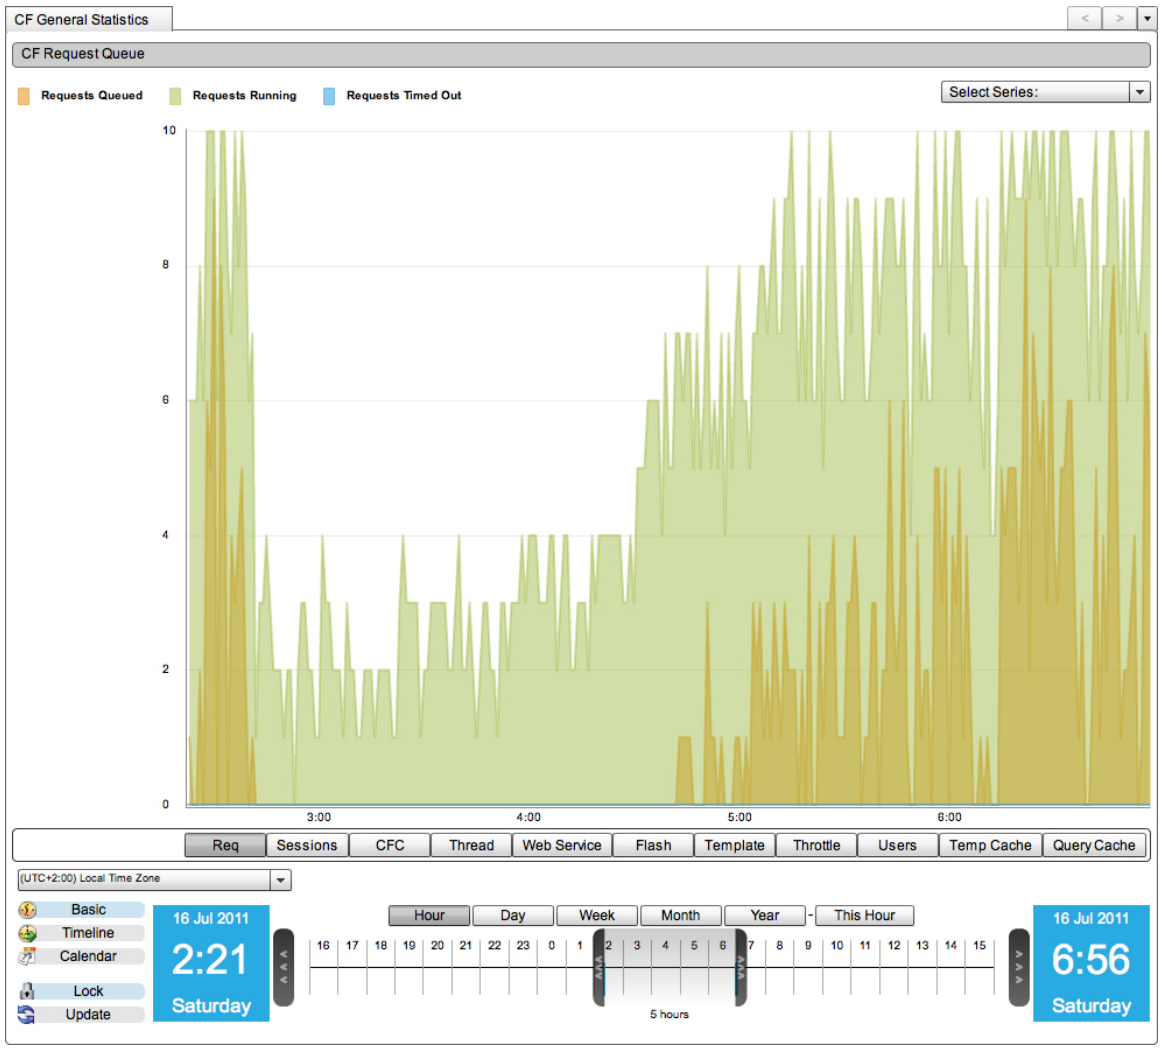

Figure 1: CF General Statistics Perspective

Back to the top

CF Scope Sizes

CF Scope Sizes

Description

This perspective shows a chart of the CF Scope Sizes within the specific period of time.

Usage

Perspective View

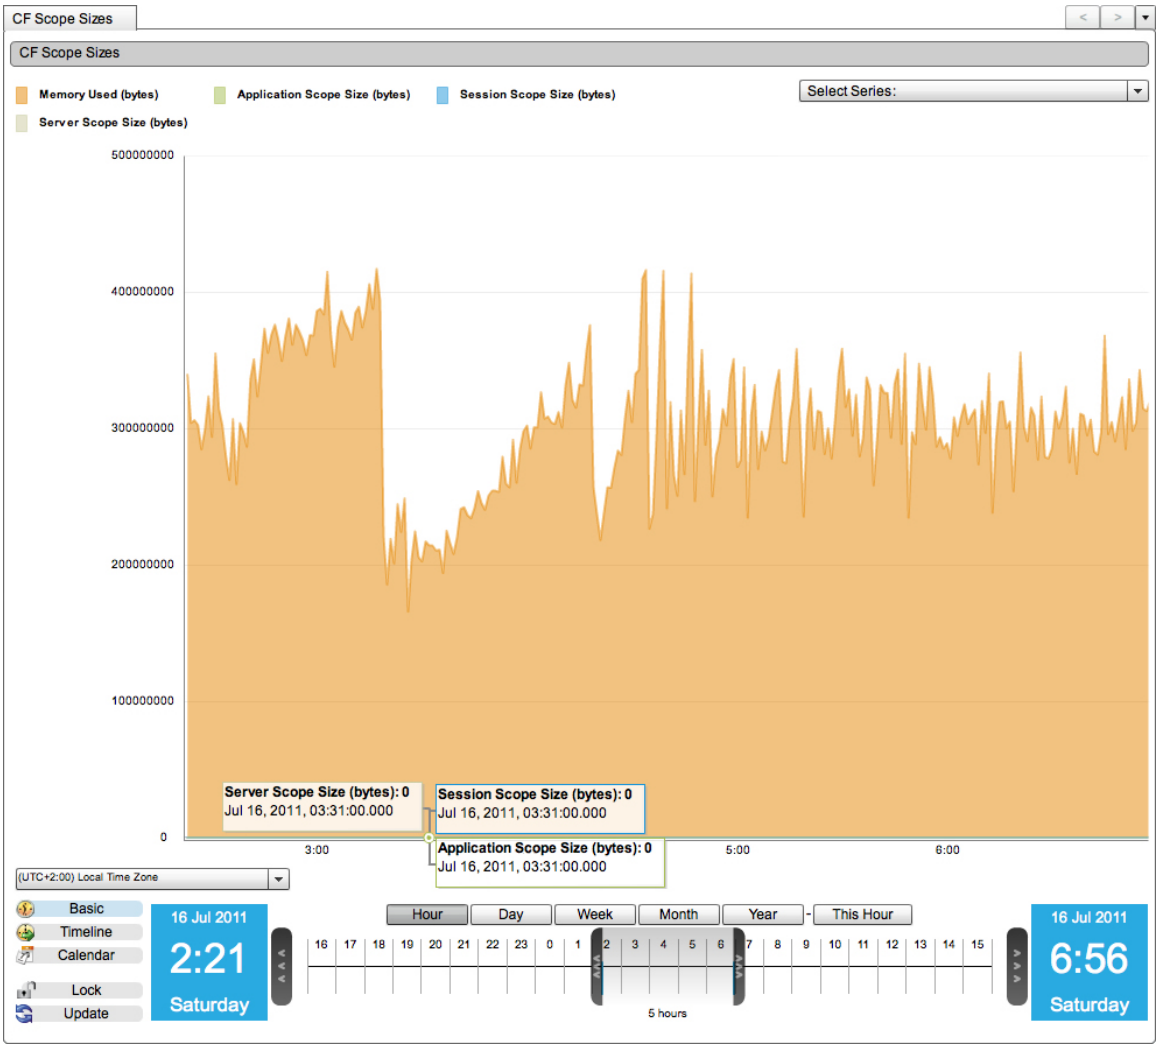

Figure 1: CF Scope Sizes Perspective

Back to the top

CF Hit Counts

CF Hit Counts

Description

This perspective shows a chart of the CF Hit Counts within the specific period of time.

Usage

Perspective View

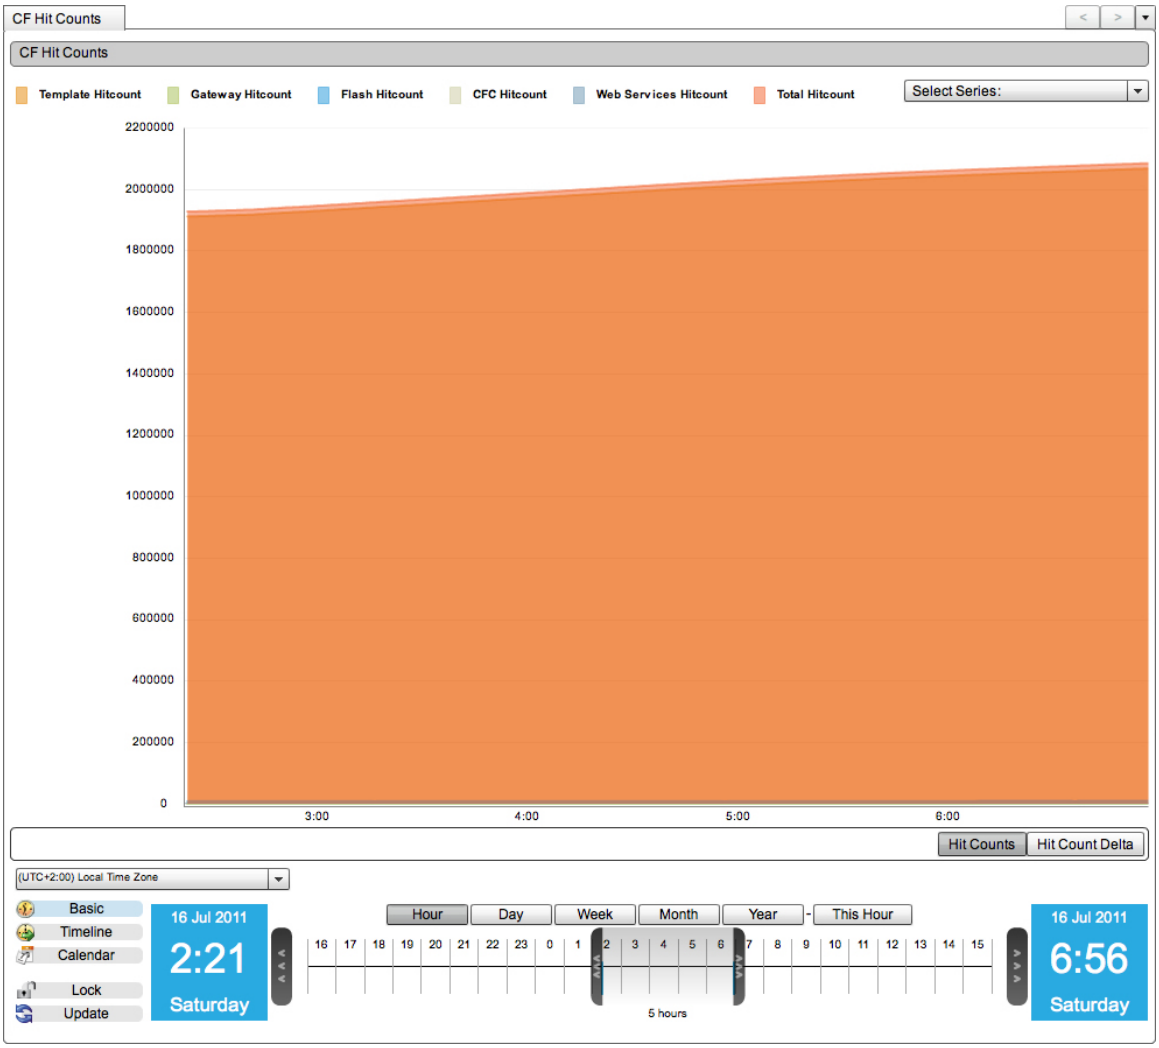

Figure 1: CF Hit Counts Perspective

Back to the top

CF DB Pool Stats

CF DB Pool Stats

Description

This perspective shows a chart of the CF DB Pool Stats Usage within the specific period of time.

Usage

Perspective View

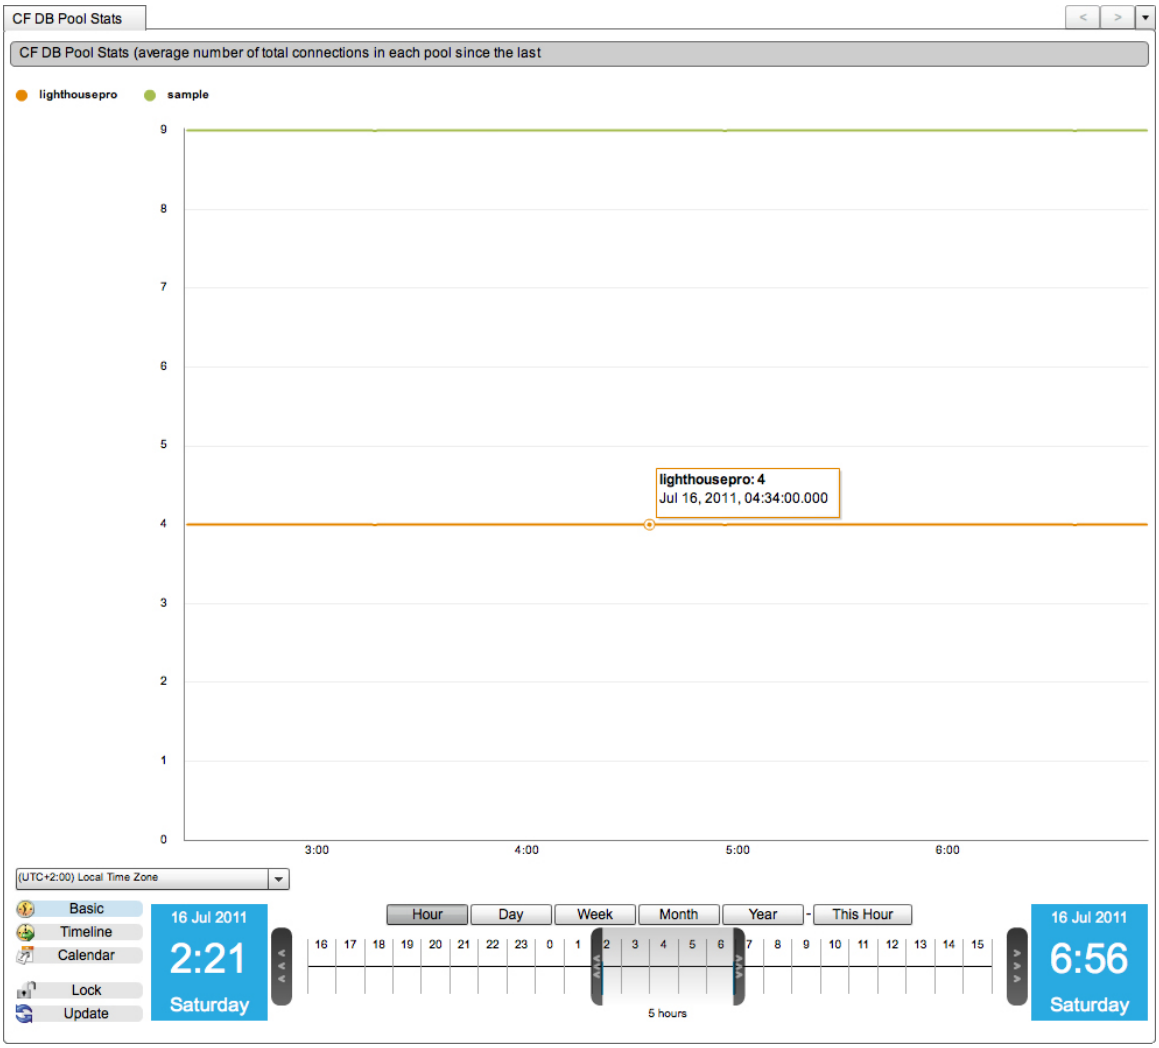

Figure 1: CF DB Pool Stats Perspective

Back to the top

Frameworks

Frameworks

Back to the top

Applications

Applications

Back to the top

FusionAnalytics Reporting

FusionAnalytics Reporting

Overview

The reporting feature of FusionAnalytics allows you to receive daily reports on the status of your server.

There are 2 different types of reports available: TAP Reports and Daily Status Reports.

> If you want to set up this feature as soon as possible, see the Quick Start Guide.

TAP Reports

The TAP Reports provide a daily application score and trend analysis, which is based on your application's Traffic, Availability and Performance (TAP).

They are designed to be less technical and are aimed at application and program managers, business owners and solution stake holders (Figure 1).

Click here to learn more about the TAP Reports.

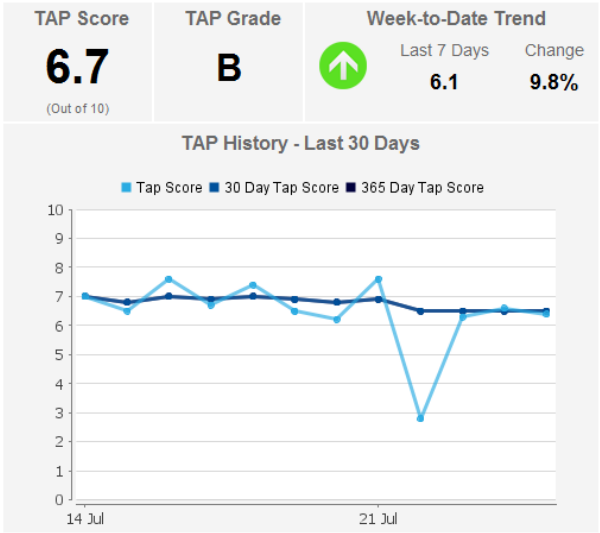

Figure 1: Part of a TAP Report

Daily Status Reports

The Daily Status Reports are intended to be used by technical developers, project leads, project managers and system administrators.

They provide a complete breakdown of your application performance, from slowest requests to system load statistics (Figure 2).

Click here to learn more about the Daily Status Reports.

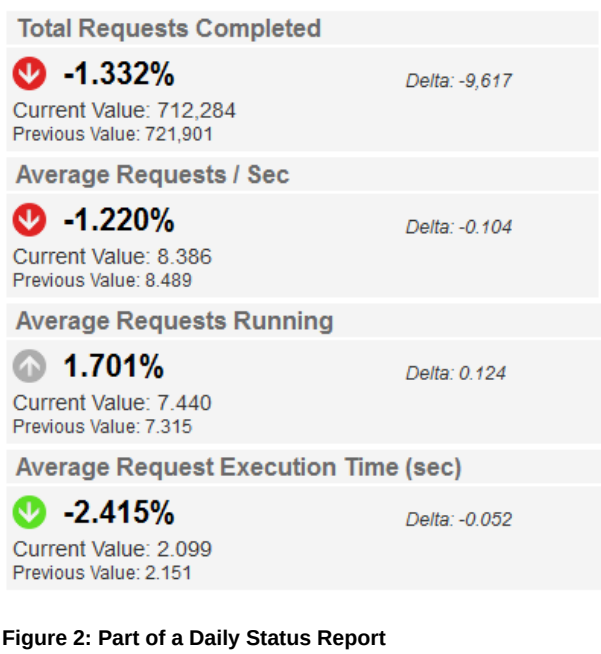

Figure 2: Part of a Daily Status Report

Next Steps

TAP Reports

- What is TAP?
- Running a TAP Report
- TAP Report Provider Arguments
- Understanding the TAP Report

Daily Status Reports

- Running a Daily Status Report
- Daily Status Report Provider Arguments
- Understanding the Daily Status Report

TAP Reports:

Running a TAP Report

TAP Report Provider Arguments

Understanding the TAP Report

Daily Status Reports:

Running a Daily Status Report

Daily Status Report Provider Arguments

Understanding the Daily Status Report

Reports straight to your inbox
Make sure you set your reports up to receive them as emails!

See the Setting Up TAP and Daily Status Reports quick guide or the **Provider Arguments** guide in each report section.

TAP Reports

TAP Reports

Overview

The TAP Reports give you a less technical view of how your server is performing (Figure 1).

Every day you will get a TAP Score and Grade for how well your server is doing. It will also compare this to your previous scores.

This score gives you a standardized comparison of applications regardless of function or hardware capability.

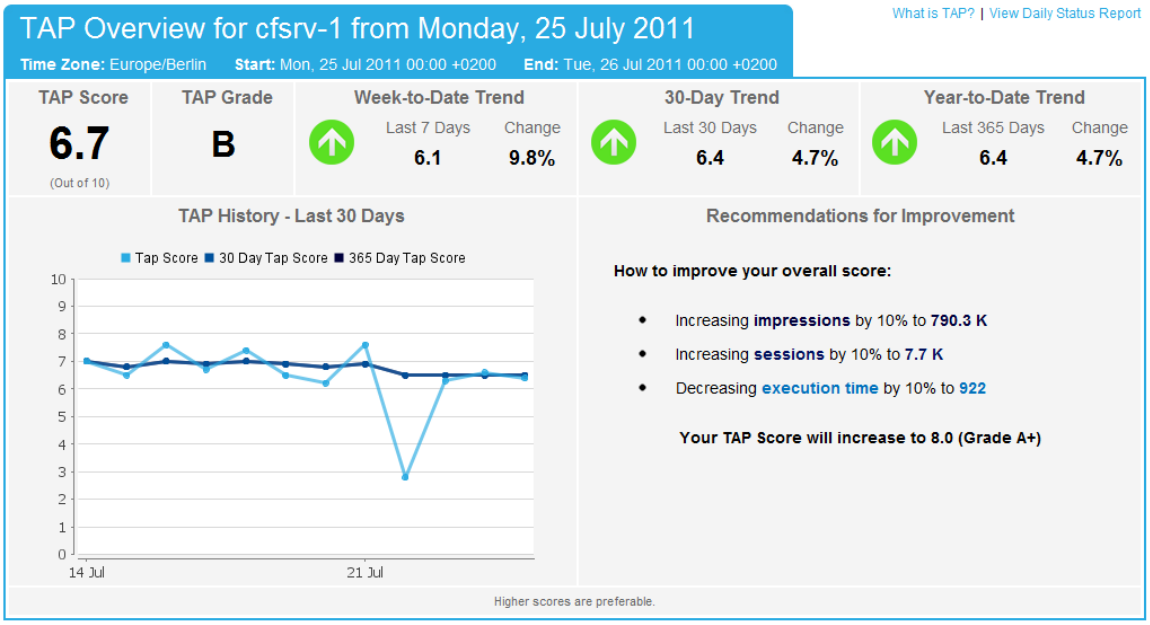

Figure 1: TAP Report Overview

If you are looking for a quick start guide for getting your reports up and running as soon as possible, see Setting Up TAP and Daily Status Reports in the Quick Start Guide.

Find out more

What is TAP?
An explanation of how the TAP Score is calculated.

Running a TAP Report
How to run a TAP Report and where to view it after it has run.

TAP Report Provider Arguments
Different arguments you can use to customize your reports.

Understanding the TAP Report
What each section of the TAP Report means and how you can use it to improve your server.

Next Steps

What is TAP?

Running a TAP Report

TAP Report Provider Arguments

Understanding the TAP Report

What is TAP?

What is TAP?

TAP is a combination of your server's Traffic, Availability and Performance.

Traffic is measured by the number of impressions and sessions.

Availability is calculated from the amount of uptime and number of server restarts.

Performance is measured by the average execution time.

TAP Score

The TAP Score is measured from 0 (worst score) to 10 (best score).

There are 5 sections, awarding up to 2 points each.

Back to the top

Impressions

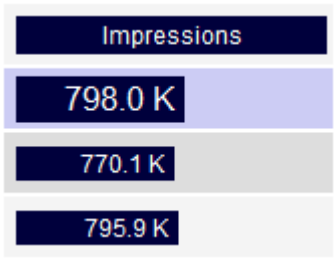	Calculated on a normalized scale from 365 days of historic data and then scaled between -4 and +4 standard deviations from the mean. • +4 standard deviations = 100% • -4 standard deviations = 0%

Back to the top

Sessions

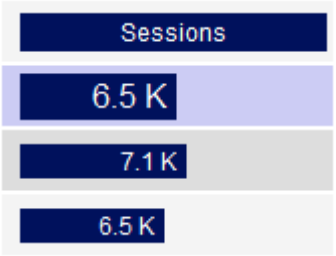	Calculated on a normalized scale from 365 days of historic data and then scaled between -4 and +4 standard deviations from the mean. • +4 standard deviations = 100% • -4 standard deviations = 0%

Back to the top

Uptime

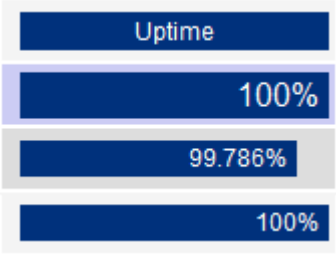	Calculated on a set exponential scale: • No downtime = 100% • 1 minute downtime = 97% • 2 minutes downtime = 94% • 3 minutes downtime = 91% • 4 minutes downtime = 89% • 5 minutes downtime = 86% • 10 minutes downtime = 74% • 20 minutes downtime = 55% • 30 minutes downtime = 41% • 60 minutes downtime = 17%

Back to the top

Restart Count

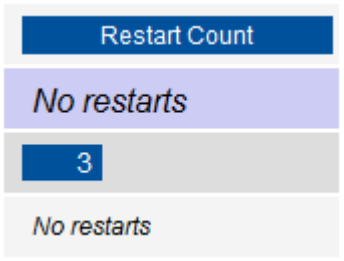	Calculated on a set exponential scale: • 0 restarts = 100% • 1 restart = 90% • 2 restarts = 81% • 3 restarts = 74% • 4 restarts = 67% • 5 restarts = 60% • 10 restarts = 37% • 20 restarts = 14% • 30 restarts = 5%

Back to the top

Execution Time

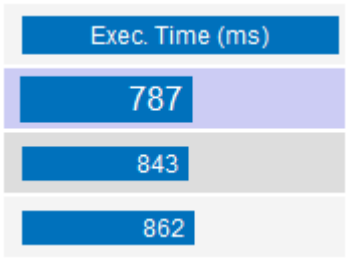	Calculated on a normalized scale from 365 days of historic data and then scaled between -4 and +4 standard deviations from the mean. As lower execution times are better, the scaling used is the inverse of that for impressions and sessions. • -4 standard deviations = 100% • +4 standard deviations = 0%

Back to the top

On This Page:

Next Steps

Running a TAP Report
How to run a TAP Report and where to view it after it has run.

TAP Report Provider Arguments
Different arguments you can use to customize your reports.

Understanding the TAP Report
What each section of the TAP Report means and how you can use it to improve your server.

Running a TAP Report

Running a TAP Report

This section covers several aspects of running a TAP Report. See the links on the right to jump to a specific section.

> If you are looking for a quick guide to getting the reports running, see Setting Up TAP and Daily Status Reports in the Quick Start Guide.

Viewing Scheduled TAP Reports

A TAP Report is automatically scheduled for each FADS application.

> ⚠ The Application must be started for the reports to be able to run.

These tasks can be found by logging into FADS and going to **Scheduled Tasks Application Tasks Summary** (Figure 1).

Figure 1: Application Scheduled Tasks

The table below gives an explanation of each column from the above picture (Figure 1).

Column Name	Description
Task Name	A description of the task; e.g. TAP Report.
Provider Name	The type of task and the application it is running on. Each application has its own TAP Report scheduled.
Duration	How long the task will run for. By default the TAP Reports will keep running indefinitely.
Schedule	When the task will run. It is recommended that the TAP Report runs every day to give the most accurate score.
Last Run	The time and date when the task last ran, either if it ran when scheduled or was manually run.
Status	Shows if the task is Scheduled, Not Scheduled or Running.

Back to the top

Changing When a TAP Report Runs

You can edit the scheduled TAP Report by clicking **Edit** (Figure 2).

This allows you to edit when the TAP Reports will run. The most suitable option is to have them running once a day, so that your receive your reports daily.

You also have the option to change if the task is active and the duration of the task, if needed.

This is also where **Provider Arguments** are entered if needed. See Customizing a TAP Report for more information.

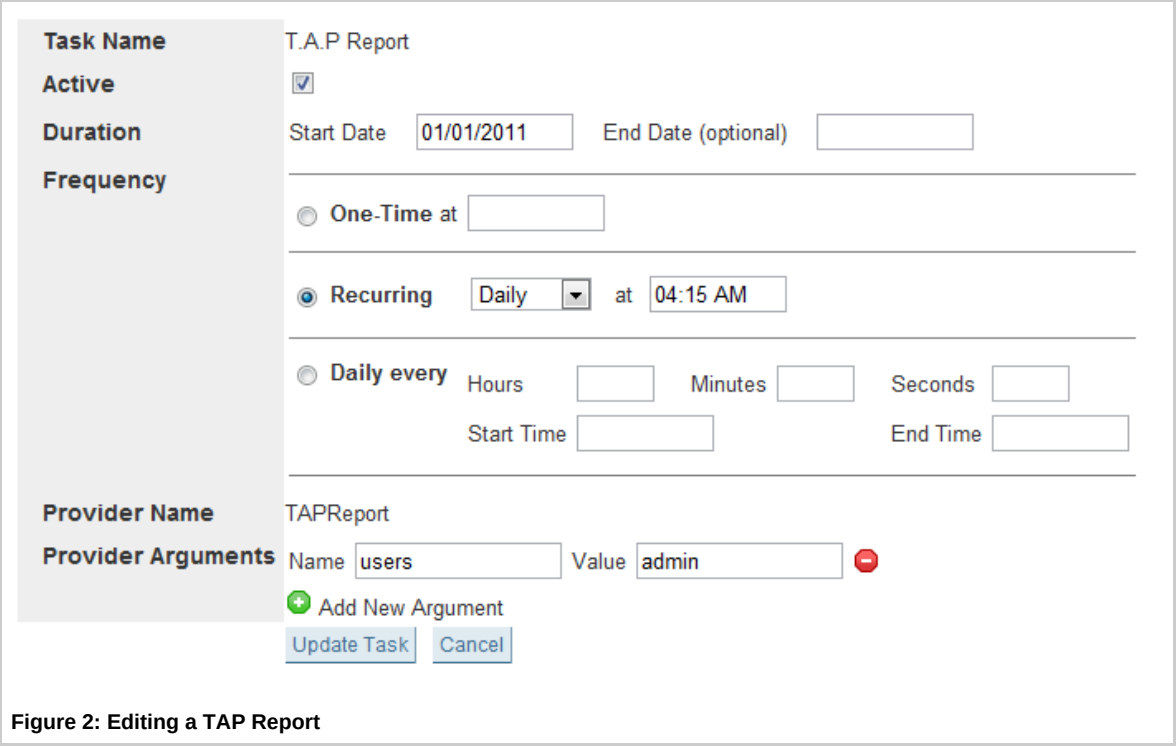

Figure 2: Editing a TAP Report

Back to the top

Running a TAP Report Instantly

You can run a TAP Report at any time by clicking **Run Now** (Figure 3).

This will generate a report for the previous day.

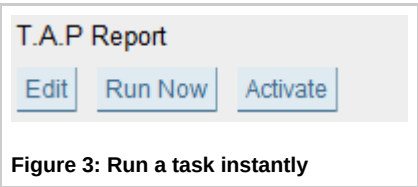

Figure 3: Run a task instantly

The Application must be started for the button to be enabled.

Back to the top

Activating and Deactivating a TAP Report

The TAP Reports in the **Application Tasks Summary** are activated by default.

If you need to deactivate TAP Reports for a particular application, you can click the **Deactivate** button (Figure 4).

You can start it again by clicking **Activate**.

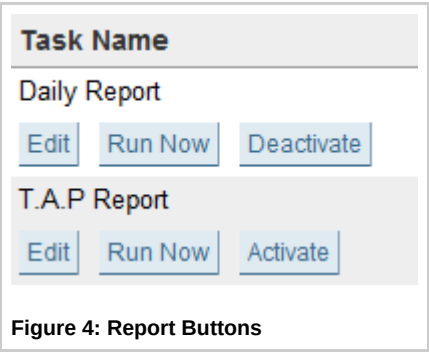

Figure 4: Report Buttons

Back to the top

Customizing a TAP Report

You can use **Provider Arguments** to customize your TAP Reports to meet your needs.

For example:

- sending reports as emails
- using a different timezone
- using HTTPS for links and images

See TAP Report Provider Arguments for more information.

Back to the top

Viewing a TAP Report

After a TAP Report has run you can view it in the client, under **Business Metrics Reports** (Figure 5).

You can double click on a report to open it.

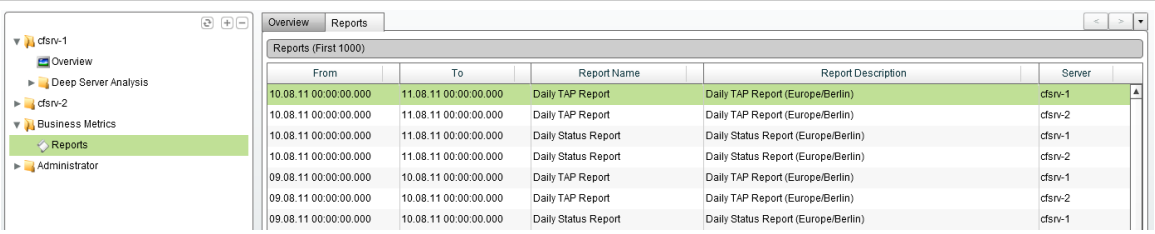

Figure 5: Reports Perspective

You can also use the **Report Filter** option, found next to the **Timeline** tab at the bottom of the screen, to find specific reports (Figure 6).

For help using the **Report Filter**, see the Reports section.

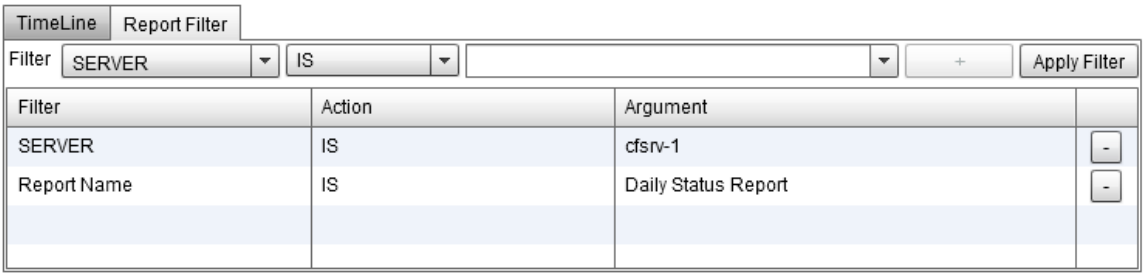

Figure 6: Report Filters

You can also configure your TAP Reports so that you can view them as emails. See the ***users*** argument in TAP Report Provider Arguments.

Back to the top

On this page:

Next Steps

What is TAP?
An explanation of how the TAP Score is calculated.

TAP Report Provider Arguments
Different arguments you can use to customize your reports.

Understanding the TAP Report
What each section of the TAP Report means and how you can use it to improve your server.

Take me back to **FusionAnalytics Reporting**.

TAP Report Provider Arguments

TAP Report Provider Arguments

> **All of these arguments are optional**
> The reports will run without any of these arguments being provided, but they can be used to enhance your reports to suit your needs.

Argument Name	Default Value	Description	Example Values
date	Previous day's date	**The date that the report will run for.** • This will override the previous day's date that is automatically provided by the scheduled task. • The value must be a UTC time in milliseconds. • The value that is entered will be rounded to midnight (the start of the day) in the specified timezone and will become the end date of the report.	1312243200000 *(Report Date: 1st-2nd August GMT)* 1310281200000 *(Report Date: 9th-10th August GMT)*
https	False	**Defines if HTTP or HTTPS will be used.** • HTTP / HTTPS is used for links and images. • The same setting will apply to reports that are emailed to users. • If the argument is not provided, HTTP is always used.	true

previousDays	False	**Collects any data that is available for previous days.**
		• If you have data from previous days imported into FusionAnalytics, but no TAP Reports for those days, you can use this argument to collect the data.
		• Using this argument will only produce one report, but the data for the previous days will be filled in.
		• **Warning:** You should not use this every time you run a TAP Report:
		  ◦ Data from previous TAP reports can automatically be retrieved, without this argument.
		  ◦ If you have a large amount of previous data, it may also take a long time to run.
			true

timezone	DataCollector's timezone	**Used to change the default timezone of the report.** • This will change the start and end dates of the report to midnight in the specified timezone. • Any valid Java timezone can be used. • If the timezone is misspelt, it will revert back to the timezone of the DataCollector. See Applications (FADC) for more information.	America/New_York Europe/Berlin GMT
users	*None*	**Used to email reports to the specified users.** • Users **must** have an email address on their account. See Users Summary (FADS) for more information. • The Mail Server **must** also be configured. See Server Settings for more information. • Multiple users can be added by separating the usernames with a comma (no space).	admin admin,user1,user2

Adding Provider Arguments

To add your Provider Arguments, you will need to go to the **Application Tasks Summary** in FADS and click the **Edit** button on the TAP Report task.

In the **Provider Arguments** section, you will see fields for the Name and Value (Figure 1).

You can add multiple arguments by clicking **Add New Argument**.

You can remove arguments by clicking .

Click **Update Task** when you have finished editing the task.

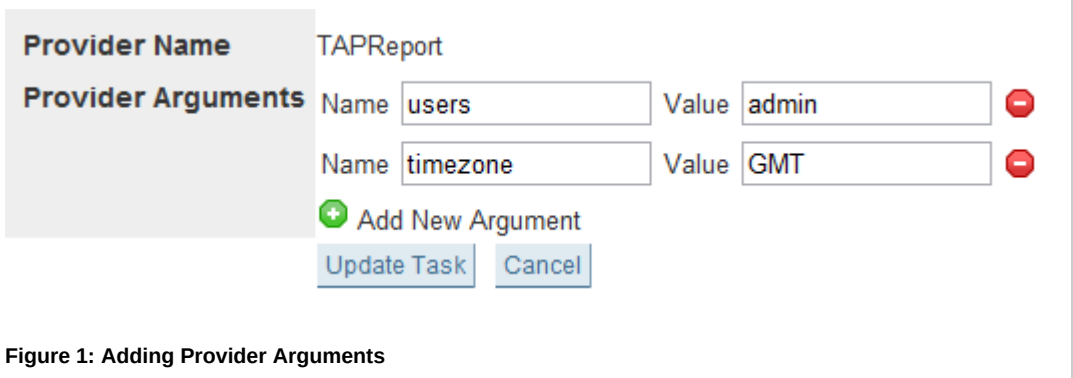

Figure 1: Adding Provider Arguments

Next Steps

What is TAP?
An explanation of how the TAP Score is calculated.

Running a TAP Report
How to run a TAP Report and where to view it after it has run.

Understanding the TAP Report
What each section of the TAP Report means and how you can use it to improve your server.

Take me back to **FusionAnalytics Reporting**.

Understanding the TAP Report

Understanding the TAP Report

This page will help you to understand the data in your TAP Report. Please see the links to the right to jump to a particular section.

In the top right corner of the report you can find 3 links:

Link Name	Description
What is TAP?	This will take you to an online description of what TAP is and how it is calculated.
View Daily Status Report	This is only displayed if there is also a Daily Status Report for the same DataCollector, date and timezone.
View in FusionAnalytics	If you are viewing the report as an email, you can click this to take you directly to this application in FusionAnalytics.

On this page:

Overview

The Overview header displays the following information (Figure 1):

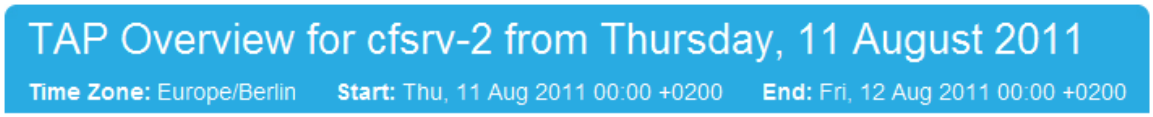

Figure 1: TAP Overview section header example

The table below gives an explanation of what each part of the header means:

Header Information	Description
DataCollector *e.g. cfsrv-2*	The DataCollector that the data belongs to.
Day *e.g. Thursday, 11 August 2011*	The date that the report is run for (from midnight to midnight of that day).
Time Zone	The default time zone is the one set on the DataCollector. See TAP Report Provider Arguments if you want to change the time zone of the report.
Start	The exact date and time that the data in the report starts from. It also includes the offset from UTC time *(e.g. +0200)*.
End	The exact date and time that the data in the report ends at. It also includes the offset from UTC time *(e.g. +0200)* .

Clicking on the Overview header will take you to a copy of the TAP Report in the Web Client.

The Overview contains the following sections:

Name	Example	Description
TAP Score	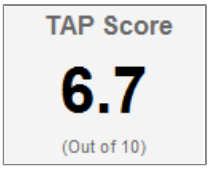	A value between 0.0 (worst score) and 10.0 (best score), this is the score that your server has received for that day. For more information on how the TAP Score is calculated, see What is TAP? .
TAP Grade	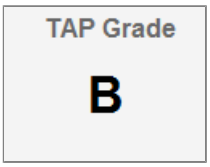	Related to the TAP Score, it is the grade that corresponds with your score: • 9.0 and above = A++ • 8.0 to 8.9 = A+ • 7.0 to 7.9 = A • 6.0 to 6.9 = B • 5.0 to 5.9 = C • 4.0 to 4.9 = D • 3.0 to 3.9 = E • 2.9 and below = F
Week-to-Date Trend	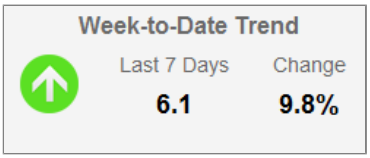	The Last 7 Days score is an average of the TAP scores from today and the 6 previous days. The arrow indicates if today's score has improved, decreased or stayed the same compared to the week's average. Change gives the percentage change from the Last 7 Days average and today's score.
30-Day Trend		The Last 30 Days score is an average of the TAP scores from today and the 29 previous days. The arrow indicates if today's score has improved, decreased or stayed the same compared to the month's average. Change gives the percentage change from the Last 30 Days average and today's score.

Year-to-Date Trend		The Last 365 Days score is an average of the TAP scores from today and the 364 previous days. The arrow indicates if today's score has improved, decreased or stayed the same compared to the year's average. Change gives the percentage change from the Last 365 Days average and today's score.
TAP History	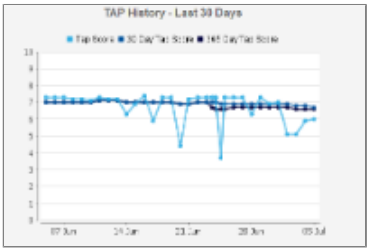 *(click to enlarge)*	The TAP History graph shows a maximum of 30 days. There are 3 lines, showing: • The TAP score for each day • The 30 day average TAP scores • The 365 day average TAP scores ⚠ Until you have more than 30 days of TAP data, the 30 and 365 day averages will be the same.
Recommendations for Improvement	 *(click to enlarge)*	The Recommendations give ideas on how you can improve your score. It will only suggest improvement in areas that can actually be improved (e.g. if Uptime is already 100%, improvement will not be suggested). You can also see the score you will achieve by making all of the improvements.

Back to the top

Example Overview

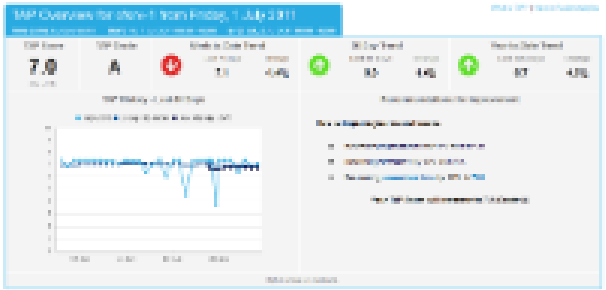

(click to enlarge)

Last 7 Days

The top row of the Last 7 Days table is highlighted; this is the data that has been generated for the day the report has been run (i.e. the previous day).

Last 7 Days contains the following sections:

Name	Example	Description
Date *e.g. Friday, 1 July 2011*	Friday, 1 July 2011 Thursday, 30 June 2011 Wednesday, 29 June 2011	The top row, highlighted in purple, is the date that the report has been generated for. The next rows show the previous 7 days, going back to the same day of the previous week (e.g. Friday 24 July 2011).

TAP Score	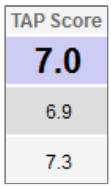	The TAP Score that your server received for each day.
30-Day Trend		This compares the TAP Score from each day with the 30-day Average score from each day. A green arrow indicates your server has improved on the 30 day average; a red arrow indicates your score has decreased and a grey arrow indicates no change.
30-Day Average	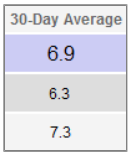	This is than average of TAP Scores from the past 30 days. It is the same as the 30-Day Trend from the Overview section.
Impressions *(Traffic)*	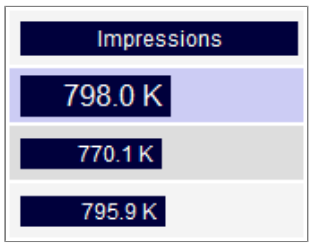	Impressions are the total number of hits received during the start and end date of the report. For Impressions, higher values are preferable. The Impressions section will receive a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the longer the bar, the better the score.
Sessions *(Traffic)*	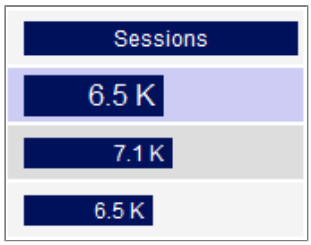	Sessions are the total number of sessions between the start and end date of the report. For Sessions, higher values are preferable. The Sessions section will receive a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the longer the bar, the better the score.
Uptime *(Availability)*	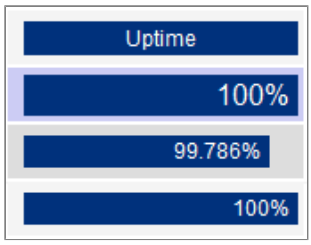	Uptime is the percentage of time that your server has been up between the start and end date of the report. If your Uptime is not 100%, you can hover your mouse over the value to see the exact amount of time your server was down. For Uptime, higher values are preferable. The Uptime section will recieve a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the longer the bar, the better the score.
Restart Count *(Availability)*	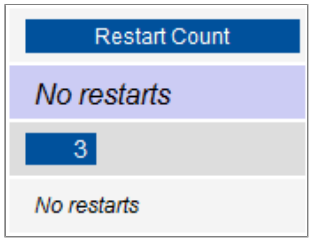	Restart Count is the total number of server restarts between the start and end date of the report. For Restart Count, lower values are preferable. The Restart Count section will receive a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the shorter the bar, the better the score.
Exec. Time *(Performance)*	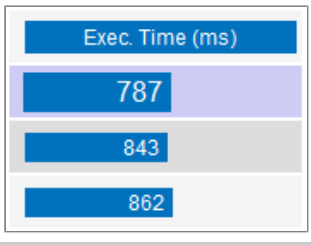	Exec. Time is the average execution time of requests between the start and end date of the report. For Exec. Time, lower values are preferable. The Exec. Time section will receive a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the shorter the bar, the better the score.

Back to the top

Example Last 7 Days

(click to enlarge)

Weekly Report

Name	Example	Description
Monday to Sunday	27 June 2011 - 3 July 2011 20 June 2011 - 26 June 2011 13 June 2011 - 19 June 2011	Each row is one week; the dates are from Monday to Sunday. The top row includes the date that the report has been generated for.
7-Day Average	7-Day Average 7.1 6.6 6.9	The 7-Day Average is an average of the TAP scores from that week (Monday to Sunday). ⚠ If the current week has not yet finished, the top row will show the average from Monday to the current day.
30-Day Trend	30-Day Trend ↑ ↓ ↓	This compares the TAP Score from each day with the 30-day Average score from each day. A green arrow indicates your server has improved on the 30 day average; a red arrow indicates your score has decreased and a grey arrow indicates no change.
30-Day Average	30-Day Average 6.9 6.9 7.0	This is than average of TAP Scores from the past 30 days. It is the same as the 30-Day Trend from the Overview section.
7-Day Average TAP Score	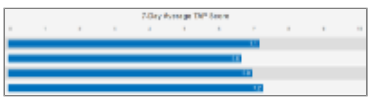 *(click to enlarge)*	This shows the 7-Day average as a bar graph, allowing you to easily compare which weeks had better TAP scores on average.

Back to the top

Example Weekly Report
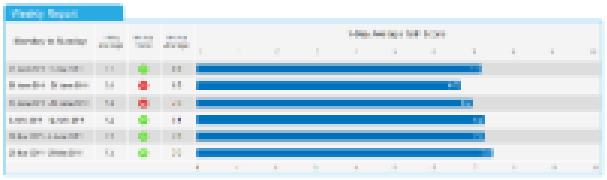
(click to enlarge)

Day of Week

Name	Example	Description
Date *e.g. Tuesday, 28 June 2011*		The top row, highlighted in purple, is the date that the report has been generated for. The next rows show the same day of the week, going back over 5 weeks.
TAP Score	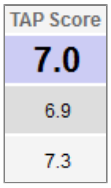	The TAP Score that your server received for each day.
Last 4 *Day* Trend *e.g. Last 4 Tuesday Trend*	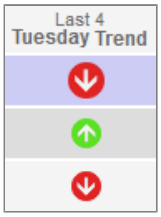	This compares the TAP Score from each day with the Last 4 *Day* Trend score from each day. A green arrow indicates your server has improved from the trend; a red arrow indicates your score has decreased and a grey arrow indicates no change.
Last 4 *Day* Average *e.g. Last 4 Tuesday Average*	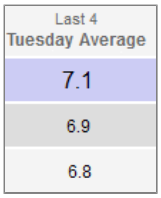	This is an average of the scores from the last 4 rows (the same day of each week).
Impressions *(Traffic)*	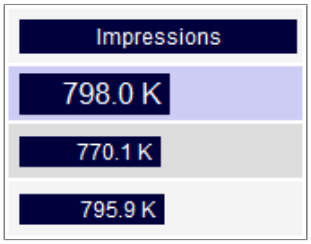	Impressions are the total number of hits received during the start and end date of the report. For Impressions, higher values are preferable. The Impressions section will receive a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the longer the bar, the better the score.
Sessions *(Traffic)*	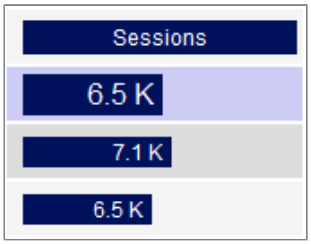	Sessions are the total number of sessions between the start and end date of the report. For Sessions, higher values are preferable. The Sessions section will receive a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the longer the bar, the better the score.

Uptime *(Availability)*	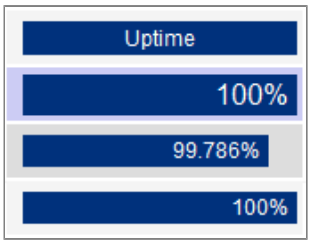	Uptime is the percentage of time that your server has been up between the start and end date of the report. If your Uptime is not 100%, you can hover your mouse over the value to see the exact amount of time your server was down. For Uptime, higher values are preferable. The Uptime section will recieve a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the longer the bar, the better the score.
Restart Count *(Availability)*	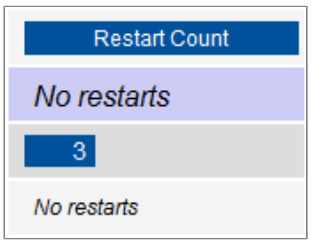	Restart Count is the total number of server restarts between the start and end date of the report. For Restart Count, lower values are preferable. The Restart Count section will receive a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the shorter the bar, the better the score.
Exec. Time *(Performance)*	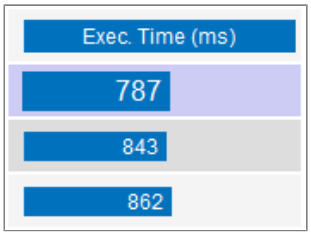	Exec. Time is the average execution time of requests between the start and end date of the report. For Exec. Time, lower values are preferable. The Exec. Time section will receive a score between 0.0 - 2.0 (see What is TAP? for more information). The size of the bar is measured on this score; the shorter the bar, the better the score.

Back to the top

Example Day of Week

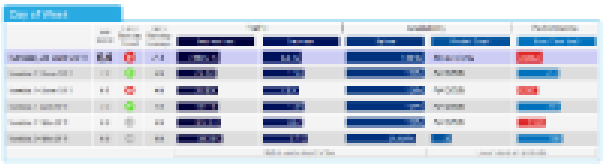

(click to enlarge)

Next Steps

What is TAP?
An explanation of how the TAP Score is calculated.

Running a TAP Report
How to run a TAP Report and where to view it after it has run.

TAP Report Provider Arguments
Different arguments you can use to customize your reports.

Daily Status Reports

Daily Status Reports

Overview

The Daily Status Reports provide detailed information about the requests, database activity, resource usage and even a breakdown of the individual applications running.

Each day you can see how your server has done compared to the previous day.

This allows you to easily analyze which areas need improvement, helping you to improve the performance of your server.

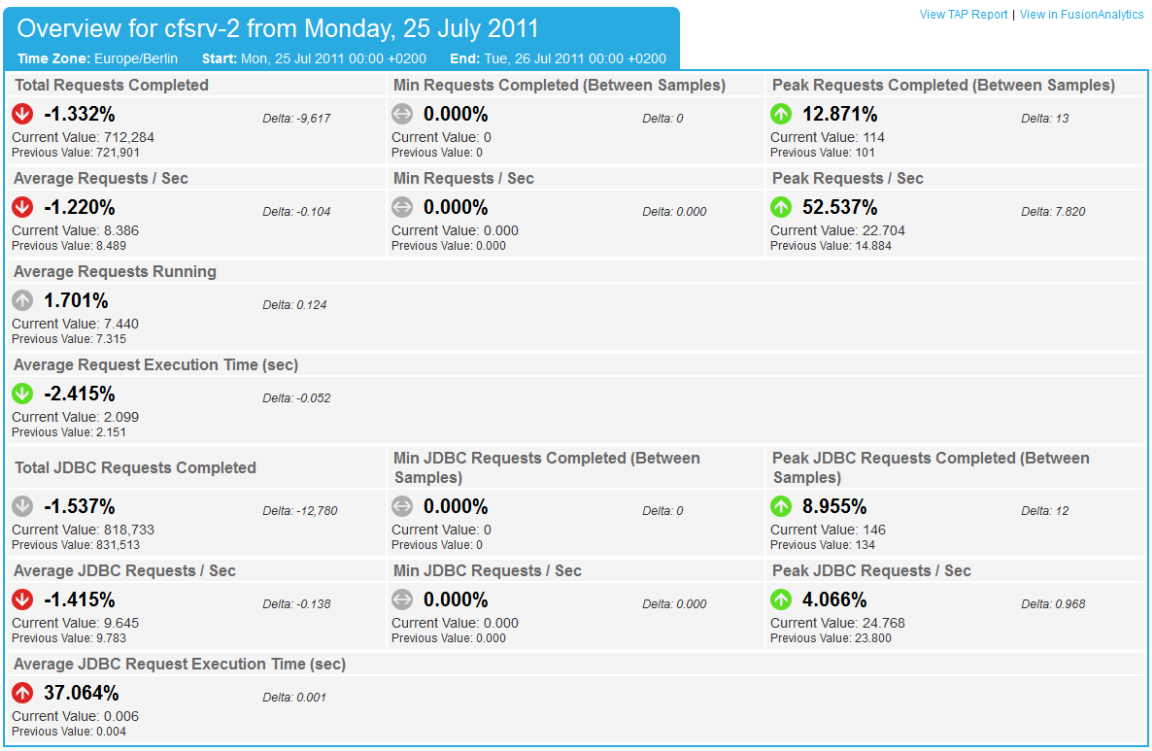

Figure 1: Daily Status Report

> If you are looking for a quick start guide for getting your reports up and running as soon as possible, see Setting Up TAP and Daily Status Reports in the Quick Start Guide guide.

Find out more

Running a Daily Status Report
How to run a Daily Status Report and where to view it after it has run.

Daily Status Report Provider Arguments
Different arguments you can use to customize your reports.

Understanding the Daily Status Report
What each section of the Daily Status Report means and how you can use it to improve your server.

Next Steps

Running a Daily Status Report

Daily Status Report Provider Arguments

Understanding the Daily Status Report

Running a Daily Status Report

Running a Daily Status Report

This section covers several aspects of running a Daily Status Report. See the links on the right to jump to a specific section.

> If you are looking for a quick guide to getting the reports running, see Setting Up TAP and Daily Status Reports in the Quick Start Guide guide.

Viewing Scheduled Daily Status Reports

A Daily Status Report is automatically scheduled for each FADS application.

⚠ The Application must be started for the reports to be able to run.

These tasks can be found by logging into FADS and going to **Scheduled Tasks Application Tasks Summary** (Figure 1).

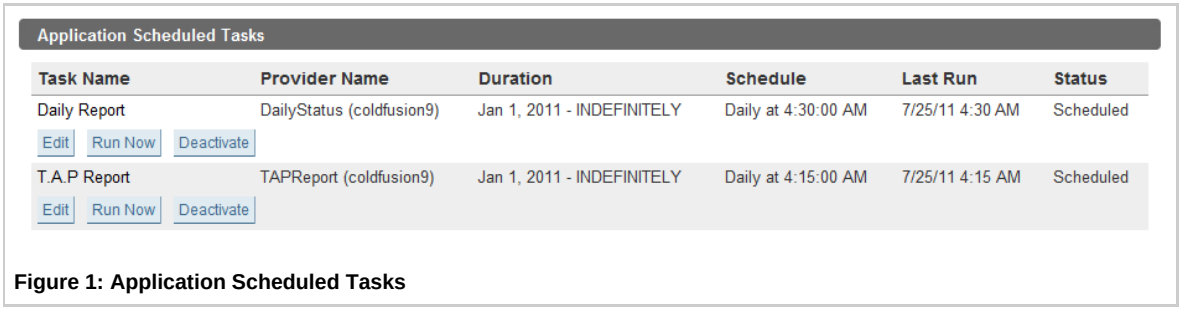

Figure 1: Application Scheduled Tasks

The table below gives an explanation of each column from the above picture (Figure 1).

Column Name	Description
Task Name	A description of the task; e.g. Daily Status Report.
Provider Name	The type of task and the application it is running on. Each application has its own Daily Status Report scheduled.
Duration	How long the task will run for. By default the Daily Status Reports will keep running indefinitely.
Schedule	When the task will run. It is recommended that you keep the Daily Status Report running every day.
Last Run	The time and date when the task last ran, either if it ran when scheduled or was manually run.
Status	Shows if the task is Scheduled, Not Scheduled or Running.

Back to the top

Changing When a Daily Status Report Runs

You can edit the scheduled Daily Status Report by clicking **Edit** (Figure 2).

This allows you to edit when the Daily Status Reports will run. The most suitable option is to have them running once a day, so that your receive your reports daily.

You also have the option to change if the task is active and the duration of the task, if needed.

This is also where **Provider Arguments** are entered if needed. See Customizing a Daily Status Report for more information.

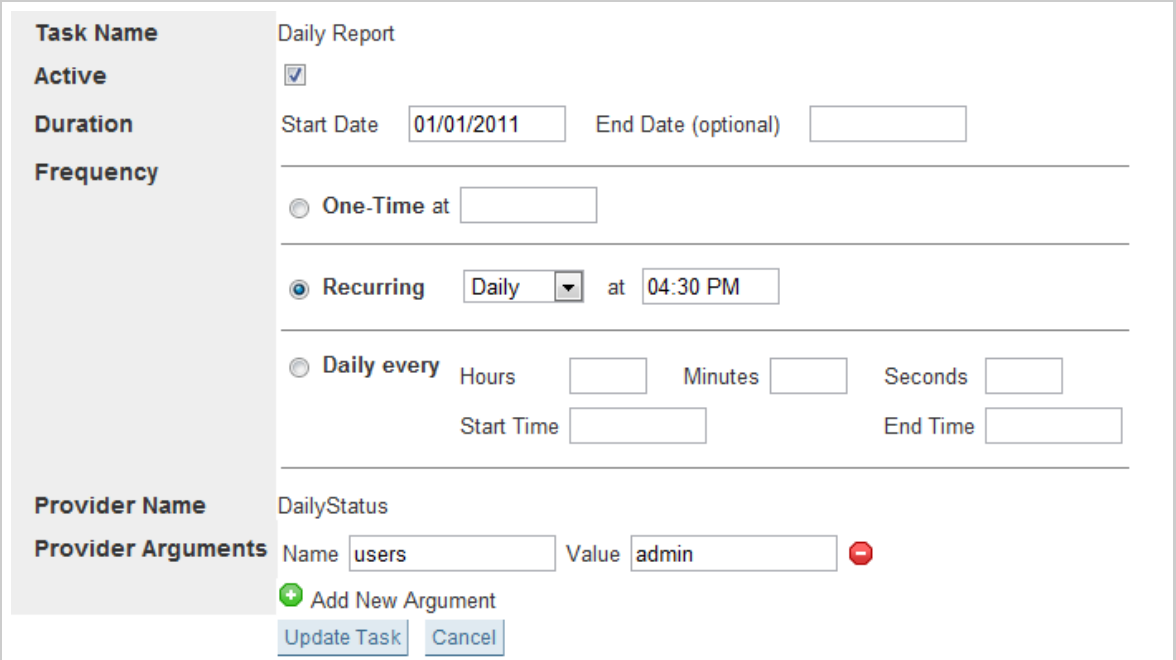

Figure 2: Editing a Daily Status Report

Back to the top

Running a Daily Status Report Instantly

You can run a Daily Status Report at any time by clicking **Run Now** (Figure 3).

This will generate a report for the previous day.

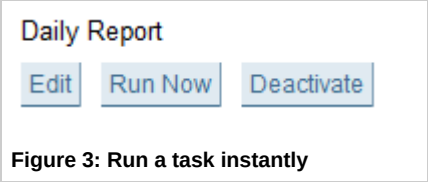

Figure 3: Run a task instantly

The Application must be started for the button to be enabled.

Back to the top

Activating and Deactivating a Daily Status Report

The Daily Status Reports in the **Application Tasks Summary** are activated by default.

If you need to deactivate the Daily Status Reports for a particular application, you can click the **Deactivate** button (Figure 4).

You can start it again by clicking **Activate**.

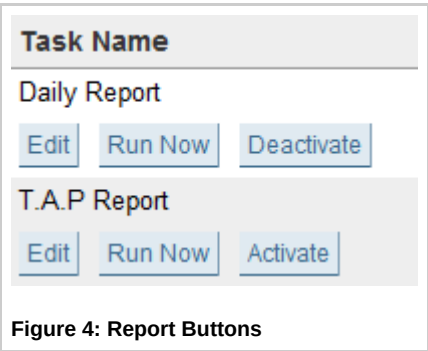

Figure 4: Report Buttons

Back to the top

Customizing a Daily Status Report

You can use **Provider Arguments** to customize your Daily Status Reports to meet your needs.

For example:

- sending reports as emails
- using a different timezone
- using HTTPS for links and images

See Daily Status Report Provider Arguments for more information.

Back to the top

Viewing a Daily Status Report

After a Daily Status Report has run you can view it in the client, under **Business Metrics Reports** (Figure 5).

You can double click on a report to open it.

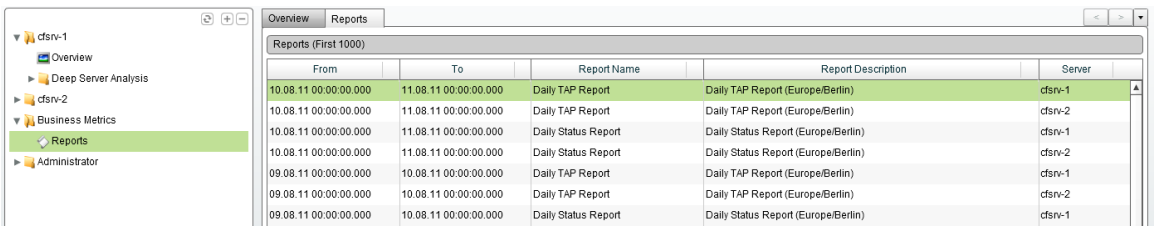

Figure 5: Reports Perspective

You can also use the **Report Filter** option, found next to the **Timeline** tab at the bottom of the screen, to find specific reports (Figure 6).

For help using the **Report Filter**, see the Reports section.

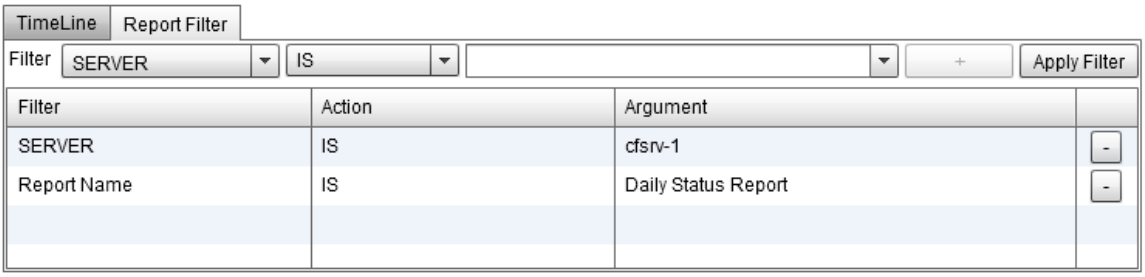

Figure 6: Report Filters

You can also configure your Daily Status Reports so that you can view them as emails. See the ***users*** argument in Daily Status Report Provider Arguments.

Back to the top

On this page:

Next Steps

Daily Status Report Provider Arguments
Different arguments you can use to customize your reports.

Understanding the Daily Status Report
What each section of the Daily Status Report means and how you can use it to improve your server.

Take me back to **FusionAnalytics Reporting**.

Daily Status Report Provider Arguments

Daily Status Report Provider Arguments

> **All of these arguments are optional**
> The reports will run without any of these arguments being provided, but they can be used to enhance your reports to suit your needs.

Argument Name	Default Value	Description	Example Values
date	Previous day's date	**The date that the report will run for.** • This will override the previous day's date that is automatically provided by the scheduled task. • The value must be a UTC time in milliseconds. • The value that is entered will be rounded to midnight (the start of the day) in the specified timezone and will become the end date of the report.	1312243200000 *(Report Date: 1st-2nd August GMT)* 1310281200000 *(Report Date: 9th-10th August GMT)*
https	False	**Defines if HTTP or HTTPS will be used.** • HTTP / HTTPS is used for links and images. • The same setting will apply to reports that are emailed to users. • If the argument is not provided, HTTP is always used.	true
timezone	DataCollector's timezone	**Used to change the default timezone of the report.** • This will change the start and end dates of the report to midnight in the specified timezone. • Any valid Java timezone can be used. • If the timezone is misspelt, it will revert back to the timezone of the DataCollector. See Applications (FADC) for more information.	America/New_York Europe/Berlin GMT
users	*None*	**Used to email reports to the specified users.** • Users **must** have an email address on their account. See Users Summary (FADS) for more information. • The Mail Server **must** also be configured. See Server Settings for more information. • Multiple users can be added by separating the usernames with a comma (no space).	admin admin,user1,user2

Adding Provider Arguments

To add your Provider Arguments, you will need to go to the **Application Tasks Summary** in FADS and click the **Edit** button on the Daily Status Report task.

In the **Provider Arguments** section, you will see fields for the Name and Value (Figure 1).

You can add multiple arguments by clicking **Add New Argument**.

You can remove arguments by clicking .

Click **Update Task** when you have finished editing the task.

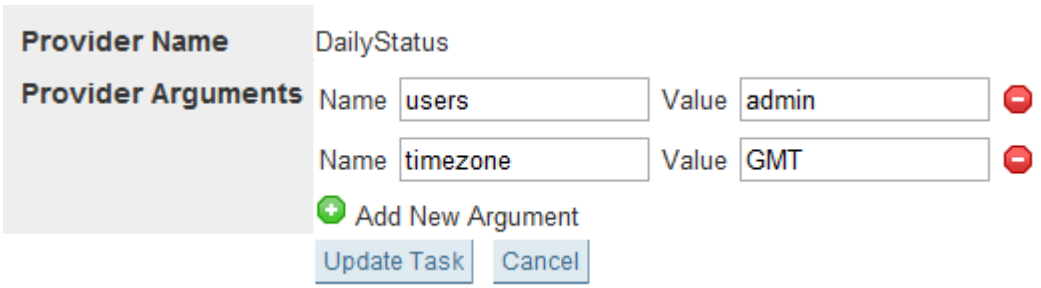

Figure 1: Adding Provider Arguments

Next Steps

Running a Daily Status Report
How to run a Daily Status Report and where to view it after it has run.

Understanding the Daily Status Report
What each section of the Daily Status Report means and how you can use it to improve your server.

Take me back to **FusionAnalytics Reporting**.

Understanding the Daily Status Report

Understanding the Daily Status Report

This page will help you to understand the data in your Daily Status Report. Please see the links to the right to jump to a particular section.

In the top right corner of the report you can find 2 links:

Link Name	Description
View TAP Report	This is only displayed if there is also a TAP Report for the same DataCollector and date.
View in FusionAnalytics	If you are viewing the report as an email, you can click this to take you directly to this application in FusionAnalytics.

This report uses arrow indicators to allow you to easily see if the performance of your server and improved or decreased.

Arrow Indicator	Description
	A green arrow indicates a good change. In some cases a lower value is preferable, e.g. Average Request Execution Time.
	A red arrow indicates a bad change. In some cases a higher value is not preferable, e.g. Average Memory Used.
	A grey arrow indicates that the change is neither good nor bad. For example, the Average Requests Running uses this type of arrow.
	This indicates no change between the current and previous values

On this page:

Overview

The Overview header displays the following information (Figure 1):

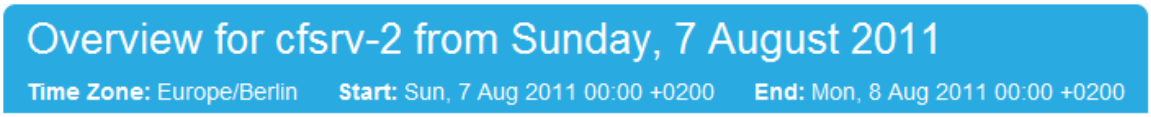

Figure 1: Overview section header example

The table below gives an explanation of what each part of the header means:

Header Information	Description
DataCollector *e.g. cfsrv-2*	The DataCollector that the data belongs to.
Day *e.g. Monday, 25 July 2011*	The date that the report is run for (from midnight to midnight of that day).
Time Zone	The default time zone is the one set on the DataCollector. See Daily Status Report Provider Arguments if you want to change the time zone of the report.
Start	The exact date and time that the data in the report starts from. It also includes the offset from UTC time (*e.g. +0200*).
End	The exact date and time that the data in the report ends at. It also includes the offset from UTC time (*e.g. +0200*).

Clicking on the Overview header will take you to the Overview perspective in the Web Client.

Overview contains the following sections:

Name	Example	Description
Total Requests Completed	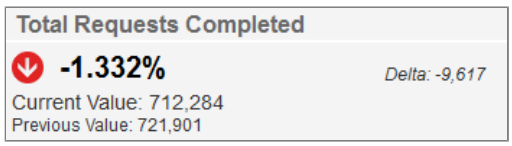	The total number of requests that have completed during the 24 hours.
Min Requests Completed (Between Samples)		The smallest number of requests completed at sampled points throughout the 24 hours.

Peak Requests Completed (Between Samples)	**Peak Requests Completed (Between Samples)** 12.871% Delta: 13 Current Value: 114 Previous Value: 101	The largest number of requests completed at sampled points throughout the 24 hours.
Average Requests / Sec	**Average Requests / Sec** -1.220% Delta: -0.104 Current Value: 8.386 Previous Value: 8.489	The average number of requests that are completed per second.
Min Requests / Sec	**Min Requests / Sec** 0.000% Delta: 0.000 Current Value: 0.000 Previous Value: 0.000	The smallest number of requests that are completed per second.
Peak Requests / Sec	**Peak Requests / Sec** 52.537% Delta: 7.820 Current Value: 22.704 Previous Value: 14.884	The largest number of requests that are completed per second.
Average Requests Running	**Average Requests Running** 1.701% Delta: 0.124 Current Value: 7.440 Previous Value: 7.315	The average number of requests running at any time.
Average Request Execution Time (sec)	**Average Request Execution Time (sec)** -2.415% Delta: -0.052 Current Value: 2.099 Previous Value: 2.151	The average execution time. in seconds, of all requests in the 24 hours.
Total JDBC Requests Completed	**Total JDBC Requests Completed** -1.537% Delta: -12,780 Current Value: 818,733 Previous Value: 831,513	The total number of JDBC requests that have completed during the 24 hours.
Min JDBC Requests Completed (Between Samples)	**Min JDBC Requests Completed (Between Samples)** 0.000% Delta: 0 Current Value: 0 Previous Value: 0	The smallest number of JDBC requests completed at sampled points throughout the 24 hours.
Peak JDBC Requests Completed (Between Samples)	**Peak JDBC Requests Completed (Between Samples)** 8.955% Delta: 12 Current Value: 146 Previous Value: 134	The largest number of JDBC requests completed at sampled points throughout the 24 hours.
Average JDBC Requests / Sec	**Average JDBC Requests / Sec** -1.415% Delta: -0.138 Current Value: 9.645 Previous Value: 9.783	The average number of JDBC requests that are completed per second.
Min JDBC Requests / Sec	**Min JDBC Requests / Sec** 0.000% Delta: 0.000 Current Value: 0.000 Previous Value: 0.000	The smallest number of JDBC requests that are completed per second.
Peak JDBC Requests / Sec	**Peak JDBC Requests / Sec** 4.066% Delta: 0.968 Current Value: 24.768 Previous Value: 23.800	The largest number of requests that are completed per second.
Average JDBC Request Execution Time (sec)	**Average JDBC Request Execution Time (sec)** 37.064% Delta: 0.001 Current Value: 0.006 Previous Value: 0.004	The average JDBC execution time, in seconds, of all requests in the 24 hours.

Back to the top

Example Overview

(click to enlarge)

Performance Charts

The Performance Charts section contains the following:

Graph	Example	Description
Request Execution Time: Average	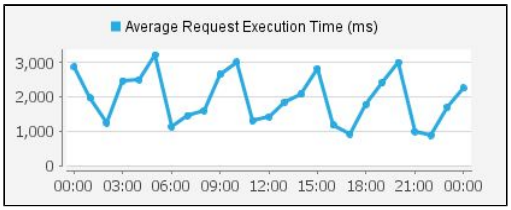	A graph showing the average execution time, in milliseconds, for the requests throughout the day. There is one point on the graph for each hour.
Request Execution Time: Peak	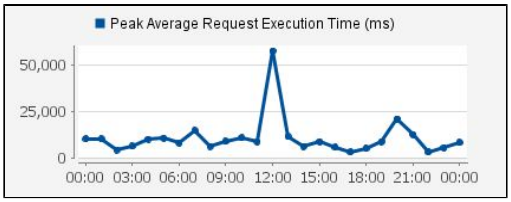	A graph showing the maximum average execution time, in milliseconds, for the requests throughout the day. There is one point on the graph for each hour.
Daily Traffic	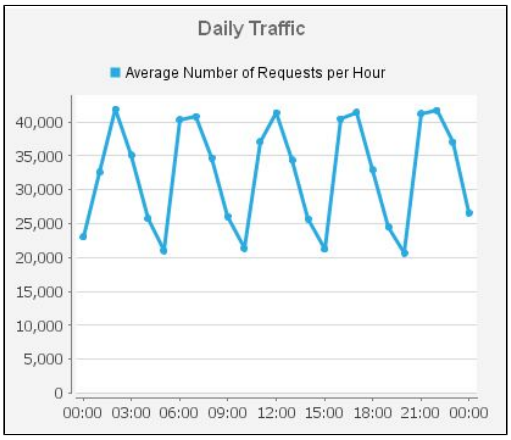	A graph showing the average number of requests per hour for the requests throughout the day. There is one point on the graph for each hour.
Database Request Execution Time: Average	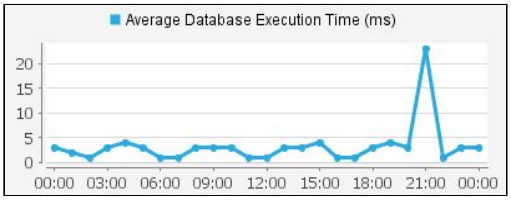	A graph showing the average JDBC execution time, in milliseconds, for the requests throughout the day. There is one point on the graph for each hour.

Database Request Execution Time: Peak	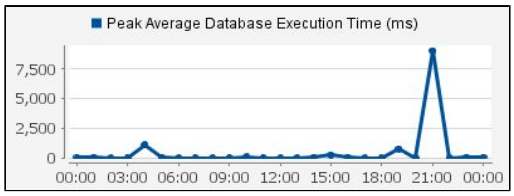	A graph showing the maximum average JDBC execution time, in milliseconds, for the requests throughout the day. There is one point on the graph for each hour.
Daily Database Traffic	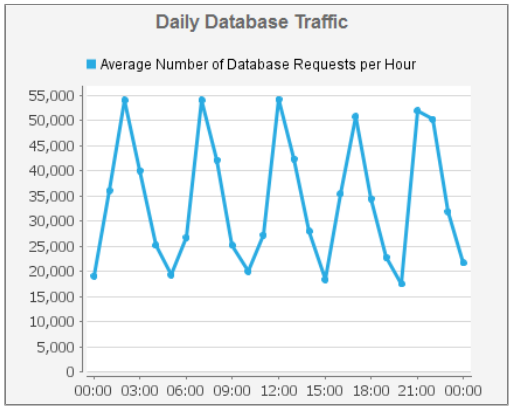	A graph showing the average number of database requests per hour for the requests throughout the day. There is one point on the graph for each hour.

Clicking on the Performance Charts header will take you to the Overview perspective in the Web Client.

Back to the top

Example Performance Charts

(click to enlarge)

System Load Statistics

The System Load Statistics section contains the following:

Name	Example	Description
System CPU Load	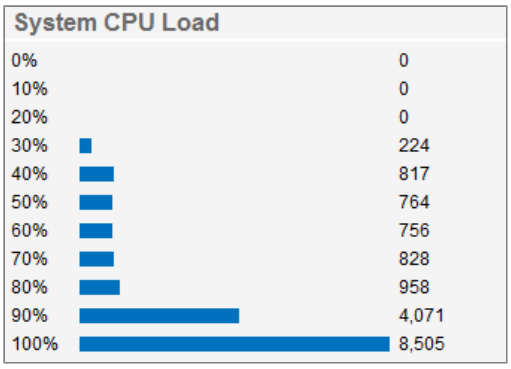	A bar graph showing the number of samples taken at each percentage load for the system CPU.

Process CPU Load	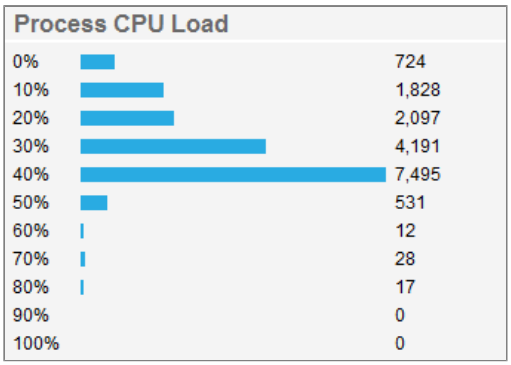	A bar graph showing the number of samples taken at each percentage load for the process CPU.
Memory Usage	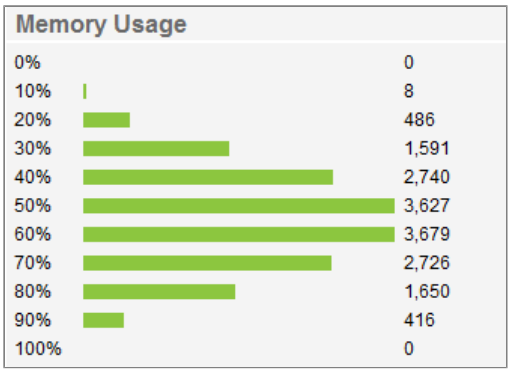	A bar graph showing the number of samples taken at each percentage load for memory.
Average System CPU (%)	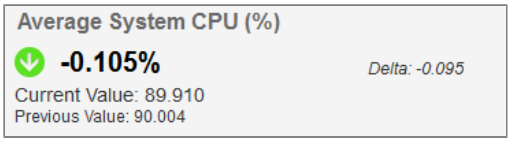	The average percentage for the amount of System CPU used.
Average Process CPU (%)	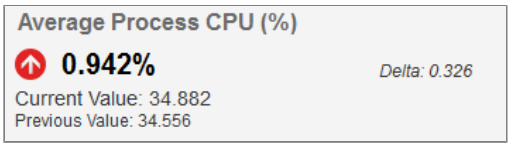	The average percentage for the amount of Process CPU used.
Average Memory (%)	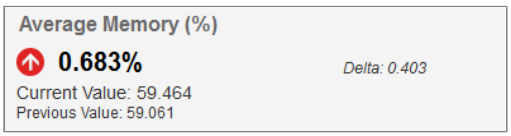	The average percentage for the amount of Memory used.

Clicking on the System Load Statistics header will take you to the CPU/Memory Usage Chart perspective in the Web Client.

Back to the top

Example System Load Statistics

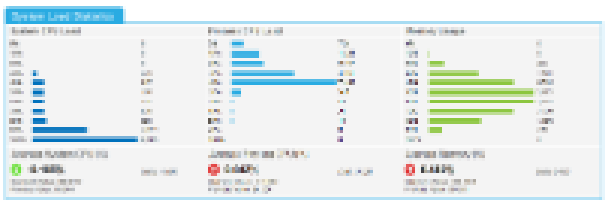

(click to enlarge)

Session Statistics

The Session Statistics section contains the following:

Name	Description

Unique Session IDs	The total number of unique session IDs used during that day. The arrow and percentage indicate the change between the values from today and yesterday.
Unique ColdFusion IDs	The total number of unique ColdFusion IDs used during that day. The arrow and percentage indicate the change between the values from today and yesterday.
Request Count	The total number of requests during that day. The arrow and percentage indicate the change between the values from today and yesterday.
Average Requests per ColdFusion ID Token	The average number of requests that occurred for each ColdFusion ID
Average Requests per JSession ID	The average number of requests that occurred for each JSession ID
Server ID	If there have been any server restarts, the server ID will be listed, with the Date/Time (see below).
Date / Time	The time of any server restarts that have occurred during that day.

Clicking on the Session Statistics header will take you to the Table of Requests perspective in the Web Client.

Back to the top

Example Session Statistics

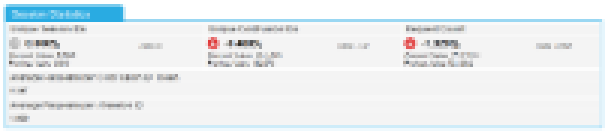

(click to enlarge)

Slowest Requests

Slowest Requests contains the following sections:

Name	Example	Description
Started	Started 15:52:01:640 01:48:55:609 18:54:14:406	The time when the request started.
Request URL	Request URL http://int0046.bbn.intergral.com:18080/apps/fads/test.cpuspike.cfm http://int0046.bbn.intergral.com:18080/apps/fads/test.cpuspike.cfm http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm	The request URL.
Execution Time (sec)	Execution Time (sec) 182.344 182.172 182.016	The length of time, in seconds, that the request took to run.

Clicking on the Slowest Requests header will take you to the Slowest Rwquests perspective in the Web Client.

You can also click **Expand / Collapse** to see more slow requests.

Back to the top

Example Slowest Requests

(click to enlarge)

Slowest Requests on Average

Slowest Requests on Average contains the following section:

Name	Example	Description
Request Count	Request Count 3 1,054 1,534	The total number of requests that have occurred for this URL.
Request URL	Request URL http://int0046.bbn.intergral.com:18080/apps/fads/test.cpuspike.cfm http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm http://int0046.bbn.intergral.com:18080/apps/ConsumejustFunction.cfm	The URL's of the requests that take the longest time to complete on average.
Execution Time (sec)	Execution Time (sec) 182.177 72.706 21.945	The average execution time in seconds for each URL.

Clicking on the Slowest Requests on Average header will take you to the Slowest Requests on Average (Grouped by URL) perspective in the Web Client.

You can also click **Expand / Collapse** to see more slow requests.

Back to the top

Example Slowest Requests on Average
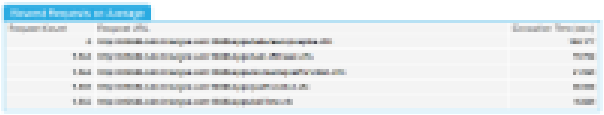
(click to enlarge)

Application Data Overview

Application Data Overview contains the following sections:

Name	Example	Description
Application Key	Application Key test.cfthread.cfm twitter flex2gateway	Each application filter that you define will be listed. To find out how to configure your own application filters, see Manage Applications.

Amount of Views	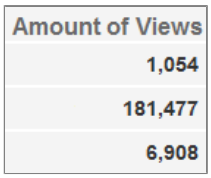	The total amount of times the pages within this filter have been viewed.
Average Execution Time (sec)	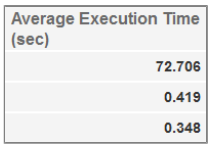	The average execution time in seconds for the requests within each application filter.
Total Execution Time	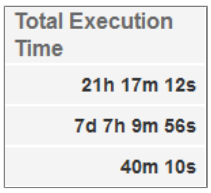	The total execution time for the requests within each application filter.
Average Bytes Sent	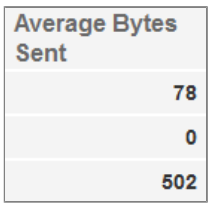	The average number of bytes for each application filter.
Total Bytes Sent	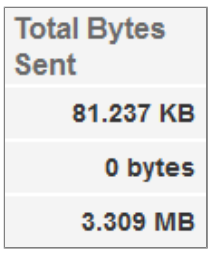	The total number of bytes sent for each application filter.
Unique Session IDs	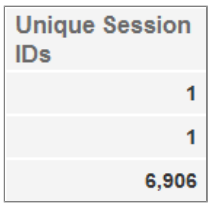	The number of unique session IDs per application.
Unique ColdFusion IDs	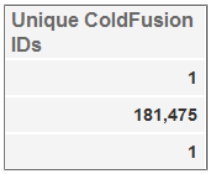	The number of unique ColdFusion IDs per application.

Back to the top

Example Application Data Overview

(click to enlarge)

Application JDBC Overview

Application JDBC Overview contains the following sections:

Name	Example	Description
Application Key	Application Key test.cfthread.cfm twitter flex2gateway	Each application filter that you define will be listed. To find out how to configure your own application filters, see Manage Applications.
Average JDBC Queries	Average JDBC Queries 1.161 0.000 2.287	The average number of database queries specifically from the requests within each application filter.
Totall JDBC Queries	Total JDBC Queries 818,750 0 560,797	The total number of database queries specifically from the requests within each application filter.
Average JDBC Time (sec)	Average JDBC Time (sec) 0.005 0.000 0.009	The average database execution time in seconds for the requests within each application filter.
Total JDBC Time	Total JDBC Time 1h 3m 25s 0s 40m 22s	The total time spent executing queries on a database for the requests within each application filter.
Average JDBC Rows	Average JDBC Rows 37.466 0.000 1.620	The average number of rows returned from the database for each application.
Total JDBC Rows	Total JDBC Rows 26,428,852 0 397,178	The total number of rows returned from the database for each application.

Back to the top

Example Application JDBC Overview

(click to enlarge)

Applications

Applications contains the following sections:

Name	Example	Description
Application Name	Applications Application Name: test.cfthread.cfm	Each application filter that you have configured will have its own sub-section within this section. To find out how to configure your own application filters, see Manage Applications.

Slowest Requests	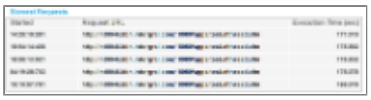 *(click to enlarge)*	This section will only show you the slowest requests for this application. It is displayed the same as the Slowest Requests section.
Slowest Requests on Average	 *(click to enlarge)*	This section will only show you requests that match the application filter, grouped by URL. It is displayed the same as the Slowest Requests on Average section.

Back to the top

Example Applications

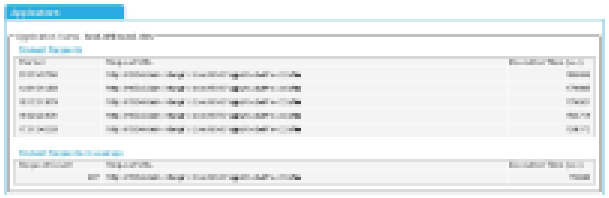

(click to enlarge)

Also See

FusionAnalytics User's Guide
See the User's Guide for more information on the different parts of FusionAnalytics.

Manage Applications
Visit this section for help configuring your own application filters.

Take me back to **FusionAnalytics Reporting**.

FAQs

FusionAnalytics for FusionReactor User's Guide

Overview

This section will include Frequently Asked Questions (FAQs) for the different sections on FusionAnalytics.

> ⚠ This page is currently growing and FAQ's will be continually added to this section of the Documentation as they come in.

Note: For further FAQs regarding FusinAnalytics please see: Frequently Asked Questions about FusionAnalytics

Next Steps

DataCollector FAQ

Desktop FAQ

FusionAnalytics Application FAQ

General FAQ

Licensing FAQ

Server FAQ

DataCollector FAQ

DataCollector FAQ

Back to the top

Desktop FAQ

Desktop FAQ

Back to the top

FusionAnalytics Application FAQ

FusionAnalytics Application FAQ

Back to the top

General FAQ

General FAQ

Back to the top

Licensing FAQ

Licensing FAQ

What is a product family?

A product family is a "group" to which applications and DataCollectors can belong to. For example, the product family of the application that comes with FusionAnalytics is "FA4FR" a.k.a "FusionAnalytics for FusionReactor". This product family is covered by a single license. If you wish to use applications belonging to a different product family, you must also install a license for that family. With the first release of Analytics only the "FusionAnalytics for FusionReactor" application is supported.

What are "constraints" in a FusionAnalytics license?

FusionAnalytics licenses are divided into 3 parts:

- Users
- Applications
- DataCollectors

Each license is for a single product family.

What is a user (as a license constraint)?

The user constraint only refers to FADS. The single user in FADC has nothing to do with licensing.

In FADS, you may add as many users as you wish. In order for this user to do anything, he must be mapped to a product family. When mapped to 1 or more applications for a family, the user is then considered a user of that product family.

In other words, whether mapped to 1 or 10 FA4FR applications, a user is considered a single user of that product family and only counts for 1 user on your license.

What is a DataCollector (as a license constraint)?

DataCollectors are only used in FADC. They are not used anywhere in FADS.

When you add additional DataCollectors to a FADC application, it is counted as a DataCollector from your license. If you add a new FADC application, by default it will have 1 DataCollector, so while you are allowed an unlimited number of FADC applications, if you have reached your limit of DataCollectors you will be not be able to upload another app (this does not effect FADS applications).

What is an Application (as a license constraint)?

The application constraint only refers to FADS applications. FADC applications do not effect this constraint.

Whenever you upload an application to FADS it counts as 1 application on your license for its product family. If you have reached your limit of applications for that family you will not be able to upload another one.

I tried to activate my license and was told I have no activations remaining

If you see this message it means you have already installed this license on the number of machines it is valid for. You must uninstall the license from another machine before you can use it again.

I didn't activate my license in time and the activation period has expired. Have I lost my configurations?

If your activation period expires there is no need to worry. Simply activate the license and your users, applications and DataCollectors will function as they did before.

My license expires soon. Can I purchase another activation for it?

If your license is going to expire soon you must purchase a new license. When you install it everything will function as it did before.

I need a bigger license. Do I have to purchase a new one?

If you need a bigger license, please contact us and we will help you upgrade. You will only have to pay the price of the upgrade, not the full price of a new bigger license.

Can I install two licenses for the same product family?

You cannot install two licenses for the same product family. You must upgrade your license if your current one is too small.

Back to the top

Server FAQ

Server FAQ

Back to the top

Feedback and Support

Feedback and Support

We welcome feedback on all our products and publications. To discover the various ways of contacting us, please use our contact page. We will address your feedback as quickly as possible.

For additional support information and material please see the Support / Knowledge Base section of the FusionAnalytics website.

Additionally you can refer to the FusionAnalytics Google Groups for further support and information.

Back to the top

Copyright

Copyright

Responsible entity for website content may be found here.