

| | |
|--|-----|
| 1. Home | 3 |
| 1.1 Installation Guides | 3 |
| 1.1.1 System Requirements | 4 |
| 1.1.1.1 Amazon Elastic Compute Cloud (Amazon EC2) | 6 |
| 1.1.2 Installing FusionAnalytics Server | 7 |
| 1.1.2.1 Prerequisites for Installation | 7 |
| 1.1.2.2 Installation | 7 |
| 1.1.2.3 Installing your license | 16 |
| 1.1.3 Installing the FusionAnalytics AIR Client | 21 |
| 1.1.4 Installing FusionReactor | 23 |
| 1.1.5 Installing FusionReactor Extensions for ColdFusion | 24 |
| 1.1.6 Uninstall a FusionAnalytics License Key | 24 |
| 1.1.7 Uninstalling FusionAnalytics | 25 |
| 1.1.8 Upgrading FusionAnalytics | 27 |
| 1.2 Quick Start Guide | 29 |
| 1.2.1 Database Configuration | 30 |
| 1.2.1.1 Configure Network Settings | 31 |
| 1.2.1.2 Create a New Database Login | 34 |
| 1.2.1.3 Create a New Database | 36 |
| 1.2.2 Accessing FusionAnalytics | 40 |
| 1.2.3 Setting Up the AIR Client | 43 |
| 1.2.4 Configuring FusionReactor for FusionAnalytics | 47 |
| 1.2.5 Setting Up Users | 49 |
| 1.2.6 Setting Up the Email Server | 52 |
| 1.2.7 Setting Up Application Filters | 55 |
| 1.2.8 Setting Up TAP and Daily Status Reports | 57 |
| 1.2.9 Setting up the SLA Reports | 60 |
| 1.3 FusionAnalytics User's Guide | 62 |
| 1.3.1 Getting Started | 64 |
| 1.3.1.1 How FusionAnalytics Is Organized | 64 |
| 1.3.1.2 Expandable and Extensible Platform | 65 |
| 1.3.1.3 Scaling and Quantization | 65 |
| 1.3.1.4 Exploring Data | 66 |
| 1.3.2 FusionAnalytics Client | 68 |
| 1.3.2.1 Getting Started with the FA Client | 69 |
| 1.3.2.2 The FusionAnalytics Web Client | 76 |
| 1.3.2.3 The FusionAnalytics AIR Client | 78 |
| 1.3.2.4 Using the Table of Contents (TOC) | 81 |
| 1.3.2.5 Date Navigation | 82 |
| 1.3.2.6 Time Zones | 86 |
| 1.3.2.7 Filtering Data | 87 |
| 1.3.2.8 Data Visualizations | 93 |
| 1.3.2.8.1 DataGrid Visualization | 93 |
| 1.3.2.8.2 Chart Base Visualization | 98 |
| 1.3.2.8.3 Time Block Chart Visualization | 103 |
| 1.3.2.8.4 Placemark Chart Visualization | 107 |
| 1.3.2.8.5 Sixway Chart Visualization | 115 |
| 1.3.2.9 Application Menus | 118 |
| 1.3.2.10 Using Tabs | 120 |
| 1.3.2.11 Using State Stores | 122 |
| 1.3.2.12 Context Menus | 124 |
| 1.3.2.13 Business Metrics | 125 |
| 1.3.2.13.1 Reports | 125 |
| 1.3.2.14 Administrator | 128 |
| 1.3.2.14.1 Manage Applications | 129 |
| 1.3.2.14.2 FusionAnalytics DataServices | 134 |
| 1.3.2.14.3 FusionAnalytics DataCollector | 135 |
| 1.3.3 FusionAnalytics Server | 136 |
| 1.3.3.1 Server Structure | 137 |
| 1.3.3.1.1 Directory Structure | 137 |
| 1.3.3.1.2 FADC Applications | 138 |
| 1.3.3.1.3 FADS Applications | 139 |
| 1.3.3.2 Server Configuration | 139 |
| 1.3.3.2.1 FusionAnalytics Log Files | 140 |
| 1.3.3.2.2 Windows Service Configuration | 140 |
| 1.3.3.2.3 FusionAnalytics DataCollector Settings | 141 |
| 1.3.3.2.4 Setting Up FusionAnalytics with HTTPS | 142 |
| 1.3.3.2.5 JVM Memory Requirements | 144 |
| 1.3.3.3 FusionAnalytics DataCollector (FADC) | 144 |
| 1.3.3.3.1 Applications (FADC) | 145 |
| 1.3.3.3.2 User (FADC) | 158 |
| 1.3.3.3.3 System (FADC) | 159 |
| 1.3.3.3.4 DCML (FADC) | 161 |
| 1.3.3.4 FusionAnalytics DataServices (FADS) | 165 |
| 1.3.3.4.1 Applications (FADS) | 166 |
| 1.3.3.4.2 Users (FADS) | 174 |
| 1.3.3.4.3 Scheduled Tasks (FADS) | 177 |

| | |
|---|-----|
| 1.3.3.4.4 System (FADS) | 181 |
| 1.3.3.4.5 APML (FADS) | 186 |
| 1.3.3.4.6 Configuring Report Images | 196 |
| 1.3.3.5 FusionAnalytics Licensing | 197 |
| 1.3.4 Move FusionAnalytics to a new server | 198 |
| 1.4 FusionAnalytics for FusionReactor User's Guide | 199 |
| 1.4.1 Using FusionAnalytics with FusionReactor | 201 |
| 1.4.2 Application Perspectives (eg. faDataCollector-1) | 201 |
| 1.4.2.1 Overview | 203 |
| 1.4.2.2 Deep Server Analysis | 204 |
| 1.4.2.2.1 Data Finder | 206 |
| 1.4.2.2.2 Request | 208 |
| 1.4.2.2.3 Database | 260 |
| 1.4.2.2.4 Memory | 271 |
| 1.4.2.2.5 CPU | 277 |
| 1.4.2.2.6 Thread | 281 |
| 1.4.2.2.7 General | 282 |
| 1.4.2.2.8 Logs | 287 |
| 1.4.2.2.9 Events | 289 |
| 1.4.2.2.10 Hot-Spots | 297 |
| 1.4.2.2.11 ColdFusion Monitor | 314 |
| 1.4.2.2.12 Applications | 323 |
| 1.4.3 FusionAnalytics Reporting | 325 |
| 1.4.3.1 TAP Reports | 326 |
| 1.4.3.1.1 What is TAP Score - and how is it calculated? | 327 |
| 1.4.3.1.2 Running a TAP Report | 329 |
| 1.4.3.1.3 TAP Report Provider Arguments | 332 |
| 1.4.3.1.4 Understanding the TAP Report | 333 |
| 1.4.3.1.5 The TAP Report configuration page | 349 |
| 1.4.3.2 Daily Status Reports | 351 |
| 1.4.3.2.1 Running a Daily Status Report | 352 |
| 1.4.3.2.2 Daily Status Report Provider Arguments | 354 |
| 1.4.3.2.3 Understanding the Daily Status Report | 355 |
| 1.4.3.2.4 The Daily Status Report configuration page | 378 |
| 1.4.3.3 SLA Reports | 379 |
| 1.4.3.3.1 Accessing the SLA Report Configuration Page | 380 |
| 1.4.3.3.2 Adding a new SLA Report Configuration | 380 |
| 1.4.3.3.3 The SLA Report Configuration | 381 |
| 1.4.3.3.4 Running the SLA Reports | 385 |
| 1.4.4 Using FusionAnalytics to analyze a cluster of servers / instances | 386 |
| 1.4.4.1 Cluster Perspectives | 387 |
| 1.5 FAQs | 388 |
| 1.5.1 General FAQ | 389 |
| 1.5.2 Server FAQ | 390 |
| 1.5.3 DataCollector FAQ | 391 |
| 1.5.4 FusionAnalytics Application FAQ | 392 |
| 1.5.5 Desktop Client FAQ | 392 |
| 1.5.6 Licensing FAQ | 395 |
| 1.5.7 SLA Report FAQ | 397 |
| 1.6 Feedback and Support | 398 |
| 1.6.1 Known Issues | 398 |
| 1.7 [FNS-1] Creating Daily Status Reports in FusionAnalytics | 399 |
| 2. Copyright | 403 |

Home

- [Introducing FusionAnalytics](#)
- [Welcome](#)
- [Feedback](#)
- [License and Copyright Information](#)

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 its 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 FusionAnalytics App from scratch. Information on customizing or building new Apps will be coming soon.

If you are **new** to FusionAnalytics, it is recommended to start by taking a look at the User Guide and once you have everything installed. 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 send an email to our support team on support@fusion-reactor.com. We will address your feedback as quickly as possible.

License and Copyright Information

- [FusionAnalytics End User License Agreement](#).
- [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](#).

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

- [Amazon Elastic Compute Cloud \(Amazon EC2\)](#)

INSTALLING FUSIONANALYTICS SERVER

- [Prerequisites for Installation](#)
- [Installation](#)
- [Installing your license](#)

INSTALLING THE FUSIONANALYTICS AIR CLIENT

INSTALLING FUSIONREACTOR

INSTALLING FUSIONREACTOR EXTENSIONS FOR COLDFUSION

UNINSTALL A FUSIONANALYTICS LICENSE KEY

UNINSTALLING FUSIONANALYTICS

UPGRADING FUSIONANALYTICS

System Requirements

System Requirements

PLEASE NOTE: Memory and Hard Disk Space requirements are **completely** dependent on the amount of metric data you wish to store and have available in FusionAnalytics. Read our Data Guidelines section.

Operating Systems



Windows

| | |
|-------------------------------|--|
| FusionAnalytics Server | Microsoft Windows Server 2016 Microsoft Windows Server 2012 Microsoft Windows Server 2008 Microsoft Windows Server 2003 Microsoft Windows 8 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 available (MIN) - Using Microsoft SQL Server 2008 R2 Express |
| Architecture | 64-BIT / 32-BIT |

Servers

ColdFusion

| | |
|-------------------------------------|--|
| Supported ColdFusion Servers | ColdFusion 2016 (64-BIT / 32-BIT) ColdFusion 11 (64-BIT / 32-BIT) ColdFusion 10 (64-BIT / 32-BIT) ColdFusion 9 (64-bit / 32-BIT) ColdFusion 8 (64-bit / 32-BIT) ColdFusion MX 7 ColdFusion MX6 |
|-------------------------------------|--|

Server Monitor

FusionReactor

On This Page:



Windows



ColdFusion



FusionReactor



Database



AIR™

Hardware

Screen Resolution

| | |
|-------------------------------|--|
| FusionReactor Editions | FusionReactor 4 (or above) Enterprise Edition (Perpetual / Subscription) FusionReactor 4 (or above) Standard Edition (Perpetual / Subscription) |
|-------------------------------|--|

Database

Microsoft

| | |
|----------------------------|--|
| SQL Server Versions | Microsoft SQL Server 2016 Microsoft SQL Server 2014 Microsoft SQL Server 2012 Microsoft SQL Server 2008 R2 Microsoft SQL Server 2008 Microsoft SQL Server 2005 Microsoft SQL Server 2012 Express Microsoft SQL Server 2008 R2 Express |
| Memory | Minimum: 1 GB Recommended: 4 GB or more Maximum: Operating system maximum Note : SQL Server Enterprise Edition supports a maximum of 2 TB of RAM or operating system maximum, whichever is lower |
| Database size | The maximum database size applies per database, excluding log files. 10 GB (MAXIMUM for Express versions) |
| CPU's | Depends on MS SQL License Processor type: Itanium processor or faster Processor speed: Recommended: 1.0 GHz or faster |
| Architecture | 64-BIT / 32-BIT |

PLEASE NOTE: Microsoft recommends that you run SQL Server 2012 on computers with the NTFS file format – please read the Hardware and Software Requirements for Installing SQL Server 2012 located on Microsoft's website for full details.

Please note that FusionAnalytics will **ONLY** support **Microsoft SQL Server** versions listed above.

Adobe AIR Runtime

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

Hardware Platform

Processor

| | |
|-------------------------------|---|
| FusionAnalytics Server | 1 gigahertz (GHz) or faster 32-bit (x86) or 64-bit (x64) processor (MIN) We recommend a 2 GHz or higher quad core processor |
|-------------------------------|---|

Screen Resolution

Resolution

| | |
|---|---|
| FusionAnalyticsAir Client FusionAnalytics Web Client | SXGA 1280 x 1024 (MIN) UXGA 1600 x 1200 or higher resolution monitor (RECOMMENDED) |
|---|---|

Next Steps

Amazon Elastic Compute Cloud (Amazon EC2)


Amazon Elastic Compute Cloud (Amazon EC2)

Amazon Elastic Compute Cloud (Amazon EC2)

Amazon charges for EC2 instances separately - and you should be aware of the pricing policies before setting up any instances.

When selecting which Amazon EC2 instance you require, it is **IMPORTANT** to note that instance requirements are **completely** dependent on the amount of metric data you wish to store and have available in FusionAnalytics. The information shown below are **ONLY GUIDELINES BASED ON ONE SPECIFIC TEST CASE**. For more information on sizing guidelines, please refer to our Data Guidelines section.

Minimum Instance Requirements for FusionAnalytics Server

| | |
|--------------------------|--|
| Operating Systems |  Windows Server 2008 R2 Base ** |
| Instance Type | Medium instance (<i>m1.medium</i>) |
| Memory | 3.75 GB (MIN) |
| Hard Disk Space | 40GB (MIN) |
| CPU Units | 2 ECUs (1 Core) |
| Architecture | 64-BIT |

**** NOTE : THE SQL DATABASE WAS NOT SETUP ON THIS TEST SERVER - ONLY FUSIONANALYTICS SERVER - The SQL Server was setup on a separate EC2 Instance**

For our test case, an Amazon EC2 instance, setup with these minimum specifications was able to manage up to 5 DataCollectors, and processing a combined total of approx. 500,000 requests per day (100,000 requests per DataCollector).

A LARGER EC2 instance will be required when using more than 5 DataCollectors and we recommend that you increase the maximum heap space of FusionAnalytics Server JVM by 128MB for each additional DataCollector.

Additional information on FusionAnalytics JVM requirements can be found [here](#).

This information is ONLY to be given as a **GUIDELINE** to show the Minimum configuration necessary to support our example application, running 5 DataCollectors / 500,000 Requests (per day)

It is the **responsibility of the customer** to configure and test their Amazon EC2 image to ensure that it will support their application running

Please note that FusionAnalytics will **ONLY** support Microsoft SQL Server 2016, Microsoft SQL Server 2014, Microsoft SQL Server 2012, Microsoft SQL Server 2008 R2, Microsoft SQL Server 2008 and Microsoft SQL Server 2005.

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

In order to get started with the installation of FusionAnalytics in your system, please make sure that the following steps have been checked.

1. You have checked the minimum system requirements for the system where you want to install FusionAnalytics. For more information please check the following link, System Requirements.
2. It's really important that you have **prepared your database**. For more information, please check out the following link, Database Configuration.
3. You have the **database login details available** with **appropriate privileges** on the target database.
4. You have the appropriate FusionAnalytics installer. The latest installer can be found in the following link; <http://www.fusion-analytics.com/download/>
5. Make sure that the JavaScript is enabled in your default browser.

NEXT STEPS

Installation

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

Installation

INSTALLING FUSIONANALYTICS SERVER

Executables

| Executable | System Architecture |
|---------------------------------------|---------------------|
| FusionAnalytics_windows_2_0_8.exe | 32-bit |
| FusionAnalytics_windows-x64_2_0_8.exe | 64-bit |

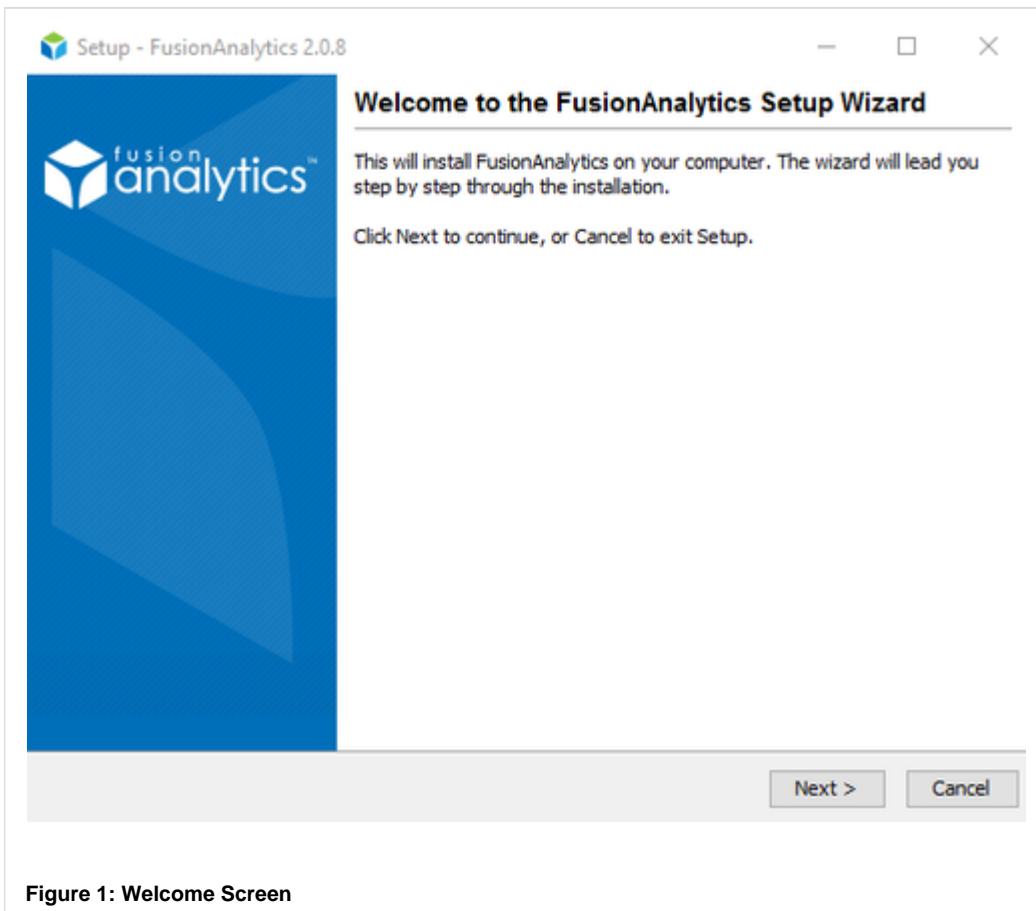
In order to get a copy of the latest executable options, please visit our official website here.

Setup

To start the installation process, double click the latest installer OR right click on the file and choose the option "Run as administrator".

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.



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.

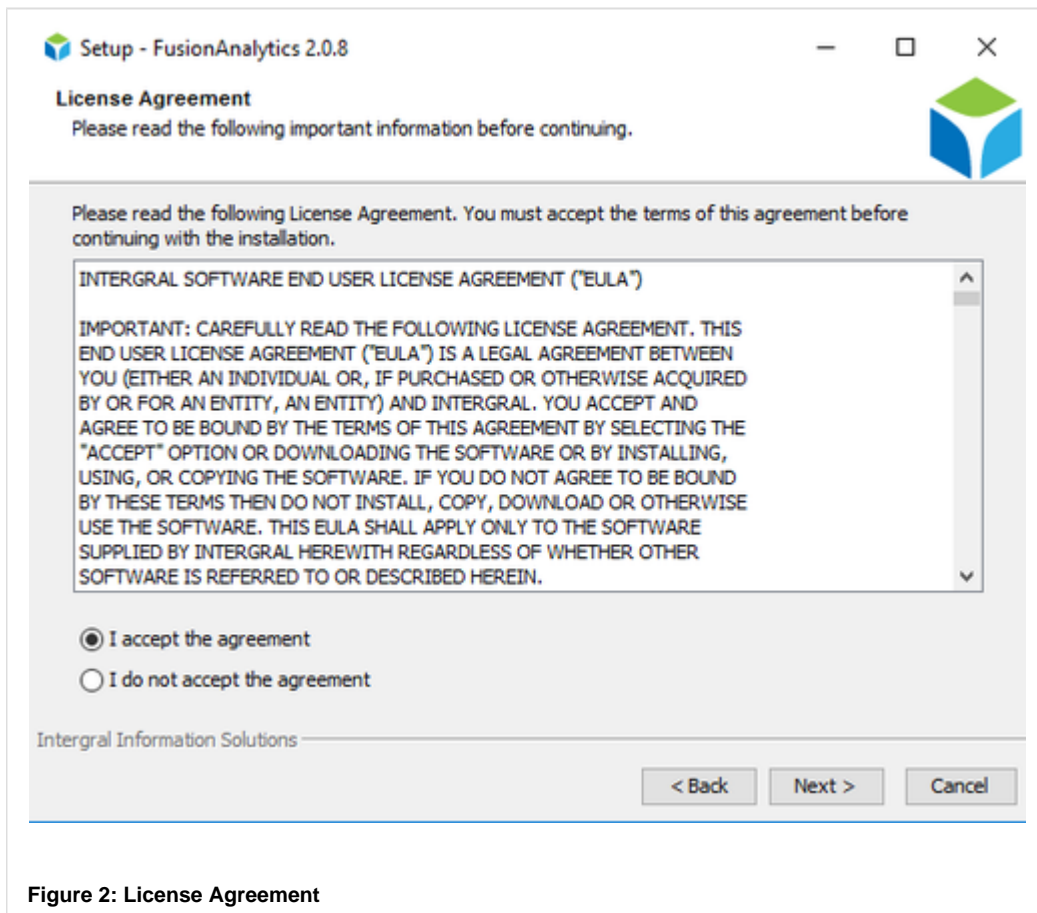


Figure 2: License Agreement

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.

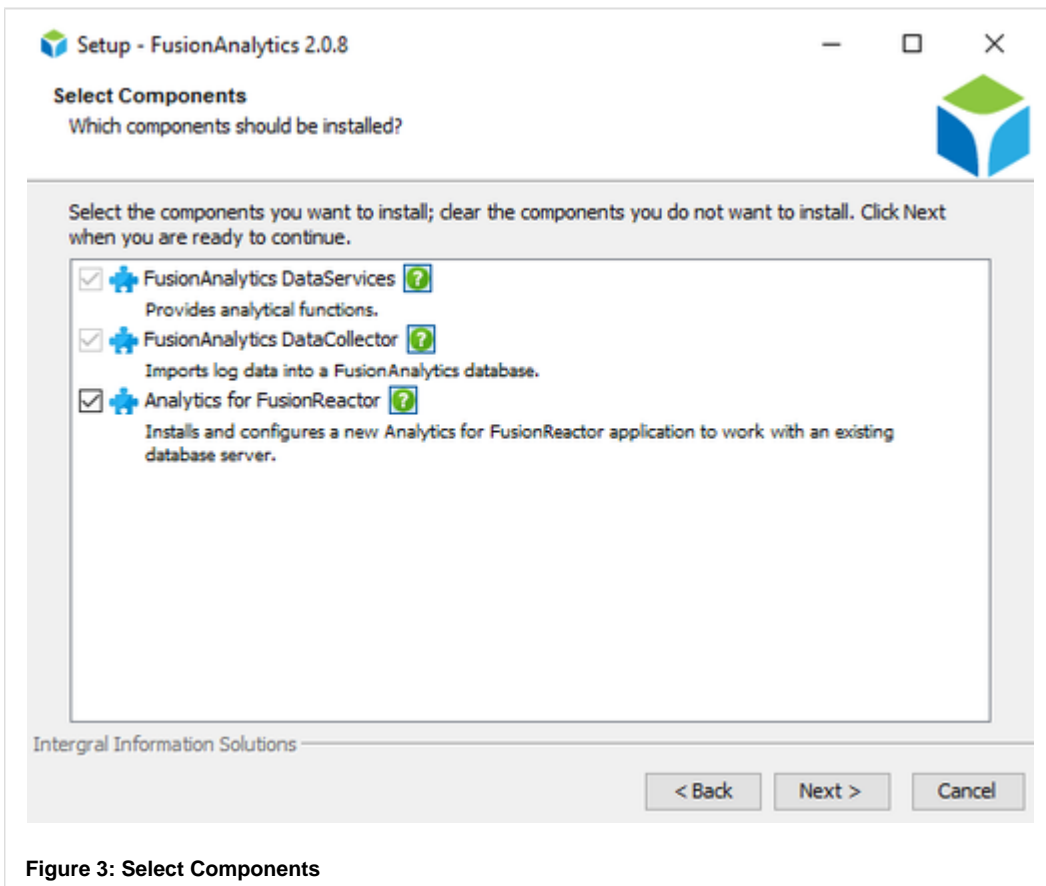


Figure 3: Select Components

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.

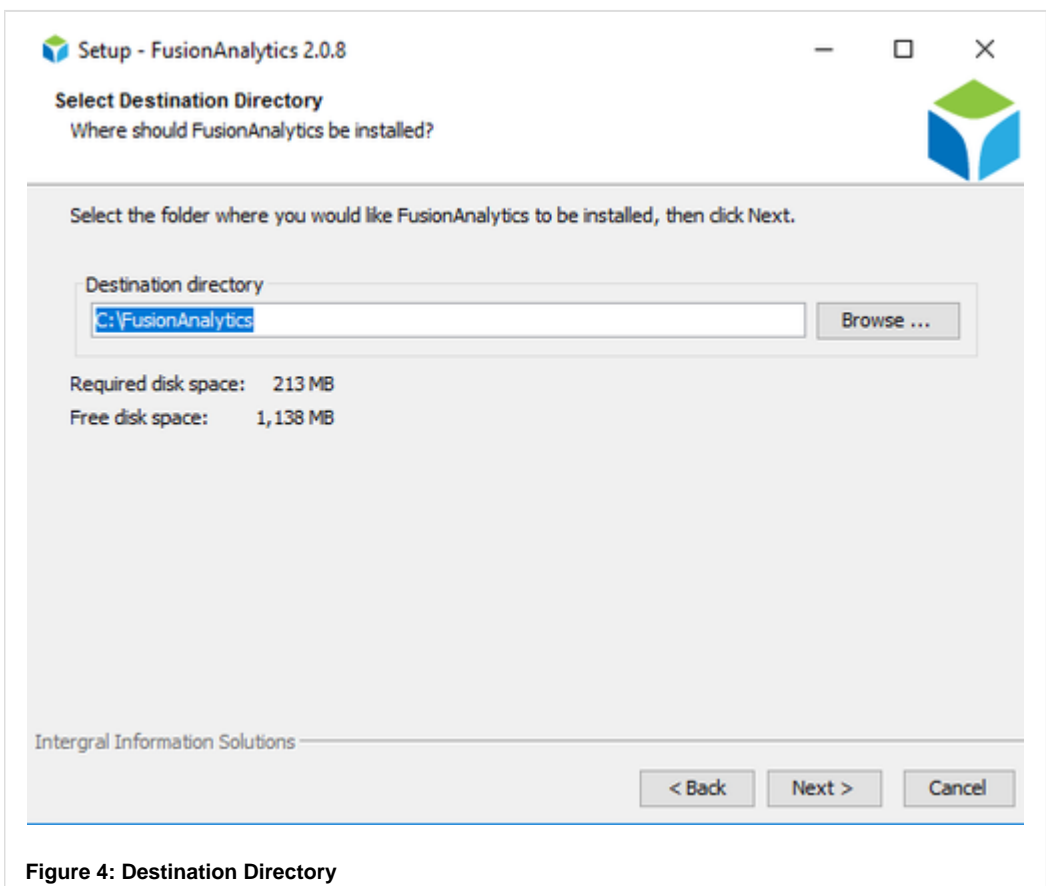


Figure 4: Destination Directory

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*.

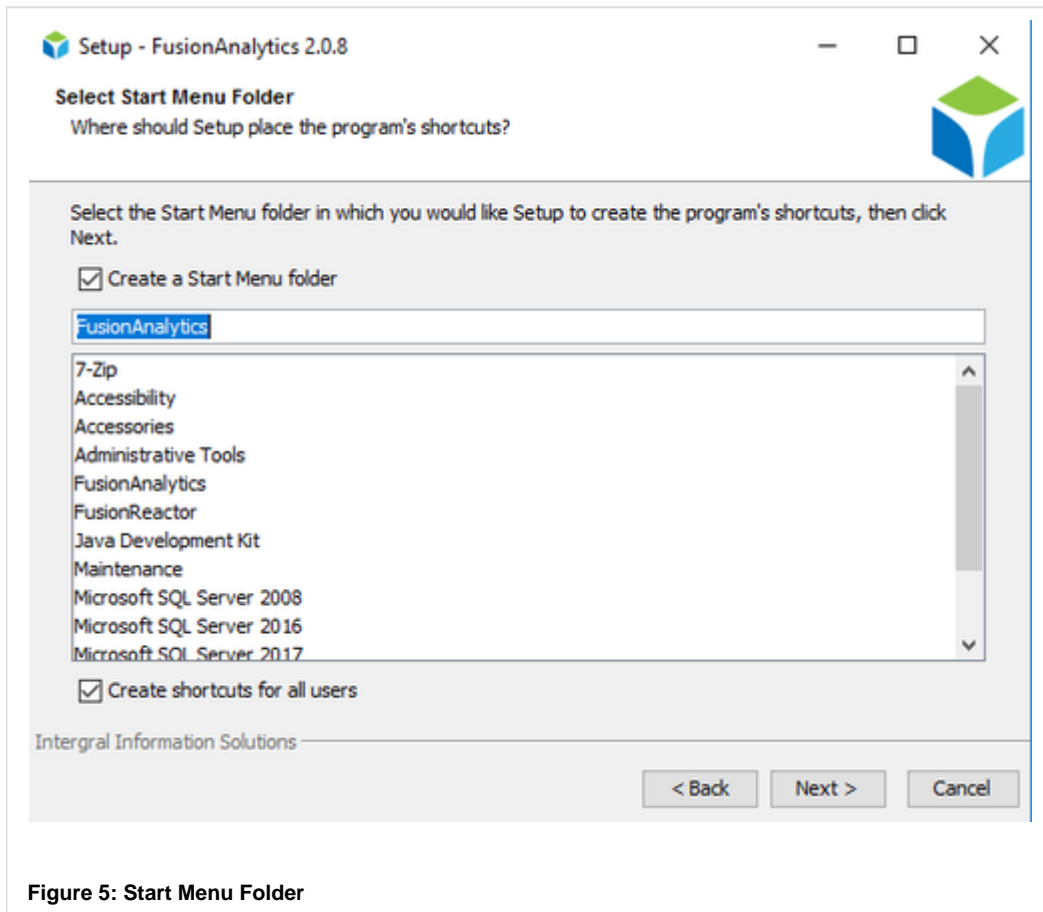


Figure 5: Start Menu Folder

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.

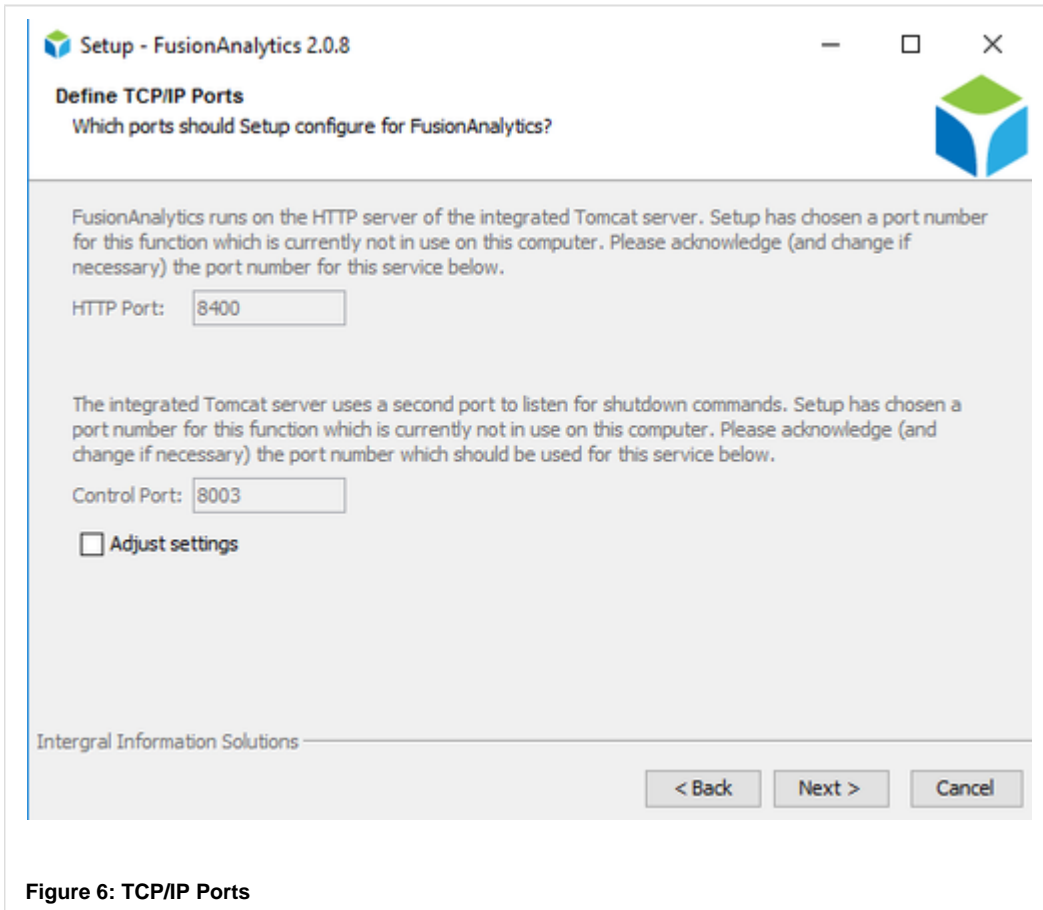



Figure 6: TCP/IP Ports

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  **Unknown macro: 'tooltip'**.

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 user account will have access to everything within FusionAnalytics.

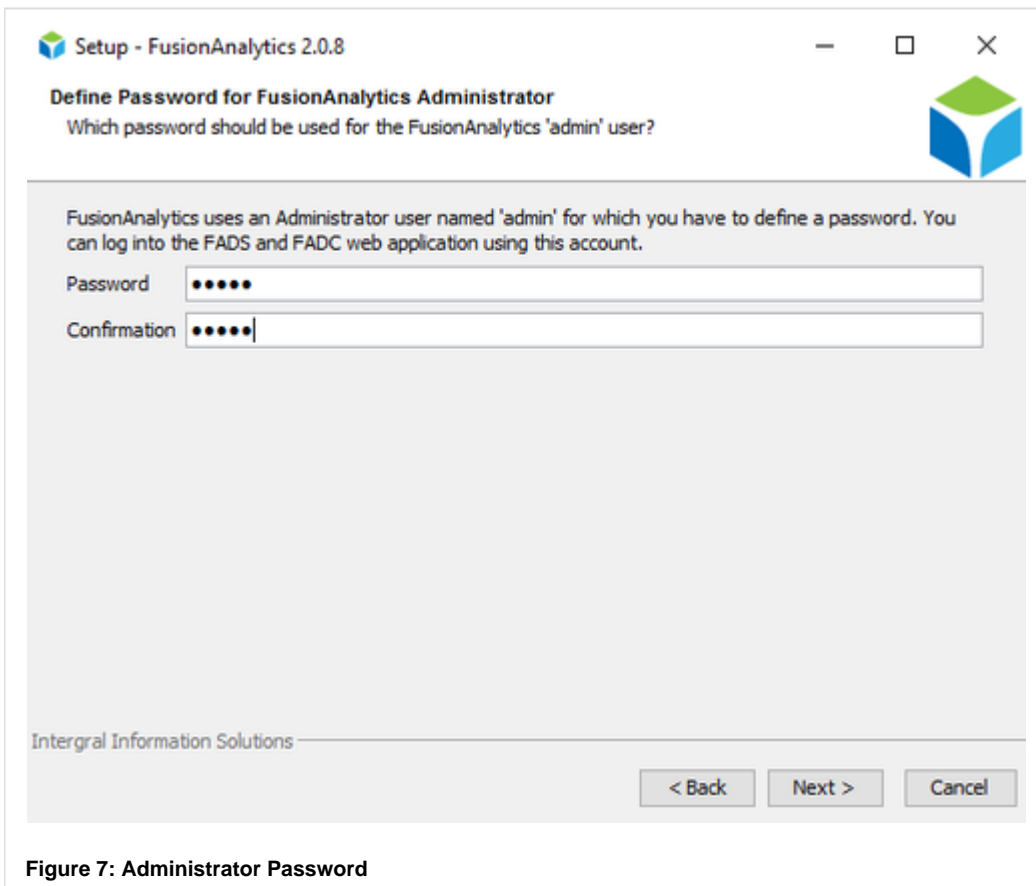


Figure 7: Administrator Password

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 **ColdFusion 11**, as this is what we will be monitoring.*

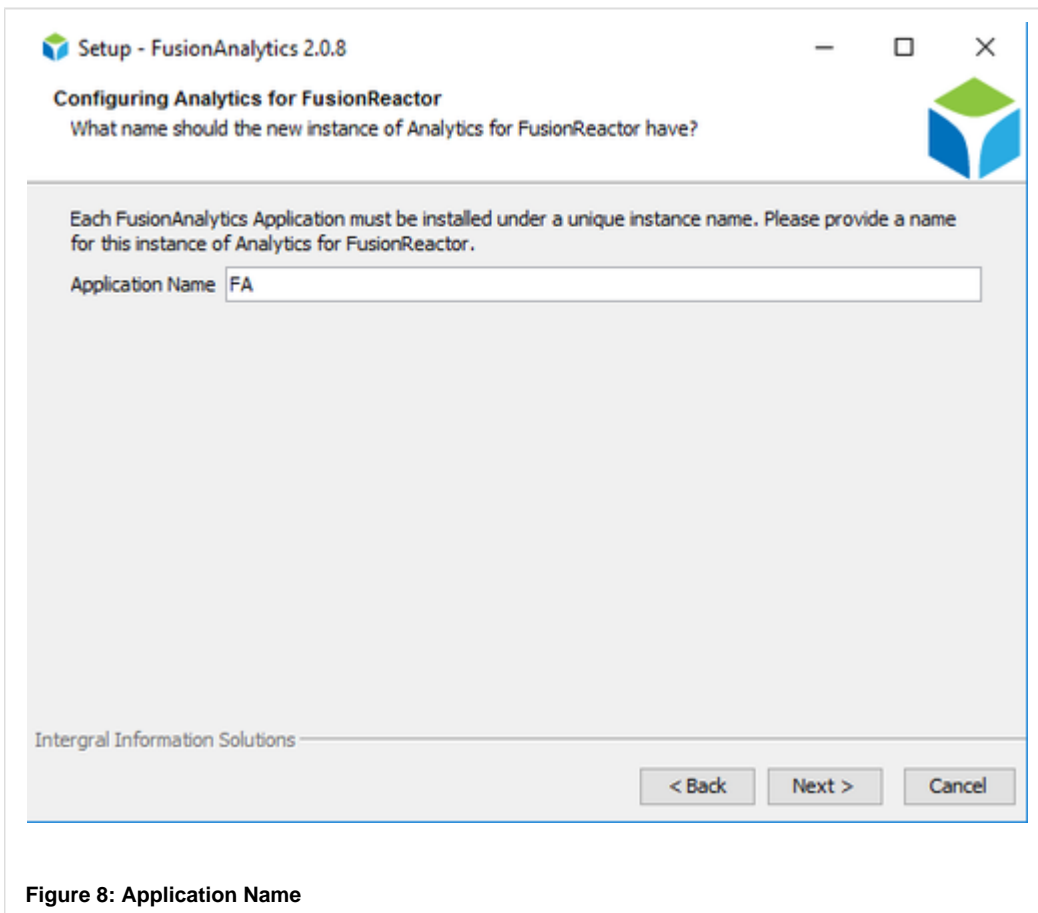


Figure 8: Application Name

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.

In FusionReactor 2.0.8 the driver options field has been added.

This by default will contain 2 options:

- loginTimeout - This reduces the login timeout from 30 seconds to 5 seconds for when a connection to the database is attempted
- sslProtocol - This allows a user to specify the TLS protocol used when connecting to an MSSQL 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.*

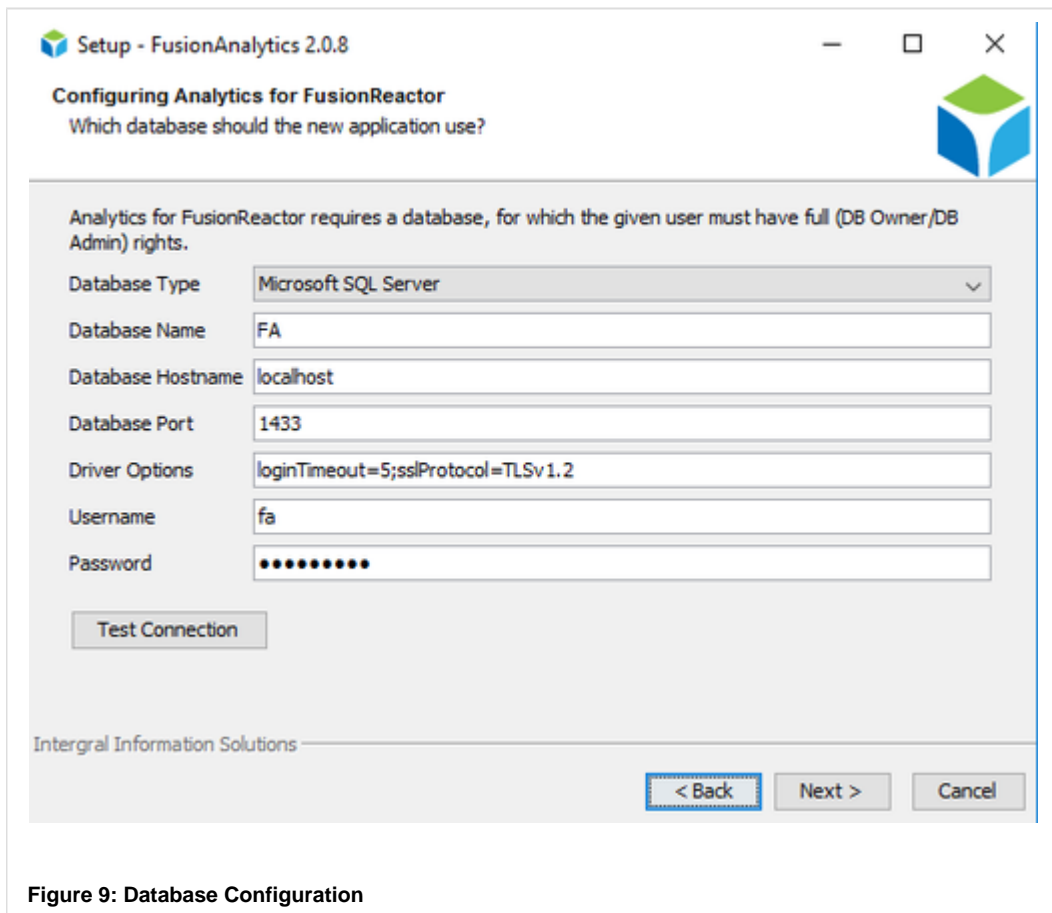


Figure 9: Database Configuration

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**.

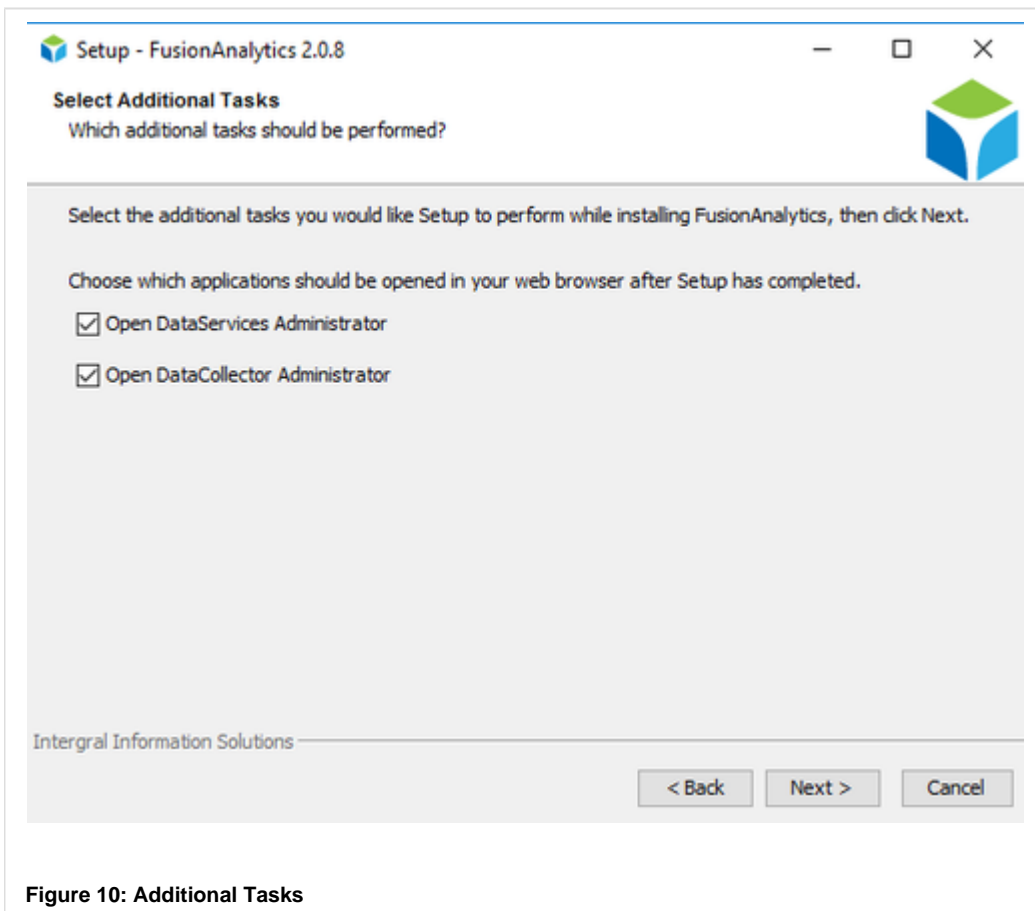


Figure 10: Additional Tasks

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.

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

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 30 days.

Note: Please keep in mind that if you have purchased an **S Pack** license key, you will only be able to monitor **ONE Application Server Instance**. For example, if you have a ColdFusion 11 and a ColdFusion 2016 installed in the same box, you have installed the latest FusionAnalytics version and you have purchased an S Pack, then you will only be able to monitor the ColdFusion 11 server **OR** the ColdFusion 2016 server. If you want to monitor both of them using the **SAME BOX**, you need to upgrade your license key.

More information about pricing can be found [here](#).

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 and by navigating to **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.

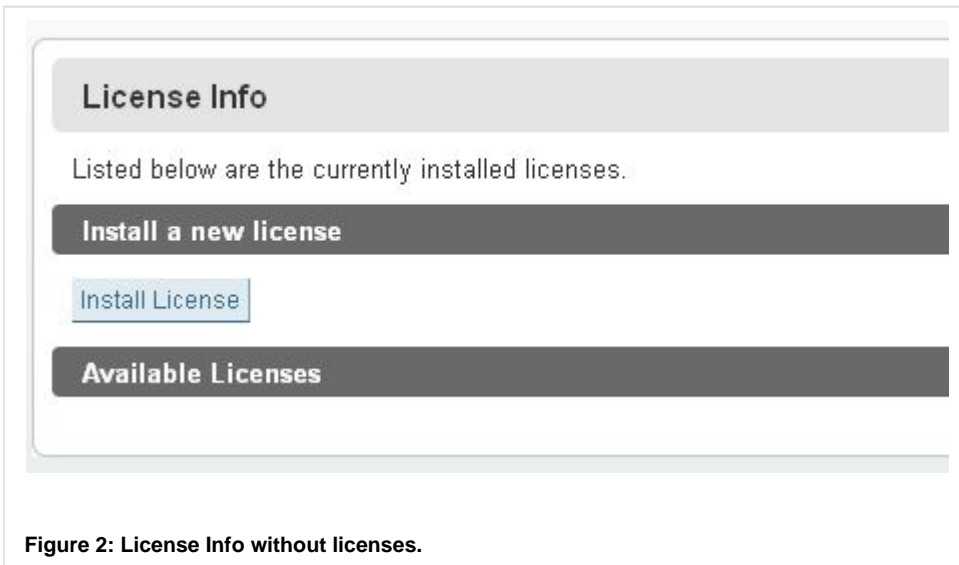


Figure 2: License Info without licenses.

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**.

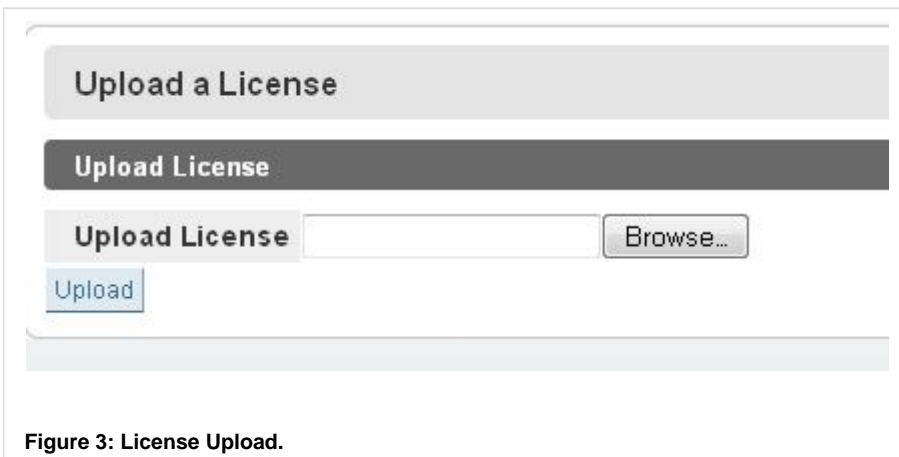


Figure 3: License Upload.

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

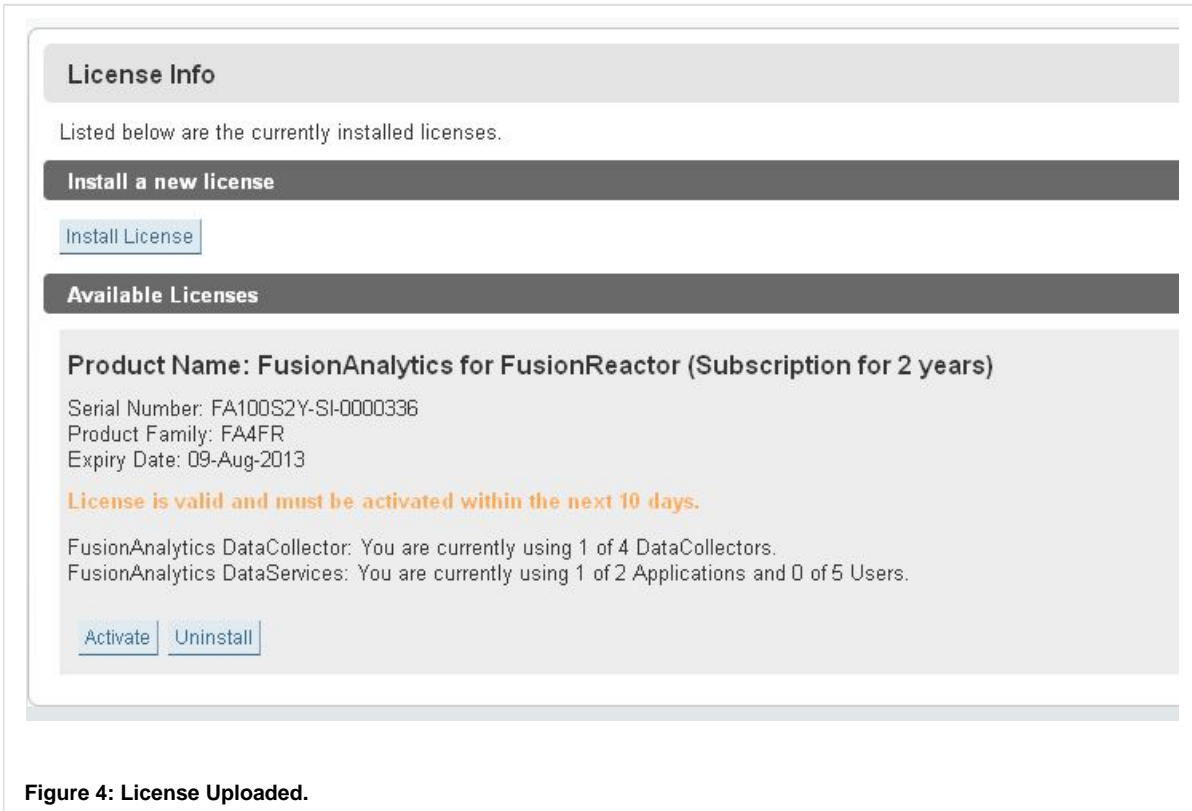


Figure 4: License Uploaded.

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 make sure that you are using a valid license file. In case you are using a valid license key please contact the FusionReactor Support Team support@fusion-reactor.com for further help.

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.

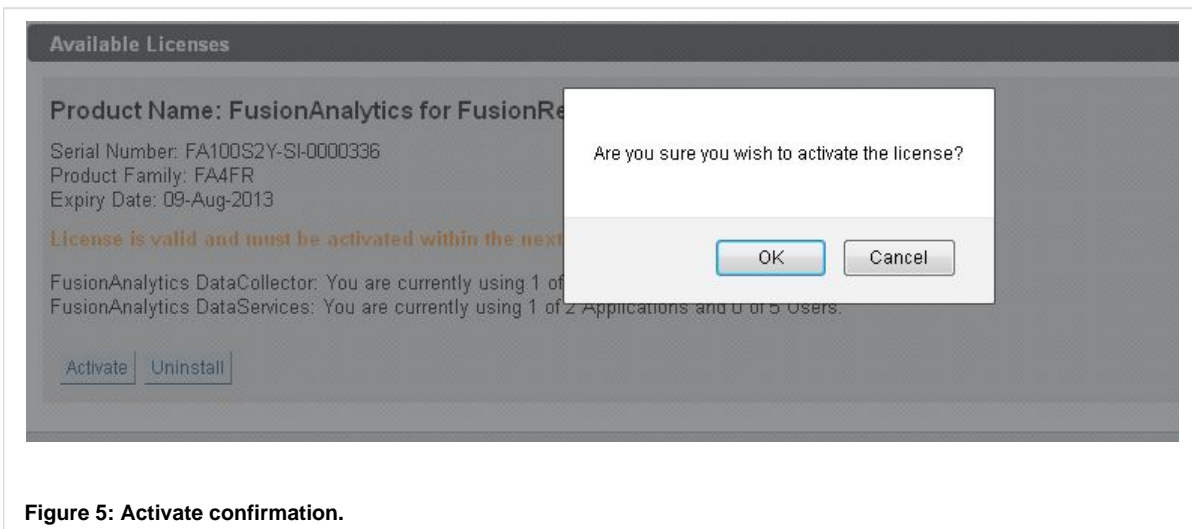


Figure 5: Activate confirmation.

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

Manual Activation

If the server running FusionAnalytics is **not connected to the Internet**, you will be taken to the manual activation page of FusionAnalytics.

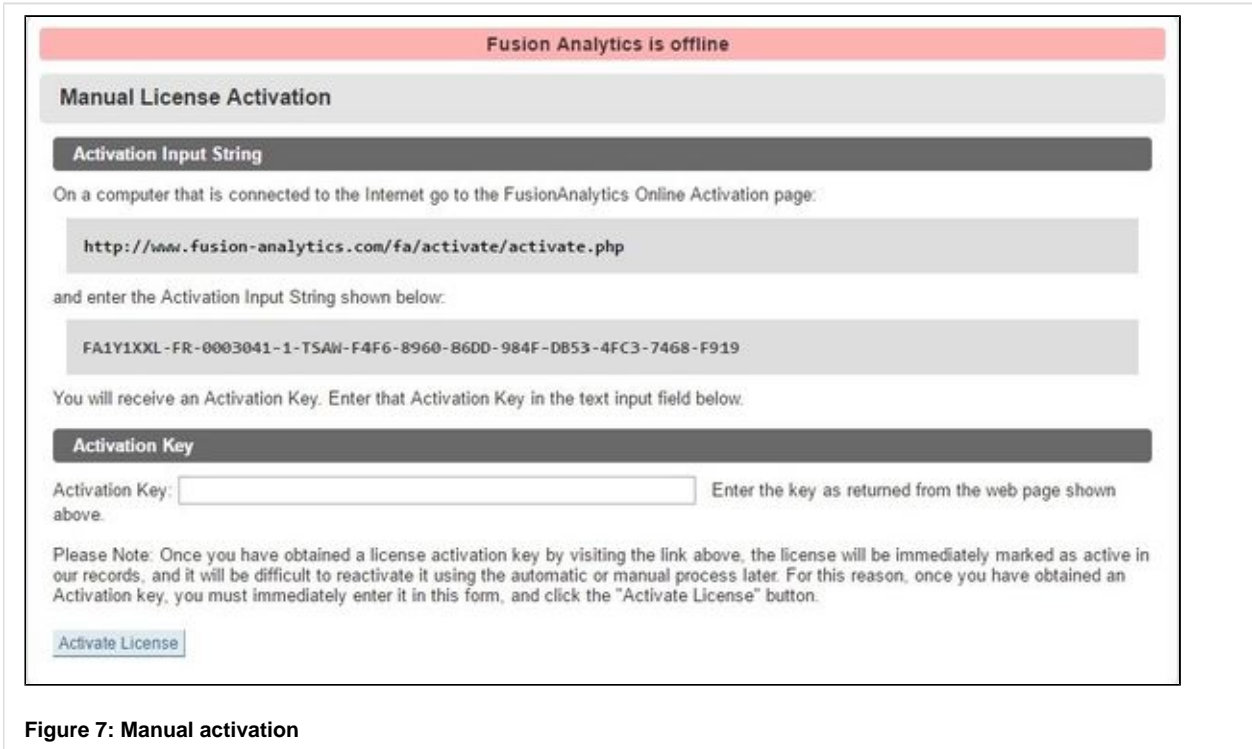
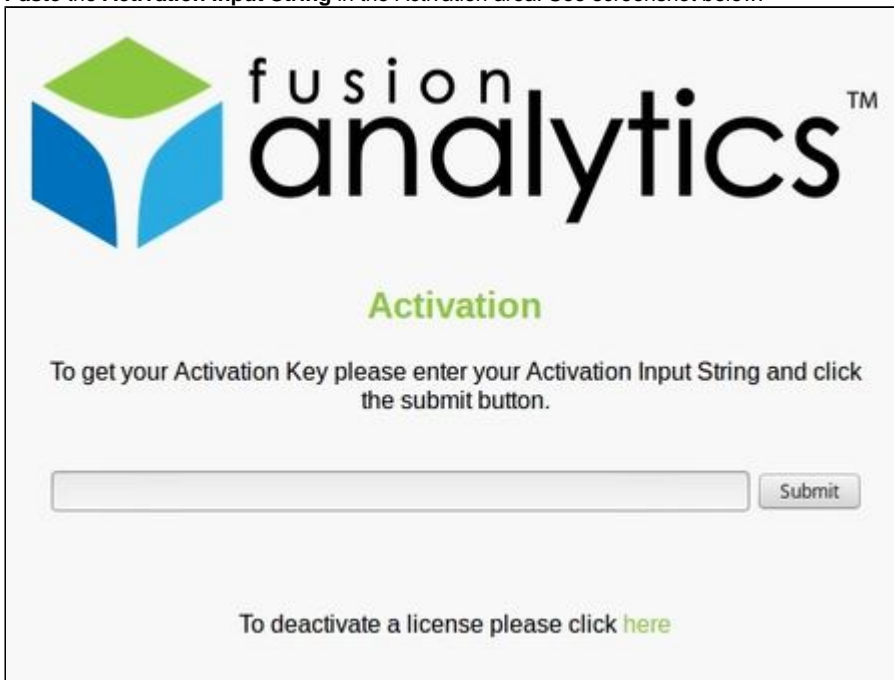


Figure 7: Manual activation

The steps below are going to guide you through the manual activation process.

1. Access the manual license activation page of FusionAnalytics **using a computer that can access the Internet**. The URL is the following, Manual License Activation.
2. Make sure that you have **copied** the whole **Activation Input String**.
3. **Paste** the **Activation Input String** in the Activation area. See screenshot below.



4. Press the **Submit** button.
5. A new page will appear and under the field with name "**Activation Key:**" you will find an Activation Key.

6. Copy the Activation Key.
7. Paste the key in the Activation Key area.

Activation Key

Activation Key: Enter the key as returned from the web page shown above.

Please Note: Once you have obtained a license activation key by visiting the link above, the license will be immediately marked as active in our records, and it will be difficult to reactivate it using the automatic or manual process later. For this reason, once you have obtained an Activation key, you must immediately enter it in this form, and click the "Activate License" button.

[Activate License](#)

8. Press the "Activate License" button.

License activated straight away

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

License activated successfully

License Info

Listed below are the currently installed licenses.

Install a new license

[Install License](#)

Available Licenses

Product Name: FusionAnalytics for FusionReactor (Subscription for 2 years)

Serial Number: FA100S2Y-SI-0000336
Product Family: FA4FR
Expiry Date: 09-Aug-2013

License is valid and activated.

FusionAnalytics DataCollector: You are currently using 1 of 4 DataCollectors.
FusionAnalytics DataServices: You are currently using 1 of 2 Applications and 0 of 5 Users.

[Uninstall](#)

Figure 6: License activated.

Your license is now fully activated.

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.

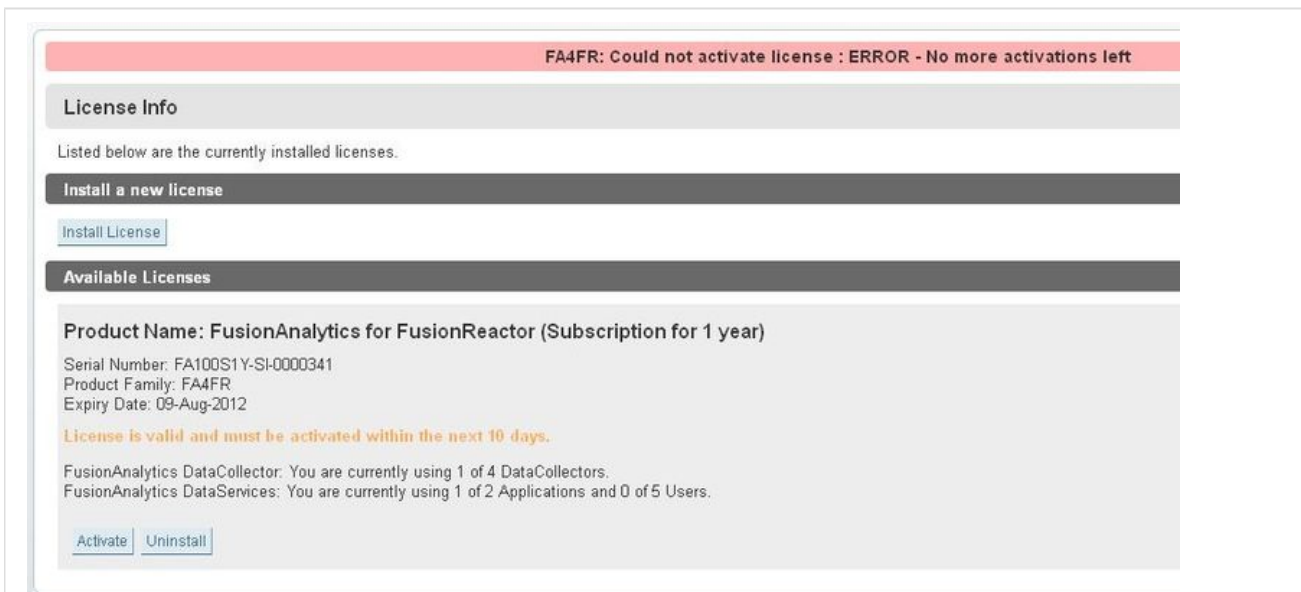


Figure 8: No activations remaining

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

Introduction

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

On this page:

About Adobe AIR

Executable

Setup

Installing the Adobe AIR Client

| Executable | System Architecture |
|-------------------------------|---------------------|
| fusionanalytics-air-2.0.7.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).

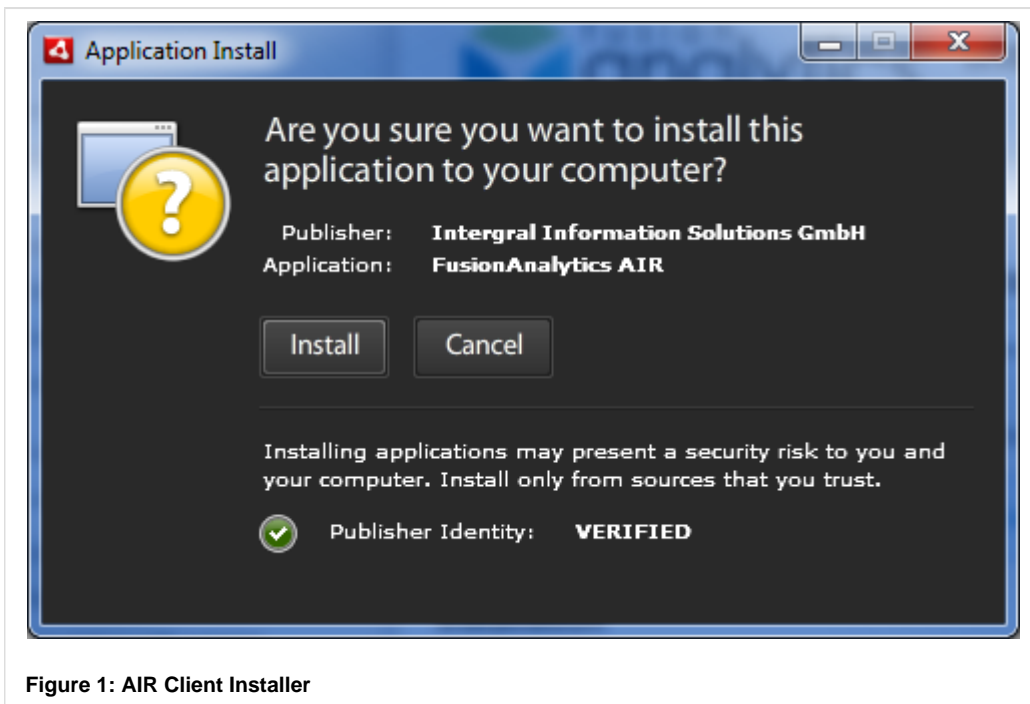


Figure 1: AIR Client Installer

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.

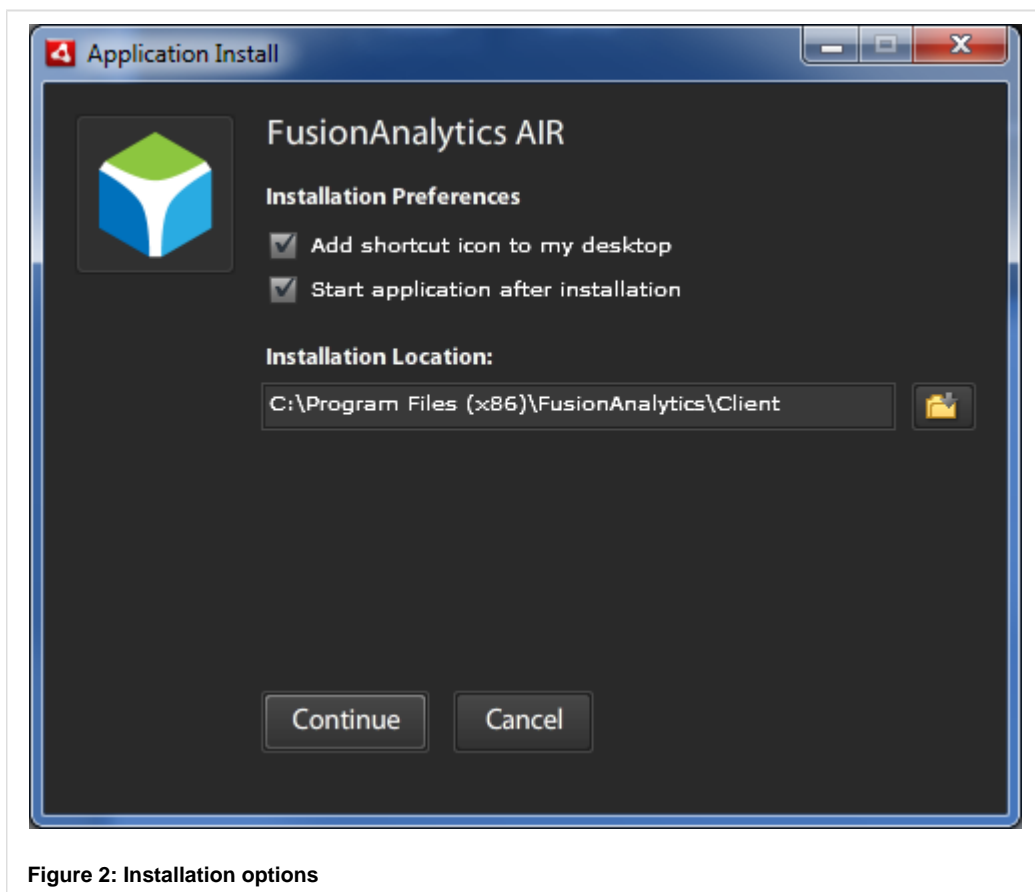


Figure 2: Installation options

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 application 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 5 or greater version. Using them together will produce highly detailed log files for your server and allow you to analyze the results.

- Check out the link, [FusionReactor Pricing](#) for more information and details about pricing. There is a 14-day trial period, enabling you to check out all the functionality of the software before purchase.
- Check out the link, [FusionReactor installation guide](#) for help on installing FusionReactor. Once installed, FusionReactor 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

Please note that this section refers to FusionReactor 4.x versions. If you are using FusionReactor 5.x and greater, there is no need to install the Extensions for ColdFusion as this is done automatically for you.

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.

For a detailed installation tutorial, please see the FusionReactor documentation: FusionReactor ColdFusion Plugin. More information about the installation process of the extensions can be found here.

Uninstall a FusionAnalytics License Key

Introduction

This section will provide readers with more information about the process of successfully uninstalling a FusionAnalytics license key.

If you want to activate a FusionAnalytics License key, please follow the link below.

- [Install a FusionAnalytics License Key](#)

Getting Started

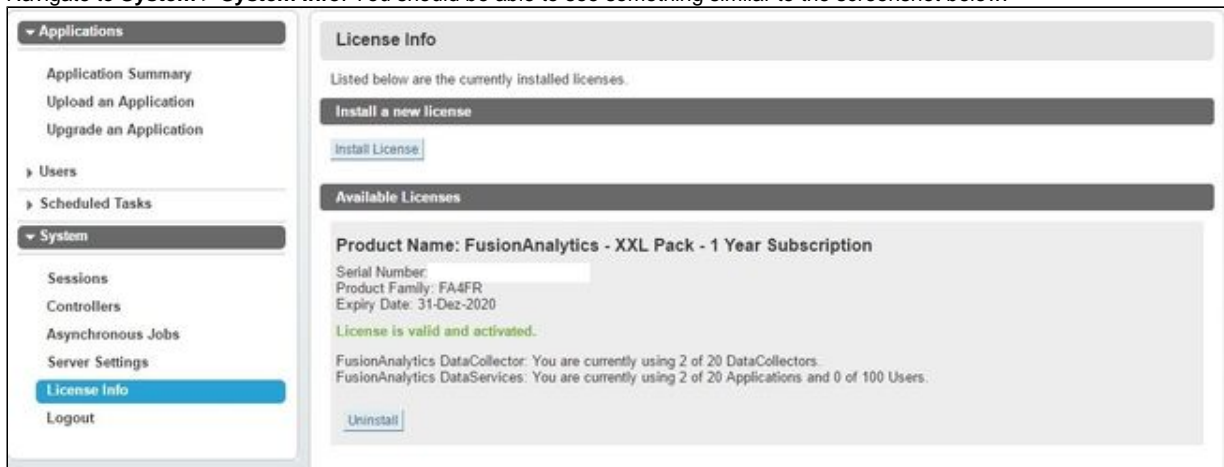
There are a number of reasons why you might want to uninstall a FusionAnalytics license key. Some of the most common reasons described below.

1. You want to monitor more Application servers than the size of your license key. For that reason, you need to uninstall the current license key and activate a new one.
2. You want to uninstall the license key from one machine in order to free an activation and use the license key on another machine.

UNINSTALLING THE LICENSE KEY

The process of uninstalling a license key in FusionAnalytics is fast and simple. In order to do that, please follow the steps listed below.

1. Access the FADS **OR** FADC application via a browser.
2. Navigate to **System > System Info**. You should be able to see something similar to the screenshot below.

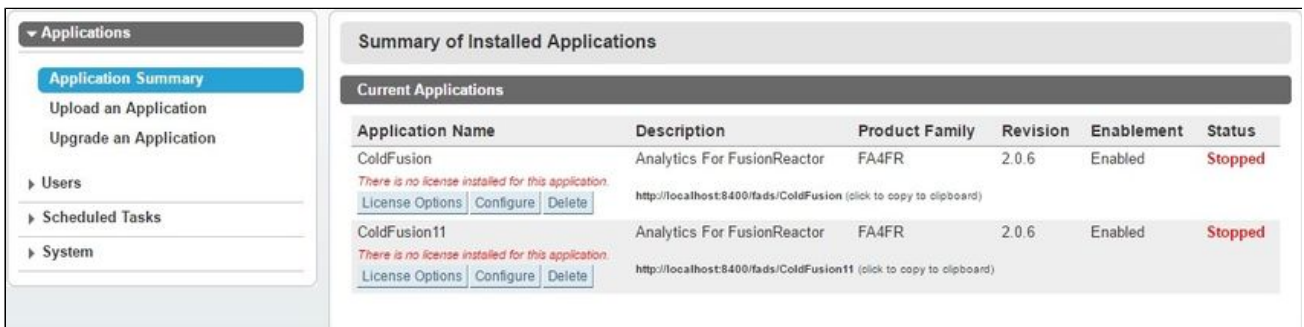


3. Click the **"Uninstall"** button.
4. In the pop-up message, click OK in order to uninstall the license key.

5. If the process has been successful, you will be able to see the following message "License Uninstalled Successfully". See screenshot below.



If you try and refresh the **FADS** page, you will be redirected by default to the "Application Summary" page and there you will be able to see the red message "**There is no license installed for this application.**" indicating that the license key has been uninstalled. See screenshot below.



Also, if the license key has been deactivated successfully, an activation seat will be available and you should be able to use the license key in another machine.

See also

- Uninstalling FusionAnalytics
- Upgrading FusionAnalytics

Uninstalling FusionAnalytics


Introduction

This section will guide readers through the process of **uninstalling** FusionAnalytics from a machine.

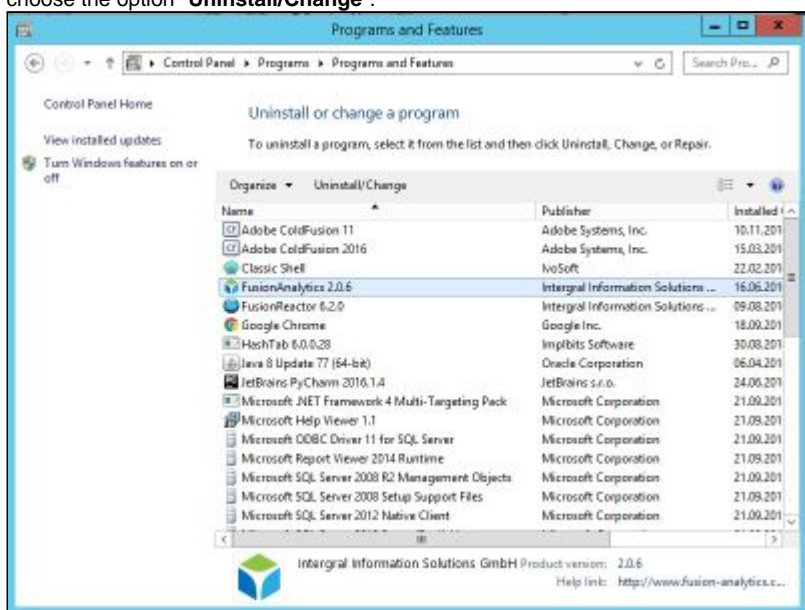
Uninstalling FusionAnalytics

If for any reason you need to uninstall FusionAnalytics, this can be done with two different ways which are described below.

1. Navigate to the FusionAnalytics directory and double click the "uninstall" icon. Alternatively, you can type in the Search bar of Windows "FusionAnalytics Uninstaller" and the uninstaller will appear.



2. You can navigate to **Control Panel > Programs > Uninstall a program**. Choose the FusionAnalytics 2.0.6 icon, right click and choose the option "**Uninstall/Change**".



Note

- **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.

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

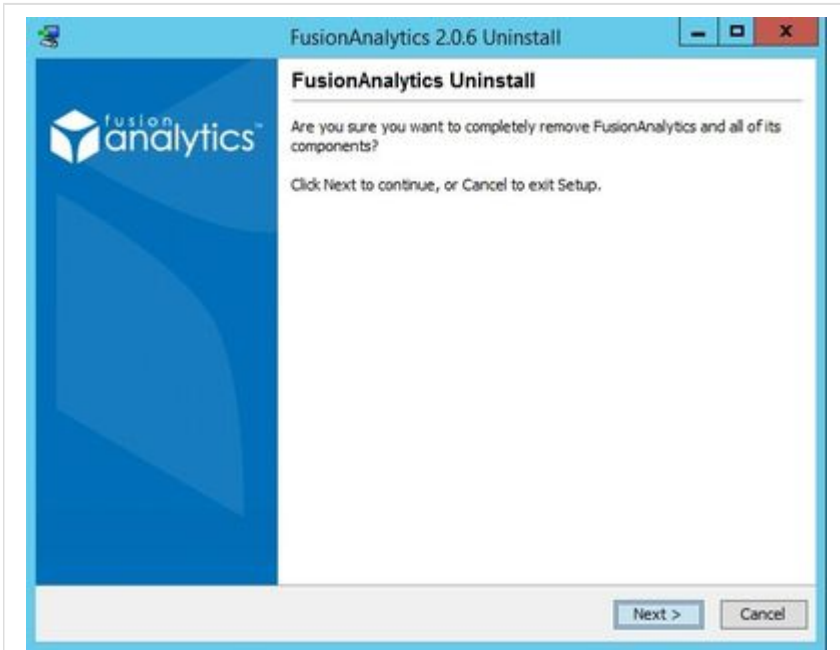


Figure 2: Uninstaller Wizard

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 the **Finish** button in order to close the uninstaller wizard.

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.

Upgrading FusionAnalytics

Introduction

This section is going to guide you through the process of Upgrading FusionAnalytics.

Upgrading FusionAnalytics

UPGRADING FUSIONANALYTICS SERVER

To upgrade your current version of FusionAnalytics (FA), simply download and run the latest FA installer (this can be downloaded from our website).

Make sure that the latest installer version is newer (greater) than your current installed version of FA.

The installer will detect your current version of FA, and will upgrade it accordingly. The setup will restart the FA services during the upgrade process and any existing applications will be kept.

However, once the upgrade is complete you will be required to manually upgrade your existing (FADS and FADC) applications, if required, therefore when the

On this page:

Upgrading FusionAnalytics Server

Upgrading FADS Applications

Upgrading FADC Applications

Server Restart Required

Whenever an APM or DCML application has been upgraded the FusionAnalytics server has to be restarted. This will ensure that any locked files can be

installer will ask which applications should be opened after the setup has been completed, select both "**Open DataServices Administrator**" and "**Open DataCollector Administrator**":

released and new files are loaded correctly.

After the update process, the **DataServices (FADS) Administrator** and **DataCollector (FADC) Administrator** will be opened in the browser.

[Back to the top](#)

UPGRADING FADS APPLICATIONS

Once FADS has been launched (if FADS is not launched please navigate to FADS and) log into the Administrator.

Before you can upgrade an application, you have to **STOP the application(s)** that you want to upgrade by clicking the "Stop" button for the relevant application.

Once the application(s) has been stopped, click the "Upgrade" button to upgrade the application:

Click the browse button and navigate to the package you would like to upgrade to:

Navigate to your FA installation Directory, and open the "Applications" folder (**FusionAnalytics > Applications**). In this folder you should see two applications (as compressed zip folders). These are the applications for the selected installed / upgraded instance of FA. Select the **DSAPP** (e.g. fusionreactor-fads-X.X.X-DSAPP.zip).

Once you have selected the application upgrade package, click the "Upgrade" button:

Once the application has been successfully upgraded, **please make sure that you restart the FusionAnalytics server before you can use the upgraded application.**

For a more detailed description regarding upgrading FADS applications, please see: [Upgrade an Application \(FADS\)](#).

[Back to the top](#)

UPGRADING FADC APPLICATIONS

Once FADC has been launched (if FADC is not launched please navigate to FADC and) log into the Administrator.

Note
Before you can upgrade an application, you have to **STOP the application(s)** that you want to upgrade by clicking the "Stop" button for the relevant application.

Once the application(s) has been stopped, click the "Upgrade" button to upgrade the application:

Click the browse button and navigate to the package you would like to upgrade to.

Navigate to your FA installation Directory, and open the "Applications" folder (FusionAnalytics > Applications). In this folder you should see two Applications (as compressed zip folders). These are the applications for the selected installed / upgraded instance of FA. Select the **DCAPP** (e.g. fusionreactor-fads-X.X.X-DCAPP.zip).

Once you have selected the application upgrade package, click the "Upgrade" button:

Once the application has been successfully upgraded, **please make sure that you restart the FusionAnalytics server before you can use the upgraded application.**

For a more detailed description regarding upgrading FADS applications, please see: Upgrade an Application (FADC).

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

SETTING UP THE SLA 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 FusionAnalytics currently ONLY supports Microsoft SQL Server.

Supported versions are: 2005, 2008, 2008 R2, 2012, 2014 and 2016

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.*

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.

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**



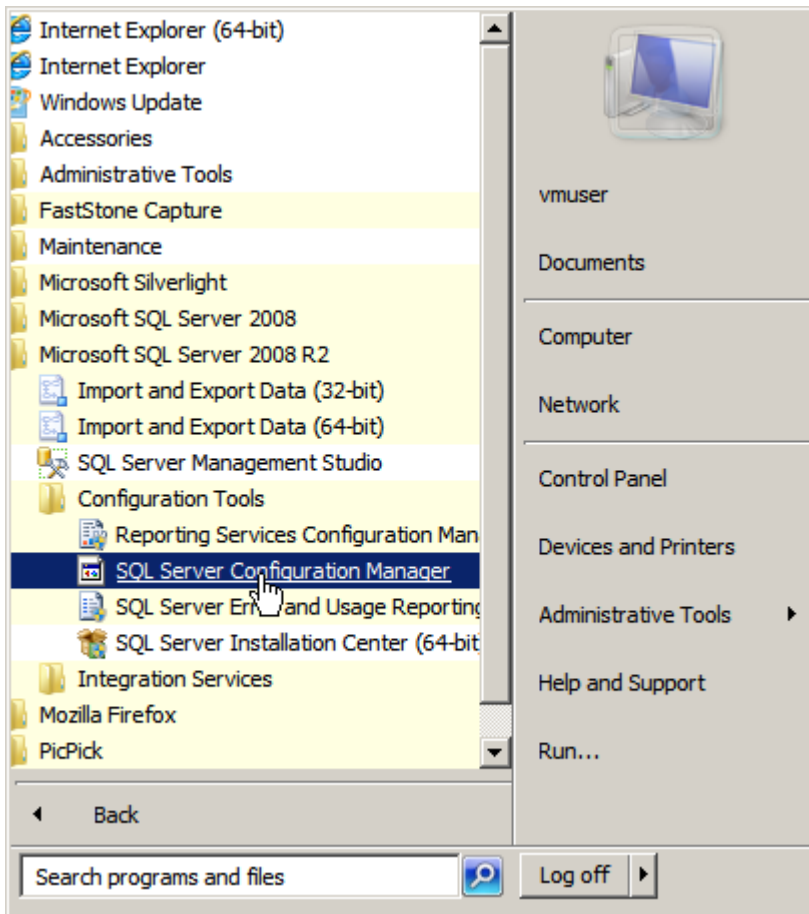


Figure 1: Starting SQL Server Configuration Manager

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

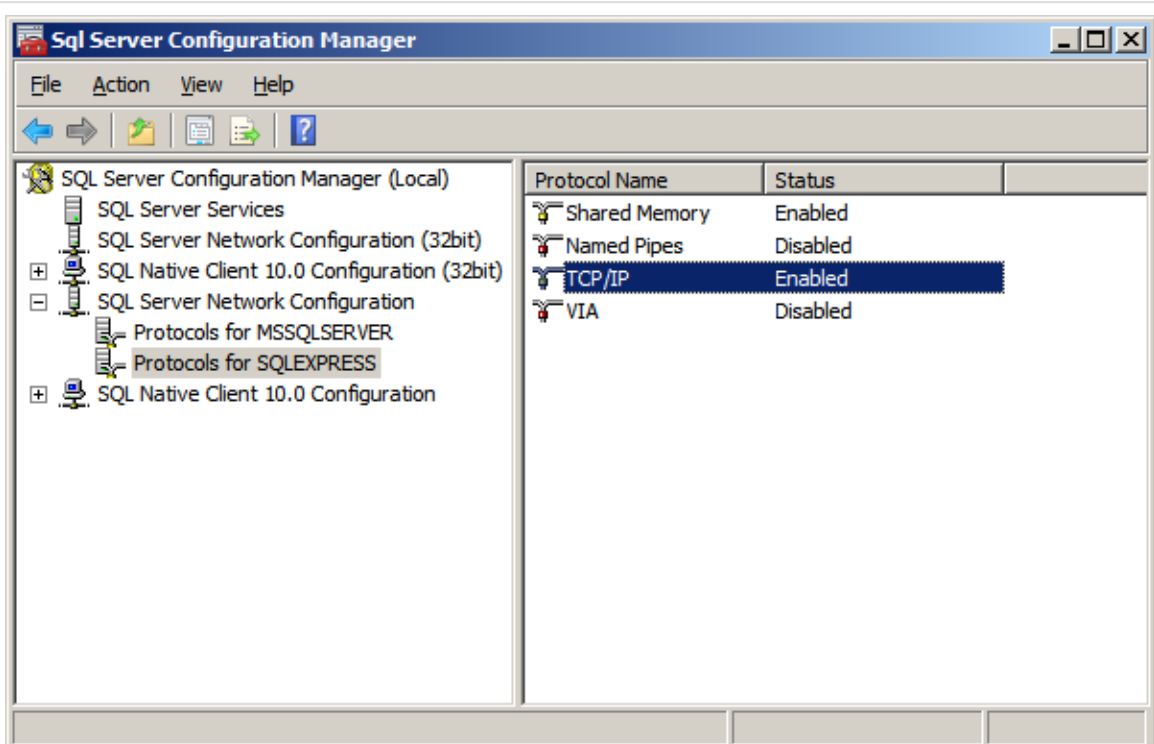


Figure 2: Configuring Protocols

In the TCP/IP Properties dialog set Enabled to Yes

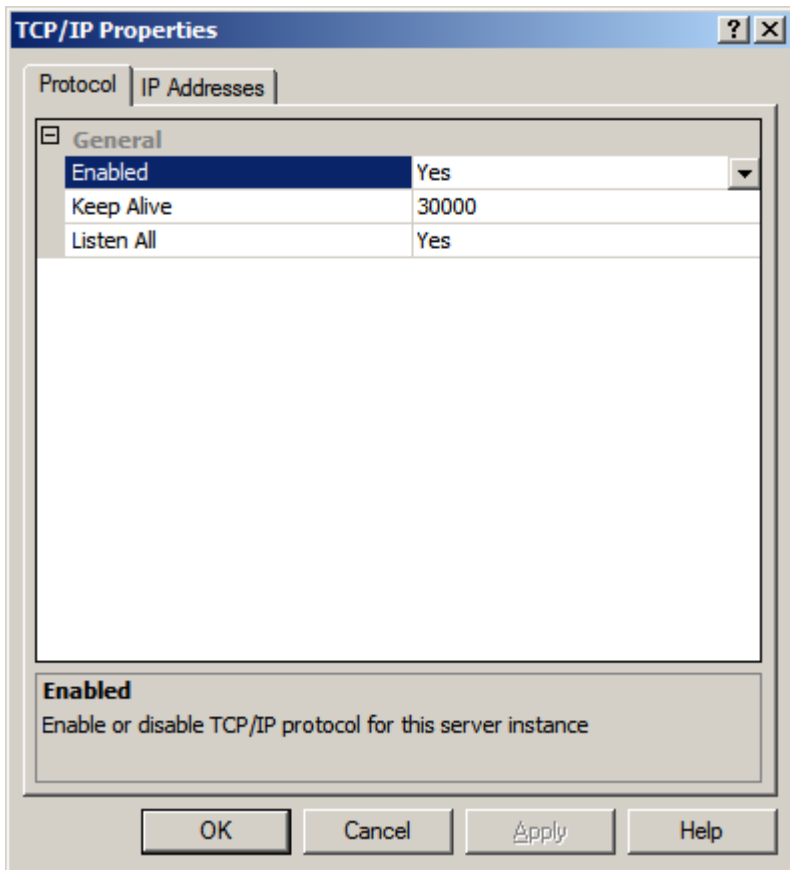


Figure 3: Enabling TCP/IP

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

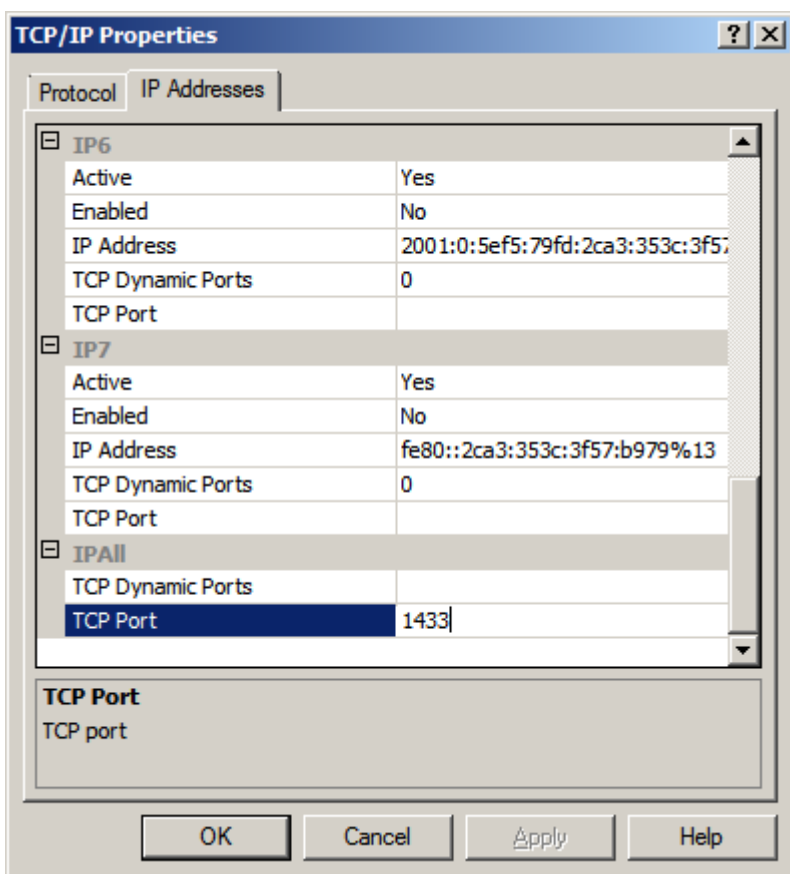


Figure 4: Configuring TCP Port

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

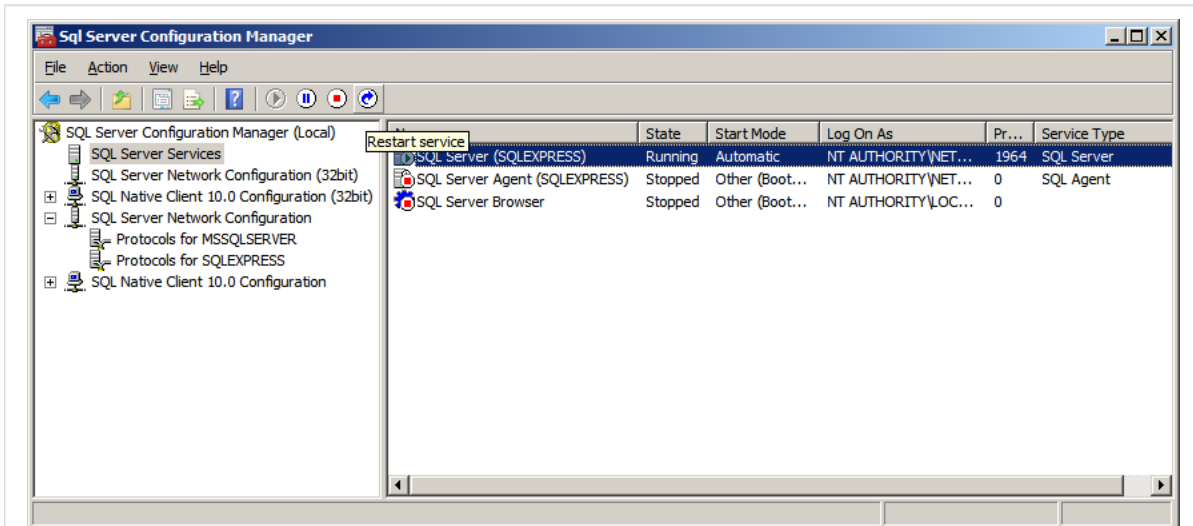


Figure 5: Restarting SQL Server

[Back to the top](#)

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

Access the MSSQL 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 login.

Enabling SQL Server Authentication

On This Page:

Enabling SQL Server Authentication

[Create a New User](#)

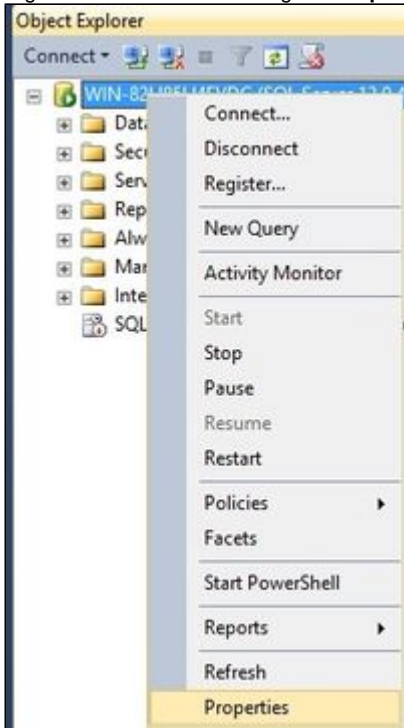
More Database Configuration Tutorials:

[Configure Network Settings](#)

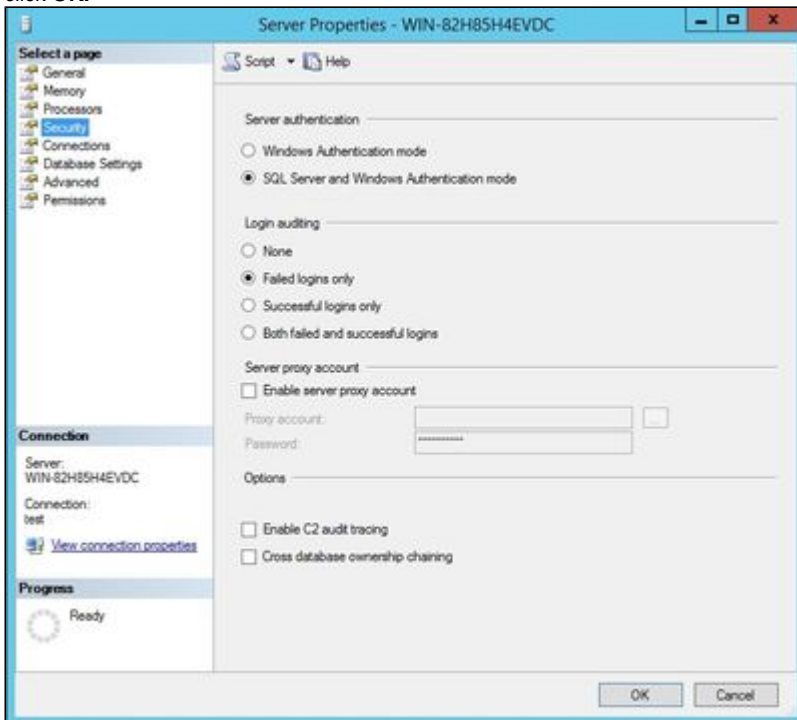
[Create a New Database Login](#)

[Create a New Database](#)

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.
i 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 the **Server Roles** page.
8. Make sure **'sysadmin'** is checked. (this is not required on SQL Server 2008 Standard RC2)
 -  This is required by FusionAnalytics to perform some of its tasks.
9. Click **OK** to create the user.

Some tasks within FusionAnalytics require the user to have **'sysadmin' permissions** (this is not required on SQL Server 2008 Standard RC2). Once your user is created:

Check if the 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

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

The steps below will guide through the process of creating a new Database which will be compatible with the System Requirements of FusionAnalytics.

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...**

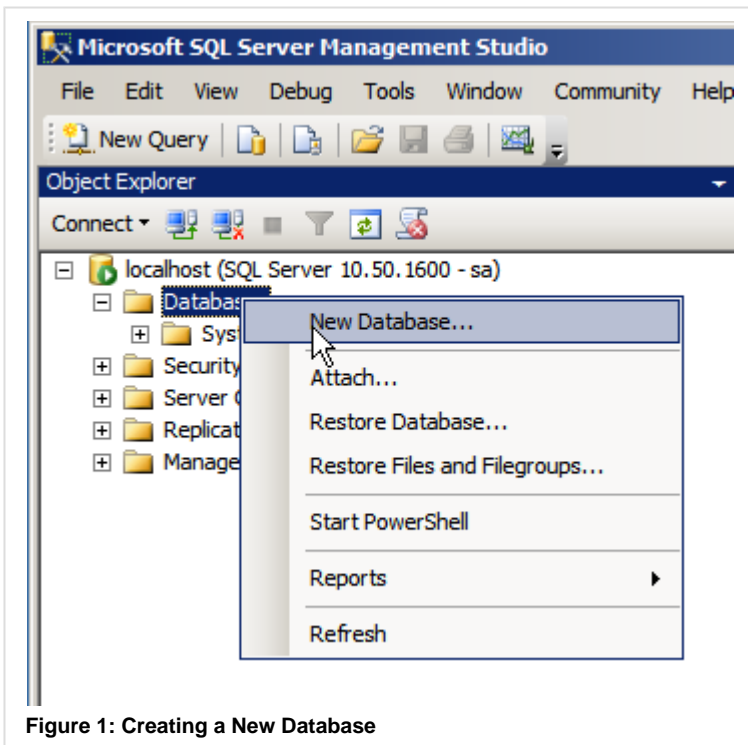


Figure 1: Creating a New Database

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

- 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.

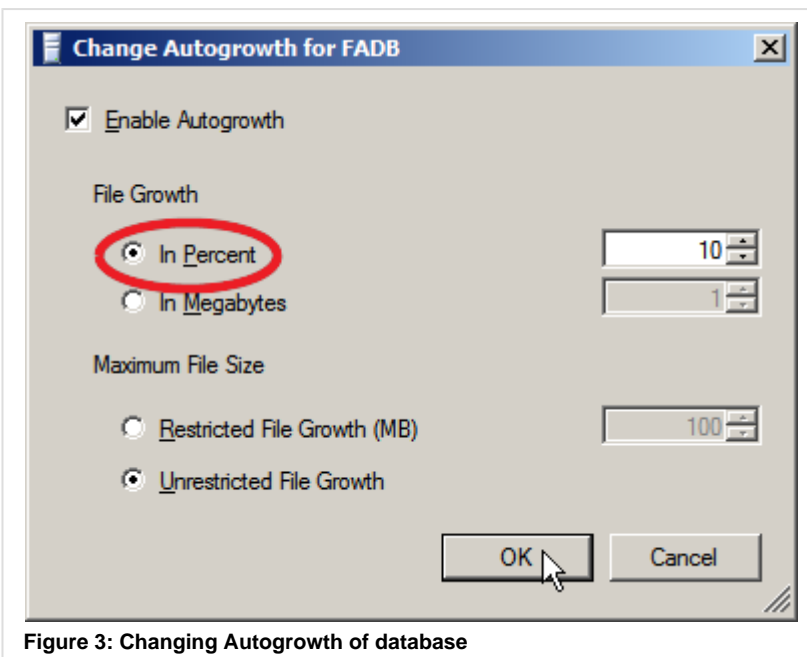


Figure 3: Changing Autogrowth of database

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

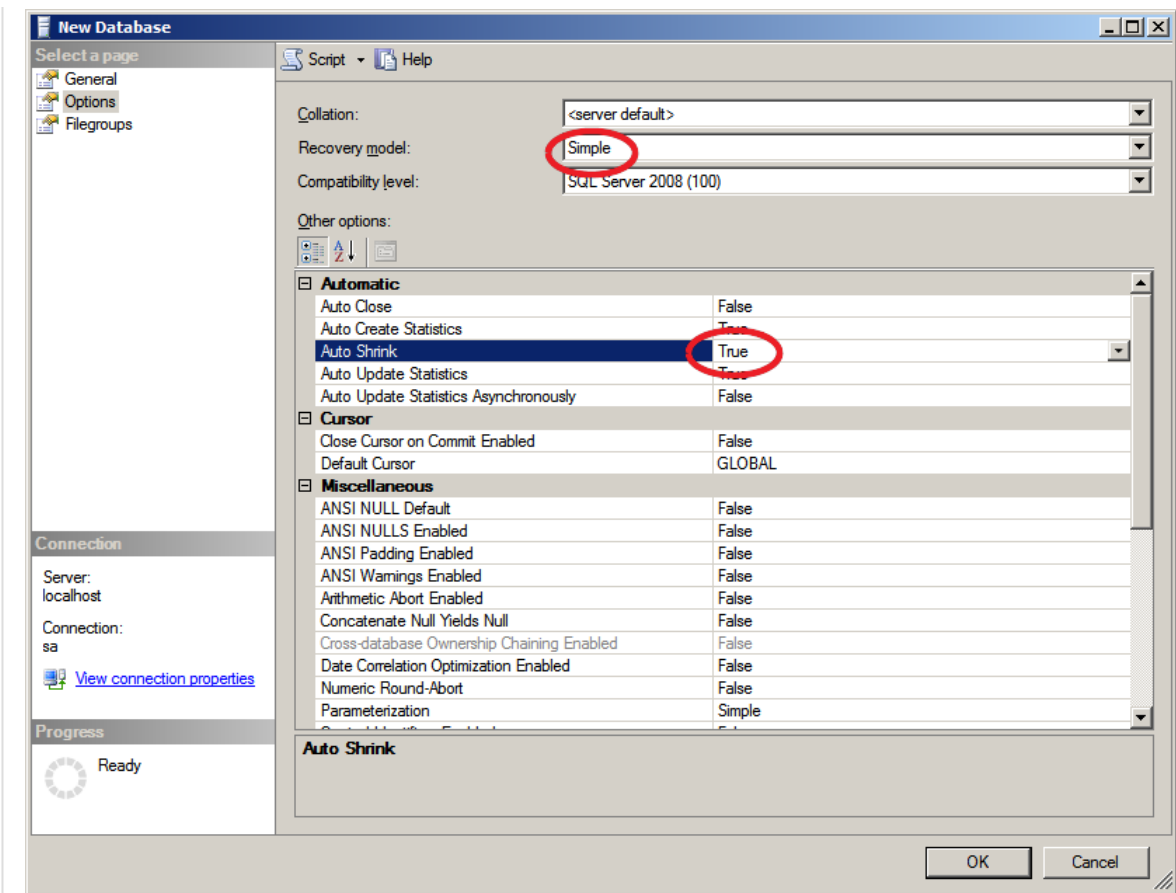


Figure 4: Configuring Recovery Model and Auto Shrink

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.



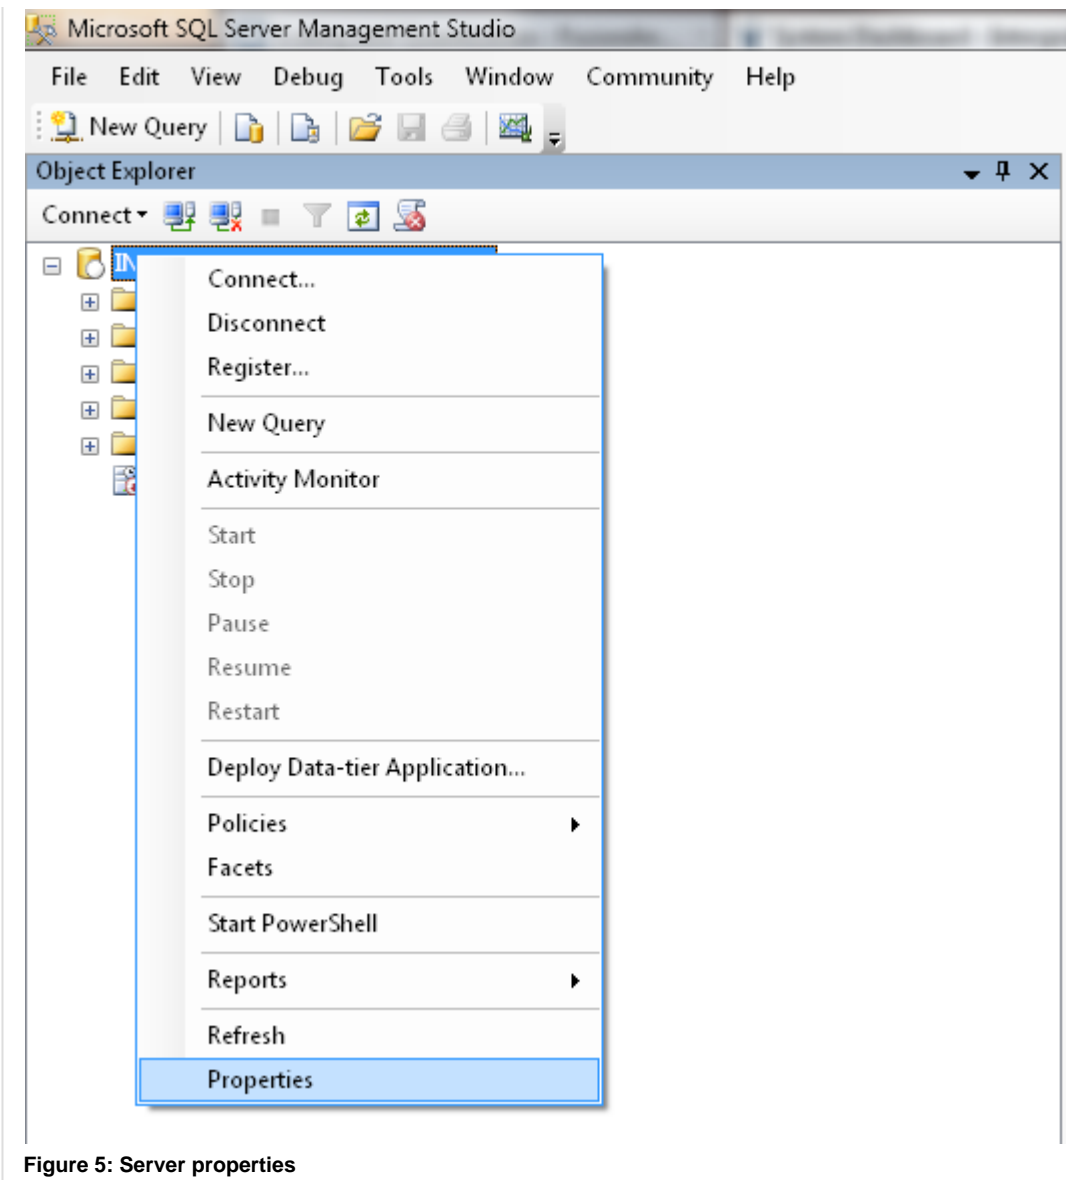



Figure 5: Server properties

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



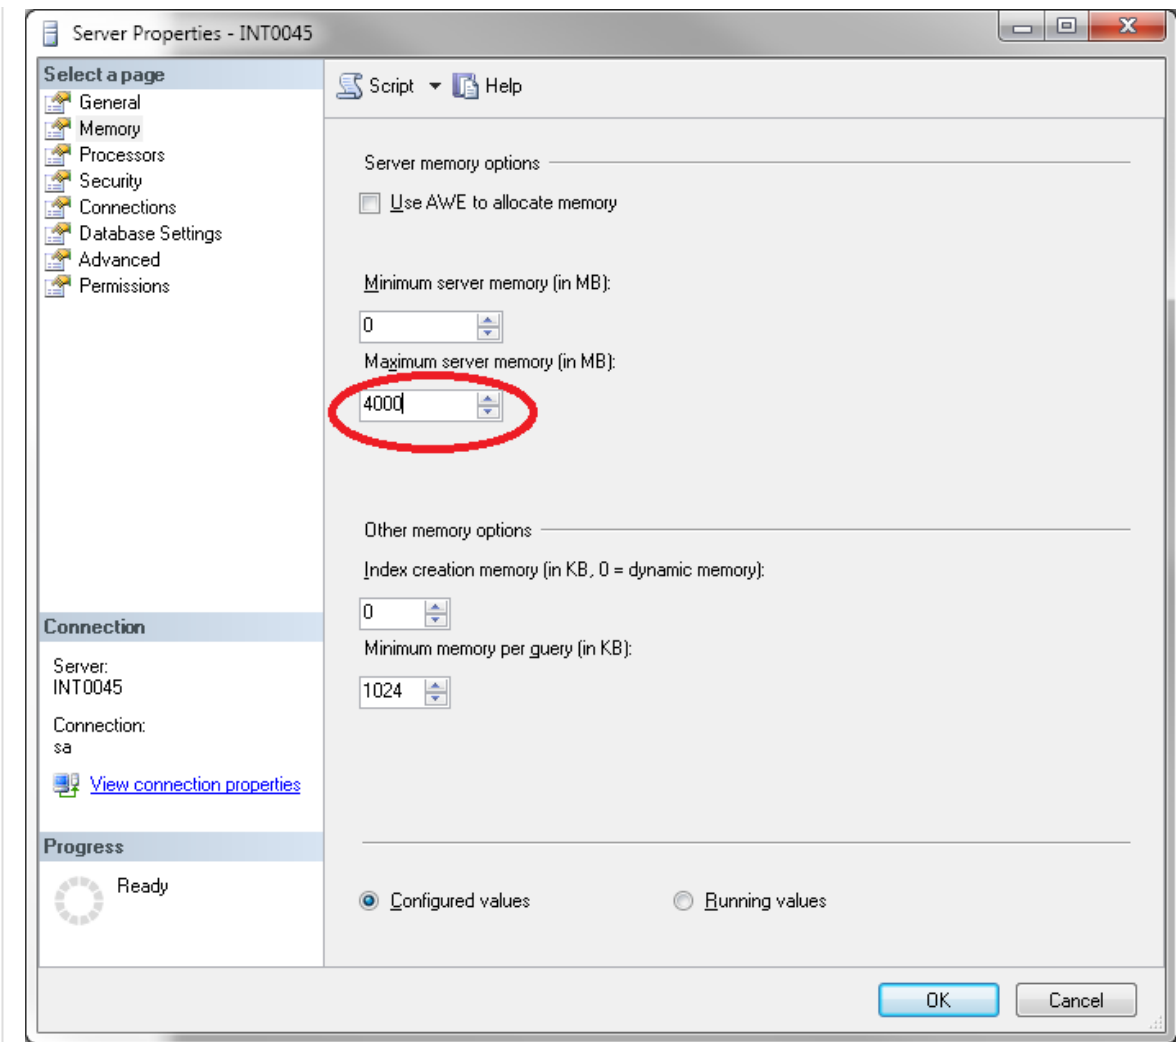


Figure 6: Allocating memory

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

[Back to the top](#)

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 the '**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>

If you are using another hostname and not the "localhost", you can change that part of the URL to the machine name that you want to access. If you changed the HTTP port number from the default (8400), then you will need to use the configured port.

[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  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.



Current Applications

| Application Name | Description |
|---|---|
| coldfusion9 Stop Restart Launch Info Scope | Analytics For FusionReactor http://127.0.0.1:8400/fads/coldfusion9 (click to copy to clipboard) |

Figure 1: Launch Application

Logging In

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

| | |
|------------------|--|
| Username: | admin |
| Password: | <i>The password you defined in the Administrator Password screen during installation.</i> |

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

[Back to the top](#)

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 port of 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, it is 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 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.

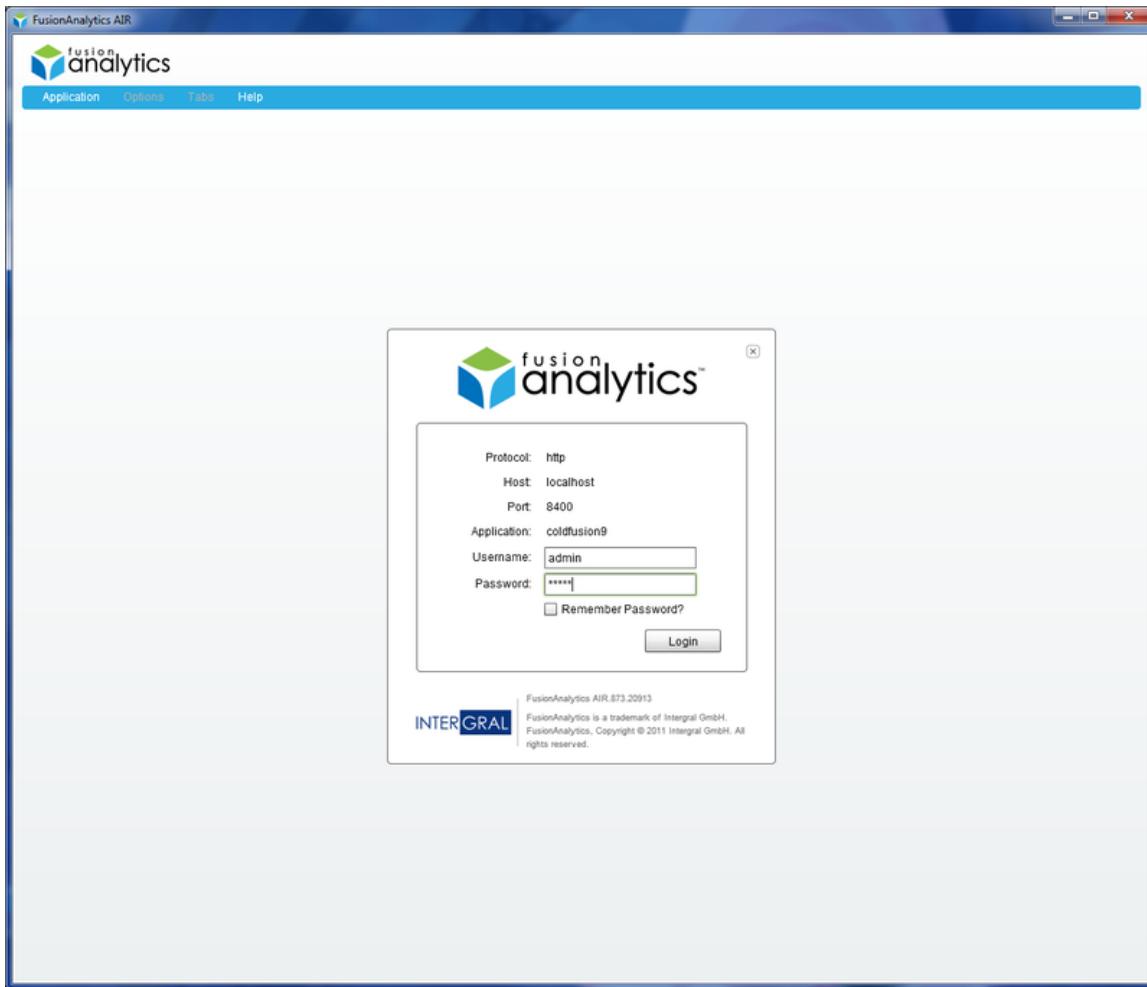


Figure 2: Login Screen

Back to the top

Alternatively

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 access the FADS application and navigate to **Applications > Application** the application that you want to view (Figure 3).



Figure 3: FADS Launch Button

Back to the top

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

Introduction

This section is going to guide you through the process of configuring FusionReactor for FusionAnalytics.

Configuring FusionReactor for FusionAnalytics

Firstly, you will need to login to the FusionReactor instance that you want to use to send data to FusionAnalytics.

To do this, navigate to **FusionReactor > Logging > Settings** (Figure 1).

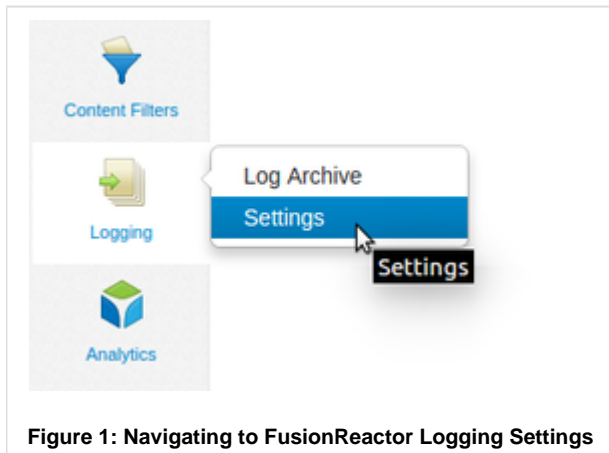


Figure 1: Navigating to FusionReactor Logging Settings

This will display the **FusionReactor Logging Settings** pop-up (Figure 2). In order to see how to configure these, go to Log Settings.

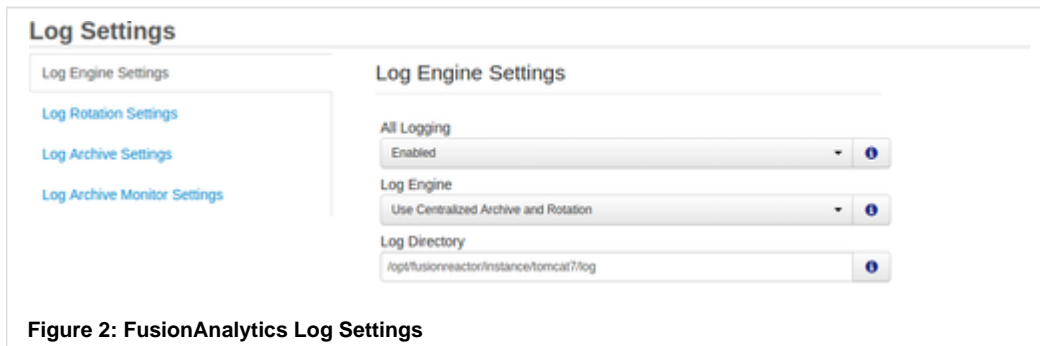


Figure 2: FusionAnalytics Log Settings

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

After the configurations above, you need to set up the **FusionAnalytics Targets**. To do this, please navigate to **Analytics > FusionAnalytics Targets** (Figure 3).

Under the displayed tab, click on the **Add New Target** button. This specifies where the log files from FusionReactor will be sent.

More Configuration Tutorials:

Accessing FusionAnalytics

Setting Up the AIR Client

Configuring FusionReactor for FusionAnalytics

Setting Up Users

Setting Up the Em

FusionAnalytics Targets

Add FusionAnalytics DataCollector Target

Add New Target

Add New Target



URL

Username

Figure 3: FusionAnalytics Targets

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 <i>coldfusion9</i> . |
| Data Collector | 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 <i>faDataCollector-1</i> . |
| 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.

If you are happy with the settings, click the **Save Target Settings** button and after that press the **Test Connection** button in order to check if the connection is okay. If the connection is okay, you should be able to see a **200 OK** message..

If you want to configure the Connect Mode, please navigate to **Analytics > FusionAnalytics Settings**.

ail
Se
rver

Set
tin
g
Up
Ap
plic
ati
on
Filt
ers

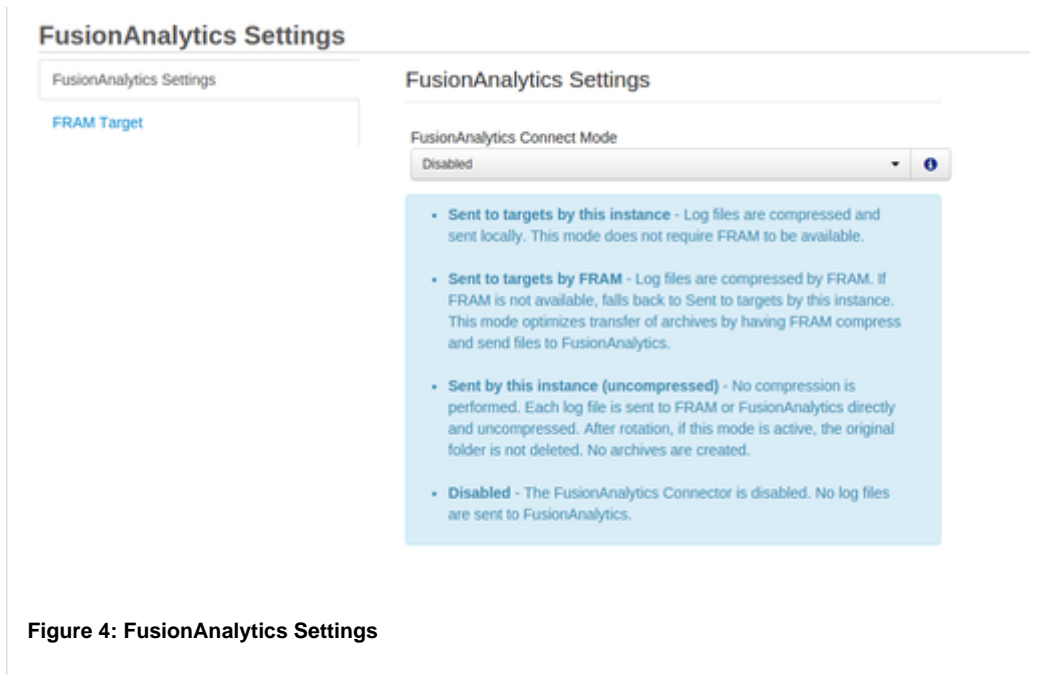
Set
tin
g
Up
TA
P
an
d
Dai
ly
Sta
tus
Re
por
ts

For
mo
re
inf
or
ma
tio
n:

[Fu
sio
nA
nA
lytic
s
Do
cu
me
nta
tion](#)

Usi
ng
Fu
sio
nA
nalytic
s
wit
h
Fu
sio
nR
ea
ctor

[Fu
sio
nR
ea
cto
r 4
Do](#)



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

Next Steps

Setting Up Users

Add more users to and allow them to view specific applications.

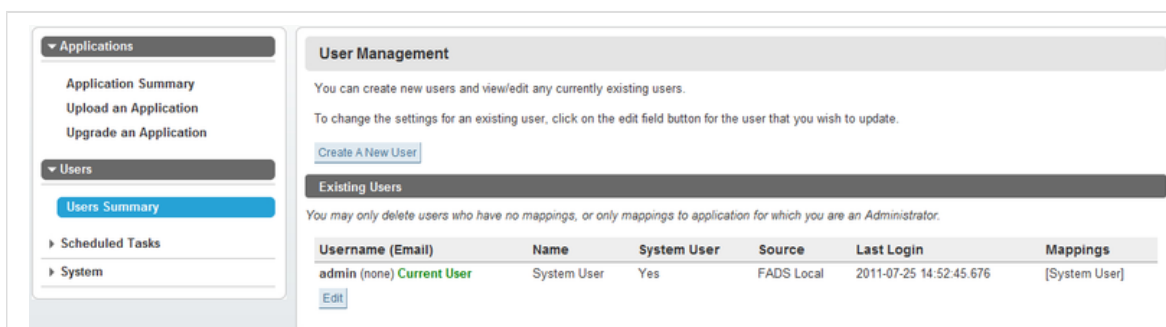
Setting Up Users

Introduction

This section is going to guide you through the process of setting up new users to your FusionAnalytics application.

Setting Up Users

You can configure more user accounts for FusionAnalytics by navigating to **Users > Users Summary**. To add a new user click the **Create** button on the screen (Figure 1).



Please note that depending the license key you obtain, you will be able to configure the appropriate number of users. For example users. More information can also be found in the link, FusionAnalytics Pricing. Additionally, you can also check the FAQ here.

Adding a New User

By the time you press the Create a New User button, you will be required to enter a new **Username** and **Password** (Figure 2). **First Name** optional.

The password strength meter can help you choose a secure password and check that the two fields match. A secure password has to con combined with numbers and special characters.

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

The screenshot shows a web interface for adding a new user. At the top, there's a 'Users' header and a 'Return to User Summary' link. Below that is a dark grey bar with the text 'Add A New User'. The form fields are as follows:

| | |
|--------------------|--|
| Username | <input type="text" value="user1"/> |
| First Name | <input type="text"/> |
| Last Name | <input type="text"/> |
| Password | <input type="password" value="*****"/> Good |
| Password (Confirm) | <input type="password" value="*****"/> Passwords Match |
| Email | <input type="text" value="user_account@company.com"/> |
| System User | <input type="checkbox"/> |

At the bottom of the form is an 'Add New User' button.

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**

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

Current Application Role Mappings

You may only modify roles for applications for which you are an administrator.

| Application | Role | Actions |
|-------------|-----------------|---|
| coldfusion9 | Administrator ▾ | Modify Delete |

Available Application Role Mappings

You may only add roles for applications for which you are an administrator.

| Application | Role | Actions |
|-------------|--------------------|--------------------------|
| cf-instance | None (No Access) ▾ | Add Role |

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

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

[Back to the top](#)

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

Introduction

This section is going to guide you through the process of setting up an Email Server with FusionAnalytics. An email server needs to be configured in case you want to receive **notifications** and **reports** from FusionAnalytics.

Setting Up the Email Server

Please note, the configuration of the email server must be completed for **both** Fusion Analytics DataCollector (FADC) and Fusion/

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

Applications

- Application Summary
- Upload an Application
- Upgrade an Application
- Upload Log Files
- Import History

User

System

- Email Settings**
- Runtime Info
- License Info
- Logout

Email Settings

These configuration settings will be used by default when sending email using FusionAnalytics DataCollector.

Email Server

From Address: fusionanalytics@company.com Notification emails will be sent from this email address.

To Address: user_account@company.com This email address will receive notification emails. Multiple recipients can be entered, separated by commas.

Mail Server: mail.company.com:25 Specify 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: FADC Email Settings

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

Applications

- Application Summary
- Upload an Application
- Upgrade an Application

Users

Scheduled Tasks

System

- Sessions
- Controllers
- Asynchronous Jobs
- Server Settings**
- Logout

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: fusionanalytics@company.com Notification emails will be sent from this email address.

To Address: user_account@company.com This email address will receive notification emails. Multiple recipients can be entered, separated by commas.

Mail Server: mail.company.com:25 Specify 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 2: FADS Email Settings

You should use an appropriate hostname or IP address for the **HTTP Server Host** and **HTTPS Server Host** above as these values are embedded in reporting emails.

- 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 commas.

These addresses do not apply to the reporting feature of FusionAnalytics (which allows you to receive TAP and Daily Status report 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 mail, for example: **company.com:25**.

By default the port number is set to 25, in case you are using another port, this can be configured like the following example: **172.0.0.1:portnumber**. If your SMTP server requires authentication, you will need to put this in front of the address: **username:password@mail.company.com:25**.

Back to the top

What Next?

Setting Up Application Filters

Configure the application filters to view specific types of request.

Setting Up Application Filters

Introduction

This section is going to guide you through the process of setting up application filters. The application filters are going to allow you to filter specific requests within the application based on their URL or user-defined parameters.

Setting Up Application Filters

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

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

In the FusionAnalytics Web Client, please navigate to **Administrator > Manage Applications** (Figure 1). Enter a name for your application **Filter** button.

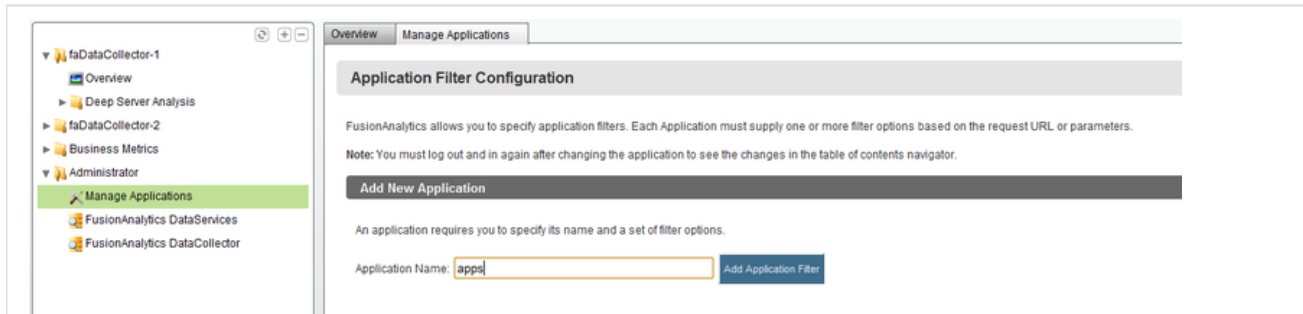




Figure 1: Application filters

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 Params** 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

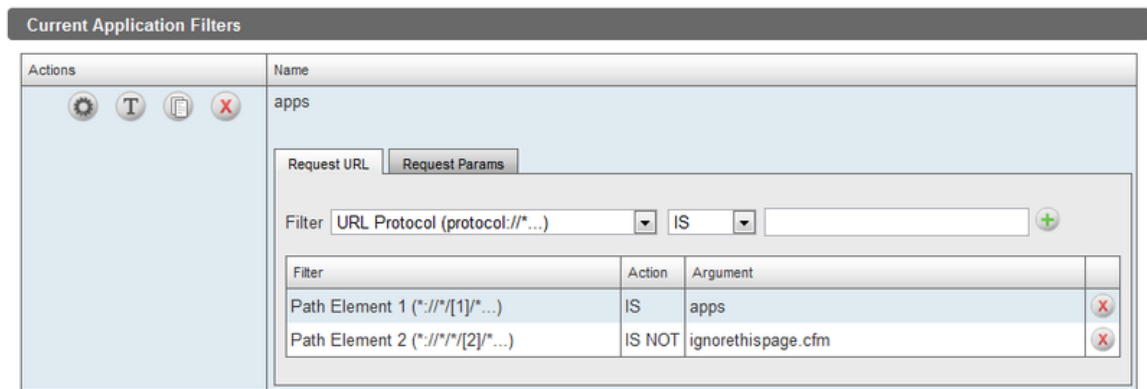


Figure 2: Newly added application in the Current Application Filters

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 **Applications**. You will see the filters that you have defined (Figure 3).

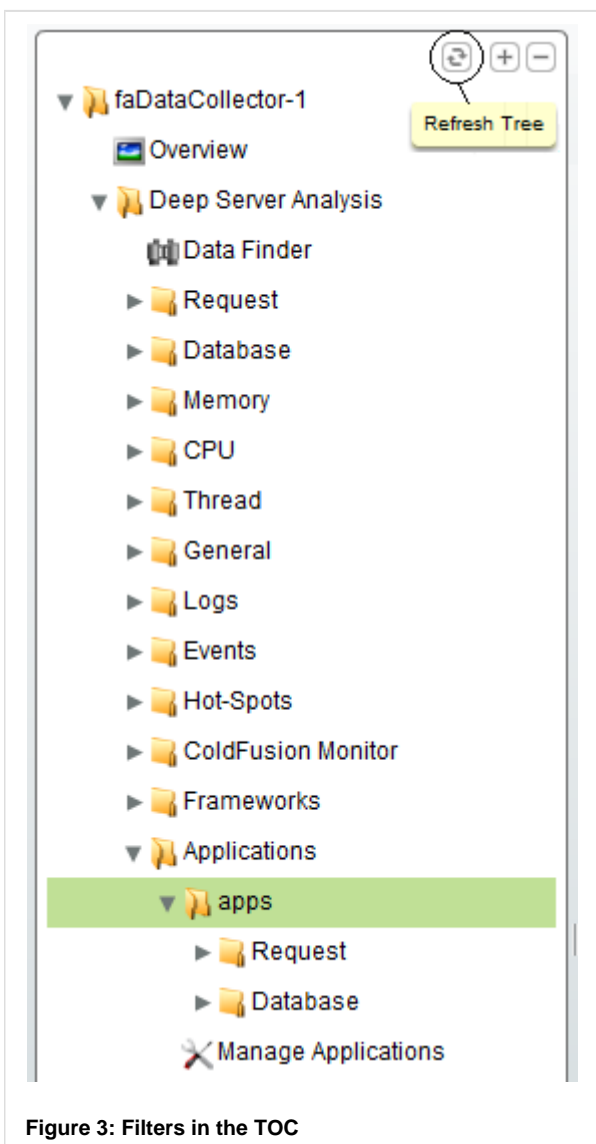


Figure 3: Filters in the TOC

Within each filter you can view a selection of charts and graphs, such as **Slowest Requests**, **Most Popular Requests** and **Slowest DB F** 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 ?Fus

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

Introduction

This section is going to guide you through the process of setting up TAP reports and Daily status reports.

FusionAnalytics supports two types of reports:

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


2. **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.

Setting Up TAP and Daily Status Reports

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. These are already configured for you with each new application.

By default, the TAP report is running every day at 4:15 AM while the Daily Status report is running every day at 4:30 AM.



| Task Name | Provider Name | Duration | Schedule | Last Run | Status |
|--------------|---------------------------|----------------------------|---------------------|-----------------|-----------|
| Daily Report | DailyStatus (coldfusion9) | Jan 1, 2011 - INDEFINITELY | Daily at 4:30:00 AM | 7/25/11 4:30 AM | Scheduled |
| T.A.P Report | TAPReport (coldfusion9) | Jan 1, 2011 - INDEFINITELY | Daily at 4:15:00 AM | 7/25/11 4:15 AM | Scheduled |


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

By the time the reports have been executed, they can be viewed in the FusionAnalytics Web Client, under the **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).



Application Scheduled Tasks

[Application Scheduled Tasks Summary](#)

Editing Application Scheduled Task




| | |
|---------------------------|---|
| Task Name | Daily Report |
| Active | <input checked="" type="checkbox"/> |
| Duration | Start Date <input type="text" value="01/01/2011"/> End Date (optional) <input type="text"/> |
| Frequency | <input type="radio"/> One-Time at <input type="text"/> |
| | <input checked="" type="radio"/> Recurring <input type="text" value="Daily"/> at <input type="text" value="04:30 PM"/> |
| | <input type="radio"/> Daily every Hours <input type="text"/> Minutes <input type="text"/> Seconds <input type="text"/> Start Time <input type="text"/> End Time <input type="text"/> |
| Provider Name | DailyStatus |
| Provider Arguments | Name <input type="text"/> Value <input type="text"/>  |
| |  Add New Argument |
| | <input type="button" value="Update Task"/> <input type="button" value="Cancel"/> |

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).



| | |
|---------------------------|--|
| Provider Name | DailyStatus |
| Provider Arguments | Name <input type="text" value="users"/> Value <input type="text" value="admin,user1"/>  |
| |  Add New Argument |
| | <input type="button" value="Update Task"/> <input type="button" value="Cancel"/> |

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](#)

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.

Setting up the SLA Reports

Introduction

This section provides a quick start guide for configuring and running SLA Reports in FusionAnalytics. For detailed information, please check out this link, full SLA Report documentation.

SLA Report Configuration - Quick Start Guide

STEP 1: ACCESSING THE SLA REPORT CONFIGURATION PAGE

To access the SLA configuration, please launch the FusionAnalytics Web Client (see FusionAnalytics Client) and navigate to **Administrator > SLA Report Configuration**.

STEP 2: ADDING A NEW REPORT

Choose a suitable report name and enter it into the text field. Click the "Add Report" button in order to start configuring your new report.

Report names cannot contain spaces nor the following special characters: ! # \$ % ^ & * () + = ~ ` [] \ ' ; , . { } | " < > ?

STEP 3: SCHEDULE AND EMAIL PREFIXES

Once you have added a new report you will be presented with the configuration screen.

Configure the options to suit your needs. For more detailed information, check out this link, The SLA Report Configuration.

On This Page:

Step 1: Accessing the SLA Report configuration page

Step 2: Adding a new report

Step 3: Schedule and email prefixes

Step 4: Email Recipients

Step 5: Domain Filters

Step 6: Thresholds

More Configuration Tutorials:

Accessing FusionAnalytics

Setting Up the AIR Client

For this tutorial we'll choose to run a report every day, and every Monday for the previous week Mon-Sun. We'll also replace the email subject with "Demo Report" and prefix it with "(S)" if our report is successful and "(F)" if not.

Press "Save Configuration" button in order to save these changes.

STEP 4: EMAIL RECIPIENTS

Next we'll configure who is going to receive our report via email.

Click the **Email Recipients** tab.

Here you have two options:

- Add a recipient from existing FADS users. For this, click the "Add from users" button to see a list of viable users. **Only users mapped to this application will be visible in this list.**
- Add a custom email address and description. Simply enter the email address and description for this address and click the "+" icon to add it as a recipient.

STEP 5: DOMAIN FILTERS

The Domain Filters tab allows you to configure which domains to either include or exclude from your report.

Assume our demo server receives requests from the following 3 domains:

- domain1.intergral.com
- domain2.intergral.com
- domain3.intergral.com

We might want to only include the first two domains on our report, or exclude the first and third domains and display everything else.

For this example, we'll leave this tab blank.

STEP 6: THRESHOLDS

This tab allows you to configure the target thresholds for your report. These values will determine whether your report has passed or failed and will effect the way the email is displayed.

The threshold targets are for a single day. If you run weekly or monthly reports, Outage and Total Request targets will be multiplied by the number of days on your report.

The report is configured with some default values, but we recommend adjusting these to match your server requirements.

For this example we will use the values shown in the image below.

Configuring
FusionReactor for
FusionAnalytics

Setting Up Users

Setting Up the Email
Server

Setting Up Application
Filters

Setting Up TAP and
Daily Status Reports

Setting up the SLA
Reports

STEP 7: SENDING THE REPORT EMAILS

You need to schedule the time at which data for the reports should be generated and the emails should be sent.

From FusionAnalytics Server (FADS) navigate to **Scheduled Tasks > Application Task Summary** to view this report task. If you like, you can click the "Run Now" button in order to generate the emails immediately.

To change the time that the emails are sent, click the **Edit** button.

WARNING

Do not change any value on this screen except the hour on the "Recurring" row. The configurations on this screen are for Daily Status and TAP reports and should not be changed for SLA.

If you wish to disable SLA reports, you should do it from the configuration shown above by choosing not to run daily, weekly or monthly reports. See here.

Next Steps

Read the full documentation for the SLA Reports: SLA Reports.

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
 - JVM Memory Requirements
- 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)
 - Disk Usage Settings (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)
 - APML (FADS)

- APM Volume Switching
- Component
- Data Provider Programs
- Layout
- Configuring Report Images
- FusionAnalytics Licensing

MOVE FUSIONANALYTICS TO A NEW SERVER

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.

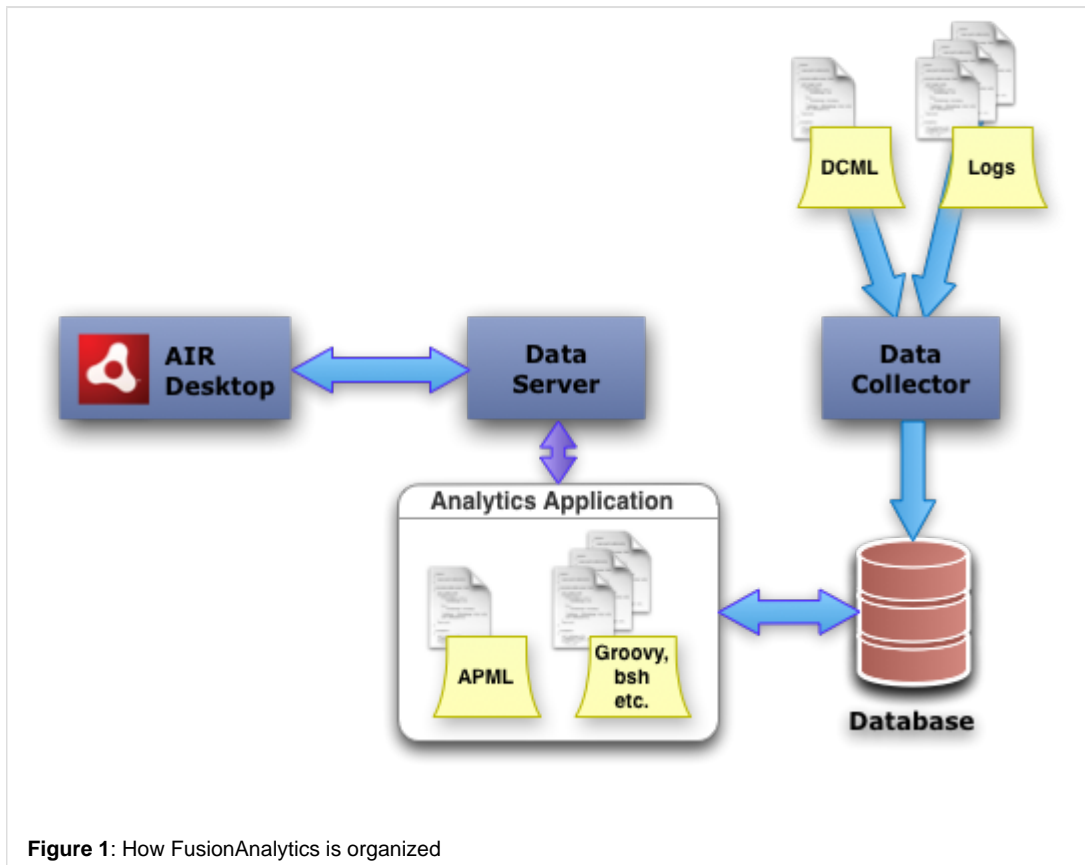


Figure 1: How FusionAnalytics is organized

On This Page:

FusionAnalytics Data Collector

FusionAnalytics Data Services

FusionAnalytics AIR

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...

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...

FusionAnalytics AIR Desktop

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

Learn more...

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...

Expandable and Extensible Platform

EXPANDABLE AND EXTENSIBLE PLATFORM

FusionAnalytics is a platform for visualizing **ANY** data. The **DataCollector** and **DataServices** do not 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.
- 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 the **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 have 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 365 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


Introduction

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 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, check out the link, Filtering Data.

DATE NAVIGATION

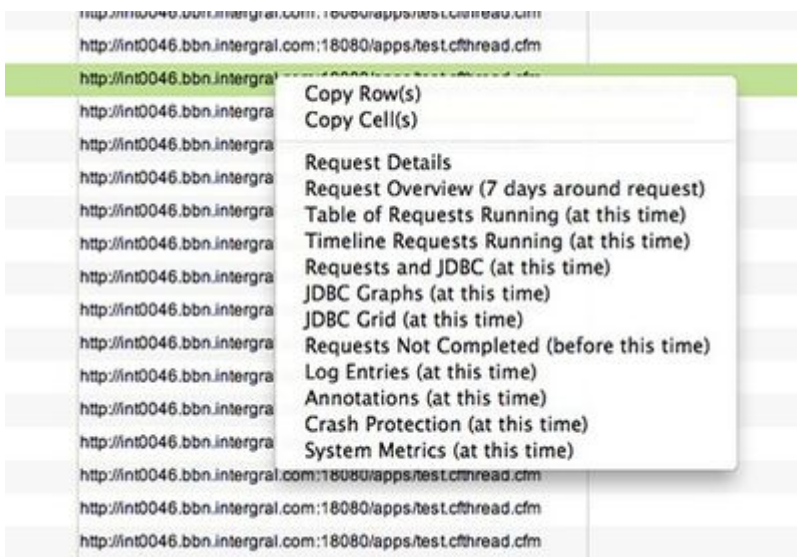


Most of the perspectives in FusionAnalytics are time based so, if you want to explore your data further, you will want to change the span of viewing. At the bottom of time-based perspectives you will see a tab named "TimeLine". By default the TimeLine will be in Basic Mode. For more information, check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

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 (the FusionAnalytics Client is running.) If you want to change your viewing Time Zone then you can select a new Time Zone from the Date Navigation area. For more information, please check out the link, Time Zones.

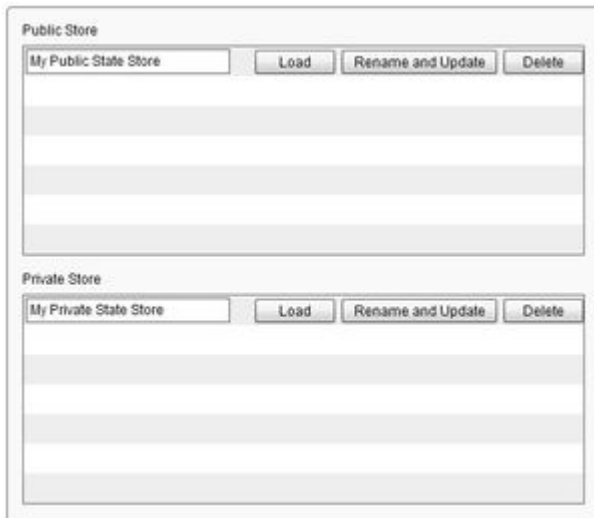
CONTEXT MENUS



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

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 have 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. The content of that tab will then be made visible. For more information, please check out the link, Using Tabs.



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 current date range as well as any currently defined filters.... For more information please see Using State Stores.

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)

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

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 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 "Overview Perspective".

The Overview Perspective is the default Perspective and (like most Perspectives in the application) is date/time-based. The perspective covers 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

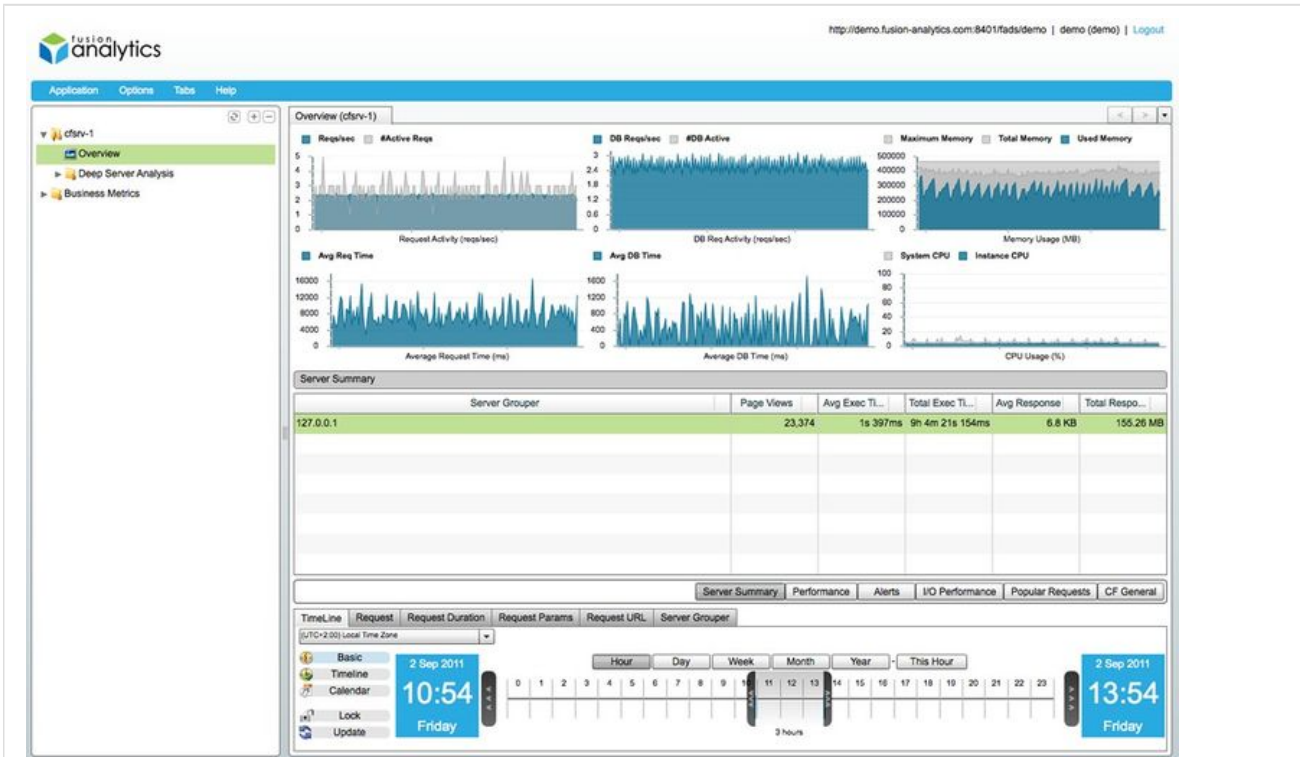


Figure 2: Default Perspective

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
- Perspective Canvas (Right side) - This is the panel where perspectives are displayed. This is where your server information is visible and 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. The 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 rare 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 **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. You can click the "Update" button and your perspective will be refreshed showing the currently selected span of time. Even if you have not clicked 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 button to toggle this feature on and off. Whilst locked, your perspective will be updated whenever you change the currently selected timespan.

Date Navigation^ts_lock.jpg!

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. You will find that the performance is slower than having the Timeline "Unlocked". In general, we recommend that you have the Timeline "Unlocked" 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 shows 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 Navigator.

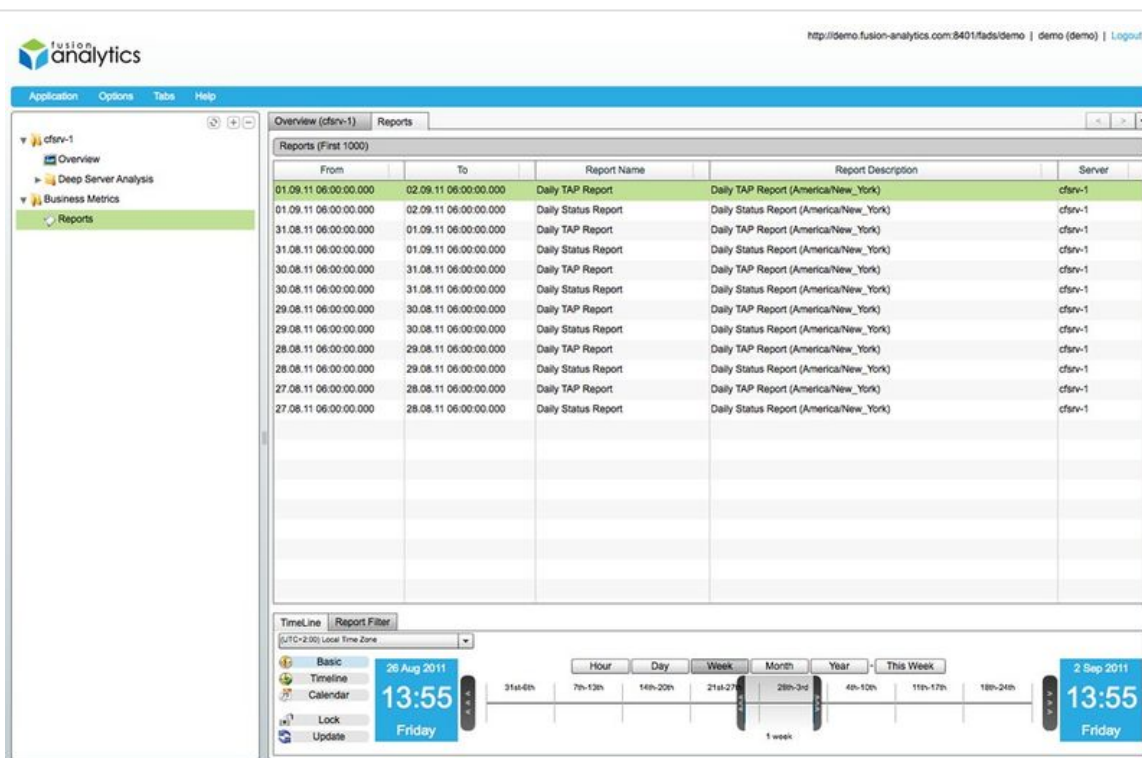


Figure 5: Reports Perspective

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 your credentials and click OK:

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

Tap Report

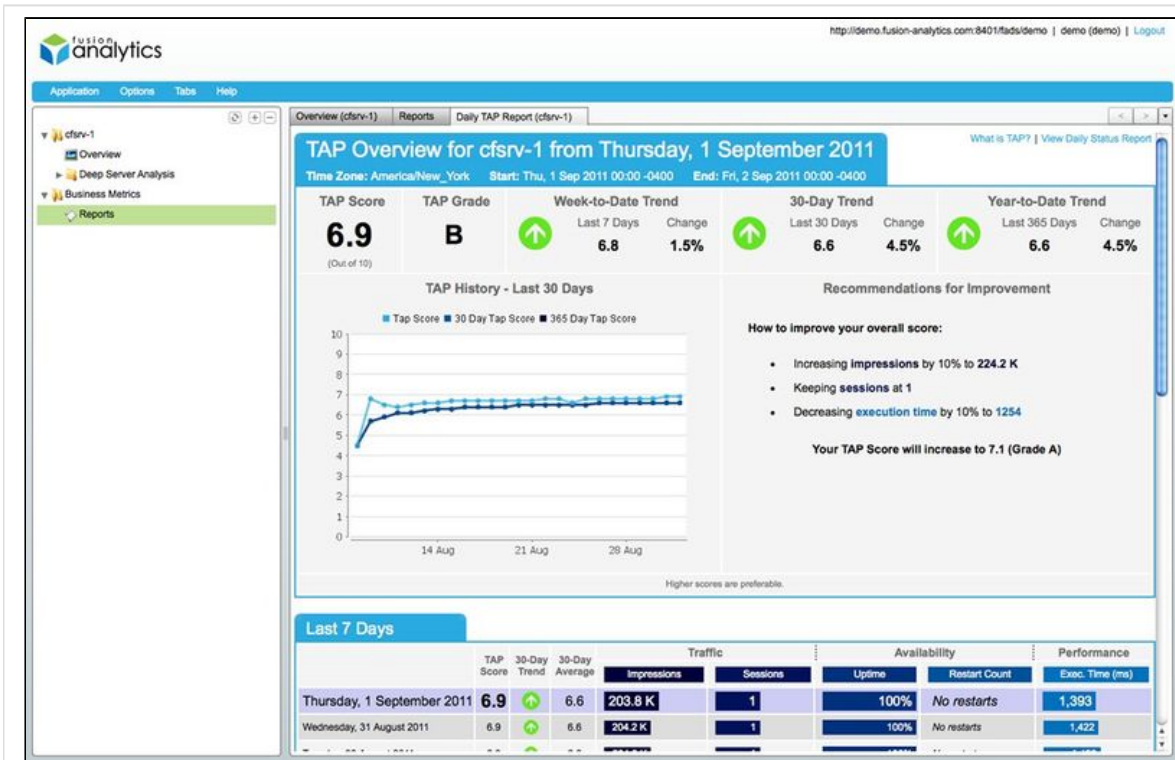


Figure 6: TAP Report

TAP Report: The objective of the Traffic, Availability and Performance (TAP) index is to give application & program managers, bu holders, a less-technical orientated metric on their servers. FusionAnalytics makes suggestions on how to improve your TAP score 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 selectin

Daily Status Report

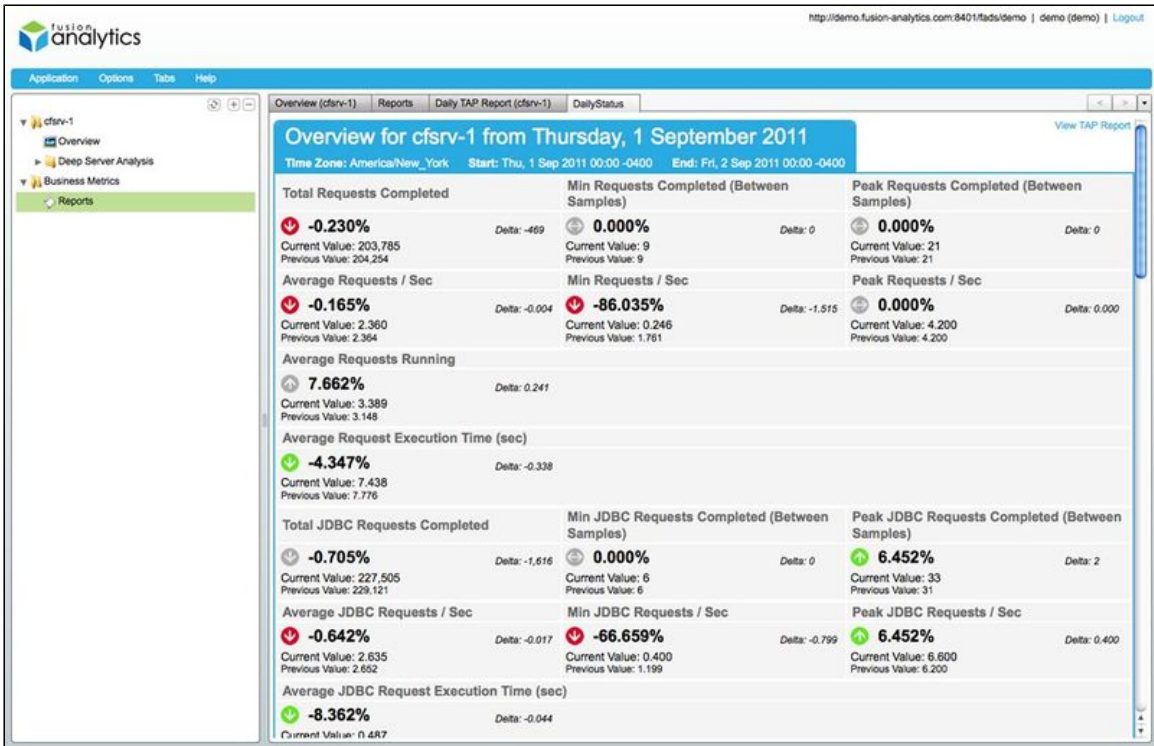


Figure 7: Daily Status Report

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

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 request times, CPU and memory information, as well as URL parameters, heap memory usage and JMX statistics. FusionAnalytics takes all of this combined visual perspectives, allowing you to visualize all transactions, SQL queries, data transfer, server events etc. taking place over time. This historical view of what happened on your server. This data is essential to support investigative analysis, continuous server analysis and at 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

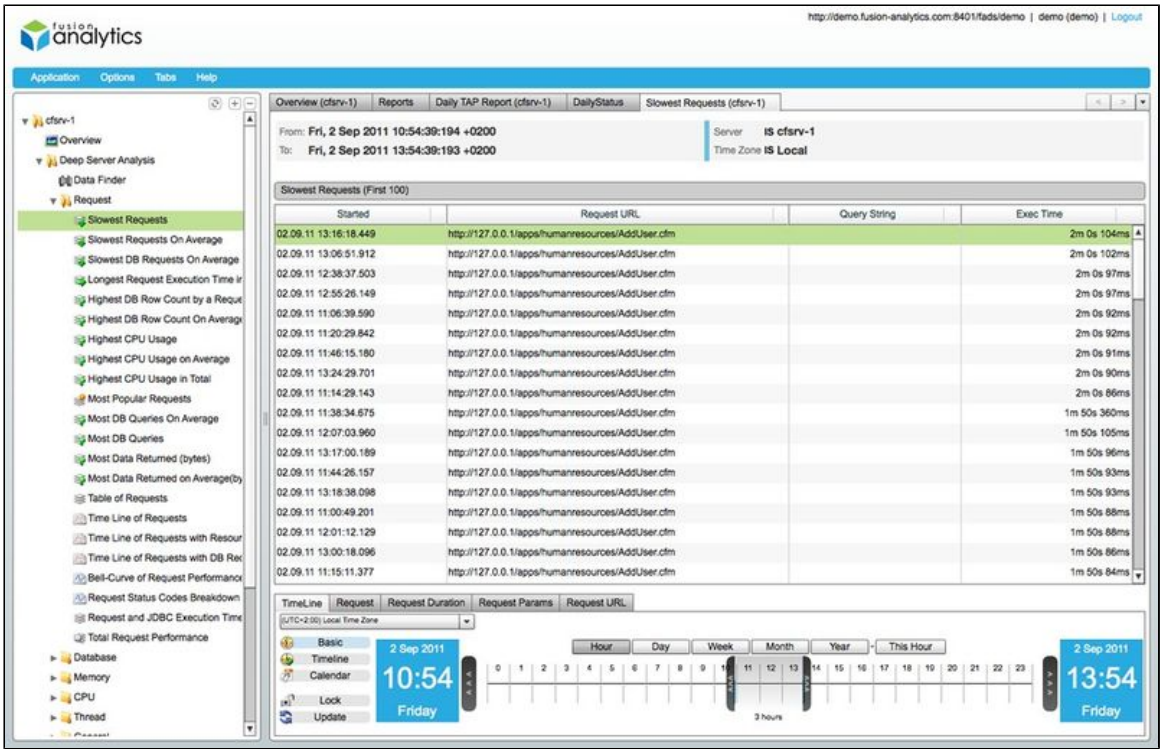


Figure 8: Slowest Requests Perspective

Within the Slowest Requests Perspective, the top 100 slowest requests (within the specified time frame) are listed allowing detailed 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 menu item.

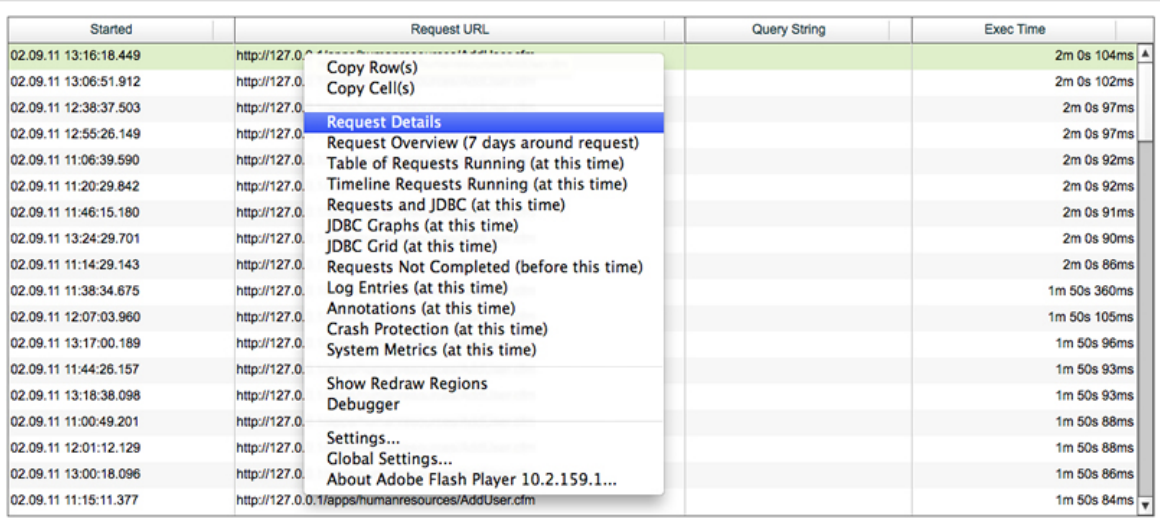


Figure 9: Slowest Requests Context Menu

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

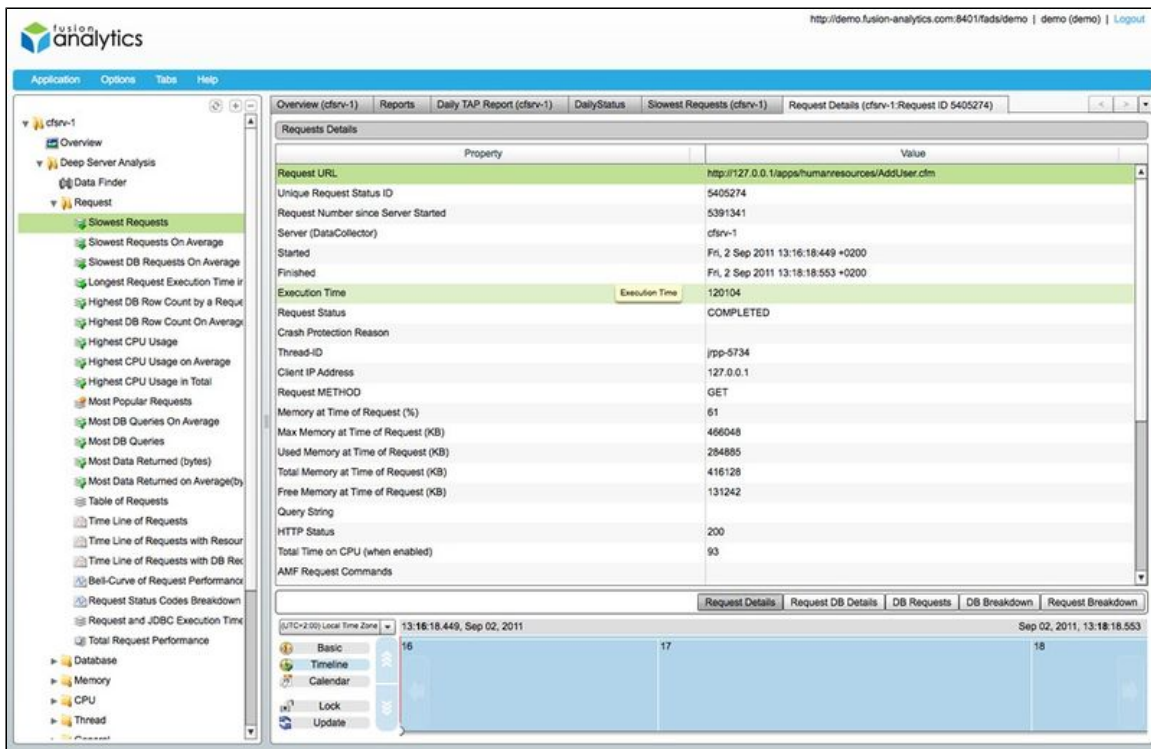


Figure 10: Request Details Perspective

The Request Details perspective shows all the details of the actual (selected) request, including URL, Start/Finish time, Status code, exec 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 accessed using the toggle buttons that are rendered below the Request Details DataGrid. For example, to see where the database time is basis relevant to the the selected request click the "DB Requests" toggle button just above the timeline.

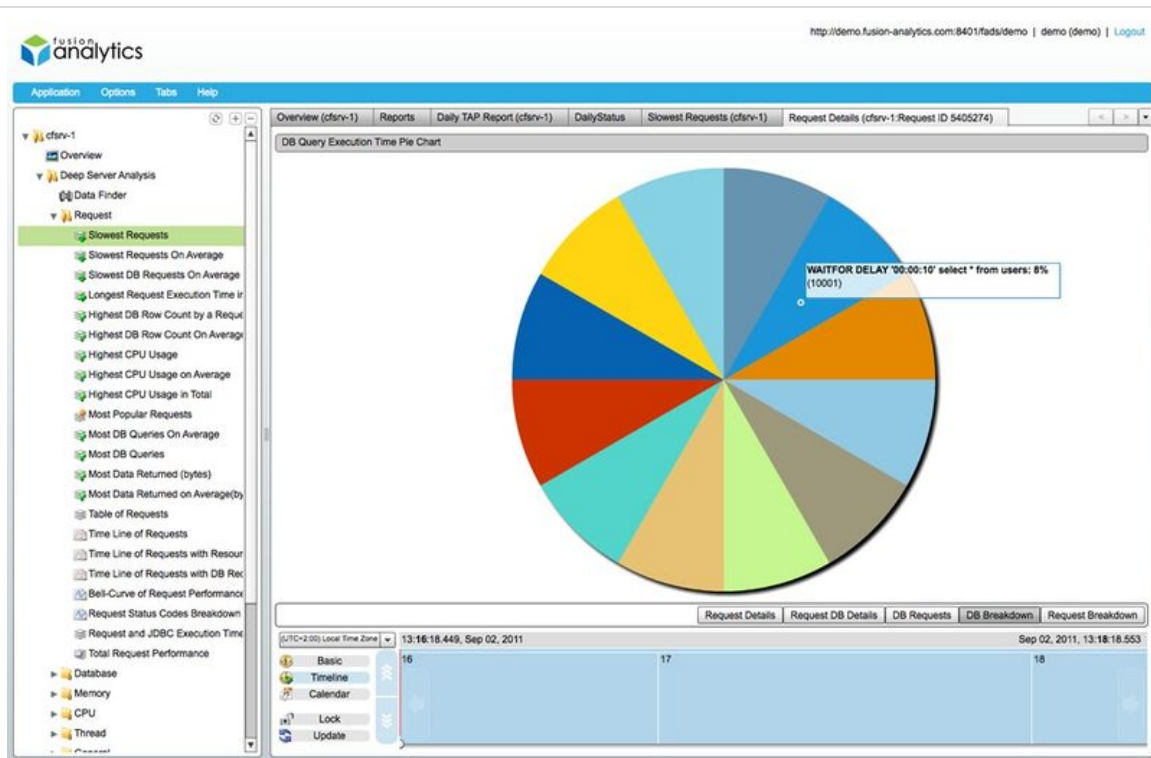


Figure 11: DB Query Execution Time

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 > Cf the CPU and Memory usage Perspective:

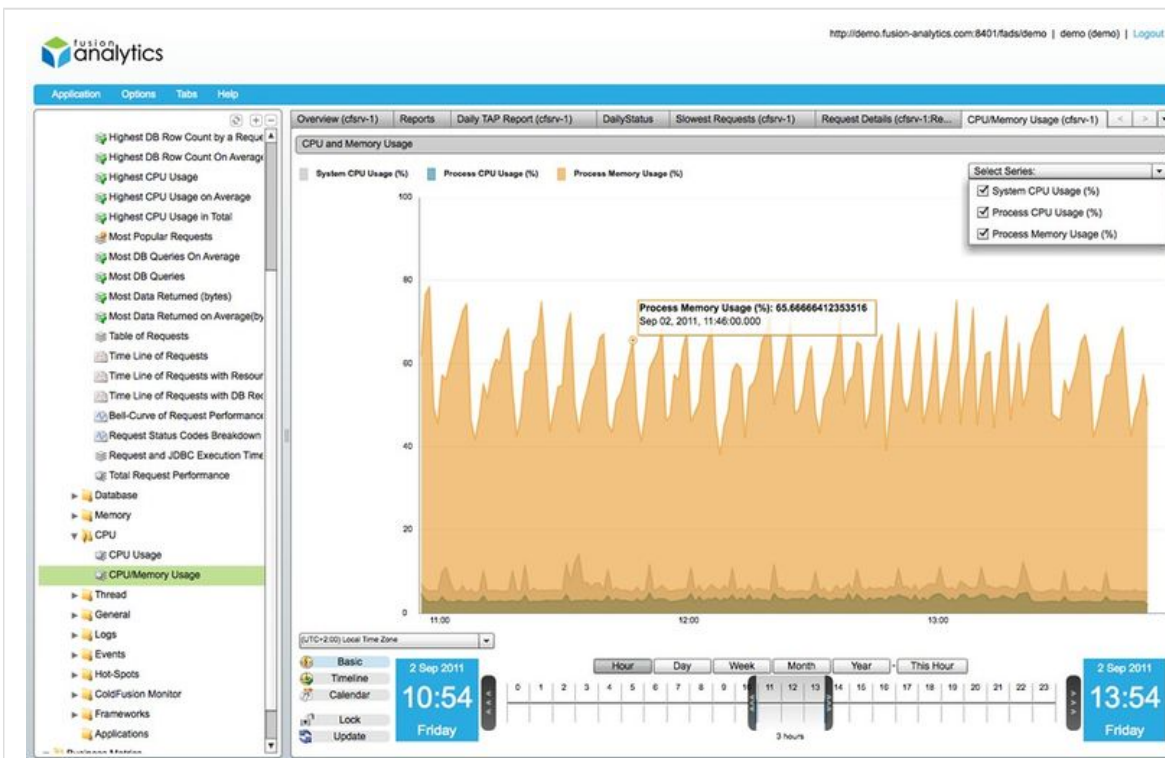


Figure 12: CPU/Memory Usage Perspective

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

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, by default you will be taken to the application overview page. For each application defined within FusionAnalytics (listed in the **Summary**) you will see a "Launch" button.

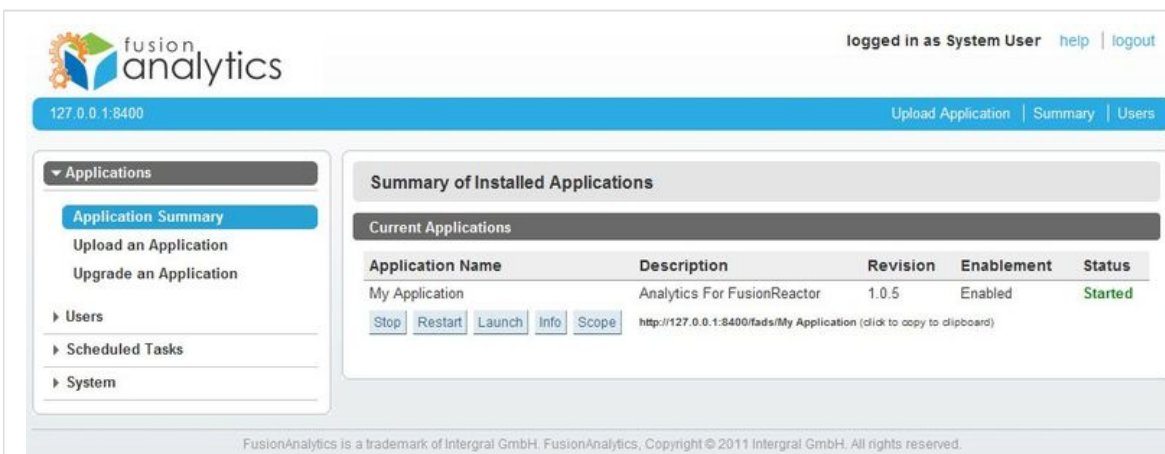


Figure 1: FusionAnalytics DataServices (FADS)

If you click on this "Launch" button, the **Web Client** will appear 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.

An Example URL of the **Web Client** is the following:

- `http://serverName:8400/fads/applicationName`

Please make sure that you change the "applicationName" from the URL above to match the name of your application.

After launching an application the first screen you will be taken to will be the Login screen. Enter your login credentials (Username and Password) and click the client:

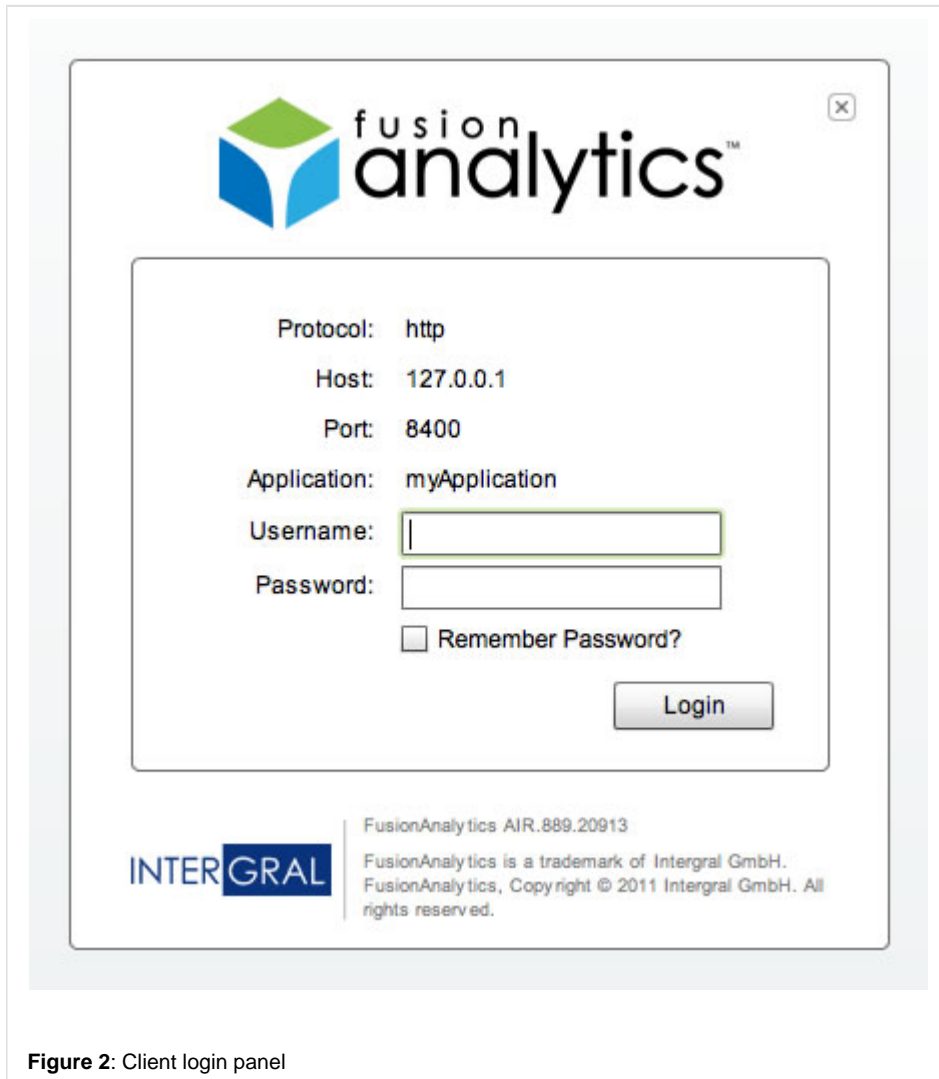


Figure 2: Client login panel

[Back to the top](#)

- 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 : ytics 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 A FusionAnalytics Application.** To get this URL you need to go to the FusionAnalytics DataServices (FADS) screen. By default you will se

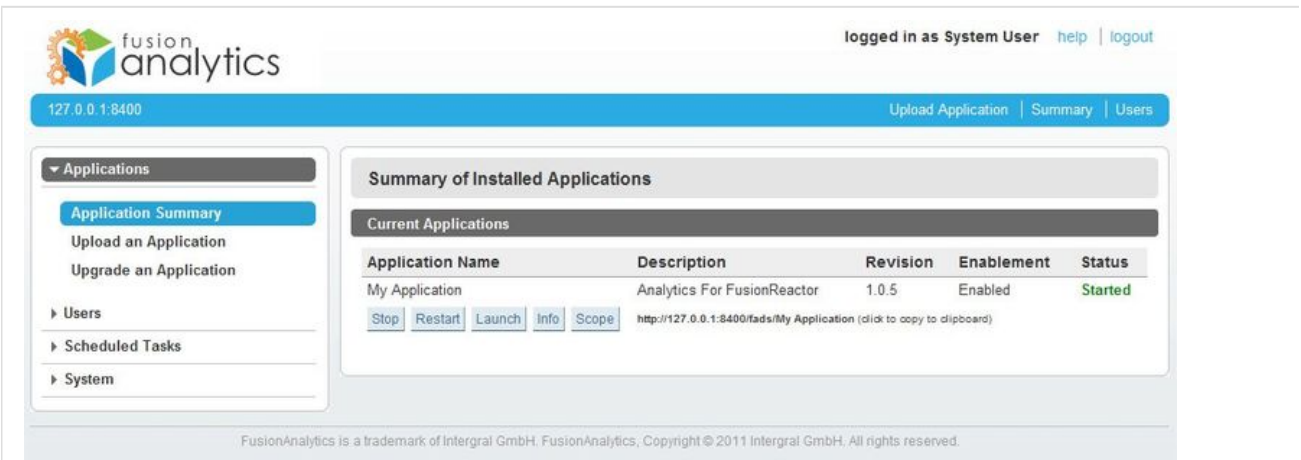



Figure 1: FusionAnalytics DataServices (FADS)

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



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 'Application'. Paste the Application URL into the text box and click the button.



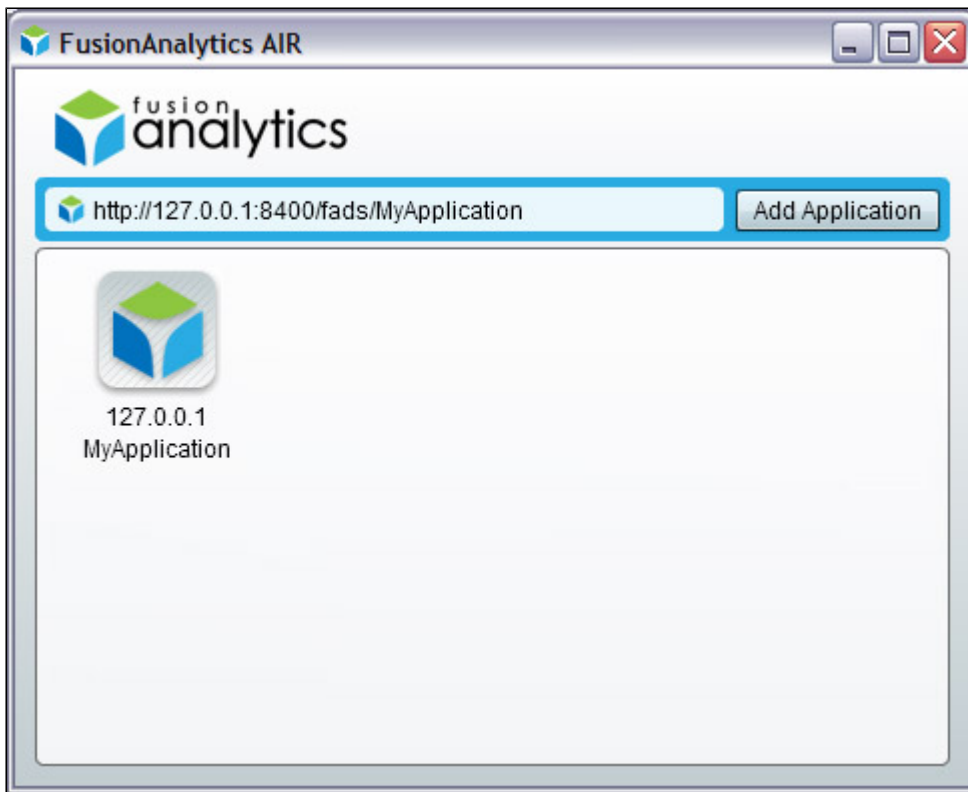


Figure 3: Paste application URL into the FusionAnalytics AIR client

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** an application link, you can right-click on the icon and the two options will appear.





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 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 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 you can change the details of this application link.



Edit Server Details [X]

Protocol:

Secure Connection:

Host:

Port:


Application:

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 Enter your Username / Password to access the rest of the client.



Protocol: http

Host: 127.0.0.1

Port: 8400

Application: My Application

Username:

Password:


 FusionAnalytics Web.794.21262
FusionAnalytics is a trademark of Intergral GmbH.
FusionAnalytics, Copyright © 2011 Intergral GmbH. All rights reserved.

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 c then the next time you log into this sever the password filled will be pre-filled so you can simply (and quickly) click the "Login" butt

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

Back to the top

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:

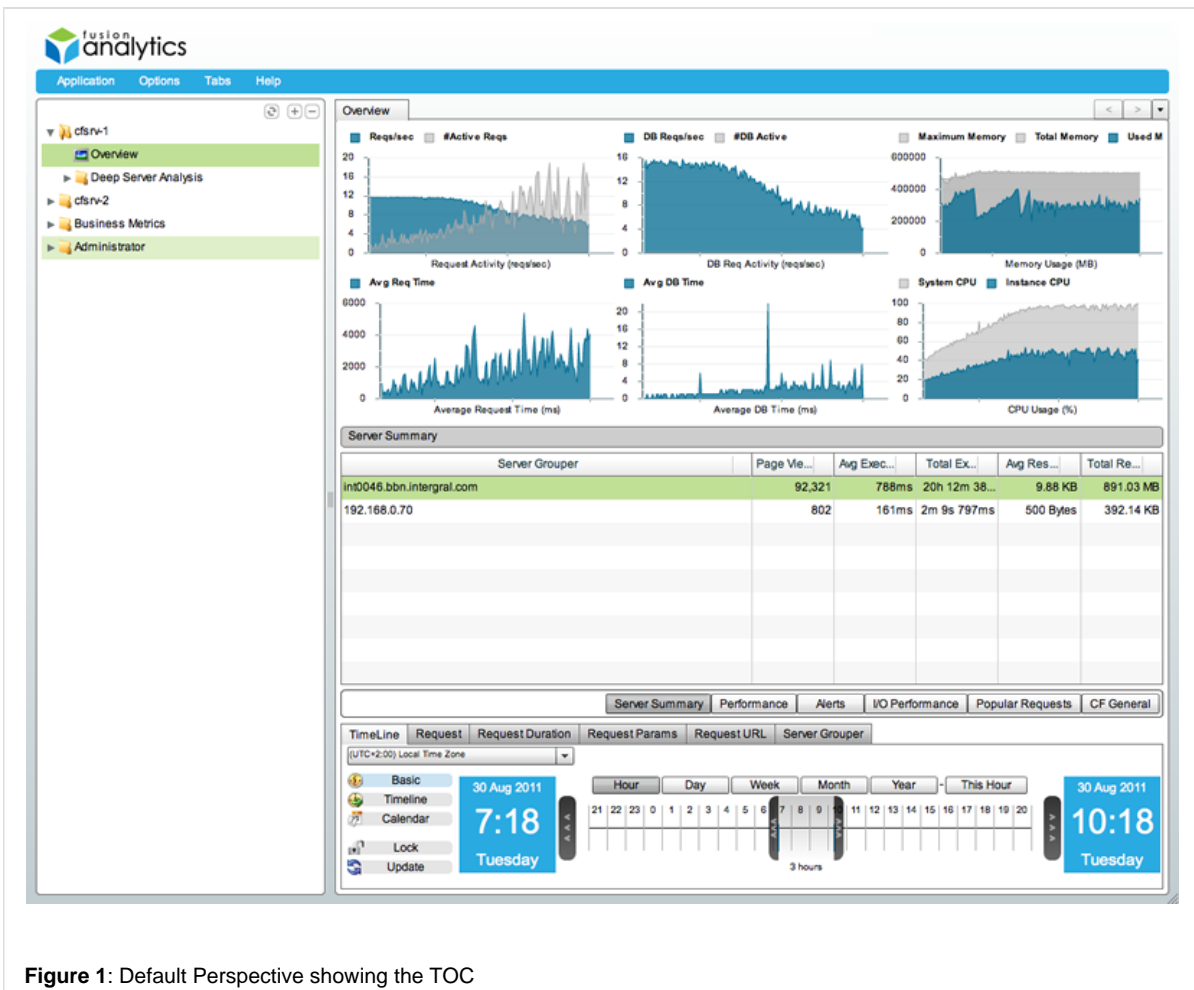



Figure 1: Default Perspective showing the TOC

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).



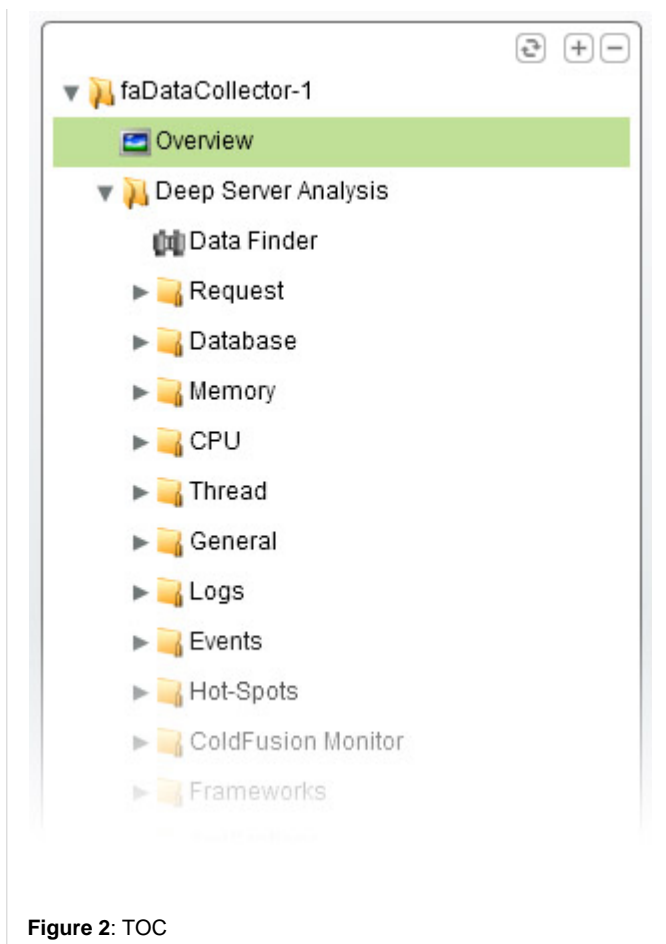





Figure 2: TOC

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 **refresh** button:

| | | |
|---|---------------------|--------------------------------------|
|  | TOC Refresh | <i>Refresh the TOC data provider</i> |
|  | Expand TOC | <i>Expand all TOC menu items</i> |
|  | Collapse TOC | <i>Collapse all TOC menu items</i> |

[Back to the top](#)

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.

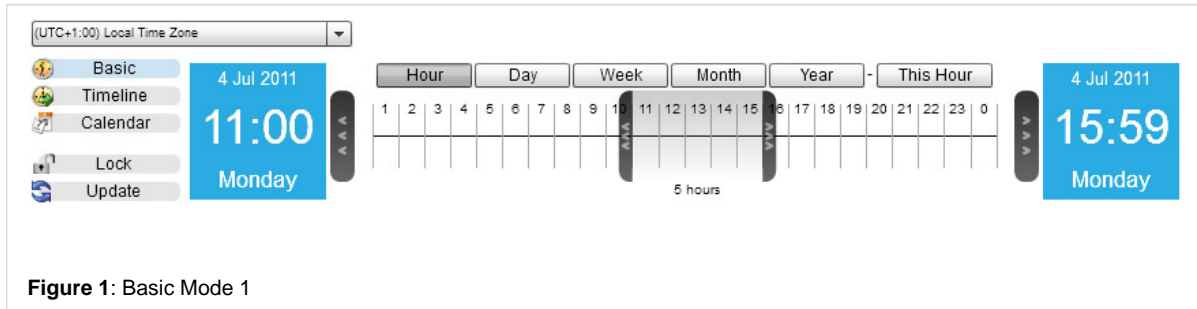


Figure 1: Basic Mode 1

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.

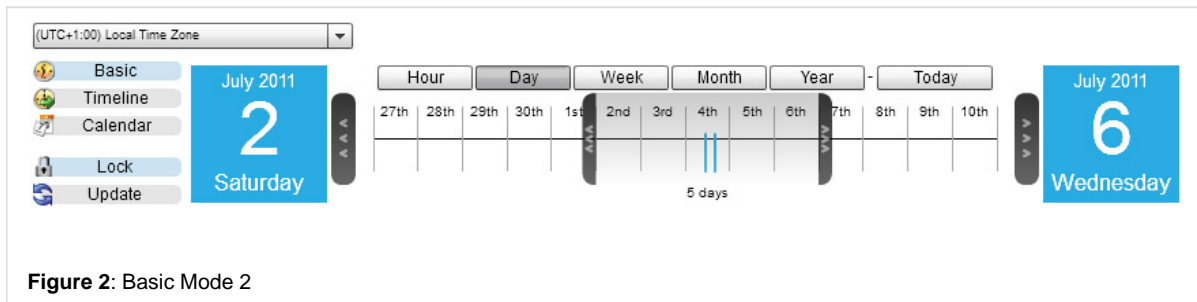


Figure 2: Basic Mode 2

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 or 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 of 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 across 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 numeric 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

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:

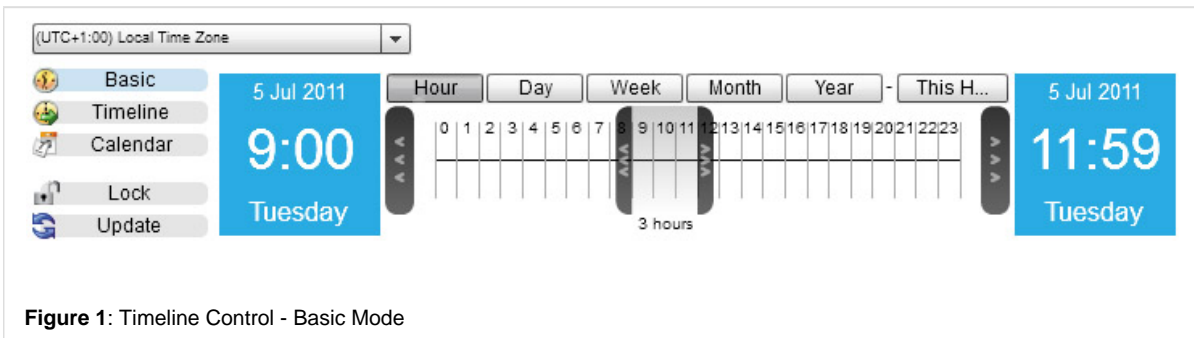


Figure 1: Timeline Control - Basic Mode

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.

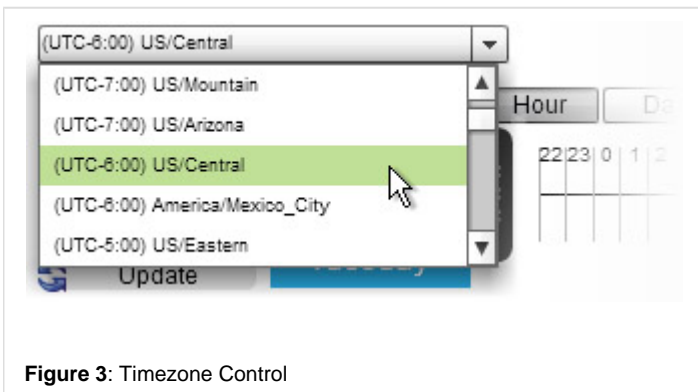


Figure 3: Timezone Control

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](#)

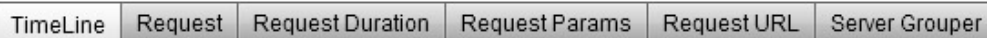
NEXT STEPS

- [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).



| | | | | | |
|----------|---------|------------------|----------------|-------------|----------------|
| TimeLine | Request | Request Duration | Request Params | Request URL | Server Grouper |
|----------|---------|------------------|----------------|-------------|----------------|

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.

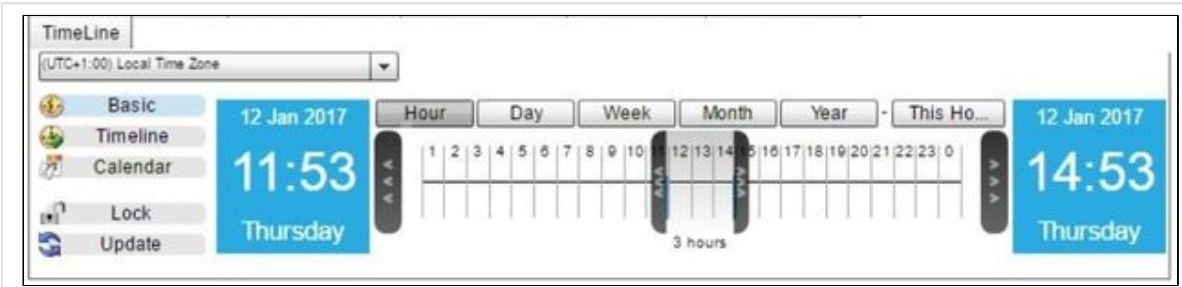


Figure 2: Timeline Tab

[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.

??

Request

Filter Memory Percent > [] + Apply Filter

| Filter | Action | Argument | |
|----------------|--------|----------|---|
| Memory Percent | > | 90 | - |
| | | | |
| | | | |

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.

Request Duration

0 60000

10078 60000

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.

Request Params

Filter fuseaction IS [] + Apply Filter

| Filter | Action | Argument | |
|------------|--------|------------|---|
| fuseaction | IS | helloworld | - |
| | | | |
| | | | |

?

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.

The screenshot shows the 'Request URL' filter configuration interface. At the top, there is a tab labeled 'Request URL'. Below the tab, there is a configuration bar with the following elements: a 'Filter' dropdown menu set to 'URL Protocol (protocol://*...)', an 'Action' dropdown menu set to 'IS', an empty 'Argument' input field, a '+' button, and an 'Apply Filter' button. Below the configuration bar is a table with the following structure:

| Filter | Action | Argument | |
|--------------------------------|--------|----------|---|
| URL Protocol (protocol://*...) | IS | https | - |
| | | | |
| | | | |

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 there 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:?

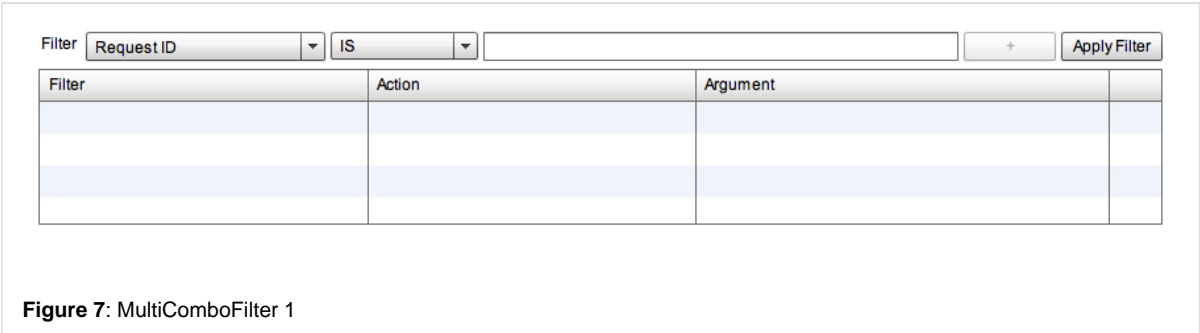


Figure 7: MultiComboFilter 1

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" :

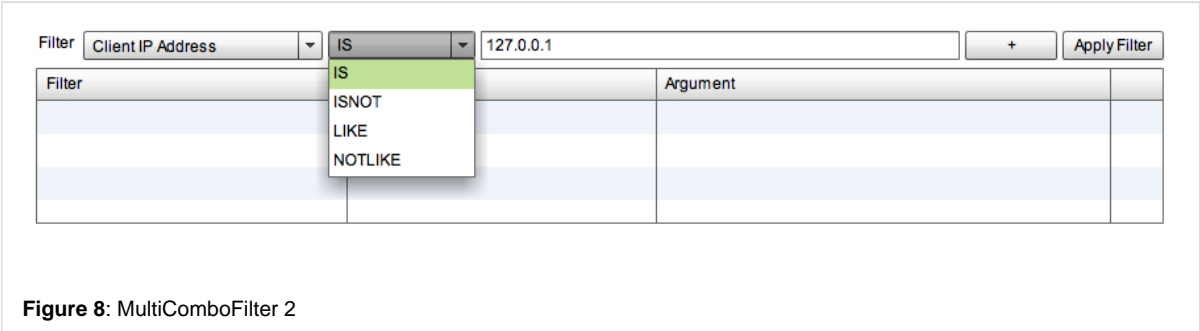


Figure 8: MultiComboFilter 2

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:

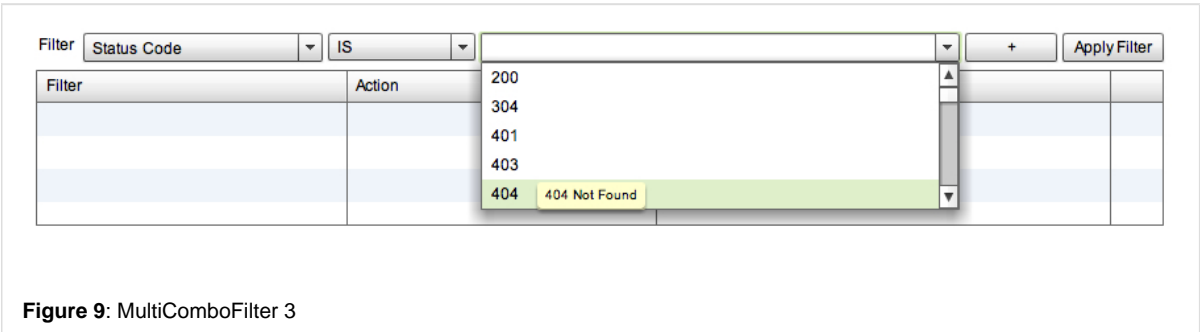


Figure 9: MultiComboFilter 3

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...) :?

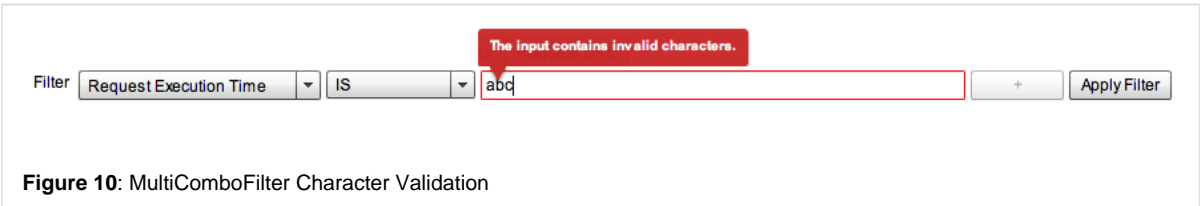



Figure 10: MultiComboFilter Character Validation



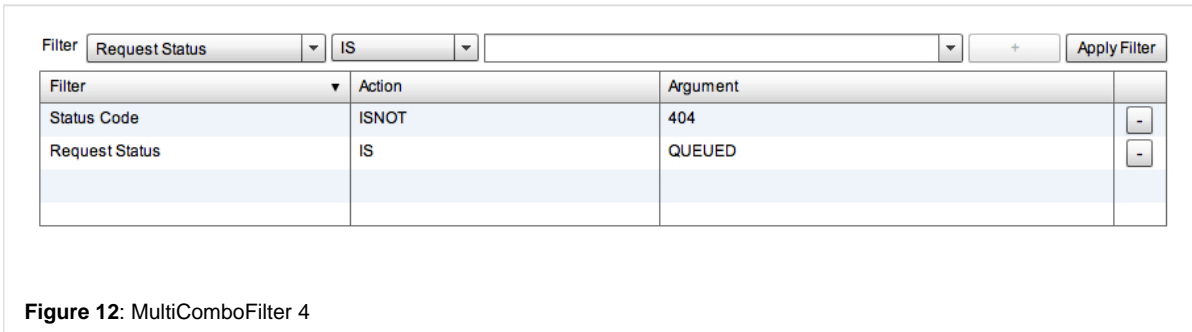
Note: The add "+" button will be disabled while the content within the Argument input is either empty or invalid (in relation to the selected Filter) to prevent any invalid filter queries against the DataBase :

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

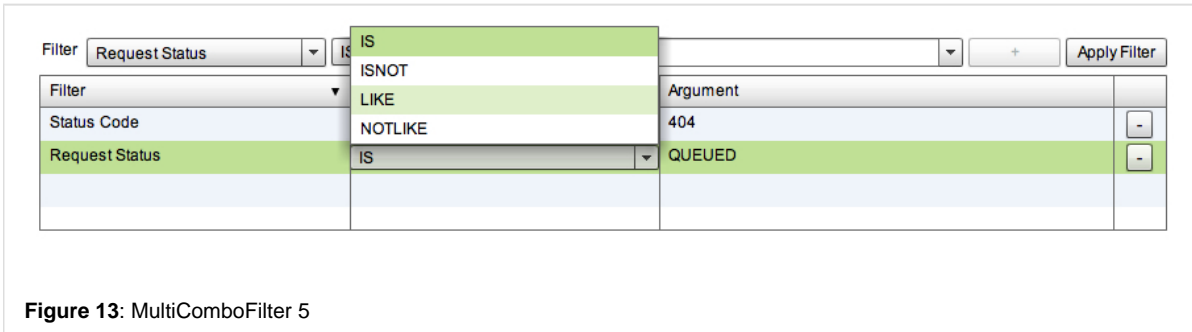
When the "Apply Filter" button is clicked only filters that have been added to the Filters Grid are applied. If the Add button is enabled (if you are adding a new filter) and the "Apply Filter" button is clicked, the new filter is both added to the Filters Grid and applied with all other filters in the 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 than 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 be further edited or deleted:



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 that 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 click the delete "-" button :?



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






Figure 14: MultiComboFilter 6

?

[Back to the top](#)

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

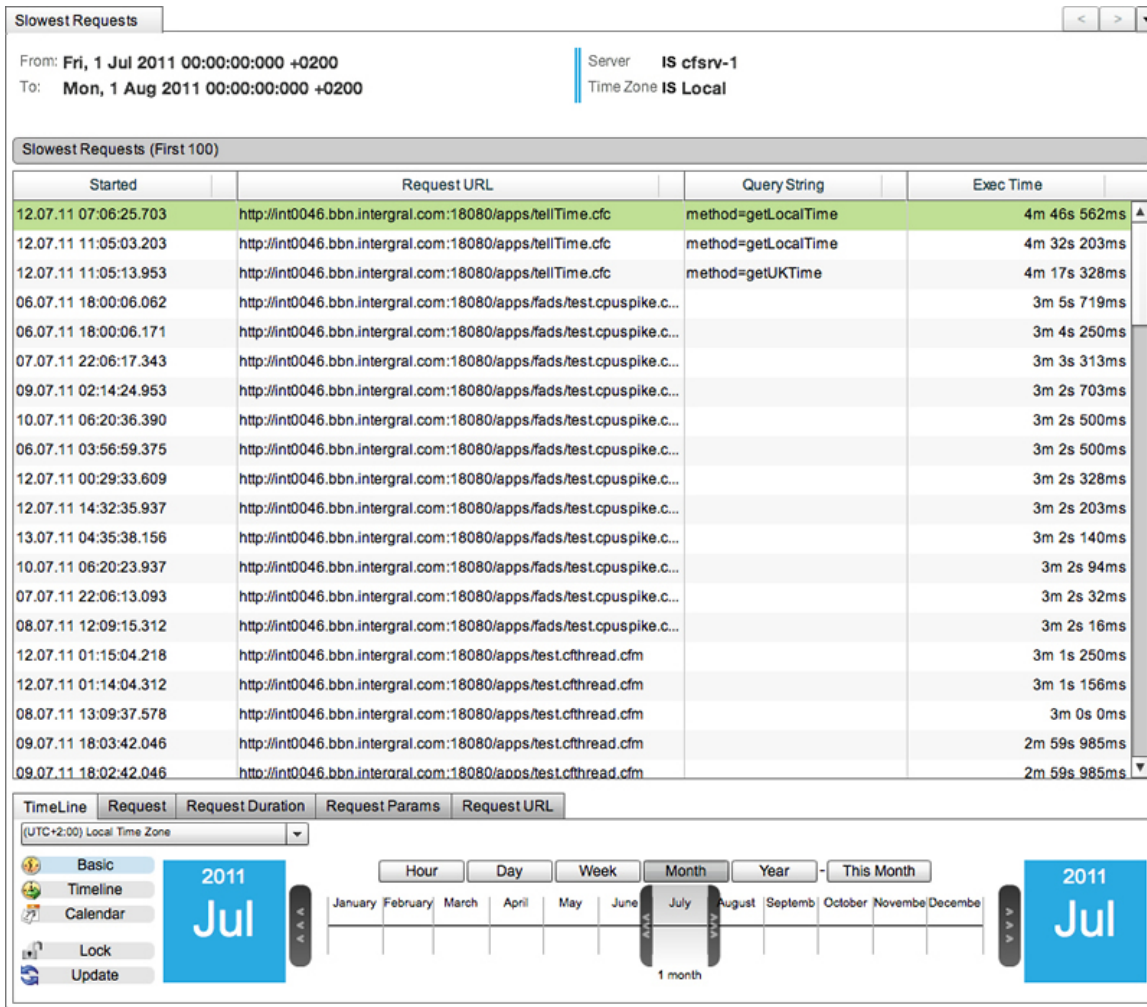
DataGrid Visualization

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 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):



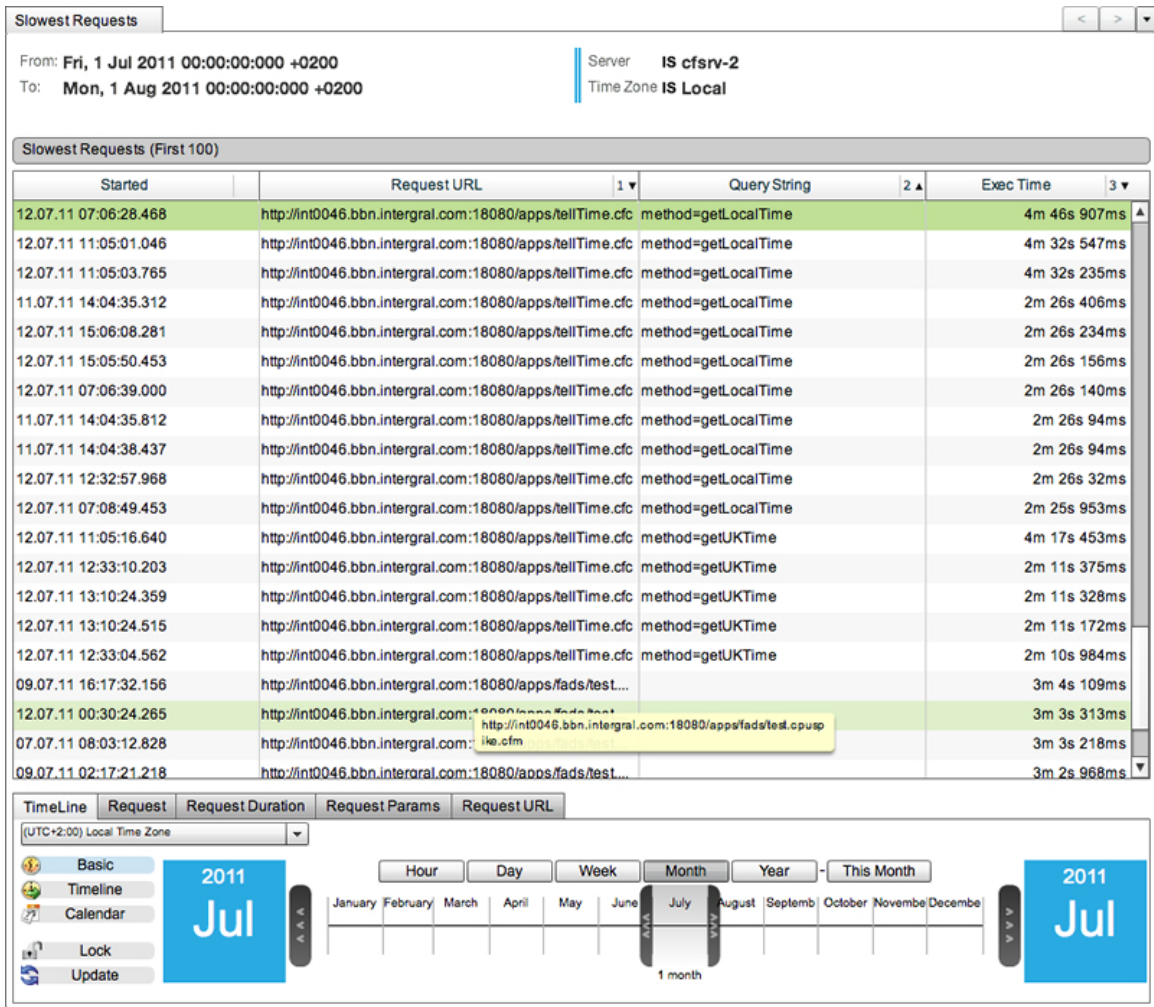


Figure 2: DataGrid Visualization - Data Sorting

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 trailing "..." 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.

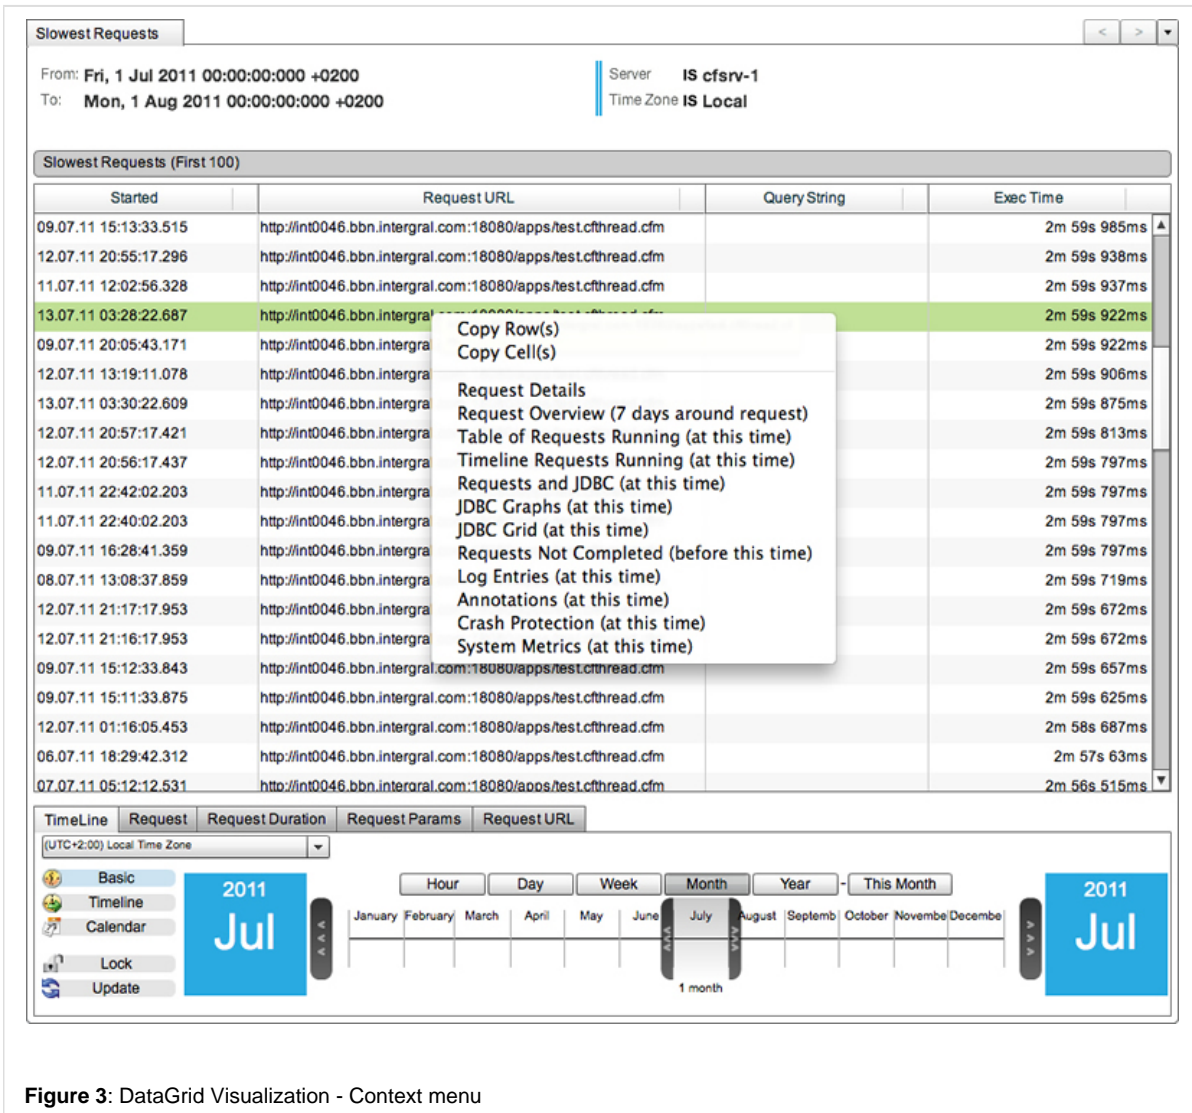


Figure 3: DataGrid Visualization - Context menu

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.:



| Started | Request URL | Query String | Exec Time |
|-----------------------|---|---------------------|--------------|
| 12.07.11 15:06:08.281 | http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc | method=getLocalTime | 2m 26s 234ms |
| 12.07.11 15:05:50.453 | http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc | method=getLocalTime | 2m 26s 156ms |
| 12.07.11 07:06:39.000 | http://int0046.bbn.intergral.com:18080/apps/tellTime.cfc | method=getLocalTime | 2m 26s 140ms |
| 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.cfm | | 3m 4s 109ms |
| 12.07.11 00:30:24.265 | http://int0046.bbn.intergral.com:18080/apps/fads/test.cfm | | 3m 3s 313ms |
| 07.07.11 08:03:10.796 | http://int0046.bbn.intergral.com:18080/apps/fads/test.cfm | | 3m 3s 218ms |
| 09.07.11 02:17:21.218 | http://int0046.bbn.intergral.com:18080/apps/fads/test.cfm | | 3m 2s 968ms |
| 12.07.11 14:33:27.687 | http://int0046.bbn.intergral.com:18080/apps/fads/test.cfm | | 3m 2s 719ms |
| 09.07.11 16:17:21.234 | http://int0046.bbn.intergral.com:18080/apps/fads/test.cfm | | 3m 2s 703ms |
| 13.07.11 04:36:30.484 | http://int0046.bbn.intergral.com:18080/apps/fads/test.cfm | | 3m 2s 375ms |
| 07.07.11 08:03:10.796 | http://int0046.bbn.intergral.com:18080/apps/fads/test.cfm | | 3m 2s 297ms |

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.cfm
cfm      3m 4s 109ms
12.07.11 00:30:24.265      http://int0046.bbn.intergral.com:
18080/apps/fads/test.cfm

```

```
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.:

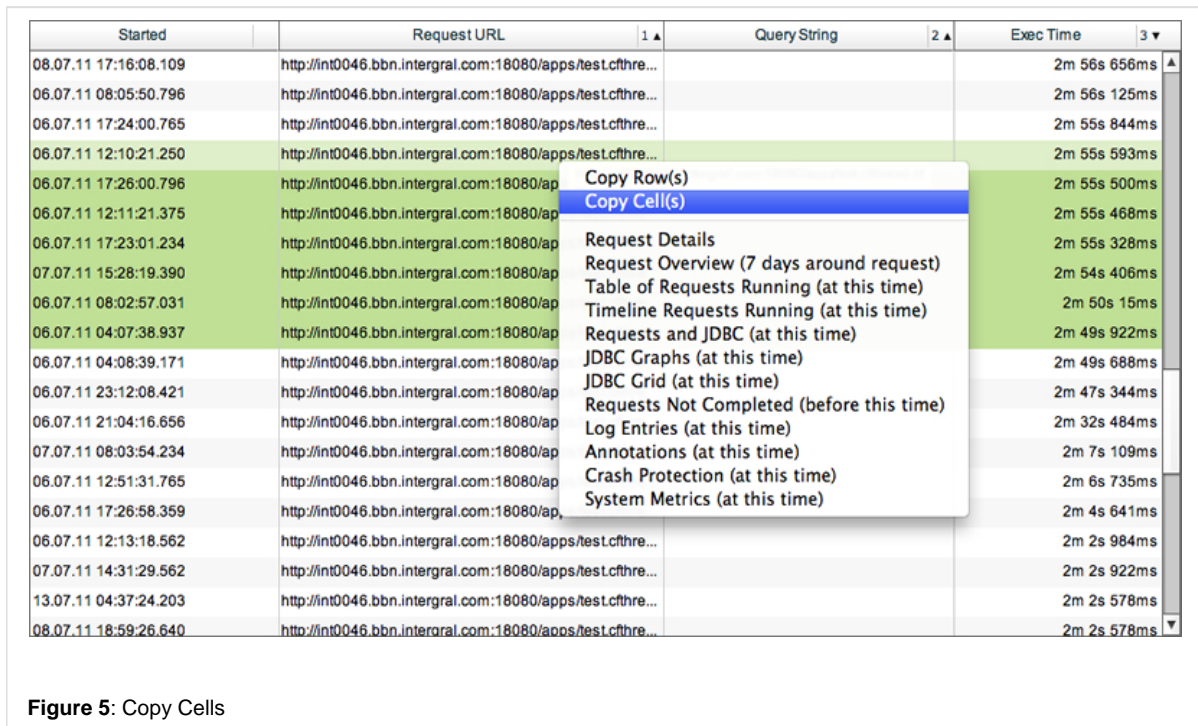


Figure 5: Copy Cells

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](#)

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:

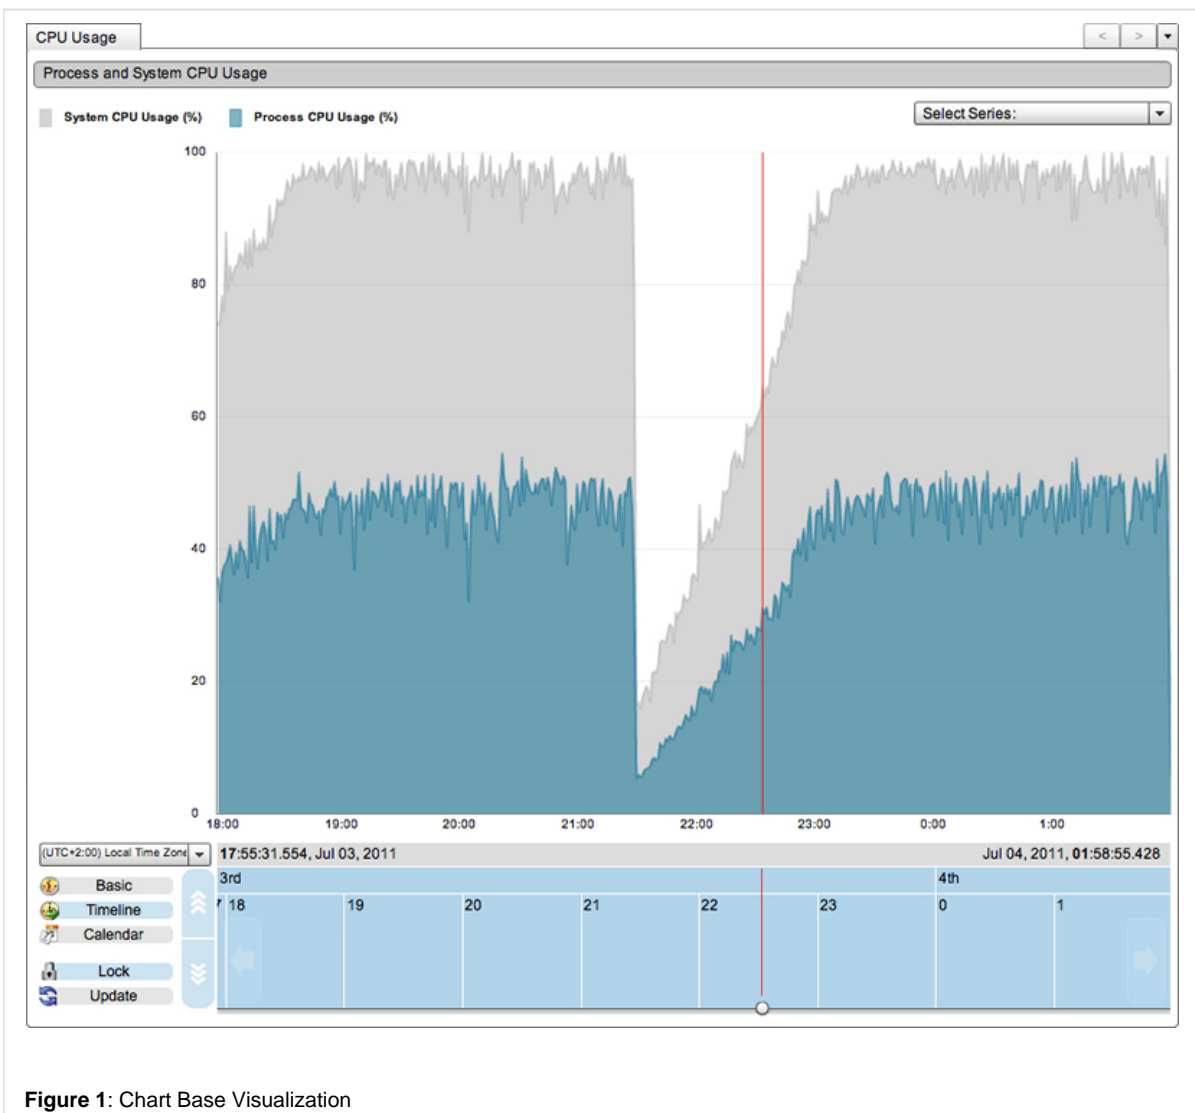


Figure 1: Chart Base Visualization

[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:

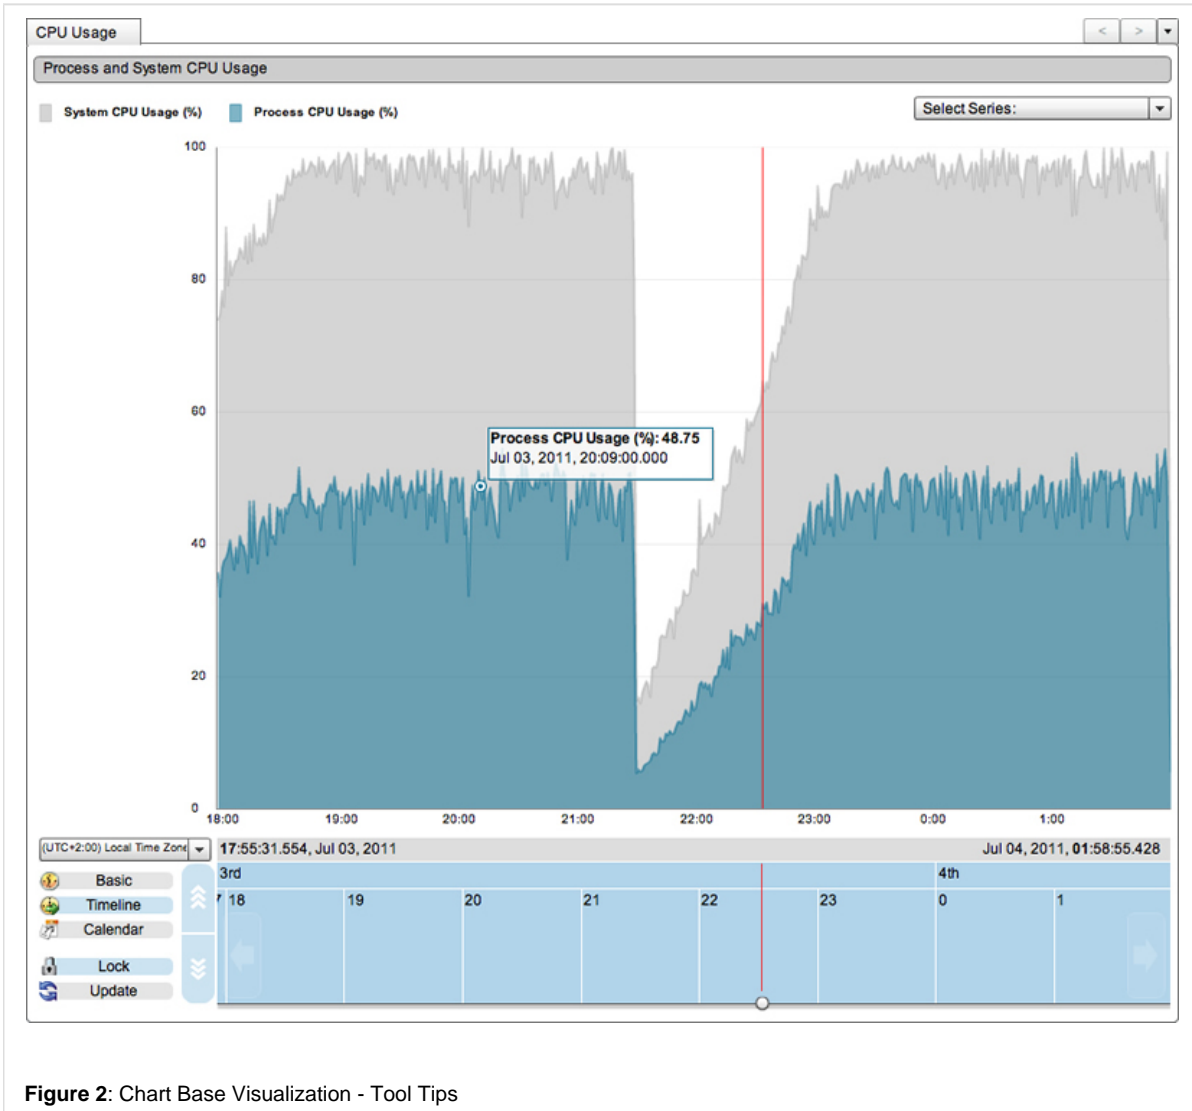



Figure 2: Chart Base Visualization - Tool Tips

[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:



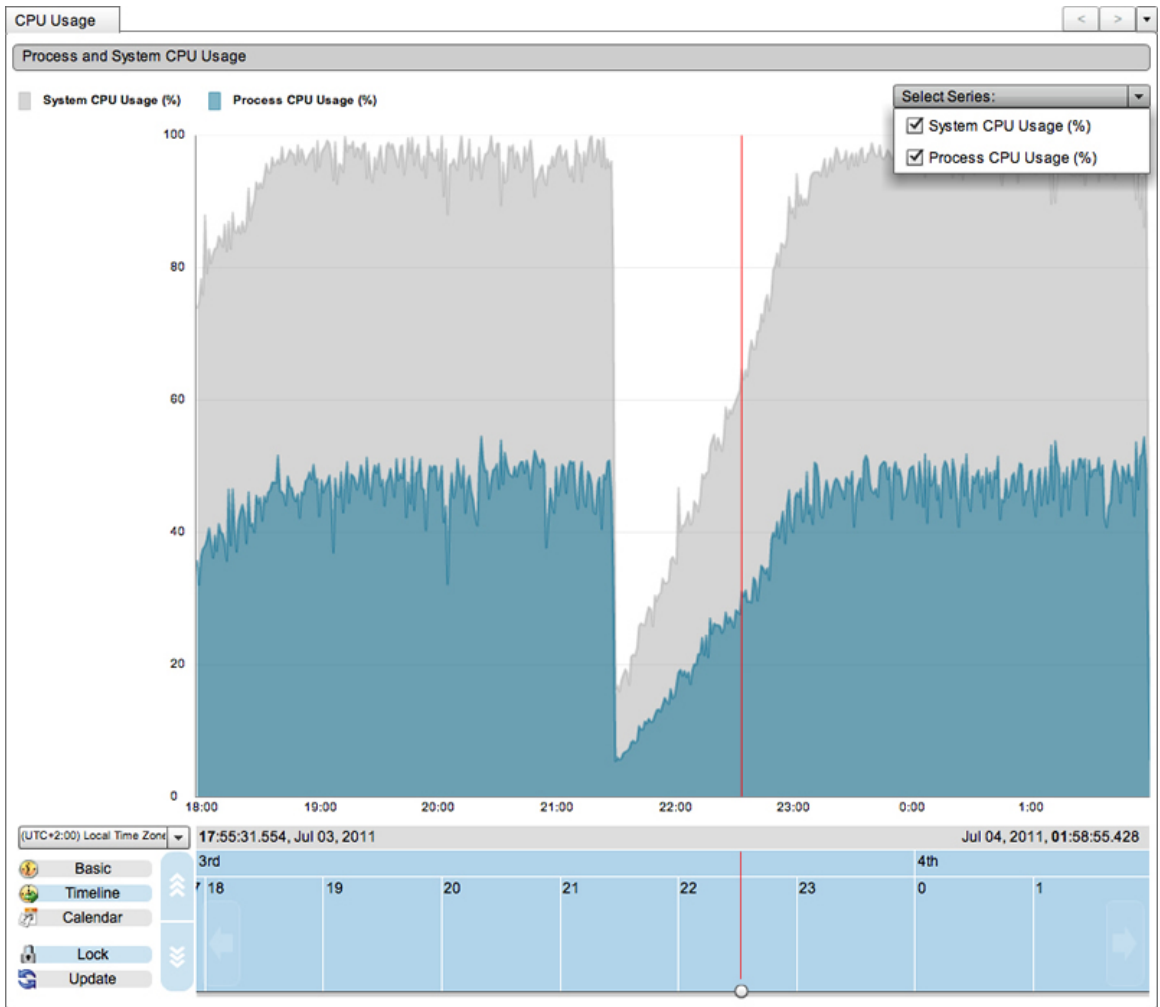


Figure 3: Chart Base Visualization - Series Select

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:

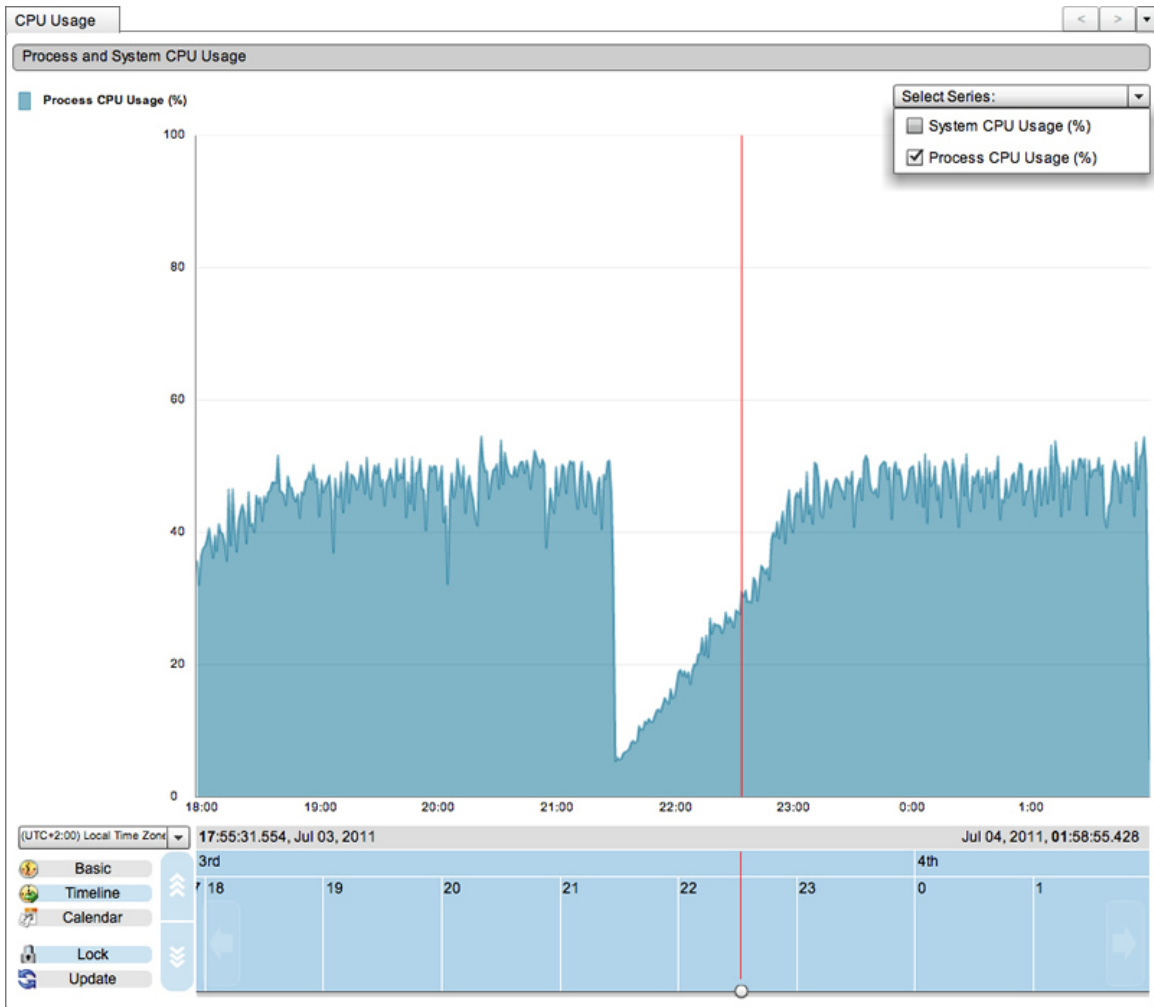


Figure 4: Chart Base Visualization - Series Select

[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 identification 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:

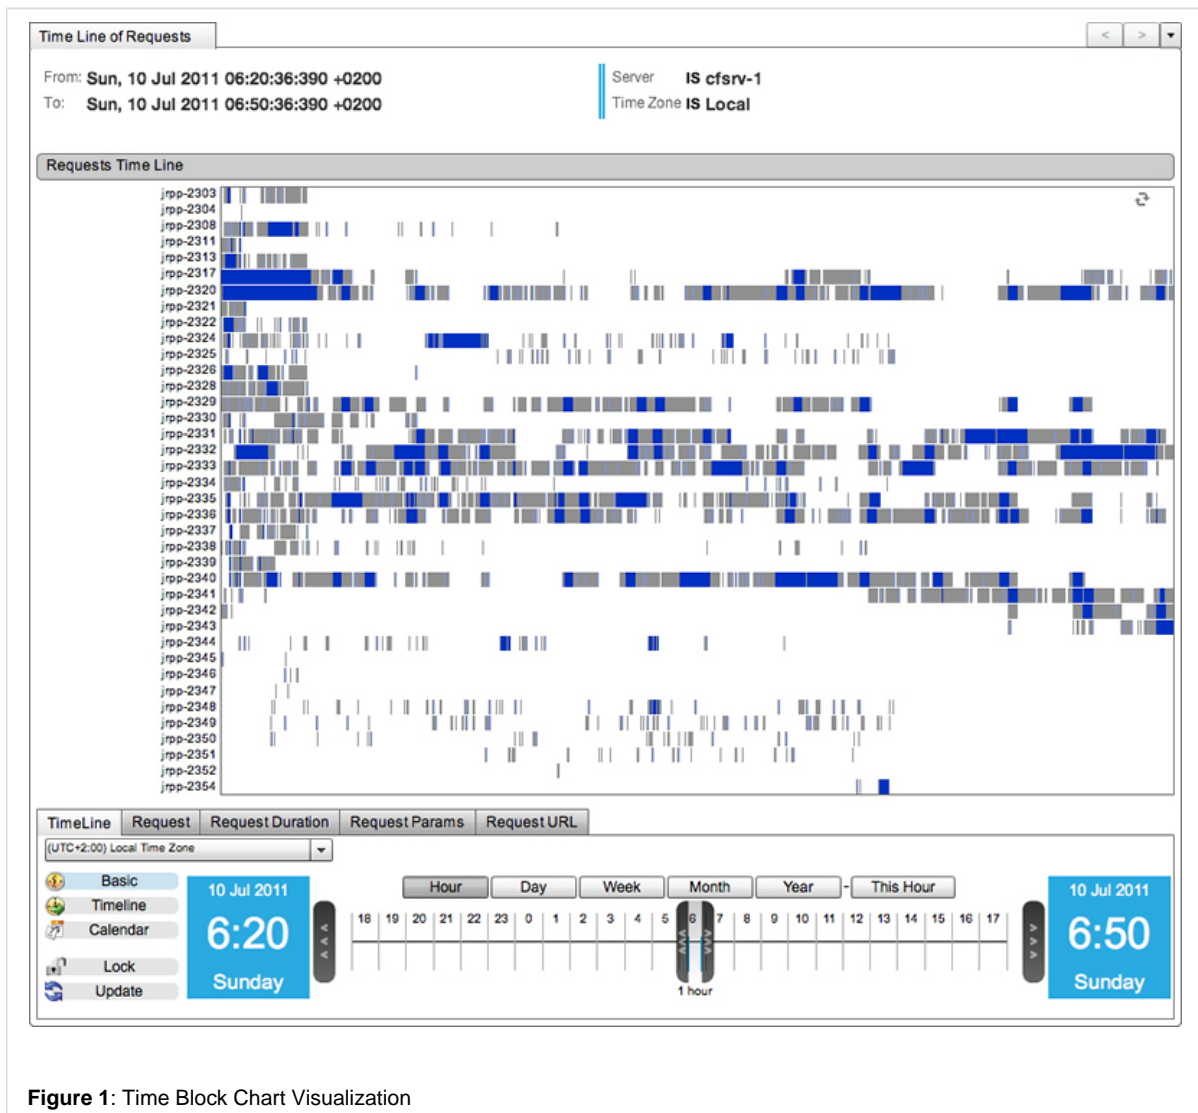


Figure 1: Time Block Chart Visualization

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:

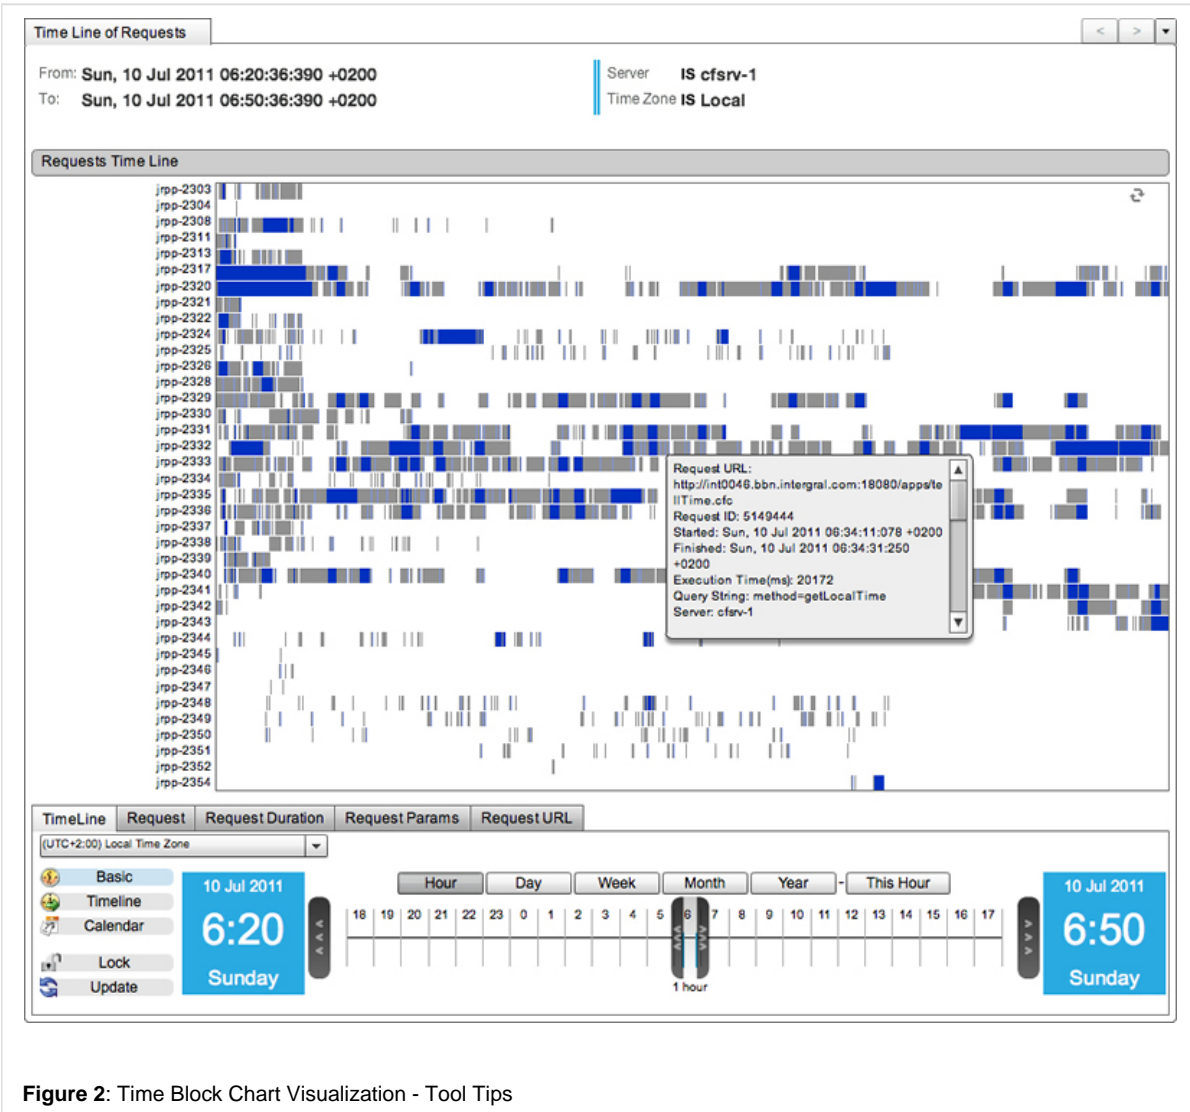


Figure 2: Time Block Chart Visualization - Tool Tips

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.

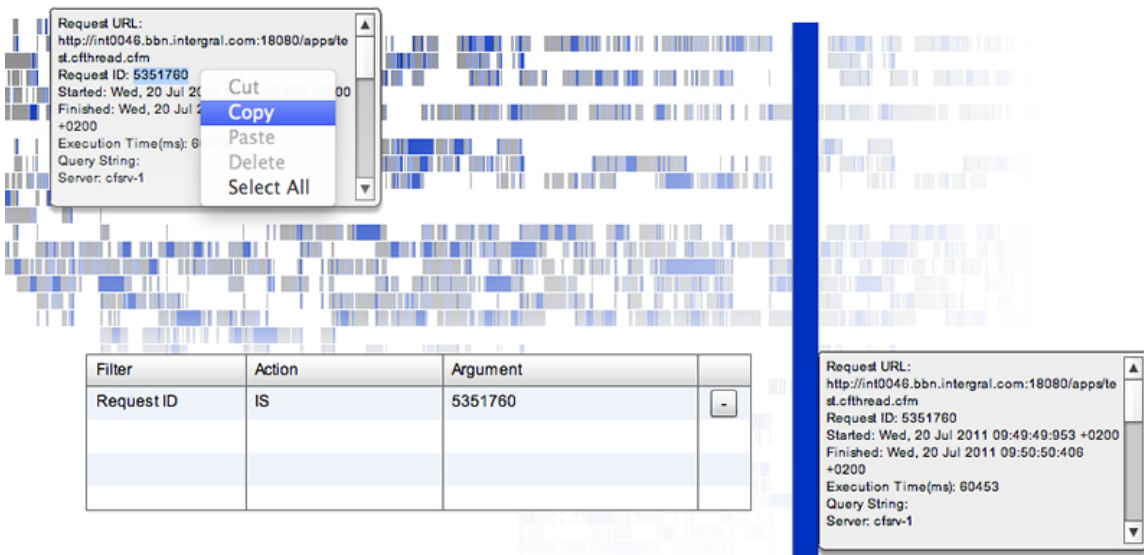


Figure 3: Time Block Chart Visualization - Copy Tool Tip

[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).

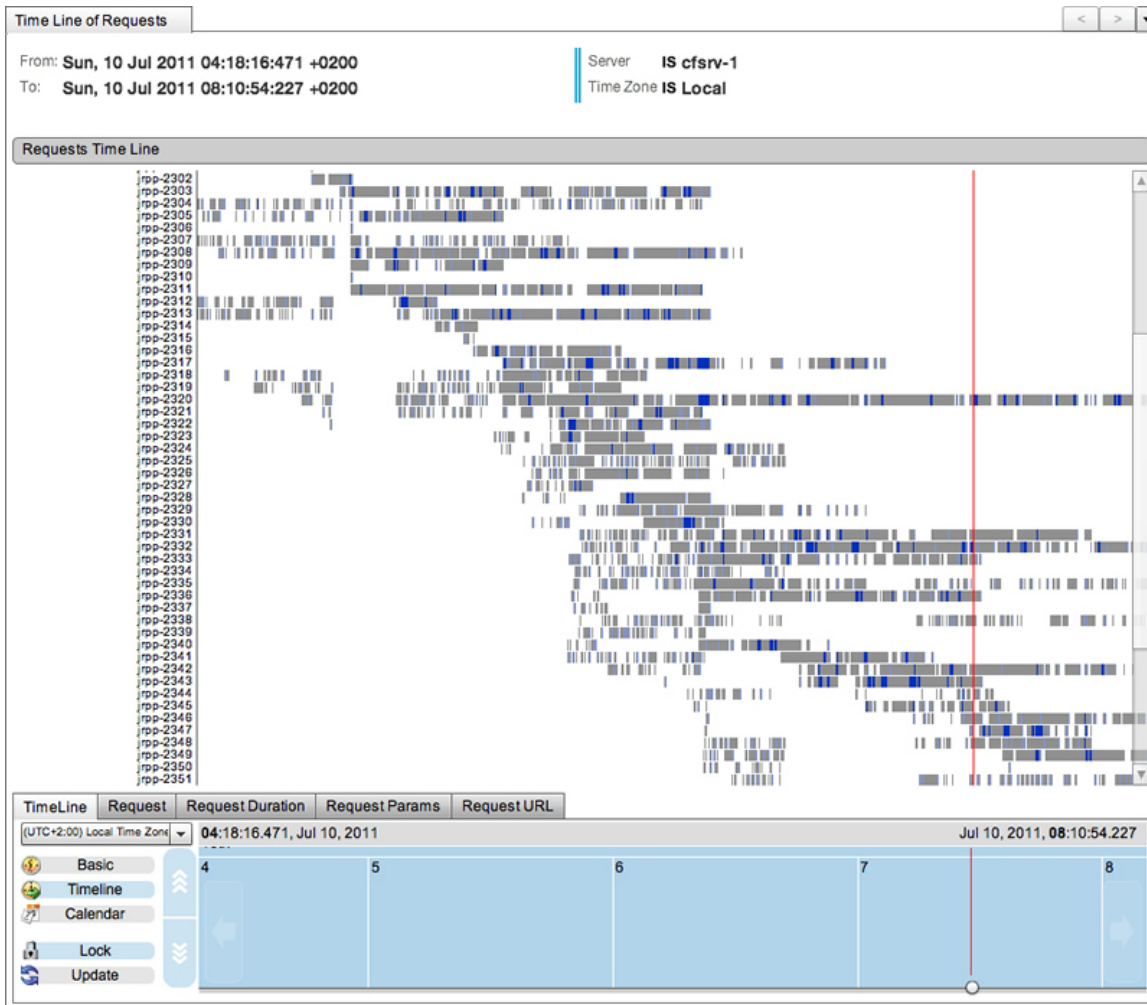


Figure 5: Time Block Chart Visualization - Chart Scrolling

[Back to the top](#)

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:

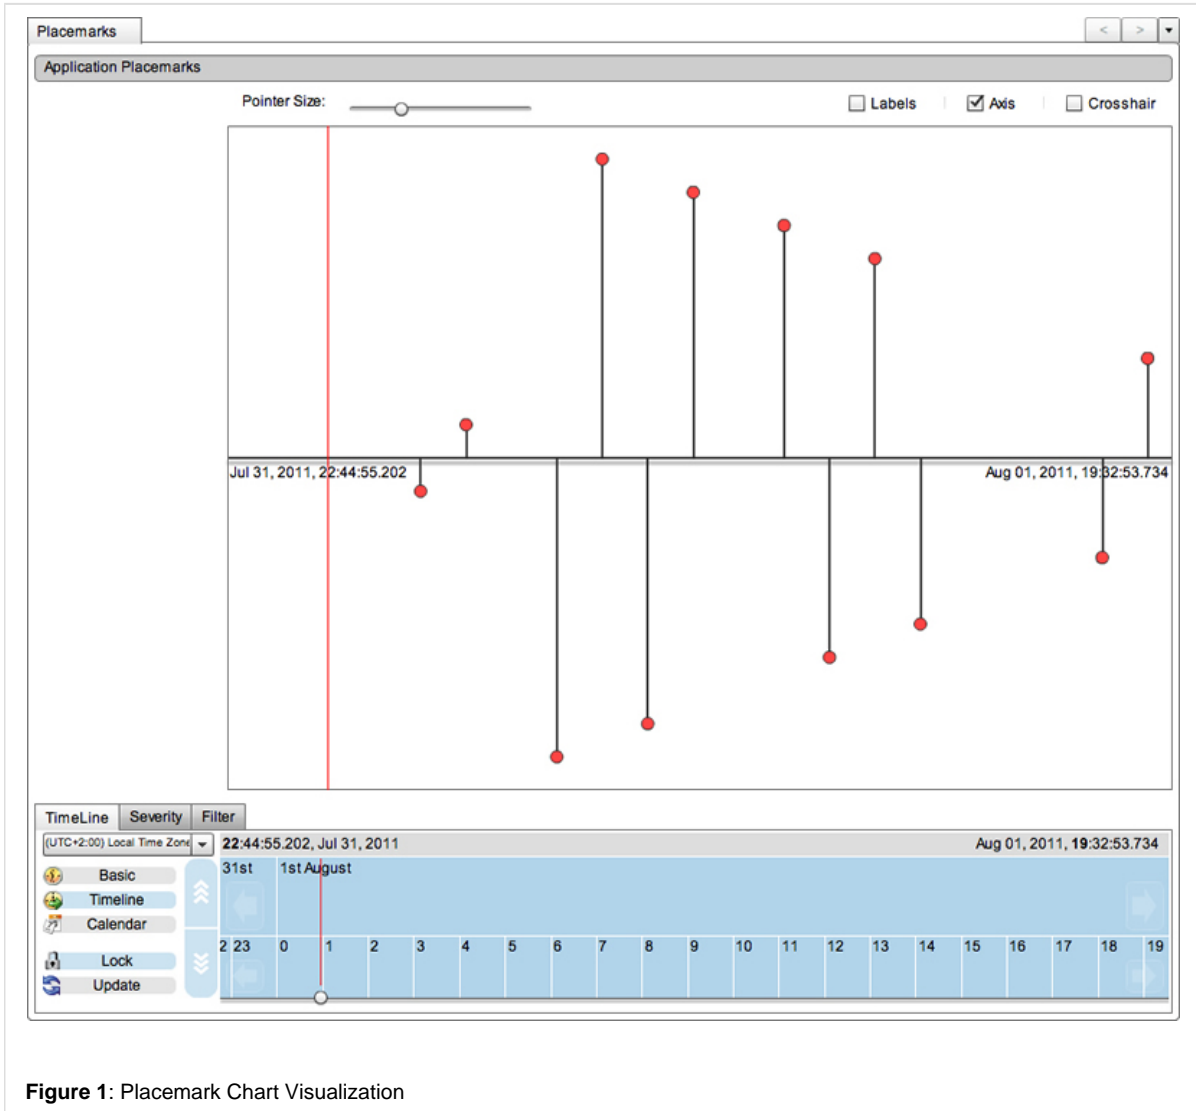



Figure 1: Placemark Chart Visualization

[Back to the top](#)

Placemark Grouping

The placemark Chart Visualization supports "grouping" allowing placemarks 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:



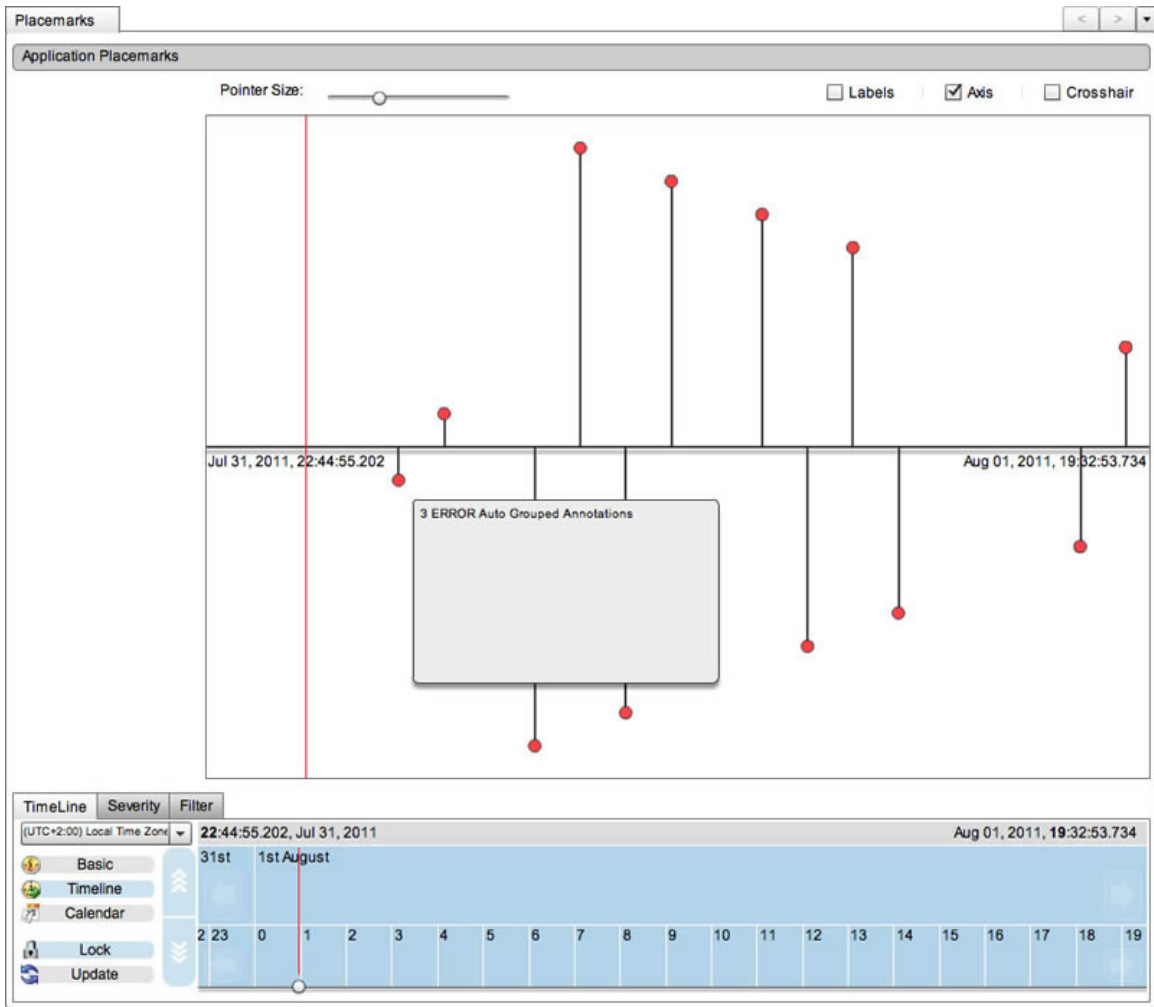


Figure 2: Placemark Chart Visualization - Grouping

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.

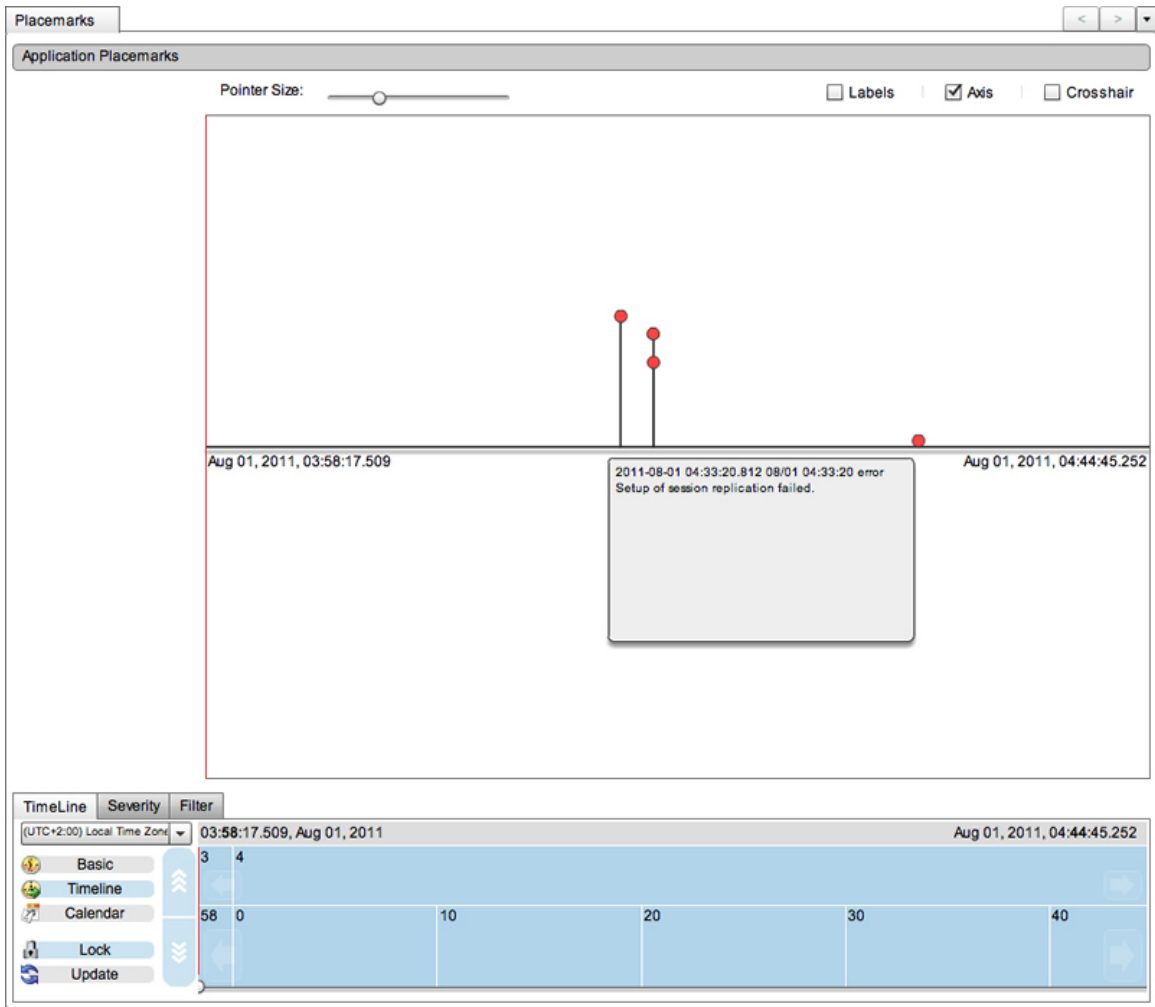


Figure 3: Placemark Chart Visualization - Grouping


Once placemark items are rendered at their 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 Placemark 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 chart's 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:



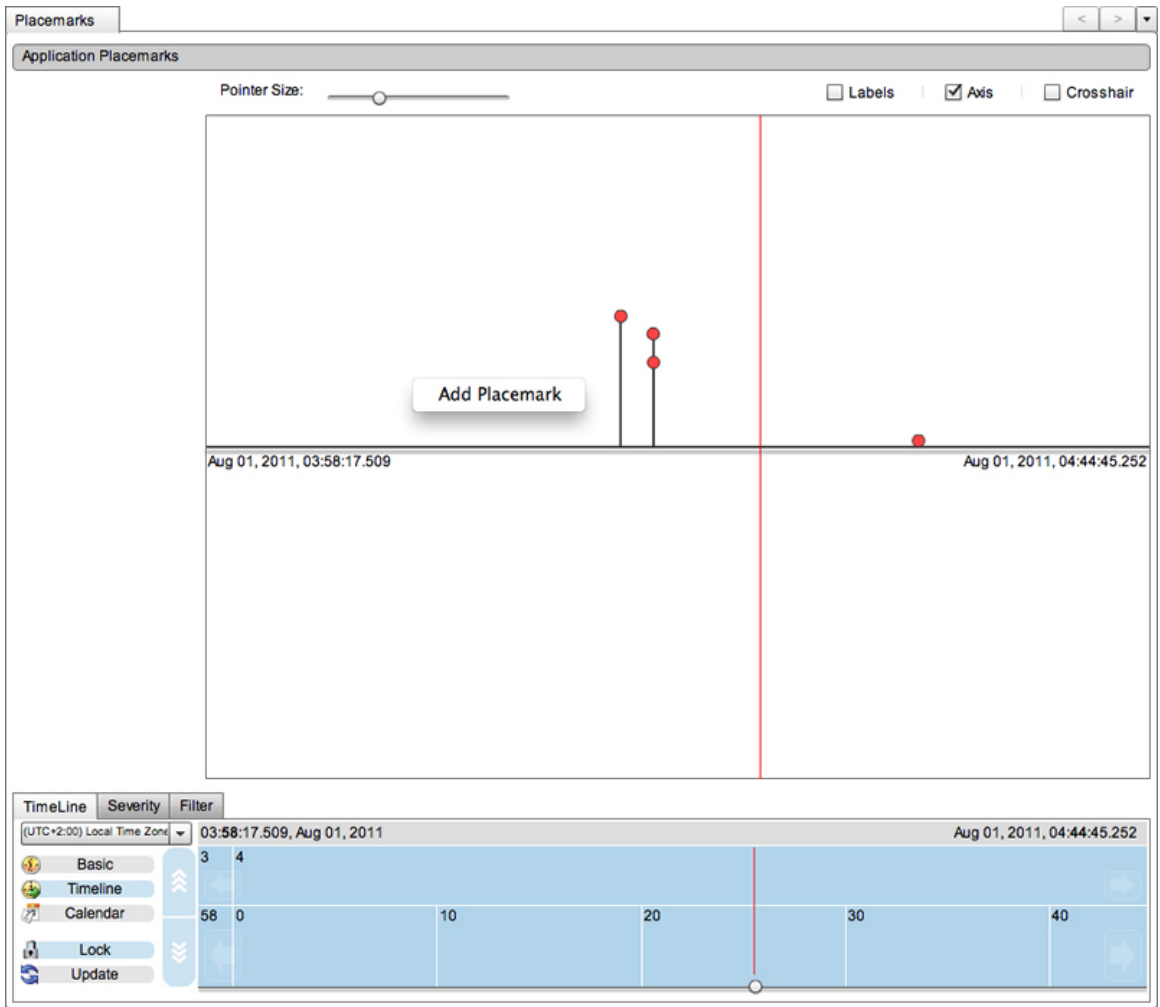


Figure 4: Placemark Chart Visualization - Add Placemark

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

The 'Add Placemark' dialog box has a title bar with a close button. Inside, there are four main sections:

- Severity:** A dropdown menu currently showing 'INFO'.
- Point Color:** A color selection area showing a green square.
- Placemark Text:** A text input field containing the text 'This is a test "info" placemark'.
- Placemark URL:** An empty text input field.

 At the bottom right of the dialog is a button labeled 'Add 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 placemark 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.

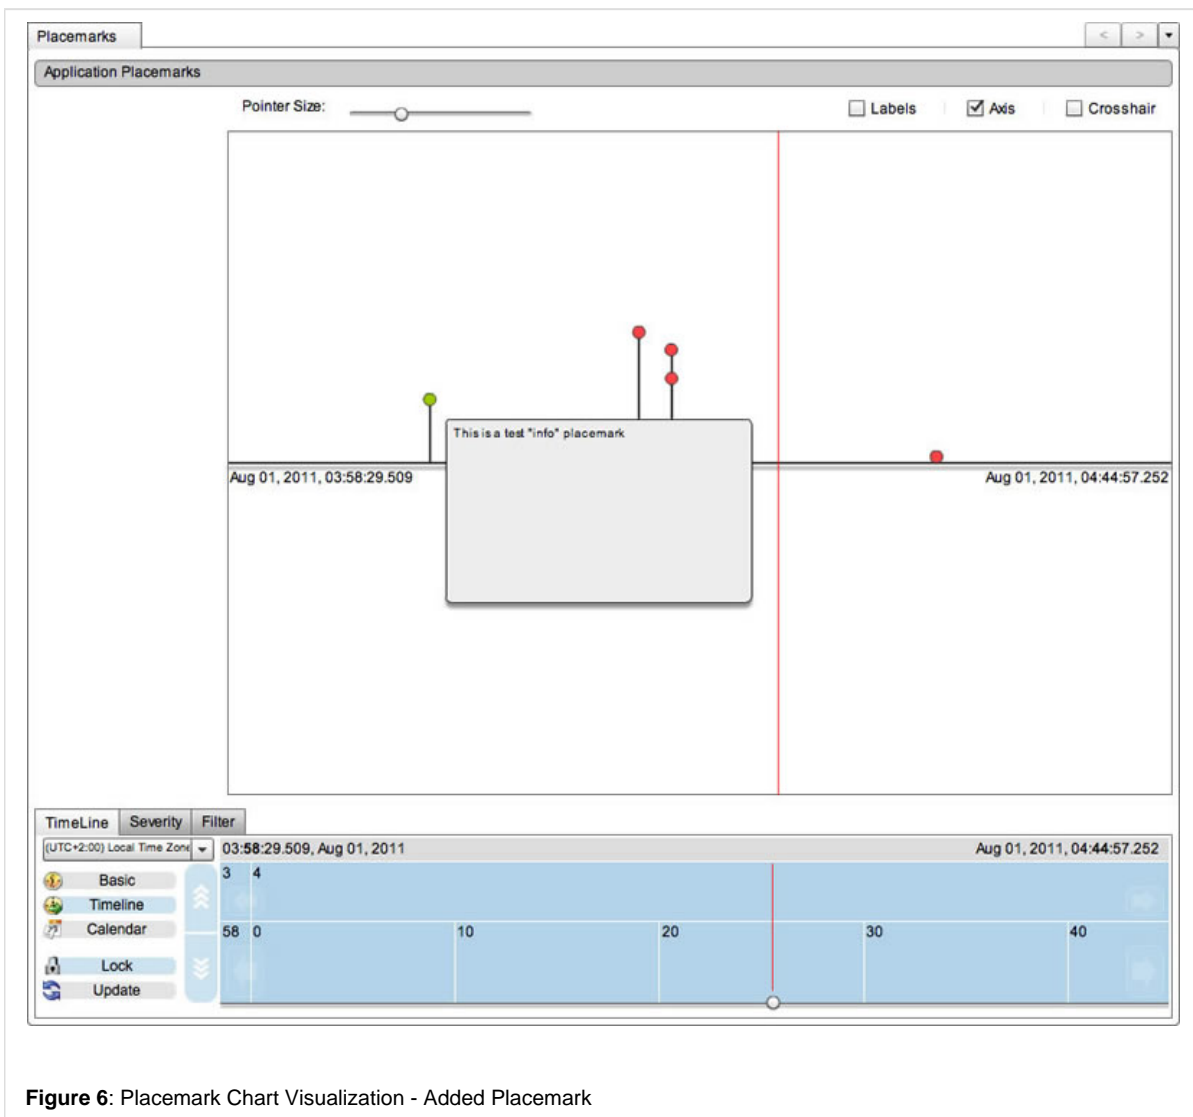


Figure 6: Placemark Chart Visualization - Added Placemark

[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:



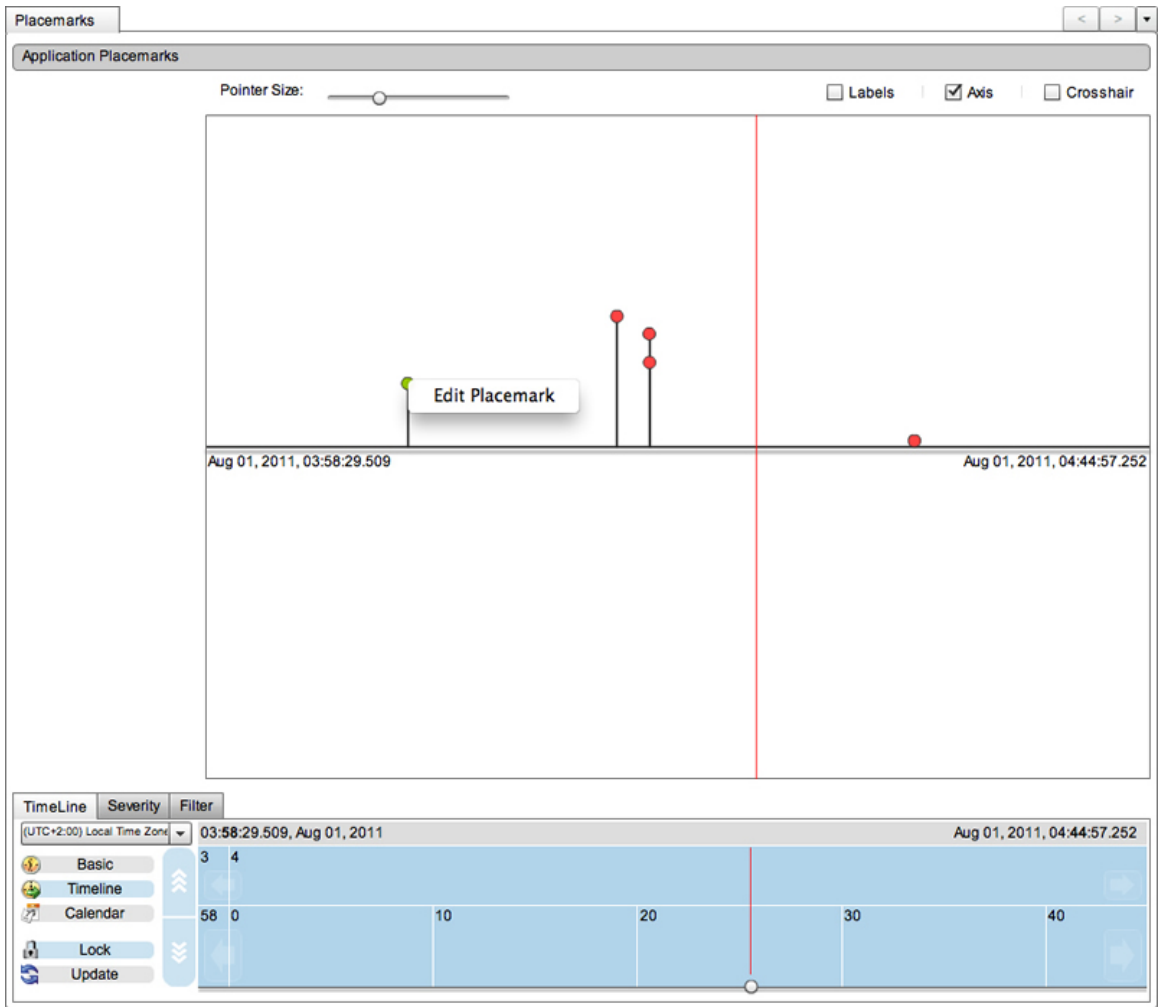


Figure 7: Placemark Chart Visualization - Edit Placemark

This will display the "Edit Placemark" dialogue:

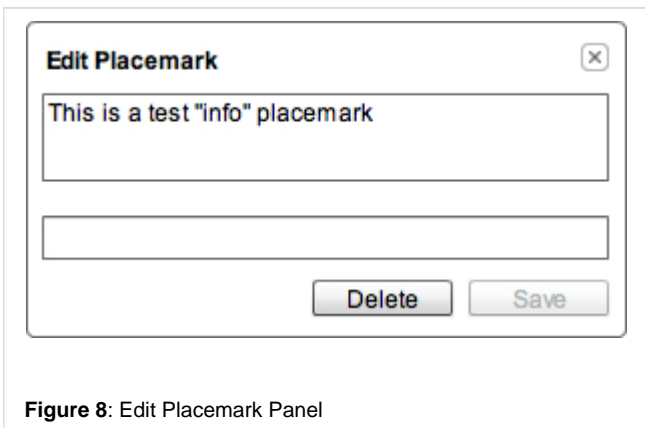


Figure 8: Edit Placemark Panel

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 whether 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 visibility 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.



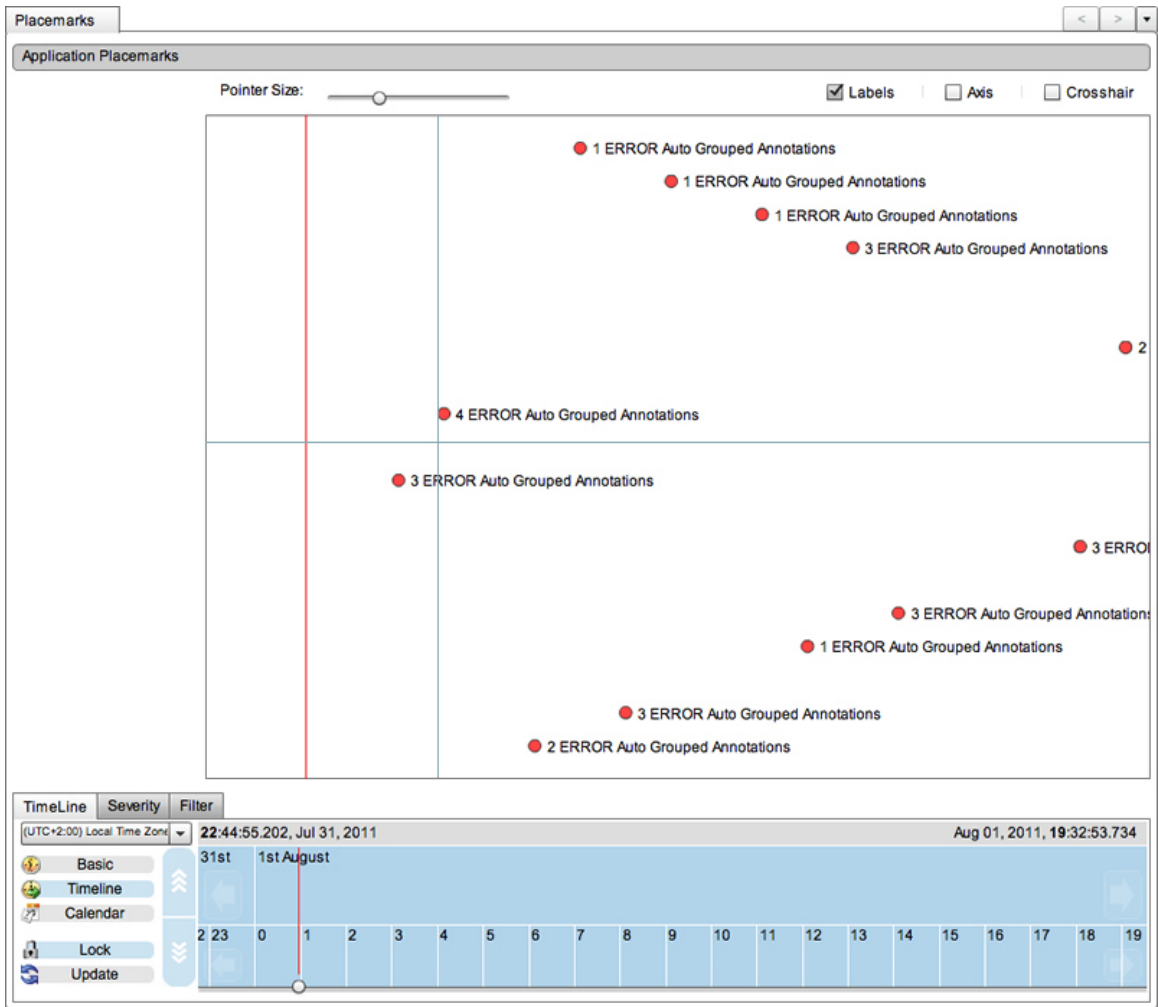


Figure 9: Placemark Chart Visualization - Crosshair

[Back to the top](#)

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:

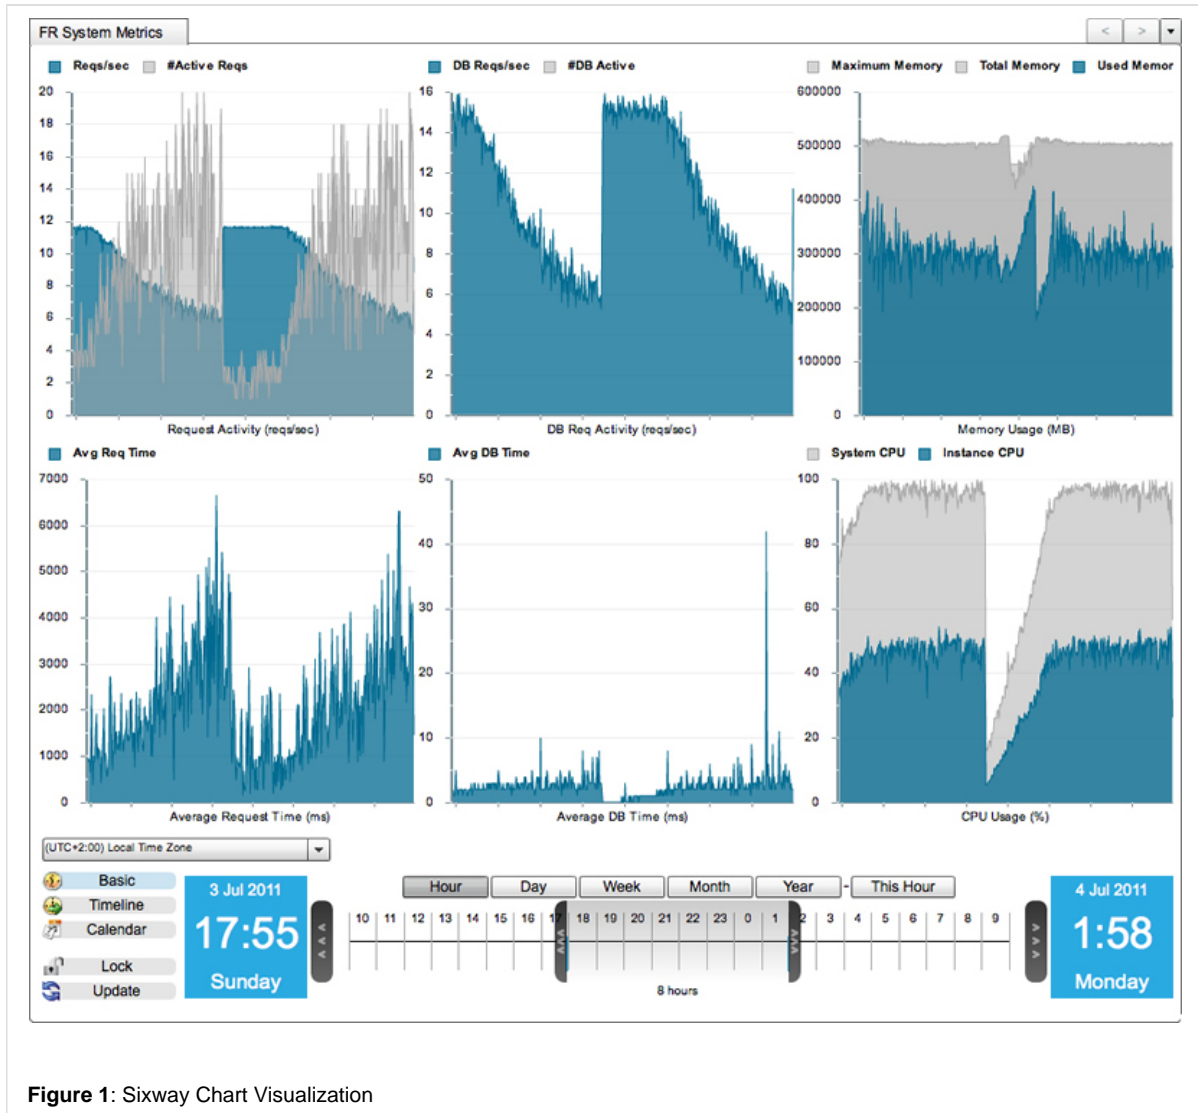



Figure 1: Sixway Chart Visualization

[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 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.



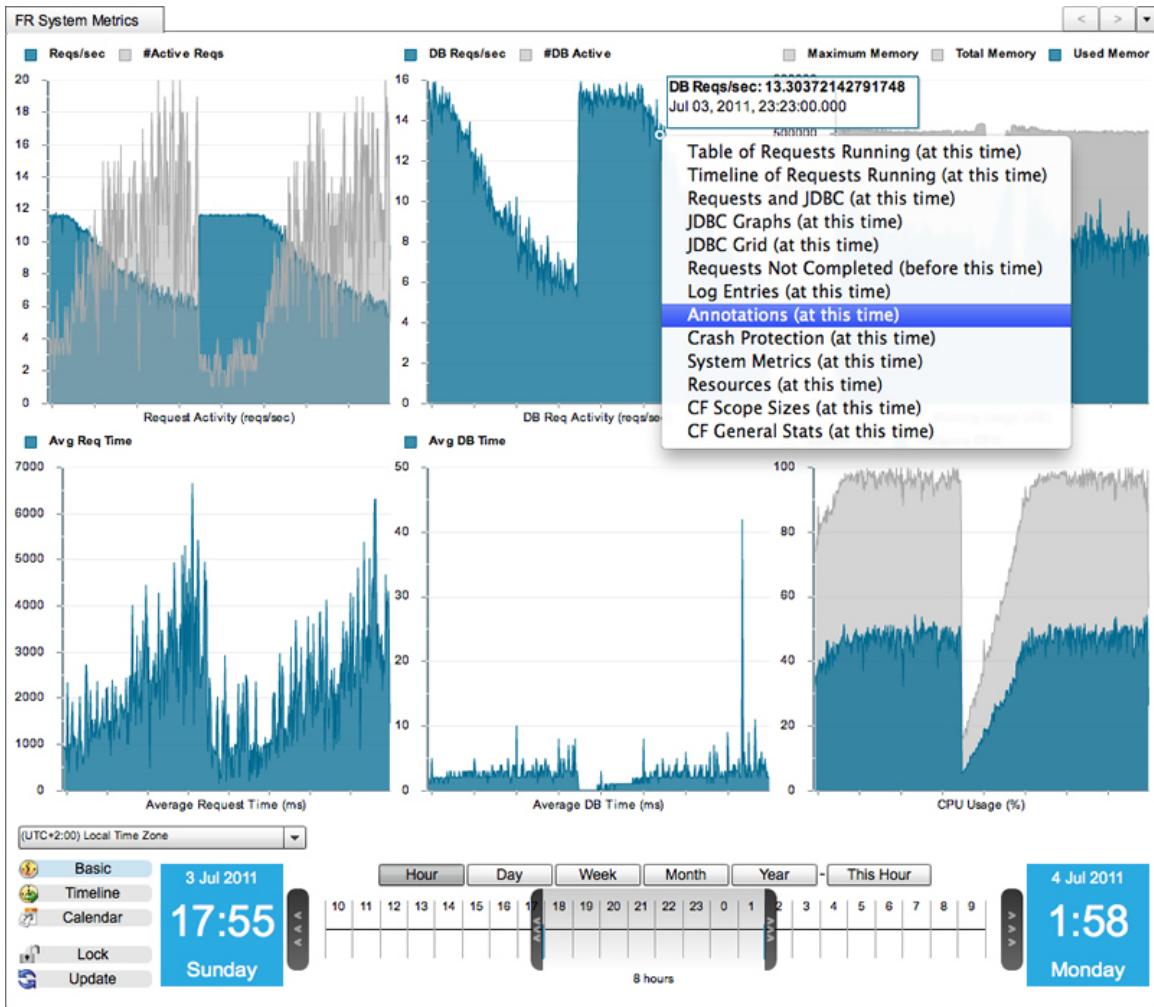


Figure 3: Sixway Chart Visualization - Context Menu

For more information about the SixWay Chart Visualization and its content / usage please see the Overview and FusionReacto r 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).

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:





Application Options Tabs Help


?

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

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. |



?

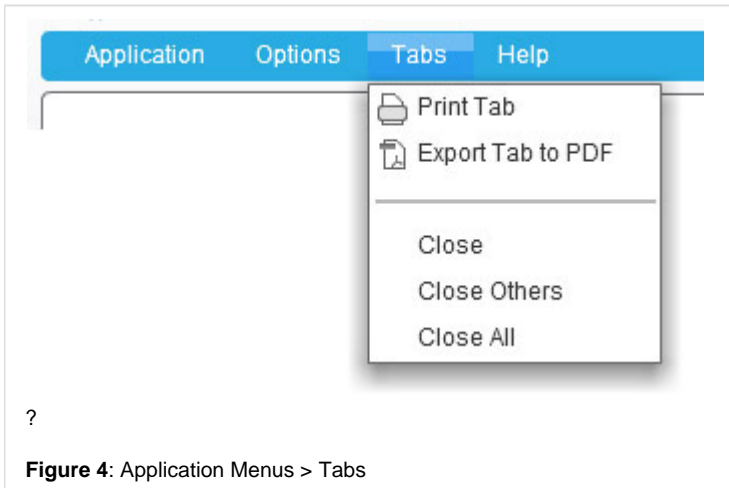
Figure 3: Application Menus > Options

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. |
|-----------|--|

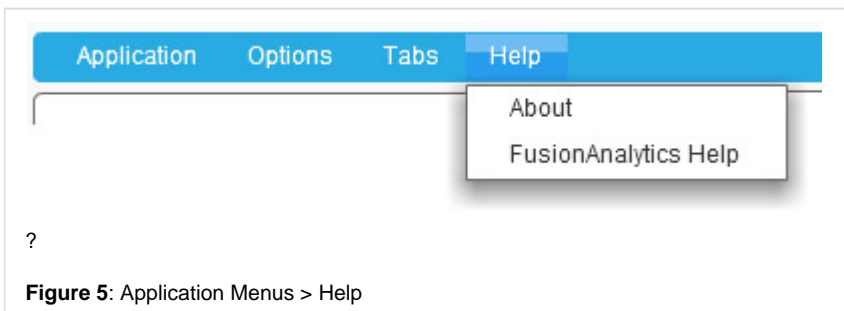
| | |
|-------------------|---|
| Expert Tab to PDF | Only available in the AIR client. If you select this menu item then an export dialog will be displayed in order to create a PDF file from the contents of the currently selected tab. |
| Close Tab | Close the currently selected tab. |
| Close Others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



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. |



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.

The app
columns
column
loaded i



?

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.



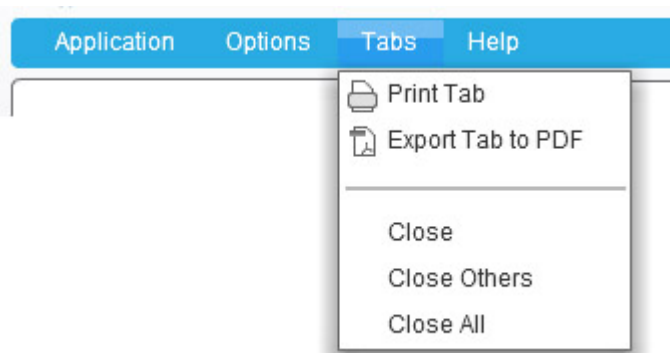


Figure 6: Printing and Exporting Tabs

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](#)

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 :

Opening
To access the
dialog box
Options



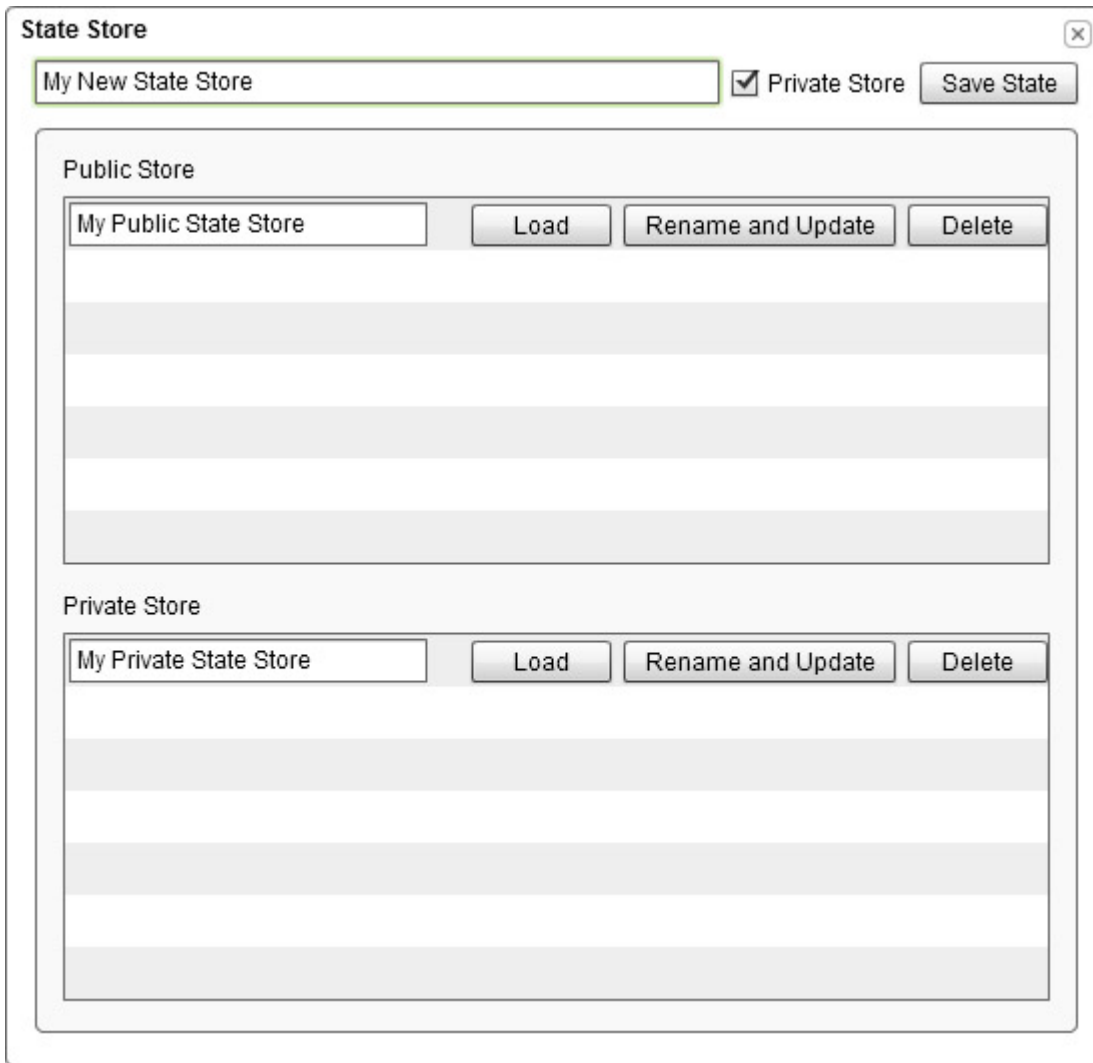


Figure 1: State Store Panel

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.

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 saved 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.



Private Store

My Private State Store

Load

Rename and Update

Delete

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.

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 Visualizations:

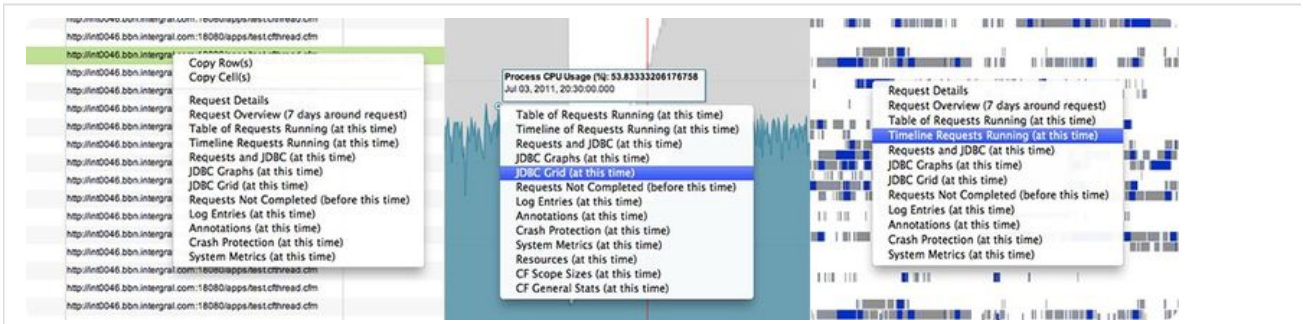



Figure 1: Application Context Menus

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 Cf Menu contain menu items that are relevant to the current Perspective and the clicked data point. The menu items are (generally) links to application relating to the selected data point and time frame:



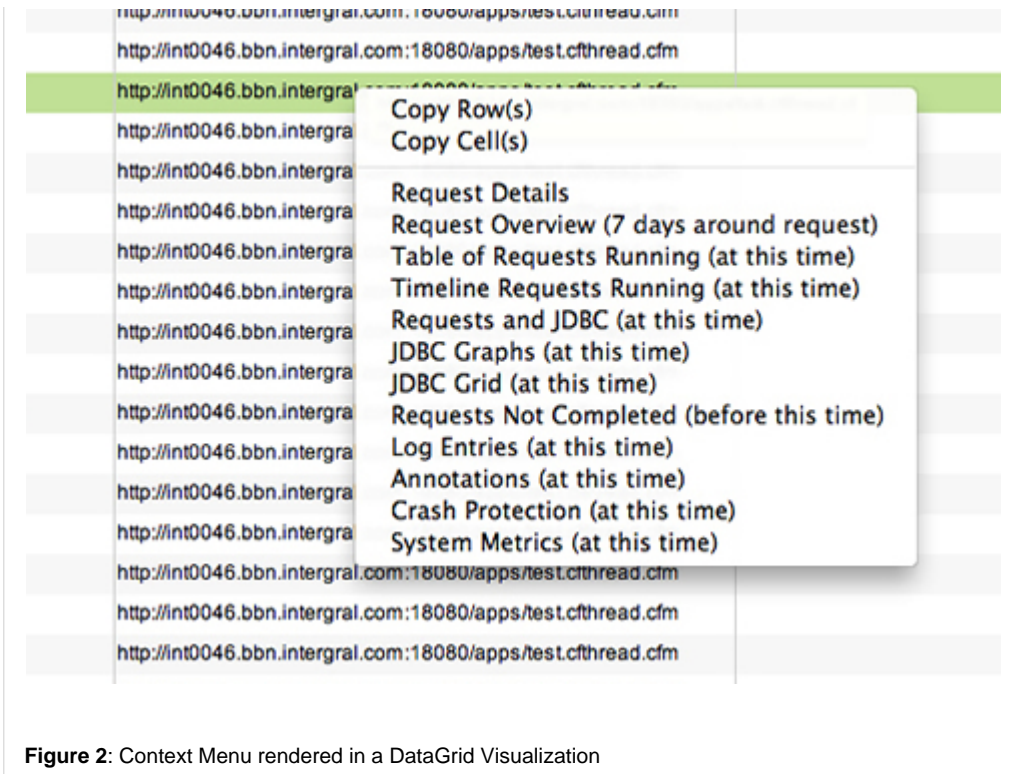


Figure 2: Context Menu rendered in a DataGrid Visualization

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)". You can use these menu items to copy content from 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) in 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) |

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

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:

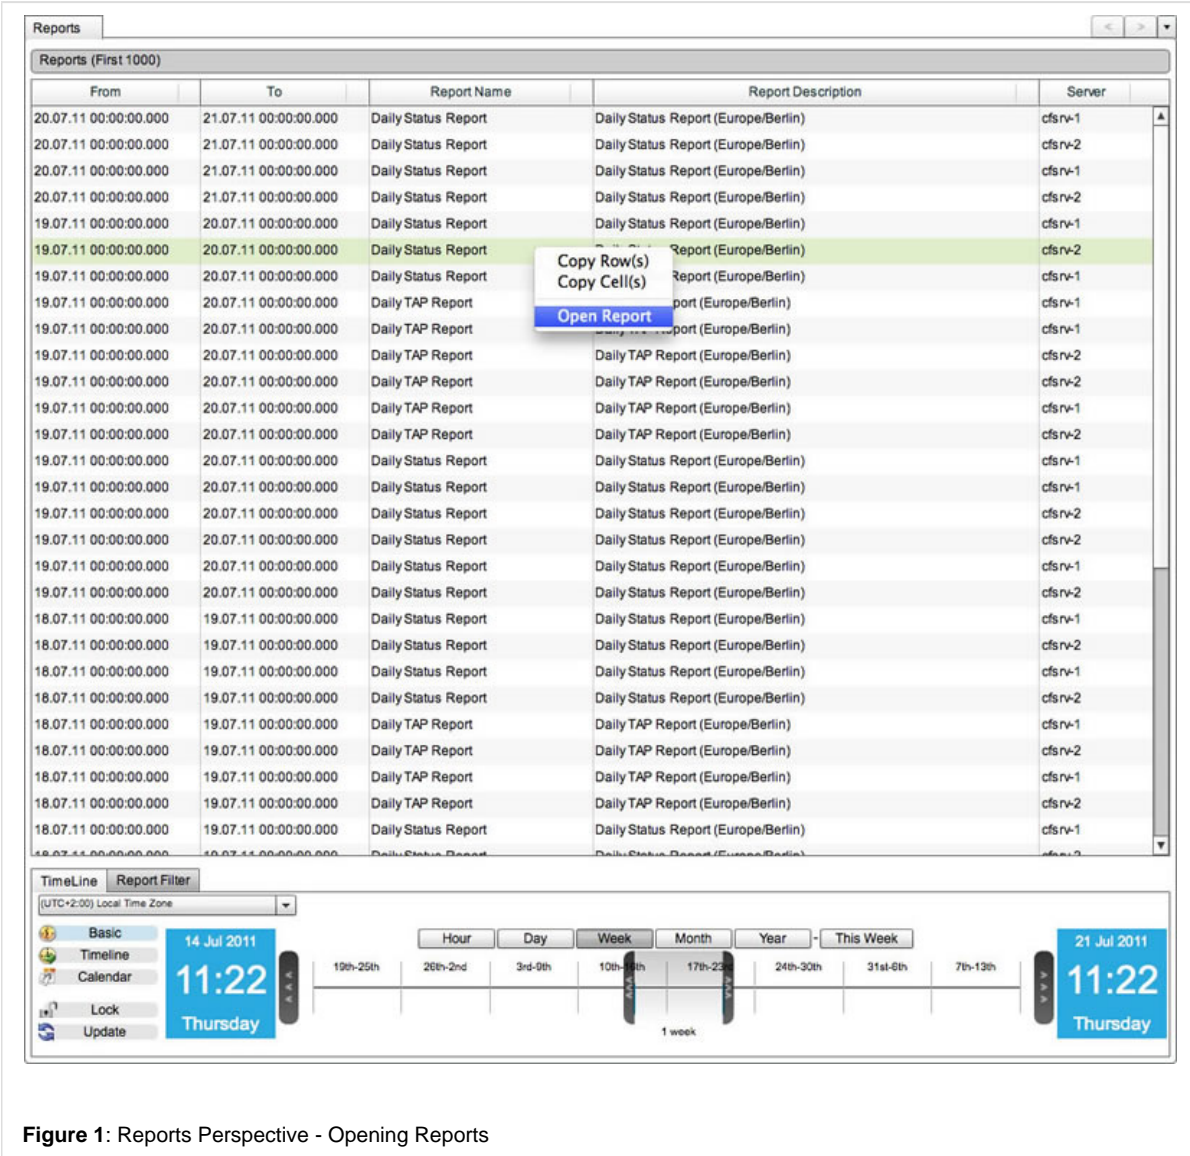



Figure 1: Reports Perspective - Opening Reports

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.



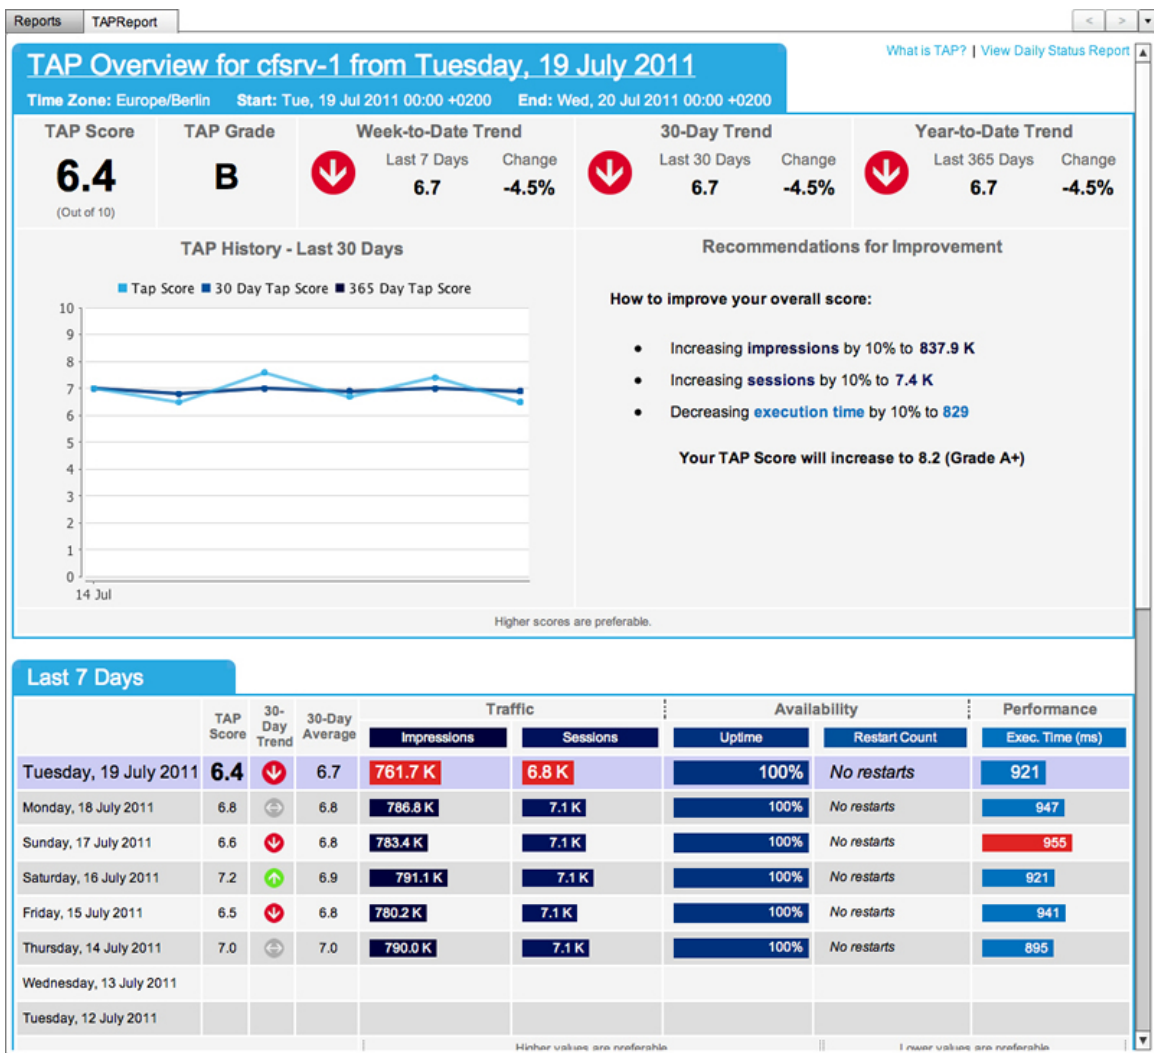


Figure 2: TAP Report

Learn more about TAP Reports

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.

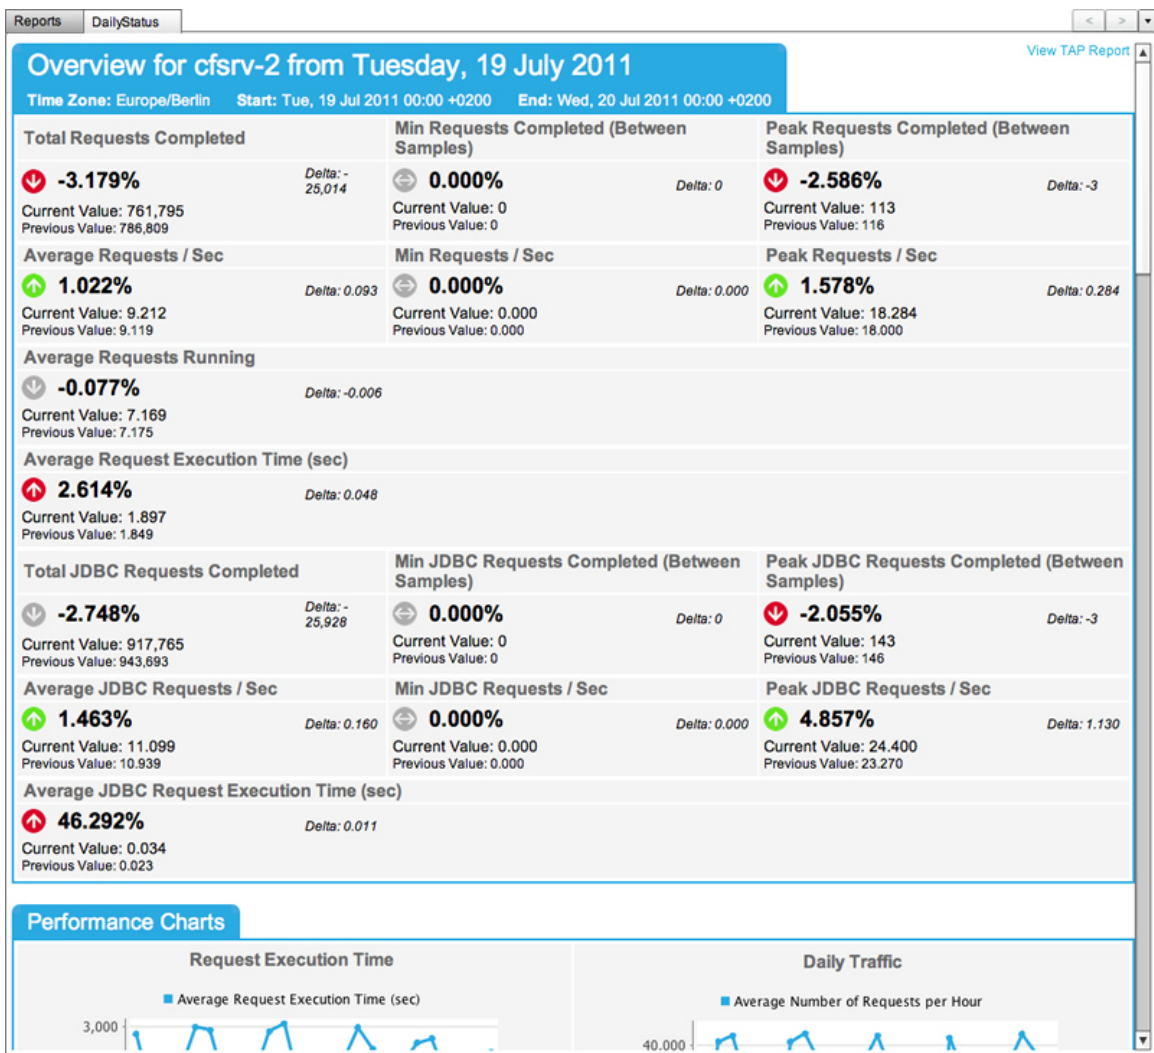


Figure 3: Daily Status Report

Learn more about Daily Status Reports

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

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:

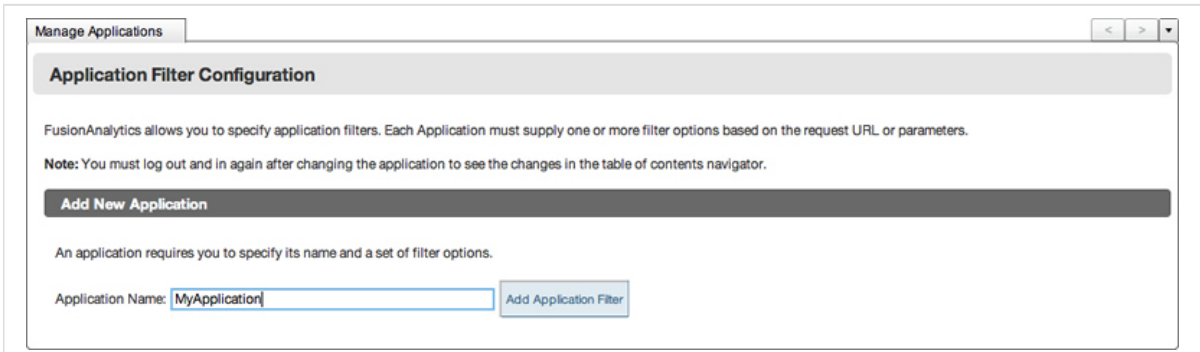


Figure 1: Adding Application Filters

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.

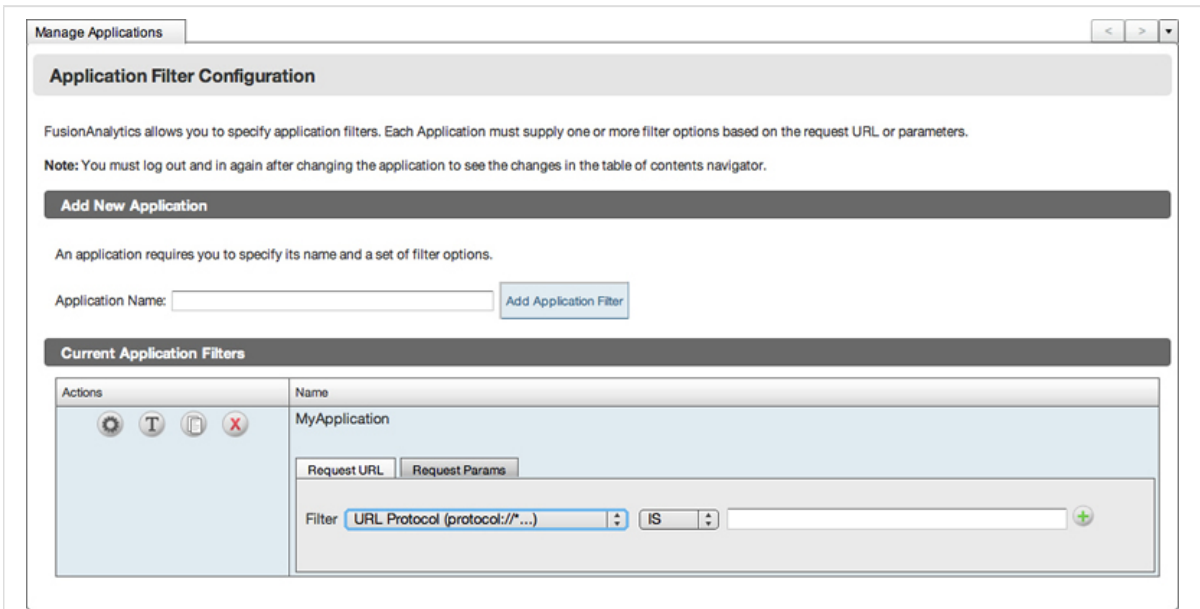


Figure 2: Added Filters

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

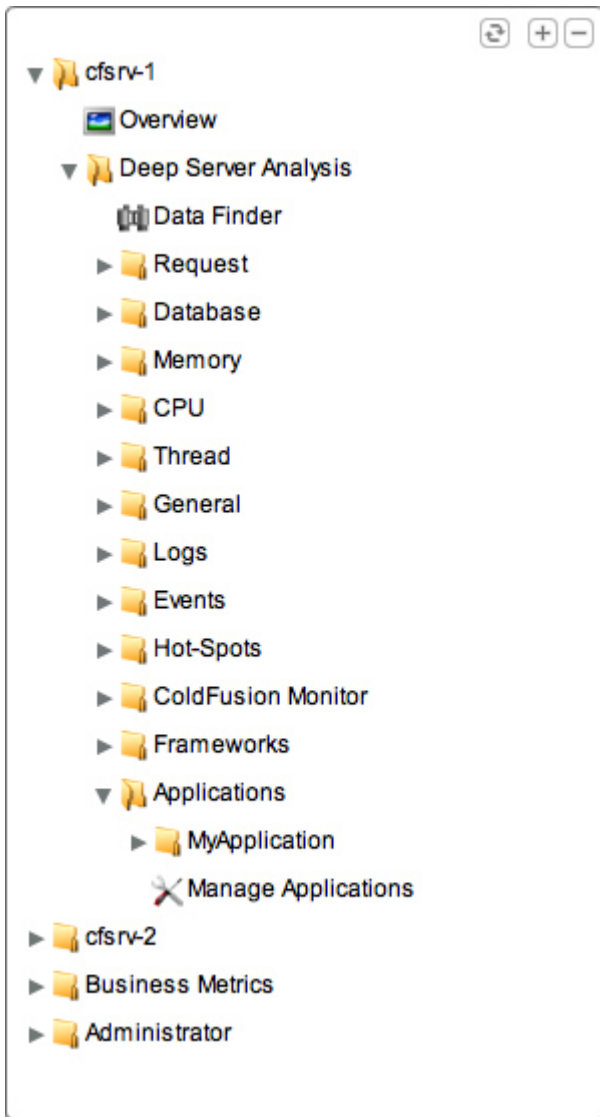


Figure 3: Added Filters listed in TOC

[Back to the top](#)

Request URL

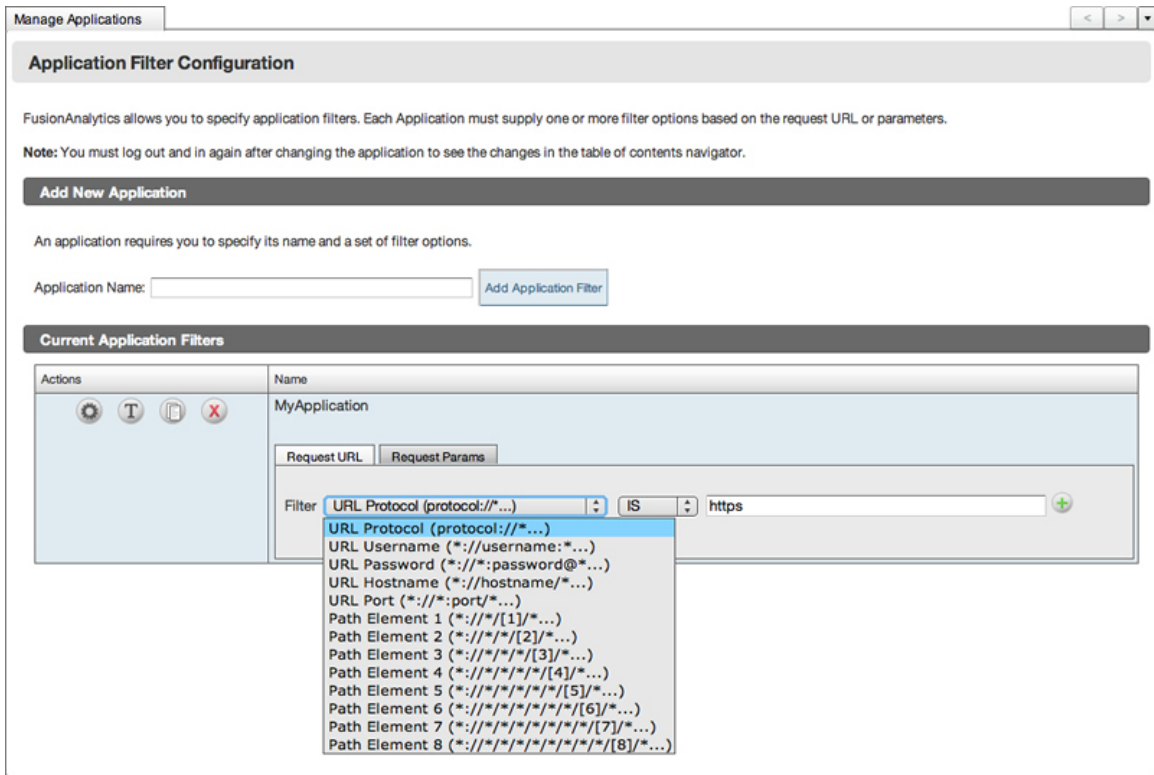



Figure 4: Adding Request URL Arguments

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.



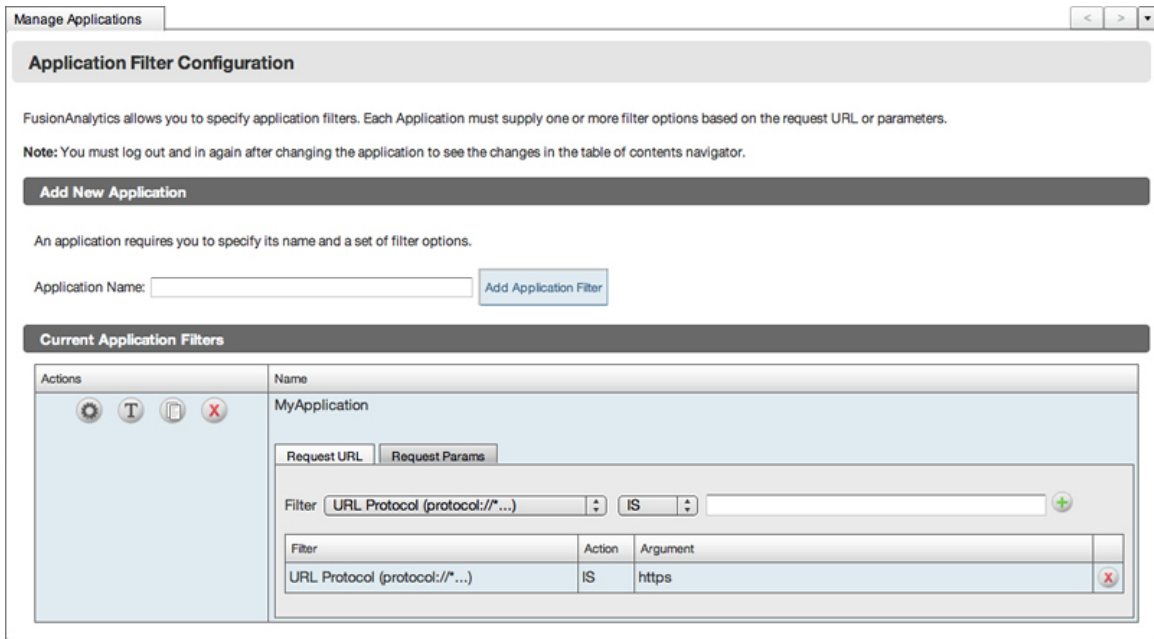


Figure 5: Request URL Arguments

Back to the top

Request Params

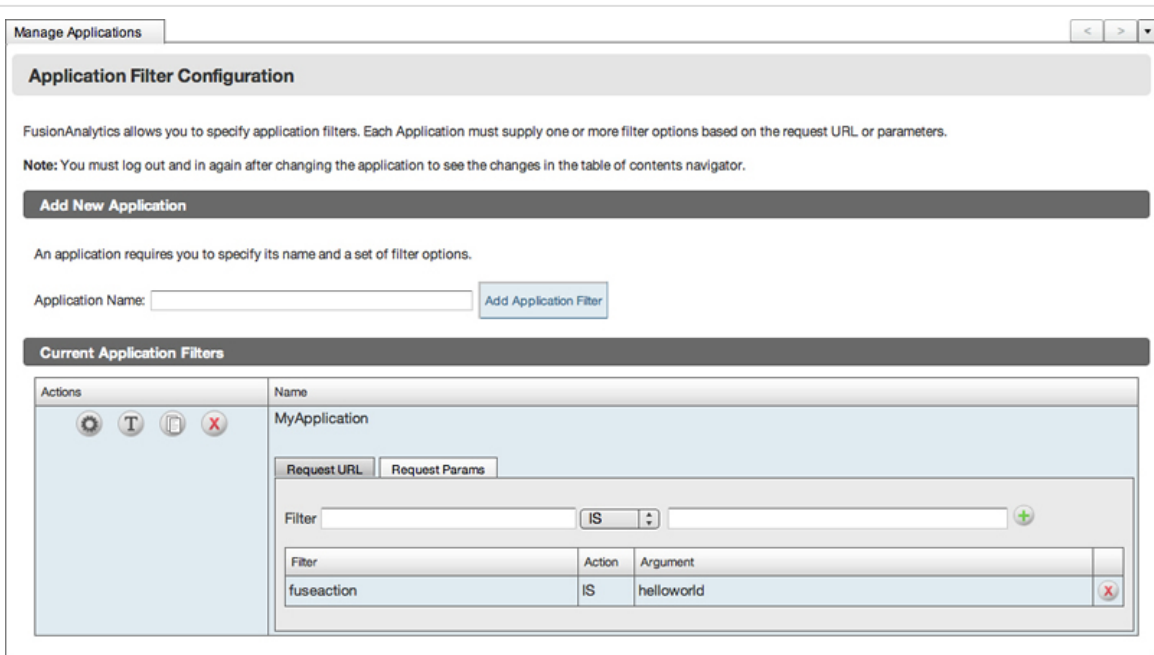


Figure 6: Request Parameter Arguments

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:



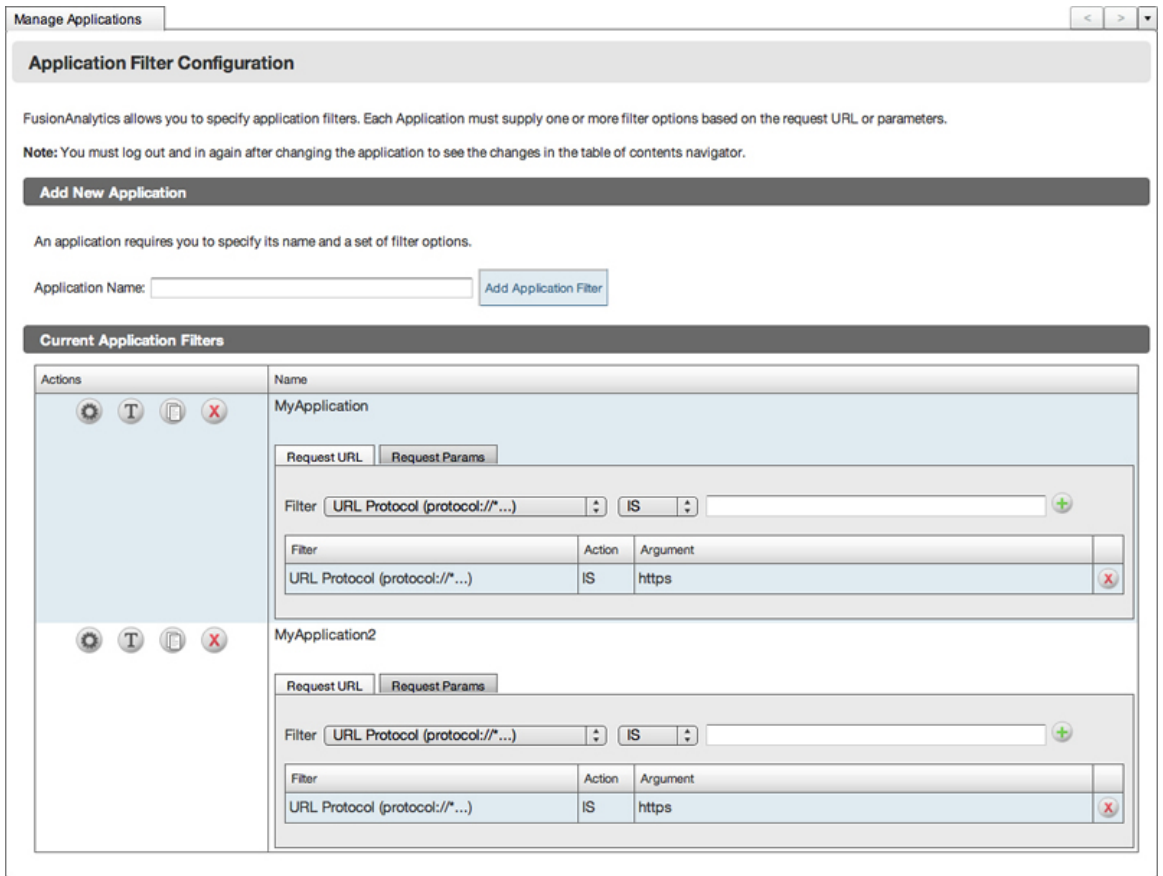


Figure 9: Copied Application Filter

See the [Setting Up Application Filters](#) page for more information

[Back to the top](#)

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 DataServices.

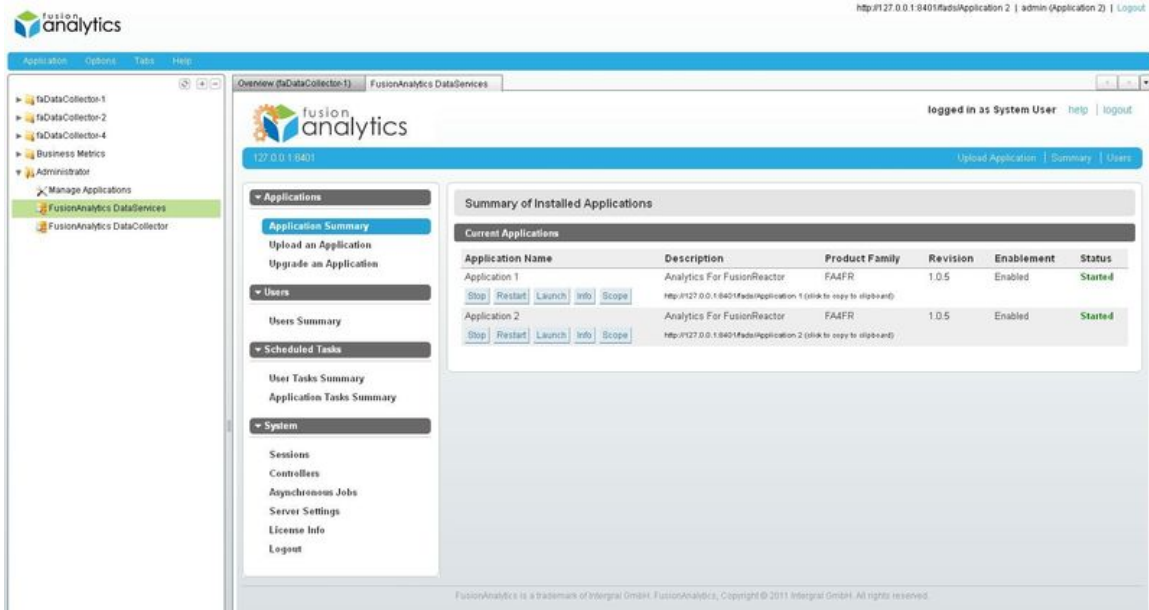


Figure 1: FADS from the client

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 DataCollector.

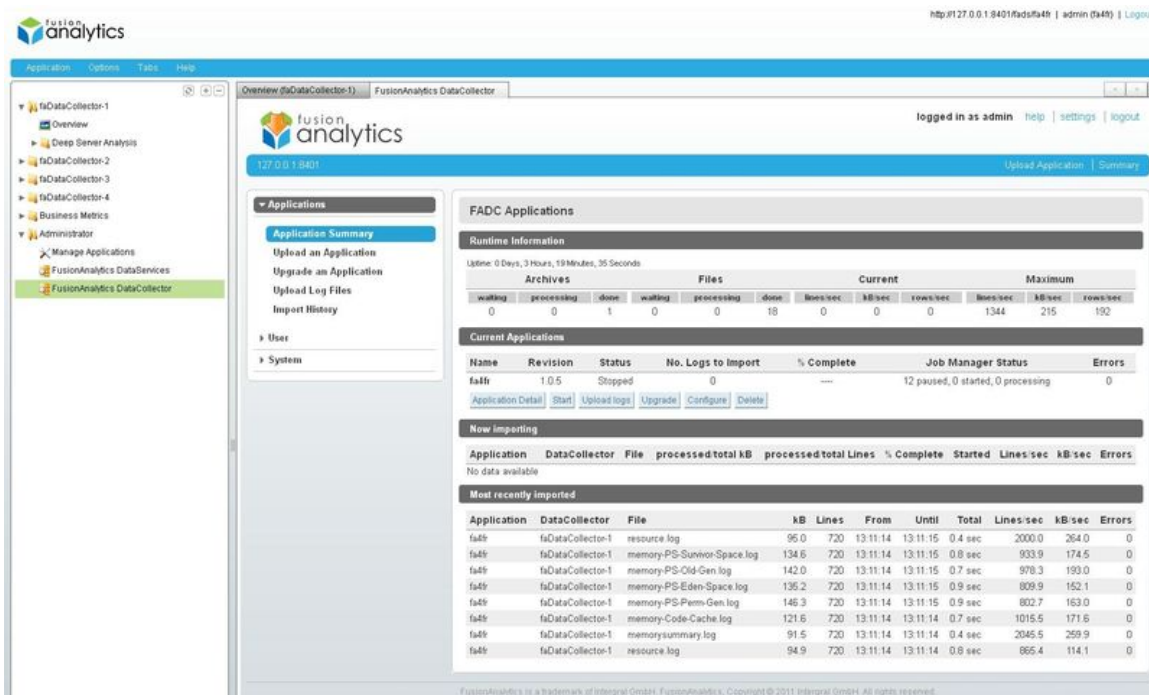


Figure 1: FADC from the client

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
- JVM Memory Requirements

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)
 - Disk Usage Settings (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)
- APM (FADS)
 - APM Volume Switching
 - Component
 - Data Provider Programs
 - Layout
- Configuring Report Images

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 APM application packages can be found and where the installed applications reside in the directory structure of the server.

NEXT STEPS

Directory Structure

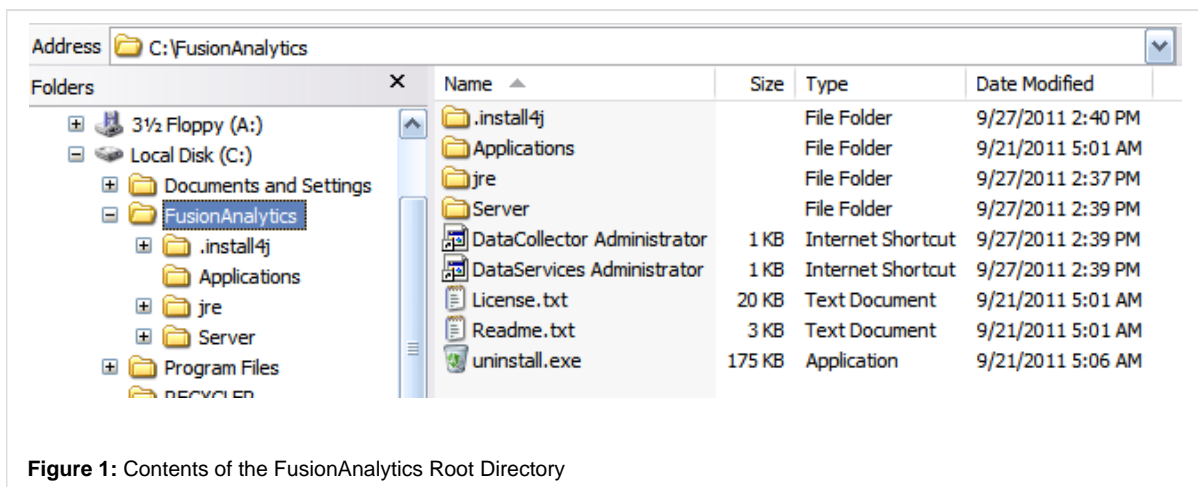
FADC Applications

FADS Applications

Directory Structure

Directory Structure

The FusionAnalytics Setup installs a Java 7 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.



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

| Description |
|-------------|
|-------------|

| Directory | |
|--------------|--|
| .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 | <code>C:/FusionAnalytics/Server/webapps/fadc</code> |
| FusionAnalytics DataServices | <code>C:/FusionAnalytics/Server/webapps/fads</code> |

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.

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:

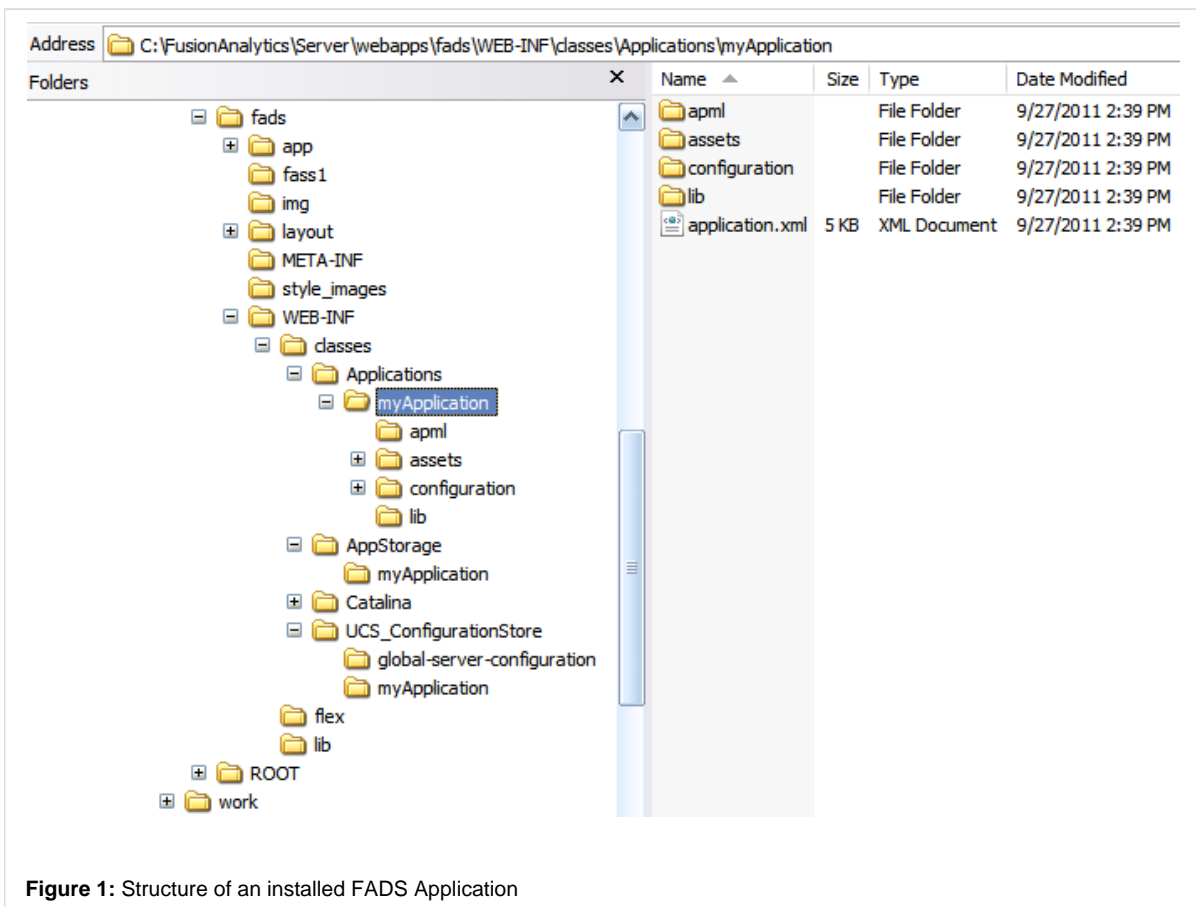


Figure 1: Structure of an installed FADS Application

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

JVM Memory Requirements

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
```

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.

for **FADS** it is defined in file

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

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

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 |

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:

```
-server
-Xms128m
-Xmx768m
-XX:MaxPermSize=256m
-Djava.awt.headless=true
-Dcom.sun.management.jmxremote.
port=3333
-Dcom.sun.management.jmxremote.
ssl=false
-Dcom.sun.management.jmxremote.
authenticate=false
-Dfusionanalytics.datacollector.
dirObserver.scanInterval=5000
-Dfusionanalytics.lm.store=C:
/FusionAnalytics/Server/data/falm
/falm.store
-Dgrails.scan.interval=3600000
```

If the default destination directory has not been changed during the setup the location of the configuration file above is

```
C:\FusionAnalytics\Server\bin\FusionAnalytics Server.
vmoptions
```

FusionAnalytics DataCollector Settings

FusionAnalytics DataCollector Settings

The FusionAnalytics DataCollector can be further configured with the following system properties.

| Name | Description | De fa ult Va lue |
|--|--|------------------------------|
| fusionanalytics. datacollector. batchSize | The number of statements in a SQL batch before committing it | 500 |
| fusionanalytics. datacollector. maxLineLength | The number of bytes a line is allowed to have before FADC rejects it as too long | 32 768 |
| fusionanalytics. datacollector. maxSize. importHistory | The maximum number of lines the Import History page displays | 100 |
| fusionanalytics. datacollector. connectionObserver Interval | 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 |

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

| | | |
|--|---|----------------------------|
| 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 | yyy y- M M- dd |
| fusionanalytics. datacollector. disableCommitStat ement | Disables database commits; useful for performance analysis only | false |
| fusionanalytics. datacollector. errorsBeforeFailure | not used | 0 |
| fusionanalytics. datacollector. importThreads | The number of import threads | #co res |
| fusionanalytics. datacollector. dirObserver. scanInterval | The number of seconds between two checks for modified /new/removed files in the application directories | 50 00 |
| fusionanalytics. datacollector. importShutdownTi meout | The number of seconds to wait for an import to stop before it will be killed | 15 |
| fusionanalytics. datacollector. jobShutdownTimeo ut | The number of seconds to wait for a job to finish before it will be killed | 15 |
| fusionanalytics. datacollector. maxLenErrorMessage | The maximum length of an error message to be written to the logs before it will be truncated | 10 24 |
| fusionanalytics. datacollector. urlencoding | | UT F-8 |
| fusionanalytics. datacollector.web. recentImportSize | The number of most recently processed log files to show in the user interface | 5 |
| com.intergral. fusionanalytics.ndc. maxPoolSize | The number of threads the FADC thread pool used to import files will accept | 8 |

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`

Sending data to FusionAnalytics

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

```
<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>

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.

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 Settings] button


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 |

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



Edit Server Details

Protocol:

Secure Connection:

Host:

Port:

Application:

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.

JVM Memory Requirements

JVM Memory Requirements

The JVM used by FusionAnalytics comes with 768MB heap and 256MB of PermGen memory configured as default. Depending on the number of applications and DataCollectors you want to use, you should increase these values so that the JVM does not run into memory issues.

The default memory configuration is sufficient to run the FA server with up to 5 DataCollectors used in your application(s). Then for every additional DataCollector we recommend to increase the maximum heap space (which is defined via the `-Xmx` argument) by 128 MB.

These value are configured in the Windows Service Configuration and can be adjusted by simply editing the values in the configuration file. To apply the changes a restart of the service is required.

Please note that if you use a 32bit version of the JVM the maximum size of the heap space is limited to ~1.5GB.

We highly recommend to use a 64bit environment for 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 Client. FADC uses

DCML

A domain specific XML language called DataCollector Markup Language which defines from which files data is

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).

imported and how this data is mapped to a relational database structure.

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)
- Disk Usage Settings (FADC)
- License Info (FADC)
- Logout (FADC)

DCML (FADC)

Applications (FADC)

Applications (FADC)

Overview

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

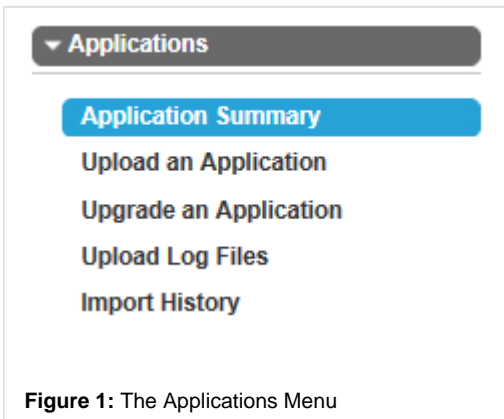


Figure 1: The Applications Menu

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 several sections:

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

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

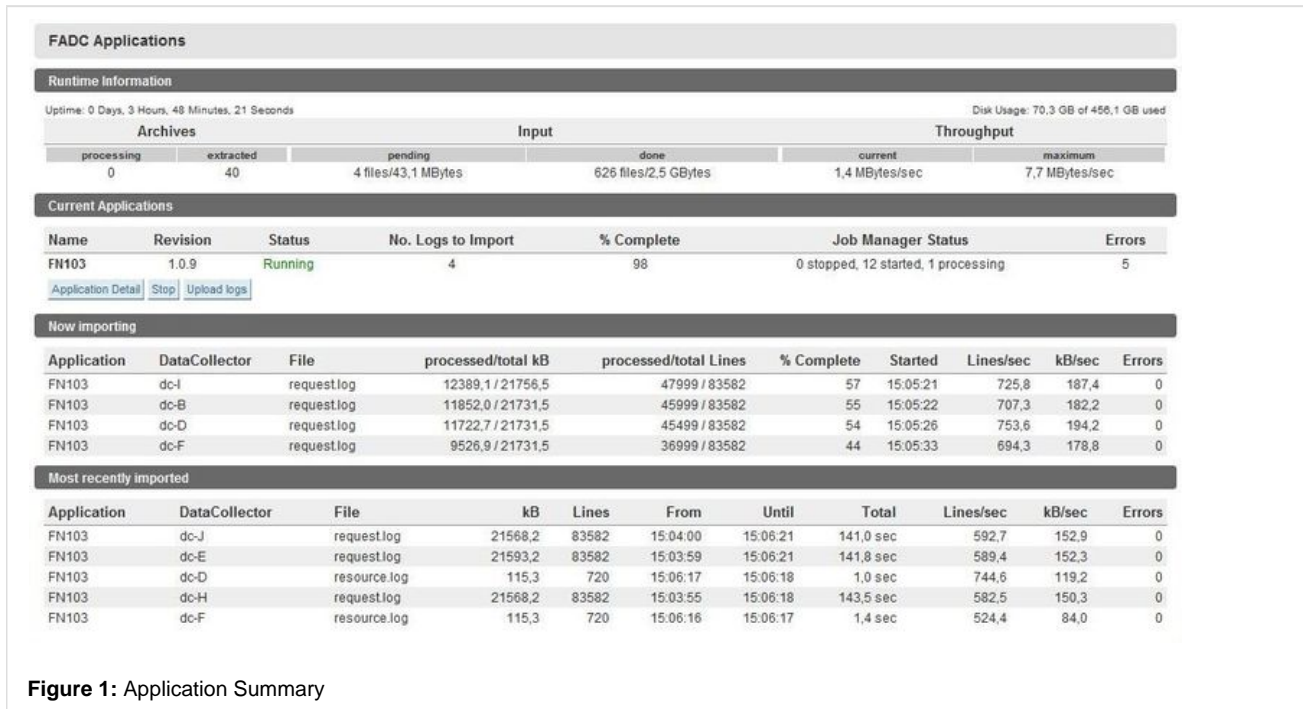


Figure 1: Application Summary

Runtime Information

This section details importing information and statistics such as uptime and disk usage, how many archives have been extracted, how many imported and the current and maximum throughput.

Current Applications

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

Data Sources

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

Data Sources

FusionAnalytics requires the data source details for the corresponding database where the information can be collected. You can fill in the details below and test if the connection is working successfully.

| | | |
|-------------------------------|---|--------------------------|
| Name | FADC | |
| Description | Analytics for FusionReactor (Collector Application) - created Tue Aug 30 11:04:56 CEST 2011 | |
| Database Type | Microsoft SQL | |
| Driver Class | com.microsoft.sqlserver.jdbc.SQLServerDriver | |
| URL | jdbc:sqlserver://localhost:1433;loginTimeout=5;databaseName=FADS | |
| User | sa | |
| Password | •• | (16 character limit) |
| Pool Size | 30 | (Concurrent connections) |
| Pool Timeout (Seconds) | 10 | (Seconds before timeout) |

[Modify](#) [Test Connection](#)

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 |

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.



Low Resolution (Quantized) data 12 month(s) ▾
 High Resolution (Detailed) data 1 week(s) ▾
 Modify Reset to default

Figure 1: Data size management

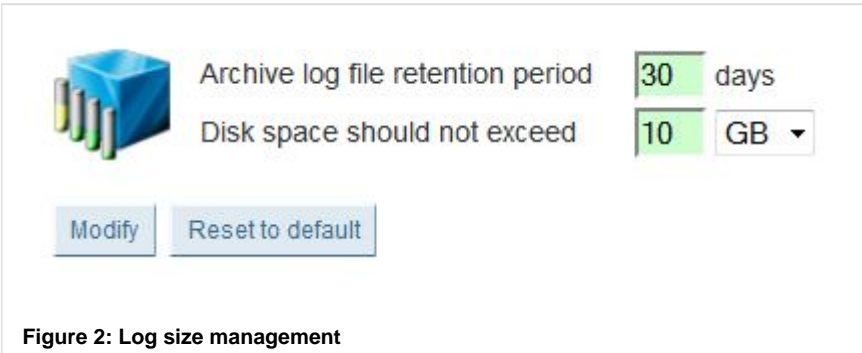
FusionAnalytics Processed Log File Management

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.

This configuration option is enabled per default. The default values are 30 days and 10GB as shown below.



Archive log file retention period 30 days
 Disk space should not exceed 10 GB ▾
 Modify Reset to default

Figure 2: Log size management

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.

| Name | Time Zone | Authentication | |
|-------------------|---------------|---|---------------------------------------|
| faDataCollector-1 | Europe/Berlin | <input checked="" type="checkbox"/> User: <input type="text" value="username"/> Password: <input type="password" value="••••••••"/> | <input type="button" value="Delete"/> |
| faDataCollector-2 | Europe/Berlin | <input type="checkbox"/> User: <input type="text"/> Password: <input type="password"/> | <input type="button" value="Delete"/> |

Fusion Analytics should delete the content from saved DataCollectors

Figure 3: DataCollectors

If you enable the 'FusionAnalytics should delete the content from saved DataCollectors' option data belonging to this DataCollector will be also deleted from the database.

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.

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.

Reset Application


This feature was introduced in version 2.0.7 of FusionAnalytics. The function provides the ability to "Reset" an application, i.e. remove all the data belonging to the application but without deleting the configuration of the application.

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. sections:

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



Application/myFR

Application Status

| Percentage Complete | Lines Processed | Total Lines |
|---------------------|-----------------|-------------|
| --- | --- | --- |

Data Input

| DataCollector | Status | Start All DataCollectors |
|-------------------|---------|--|
| faDataCollector-1 | Stopped | DataCollector Detail Start |

[http://127.0.0.1:9400/fadc/import/myFR/faDataCollector-1 \(click to copy to clipboard\)](#)

[Show/Hide DataCollectors](#)

Data Management

There are 12 available Jobs [Start all Jobs](#)[Show/Hide Jobs](#)

Database Details

| | |
|---------------|---------------------------------|
| Database Name | FADB |
| Database Type | Microsoft SQL Server 10.50.1600 |
| Database Size | 5486.75 MB |

Figure 1: Application Detail

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 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 having clicked on the [Show/Hide DataCollectors](#) button. In addition to the name of the DataCollector its import URL as used by Fusi [ataCollector Detail](#) button opens the DataCollector detail page.

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 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

A click on the [Job Detail](#) button opens a new page where you can edit the cron pattern that defines when the jobs executes.

Edit Job/Update Request Status

Update Request Status Configuration

Cron Pattern

[Submit changes](#)

Figure 2: Edit Cron Pattern of Job

Next Steps

- [Application Configuration \(FADC\)](#)
- [Upload an Application \(FADC\)](#)
- [Upgrade an Application \(FADC\)](#)

[UPLOAD AN APPLICATION \(FADC\)](#)

Introduction

This section is going to guide you through the process of uploading/creating an application step-by-step with the use of screenshots as well as with detailed explanation.

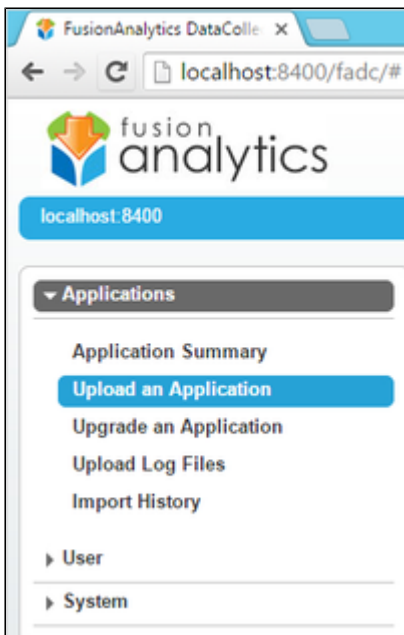
Getting Started

Step 1: Access the FADC application

First step is to access the FADC application via the Browser. Depending on the port that was configured during the installation process, you will be able to access the FADC application by typing **localhost:8400/fadc**. If you have configured the FADC application in another port, then you should use the configured port.

Step 2: Locate the Upload an Application option

In order to locate the "**Upload an Application**" option, navigate to **Applications > Upload an Application**.



Step 3: Upload an Application

When you click the "Upload an Application" option, a new screen is going to appear. See screenshot below.

- In the "Application Name" field choose a name of your preference. However, be aware that the name of the application is **case sensitive** and can **ONLY** use the options listed below.
 - - a - z
 - - A - Z
 - - 0 - 9
 - - dash (-)
 - - underscore (_)
 - - whitespace
- In the "DCML Application Package" field click the **Browse** button and navigate to the following directory: **C:\FusionAnalytics\Applications**. In the directory you should be able to find two zip folders, the **fusionreactor-fadc-2.0.6-DCAPP** and the **fusionreactor-fads-2.0.6-DSAPP**. For this application you need to choose the **fusionreactor-fadc-2.0.6-DCAPP** zip folder.
- Click the "Upload" button in order to upload the new application.

If you encounter any issues with uploading an application, please double check that the application name is using the options listed below.

Step 4: FusionAnalytics for FusionReactor Data Management Configuration

If the upload of the application was successful, a new screen is going to appear and is related with the Data Management configuration. See screenshot below.

In that step, you should configure your database details. When the database configuration is completed, click the **"Test Connection"** button in order to see if the connection was successful.

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 start to upgrade an application from one of the following 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.

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.

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

Figure 1: Select the application you wish to upgrade

Step 3: Starting the Upgrade

- Click Upgrade and FusionAnalytics will now make the necessary changes to your database and application configuration for the new version.

Please be patient during the upgrade process. The time taken for this process to complete will vary depending on the size of your database!

Step 4: Restart Server

Finally after the application upgrade has finished you have to restart the FusionAnalytics server. This completes the upgrade process.

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.

Upload log files to an Application

Upload a log or ZIP file

You can upload a single log file or a ZIP file containing log files. Zipped log files can remain in their original directory structure, so you can simply zip an existing directory with logs and then upload the zipped directory.

File **Upload** (300 MB max)

DataCollector faDataCollector-1 ▼

Submit

Figure 1: Manually uploading logs





IMPORT HISTORY (FADC)

Import History (FADC)

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

| File Name | Import Date | Import Status |
|-----------|-------------|---------------|
|-----------|-------------|---------------|

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

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.

Change Password

Current Password: [masked] Please enter the current password

New Password: [masked] Enter your new password

Confirm New Password: [masked] Confirm new password

[Update Password](#)

Figure 1: Changing password

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)

Disk Usage Settings (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).

Email Settings

These configuration settings will be used by default when sending email using FusionAnalytics DataCollector.

Email Server

From Address: fusionanalytics@company.com Notification emails will be sent from this email address.

To Address: user_account@company.com This email address will receive notification emails. Multiple recipients can be entered, separated by commas.

Mail Server: mail.company.com:25 Specify 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: 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**.

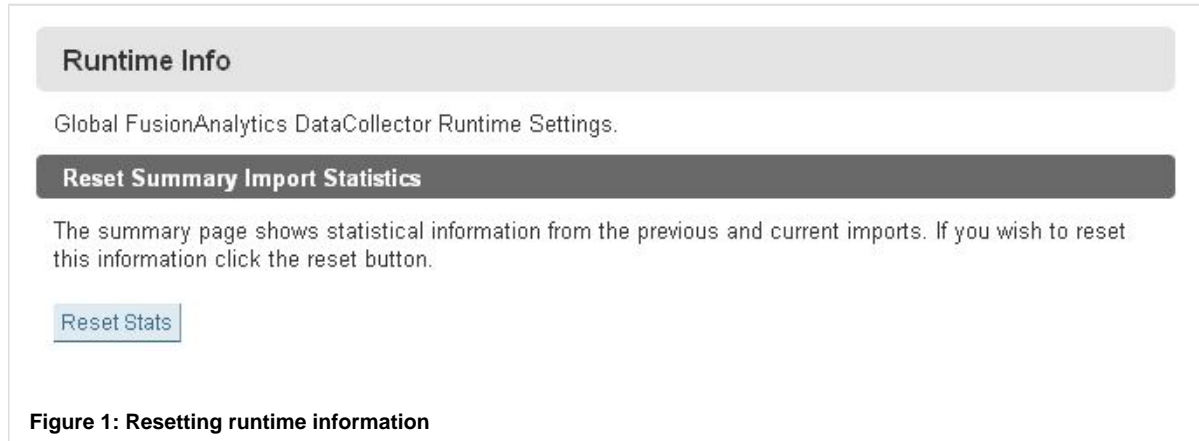
By default the port number is set to 25, but you can specify a port number after the address if needed; e.g. **127.0.0.1:587**.

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

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.



The screenshot shows a web interface for 'Runtime Info'. At the top, there is a header 'Runtime Info' in a grey bar. Below it, the text reads 'Global FusionAnalytics DataCollector Runtime Settings.' Underneath, there is a dark grey button labeled 'Reset Summary Import Statistics'. Below the button, a paragraph explains: 'The summary page shows statistical information from the previous and current imports. If you wish to reset this information click the reset button.' At the bottom of this section, there is a light blue button labeled 'Reset Stats'.

Figure 1: Resetting runtime information

DISK USAGE SETTINGS (FADC)

Disk Usage Settings

On this page you can configure if and how FusionAnalytics should handle low disk space situations. Per default the **Free Disk Space Threshold** is set to 20% of the total disk space. Changing this value does update the numbers in the description so that you can immediately see the impact of your change.

Please note that you must have configured email settings beforehand to be able to receive notification emails.

If you have enabled the **Reject Imports option** and there is not enough free disk space available new files added manually or sent by a FusionReactor instance will be rejected and deleted from the file system. A warning message for each rejected file will be written to the logs.

Disk Usage Settings

These configuration settings control how FusionAnalytics DataCollector handles low disk space situations when new data is uploaded or copied to an import directory.

Disk Usage

Total size of disk is 255.9 GB, 157.8 GB available, 106.6 GB remaining before FusionAnalytics DataCollector will report a low disk situation.

| | | |
|------------------------------------|-------------------------------------|--|
| Free Disk Space Threshold | <input type="text" value="20"/> % | If the free disk space becomes less then the given percentage value of the total disk space execute the actions selected below. |
| Send Email | <input checked="" type="checkbox"/> | Send a notification email to the addresses configured in the Email Settings if the available disk space became less than the configured value. |
| Notification Email Interval | <input type="text" value="1"/> | Minimum number of minutes between notification emails. Increase the interval if you do not want your mailbox filled with notification emails. |
| Reject Imports | <input type="checkbox"/> | Reject incoming files if the available disk space is less than the configured value (51.2 GB). |

[Update Settings](#)

LICENSE INFO (FADC)

License Info (FADC)

The license info page can be found in **System > License Info** and 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). | 31-Dec-2017 |
| 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 5 of 20 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.

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.

DCML (FADC)

DCML (FADC)

Overview

DataCollector Markup Language

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' />
```

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

The FADC DCML Database Documentation is available here

```

        <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='(?:)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.

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)

APML (FADS)

- APML Volume Switching
- Component
- Data Provider Programs
- Layout

Configuring Report Images

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 vary depending on whether it is started or stopped.

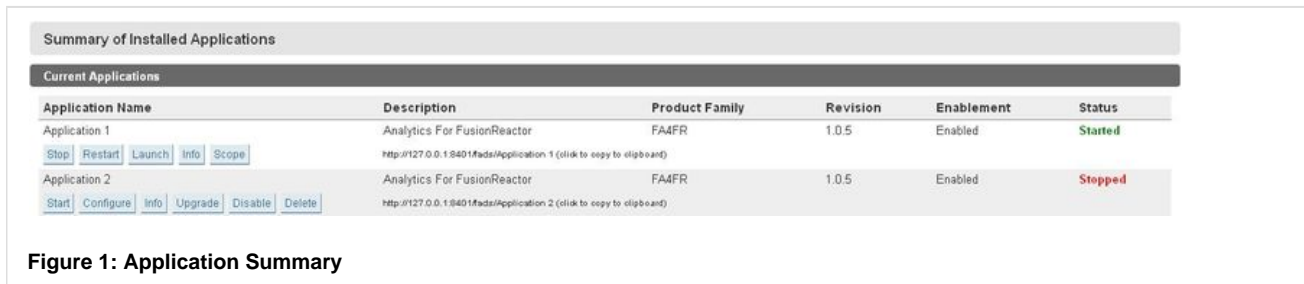


Figure 1: Application Summary

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 will be displayed.

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)

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 configured.

| Name | Description |
|---------|---|
| Stop | Stops the application. This will remove all scheduled jobs from the scheduler, preventing them from running until the application is restarted. |
| 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 status 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

| Name/Value Elements | | | |
|---------------------|-----------------------|--------|--------|
| OnUpgrade | EventProvider | Modify | Delete |
| DateFormat | DD.MM.YY JJ:NN:SS.TTT | Modify | Delete |
| OnStop | EventProvider | Modify | Delete |
| OnStart | EventProvider | Modify | Delete |
| APML | apml | Modify | Delete |
| | | Add | |

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 pre-configured for you and should not be changed.

It is worth stating that in version 2.0.7 of FusionAnalytics the "DateFormat" field was added as a way to help users to change the date format used in the various tables in the UI. For example, if you want to change the display date format from the European style to an American style then you can simply change the "DateFormat" field to be "MM.DD.YY JJ:NN:SS.TTT" and afterwards click the "Modify" button in order to save the changes. If the changes have been successfully saved, a green message will appear on the top of your screen. The message would look like the screen shot below.

Name/Value element modified successfully

Setting the "DateFormat" will only work correctly if the APML files (precisely the value of the "DateFormat" element in these files) of the application have not been manually altered.

Application Scheduled Tasks

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

| Application Scheduled Tasks | | |
|-----------------------------|-------------------|---------------------|
| Name | Schedule | |
| Daily Report | Daily at 04:30:00 | Activate Deactivate |
| T.A.P Report | Daily at 04:15:00 | Activate Deactivate |

Figure 2: Activating and Deactivating scheduled tasks.

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.

Data Sources

| | |
|-------------------------------|---|
| Name | frdb <small><i>This data source is required and may not be renamed or deleted.</i></small> |
| Description | This datasource must point to data generated by the Analytics for FusionReactor (Collector Application) |
| Database Type | Microsoft SQL ▾ |
| Driver Class | com.microsoft.sqlserver.jdbc.SQLServerDriver |
| URL | jdbc:sqlserver://127.0.0.1:1433;loginTimeout=5;databaseName=FADS |
| User | admin |
| Password | ●●●●●●●●●● |
| Pool Size | 40 |
| Pool Timeout (Seconds) | 30 |

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 automatically create all of the required tables and database objects. This will then allow you to launch your FADS application.

Application Detail (FADS)

Application Detail (FADS)

There are three sections to the application detail page:

- Summary

On this 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 |

Manifest Data

Data Dumps

| | | |
|---------------|------------------|---|
| 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)

Introduction

This section is going to guide you step-by-step through the process of uploading an application in FADS. Additionally, this documentation contains detailed explanation of the each step alongside informative screenshots.

Getting Started

Step 1: Access the FADS application

The first step is to access the FADS application via the browser. If you have kept the default configurations during the installation process of FusionAnalytics, then you should be able to access the FADS application by typing: **localhost:8400/fads**

Step 2: Locate the Upload an Application option

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



Step 3: Upload an application

When you click the "Upload an Application" option a new tab is going to appear. See screenshot below.

Upload APML Application

To add a new application choose an APML application package file, provide a unique deployment key for the new application and click on the Upload button. Different capitalization of keys is not considered unique. Deployment keys cannot contain the slash character (neither '/' nor '\'), the semicolon or the colon in any position; nor can they be named with a single dot ('.') or at ('@').

Upload Application Form

Application Deployment Key (must be unique)

APML Application Package

- In the "Application Deployment Key" choose a name of your preference. However, be aware that the Application name **cannot** contain **special characters** like (';', ':', '.', ':', '@') or any **slash characters** like ('\\', '/') or have the **same name as an existing application**.
- In the "APML Application Package" field click the **Browse** button and navigate to the following directory: **C:\FusionAnalytics\Applications**. In the directory you should be able to find two zip folders, the **fusionreactor-fadc-2.0.6-DCAPP** and the **fusionreactor-fads-2.0.6-DSAPP**. For this application you need to choose the **fusionreactor-fads-2.0.6-DCAPP** zip folder.
- Click the "**Upload**" button in order to upload the new application.

If you encounter any issues with uploading an application, please double check that the application name is using the options listed below.

Application Configuration (FADS)

If the application has been successfully updated, a new tab is going to appear and is related with the Application configuration.

For more information about the application configuration, please check out the link, [Application Configuration \(FADS\)](#).

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

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 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.

Upgrade APML Application

To upgrade an application choose the APML application package file that contains the upgraded application and then click the upgrade button

Upgrade Application Form

Choose the application to upgrade Please stop the Application 2 before upgrade ▾

APML Application Package

Browse...

Upgrade

Figure 1: Select the application you wish to upgrade

Back to the top

Step 3: Upgrading the application

Next press "Upgrade" and you will be taken to the application configuration screen in order to configure the upgraded application. See [Ap1](#) more information on configuring applications.

Analytics For FusionReactor

Instance: 1

Name/Value Elements

| Name | Value | Action |
|-----------|---------------|---------------|
| OnUpgrade | EventProvider | Modify Delete |
| OnStop | EventProvider | Modify Delete |
| OnStart | EventProvider | Modify Delete |
| APML | apml | Modify Delete |

Application Scheduled Tasks

| Name | Schedule | Action |
|--------------|------------------------|---------------------|
| Daily Report | Daily at 10:10:00 AM | Activate Deactivate |
| T.A.P Report | Monthly at 10:20:00 AM | Activate Deactivate |

Data Sources

Name: fdb
Description: This datasource must point to data generated by the Analytics for FusionReactor (Collector Application)

Database Type: Microsoft SQL

Driver Class: com.microsoft.sqlserver.jdbc.SQLServerDriver

URL: jdbc:sqlserver://192.0.0.143;loginTimeout=5;databaseName=RC6-3

User: sa

Password: *****

Pool Size: 50

Pool Timeout (Seconds): 50

Modify Test Connection

Name: _____

Description: _____

Database Type: Select DB

Driver Class: _____

URL: _____

User: _____

Password: _____

Pool Size: _____

Pool Timeout (Seconds): _____

Add Test Connection

Figure 2: Application Configuration page

Step 4: Restart Server

Finally after the application upgrade has finished you have to restart the FusionAnalytics server. This completes the upgrade process.

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

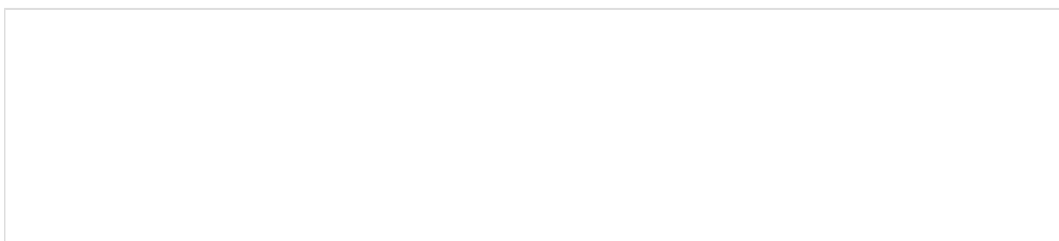
This section details creating, updating and mapping users as well as explaining the FusionAnalytics DataServices users interface.

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:



Add A New User

Username

First Name

Last Name

Password

Password (Confirm)

Email

[Add New User](#)

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:

Existing Users

You may only delete users who have no mappings, or only mappings to application for which you are an Administrator.

| Username (Email) | Name | System User | Source | Last Login | Mappings |
|--|---------------|-------------|--------------------|-------------------------|---------------|
| admin (sysadmin@company.com) Current User | System Admin | Yes | FADS Local | 2011-08-26 09:33:23.347 | [System User] |
| user () | System User | No | FADS Administrator | Never | 1 mappings |
| impJohn (smith.J@company.com) | John Smith | No | FADS Administrator | Never | No mappings |
| impClaire (jonesC@company.com) | Claire Jones | No | FADS Administrator | Never | No mappings |
| impDavid (edwardsD@company.com) | David Edwards | No | FADS Administrator | Never | No mappings |

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:



Available Application Role Mappings

You may only add roles for applications for which you are an administrator.

| Application | Role | Actions |
|---------------|--------------------|----------|
| Application 1 | None (No Access) ▾ | Add Role |
| Application 2 | None (No Access) ▾ | Add Role |

Figure 1: Available Mappings

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

Available Application Role Mappings

You may only add roles for applications for which you are an administrator.

| Application | Role | Actions |
|---------------|--------------------|----------|
| Application 1 | Administrator ▾ | Add Role |
| Application 2 | None (No Access) ▾ | 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 wish to edit a mapping for you should see the following screen:

Current Application Role Mappings

You may only modify roles for applications for which you are an administrator.

| Application | Role | Actions |
|---------------|-----------------|---------------|
| Application 1 | Administrator ▾ | Modify Delete |

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

Create A New User Scheduled Task

| | |
|---------------------------|--|
| Task Name | <input type="text"/> |
| Active | <input checked="" type="checkbox"/> |
| Duration | Start Date <input type="text"/> End Date (optional) <input type="text"/> |
| Frequency | <input checked="" type="radio"/> One-Time at <input type="text"/> |
| | <input type="radio"/> Recurring <input type="text" value="Daily"/> at <input type="text"/> |
| | <input type="radio"/> Daily every Hours <input type="text"/> Minutes <input type="text"/> Seconds <input type="text"/> Start Time <input type="text"/> End Time <input type="text"/> |
| Provider Name | <input type="text" value="AnnotationADGProvider"/> |
| Provider Arguments | Name <input type="text"/> Value <input type="text"/> <input type="button" value="−"/> |
| | <input type="button" value="+"/> Add New Argument |
| | <input type="button" value="Add New Task"/> <input type="button" value="Cancel"/> |

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.

Create A New User Scheduled Task

Task Name

you must provide a task name

Active

Duration

Start Date

End Date (optional)

year must be within range current year - 2099
month must be within range 1 - 12
day must be within range 1 - 31

Frequency

One-Time at

minute must be within range 0 - 59
hour must be within range 0 - 23

Recurring Daily ▼ at

Daily every

Hours

Minutes

Seconds

Start Time

End Time

Provider Name

Provider Arguments

Name

Value

✖

+ Add New Argument

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

Editing Application Scheduled Task

Task Name Daily Report

Active

Duration Start Date 01/01/2011 End Date (optional)

Frequency

One-Time at

Recurring Daily at 04:30 PM

Daily every Hours Minutes Seconds
Start Time End Time

Provider Name DailyStatus

Provider Arguments Name Value

Add New Argument

Update Task Cancel

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 its 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



| Name | Schedule | Activate | Deactivate |
|--------------|---------------------|----------|------------|
| Daily Report | Daily at 4:30:00 PM | Activate | Deactivate |
| T.A.P Report | Daily at 4:15:00 PM | Activate | Deactivate |

Figure 2: Application Scheduled Tasks

- visiting the Application Tasks Summary page



| Task Name |
|---|
| Daily Report |
| Edit Run Now Deactivate |
| T.A.P Report |
| Edit Run Now Deactivate |

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 started and then you terminate 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 have two 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:

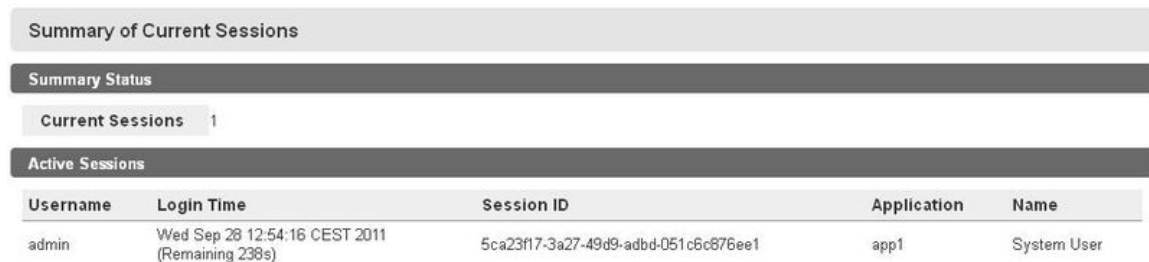


Figure 1: Session Summary

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, I User's Name



The screenshot shows a web interface with a header 'Summary of Current Sessions' and a sub-header 'Summary Status'. Below the sub-header, there is a section 'Current Sessions' with a count of '1'. Underneath is a table titled 'Active Sessions' with the following data:

| Username | Login Time | Session ID | Application | Name |
|----------|---|--------------------------------------|-------------|-------------|
| admin | Wed Sep 28 12:54:16 CEST 2011 (Remaining 238s) | 5ca23f17-3a27-49d9-adbd-051c6c876ee1 | app1 | System User |

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.



The screenshot shows a web interface with a header 'Application Name' and a sub-header 'Description'. Below the sub-header, the text reads 'app1' and 'Analytics For FusionReactor'. Below the text, there are buttons for 'Stop', 'Restart', 'Launch', 'Info', and 'Scope'. Below the buttons, the URL 'http://127.0.0.1:6666/fads/app1 (click to copy to clipboard)' is displayed.

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

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 contained components. E.g., the Cache Controller which is responsible for intelligently storing and retrieving large amounts of data very quickly. A menu item from the "System" section in the main Table of Contents will display all the controllers available within DataServices. Some of them allow you to perform some action. These actions vary from clearing records, dumping the state to the screen and also stopping/starting in certain circumstances, you shouldn't need to change anything in the Controllers page; however you may be directed to obtain information from the support request.

| All Controllers | | |
|--|--|---|
| Name | Implementation | Status |
| Data Persistence Controller | com.intergral.fads.impl.persistence.PersistenceControllerImpl | Ready |
| State Management Controller | com.intergral.fads.impl.state.StateControllerImpl | Ready Dump 4 storage element(s). |
| Universal Configuration - Adapter Controller | com.intergral.fads.impl.uc.UniversalConfigurationAdapterControllerImpl | Ready 4 configuration(s) loaded. |
| External Object Cache Controller | com.intergral.fads.impl.eoc.ExternalObjectCacheControllerImpl | Ready Dump Clear Cache Status: STARTED |
| Cache Controller | com.intergral.fads.impl.cache.CacheControllerImpl | Ready Stop Dump Clear 0 cache(s) in memory. |
| Session Controller | com.intergral.fads.impl.session.SessionControllerImpl | Ready 2 session(s) active. |
| Analytics Controller | com.intergral.fads.impl.ac.AnalyticsControllerImpl | Ready |
| Scheduled Tasks Controller | com.intergral.fads.impl.quartz.ScheduledTaskControllerImpl | Ready |
| Credential and Security Controller | com.intergral.fads.impl.credential.CredentialControllerImpl | Ready |
| Application Management Controller | com.intergral.fads.impl.app.ApplicationControllerImpl | Ready Dump 2 application(s) loaded. |
| Blaze Communications Controller | com.intergral.fads.impl.blaze.BlazeControllerImpl | Ready |
| Google Visualizations Data Controller | com.intergral.fads.impl.gvdc.GVDataControllerImpl | Ready |
| Asynchronous Job Controller | com.intergral.fads.impl.async.AsynchronousJobControllerImpl | Ready Clear Pool Size: 10, Active: 0, HWM: 10, Tasks Queued: 0, Aggressive Streamed: 57, OK: 56, Can: 1, Fail: 0 (Reported: 57) Executor Refreshes: 1 |
| Utility Controller | com.intergral.fads.impl.utils.UtilityControllerImpl | Ready |

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, it 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.

When DataServices receives the job, it queues it, executing it when it has available resources. With many jobs being queued and executed, it is 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 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. 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.integral.fads.async.jobs.ExecuteDataProviderJob com.integral.fads.async.receivers.ExecuteCallbackReceiver | [app]=[RC8 (Analytics For FusionReactor)], component=com.integral.fads.model.component.DefaultServerComponent@569dcb58, dp=null, perspective=[Perspective] Overview (Server Summary), provider=Provider: FRSystemMetricsProvider apmi/FRSystemMetricsProvider.groovy] |
| 28 Sep. 11 11:26:27.811 11:26:27.834 11:26:27.875 | OK | 23 | 41 | 64 | admin | com.integral.fads.async.jobs.ExecuteDataProviderJob com.integral.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.integral.fads.async.jobs.ExecuteDataProviderJob com.integral.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.integral.fads.async.jobs.ExecuteDataProviderJob com.integral.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.integral.fads.async.jobs.ExecuteDataProviderJob com.integral.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.integral.fads.async.jobs.ExecuteDataProviderJob com.integral.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 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 where notification emails are sent from and to as well as adding a mail server to process these requests. (This is critical configuration for C

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: This is the hostname from where this application is served from. This field is required.

HTTP Server Port: 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: This is the hostname from where this application is served from for https (secure) connections. This field is required.

HTTPS Server Port: 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: Specify 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

LICENSE INFO (FADS)

License Info (FADS)

The license info page can be found in **System > License Info** and 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). | 31-Dec-2017 |
| 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 5 of 20 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](#)

APML (FADS)

APML (FADS)

Overview

This document introduces **APML**, the **Analytical Program Markup Language**, a system for describing analytical applications.

APML is an XML grammar for describing software which performs analytical functions. APML, and the auxiliary files which comprise an APML Application (also known as a DataServices Application) are loaded by FusionAnalytics to answer demands for data from the FusionAnalytics client software (either web or AIR-based).

APML is used to describe the three main components of a FusionAnalytics application:

- **Perspectives**, which are made up of one or more
- **Views**, whose data is supplied by
- **Data Providers**

Perspectives are visible in the client as separate windows. The views are contained within them, and the calls to the individual Data Provider client framework.

In Figure 1, the CPU Usage perspective is visible. The perspective consists of two views:

- A CPU Line Chart showing two peaks
- A timeline controller

The timeline is simply another view, and all that is required to 'wire up' these components is that the timeline be present in the perspective. The client framework will take care of handling the internal message bus, ensuring that all interested parties are scrolled in sync with the ti

Here, for example, is the APML used to construct the perspective above:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>

<Module xmlns="http://www.intergral.com/FADS/v1/APML">
  <Perspectives>
    <Perspective id="CPUUsage">
      <Name>CPU Usage</Name>
      <Description>Shows details of all resource usage</Description>
      <Layout type="com.intergral.fusionanalytics.view.layout.layoutM
SimpleFlowLayoutManager">
        <VBox>
          <ViewRef ref="CPULC" title="CPU"/>
          <ViewRef ref="Global.TimeLine" height="145"/>
        </VBox>
      </Layout>
    </Perspective>
  </Perspectives>

  <Views>
    <View id="CPULC">
      <Name>CPU Usage</Name>
      <Description>Process and System CPU Usage</Description>
      <Component ui="com.intergral.fads.model.renderer.ChartBaseVisua
        <DataProviderRef ref="CPULC"/>
        <Argument name="Title">Process and System CPU Usage</Argumen
        <Argument name="Background.Color">0xffffffff</Argument>
        ...further arguments omitted for brevity...
      </Component>
    </View>
  </Views>

  <DataProviders>
    <DataProvider id="CPULC">
      <Description>CPULC</Description>
      <Arguments>
        <Argument name="server" bindingname="Server.FilterBy.server"
        <Argument name="start" bindingname="Global.TimeLine.start"/>
        ...further arguments omitted for brevity...
      </Arguments>
      <Provider>
        <Program language="groovy" return="response" program="apml
/SeriesDataProvider.groovy"/>
      </Provider>
    </DataProvider>
  </DataProviders>
</Module>
```

Listing 1: APML for the CPU Usage Perspective

The APML, application manifest and Data Provider samples are all taken from our Analytics for FusionReactor DataServices application. This application ships with FusionAnalytics (it is optionally preinstalled by the installer). The zip file containing the application is installed in the FusionReactor applications folder, and is called fusionreactor-fads-1.0.0-DSAPP.zip. Take a look at the files yourself.

Next Steps

Data Provider Programs

Learn about how APML uses a Data Provider to specify how data is retrieved from the database

Layout

Learn about using the Layout tag.

Component

Learn more about the View Component and the arguments used by it.

APML Volume Switching

Customize the component to your needs using switch/case/default structuring.

APML VOLUME SWITCHING

APML Volume Switching

Within an APML **View** tag, at least one **Component** must be specified. The Component is the visualization used in the Client to display the data points returned from the DPP.

In order that components may be switched based on data volumes, a simple **Switch/Case/Default** structure is supported.

Listing 1 shows a snippet of APML used in the Analytics for FusionReactor "DataFinder" perspective. A different component may be selected based on the result of the execution of a Volume Provider (VP), in this case GPRequestProviderVP. The Volume Provider is a DPP, named in the SwitchConditionProvider tag, which simply returns a single integer representing the number of data points which would be selected if the corresponding DPP were run.

```
<View id="DataFinder">
  <Name>DataFinder</Name>
  <Description>DataFinder<
/Description>
  <Switch>
    <SwitchConditionProvider ref="
GPRequestProviderVP" />

    <Case to="86400000" resolution="
HOURL">
      <Component ui="com.intergral.
fusionanalytics.view.itemRenderers.
visualisations.TimeBlockChartBitmap">
        <DataProviderRef ref="
PagedRequestDataFinderProvider" />
        <Argument name="Title"
>DataFinder</Argument>
        ... further arguments
omitted for brevity ...
      </Component>
    </Case>
```

```

    <Case to="2592000000"
resolution="DAY">
    <Component ui="com.intergral.
fusionanalytics.view.itemRenderers.
visualisations.TimeBlockChartBitmap">
    <DataProviderRef ref="
PagedRequestDataFinderProvider"/>
    <Argument name="Title"
>DataFinder (Day)</Argument>
    ... further arguments
omitted for brevity ...
    </Component>
</Case>

    <Default resolution="YEAR">
    <Component ui="com.intergral.
fusionanalytics.view.itemRenderers.
visualisations.TimeBlockChartBitmap">
    <DataProviderRef ref="
PagedRequestDataFinderProvider"/>
    <Argument name="Title"
>DataFinder (Year)</Argument>
    ... further arguments
omitted for brevity ...
    </Component>
</Default>
</Switch>
</View>

```

Listing 1: View with Switch/Case/Default Structure

The switch processor then selects a component based on that result. In this way, FusionAnalytics is able to select a visualization component based on the volume of data being rendered.

In Listing 1, you can see the same **DataProvider** is being run (**PagedRequestDataFinderProvider**), but the arguments differ. This allows the component to be customized based on the volume of data selected for display.

The DPP Provider tag may also be switched in a similar way. The Component will remain unswitched, but the Provider tag is placed within a Switch/Case/Default structure, and a Volume Provider is run prior to the switch.


Although the component will not change, the provider (and the arguments supplied within the Provider tag) can be different, customized to the volume of data required for rendering.

COMPONENT

Component

Within an APML **View** tag, at least one **Component** must be specified. The Component is the visualization used in the Client to display the data points returned from the DPP. There are many visualization components supplied within the client.

The component used by the CPU Usage perspective given in the example, is listed below:



```

<Component ui="com.intergral.fads.model.renderer.ChartBaseVisualization"
  <DataProviderRef ref="CPULC"/>
  <Argument name="Title">Process and System CPU Usage</Argument>
  <Argument name="Background.Color">0xffffffff</Argument>
  <Argument name="Chart.Maximum">100</Argument>
  <Argument name="start" bindingname="Global.TimeLine.start"/>
  <Argument name="end" bindingname="Global.TimeLine.end"/>
  <Argument name="timezone" bindingname="Global.TimeLine.timezone"/>
  <Argument name="Series.Count">3</Argument>
  <Argument name="Series.1.ID">qseries.resource.system_cpu_pct</Argument>
  <Argument name="Series.1.Type">AreaSeries</Argument>
  <Argument name="Series.1.Label">System CPU Usage (%)</Argument>
  <Argument name="Series.1.Relation">qseries.resource.system_cpu_pct</Arg
  <Argument name="Series.1.X">DT</Argument>
  <Argument name="Series.1.Y">system_cpu_pct</Argument>
  <Argument name="Series.1.Color">0xaaaaaa</Argument>
  <Argument name="Series.1.Alpha">0.5</Argument>
  <Argument name="Series.1.Visible">>true</Argument>
  <Argument name="Series.2.ID">qseries.resource.process_cpu_pct</Argument>
  <Argument name="Series.2.Type">AreaSeries</Argument>
  <Argument name="Series.2.Label">Process CPU Usage (%)</Argument>
  <Argument name="Series.2.Relation">qseries.resource.process_cpu_pct</Ar
  <Argument name="Series.2.X">DT</Argument>
  <Argument name="Series.2.Y">process_cpu_pct</Argument>
  <Argument name="Series.2.Color">0x006b91</Argument>
  <Argument name="Series.2.Alpha">0.5</Argument>
  <Argument name="Series.2.Visible">>true</Argument>

  <Argument name="Series.Selected" bindingname="CPULC.Series.Selected"/>
  <Argument name="Timeline.Cursor">true</Argument>
  <Argument name="Timeline.Cursor.Color">0xFF0000</Argument>
  <Argument name="Timeline.Cursor.Weight">1</Argument>
  <Argument name="Timeline.Cursor.Alpha">0.9</Argument>
</Component>

```

Visualizations

In the declaration of the component, you may have noticed the following **ui** attribute:

```

<Component ui="com.intergral.fads.model.renderer.ChartBaseVisualization"

```

This attribute is used to specify how the data will be rendered to the screen. In the example, **ChartBaseVisualization** is used, to create a

Figure 1: ChartBaseVisualization Area Graph

Other types of graph can be created using **ChartBaseVisualization**, by specifying a **Type** argument within the component, like so:

```
<Argument name="Series.1.Type">AreaSeries</Argument>
```

Some of the type arguments provided by the client are:

- AreaSeries - creates and Area Graph like the one in Figure1.
- LineSeries - creates a Line graph
- BarSeries - creates a Bar Chart

Other Visualisations:

AdvancedDataGridRenderer can be used to render the data in a table such as this:

Figure 2: A visualisation of a Data Grid

PieChartRenderer can be used to render data in a Pie Chart like this:

Figure 3: A visualisation of a Pie Chart

TimeBlockChartBitmap can be used to create a TimeLine chart, which shows each data point as a block, at its respective point in time. Here is a TimeLine that shows each of the requests that happened within an hour period:

Figure 4: A visualisation of a TimeLine

Arguments

Within the component there are several Arguments. These are used to define different properties for how the data is displayed. We will now cover some of the arguments used to create the **ChartBaseVisualization** and **AdvancedDataGridRenderer**, and their purpose. Only the essential arguments will be covered, the more advanced arguments will be omitted.

Where arguments are defined for a specific series (i.e. Series.1.Type), the number defines which series the argument relates to. The same arguments can be included for each additional series that is added to the graph, as long as the series number is incremented and

| Argument | Purpose |
|--|---|
| <code><DataProviderRef ref="CPULC" /></code> | This provides a reference to the Data Provider which gets the data |
| <code><Argument name="Title">Process and System CPU Usage</Argument></code> | Sets the title that is used for the View, and is seen in the grey bar |
| <code><Argument name="Background.Color">0xffffffff</Argument></code> | Sets the background colour of the graph using a hex colour value to see the data when the background is white. |
| <code><Argument name="Chart.Maximum">100</Argument></code> | Only for ChartBaseVisualization to set the max height of the Y |
| <code><Argument name="start" bindingname="Global.TimeLine.start" /></code> | This is a required argument, used to link the view to the TimeLine |
| <code><Argument name="end" bindingname="Global.TimeLine.end" /></code> | This is a required argument, used to link the view to the TimeLine |
| <code><Argument name="timezone" bindingname="Global.TimeLine.timezone" /></code> | This is a required argument, used to link the view to the TimeLine |
| <code><Argument name="Series.Count">3</Argument></code> | Only for ChartBaseVisualization . Sets how many series you want. For example, there are 2 series; Process Memory and System Memory |
| <code><Argument name="Series.1.Type">AreaSeries</Argument></code> | Only for ChartBaseVisualization to set the type of graph that you want |
| <code><Argument name="Series.1.Label">System CPU Usage (%)</Argument></code> | Only for ChartBaseVisualization . Gives a label to the series with the data |
| <code><Argument name="Series.1.Relation">qseries.resource.system_cpu_pct</Argument></code> | Only for ChartBaseVisualization . Sets the relation from where the data is returned by the data provider |
| <code><Argument name="Series.1.X">DT</Argument></code> | Only for ChartBaseVisualization . Sets the units that will be used for the X-axis |
| <code><Argument name="Series.1.Y">system_cpu_pct</Argument></code> | Only for ChartBaseVisualization . Sets the units that will be used for the Y-axis |
| <code><Argument name="Series.1.Color">0xaaaaaa</Argument></code> | Only for ChartBaseVisualization . Sets the colour that will be used for the series |
| <code><Argument name="Series.1.Alpha">0.5</Argument></code> | Only for ChartBaseVisualization . Sets the transparency of the series |
| <code><Argument name="Series.Selected" bindingname="CPULC.Series.Selected" /></code> | Only for ChartBaseVisualization . Defines a Series Selector driver to turn on/off a series while viewing the graph. |
| <code><Argument name="Columns.Count">4</Argument></code> | Only for AdvancedDataGridRenderer . Sets how many columns to display |
| <code><Argument name="Columns.1.Field">STARTED</Argument></code> | Only for AdvancedDataGridRenderer . Sets the Data you want to display in the first column that is returned by the Data Provider. |
| <code><Argument name="Columns.1.Label">Started</Argument></code> | Only for AdvancedDataGridRenderer . Sets a label for the column to be displayed. |
| <code><Argument name="Columns.1.Type">Date</Argument></code> | Only for AdvancedDataGridRenderer . Sets the dataType for the column. |
| <code><Argument name="Columns.1.Format">DD.MM.YY JJ:NN:SS.TTT</Argument></code> | Only for AdvancedDataGridRenderer . If you want to display the data in a specific format. Otherwise, it is not needed. |

| | |
|---|---|
| <code><Argument name="Columns.1.Align">left</Argument></code> | Only for AdvancedDataGridRenderer . Set how you want the d |
| <code><Argument name="Columns.1.Width">20</Argument></code> | Only for AdvancedDataGridRenderer . Set the width of the col is allocated on the page. |

DATA PROVIDER PROGRAMS

Data Provider Programs

You may have noticed in our example, that the `DataProvider` used in the example (`CPUMemLCProvider`) does not actually specify how data is actually retrieved from the database.

That is the work of the `DataProvider Program` or `DPP`.

It is possible to use an XML `<![CDATA[xyz]]>` construct to actually place the program inline in the XML text; however this is not recommended and we will not discuss it here.

Placing the program in a separate file enables it to be reused, tracked in source control, and edited independently of the `APML`.

The program is specified as follows:

```
<Program language="groovy" return="
response" program="apml/CPUMemLCProvider.
groovy" />
```

The tag above specifies the language, return variable and program source for the provider.

Supported Languages

`DataProvider Programs (DPPs)` may be supplied in two Java-like languages:

- Groovy: <http://groovy.codehaus.org/>
language="groovy"
- BeanShell: <http://www.beanshell.org/>
language="bsh"

These languages are supported in `FusionAnalytics` using abstractions called **Execution Engines**. These engines take care of injecting variables into the program's environment, invoking the relevant language compiler or interpreter, and extracting the results.

We recommend the Groovy language, since this is a very full-featured and mature environment and lends itself very well to the development of `DPPs`.

Return Variable

The `return="response"` attribute in the example `Program` tag specifies the variable which the execution engine will use when the program completes, in this case the variable called **response**.

Program Source Code

Finally, the `program="apml/CPUMemLCProvider.groovy"` attribute instructs the execution engine from where the source code for the program can be obtained; in this case in the folder **apml**, in a file called **CPUMemLCProvider.groovy**.

The contents of the groovy source file are presented in Figure 1. We are using a **ResourceQueryBuilder** object to construct the SQL being run against the database.

On This Page:

Data Provider Programs
Supported Languages
Return Variable
Parameter Binding

```

void go()
{
    Util.recordStartAndEnd(session,
argument);
    response = ResourceQueryBuilder.select
(nature, argument, datafactory, 10000)
    {
        def resolution = Resolution.
getResolution(argument.get("start"),
argument.get("end"));
        it.column("r.DT as DT");
        if (resolution == 0)
        {
            it.column("r.PROCESS_CPU_PCT as
PROCESS_CPU_PCT");
            it.column("r.SYSTEM_CPU_PCT as
SYSTEM_CPU_PCT");
            it.from("RESOURCE r");
        }
        else
        {
            it.column("r.AVG_PROCESS_CPU_PCT
as PROCESS_CPU_PCT");
            it.column("r.AVG_SYSTEM_CPU_PCT as
SYSTEM_CPU_PCT");
            it.from("RESOURCE_Q r");
            it.where("r.
RESOLUTION=${resolution}");
        }
        it.addArgumentBasedFilters(argument.
get("start"), argument.get("end"));
        it.orderBy("r.DT");
    }
}
go();

```

Listing 1: Data Provider Written in Groovy

DataProvider Programs may return several object types: in this case we are returning the result of the ResourceQueryBuilder, which is a Groovy SQL result set.

Because Analytics for FusionReactor supports quantization, a **Resolution** object has been written :

```

Resolution.getResolution(argument.get
("start"), argument.get("end"));

```

This Resolution object automatically selects the correct resolution of data based on the start and end (which comes from the timeline control) times.

DPP Variable Injection

Several arguments are injected into the Groovy (and Beanshell) environment before execution.

We will not enumerate them completely here, but a few are:

Parameter Binding

You may have noticed in Figure 2 the following DataProvider arguments:

```
<Argument name="start" bindto="Global.
TimeLine.start" bindingname="dp.start" />
<Argument name="end" bindto="Global.
TimeLine.end" bindingname="dp.end" />
```

These arguments specify how data value changes in the client are passed into the DPP. The **name** attribute specifies the name of the injected variable, the **bindto** attribute specifies the event used as the source of the argument, and the **bindingname** specifies the name of this argument binding in the client. These variables appear in the argument variable scope (see table above).

The client software manages an event model comprising event sources and bindings, and calls interested Components appropriately when the bound variables change. In most cases, the Component will pass on the request to the DataProvider1. Any Component can raise such events, and any Component may listen for them using bindings, provided the name of the event is known.

Component

The Component may also specify bound variables in the same way. These declare an interest in the bound variable at the component level. When the bound variable changes, the Component will be notified, and can take whatever action it deems appropriate.

LAYOUT

Layout

The Layout tag, within APML, is used to define how the different elements of the perspective will be laid out on the page. The FusionAnalytics Client software makes a number of different layout engines available, which may be embedded. This allows quite cor

Here is an example of some of the Layout tags within APML:

```
<Layout type="com.intergral.fusionanalytics.view.layout.layoutManager.Si
  <VBox>
    <Layout type="com.intergral.fusionanalytics.view.layout.layoutManag
      <ViewRef ref="CPULC" title="CPU" />
      <ViewRef ref="ResourceMemoryAC" title="Process Memory Usage" />
    </Layout>
    <ViewRef ref="Global.TimeLine" height="145" />
  </VBox>
</Layout>
```

Listing 1: Layout tag within the Perspective.

We can see two layout managers in use here:

1. **SimpleFlowLayoutManager** which, in this case, uses a **VBox** (a box which lays its contents out vertically, expanding horizontally)
2. **IconSelectorManager** which is being used to toggle two further views, Process and System CPU Usage and Process Memory U:

The final view, the **Timeline**, belongs to the **SimpleFlowLayoutManager's VBox**, so it will appear beneath the view laid out by the **IconSelectorManager**.

The layout tags used in listing 1 construct the following perspective

Figure 1: Perspective constructed from Listing 1

More Layout Tags

- **HorizontalGridLayoutManager** - which lays its contents out in a grid, optionally allowing them to extend vertically beyond their e:
- **ComboBoxSelectorLayoutManager** - used to toggle further views, similar to IconSelectorManager, using a drop down combo b
- **HBox** - a box which lays its contents out horizontally, expanding vertically to fit its extents

Configuring Report Images

Configuration Report Images

When viewing FusionAnalytics reports which have been delivered via email, you may see a request to log in to your FusionAnalytics Data Services server.

This happens because the images which comprise part of the reports are served directly from the FADS system itself.

If you are seeing a request to log in, you have not obtained a cookie for the FADS server on the address from where images are being served.

There are three solutions to this problem.

Log in

When the login dialog appears, log in as a FADS user who is able to view the images.

Adjust the security settings of the application

You can set your FADS application not to require authentication to view images.

Warning: If you disable this protection, anyone will be able to view these report images, regardless of whether they are a valid FADS user or not.

1. Navigate on your server to FusionAnalytics\Server\webapps\fads\WEB-INF\classes\Applications\YOUR_APP_NAME and open application.xml with a text editor. This editor must be capable of reading and writing UTF-8 Unix files. Windows Notepad is not able to do this - we recommend Notepad++.
2. In this file, locate the ApplicationWebServer block and carefully remove the Path whose pathattribute is /_Persistent.

You should be left with:

```
<ApplicationWebServer>
  <PathRules>
    <Path path="/configuration">
      <Protected/>
    </Path>
  </PathRules>
</ApplicationWebServer>
```

3. Restart the application by logging in to the Data Services webapp (usually located on port 8400), and restarting the data service itself using the Restart button.

Caveat: If the application is upgraded or reinstalled, this procedure must be repeated.

Change the domain name used to serve image links

The problem can also be solved by changing the host name used to serve image links. If you ordinarily log in to the FusionAnalytics Data Service webapp using - for example - http://fred.foo.com:8400/, then you should already have an authentication cookie for this domain. If the system has several aliases, images may be being served from a different domain name and the login dialog will appear. If this is the case, the domain name used in links can be changed:

1. Log in to the Data Services webapp (usually located on port 8400). Open the **System** section and click on the **Server Settings** e ntry.

2. Under the settings **HTTP Server Host** and **HTTPS Server Host**, change the value to the domain name you usually use.
3. Click **Update Server Settings**
 - This will cause generated links to use the new domain name, for which you already have an authentication cookie, and the login dialog should no longer appear.

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.

On this page:

How FusionAnalytics is Licensed

Instances in relation to licensing

FusionAnalytics Users

FusionAnalytics DataCollectors

FusionAnalytics DataCollectors in relation to licensing

I have more than 20 instances to analyze -- what can I do?

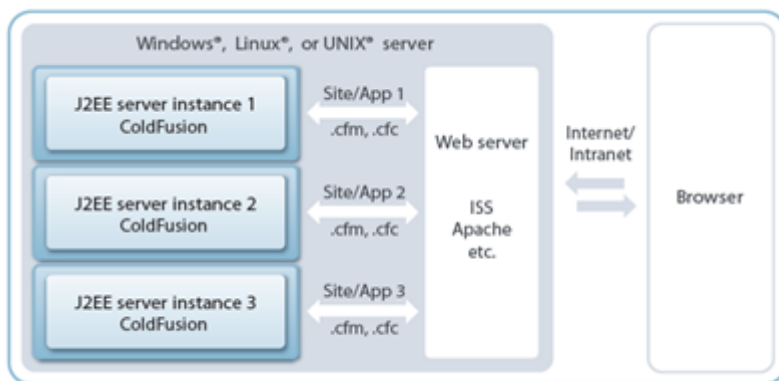


Figure 1: Instances

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

Move FusionAnalytics to a new server

Introduction

This documentation will provide readers with more information about the best way of moving a FusionAnalytics server to a new environment. For this scenario, we will assume that the **database will remain the same** and only the FusionAnalytics server will be transferred to a new environment.

Getting Started

Disable the license key in use

In order to move FusionAnalytics to a new server, please follow the steps mentioned below and make sure that the license key currently used has been disabled.

1. Access the FADC and the FADS application via the browser.
2. Navigate to **System > License Info**.
3. Click the **'Uninstall'** button in order to uninstall the license key from that server.

The steps mentioned above are vital as you will have a server seat available to activate on the new server.

Firewall configuration

In order to avoid any data loss and before the installation of FusionAnalytics, please make sure that a Firewall policy/configuration on the new server is **not blocking traffic between the server and the configured database**. More information about the **prerequisites** can be found in the following link; <https://docs.fusion-reactor.com/display/FA207/Prerequisites+for+Installation>.

Transfer existing configuration to the new server

In order to keep the existing configuration of the FusionAnalytics server, please follow the steps mentioned below.

1. Make a copy of the FusionAnalytics directory somewhere in your environment.
2. Access the new server.
3. Run the FusionAnalytics installer using the appropriate privileges. More information about the installation process can be found in the following link; FusionAnalytics Installation.
4. Make sure the installation process was successful and that you are able to locate the 'FusionAnalytics Server' windows service.
5. Stop the 'FusionAnalytics Server' service.
6. Replace the FusionAnalytics folder with the one copied in step 1.
7. Start the 'FusionAnalytics Server' service.
8. Make sure that you are able to access the FADC and FADS applications via the browser and with the existing configurations.

FusionReactor configuration changes

Because the FusionAnalytics server has been transferred to a new environment, you need to make changes in the **'FusionAnalytics Target'** and in the **'FusionAnalytics Settings'** found in the FusionReactor UI. For more information, please check out the links below.

- <https://docs.fusion-reactor.com/display/FR62/FusionAnalytics+DataCollector+Targets>

- <https://docs.fusion-reactor.com/display/FR62/FusionAnalytics+Settings>

If FusionReactor was not configured with FusionAnalytics and you have an hour downtime, you have the ability to manually add the logs in the datacollector. This can be done by pressing the 'Transfer All' button found under **Analytics > Connector Status**.

Test connection between FusionReactor and FusionAnalytics

Most connectivity issues between FusionReactor and FusionAnalytics revolve around the server where FusionReactor is running, not being allowed or able to connect to the FusionAnalytics server. Please make sure that the connection between FusionReactor and the FusionAnalytics server is successful. This can be achieved by either clicking the 'Test' button found in the 'FusionAnalytics DataCollector Targets' area **OR** by using a Web Browser directly on the server that is running FusionReactor and seeing if you can open FADC (FusionAnalytics DataCollector) from that server.

More information about the 'FusionAnalytics DataCollector Targets' configuration can be found in the following link; <https://docs.fusion-reactor.com/display/FR62/FusionAnalytics+DataCollector+Targets>.

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)

- 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
 - Most Hits By IP Address
 - Longest Request Execution Time By IP Address
 - Aggregate Request List
 - Request Breakdown
 - Request Details

- 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)
 - Requests Restart Coverage
 - Database Hot-Spots
 - Slowest Database Requests (Hot-Spots)
- ColdFusion Monitor
 - CF General Statistics
 - CF Scope Sizes
 - CF Hit Counts
 - CF DB Pool Stats
- Applications

FUSIONANALYTICS REPORTING

- TAP Reports
 - What is TAP Score - and how is it calculated?
 - Running a TAP Report
 - TAP Report Provider Arguments
 - Understanding the TAP Report
 - The TAP Report configuration page
- Daily Status Reports
 - Running a Daily Status Report
 - Daily Status Report Provider Arguments
 - Understanding the Daily Status Report

- The Daily Status Report configuration page
- SLA Reports
 - Accessing the SLA Report Configuration Page
 - Adding a new SLA Report Configuration
 - The SLA Report Configuration
 - Report Configuration Tab
 - Email Recipients Tab
 - Domain Filters Tab
 - Thresholds Tab
 - Running the SLA Reports

USING FUSIONANALYTICS TO ANALYZE A CLUSTER OF SERVERS / INSTANCES

- Cluster Perspectives

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
 - Most Hits By IP Address
 - Longest Request Execution Time By IP Address
 - Aggregate Request List
 - Request Breakdown
 - Request Details
- 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)
 - Requests Restart Coverage
 - Database Hot-Spots
 - Slowest Database Requests (Hot-Spots)
- ColdFusion Monitor
 - CF General Statistics
 - CF Scope Sizes
 - CF Hit Counts
 - CF DB Pool Stats
- Applications

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. 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 |
| 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, Throt Q, Template, Users, Temp Cache or Query Cache |

Overview - Perspective View

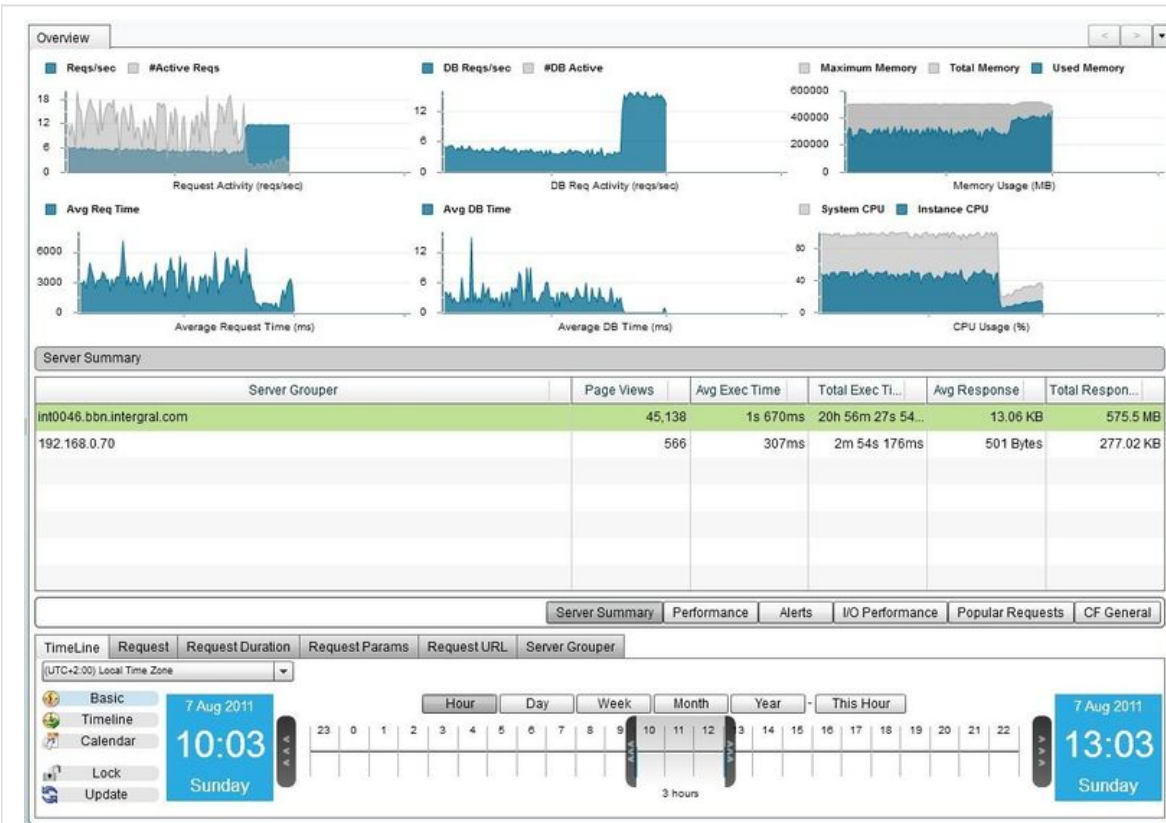


Figure 2: Overview Perspective

[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
- Most Hits By IP Address
- Longest Request Execution Time By IP Address
- Aggregate Request List
- Request Breakdown
- Request Details

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)
 - Requests Restart Coverage
- Database Hot-Spots
 - Slowest Database Requests (Hot-Spots)

ColdFusion Monitor

- CF General Statistics
- CF Scope Sizes
- CF Hit Counts
- CF DB Pool Stats

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 show 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.

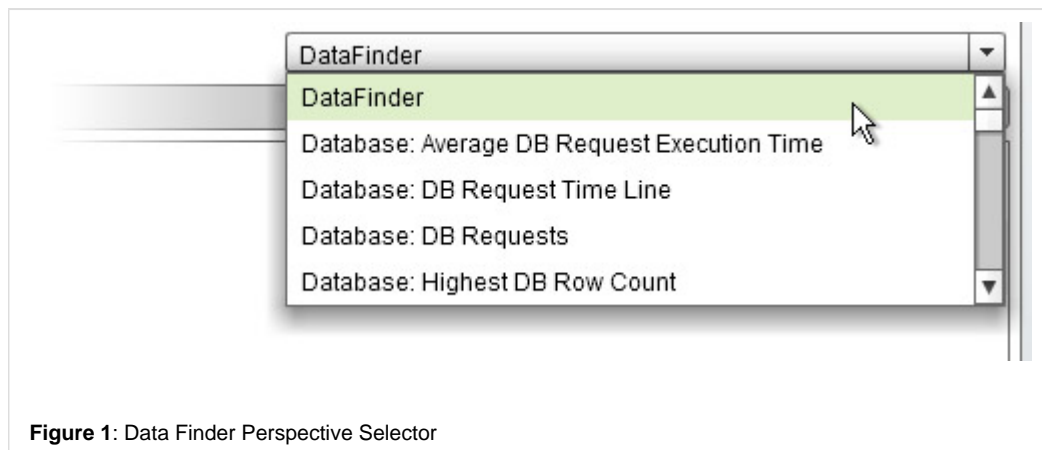


Figure 1: Data Finder Perspective Selector

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



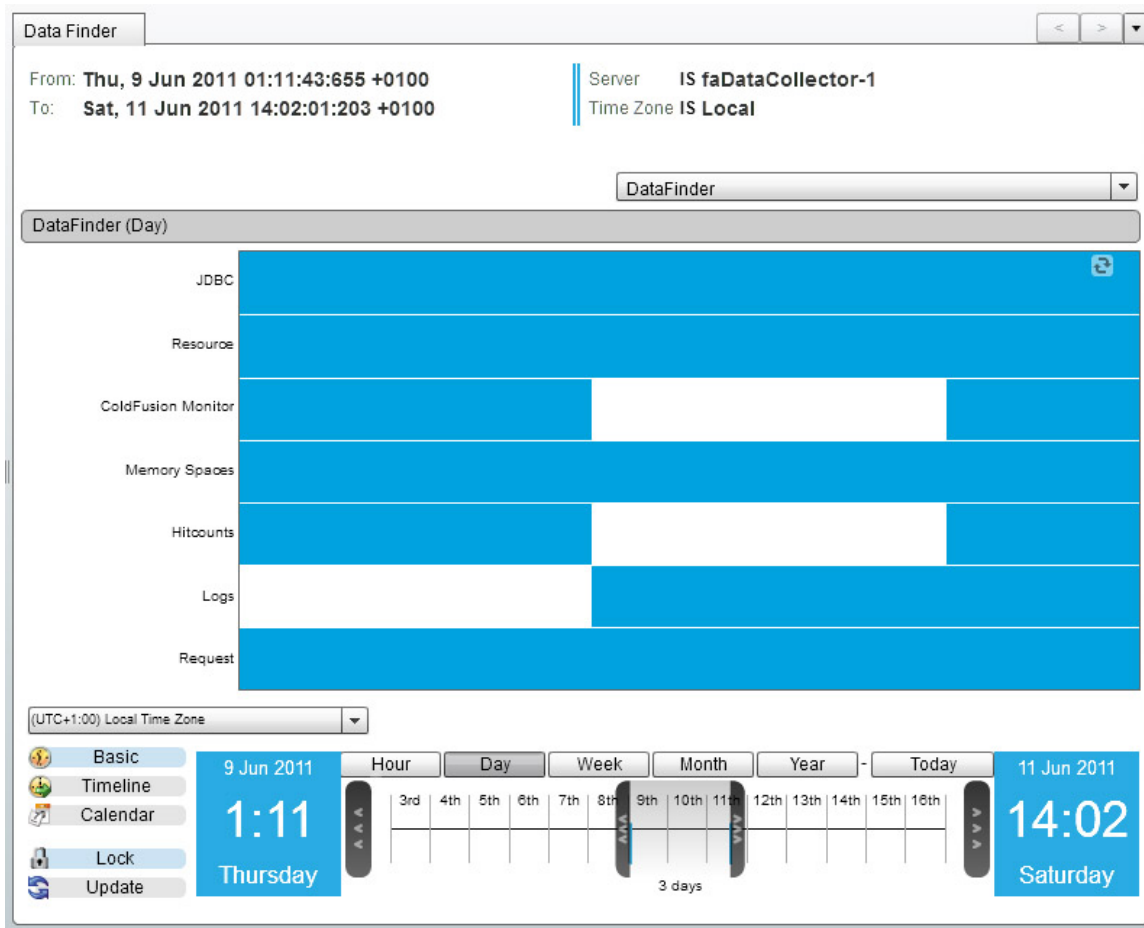


Figure 2: Data Finder Perspective

[Back to the top](#)

Request

Request

Overview

Next Steps

Slowest Requests

Slowest Requests 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

Most Hits By IP Address

Longest Request Execution Time By IP Address

Aggregate Request List

Request Breakdown

Request Details

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 shown, click on one of the table column headings. Use the Date Navigation in order to change covers.

You can right click on any of the requests in the table to bring up a context menu. This will allow you to navigate to other useful perspectives Requests and System Metrics.

You can also view Request Details, to see specific details of that request, such as the exact end time, the amount of memory it used, and

Perspective View

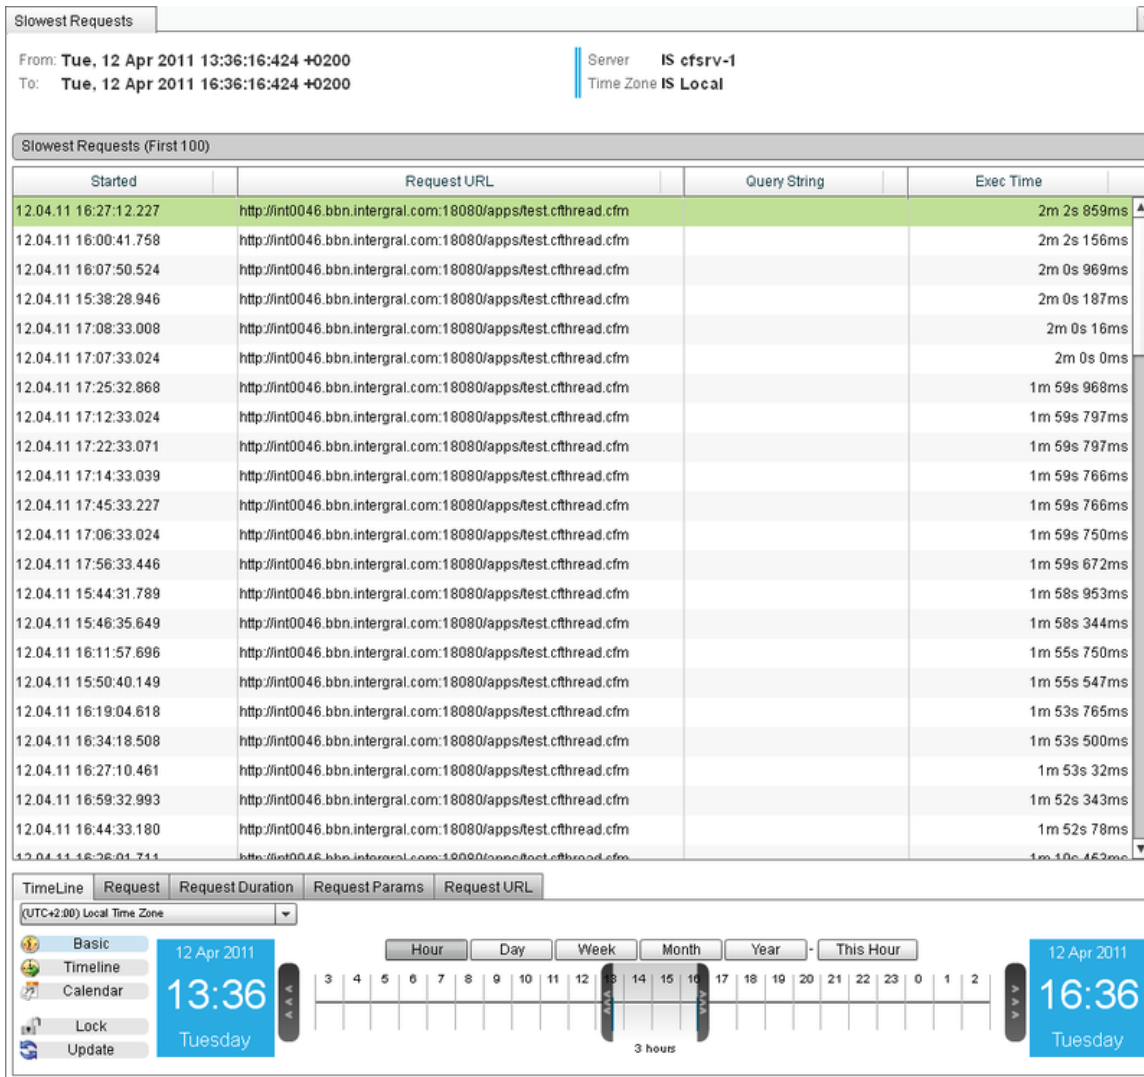


Figure 1: Slowest Requests Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Slowest Requests fr It shows the slowest requests from the entire cluster rather than just one server. You can see which server each request came from in the Name.

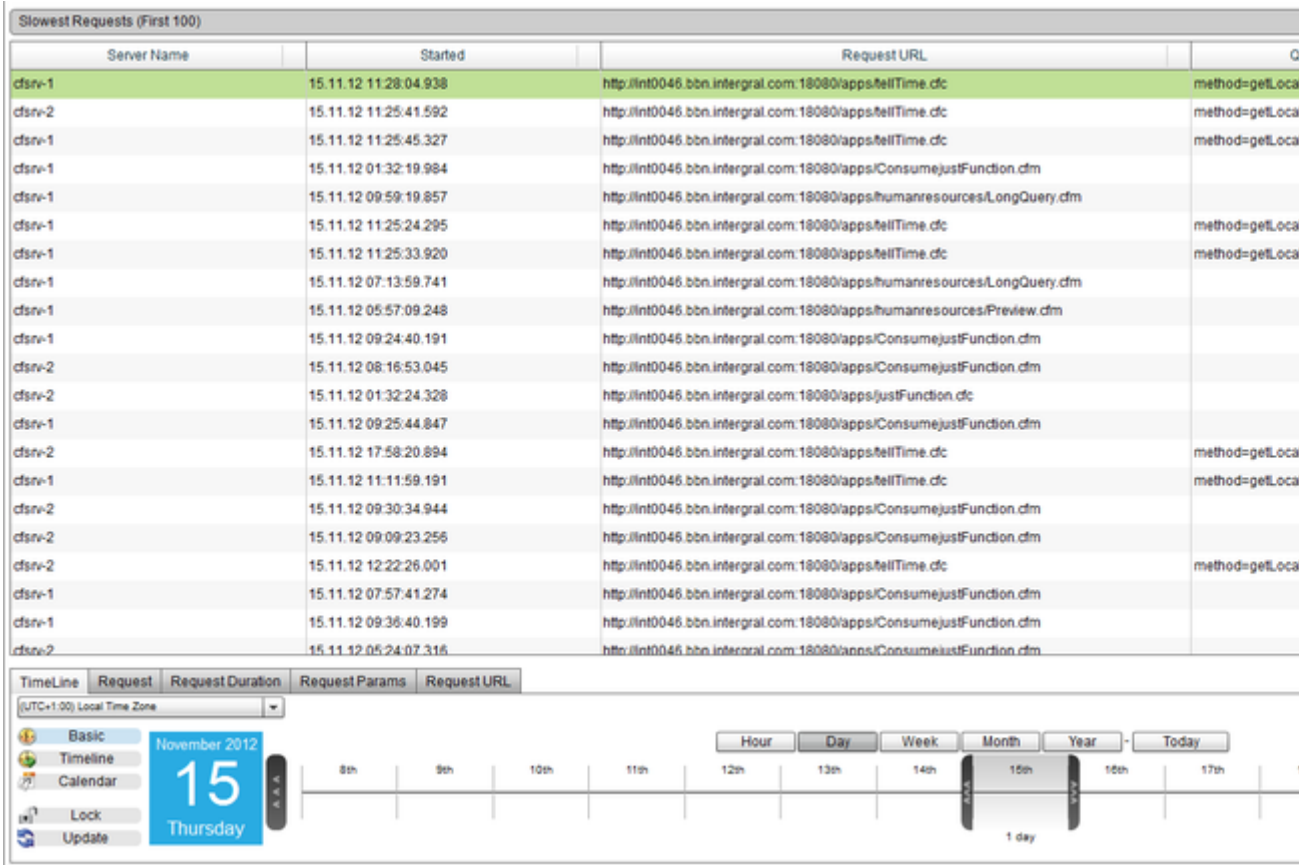


Figure 2: Cluster Version of Slowest Requests Perspective

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 shown, click on one of the table column headings. Use the Date Navigation in order to char covers.

Perspective View



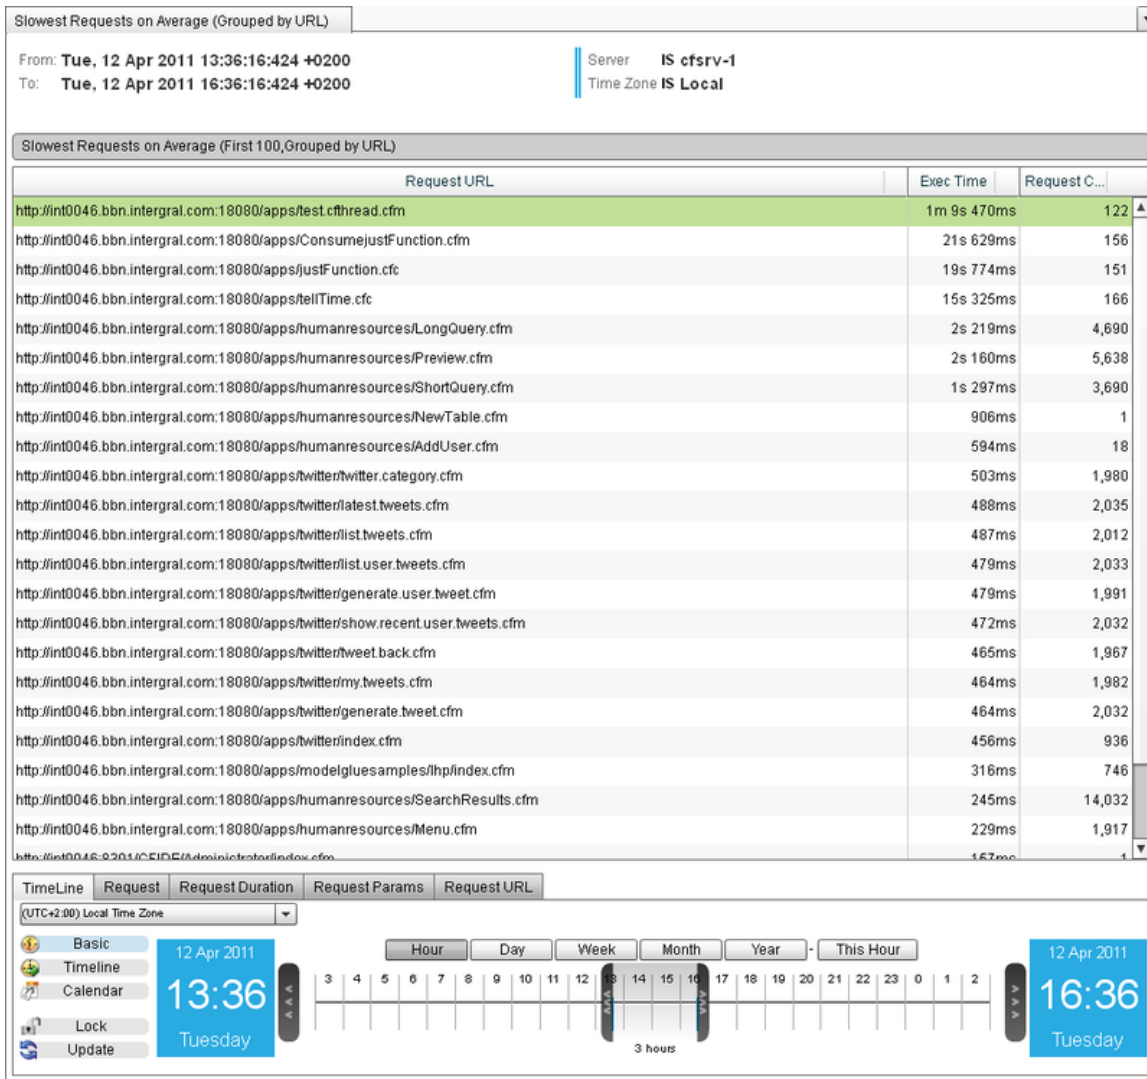


Figure 1: Slowest Requests on Average Perspective

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Slowest Requests On It shows the First 100 Requests, ordered by slowest Average Exec Time for the whole cluster. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

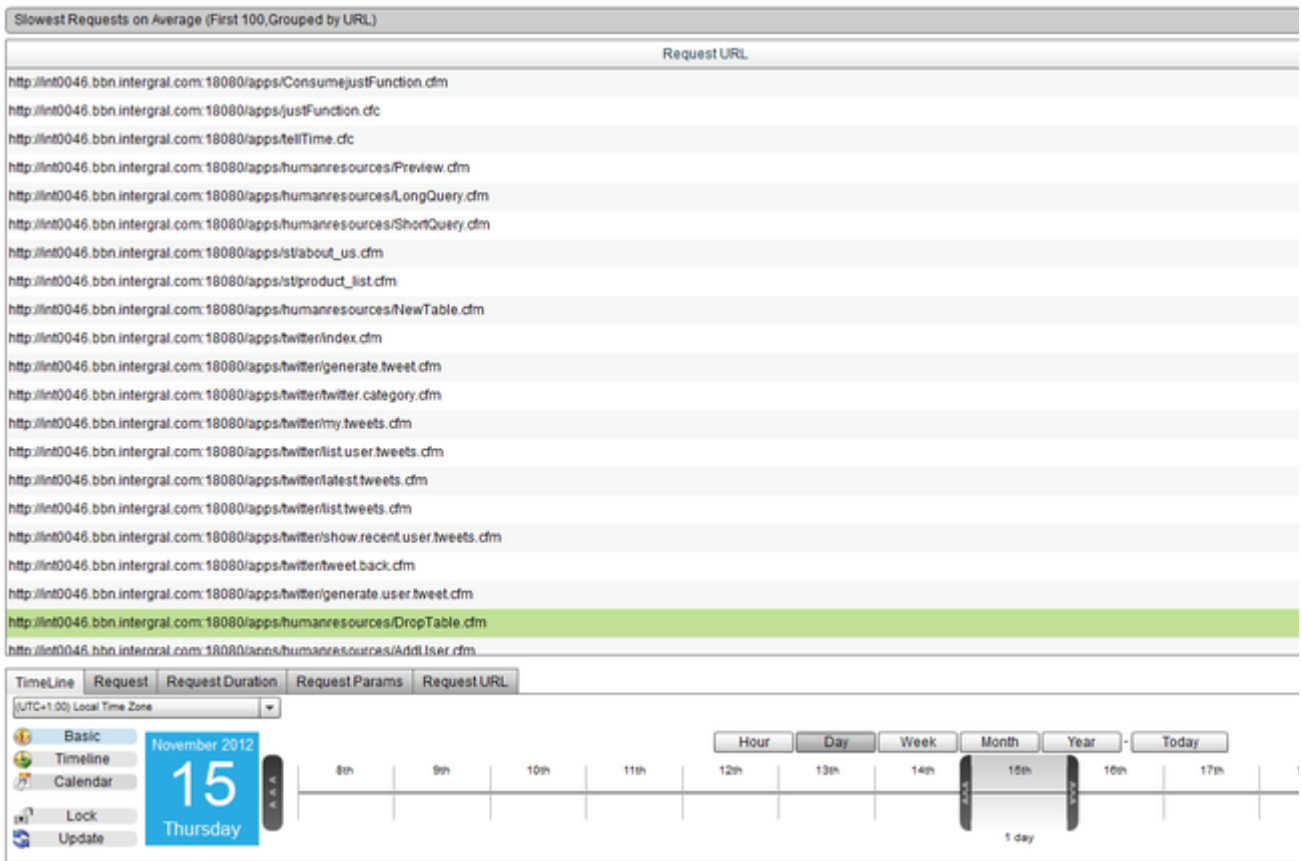


Figure 2: Cluster version of Slowest Requests On Average perspective

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 shown, click on one of the table column headings. Use the Date Navigation in order to change covers.

Perspective View



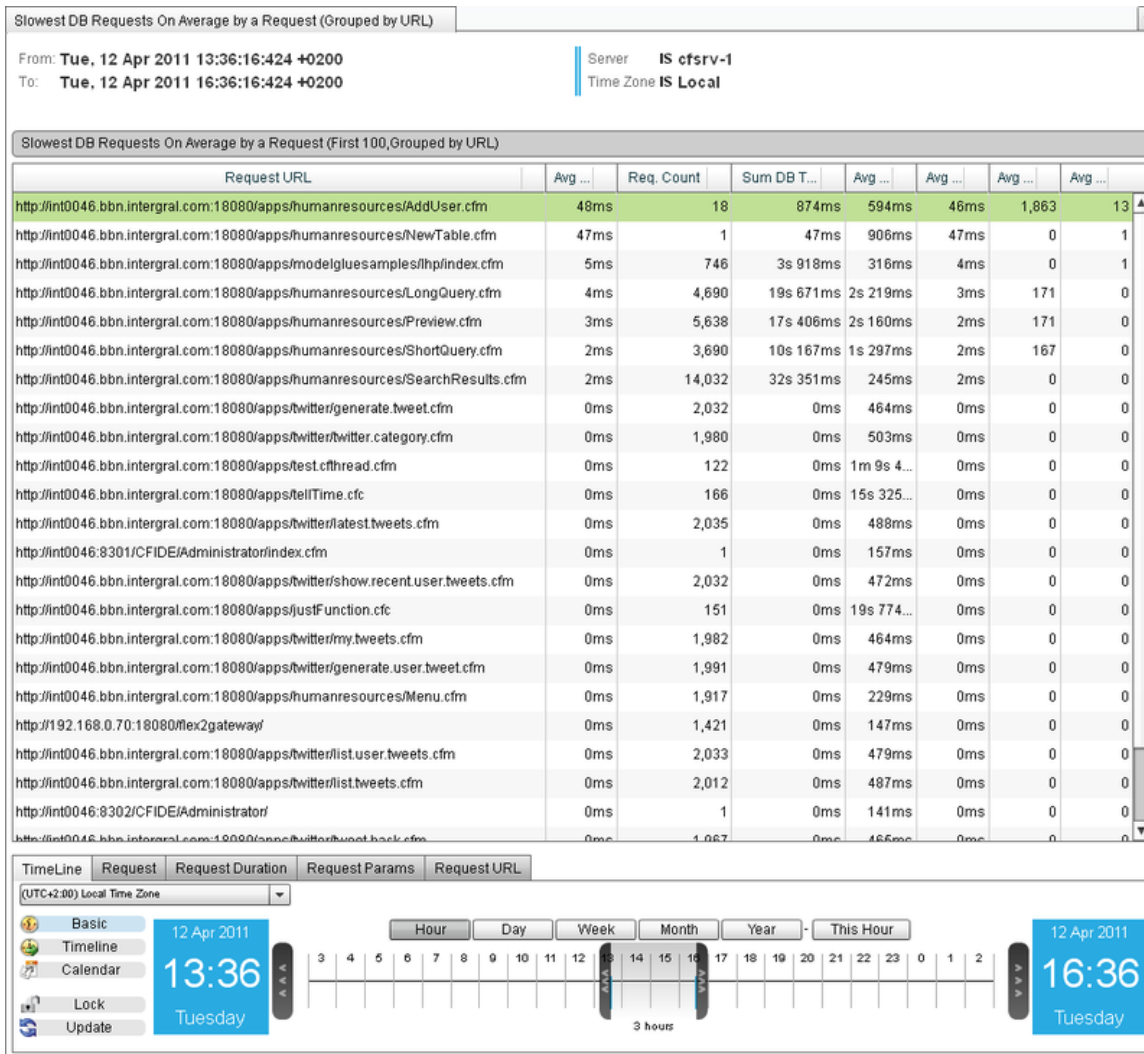


Figure 1: Slowest DB Requests On Average Perspective

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Slowest DB Requests It shows the First 100 Requests, ordered by slowest Average DB Exec Time for the whole cluster. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

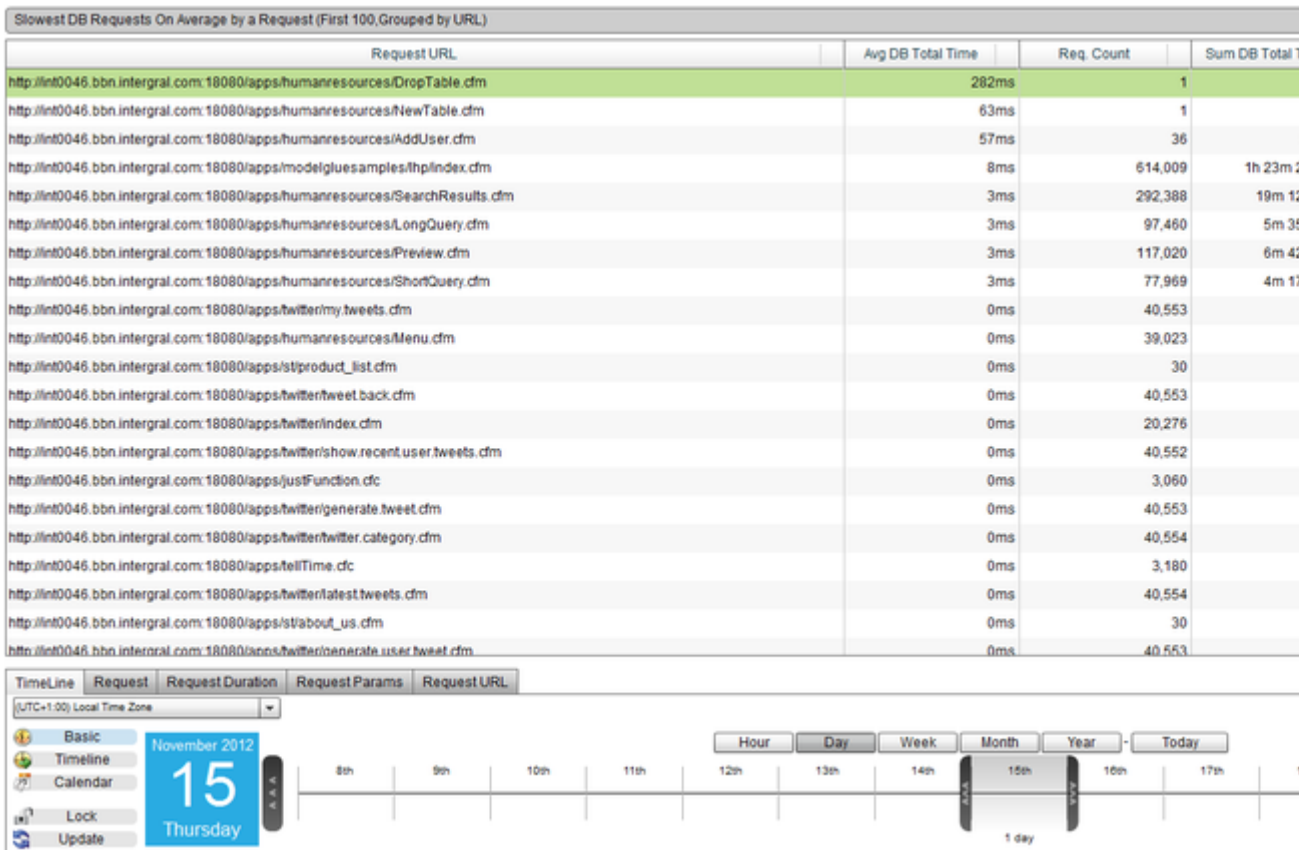


Figure 2: Cluster Version of Slowest DB Requests On Average perspective

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. This is an important perspective for analyzing the performance of an application or a server. The total execution time is the amount of time that the request actually spent running on the server. This perspective can help you analyze:

- Fast request but runs so often that they use they have the highest total execution time
 - In this case optimizing even a small amount of time out the request can reap big benefits simply because of multiplying the request is run
- Slow Requests take so long to execute that they use more execution time than the most popular requests
 - In this case look to optimize the request heavily. Consider caching strategies and when the requests run to reduce the impact
- Requests that vary in performance heavily
 - It's possible that a request typically runs in an acceptable time, but sometimes seems to spike a take much longer. These requests can cause the server to be unresponsive.
 - Identify why the request sometimes take a long time to run. Look at external resources such as to Databases and other network resources
 - Check the request parameters on slow requests. It could be that the request only runs slowly when passed specific parameters
 - Check the timing to see if the request only runs slowly at certain times and look at what other activity the server or application is doing
 - Check for locking issues with other requests or resources. It is common that a request may execute for a long time, holding a lock that conflict with other requests running at the same 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 characterize the data.

Perspective View

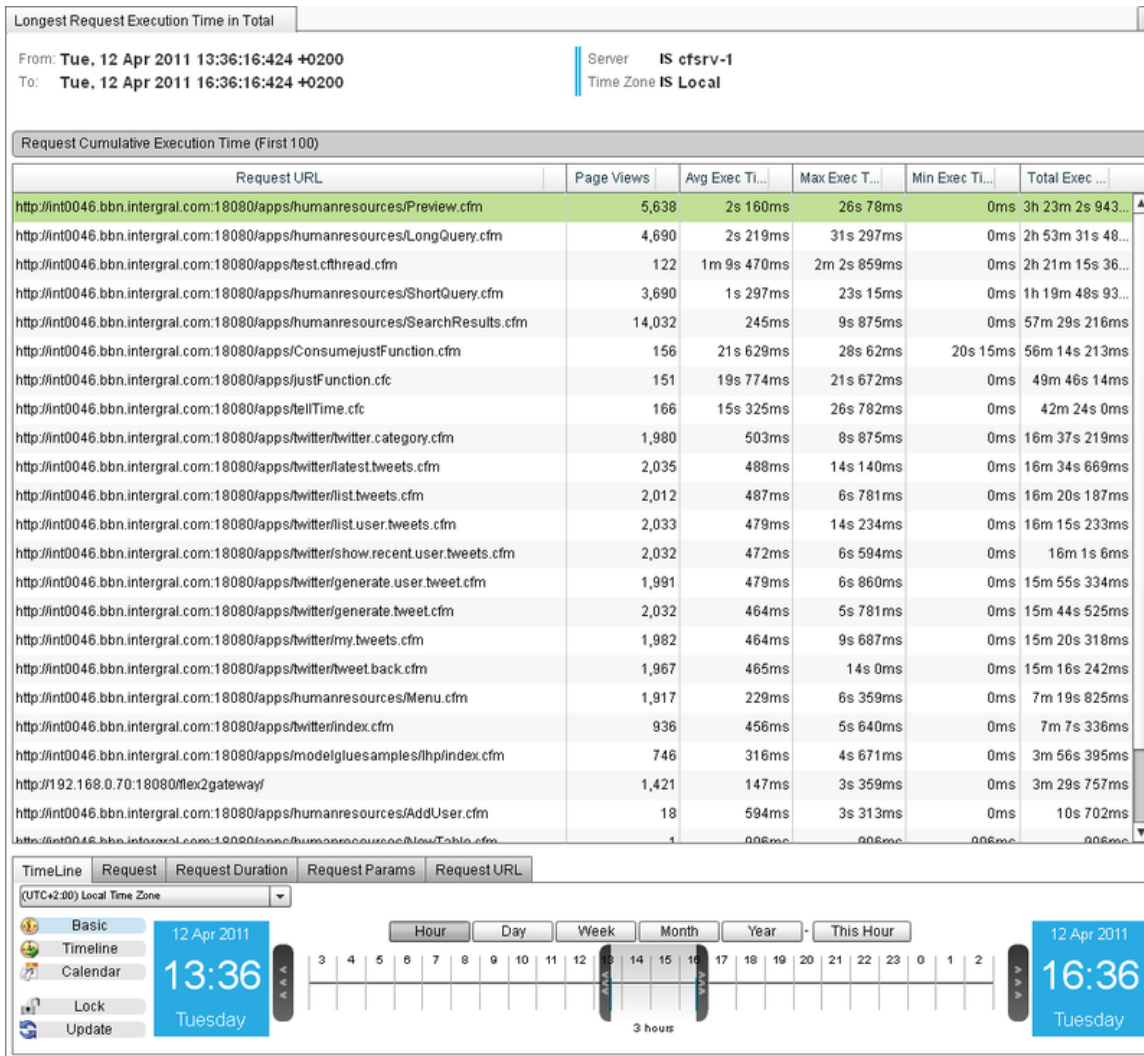


Figure 1: Longest Request Execution Time in Total Perspective

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Longest Request Execution Time in Total.

It shows the First 100 Requests, ordered by Total Exec Time for the whole cluster.

The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

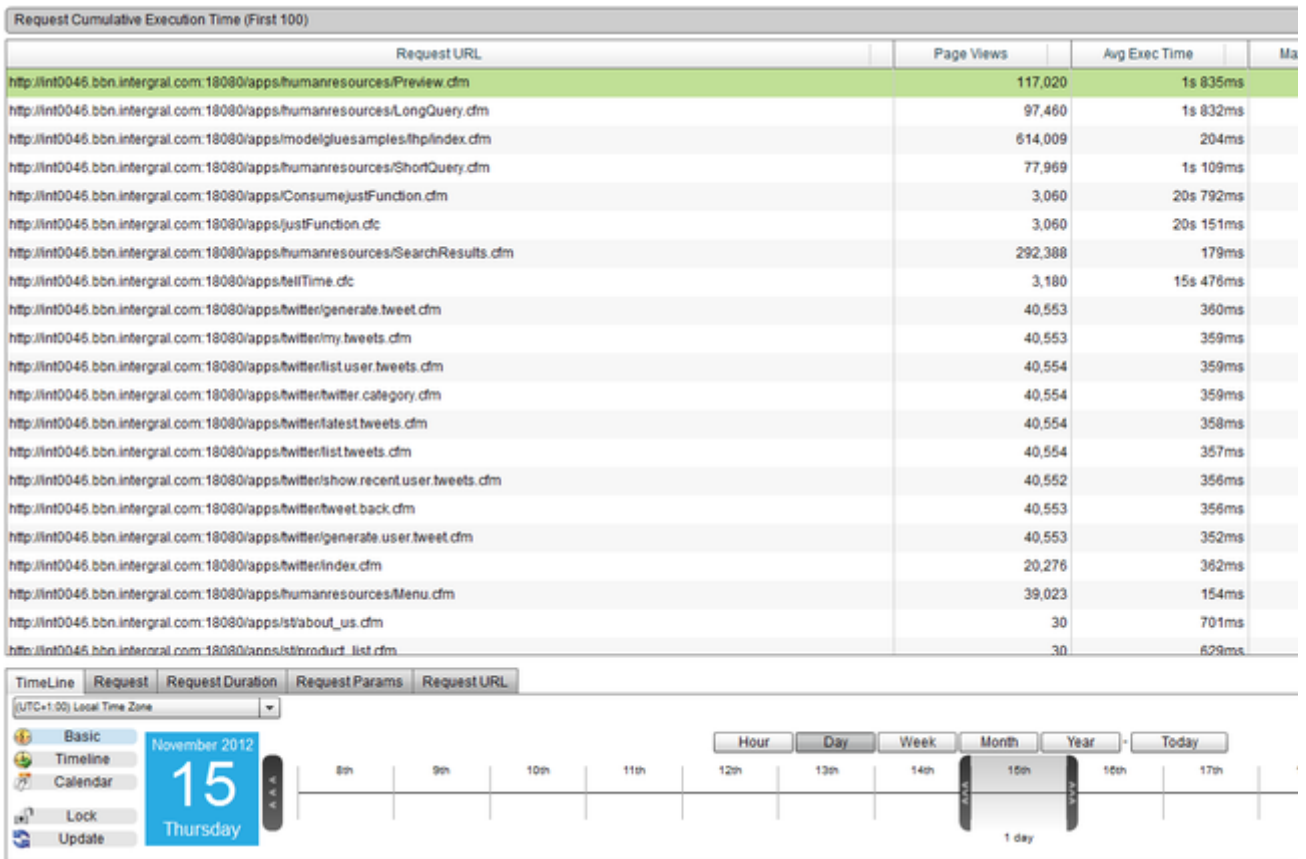


Figure 2: Cluster Version of the Longest Request Execution Time in Total perspective

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 rep database by a request. Knowing the number of rows returned from the database can help you to identify many potential issues and give yr applications. Generally fetching rows from the database takes time, resource and can induce network overhead if the database is located database are stored in memory and this can increase the overall memory usage of your application or potentially create spikes in memory

Using this report you should look to fix, tune or optimize the requests that return most rows from the database. This will typically increase 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



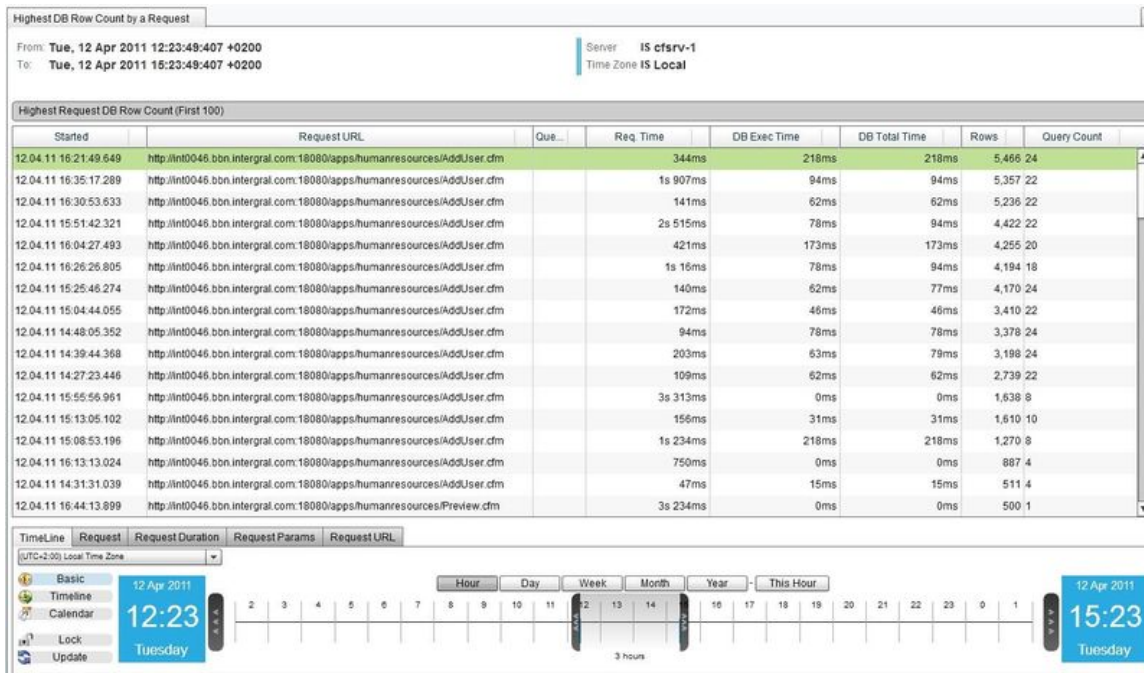


Figure 1: Highest DB Row Count by a Request Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Highest DB Row Count Comparison Grid.

It shows the requests from the entire cluster, which fetched the most rows of data from the database. You can see which server each request which displays Server Name.

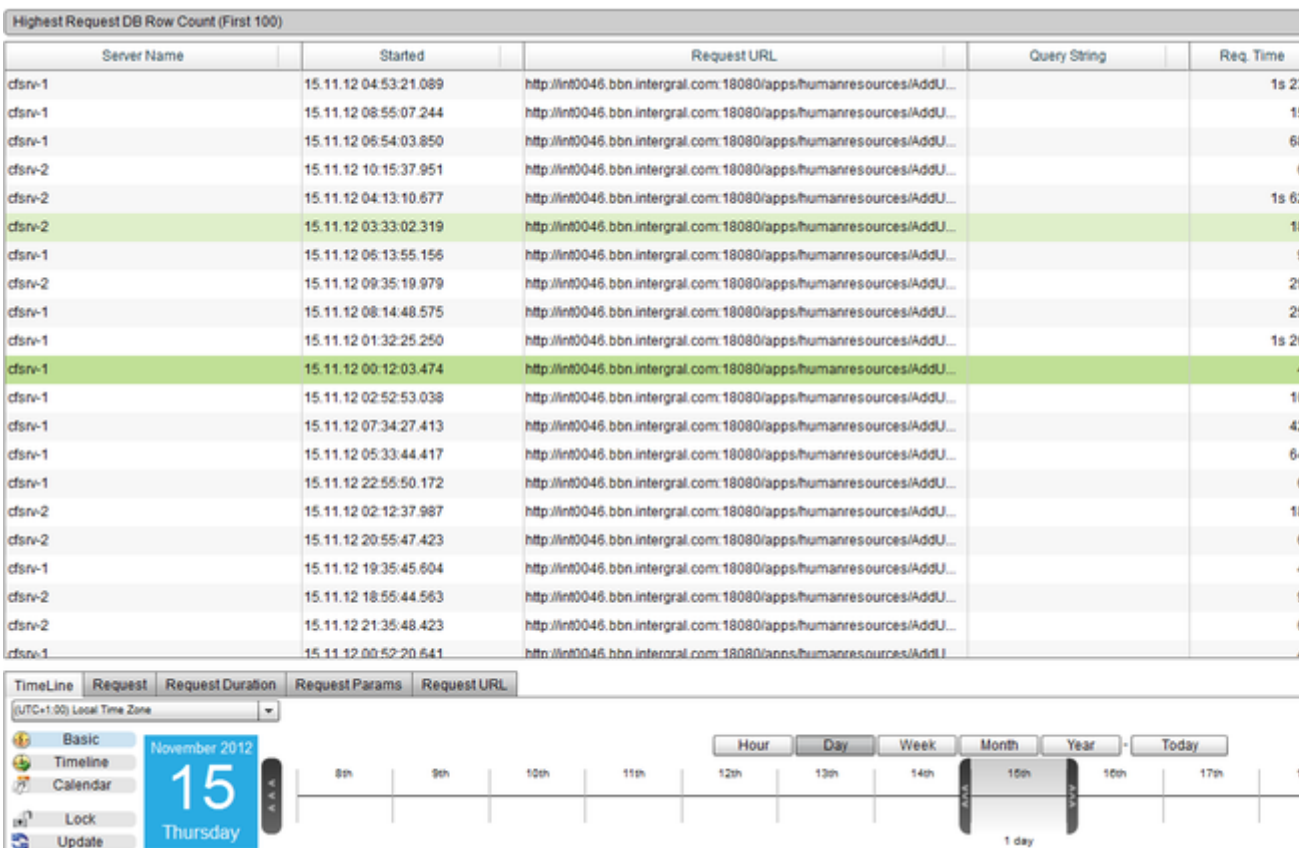


Figure 2: Cluster version of the Highest DB Row Count by a Request perspective

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 identify which requests fetch the most rows of data from the database. It fetched from the database per request run grouped them by their URL. Knowing the number of rows returned from the database can help insight into the resource usage and scale of your applications. Generally fetching rows from the database takes time, resource and can hit a remote server. Often rows fetched from the database are stored in memory and this can increase the overall memory usage of your app.

Using this report you should look to fix, tune or optimize the requests that return most rows from the database. This will typically increase 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

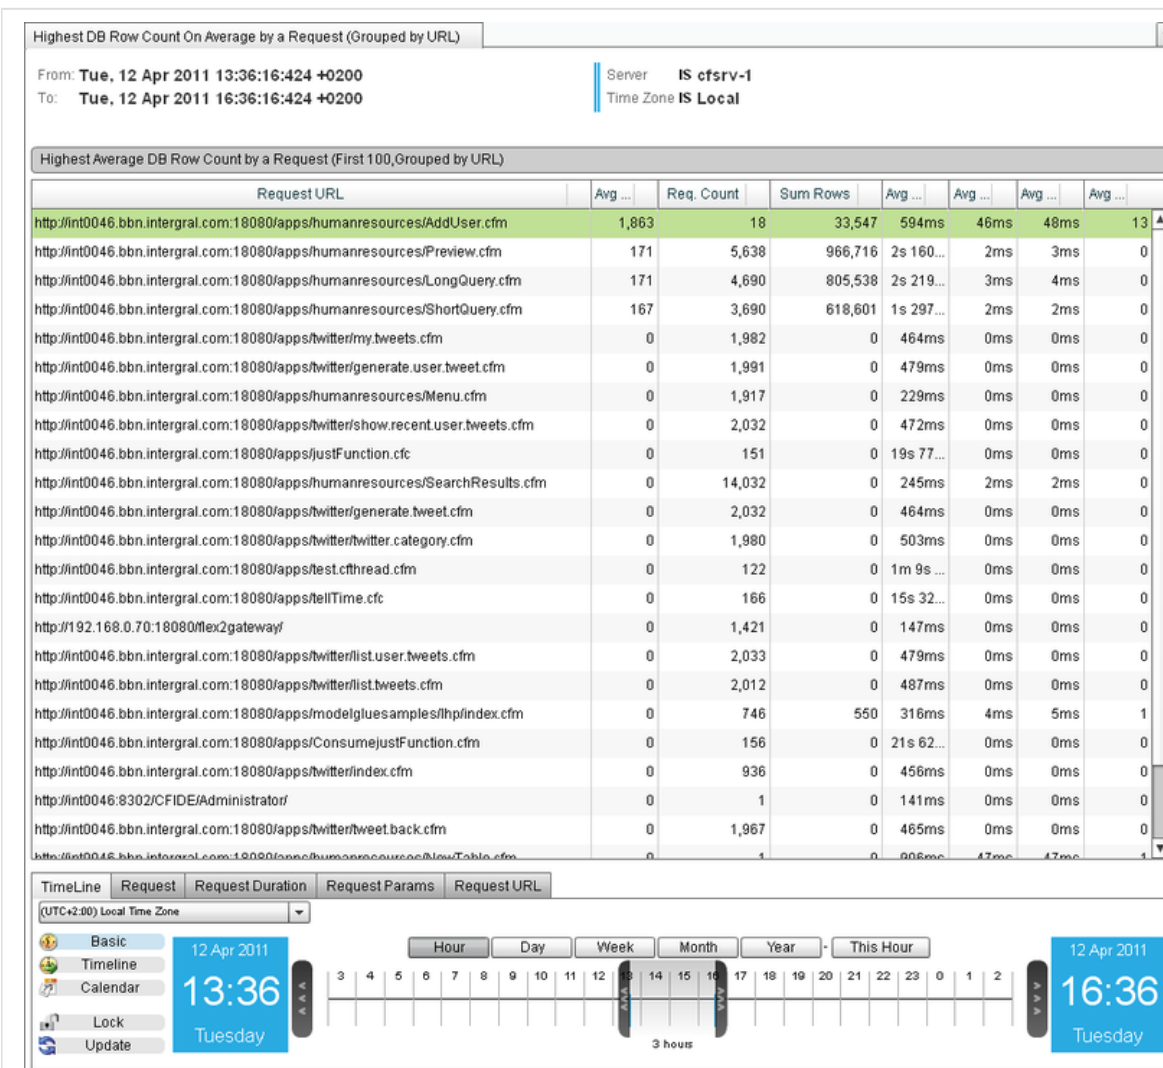


Figure 1: Highest Total DB Row Count On Average Perspective

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Highest DB Row Count View Grid.

It shows the First 100 Requests that have the highest average number of rows fetched from the database. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

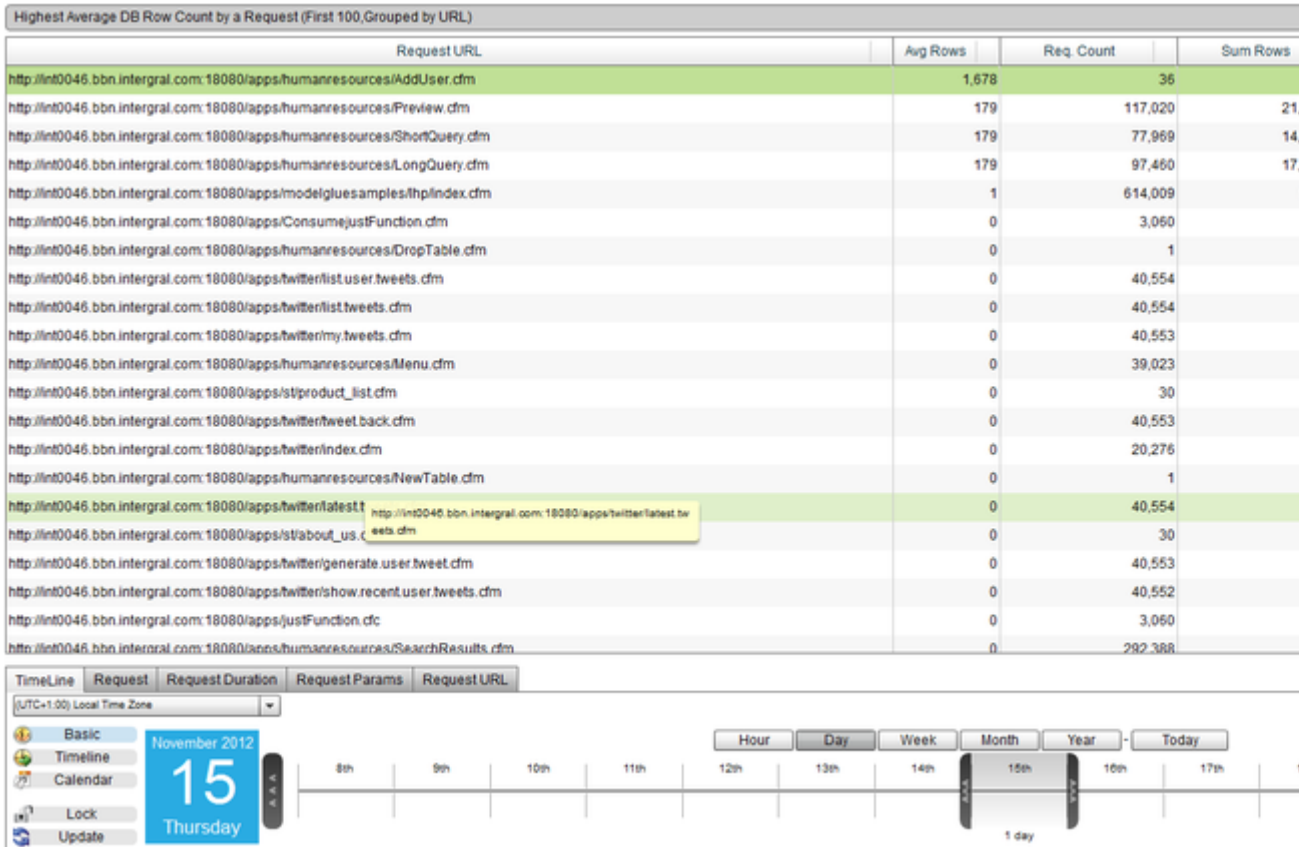


Figure 2: Cluster version of the Highest DB Row Count On Average by a Request perspective

Back to the top

HIGHEST CPU USAGE

Highest CPU Usage

Description

This perspective shows a table of the top 100 individual requests by the amount of time spent on the CPU (CPU Time) within the specific t

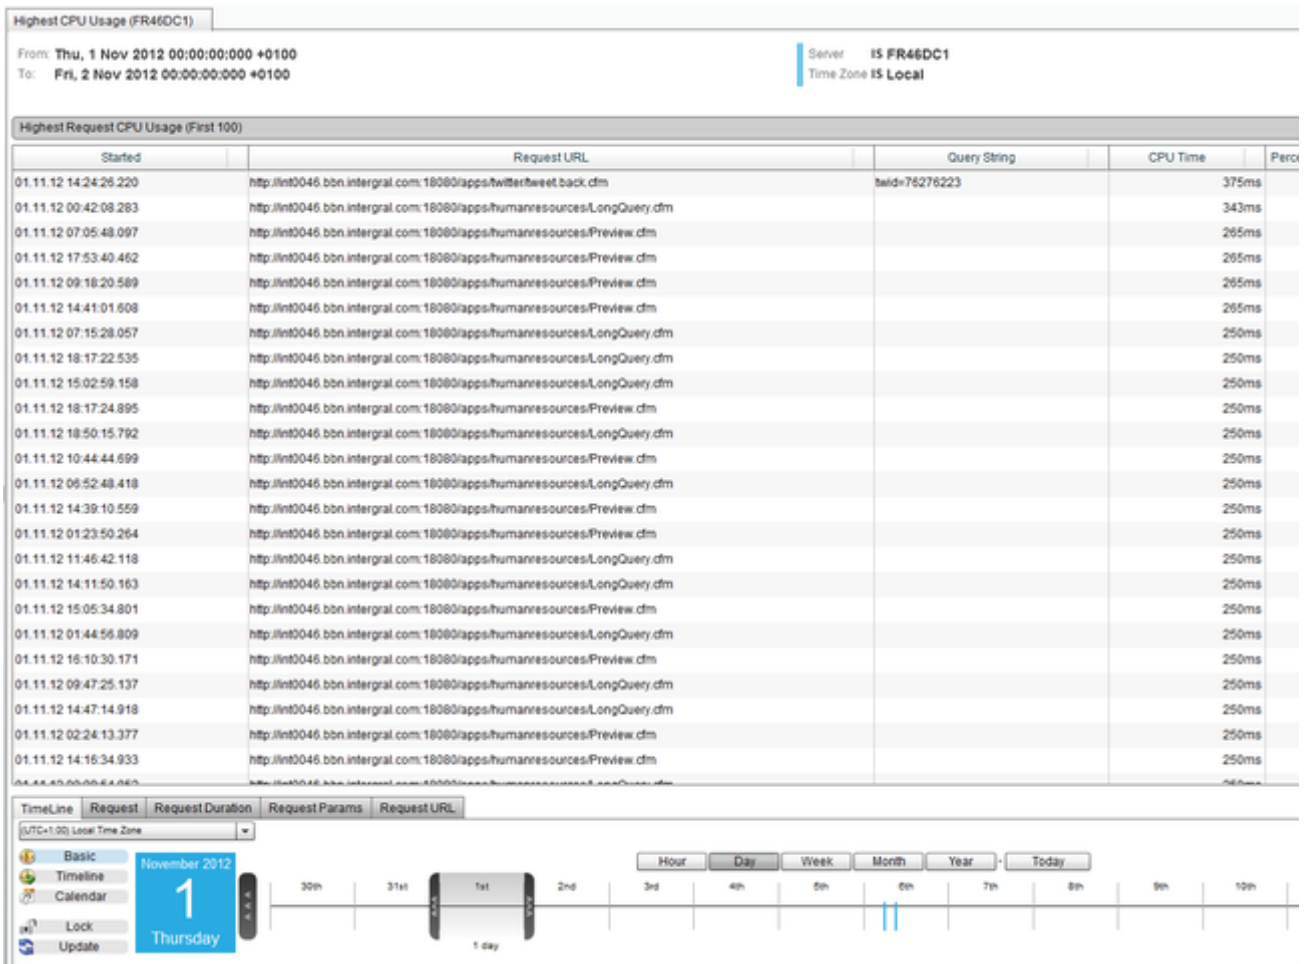
Usage

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 chan covers.

You can right click on any of the requests in the table to bring up a context menu. This will allow you to navigate to other useful perspective Requests and System Metrics.

You can also view Request Details, to see specific details of that request, such as the exact end time, the amount of memory it used, and

Perspective View



Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Highest CPU Usage. It shows the top 100 requests from the entire cluster, by the amount of time spent on the CPU (CPU Time). You can see which server each request came from in the leftmost column of the table, which displays Server Name.

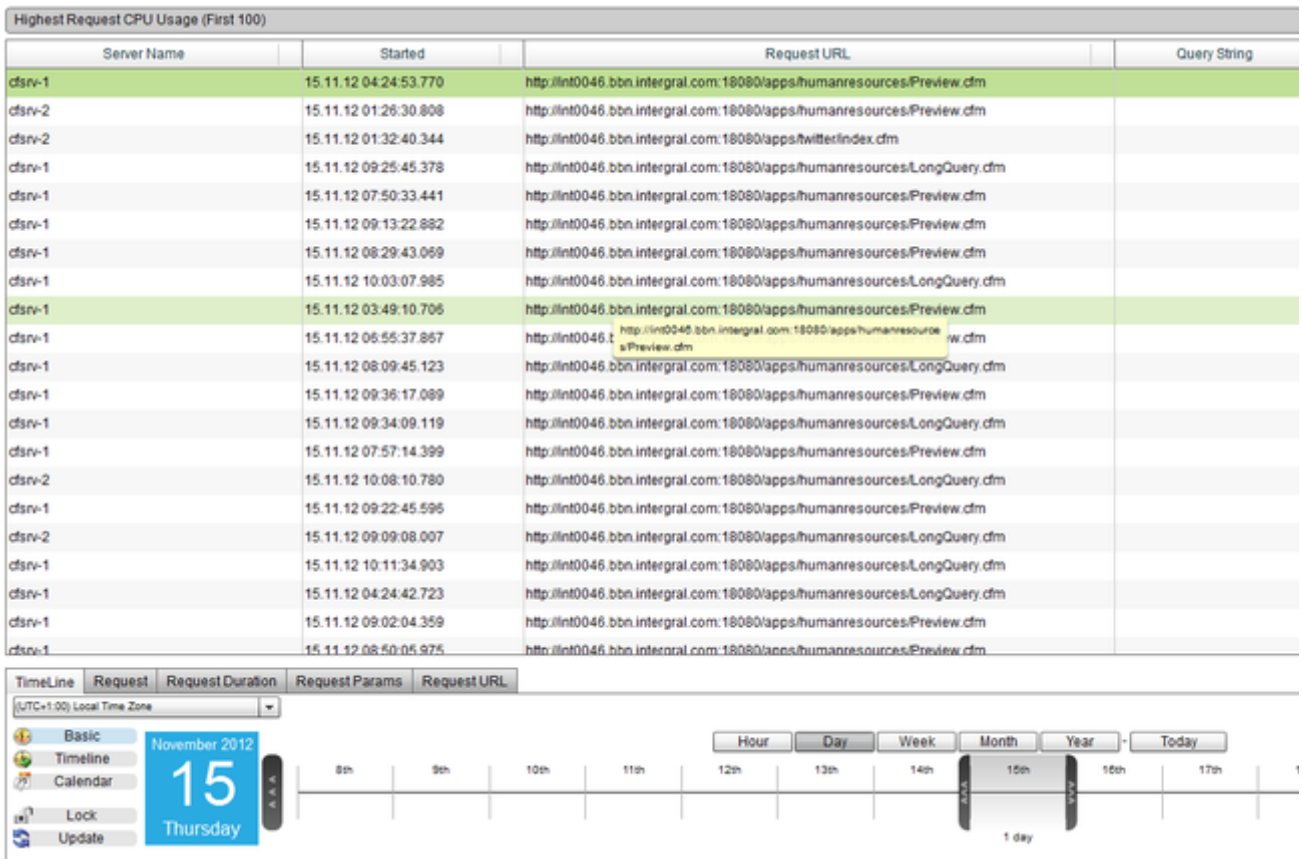


Figure 2: Cluster version of the Highest CPU Usage perspective

Back to the top

HIGHEST CPU USAGE ON AVERAGE

Highest CPU Usage on Average

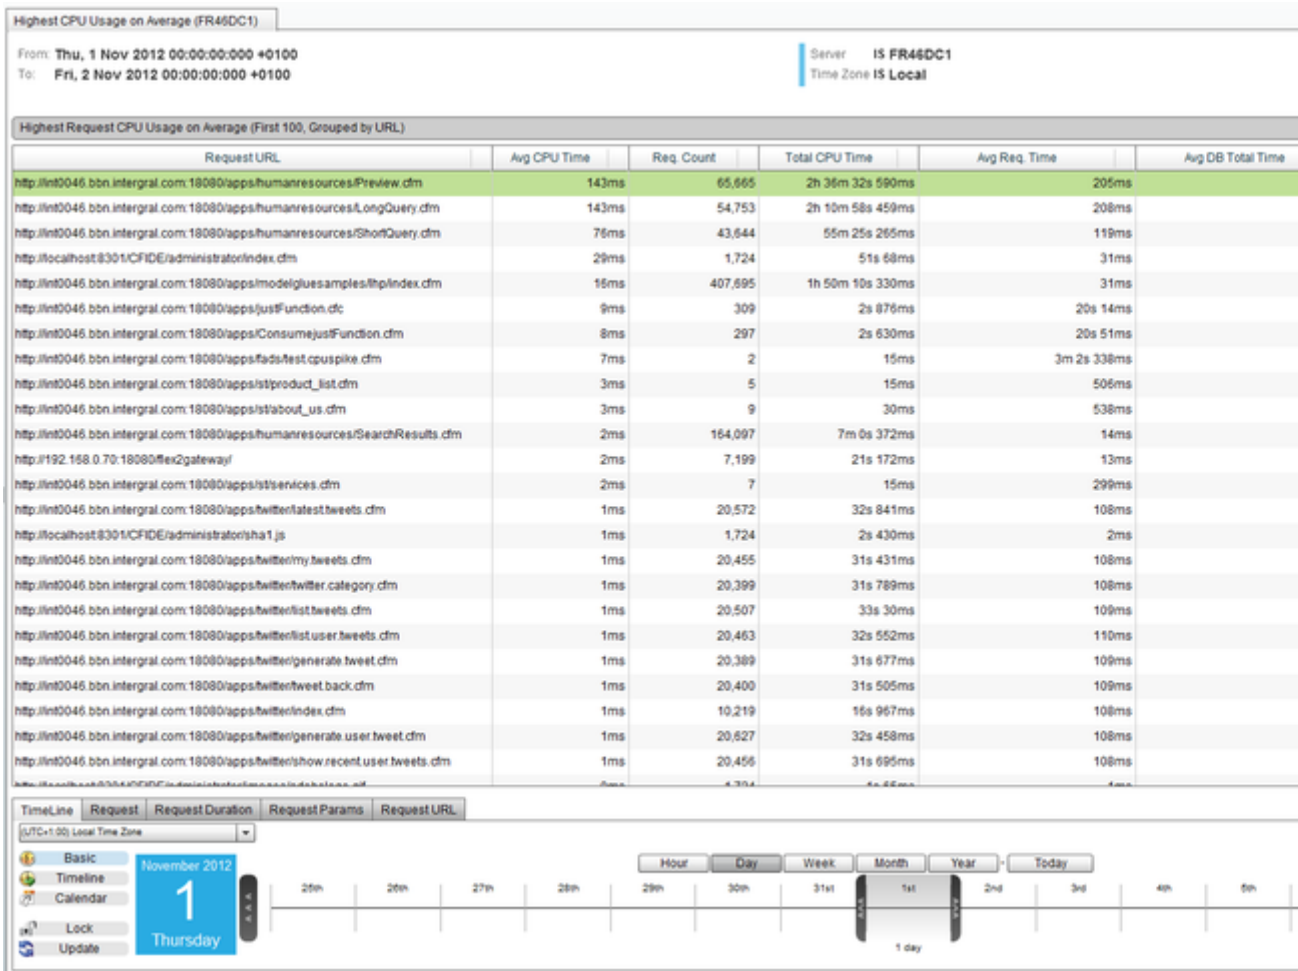
Description

This perspective shows a table of the top 100 requests by their average time spent on the CPU(Avg CPU Time) within the specific period

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 char covers.

Perspective View



Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Highest CPU Usage c It shows the First 100 Requests for the entire cluster, ordered by the Average Time spent on the CPU (Avg CPU Time). The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

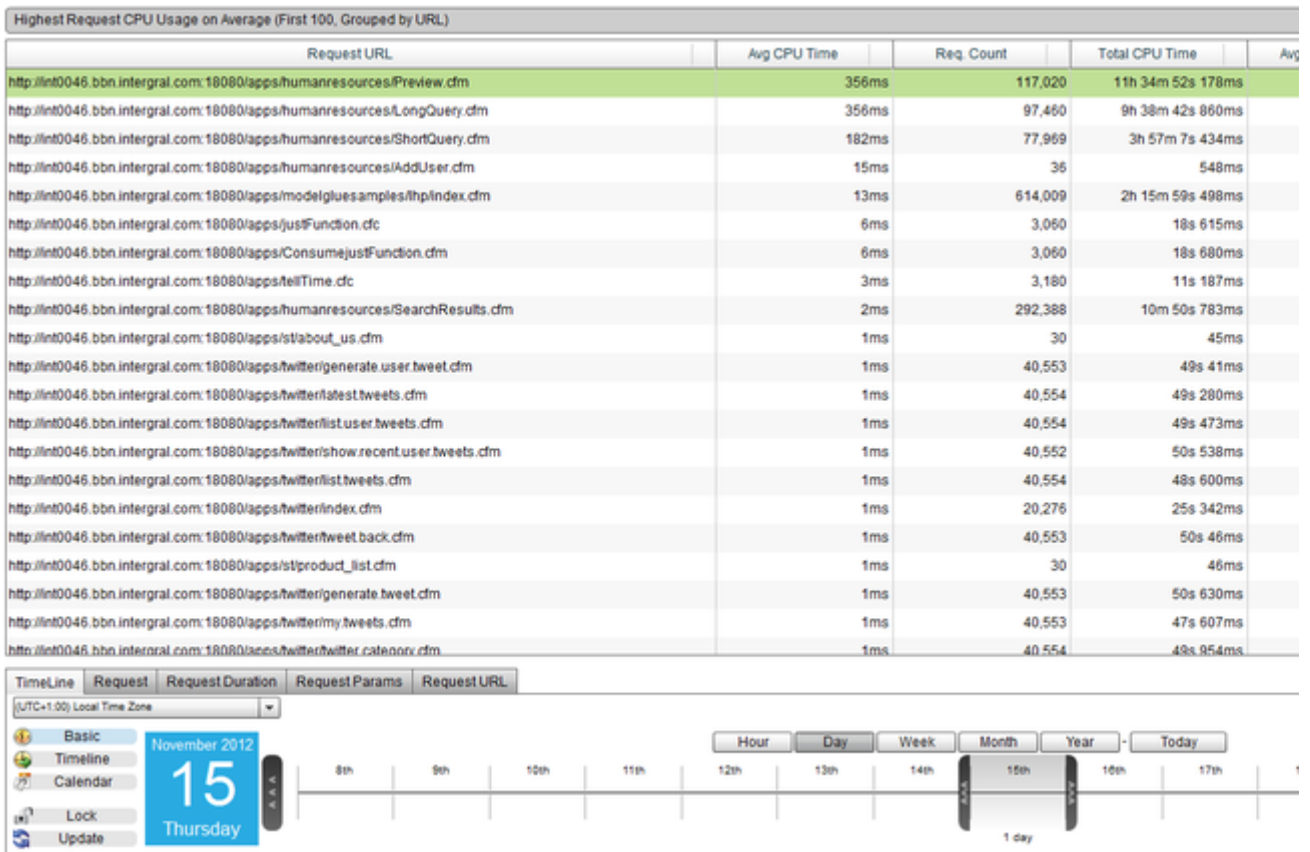


Figure 2: Cluster version of the Highest CPU Usage on Average perspective

Back to the top

HIGHEST CPU USAGE IN TOTAL

Highest CPU Usage in Total

Description

This perspective shows a table of the top 100 requests by the total amount of time they have spent using the CPU (Total CPU Time) during

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 covers.

Perspective View

Highest CPU Usage in Total (FR46DC1)

From: Thu, 1 Nov 2012 00:00:00:000 +0100
 To: Fri, 2 Nov 2012 00:00:00:000 +0100

Server: IS FR46DC1
 Time Zone: IS Local

Highest Request CPU Usage in Total (First 100, Grouped by URL)

| Request URL | Total CPU Time | Avg CPU Time | Req. Count | Avg Req. Time |
|--|------------------|--------------|------------|---------------|
| http://int0046.bbn.integral.com:18080/apps/humanresources/Preview.cfm | 2h 36m 32s 590ms | 143ms | 65,665 | 205ms |
| http://int0046.bbn.integral.com:18080/apps/humanresources/LongQuery.cfm | 2h 10m 58s 459ms | 143ms | 54,753 | 208ms |
| http://int0046.bbn.integral.com:18080/apps/model/gluesamples/ftpIndex.cfm | 1h 50m 10s 330ms | 16ms | 407,695 | 31ms |
| http://int0046.bbn.integral.com:18080/apps/humanresources/ShortQuery.cfm | 55m 25s 265ms | 76ms | 43,844 | 119ms |
| http://int0046.bbn.integral.com:18080/apps/humanresources/SearchResults.cfm | 7m 0s 372ms | 2ms | 164,097 | 14ms |
| http://localhost:8301/CFIDE/administrator/index.cfm | 51s 68ms | 29ms | 1,724 | 31ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/list/tweets.cfm | 33s 30ms | 1ms | 20,507 | 109ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/latest/tweets.cfm | 32s 841ms | 1ms | 20,572 | 108ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/list/user/tweets.cfm | 32s 552ms | 1ms | 20,463 | 110ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/generate/user/tweet.cfm | 32s 458ms | 1ms | 20,627 | 108ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/category.cfm | 31s 789ms | 1ms | 20,399 | 108ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/show/recent/user/tweets.cfm | 31s 695ms | 1ms | 20,456 | 108ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/generate/tweet.cfm | 31s 677ms | 1ms | 20,389 | 109ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/tweet/back.cfm | 31s 505ms | 1ms | 20,400 | 109ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/my/tweets.cfm | 31s 431ms | 1ms | 20,455 | 108ms |
| http://192.158.0.70:18080/flux2/gateway/ | 21s 172ms | 2ms | 7,199 | 13ms |
| http://int0046.bbn.integral.com:18080/apps/humanresources/Menu.cfm | 19s 391ms | 0ms | 21,791 | 6ms |
| http://int0046.bbn.integral.com:18080/apps/twitter/index.cfm | 16s 967ms | 1ms | 10,219 | 108ms |
| http://int0046.bbn.integral.com:18080/apps/justFunction.cfc | 2s 876ms | 9ms | 309 | 20s 14ms |
| http://int0046.bbn.integral.com:18080/apps/Consume/justFunction.cfm | 2s 630ms | 8ms | 297 | 20s 51ms |
| http://localhost:8301/CFIDE/administrator/sha1.js | 2s 430ms | 1ms | 1,724 | 2ms |
| http://localhost:8301/CFIDE/administrator/images/loginbackground.jpg | 1s 410ms | 0ms | 1,724 | 1ms |
| http://localhost:8301/CFIDE/administrator/images/adobe/logo.gif | 1s 65ms | 0ms | 1,724 | 1ms |
| http://localhost:8301/CFIDE/administrator/images/spacer.gif | 360ms | 0ms | 1,724 | 0ms |



Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Highest CPU Usage i It shows the First 100 Requests for the entire cluster, ordered by the total amount of time they have spent using the CPU. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

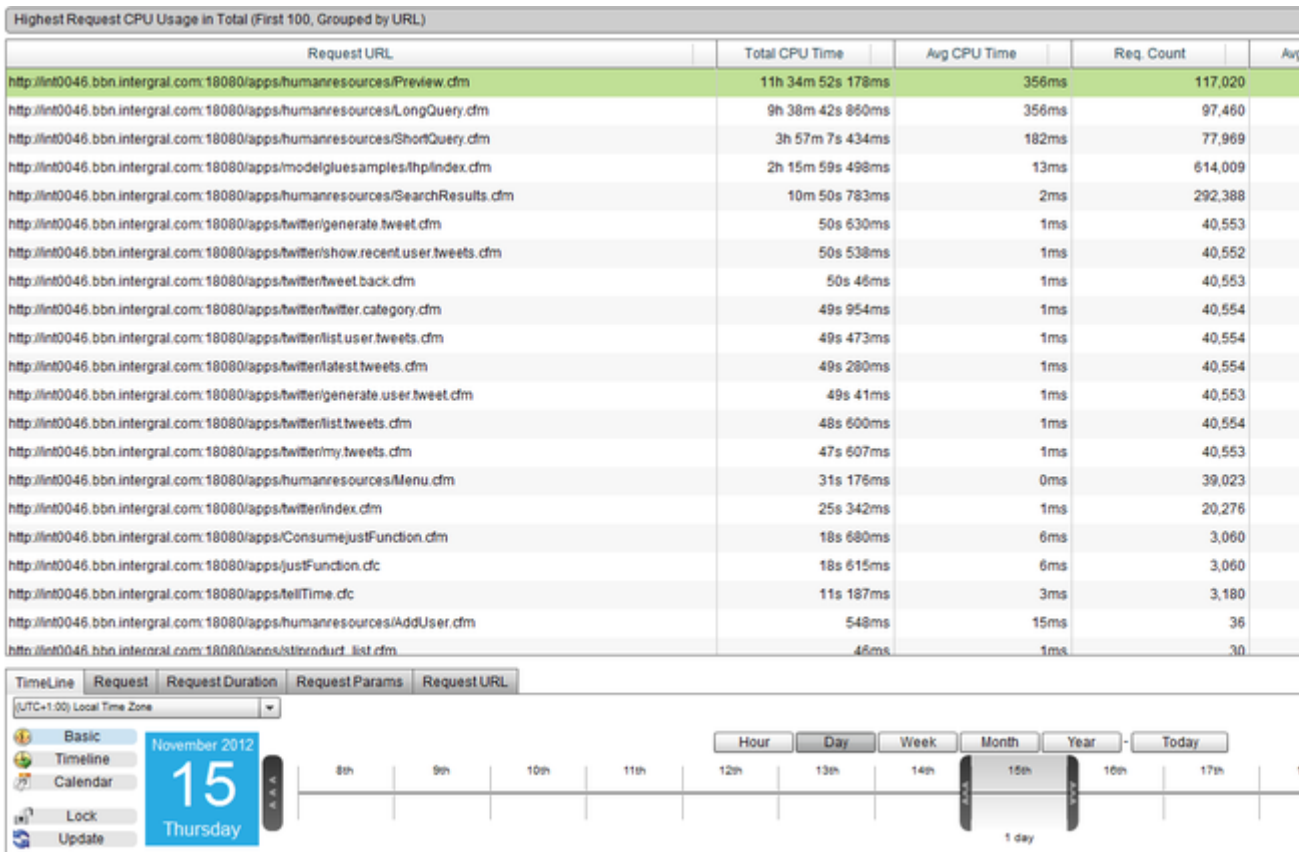


Figure 2: Cluster version of the Highest CPU Usage in Total perspective

Back to the top

MOST POPULAR REQUESTS

Most Popular Requests

Description


The most popular requests perspective shows which requests ran the most against the server within a given time frame. The information is which requests are run the most. The most popular requests are good candidates for optimization on both request size and execution time. Execution time of the most popular pages may result in a large improvement in the user response time. It also can result in lower resource server is able to handle.

Usage

To change the order in which the requests are shown in the table, click on one of the table column headings. Use the Date Navigation in perspective covers.

You can right click any of the requests to bring up a context menu. From there, you can navigate to an Overview for that request for differ

Perspective View



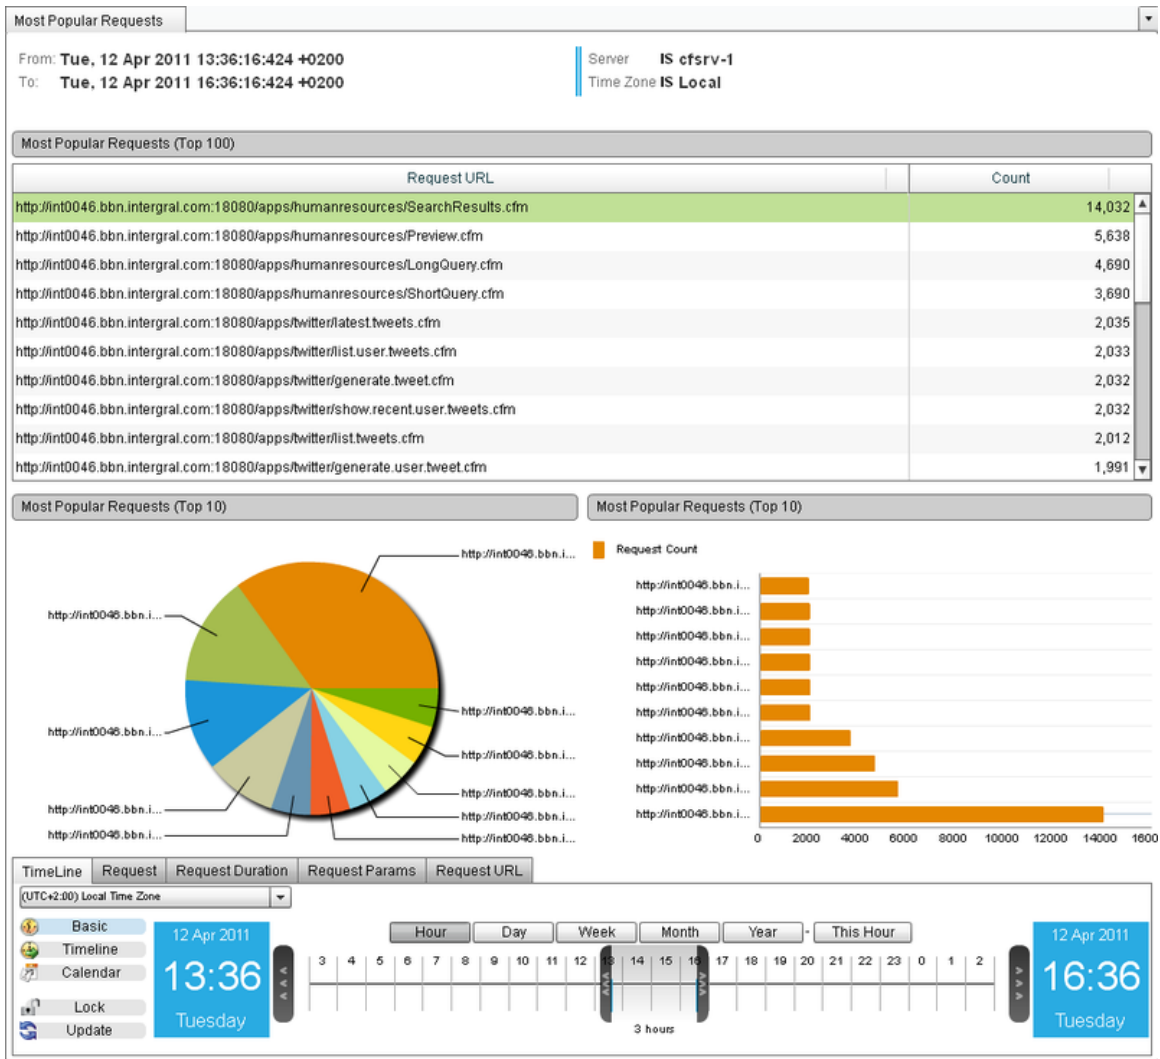


Figure 1: Most Popular Requests Perspective

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Most Popular Reques It shows which requests ran the most against the whole cluster within a given time frame. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

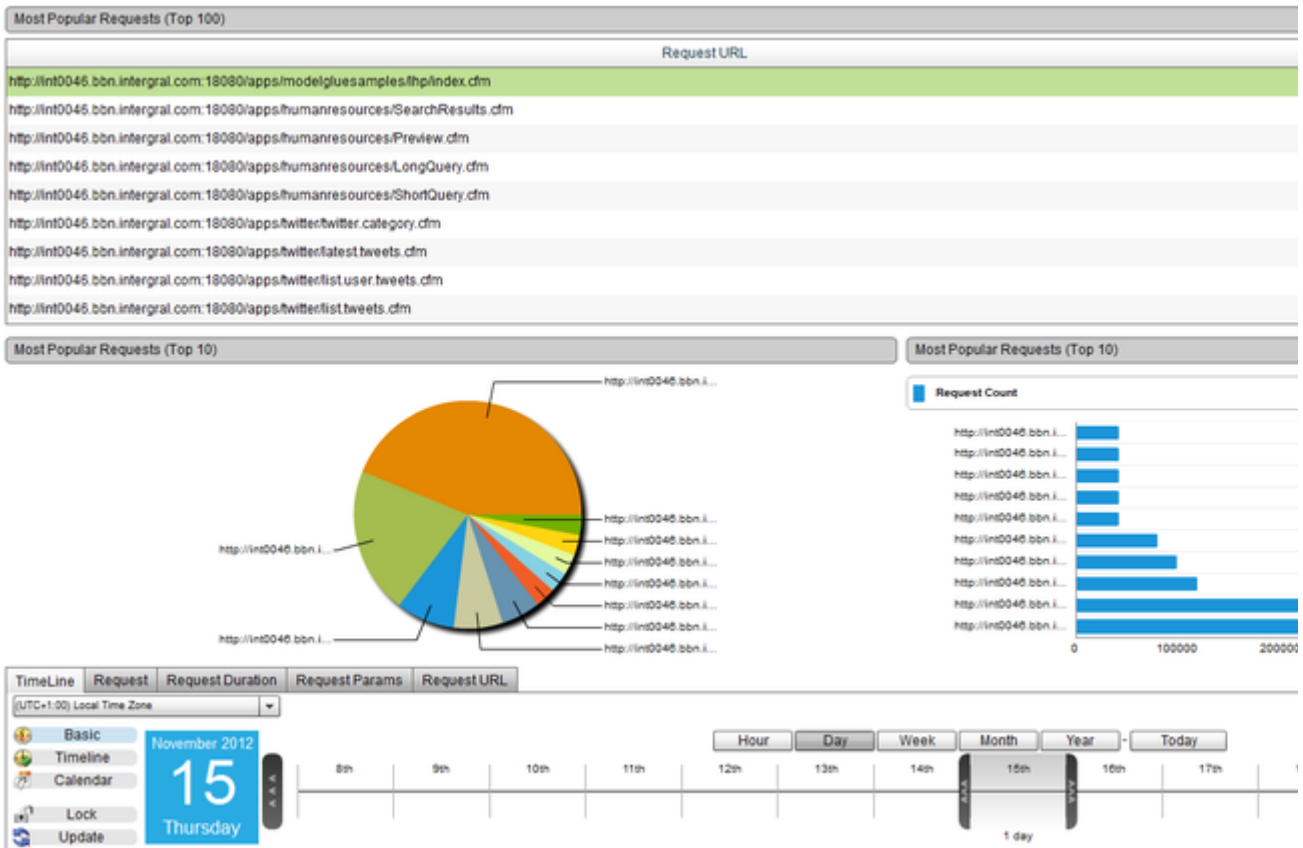


Figure 2: Cluster version of the Most Popular Requests perspective

Back to the top

MOST DB QUERIES ON AVERAGE

Most DB Queries On Average


Description

This perspective shows a table of the top 100 Requests by Average Query Count (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 covers.

Perspective View



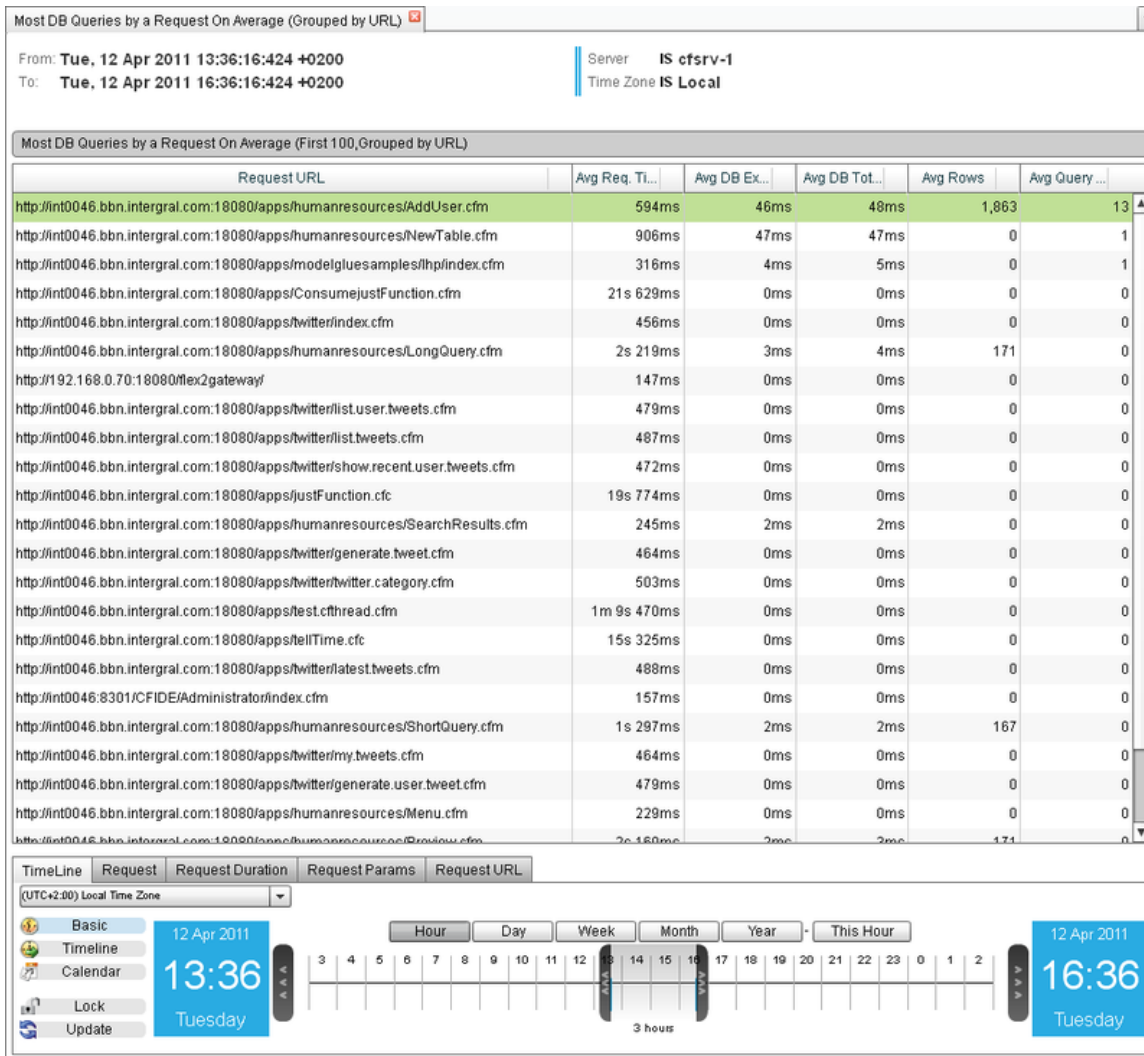


Figure 1: Most DB Queries On Average Perspective

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Most DB Queries On / It shows the First 100 Requests for the entire cluster, ordered by the Average Query Count. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster

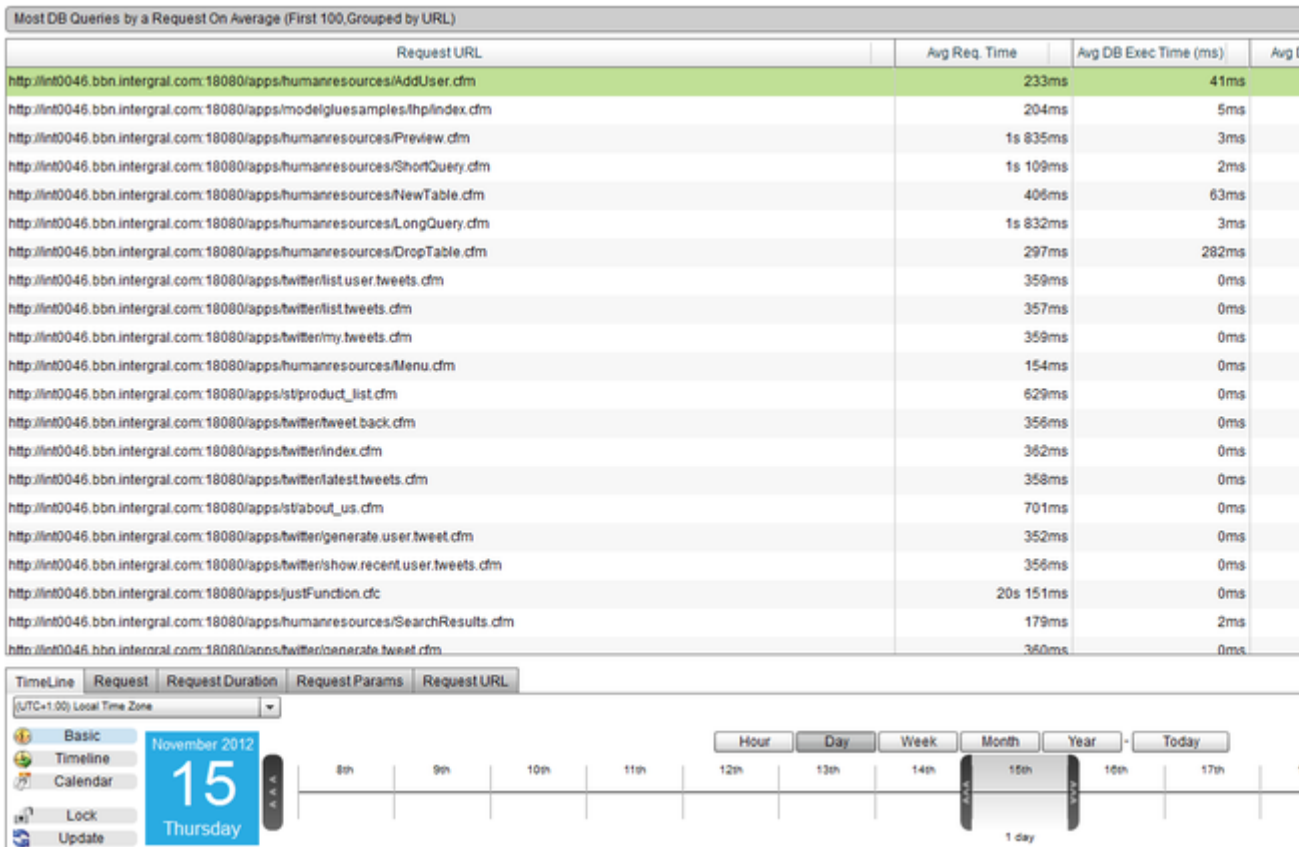


Figure 2: Cluster version of the Most DB Queries On Average perspective

Back to the top

MOST DB QUERIES

Most DB Queries

Description

This perspective shows a table of the top 100 Requests by Query Count 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 covers.

You can right click on any of the requests in the table to bring up a context menu. This will allow you to navigate to other useful perspectives Requests and System Metrics.

You can also view Request Details, to see specific details of that request, such as the exact end time, the amount of memory it used, and

Perspective View



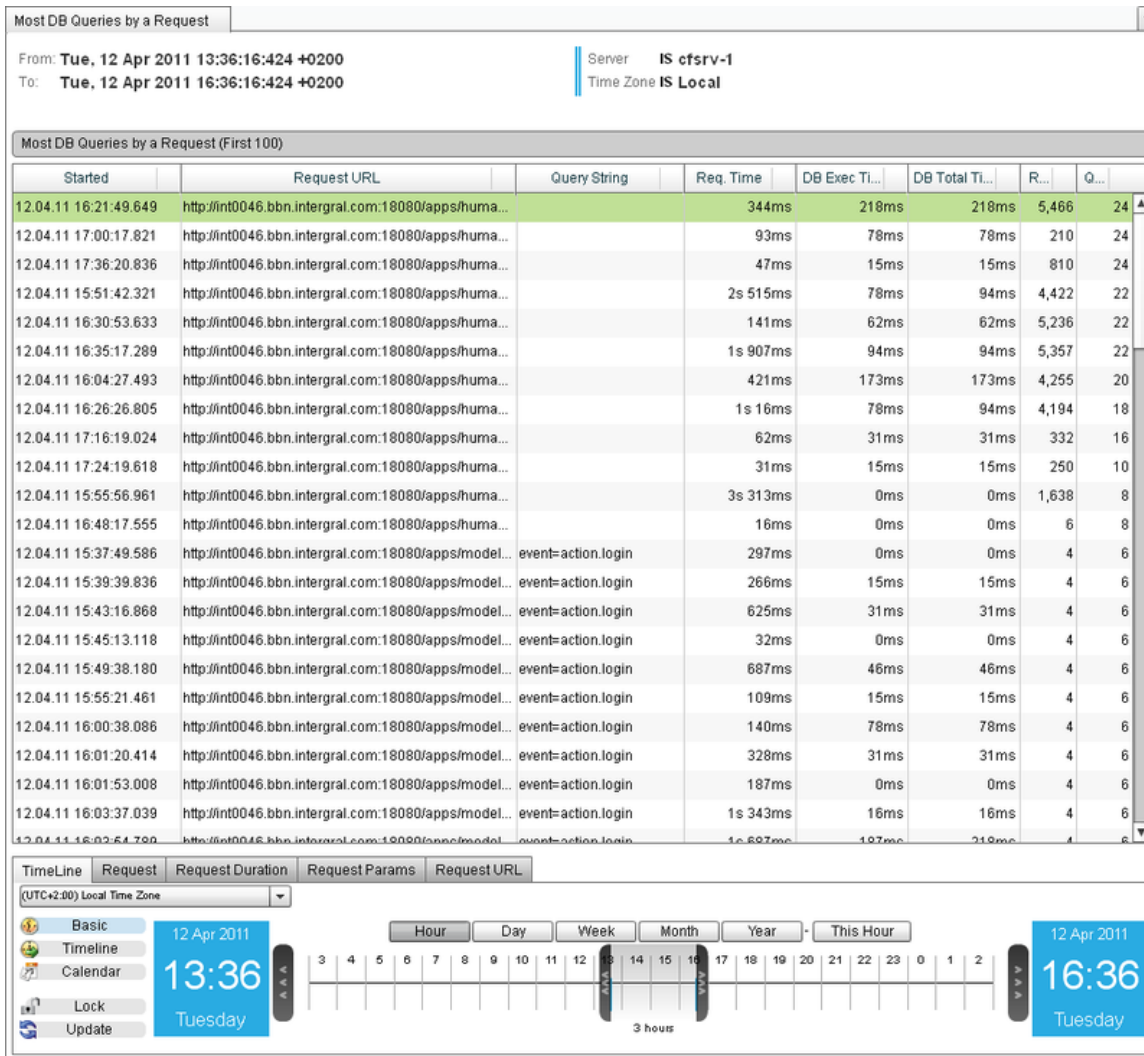


Figure 1: Most DB Queries Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Most DB Queries for It shows the top 100 requests from the entire cluster, by the Query Count of the requests. You can see which server each request came from in the leftmost column of the table, which displays Server Name.

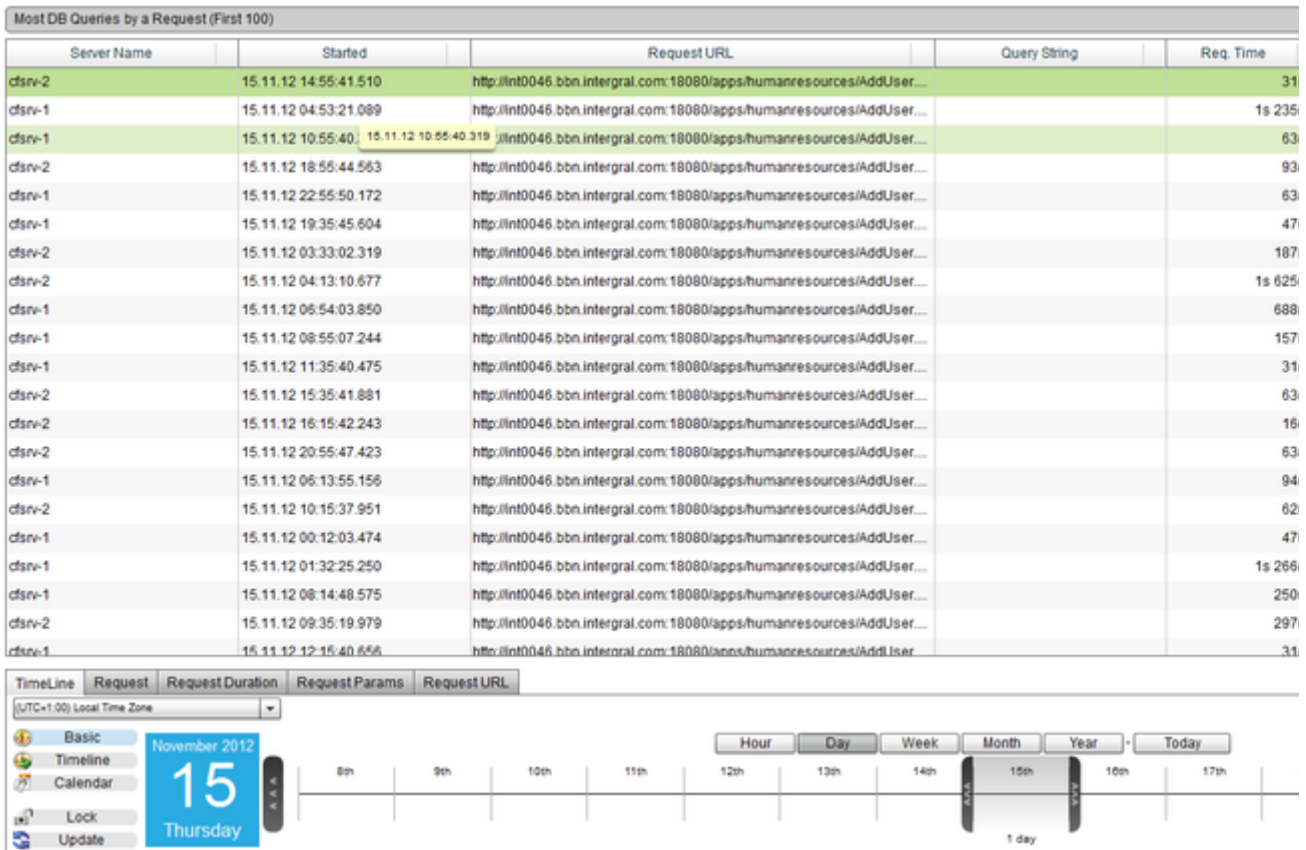


Figure 2: Cluster version of the Most DB Queries perspective

Back to the top

MOST DATA RETURNED (BYTES)

Most Data Returned (bytes)

Description


Largest Size (bytes) is a report to help you identify which request returned the most bytes in its response to the client. Generally the more content (bytes) to be returned are buffered or stored in memory until a request is complete, which can cause large quantities of flush data to the client each time the buffered response reaches a certain size (configurable in many cases) but often the request processing response being sent. This can cause performance degradation in requests and thread starvation. It can also be used to create a Denial of Service by 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 of data. This will typically increase application memory, CPU and network traffic.

Usage

- Discover which requests return the most bytes to the client
- Identify unexpected results - requests with poorly coded
- Use the information to find bugs in requests that cause too 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



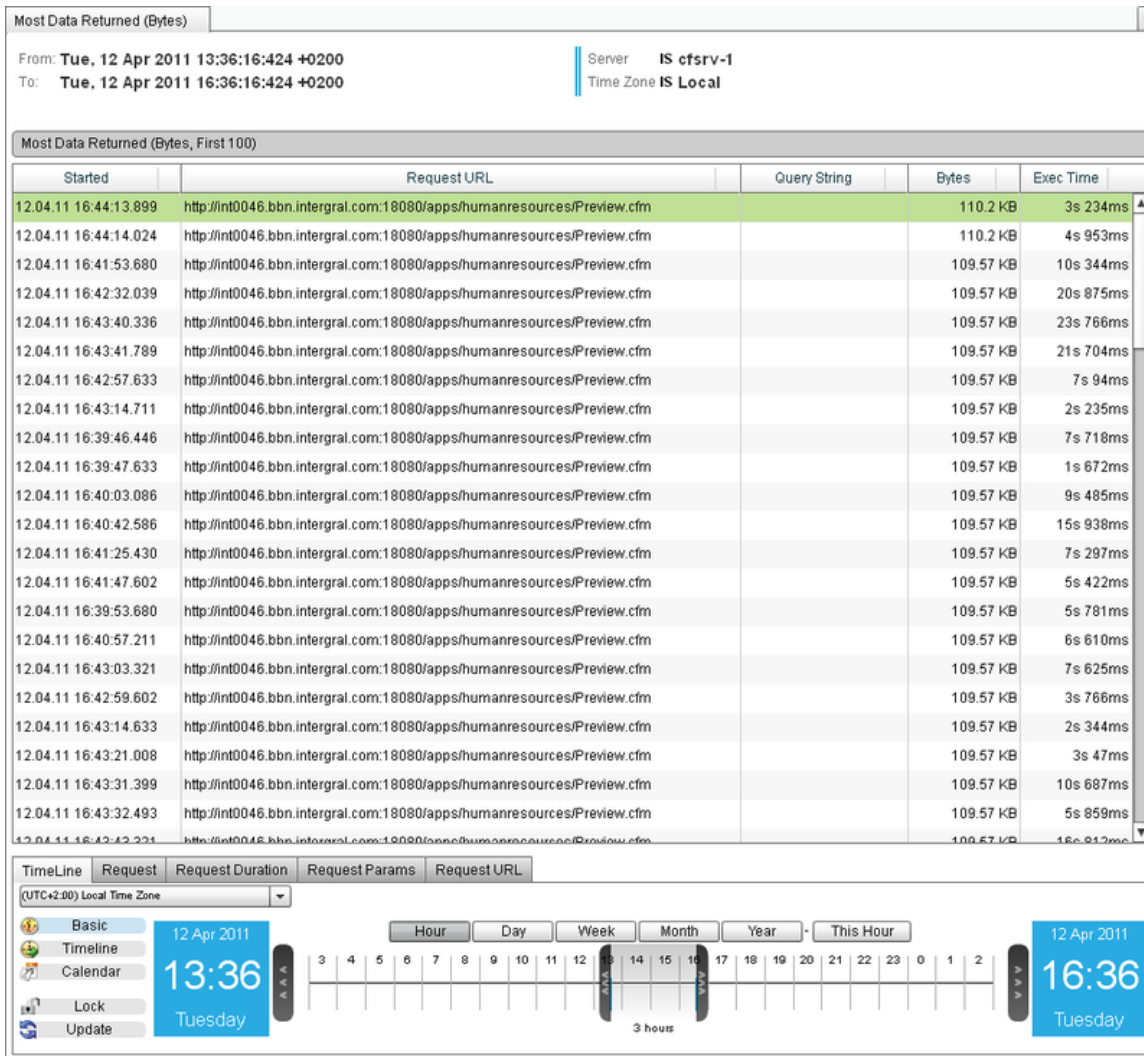


Figure 1: Most Data Returned (bytes) Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Most Data Returned (It shows the top 100 requests from the entire cluster, that returned the most bytes in it's response to the client. You can see which server each request came from in the leftmost column of the table, which displays Server Name.

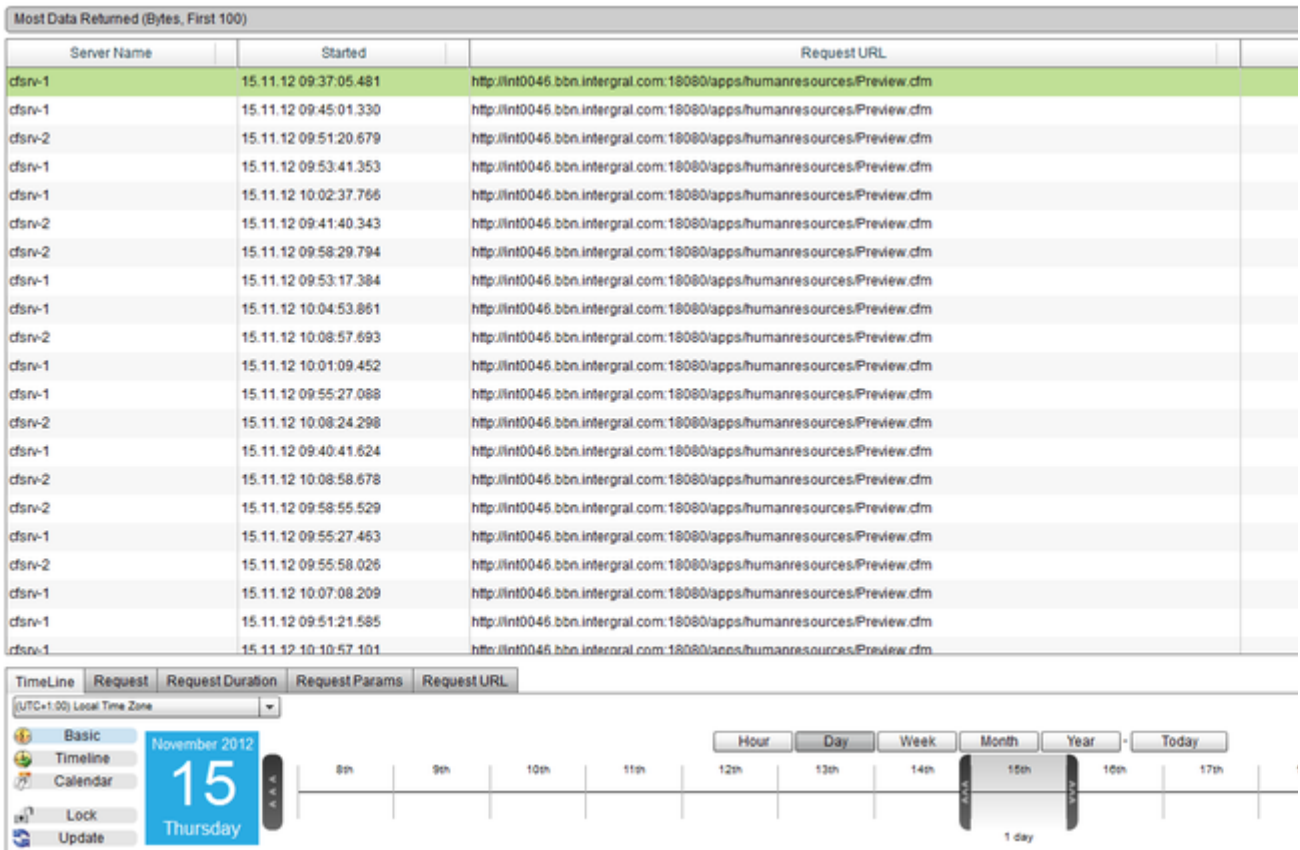


Figure 2: Cluster version of the Most Data Returned (bytes) perspective

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 covers.

Perspective View



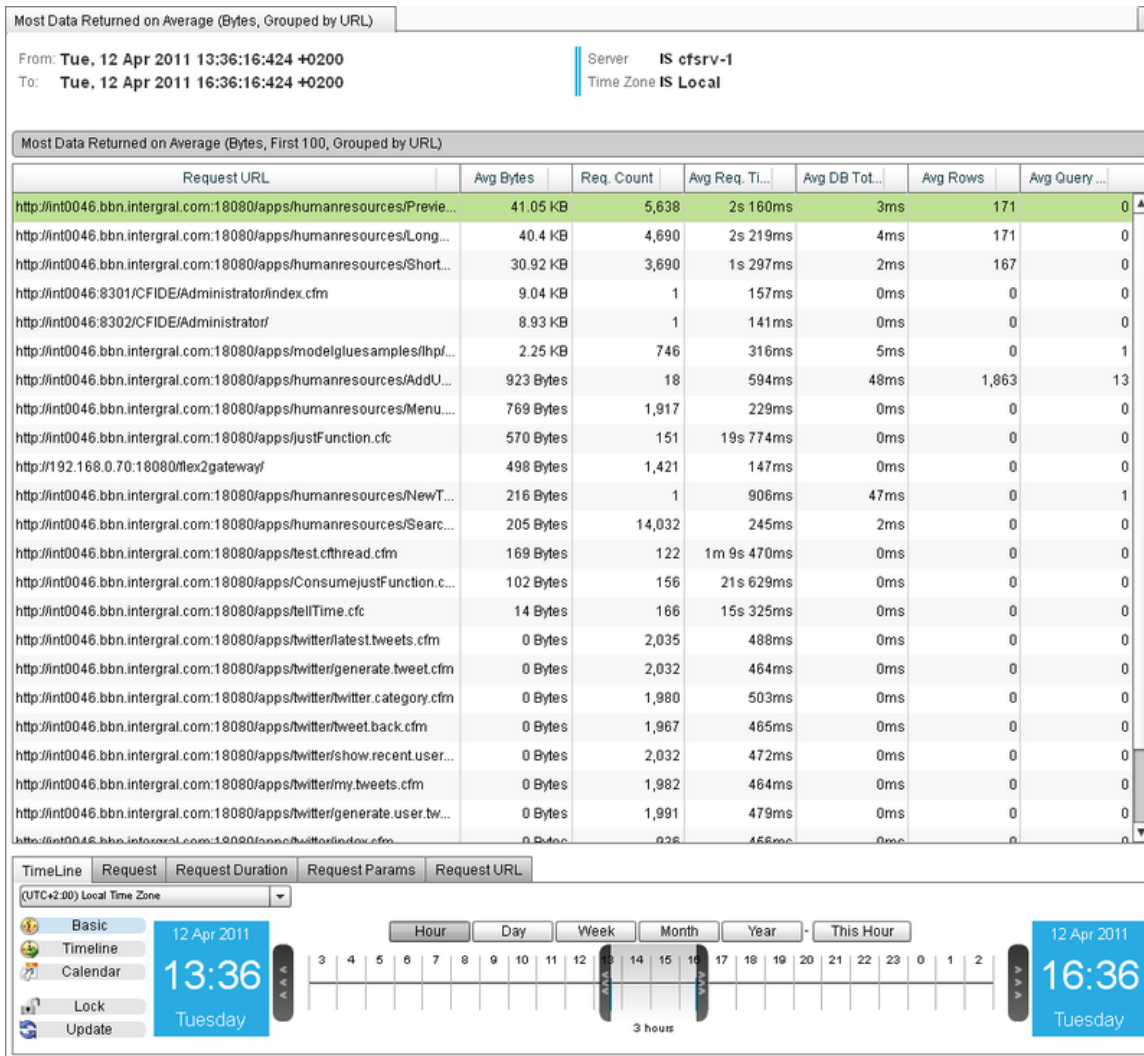


Figure 1: Most Data Returned on Average Perspective

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Most Data Returned c rid. It shows the First 100 Requests (grouped by URL) for the entire cluster, ordered by the Average Number of Bytes returned by the request. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

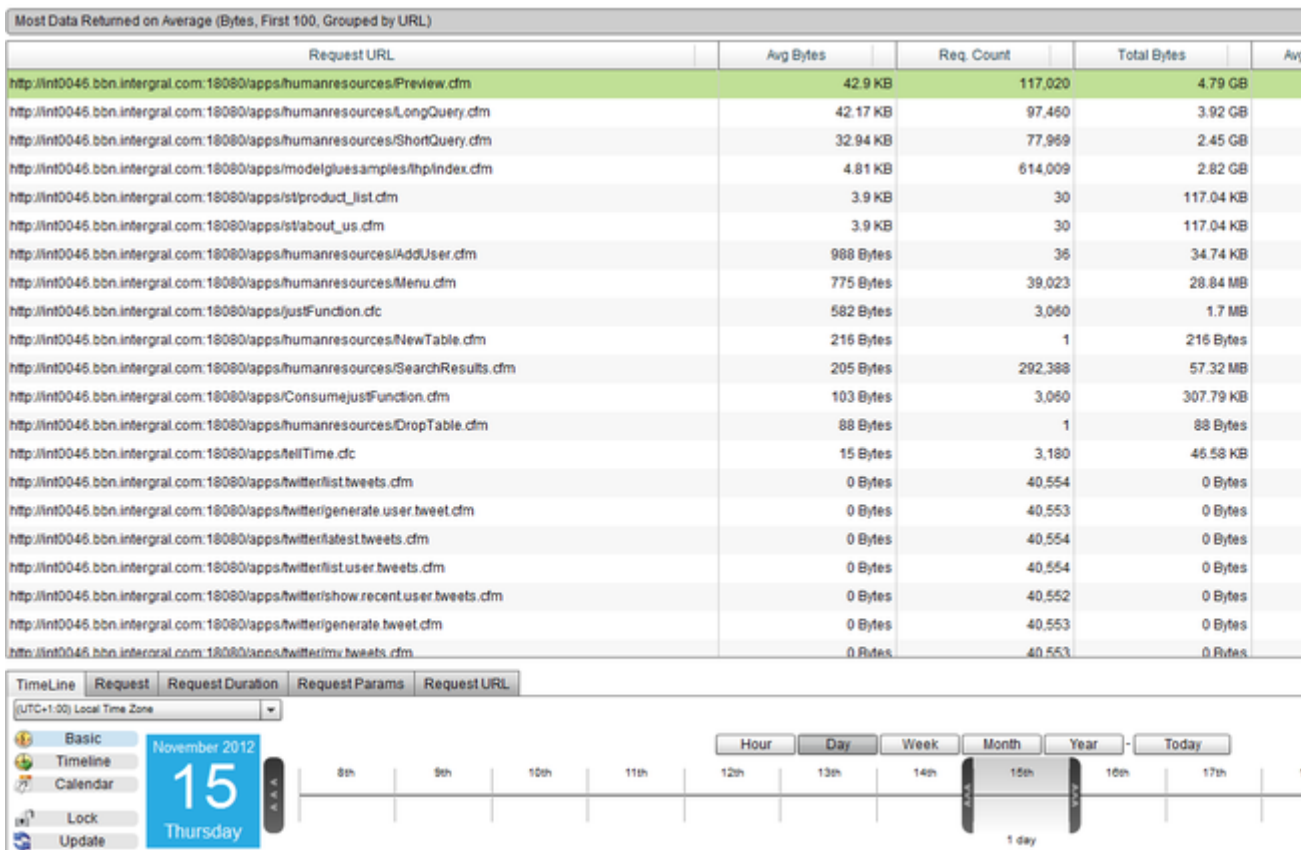


Figure 2: Cluster version of the Most Data Returned on Average (bytes) perspective

Back to the top

TABLE OF REQUESTS

Table of Requests


Description

This perspective shows a table of the top 1000 Requests, ordered by Start Time, 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 char covers.

Perspective View



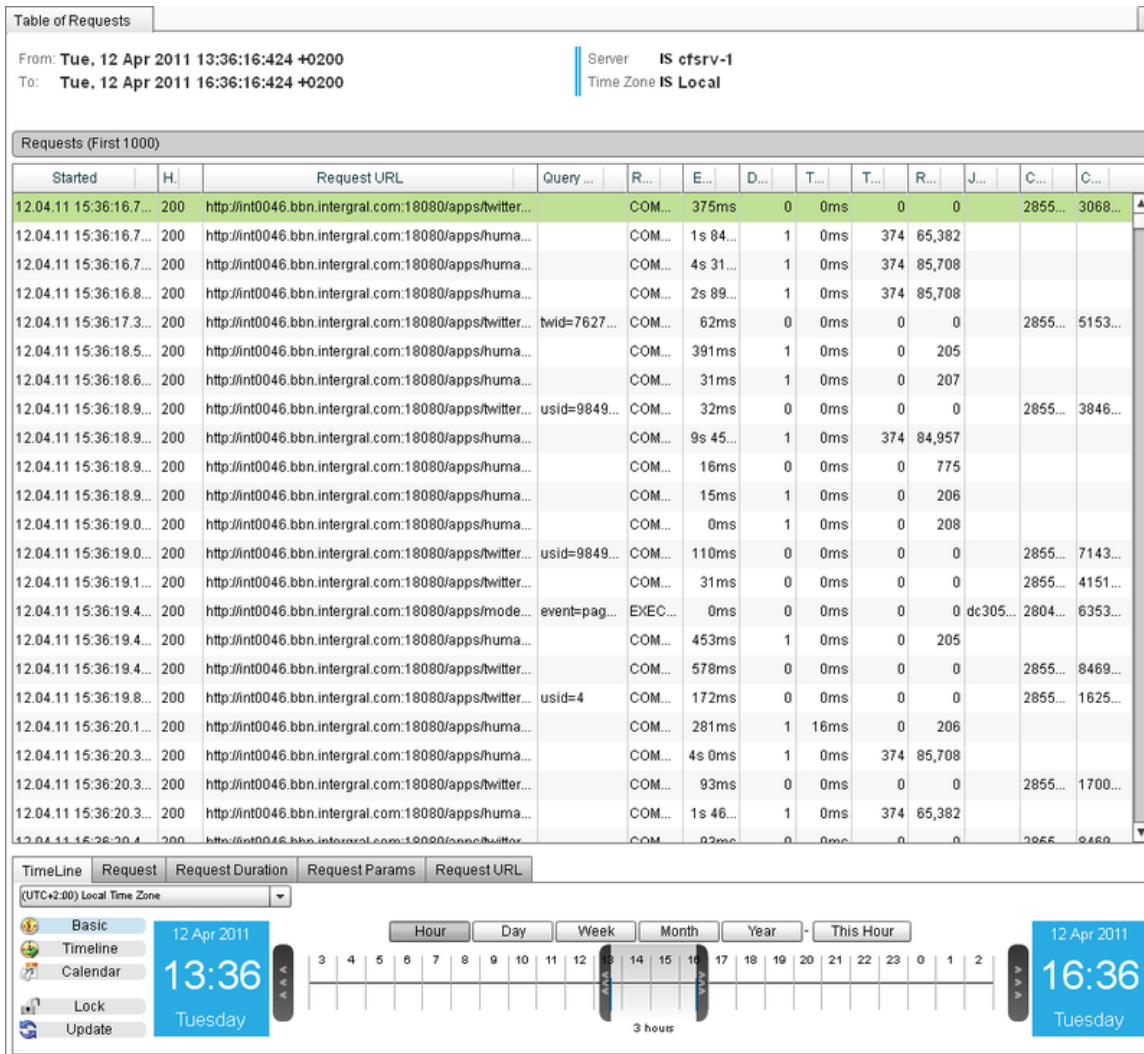


Figure 1: Table of Requests Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Table Of Requests frc It shows the top 1000 requests from the entire cluster, ordered by Start Time. You can see which server each request came from in the leftmost column of the table, which displays Server Name.

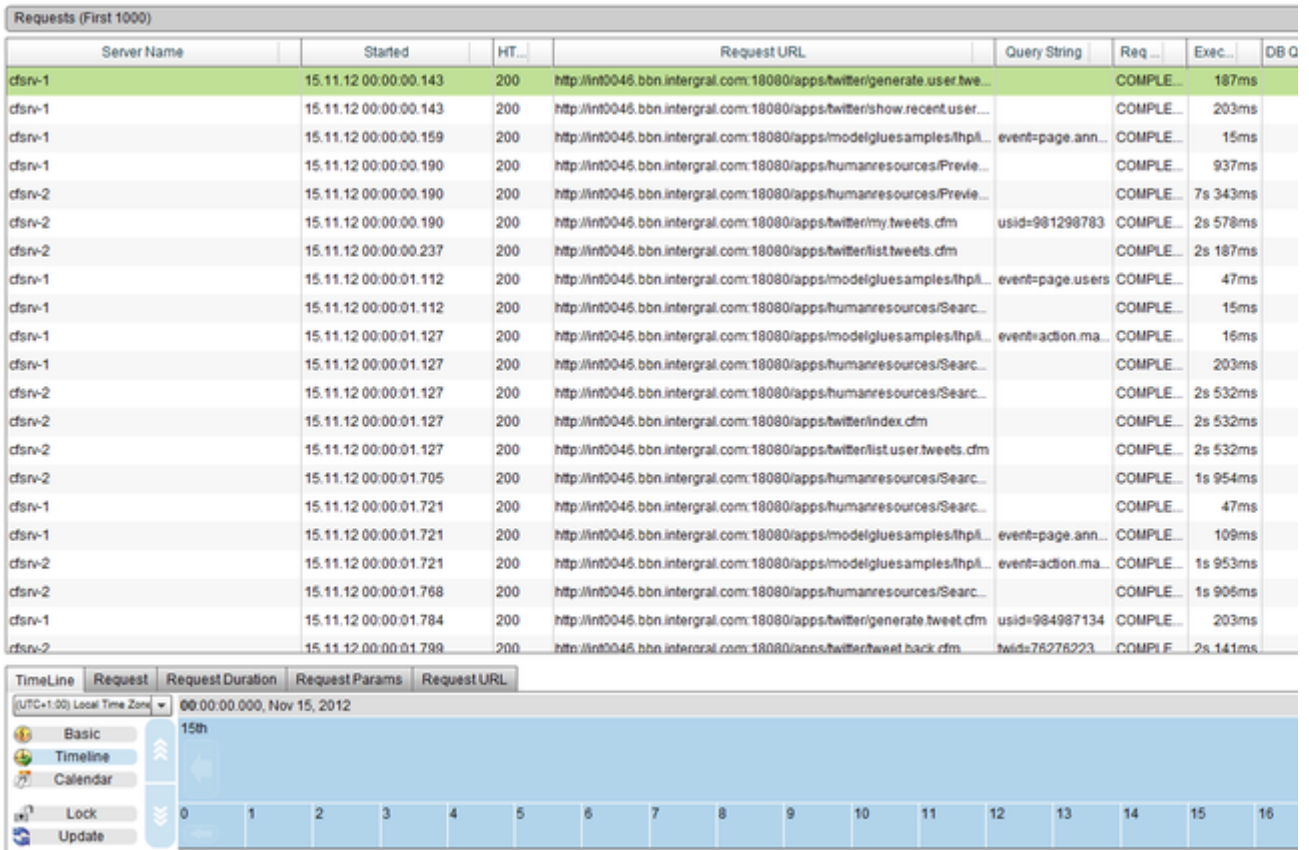


Figure 2: Cluster version of the Table Of Requests perspective

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

Use the Date Navigation in order to change the period of time over which this perspective covers. You can right click on any of the requests to bring up a context menu. This will allow you to navigate to other useful perspectives for the server and System Metrics. You can also view Request Details, to see specific details of that request, such as the exact end time, the amount of memory it used, and ran.

Perspective View



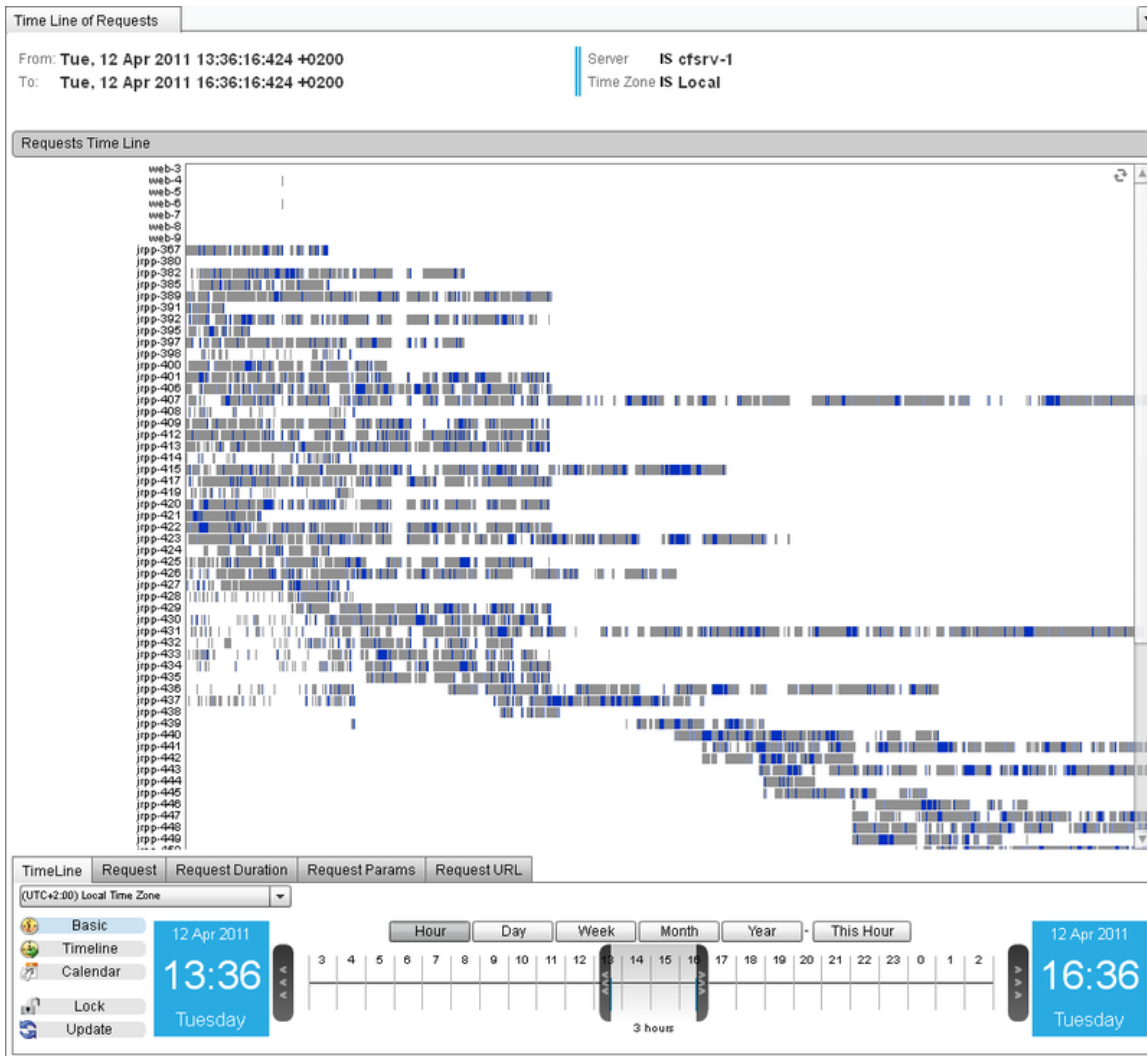


Figure 1: Time Line of Requests Perspective

Back to the top

TIME LINE OF REQUESTS WITH RESOURCES

Time Line of Requests with Resources

Description

This perspective shows a Time Block Chart of Requests, along with a Table of the top 1000 requests, ordered by Start Time, as well as the Server, all within the specific period of time.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. You can right click on any of the requests in the table or the timeline to bring up a context menu. This will allow you to navigate to other views e.g. JDBC Graphs and System Metrics. You can also view Request Details, to see specific details of that request, such as the exact end time, the amount of memory it used, and ran.

To move between the three different views, click on the heading for the view you want to see e.g. "Requests Time Block Chart" and that view. You can switch between these different perspectives very easily this way.

Perspective View

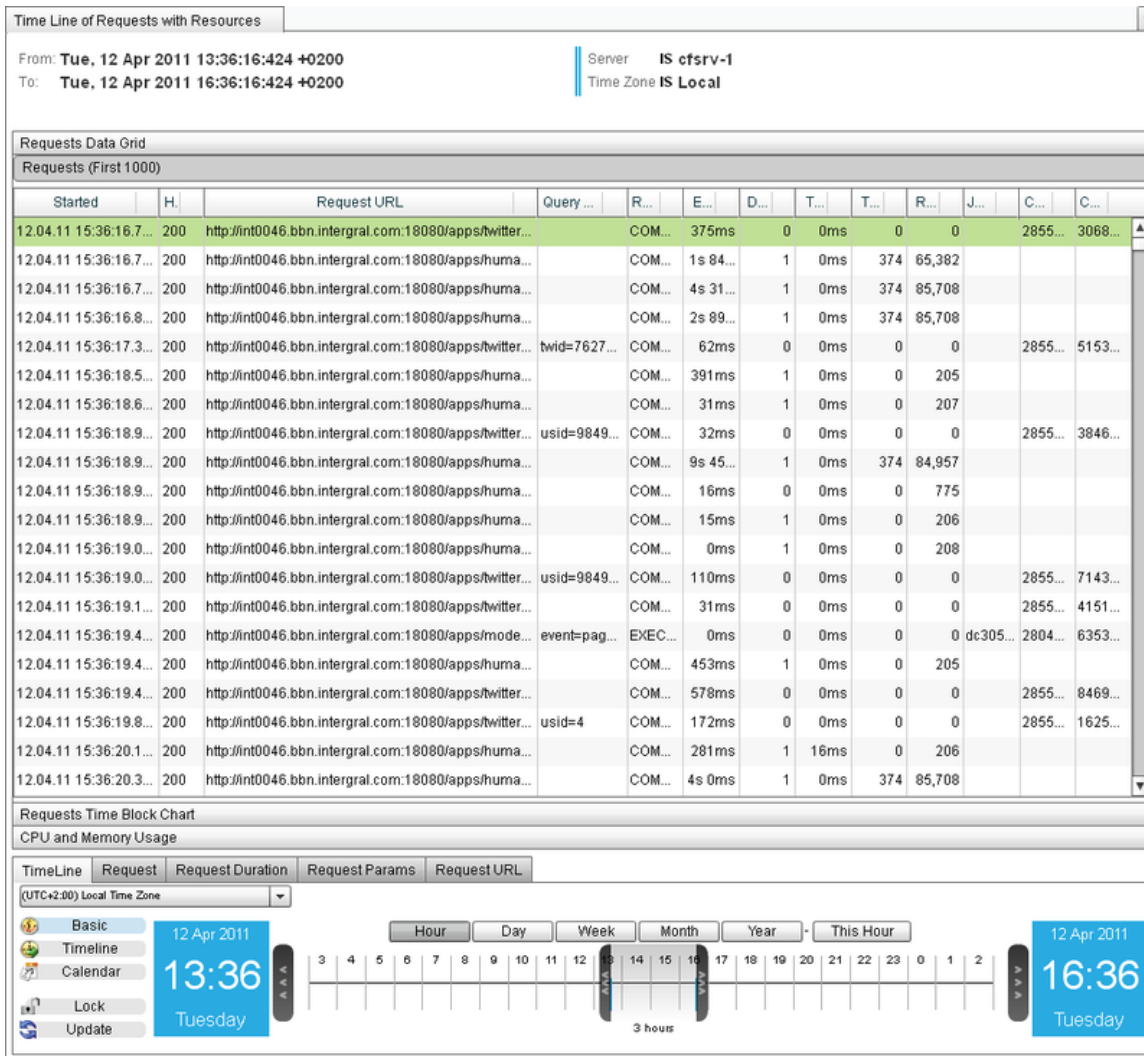


Figure 1: Time Line of Requests with Resources Perspective

Back to the top

TIME LINE OF REQUESTS WITH DB REQUESTS

Time Line of Requests with DB Requests

Description

This perspective shows a Time Block Chart of Requests, alongside another Time Block Chart which shows the JDBC Requests.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. You can right click on any of the requests to bring up a context menu. This will allow you to navigate to other useful perspectives for the se and System Metrics. You can also view Request Details, to see specific details of that request, such as the exact end time, the amount of memory it used, and ran

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

Use the Date Navigation in order to change the period of time over which this perspective covers.

Perspective View



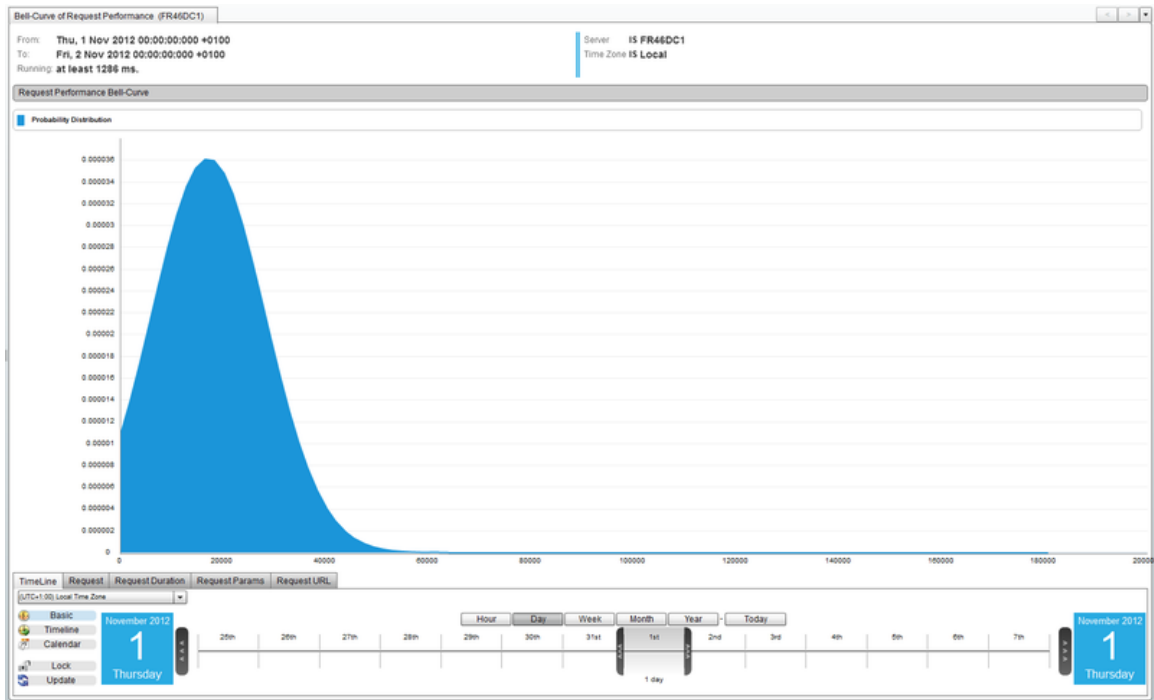


Figure 1: Bell-Curve of Request Performance Perspective

Cluster Perspective

The cluster version of this perspective takes request data from across the cluster to form a Bell Curve for the probability distribution of requests. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

Back to the top

REQUEST STATUS CODE BRAKEDOWN

Request Status Code Brakedown

Description

This perspective shows a breakdown of the Status Codes returned by all the requests that were executed during the specified time period to help visualize which status codes are occurring the most.

The Pie Chart and the Bar Chart are different ways to display the count of the status codes, and the table shows more information such as the time taken to execute that Status Code.

Usage

To change the order in which the requests are shown in the table, click on one of the table column headings. Use the Date Navigation in this perspective covers.

You can right click any of the requests to bring up a context menu. From there, you can click "Show Requests", to show a Table of the First returned that Status Code.

Perspective View

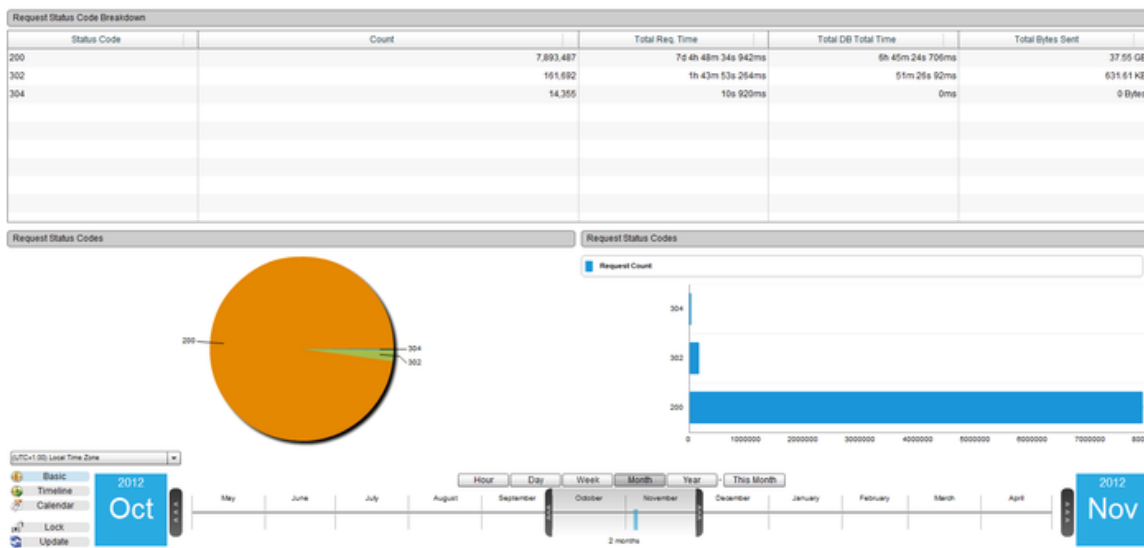


Figure 1: Request Status Code Brakedown Perspective

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Most Popular Reques It shows a breakdown of the Status Codes returned by all the requests for the whole cluster within a given time frame. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

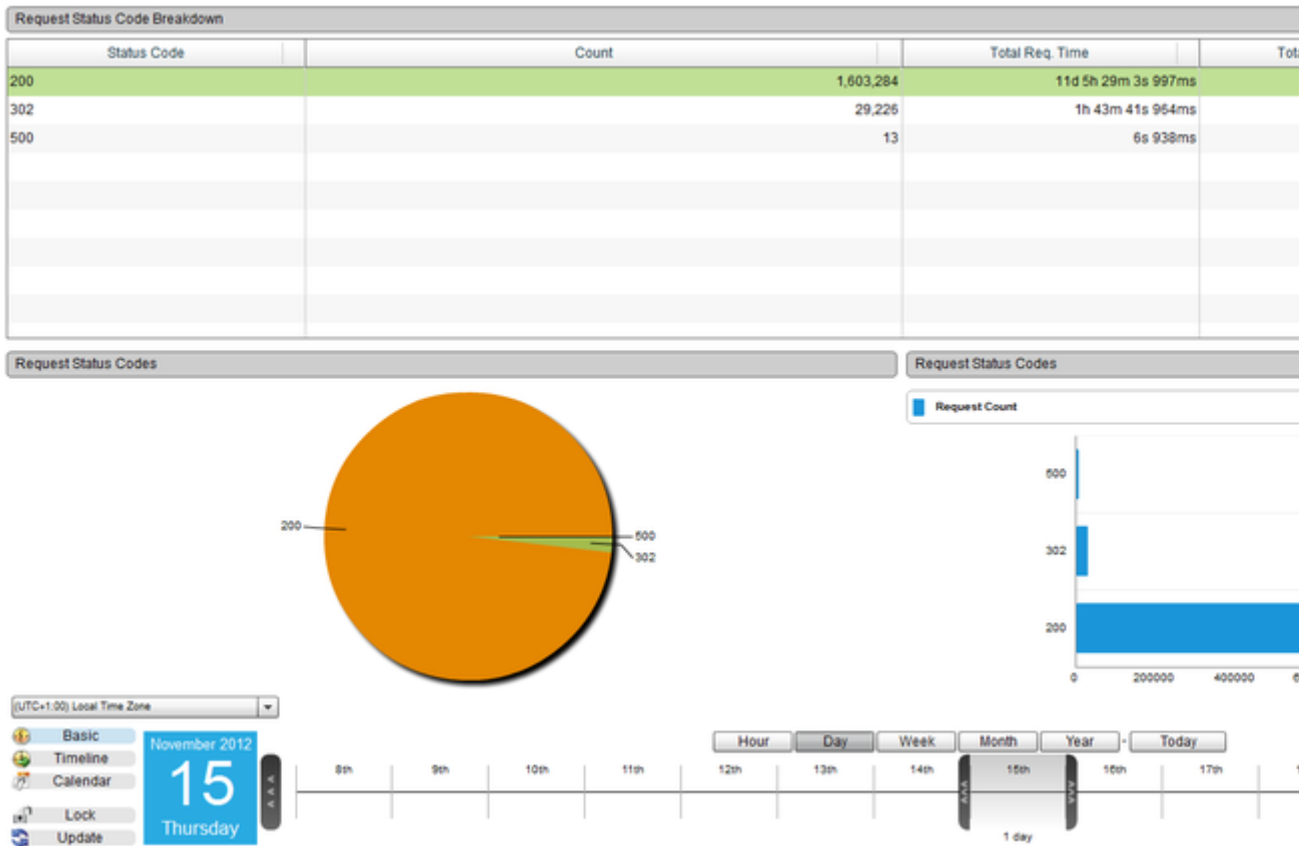


Figure 2: Cluster version of the Request Status Code Brakedown perspective

Back to the top

Request and JDBC Execution Time

Description

This perspective shows an area chart of the Average Request Execution Time and the Average Execution Time of Database Requests, wi

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers.

You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining :

Perspective View

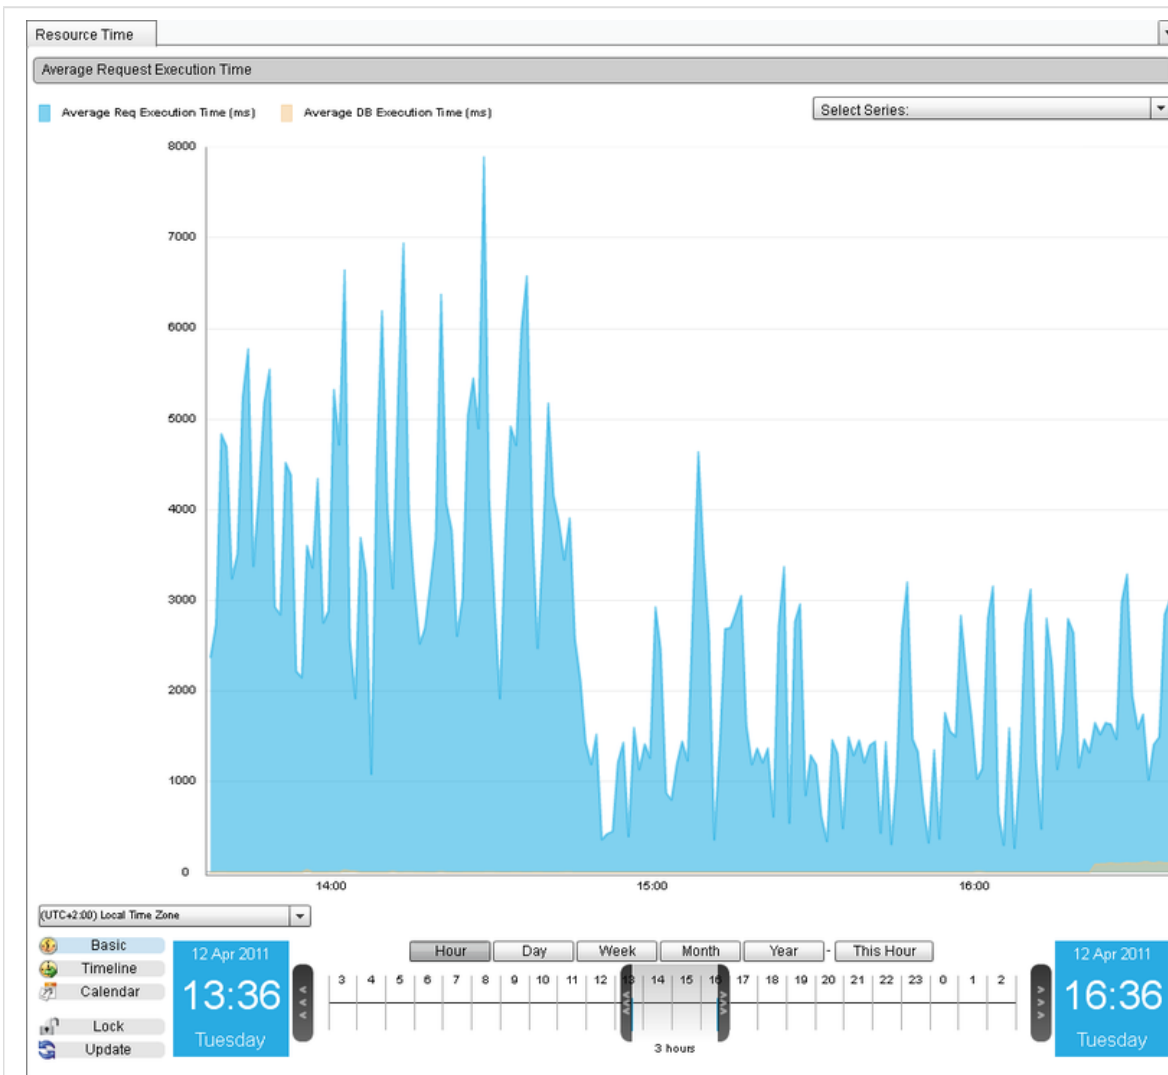


Figure 1: Request and JDBC Execution Time Perspective

Cluster Perspective

The cluster version of this perspective now has two different views that can be switched between using the drop down box in the uppermo. The first view is a Server Comparison graph view showing two separate graphs. The first graph showing the Average Request Execution. The second graph shows the Average DB Execution Time for each individual server within the cluster.

The second view is a Cluster View showing one Area graph, which displays the Average Request Exec Time and the Average DB Exec Ti




Figure 2: Server Comparison View for Request and JDBC Execution Time perspective




Figure 3: Cluster View for Request and JDBC Execution Time perspective

[Back to the top](#)

TOTAL REQUEST PERFORMANCE

Total Request Performance

Description


This perspective has four different views that you can switch between using the drop down box in the uppermost right hand corner of the page. The four different views are:

- **Totals** - This shows two area charts. The first shows the total number of requests that were started and finished during the specific time period, and the second shows the total number of Database requests that were finished during the same time period.
- **Load** - This shows two area charts. The first shows the number of requests that were running at the time, alongside the number of requests that were completed at the time. The second chart shows the number of Database requests that were completed at the time.
- **Average** - This shows a chart that displays the Average Request Execution Time alongside the Average Execution Time of Database requests.
- **Per(sec)** - This shows an area chart that displays the Number of Requests Per Second and the number of Database Requests Per Second.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to view the remaining series.

Perspective View



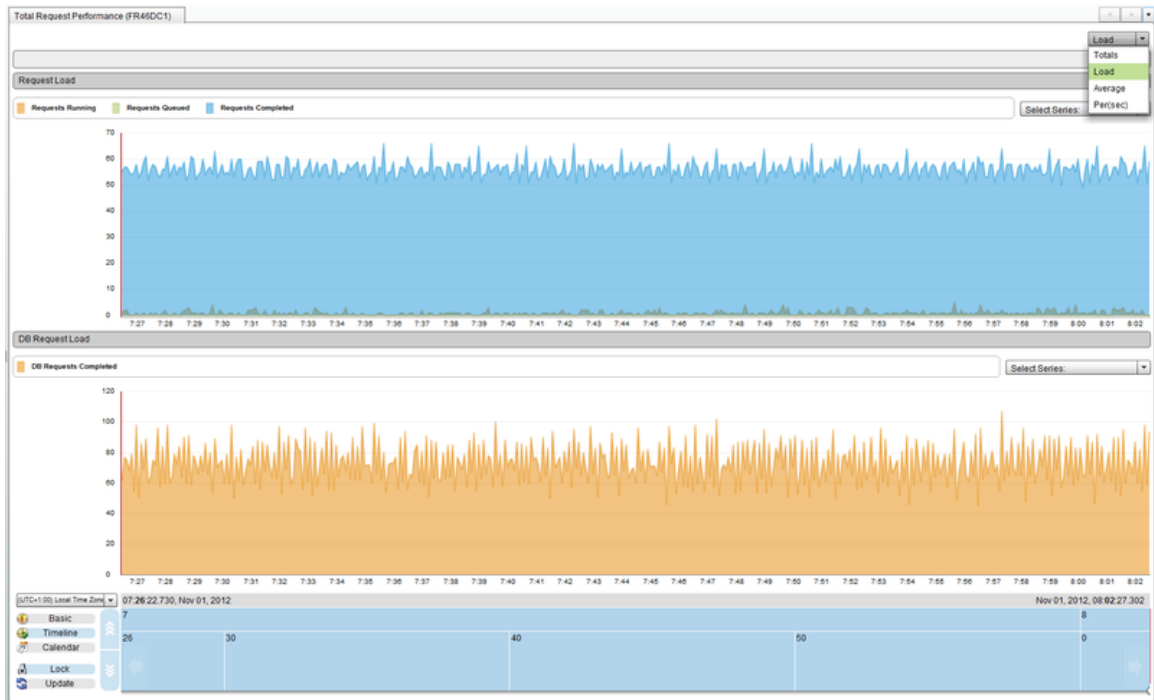


Figure 1: Total Request Performance

Cluster Perspective

The cluster version of this perspective now has multiple different views that can be switched between using the drop down box in the upper left. Each of the views present in the Server version of the perspective have each been given a Cluster View and a Server Comparison for the




Figure 2: Example Server Comparison view from Total Request Performance perspective

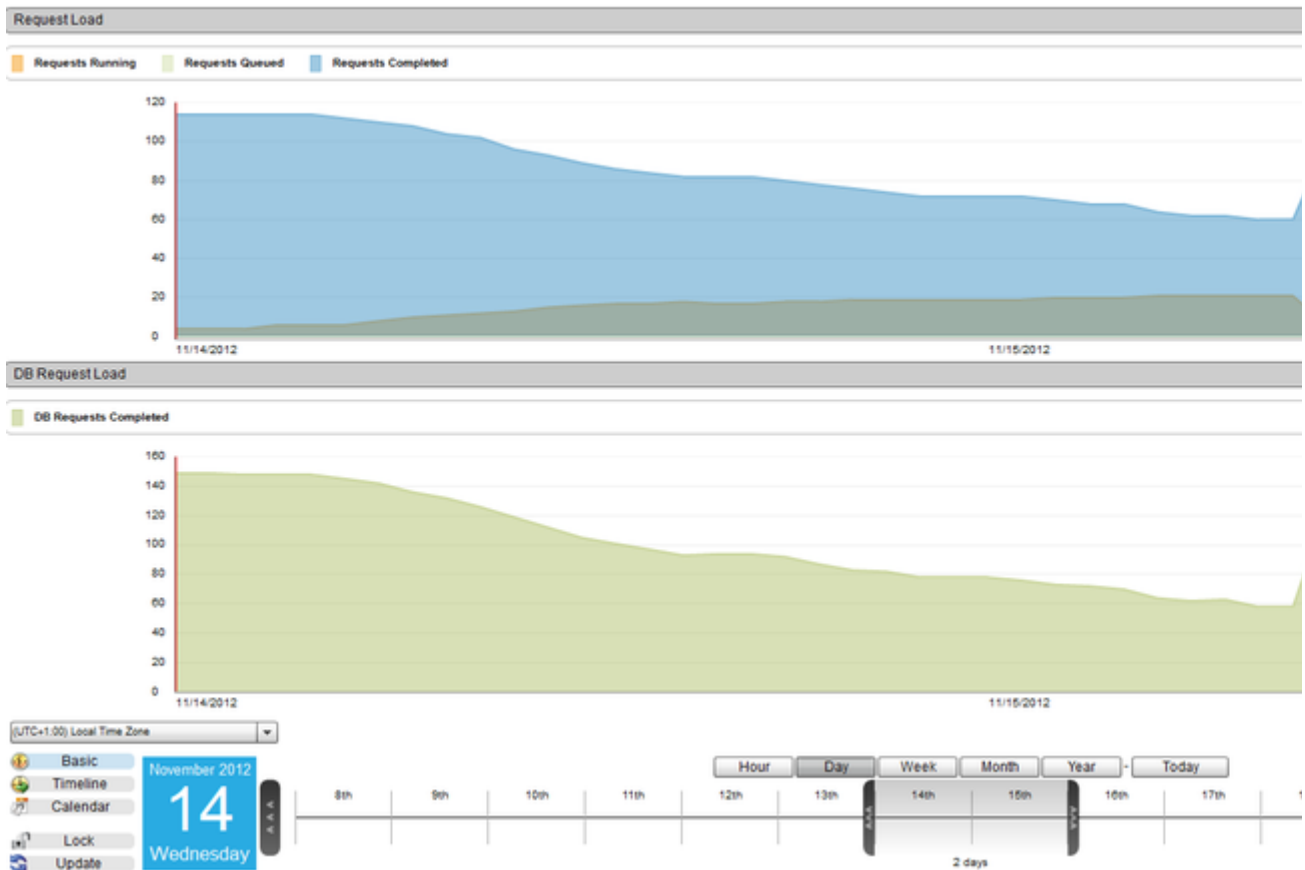


Figure 2: Example Cluster View from Total Request Performance perspective

[Back to the top](#)

MOST HITS BY IP ADDRESS

Most Hits By IP Address

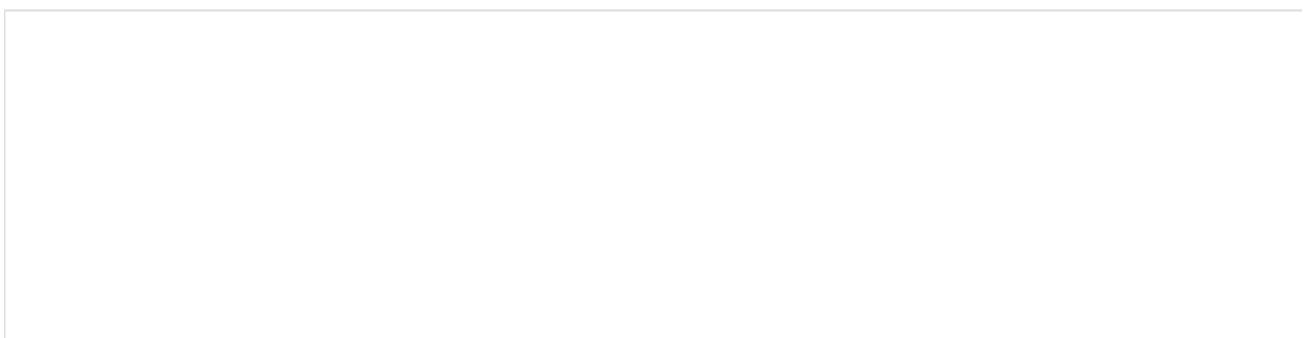
Description

This perspective shows a table of the top 100 IP Addresses making the most hits (within the specific period of time). Using this information most requests; often these may be Spiders, Bots and potentially even Denial of Service attacks. You may consider blocking or throttling re performance issues caused by particularly high traffic from them. It's a good idea to review this information around the periods where a sei

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 chan covers.

Perspective View



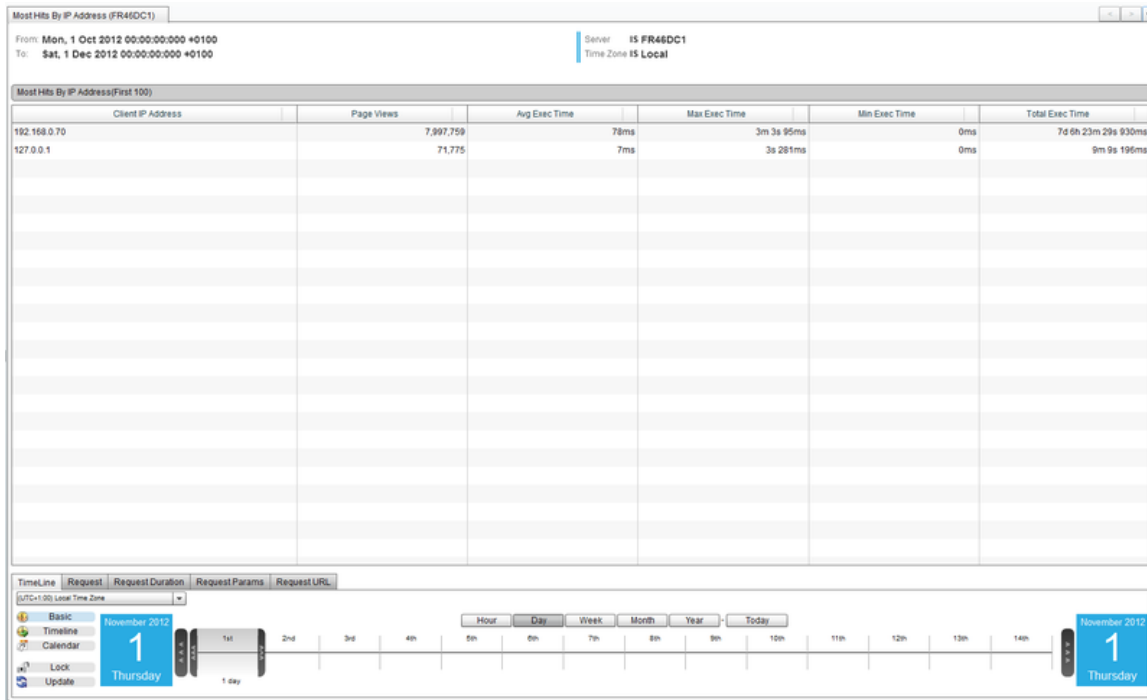


Figure 1: Most Hits By IP Address

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Most Hits By IP Address, shows the top 100 IP addresses making the most hits (within the specific period of time) for the whole cluster. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

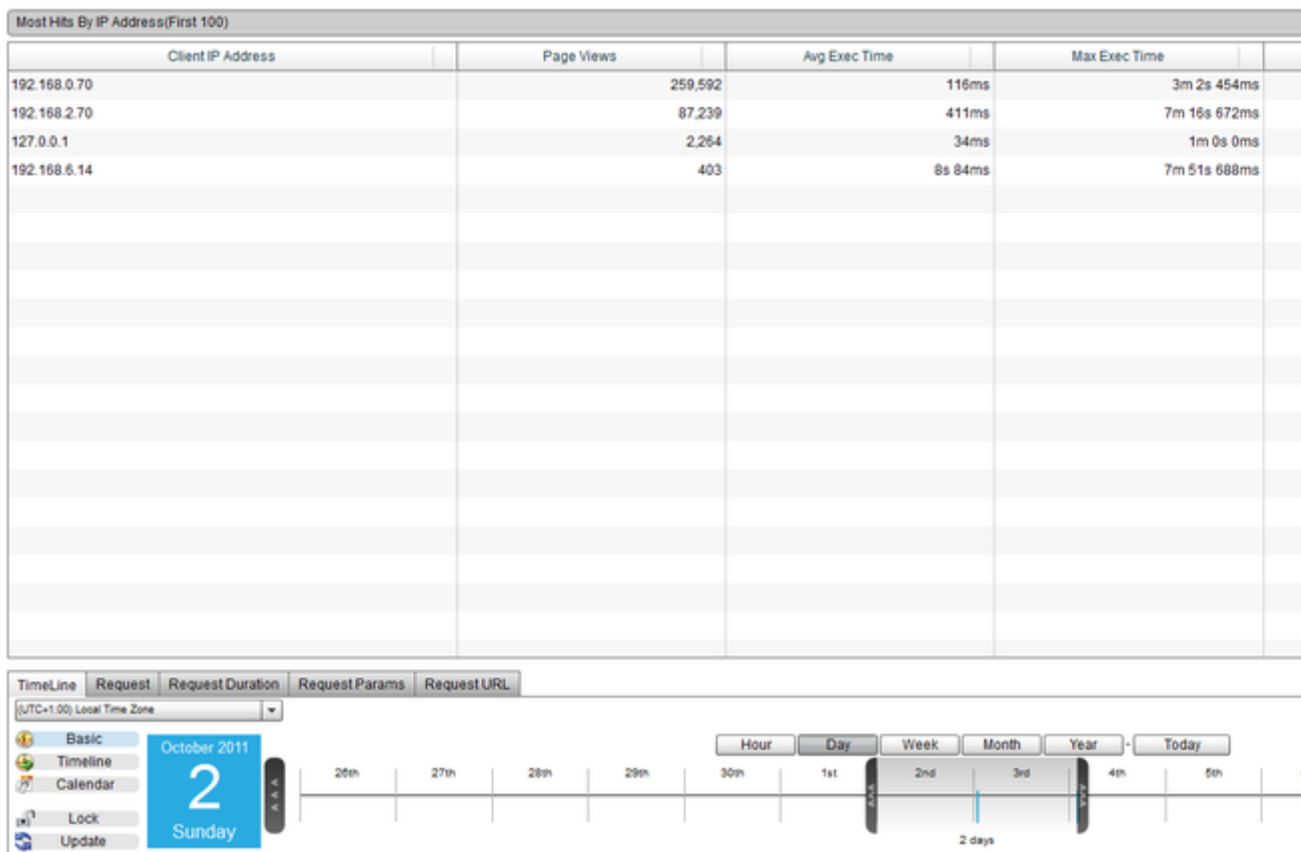


Figure 2: Cluster version of the Most Hits By IP Address perspective

[Back to the top](#)

LONGEST REQUEST EXECUTION TIME BY IP ADDRESS

Longest Request Execution Time By IP Address

Description

This perspective shows a table of the top 100 IP Addresses with the longest total request execution time (within the specific period of time) Addresses spending the most time running requests; often these may be Spiders, Bots and potentially even Denial of Service attacks. You can identify specific IP Addresses if you have performance issues caused by particularly high traffic from them. It's a good idea to review this information.

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 time covers.

Perspective View

| Client IP Address | Page Views | Avg Exec Time | Max Exec Time | Min Exec Time | Total Exec Time |
|-------------------|------------|---------------|---------------|---------------|---------------------|
| 192.168.0.70 | 1,184,856 | 968ms | 3m 2s 812ms | 0ms | 13d 6h 38m 2s 649ms |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |

Figure 1: Longest Request Execution Time By IP Address

Cluster Perspective

The cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Request > Longest Request Execution View Grid. It shows the top 100 IP Addresses with the longest total request execution time (within the specific period of time) for the whole cluster. The perspective will look the same as the Server version, but the data will be calculated for the entire cluster.

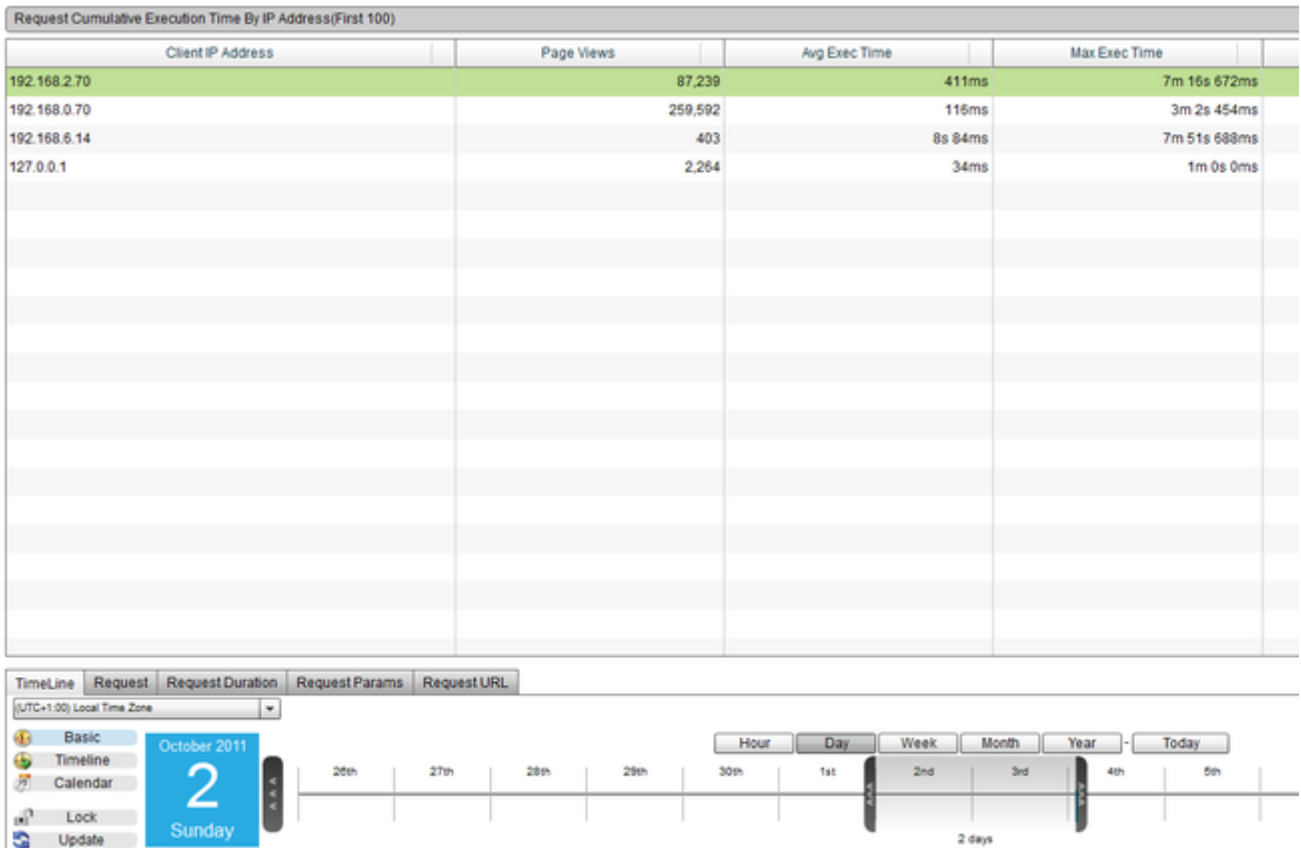


Figure 2: Cluster version of the Longest Request Execution Time By IP Address perspective

[Back to the top](#)

AGGREGATE REQUEST LIST

Aggregate Request List

Description

When viewing a grid that displays aggregate statistics for a group of requests, such as "Slowest Requests on Average" or "Longest Request on Average", you can click on any of the request URL's to bring up the Context Menu. From there you can navigate to the "Request List" perspective. This perspective will show a list of matching requests for the specified time period. The list will be sorted relevant to the original perspective. For example, if viewing "Slowest Requests on Average", the list will be sorted by Execution Time. Or, if viewing "Highest CPU Usage on Average", the list will be sorted by CPU Usage.

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 time period covered.

You can right click on any of the requests in the table to bring up a context menu. This will allow you to navigate to other useful perspectives such as Request Details and System Metrics.

You can also view Request Details, to see specific details of that request, such as the exact end time, the amount of memory it used, and the amount of CPU it used.

Perspective View

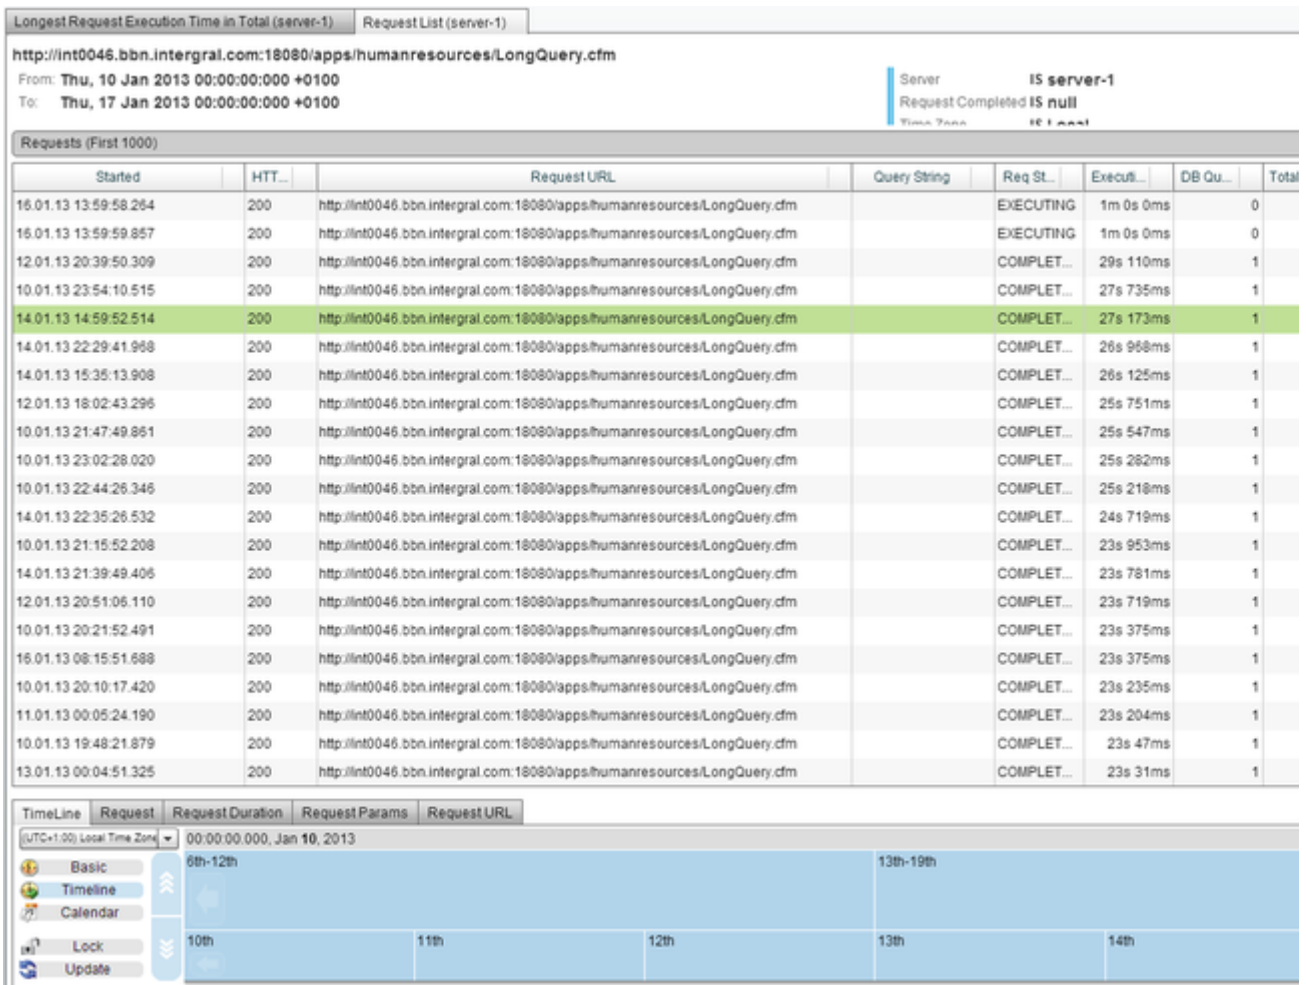


Figure 1: Aggregate Request List perspective

REQUEST BREAKDOWN

Request Breakdown (for a collection of requests)

Description

This perspective allows you to breakdown a collection of requests to see where the collection spend their execution time. This information time the collection of requests spends in each area.

Also provided on this page, is a table showing the minimum, average and maximum for a variety of metrics - such as DB query & row count final response sizes

In addition, this perspective offers request collection JDBC statistics, pin-pointing the "top 5" SQL statements in terms of cumulative execution metrics. Each of these show a sample SQL statement with the file path & line number to expedite your improvement process.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. All requests within the time period that n To change the order that the data in the table is shown, click on one of the table column headings to toggle between ascending and descending

To navigate to the JDBC statistics, select "JDBC statistics" from the drop down box in the uppermost right hand corner of the perspective.

NOTE: The loading times for this perspective can vary greatly depending on the number of requests that are included in the data set. It can take a few minutes to load due to the complex calculations that are being made.

Perspective View

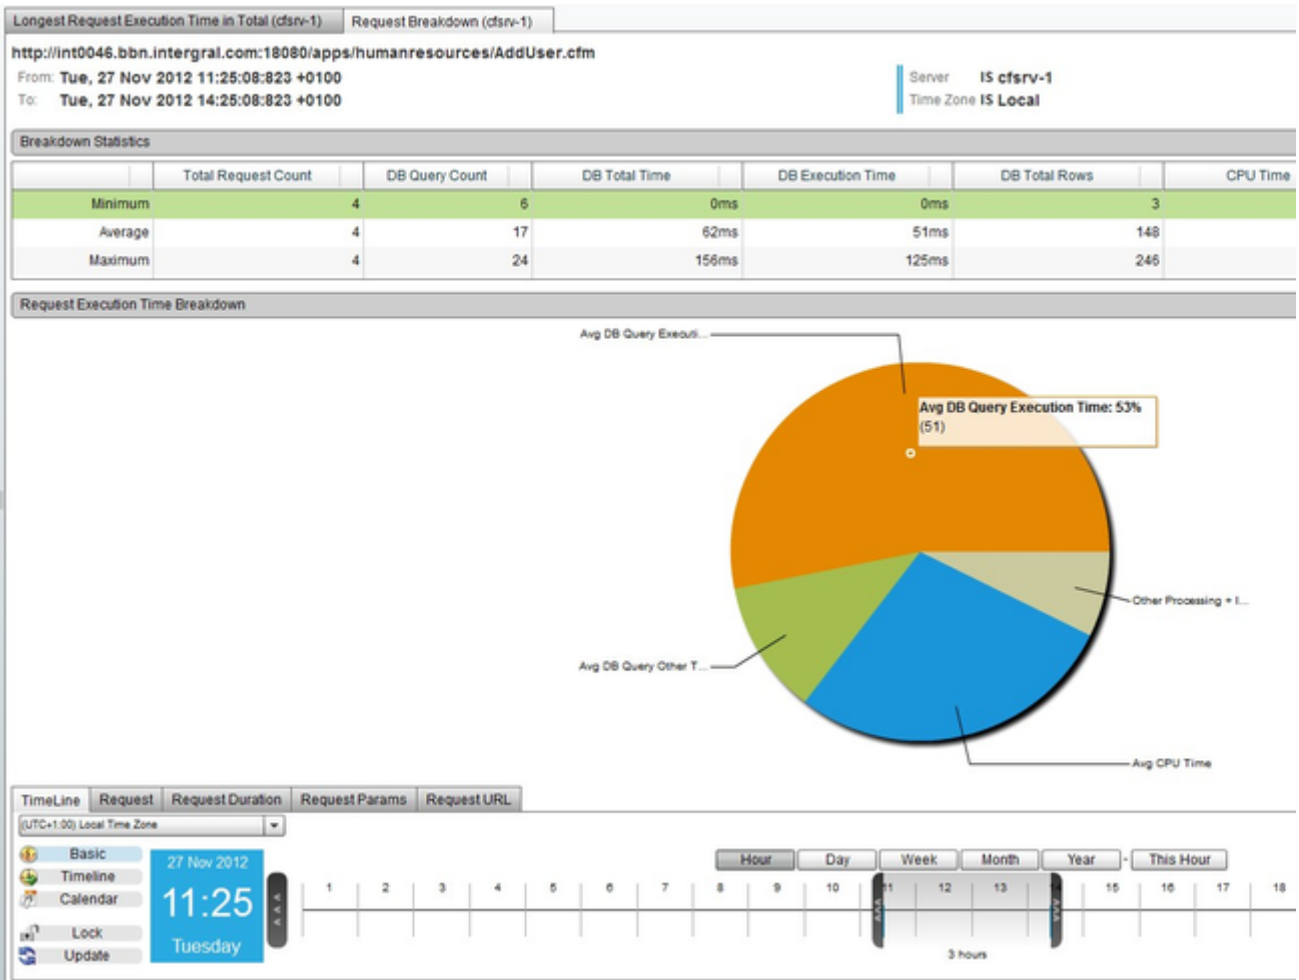


Figure 1: Request Breakdown -> Request Statistics view

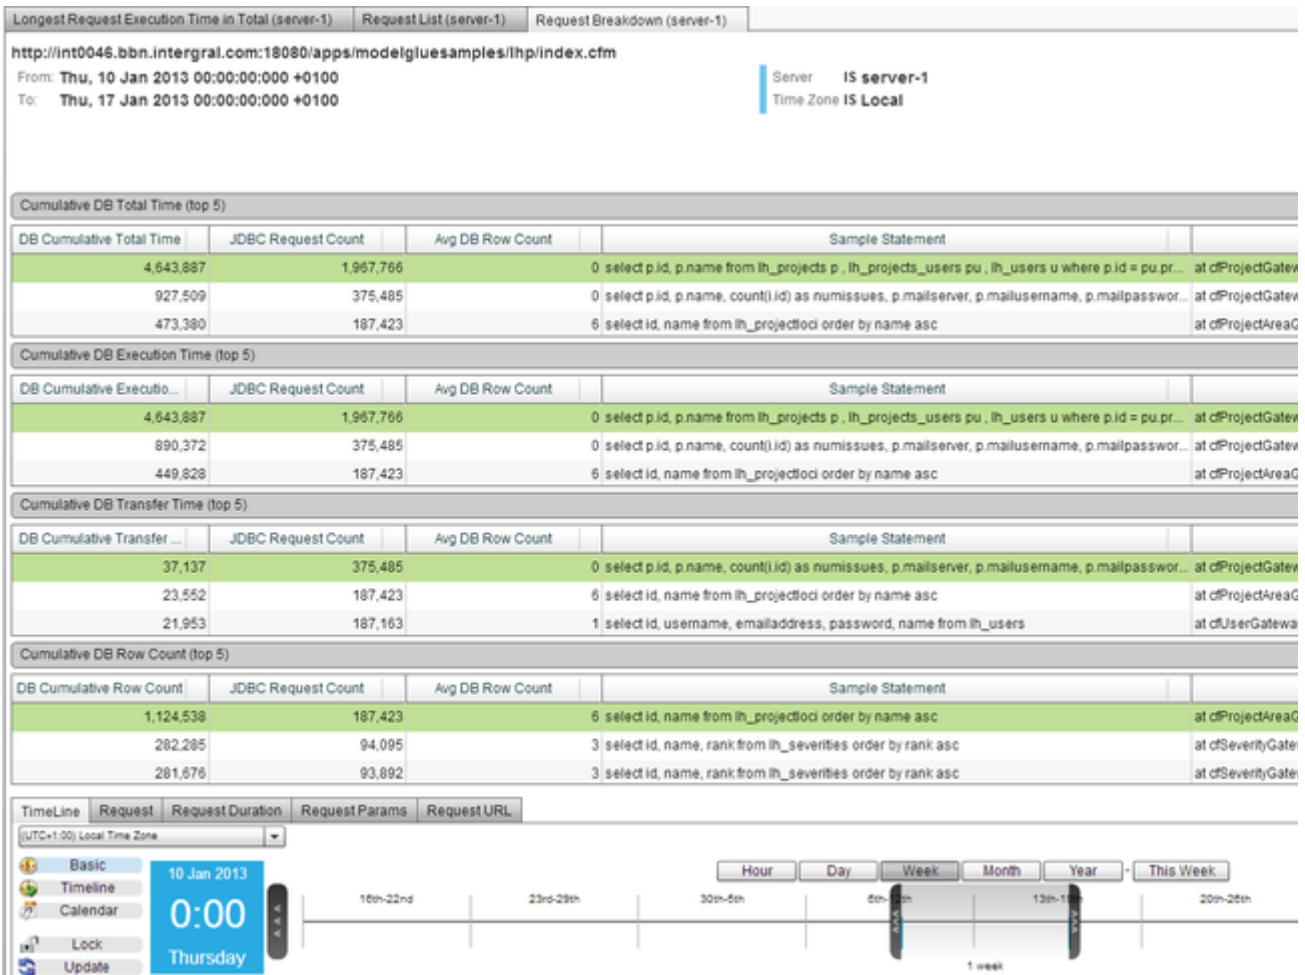


Figure 2: Request Breakdown -> JDBC Statistics
 REQUEST DETAILS

Request Details

Description

When viewing a grid that displays individual requests, such as "Slowest Requests" or "Highest CPU Usage", you can right click on any one you can navigate to the Request Details perspective.

This perspective has five different views that you can switch between using the buttons located above the timeline.

The five different views are:

- Request Details - Shows all of the details of an individual request, on one easy to read page.
- Request DB Details - This displays list of all the queries that were performed by the request.
- DB Requests - Shows a timeline of all the queries that were performed by the request.
- DB Breakdown - Displays a pie chart showing a breakdown of query execution time across all the queries that were performed by the request.
- Request Breakdown - Displays a pie chart which shows a breakdown where the request spends its time (CPU, network/IO, DB, et

Usage

To change the order in which the details are shown in the tables, click on one of the column headings to toggle between ascending and descending. You can right click on any table row or cell, and you will be able to copy the details.

Perspective View

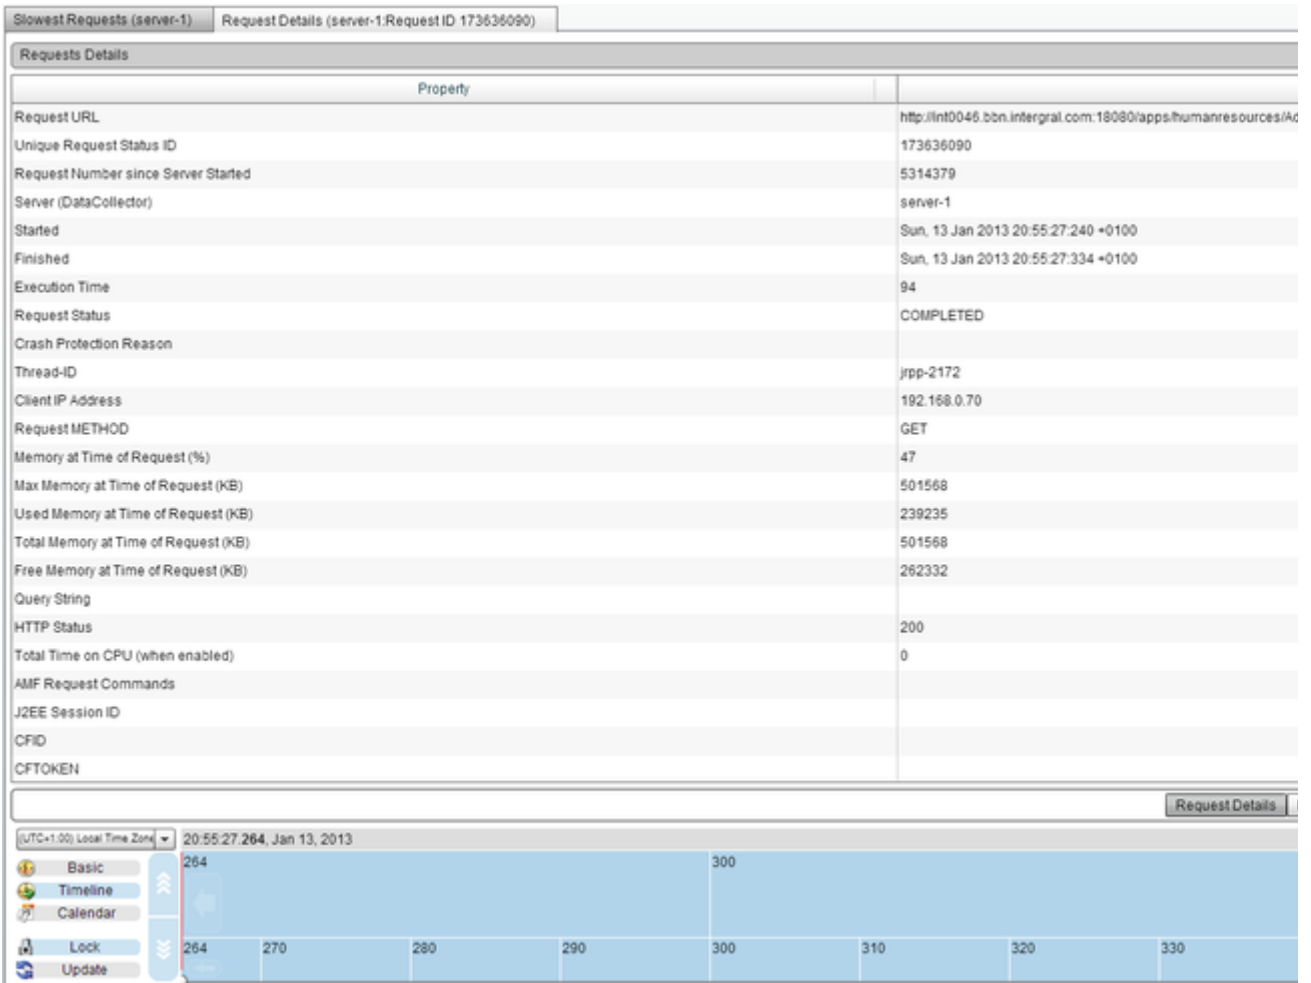


Figure 1: Request Details view

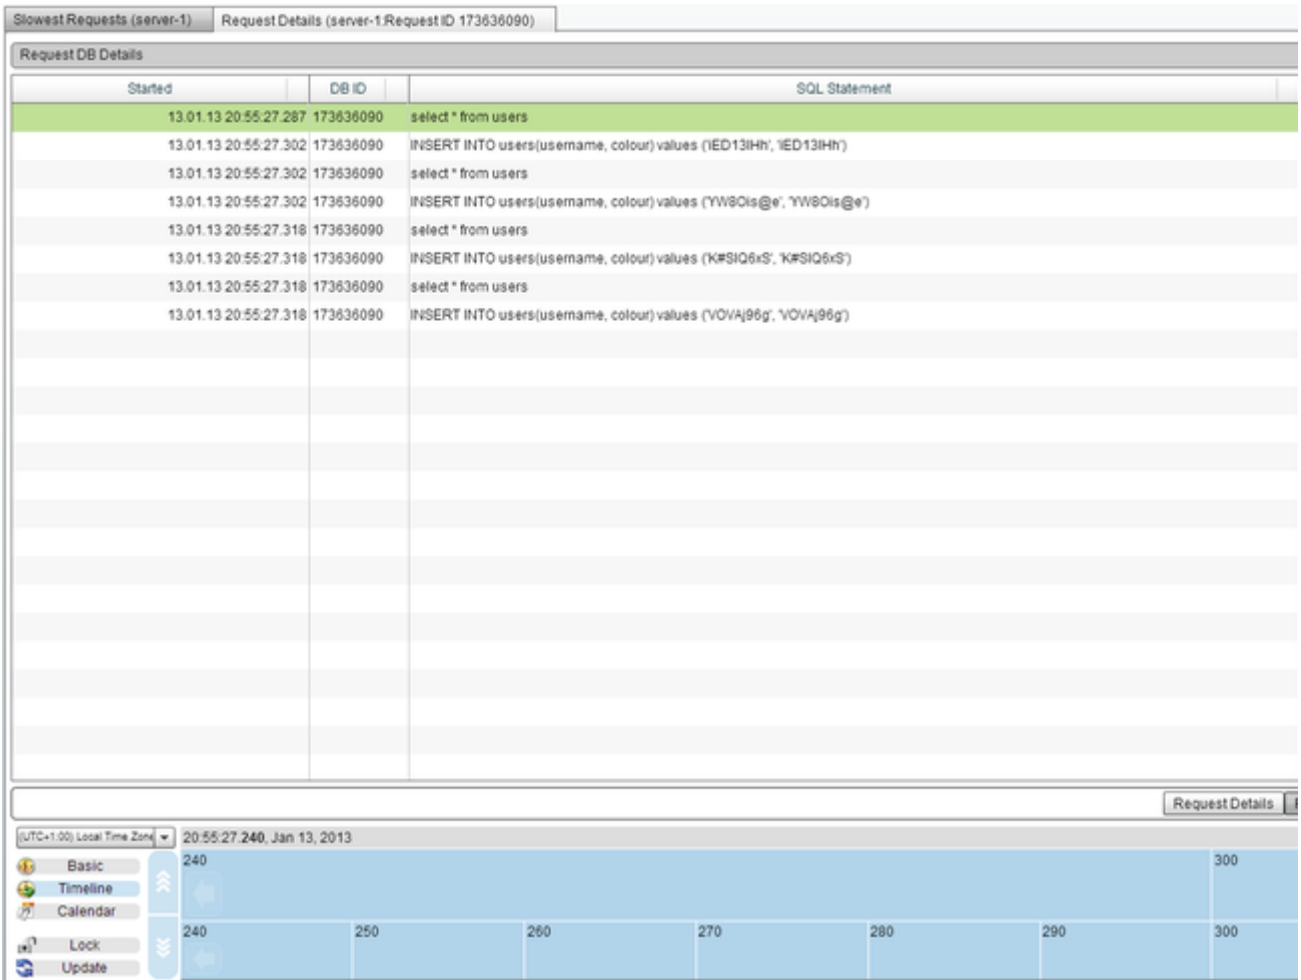


Figure 2: Request DB Details view

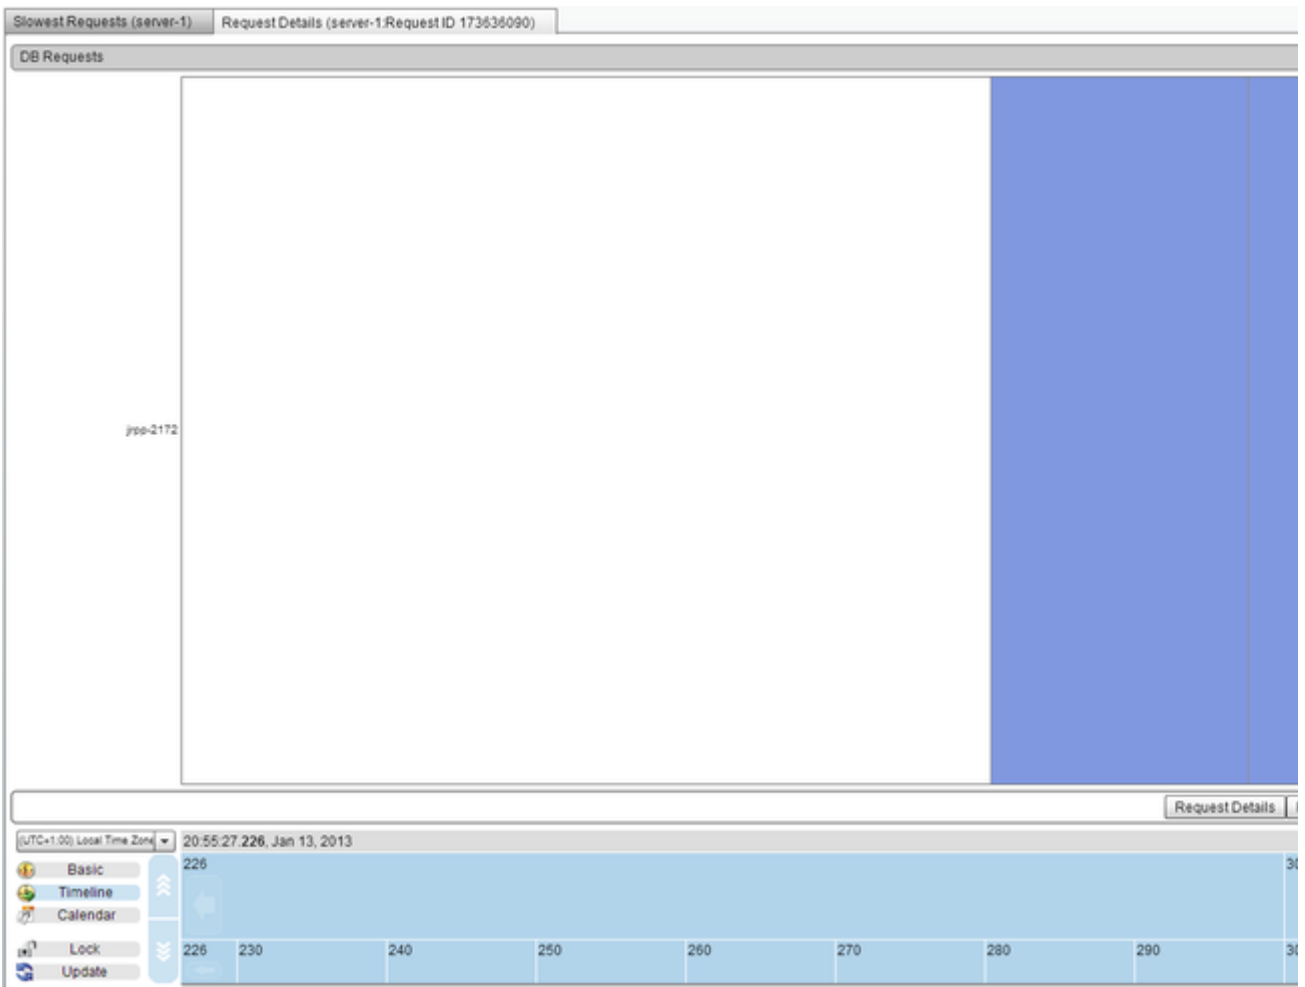


Figure 3: DB Requests view




Figure 4: DB Breakdown view




Figure 5: Request Breakdown

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 first 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 char covers.

Perspective View

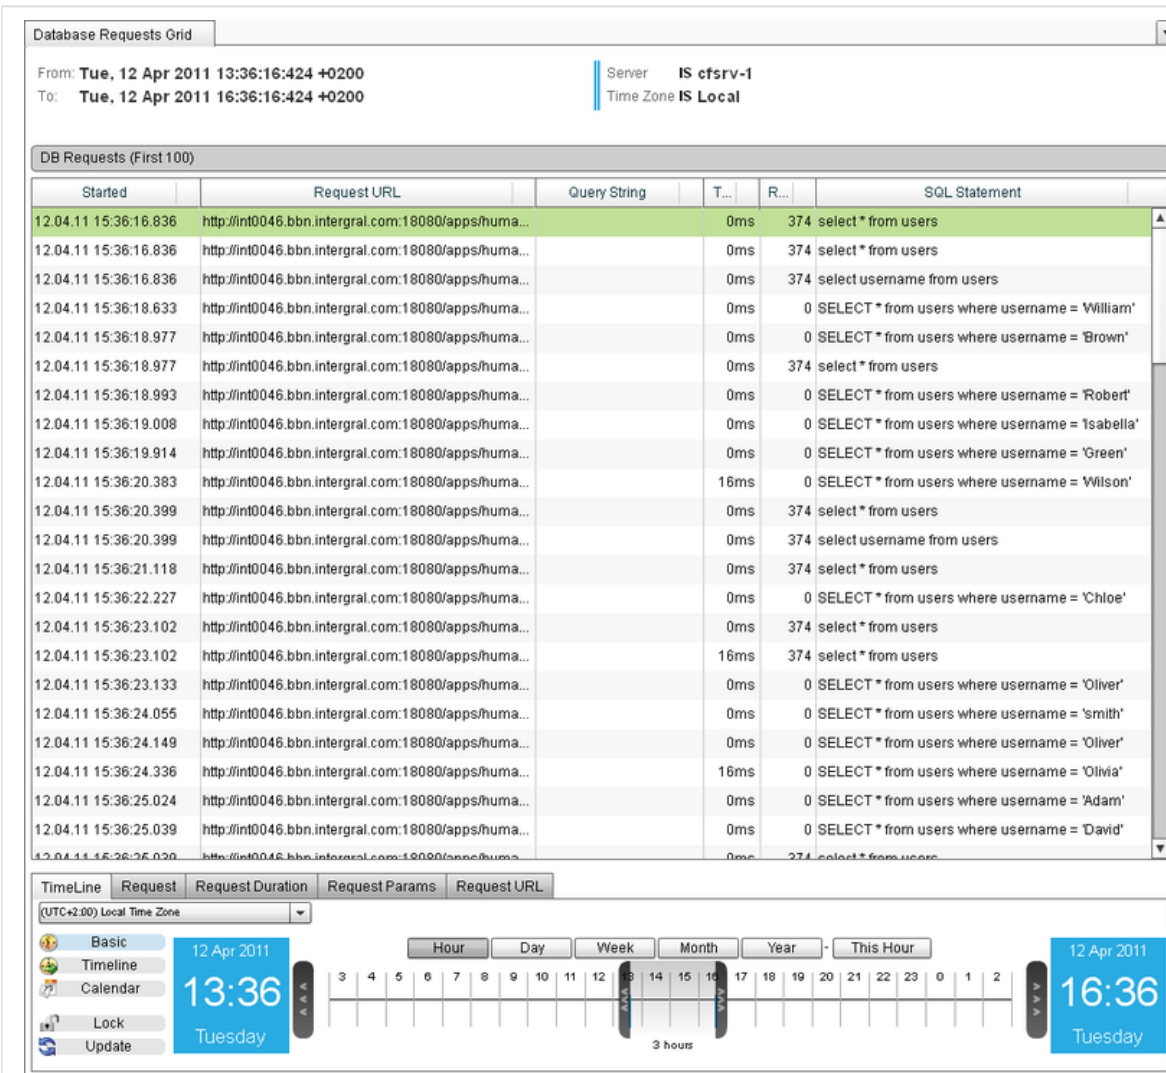


Figure 1: Database Requests Grid Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Database > Database Requests It shows the first 100 database requests from the entire cluster, ordered by Start Time. You can see which server each request came from Server Name.

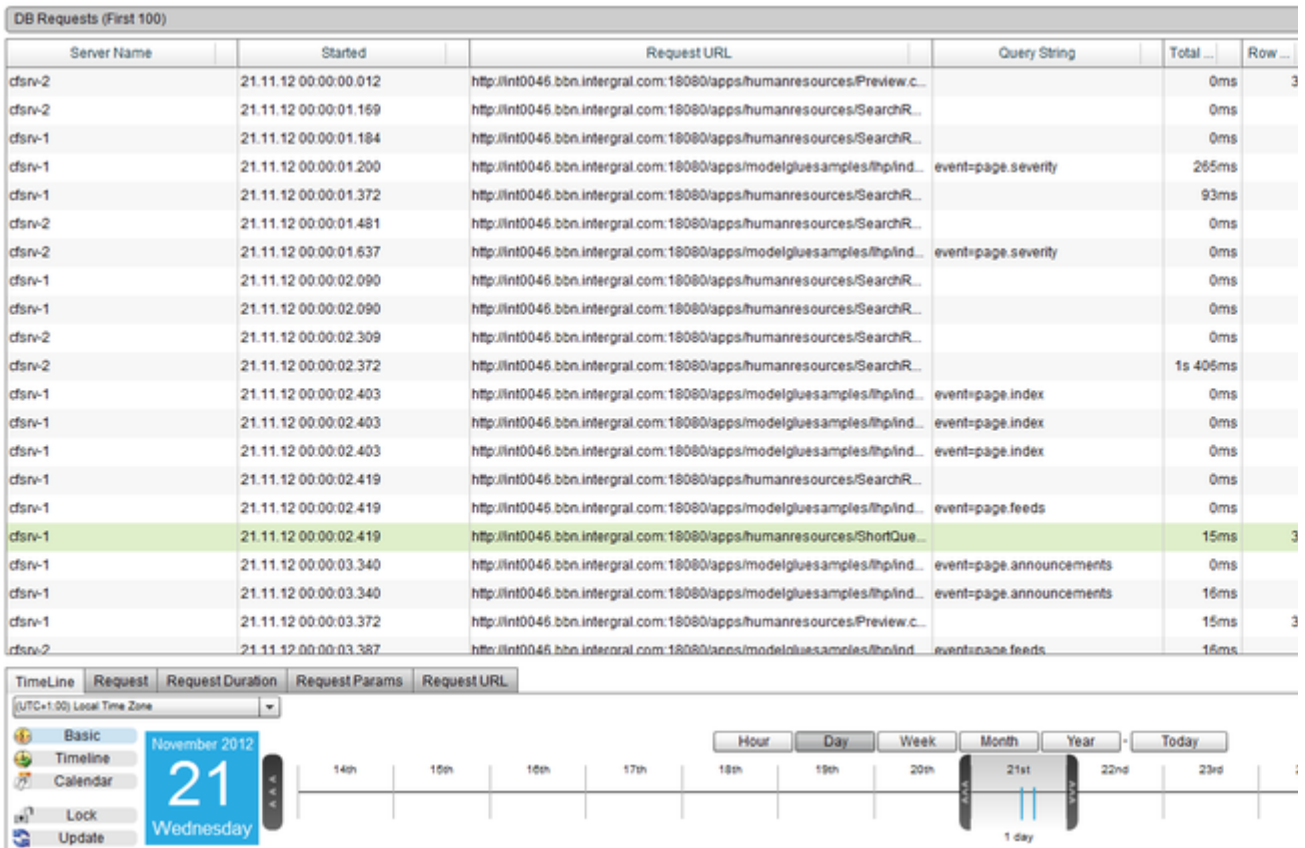


Figure 2: Cluster version of the Database Requests Grid perspective

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 can be used to help you identify many potential issues and give you ideas on how to fix them. Generally fetching rows from the database takes time, resource and can induce network overhead if the database is located far from the application. If the database is 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 performance and reduce 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



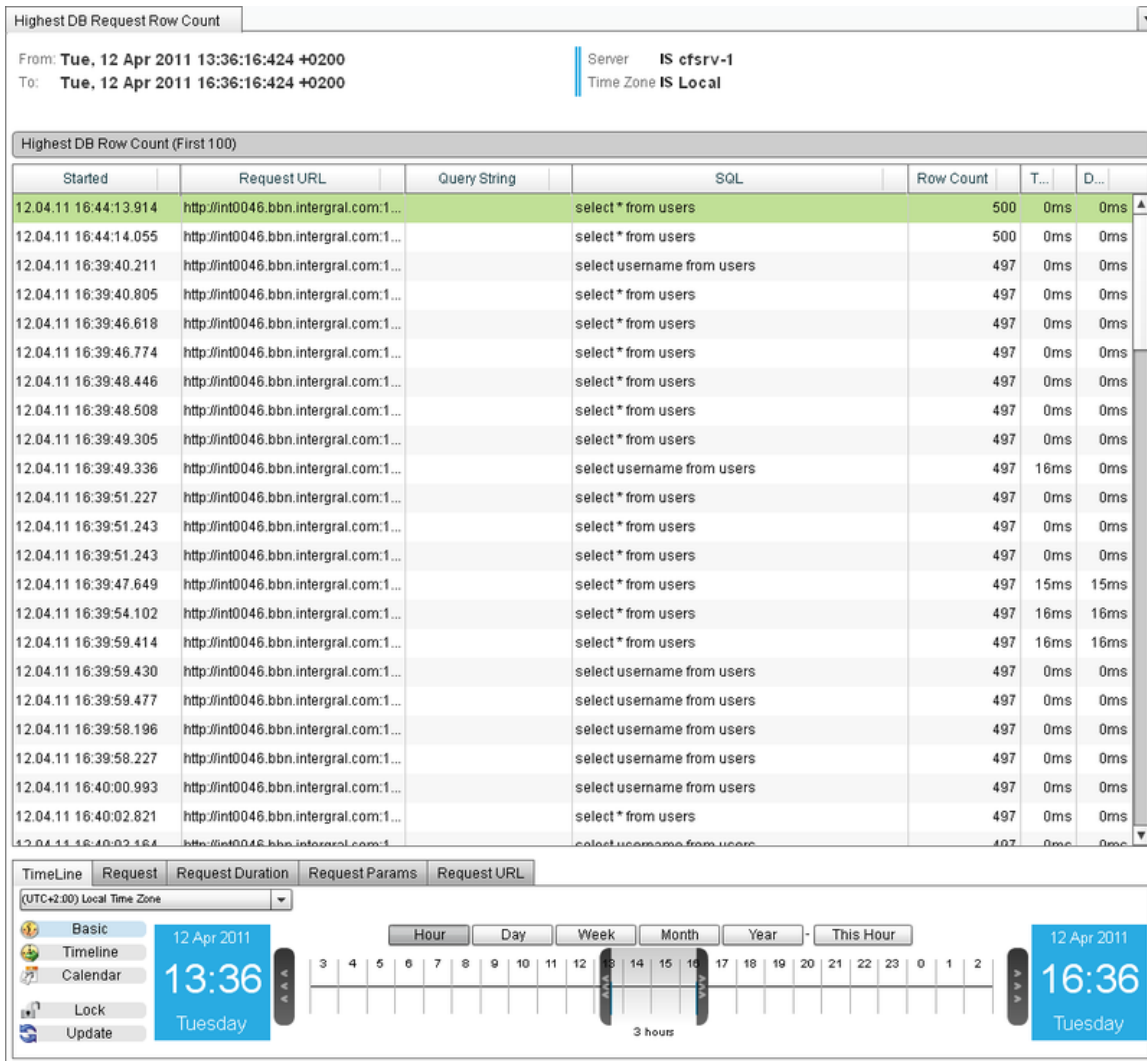


Figure 1: Highest DB Request Row Count Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Database > Highest DB Request Row Count.

It shows which request from the entire cluster fetched the most rows of data from the database. You can see which server each request came from. The tool displays Server Name.

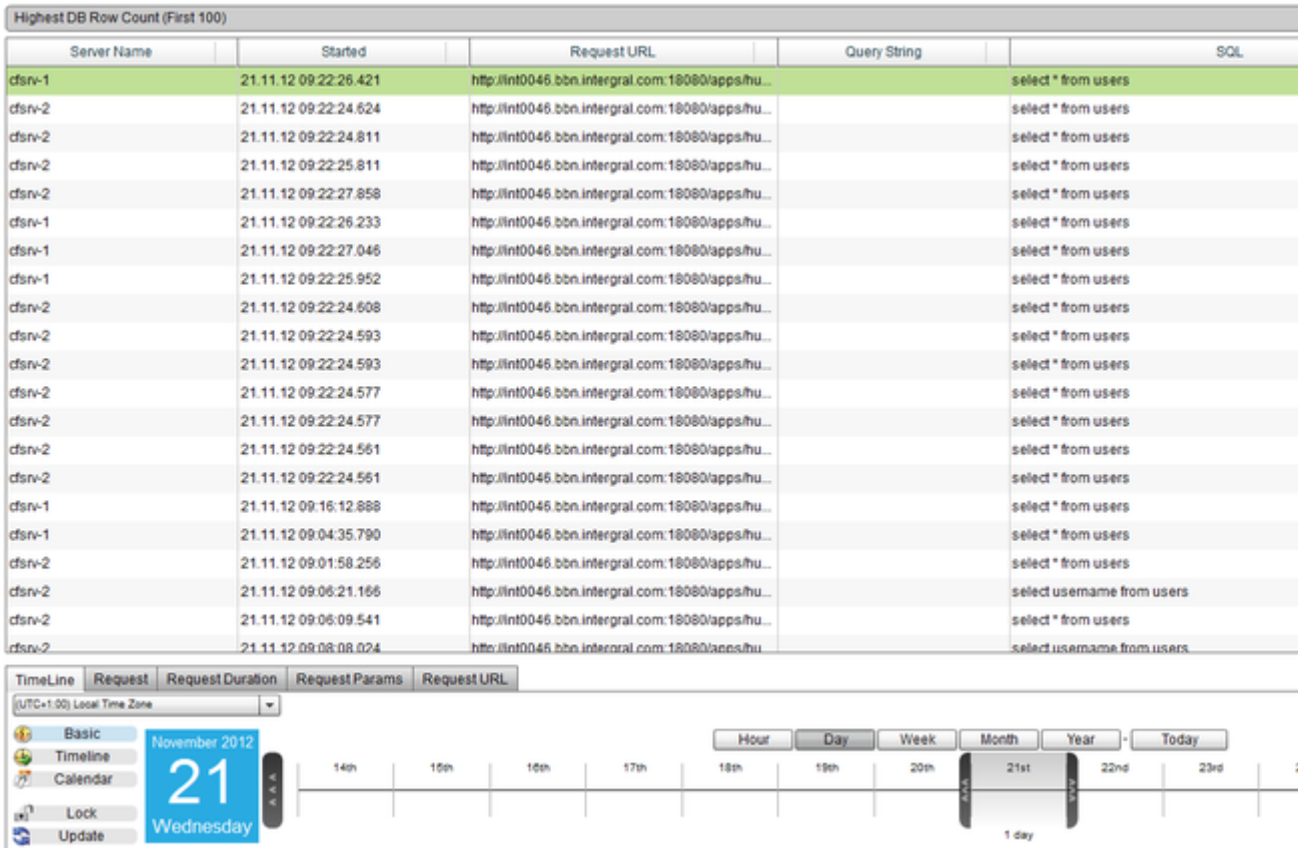


Figure 2: Cluster version of the Highest DB Request Row Count perspective

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 covers.

Perspective View



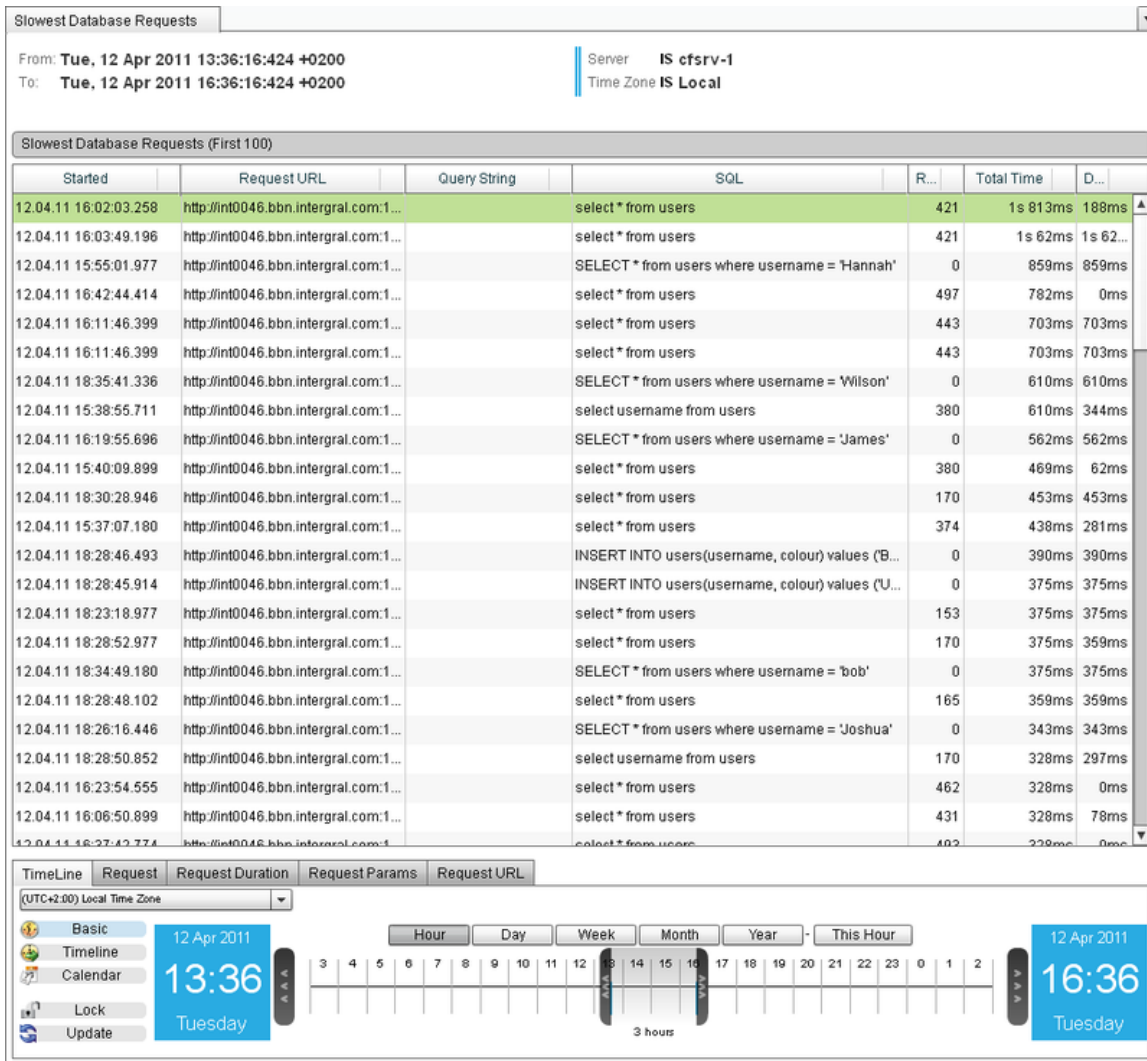


Figure 1: Slowest Database Requests Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Database > Slowest Database F It shows the top 100 slowest database requests from the entire cluster rather than just from one server. You can see which server each re which displays Server Name.

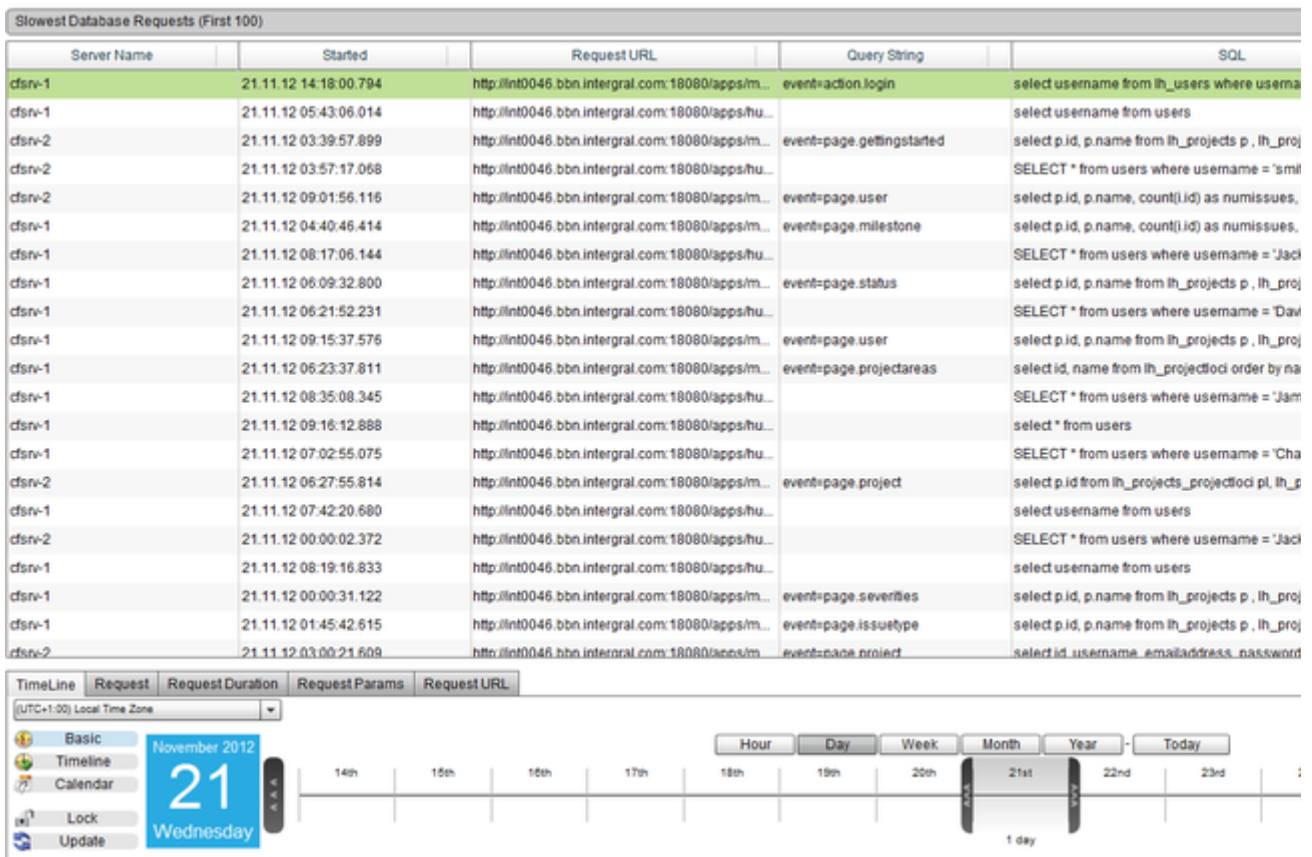


Figure 2: Cluster version of the Slowest Database Requests perspective

Back to the top

DATABASE USAGE GRAPHS

FusionAnalytics DataServices

Description

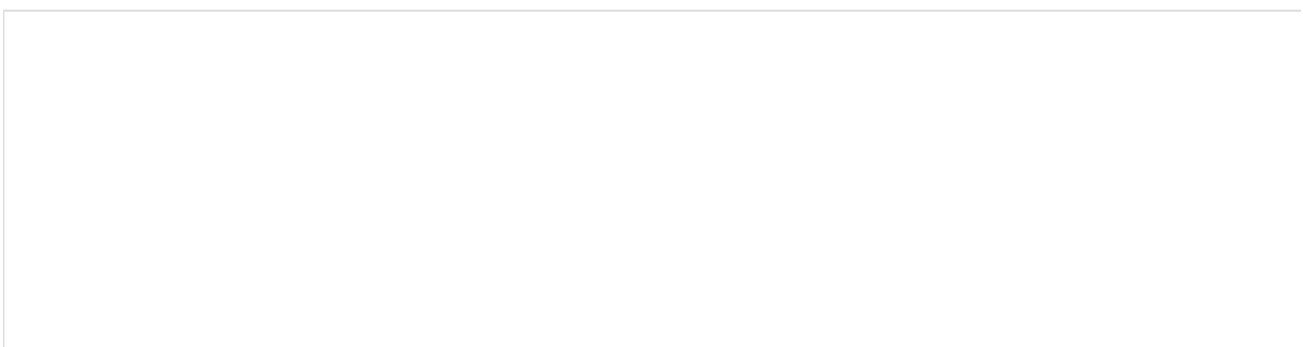
This perspective has three different views that you can switch between using the drop down box in the uppermost right hand corner of the The three different views are:

- DB Exec Time - This shows an area chart which displays the Average Execution Time of Database requests within the specified t
- DB Completed - This displays a chart showing the number of Database requests completed at the time.
- DB Per Second - This displays an area chart showing the number of Database Requests Per Second.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining :

Perspective View






Figure 1: Database Usage Graphs Perspective

Cluster Perspective

The cluster version of this perspective keeps the same three views as the server perspective, however, each view has a Server Comparison side by side, and a Cluster View graph and aggregate of the data across the whole server. The first view of the perspective can be seen in figure

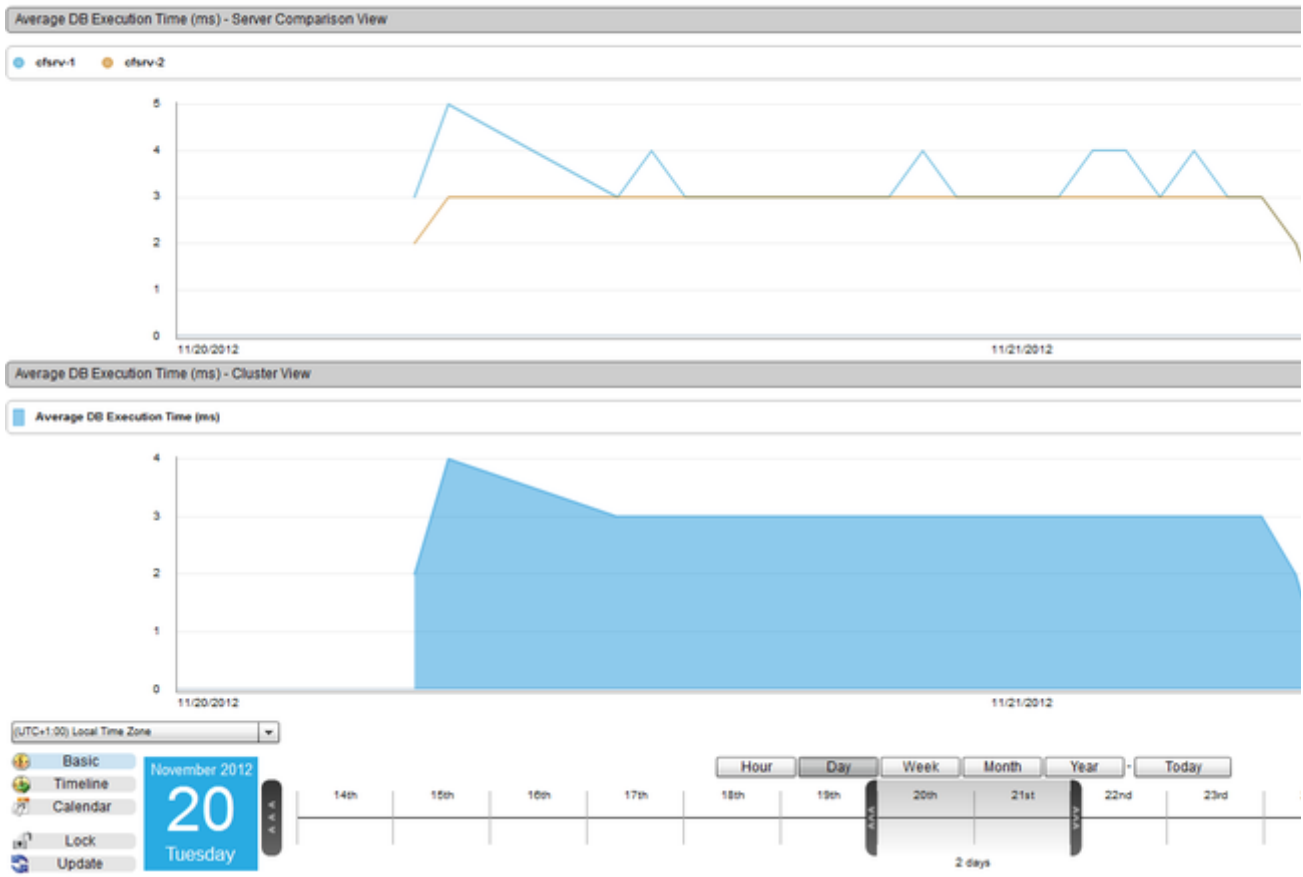


Figure 2: Cluster version of the Database Usage Graphs perspective

[Back to the top](#)

REQUESTS AND JDBC EXECUTION TIME

Request and JDBC Execution Time


Description

This perspective shows an area chart of the Average Request Execution Time and the Average Execution Time of Database Requests, wi

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining :

Perspective View



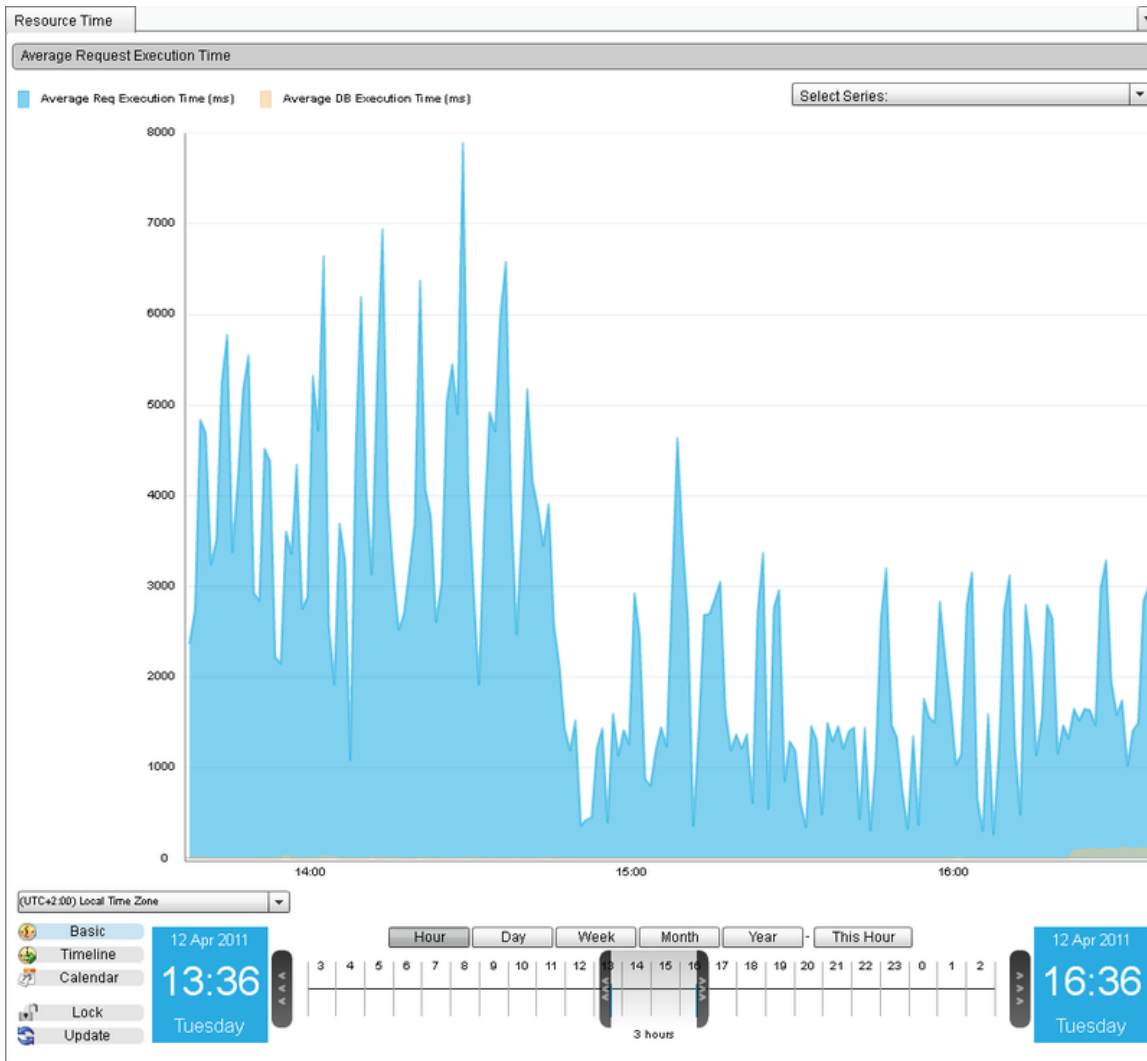


Figure 1: Request and JDBC Execution Time Perspective

Cluster Perspective

The cluster version of this perspective now has two different views that can be switched between using the drop down box in the uppermo. The first view is a Server Comparison graph view showing two separate graphs. The first graph showing the Average Request Execution. The second graph shows the Average DB Execution Time for each individual server within the cluster.

The second view is a Cluster View showing one Area graph, which displays the Average Request Exec Time and the Average DB Exec Ti




Figure 2: Server Comparison View for Request and JDBC Execution Time perspective

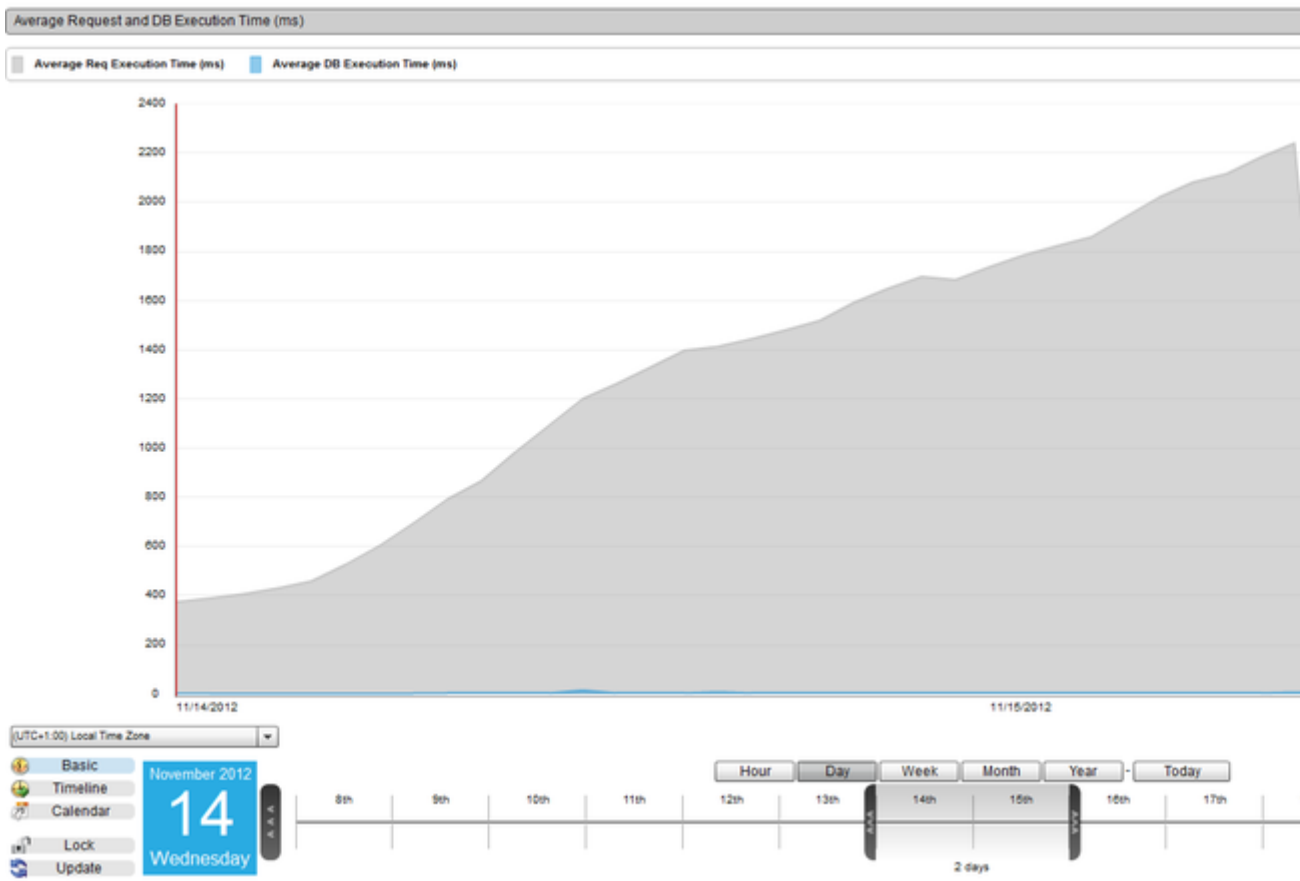


Figure 3: Cluster View for Request and JDBC Execution Time perspective

[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 has two different views that you can switch between using the drop down box in the uppermost right hand corner of the p. The first view shows a chart of the Used Memory and Maximum Memory in kb within the specific period of time. The second view shows the Used Memory and Maximum Memory in % usage within the specific time period.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining :

Perspective View

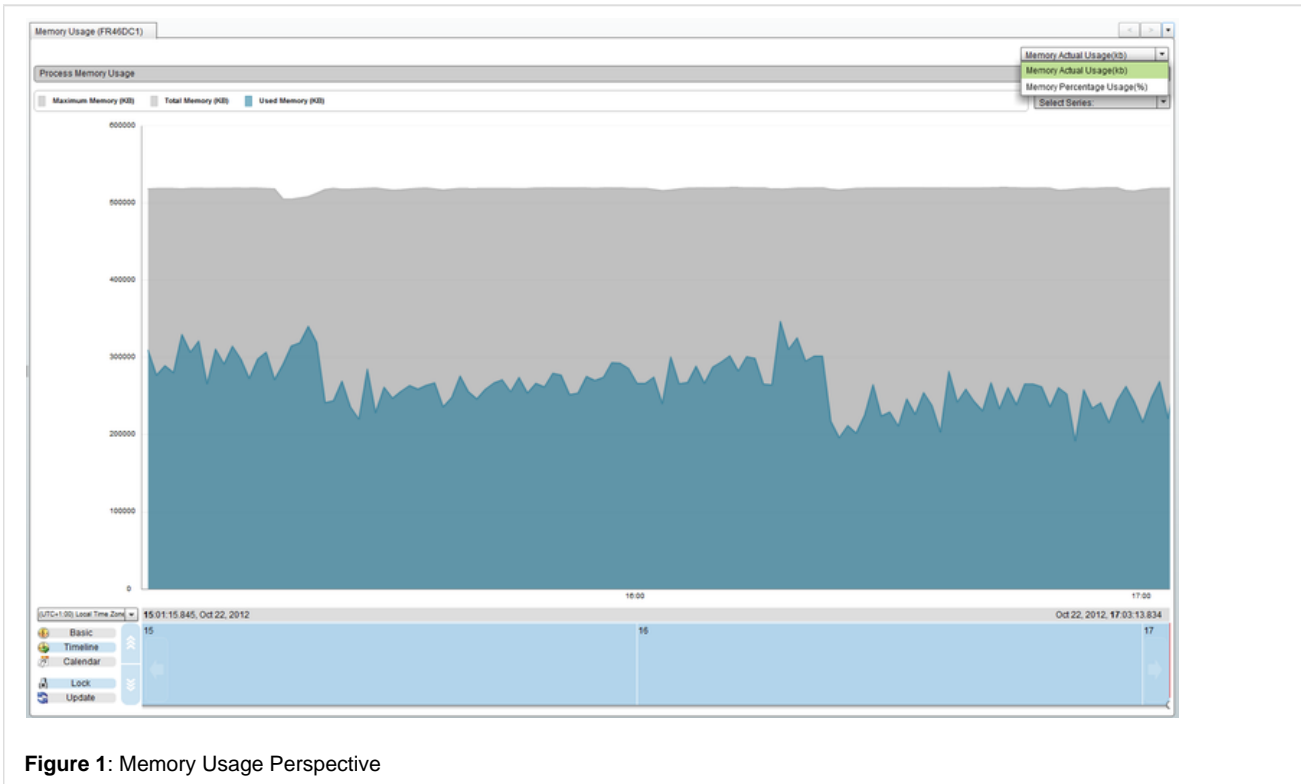


Figure 1: Memory Usage Perspective

Cluster Perspective

The cluster version of this perspective now has three different views that can be switched between using the drop down box in the upper right. Memory Actual Usage and Memory Percentage Usage both have a Server Comparison graph view showing the Memory Usage for each in the same graph.

The final view is a Cluster View showing one Area graph, which displays the Average Memory Usage of the whole cluster.



Figure 2: Cluster version of the Memory Usage perspective

[Back to the top](#)

MEMORY SPACES

Memory Spaces

Description

This perspective is used to analyze different aspects of memory such as Code-Cache, and displays a chart of the memory usage within the components.


You can switch between the different views that the perspective has to offer by using the drop down box in the uppermost right hand corner.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers.

You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to view the remaining series.

Perspective View



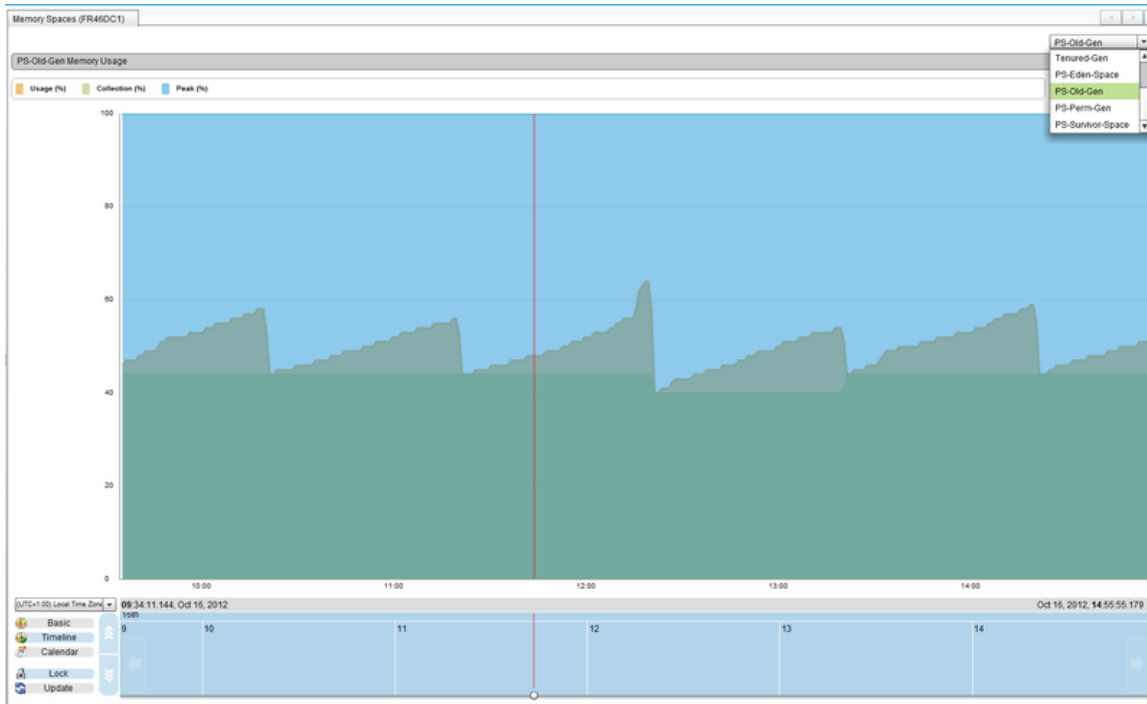


Figure 1: Memory Spaces Perspective

Cluster Perspective

Each of the views available in the server version of the perspective has been given a Server Comparison view for easy comparison between a more generalized view of the performance of the cluster .
 You can switch between the different views that the perspective has to offer by using the drop down box in the uppermost right hand corner



Figure 2: Cluster version of the Memory Spaces perspective

[Back to the top](#)

MEMORY HEAP NON-HEAP SUMMARY

Memory Heap Non-Heap Summary


Description

This perspective has three different views that you can switch between using the drop down box in the uppermost right hand corner of the view. The first view shows a chart of the Memory Heap/Non-Heap Actual Usage in bytes within the specific period of time. This view is offered in the Timeline view. The final view shows the Memory Heap/Non-Heap Actual Usage % Usage.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to view the remaining series.

Perspective View






Figure 1: Memory Heap Non-Heap Summary Perspective

Cluster Perspective

The **Heap/Non-Heap Actual Usage** view from the server version of the perspective has been given a Server Comparison view to see the side by side on the same graph, and also a Cluster View which shows an aggregate of the data across the whole cluster. The **Heap/Non-Heap Percentage Usage** view from the server version of the perspective has been given **two** Cluster views, one for a stack has a Server Comparison view to see the percentage usage for each individual server within the cluster.

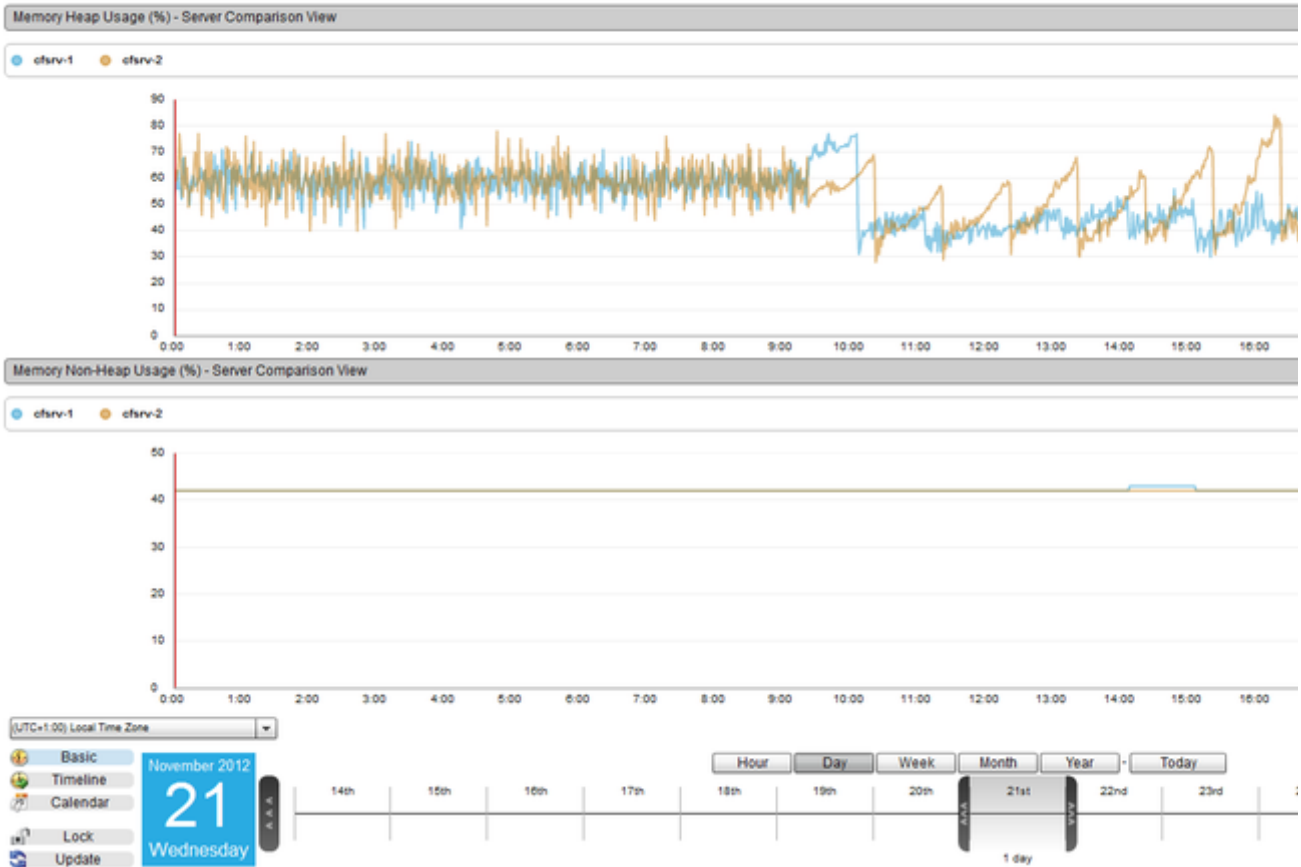


Figure 2: Cluster version of the Memory Heap Non-Heap Summary perspective.

[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

Use the Date Navigation in order to change the period of time over which this perspective covers.

You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining :

Perspective View

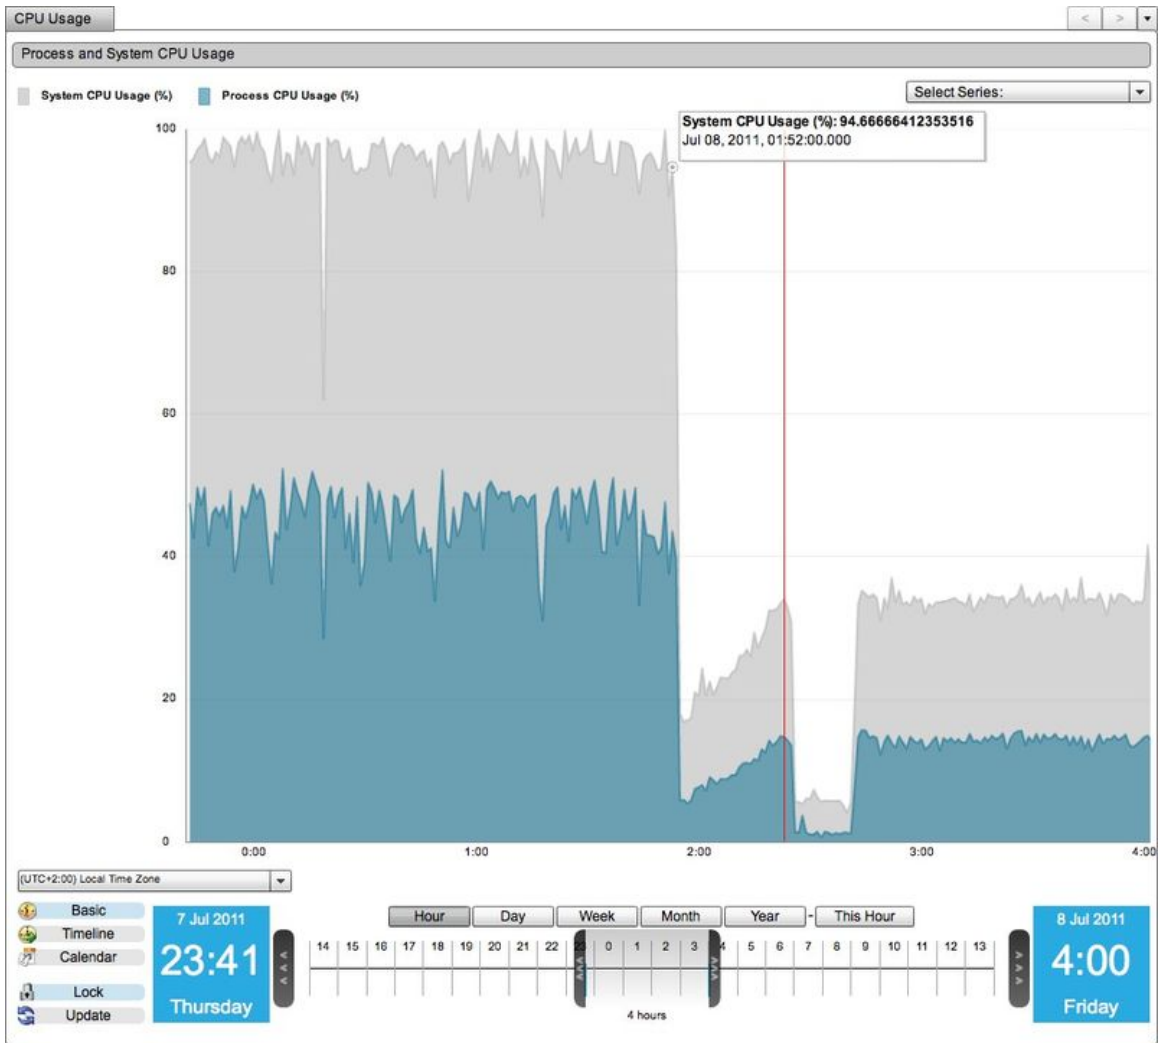


Figure 1: CPU Usage Perspective

Cluster Perspective

The cluster version of this perspective now has two different views that can be switched between using the drop down box in the uppermo. There is a Server Comparison graph view showing the CPU Usage for each individual server within the cluster, side by side on the same g. There is also a Cluster View showing one Area graph, which displays the Average CPU Usage of the whole cluster.



Figure: Cluster version of the CPU Usage perspective

[Back to the top](#)

CPU/MEMORY USAGE

CPU/Memory Usage

Description


This perspective shows a chart of the System CPU Usage(%) alongside the Process CPU Usage(%) and Process Memory Usage(%) with

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers.

You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining :

Perspective View






Figure 1: CPU Memory Usage Perspective

Cluster Perspective

The cluster version of this perspective is the same as the cluster version of the CPU Usage perspective, except that the Server Comparison Process Memory Usage(%) for each server, and the Cluster View has an extra series showing the Process Memory Usage(%).



Figure 2: Cluster version of the CPU/Memory Usage perspective

[Back to the top](#)

Thread

Thread

Overview

Next Steps

Thread Activity

THREAD ACTIVITY

Thread Activity

Description

The Thread Activity perspective shows a Time Block Chart of the Activity of all threads during the specified time period. The state of each thread is represented by different colors. The color red means that the thread was blocked by something. The color green means that the thread is active and running normally, and the color blue means that the thread is waiting or idle.

This perspective does not appear in the Cluster folder.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. You can right click on any of the requests to bring up a context menu. This will allow you to navigate to other useful perspectives for the selected Metrics.

Perspective View




Figure 1: Thread Activity perspective

[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 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 each graph. You can click on the garbage can on the Memory Demand graph to request the system garbage collector.

Perspective View

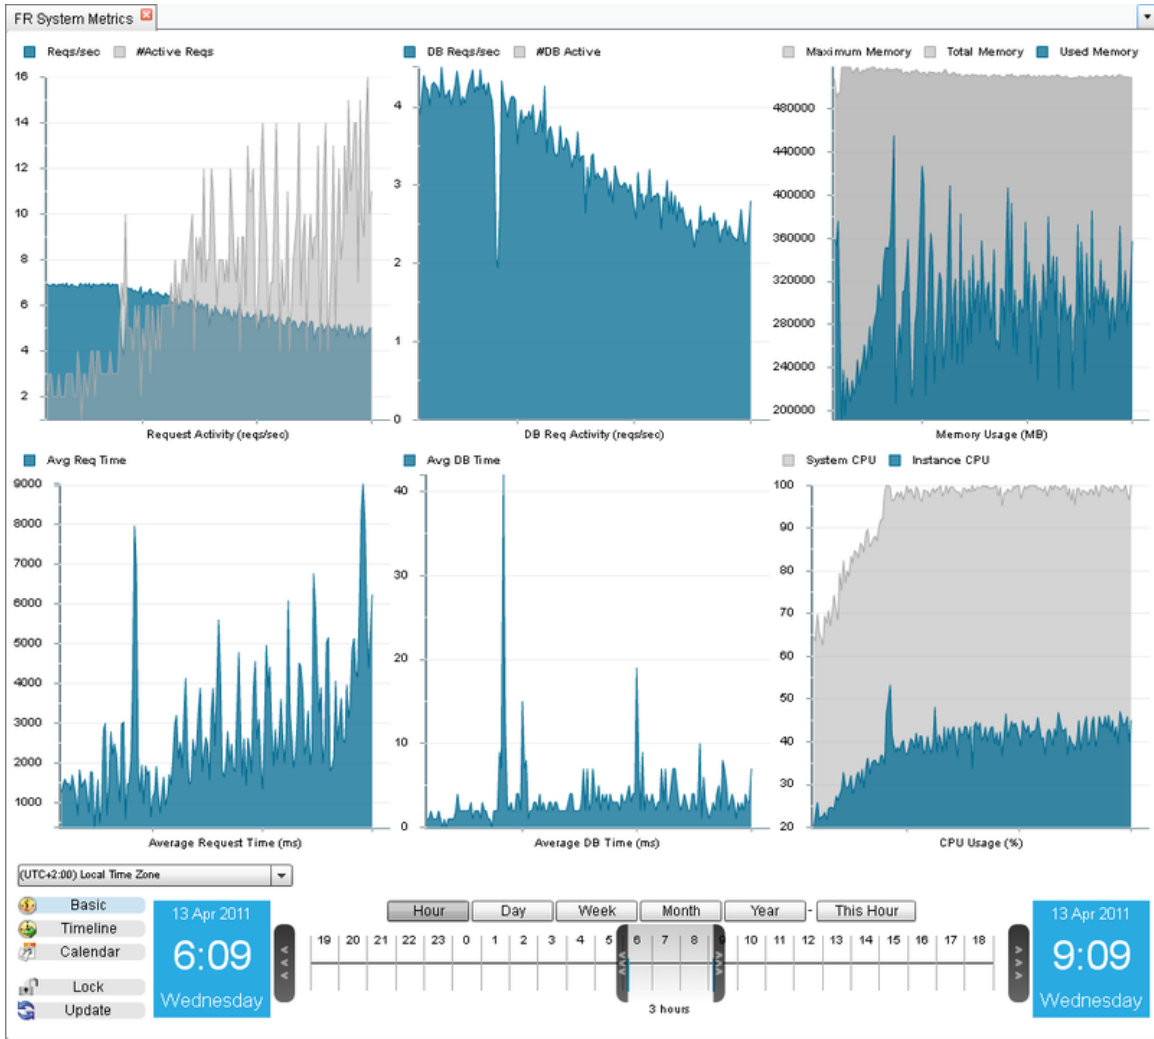


Figure 1: FusionReactor System Metrics Perspective

Cluster Perspective

The cluster version of this perspective again shows the 6 graphs: Request Load, Request Time, JDBC Load, JDBC Request Time, Memory graph will be an aggregate of the data for all servers in the cluster. Therefore the perspective will look almost the same as the server versi

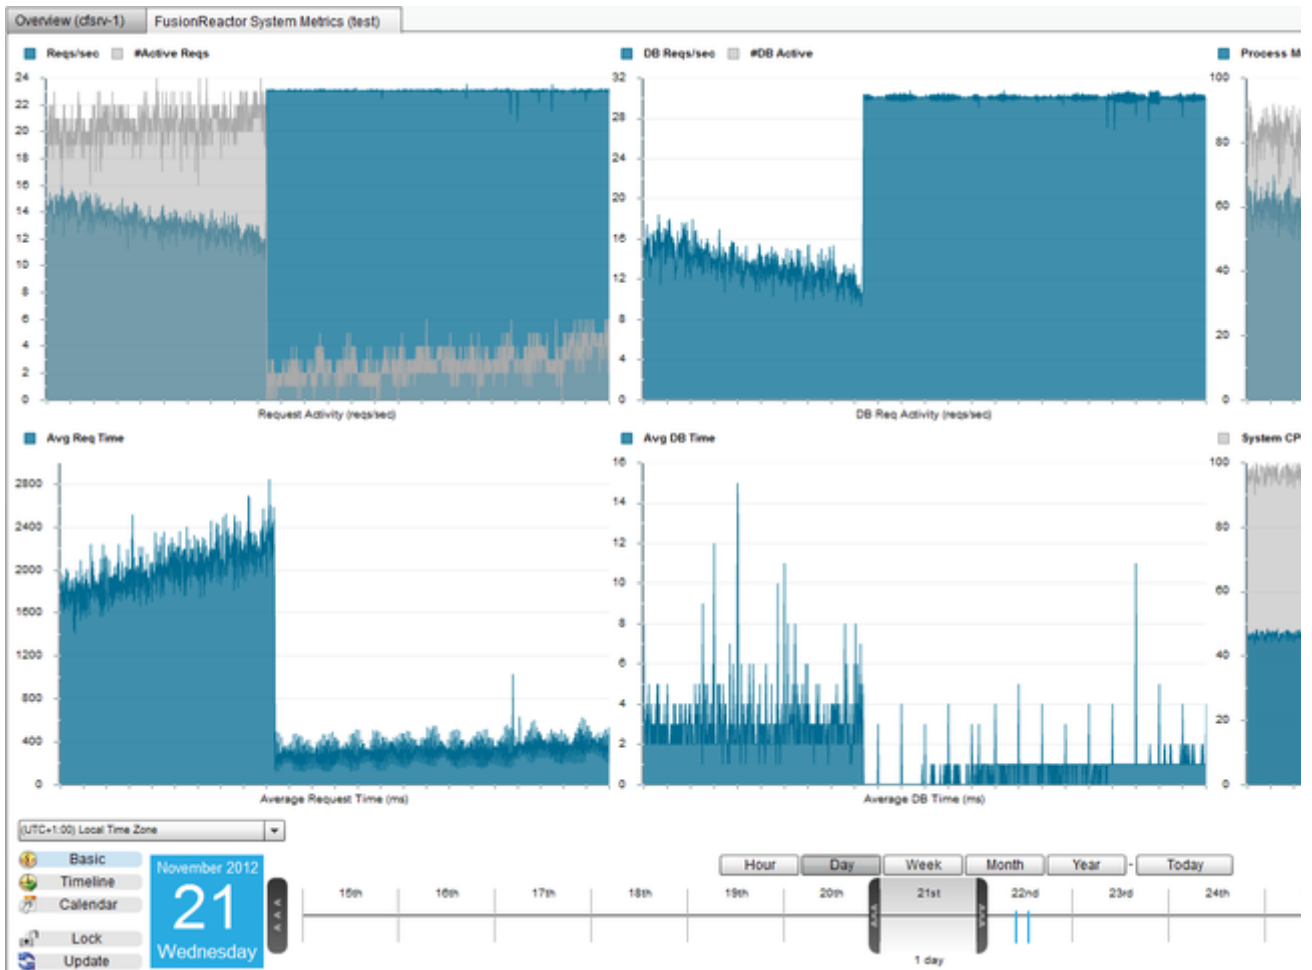


Figure 2: Cluster version of the FusionReactor System Metrics perspective

[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 perspective covers.

Perspective View

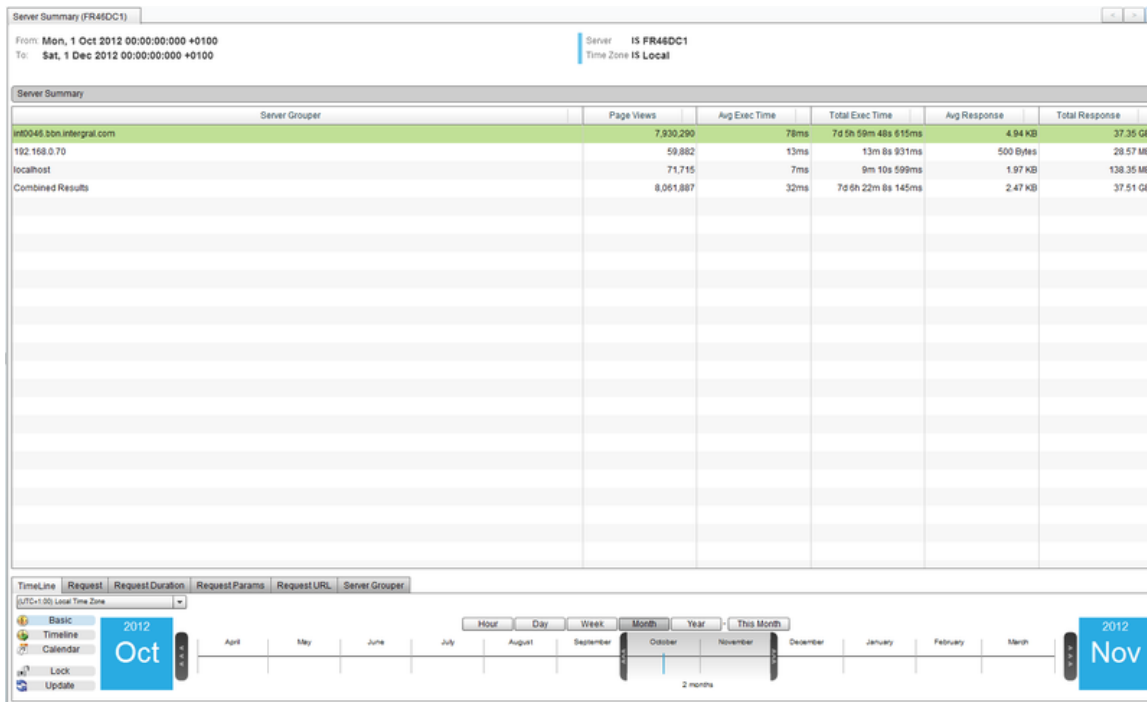


Figure 1: Server Summary Perspective

Cluster Perspective

The cluster version of this perspective shows the same perspective as the server version(above), however the data will be calculated using the cluster to form an aggregate. The Cluster perspective will look and work the same as the server version of the perspective.

Back to the top

RESOURCE USAGE DATA GRID

Resource Usage Data Grid

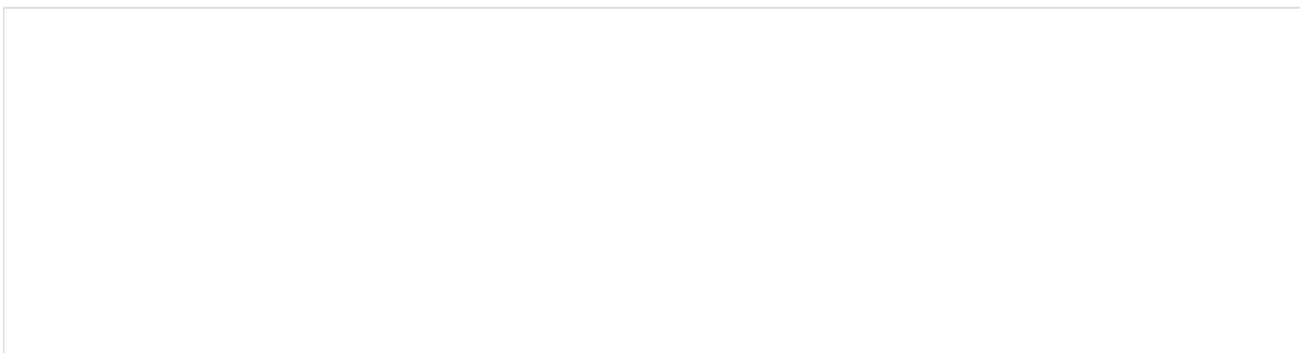
Description

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 covers.

Perspective View



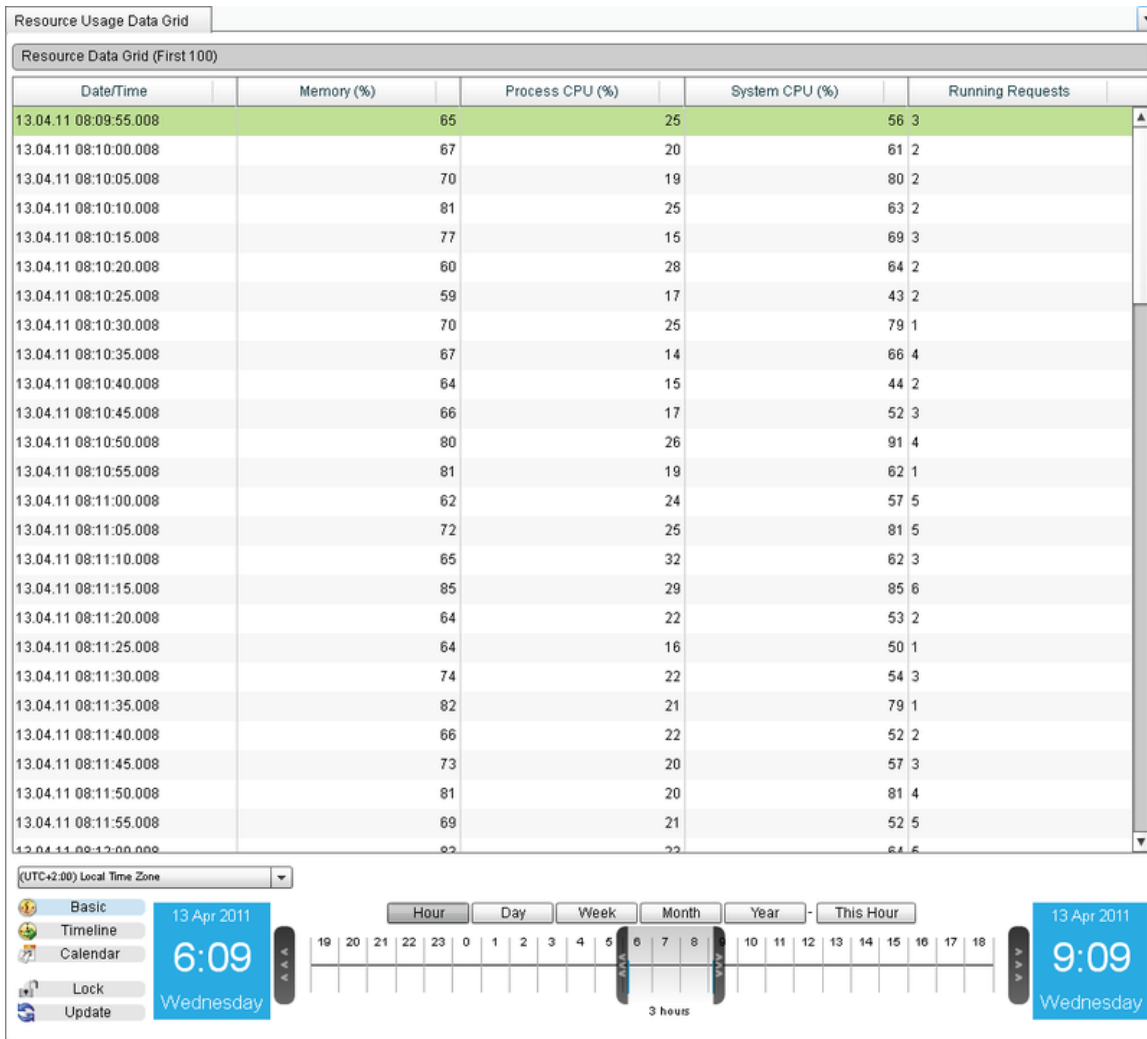


Figure 1: Resource Usage Data Grid Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > General > Resource Usage Data: It shows a table of the usage of different resources, such as memory and cpu, at different times. This data is recorded every 5 seconds for You can see which server the data came from in the leftmost column of the table, which displays Server Name.

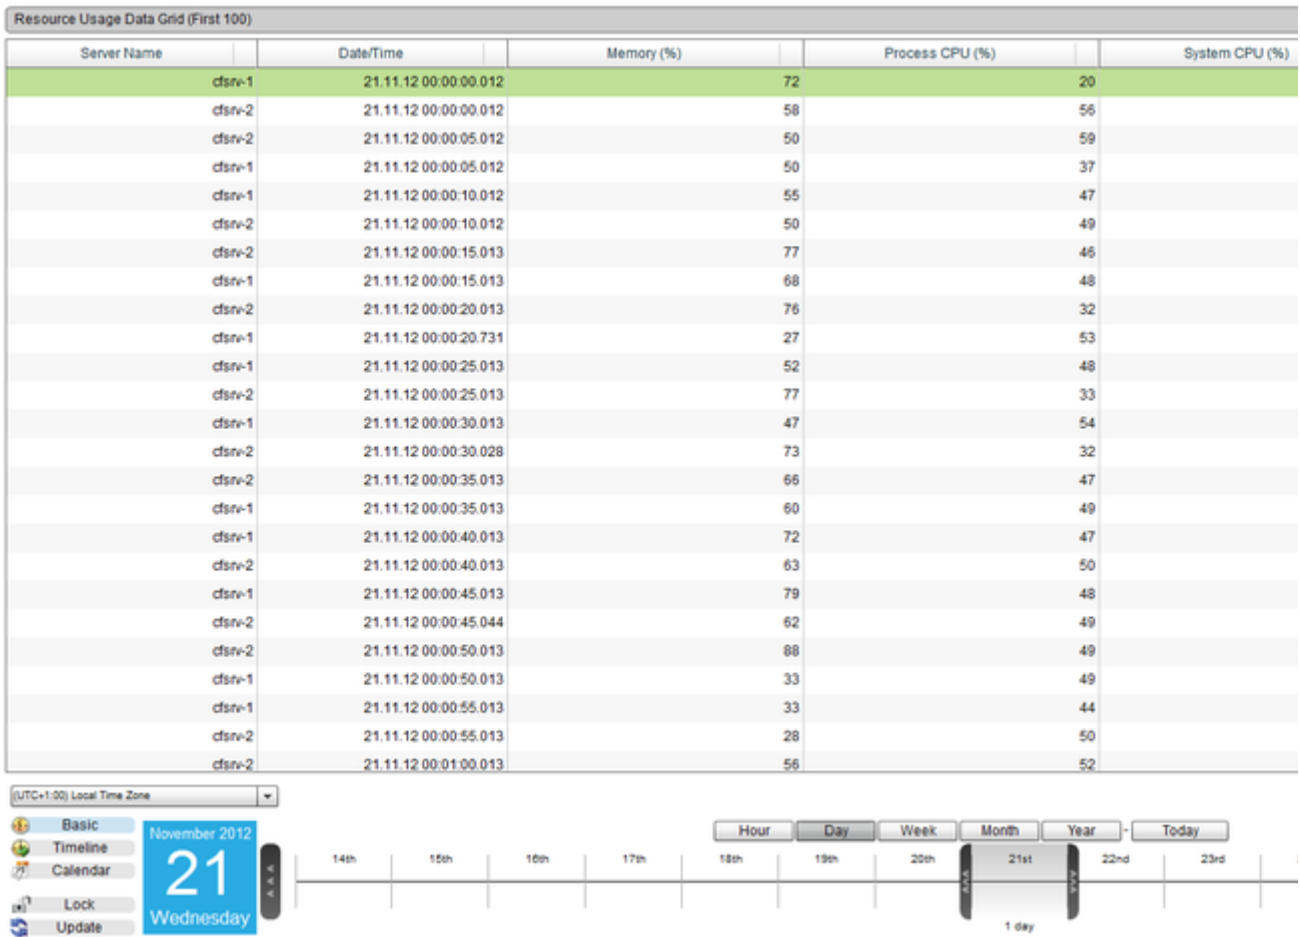


Figure 2: Cluster version of the Resource Usage Data Grid perspective

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, ordered by Date.

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 time covers.

Perspective View



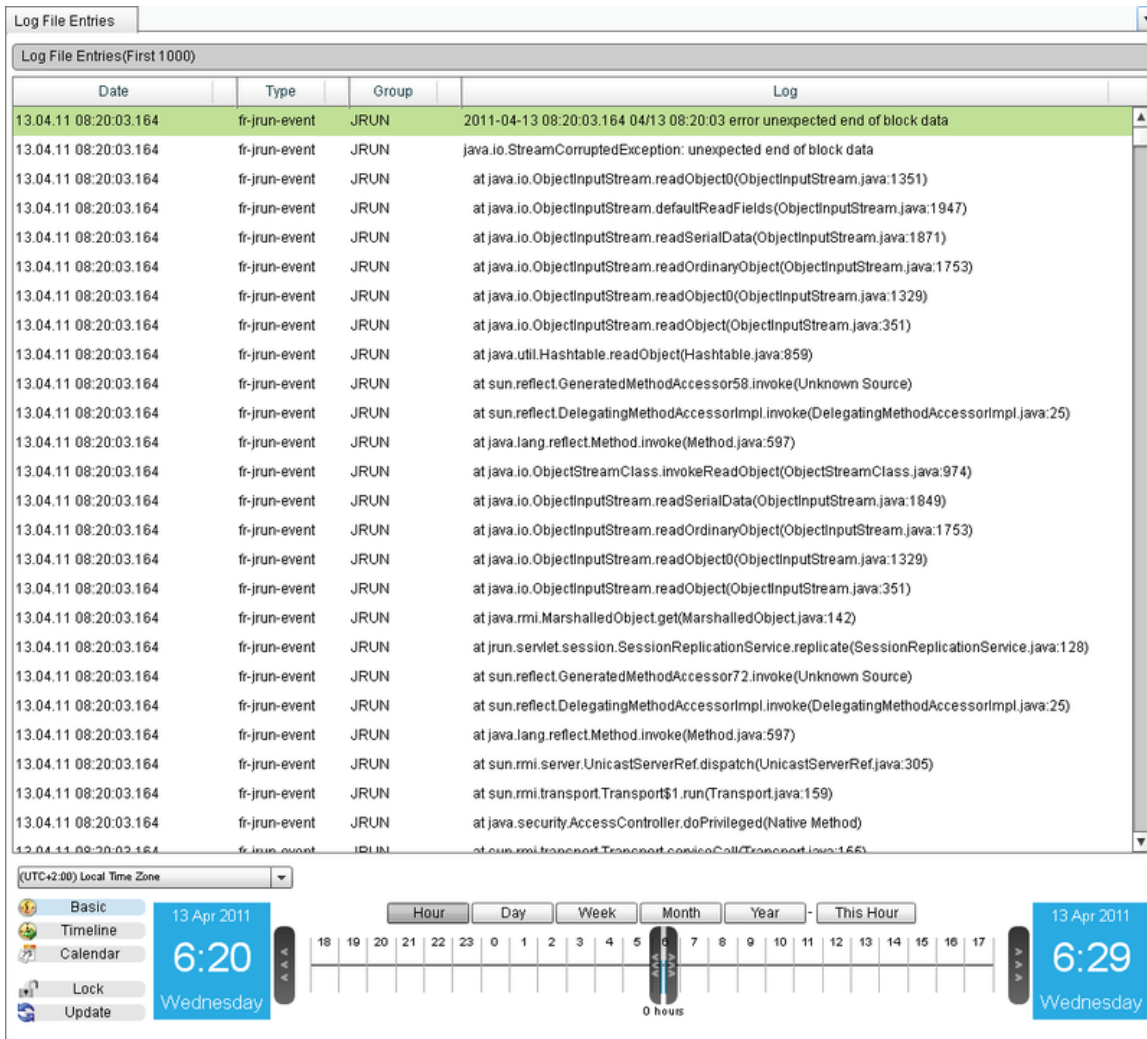


Figure 1: Log File Entries Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Logs > Log File Entries from the It shows the first 100 Log File Entries, ordered by Date, from the entire cluster. You can see which server each Log File Entry came from in its Server Name.

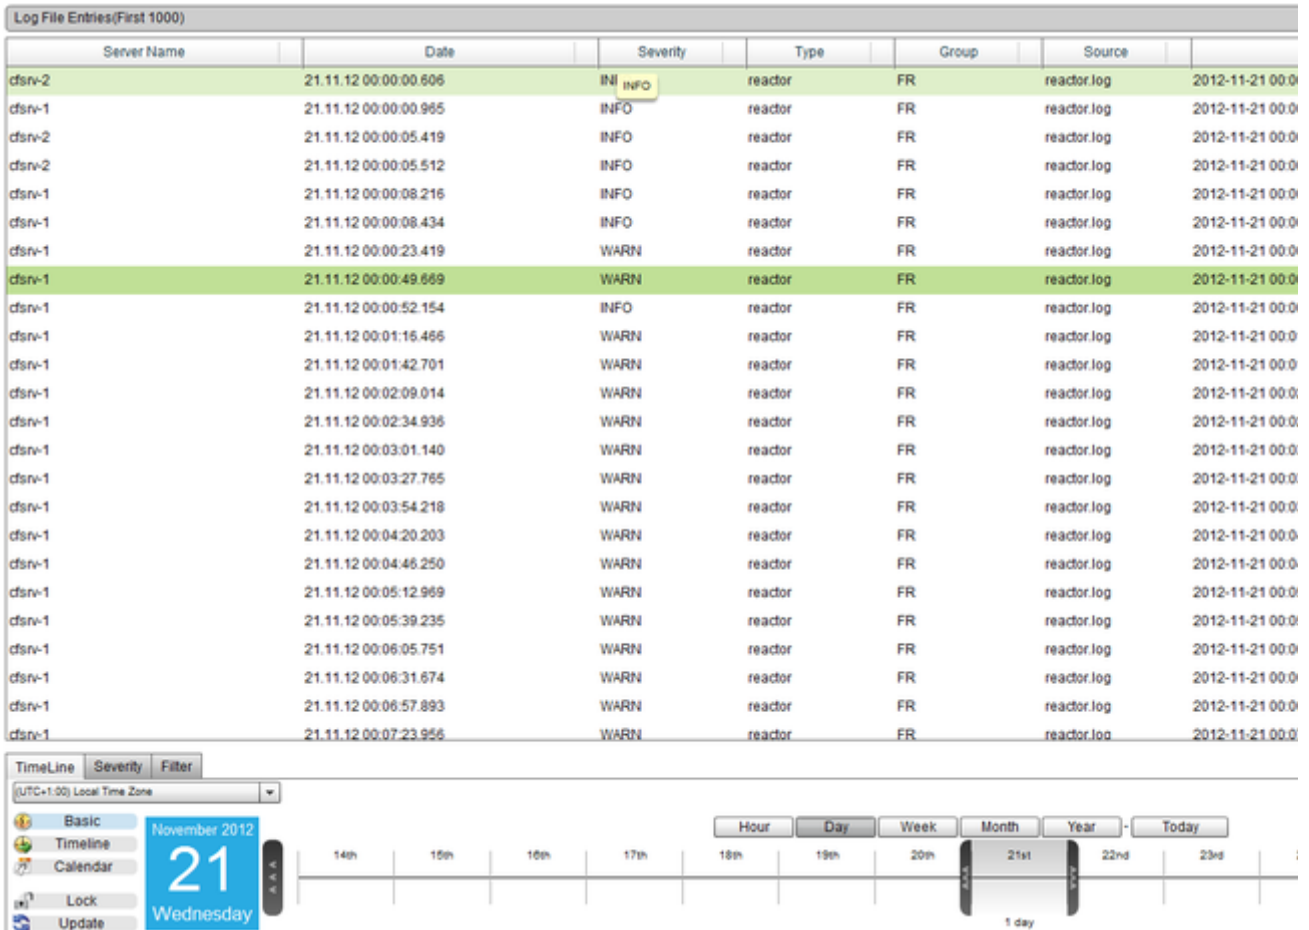


Figure 2: Cluster version of the Log File Entries perspective

Back to the top

Events

Events

Overview

Next Steps

Annotations

Placemarks

Crash Protection

Resources with Annotations

ANNOTATIONS

Annotations

Description

This perspective shows a table of the Annotations that were created within the specified time period and a Time Block Chart of when these

Usage

To change the order in which the annotations are shown in the table, click on one of the table column headings. Use the Date Navigation i perspective covers.

By hovering the mouse over the Annotation on the timeline, its details will be displayed in a popout box.

Perspective View

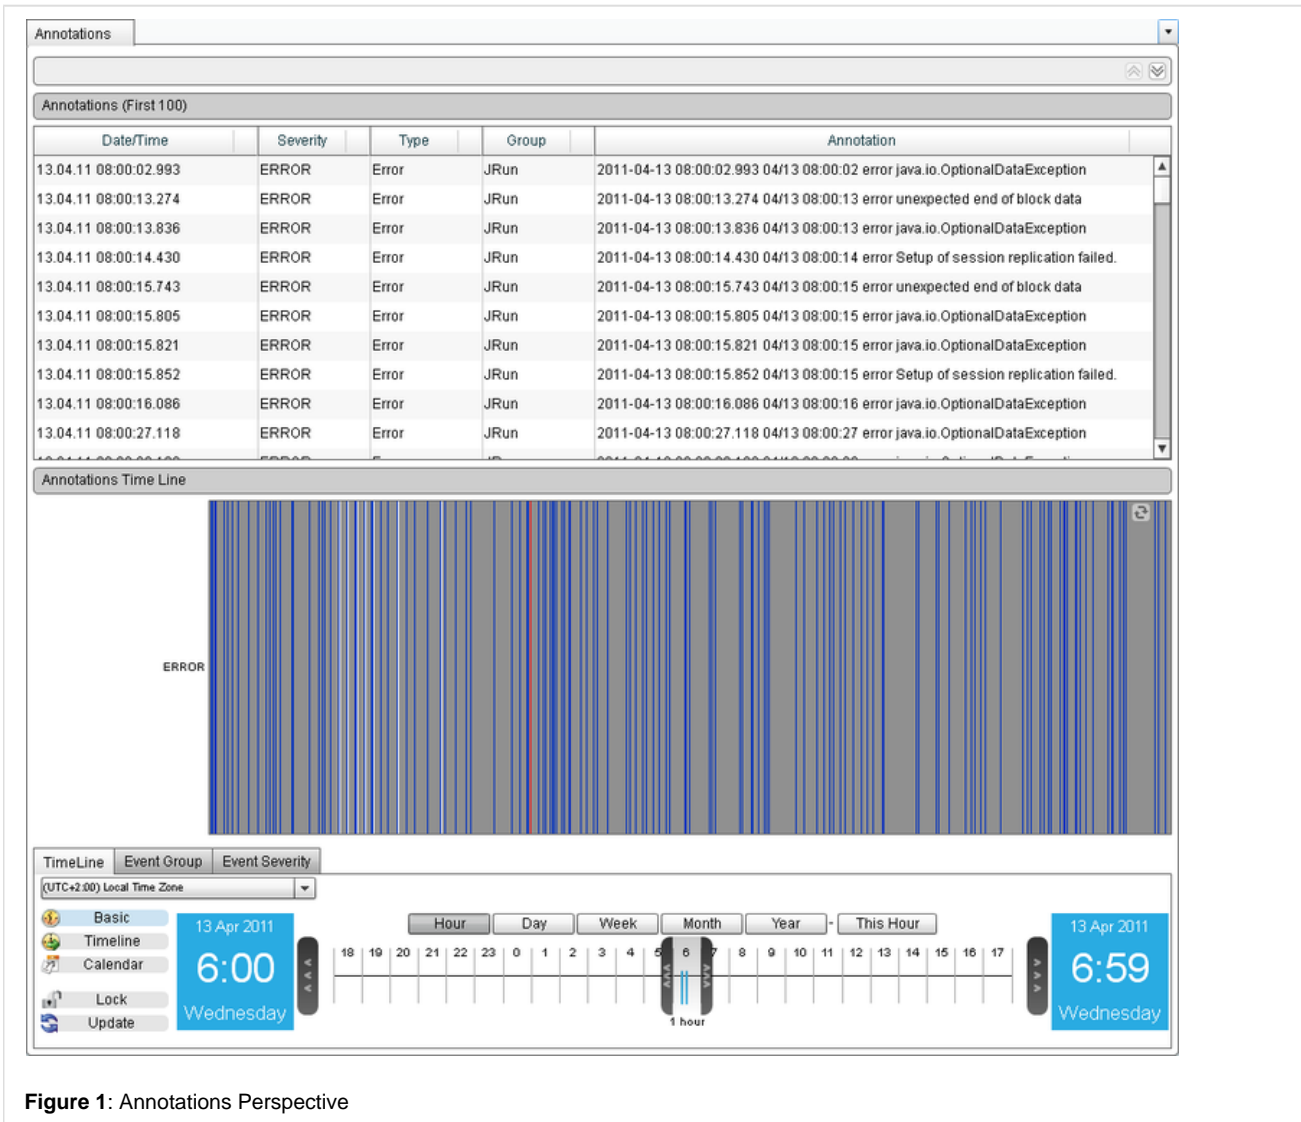


Figure 1: Annotations Perspective

Cluster Perspective

The cluster version of this perspective shows a table of the Annotations that were created within the specified time period, across the whole annotations occurred.

You can see which of the servers the Annotation relates to, by viewing the leftmost column in the table (named Server Name) or by hovering a popup is displayed.

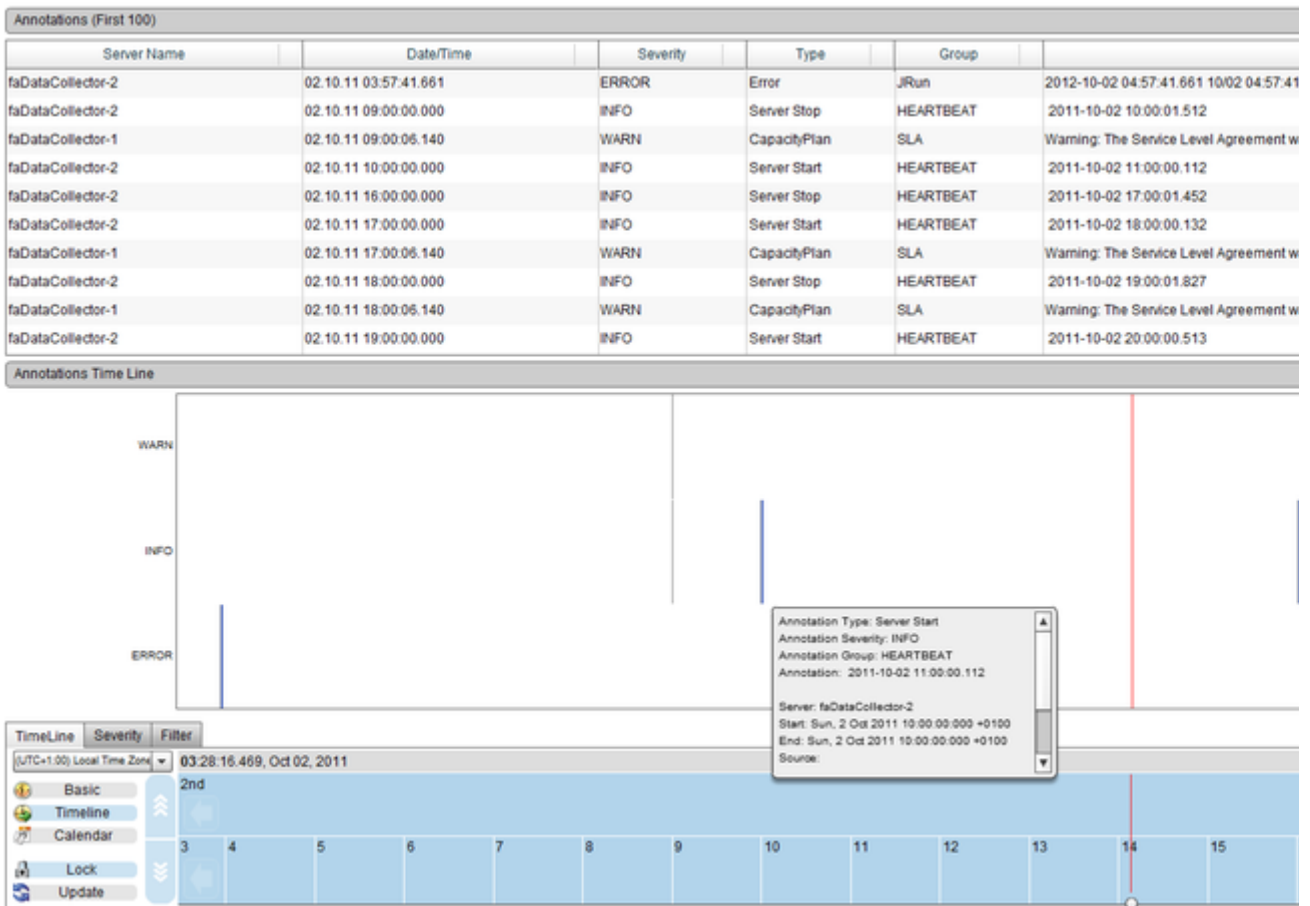


Figure 2: Cluster version of the Annotations perspective

Back to the top

PLACEMARKS

Placemarks

Description


This perspective shows a chart containing the Application Placemarks within the specific period of time.

Usage

When there are a lot of Placemarks within a certain period of time, they will be automatically grouped into blocks that are easier to read and When you zoom in to the auto grouped block, the Placemarks will ungroup and you will be able to see each individual one(see Fig.1).

You can see details of the Placemark by hovering the mouse over it and can make edits to the Placemark, such as the name and description.

Perspective View



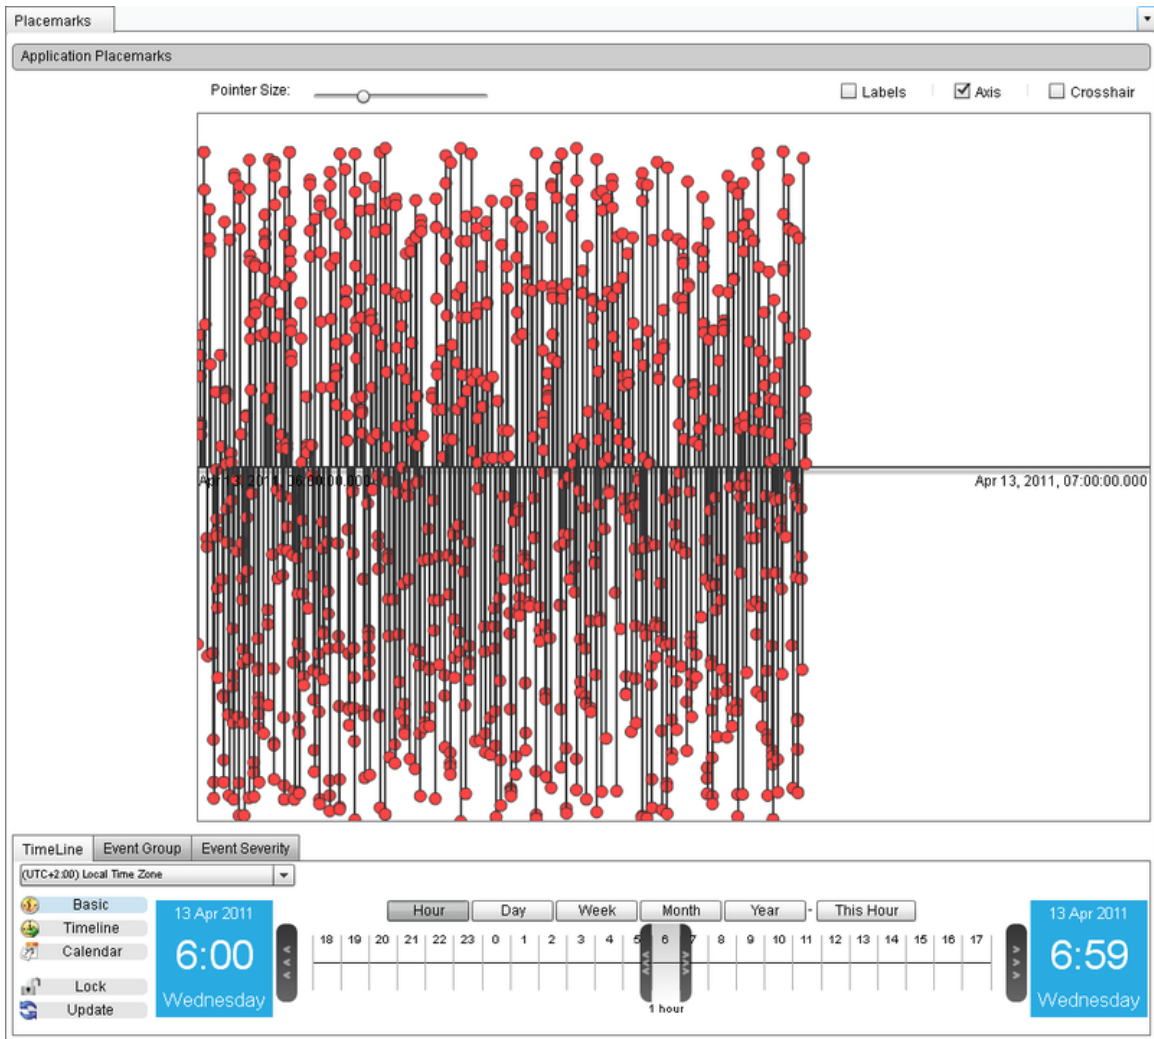


Figure 1: Placemarks Perspective

Cluster Perspective

The cluster version of this perspective will show all the Placemarks for all the servers in the cluster, within the specified time period. The placemarks will be grouped and displayed in the same way, and you will be able to see which of the servers the placemark relates to

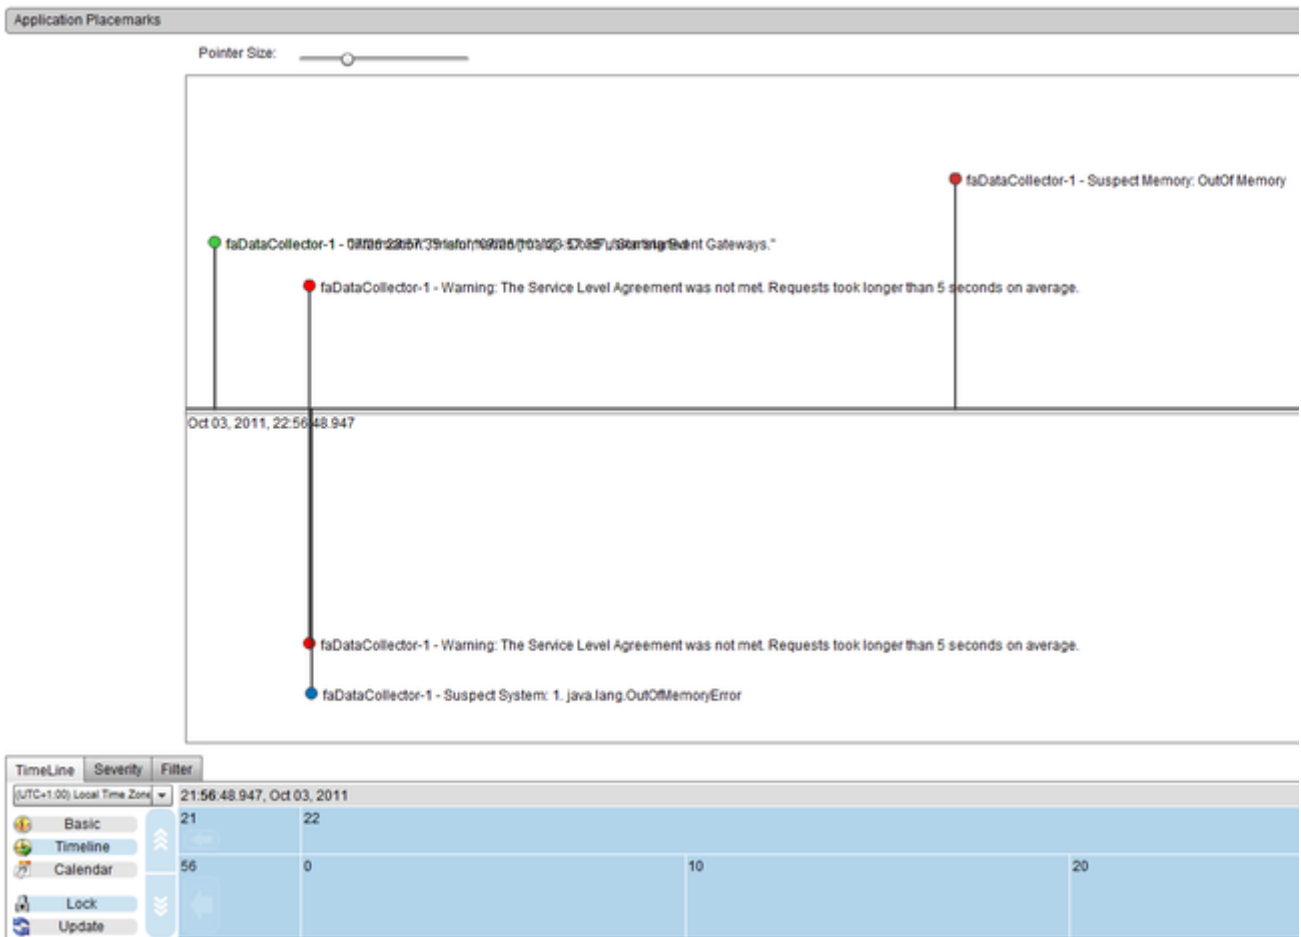


Figure 2: Cluster version of the Placemarks perspective

[Back to the top](#)

CRASH PROTECTION

Crash Protection

Description

This perspective shows the first 1000 Crash Protection notifications that occurred within the specified time period, and the reasons for the

Usage

To change the order in which the data is shown, click on one of the table column headings. Use the Date Navigation in order to change the You can right click on any of the rows in the table to bring up a context menu. This will allow you to navigate to other useful perspectives for and System Metrics.

Perspective View

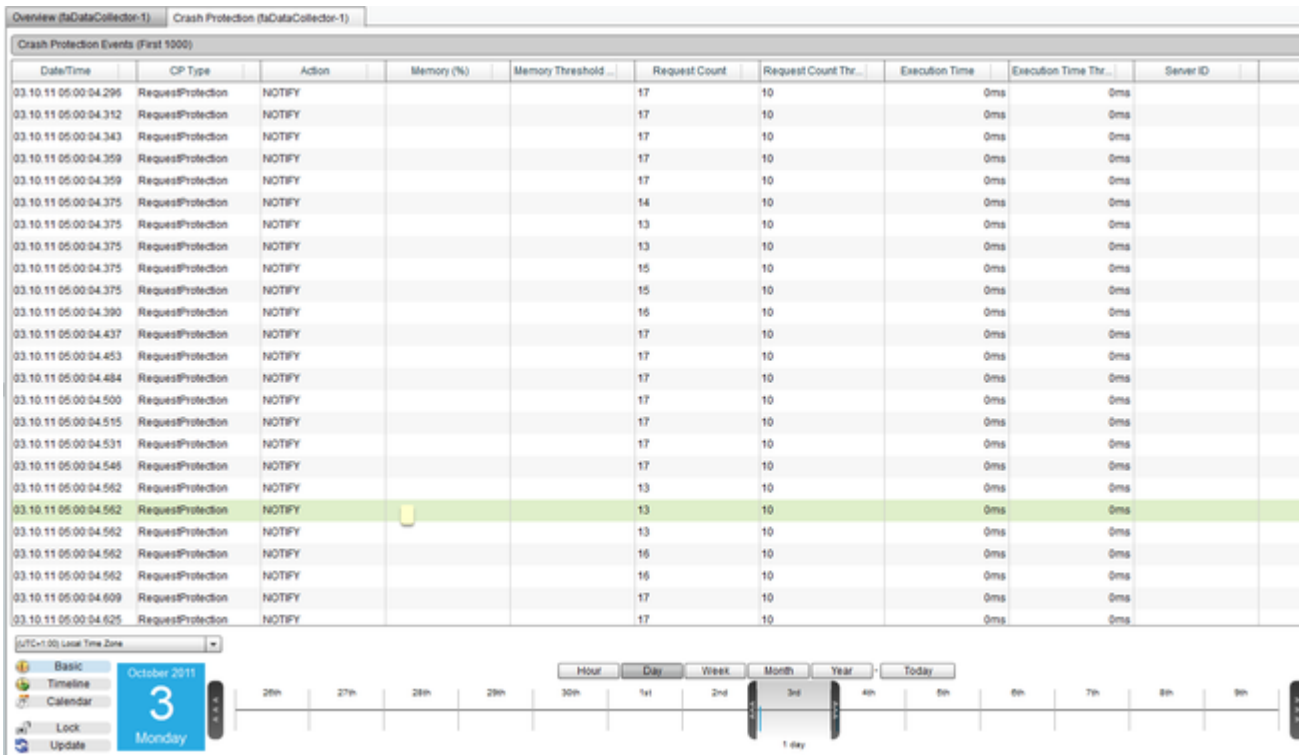


Figure 1: Crash Protection Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Events > Crash Protection from It shows the first 1000 Crash Protection events from the entire cluster, ordered by Date/time. You can see which server each event came from and it displays Server Name.

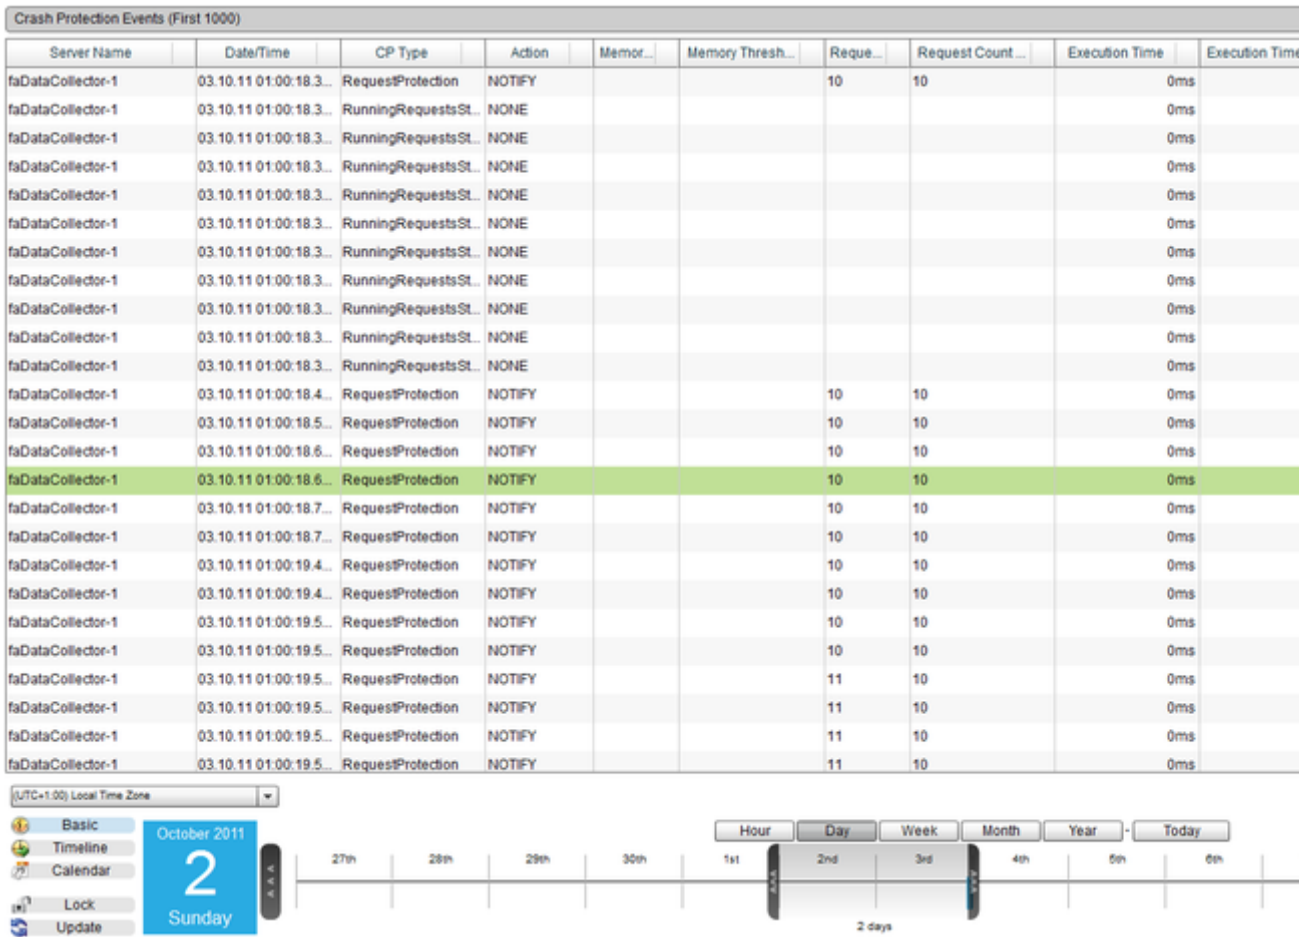


Figure 2: Cluster version of the Crash Protection perspective

[Back to the top](#)

RESOURCES WITH ANNOTATIONS

Resources with Annotations

Description

This perspective has three different views that you can switch between using the drop down box in the uppermost right hand corner of the. The first view shows a chart of the Process and System CPU Usage within the specific period of time.

The second view shows a chart of the Used Memory and Maximum Memory in kb within the specific period of time.


The third view shows a Time Block Chart of the annotations that occurred during this time period.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers.

You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining ; By hovering the mouse over the Annotation on the timeline, its details will be displayed in a popout box.

Perspective View






Figure 1: Resources with Annotations Perspective

Cluster Perspective

The cluster version of this perspective has all the views from the Cluster version of the CPU Usage, Memory Usage and Annotations persp. You can switch between the views using the drop down box in the uppermost right hand corner of the perspective




Figure 2: Cluster version of the 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 \(Hot-Spots\)](#)
- [Requests Restart Coverage](#)

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

This perspective shows a table of the first 100 times that the Memory of the server reached 100% usage.

To change the order in which the data is shown, click on one of the table column headings. Use the Date Navigation in order to change the You can right click on any of the rows in the table to bring up a context menu. This will allow you to navigate to other useful perspectives for and System Metrics.

[Perspective View](#)

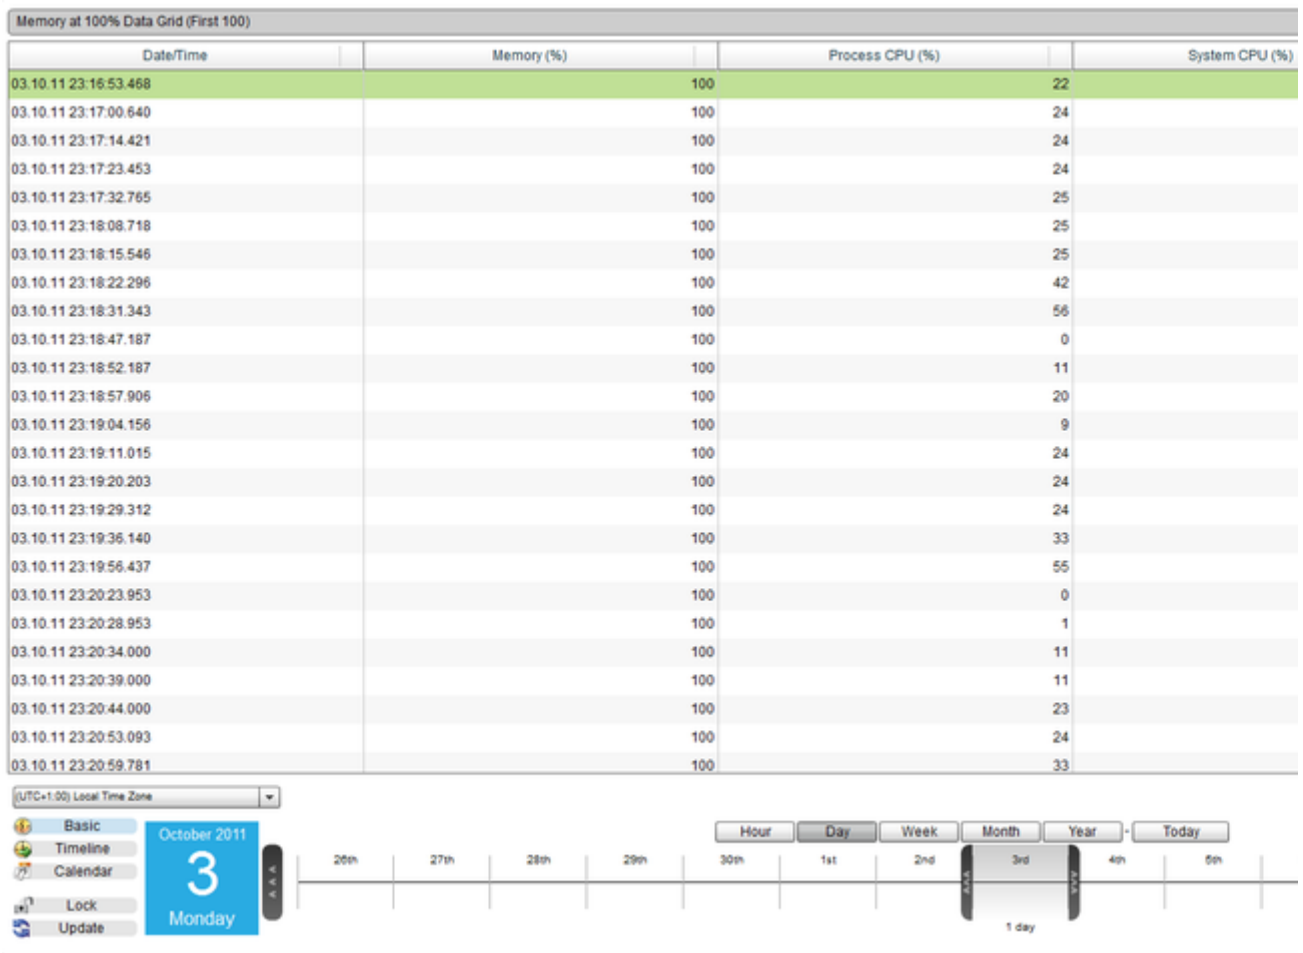


Figure 1: Memory at 100% Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Hot-Spots > Memory Hot-Spots > Comparison Grid.

It shows the first 100 times that the Memory of the any of the servers in the cluster reached 100% usage. You can see which server the data belongs to which displays Server Name.

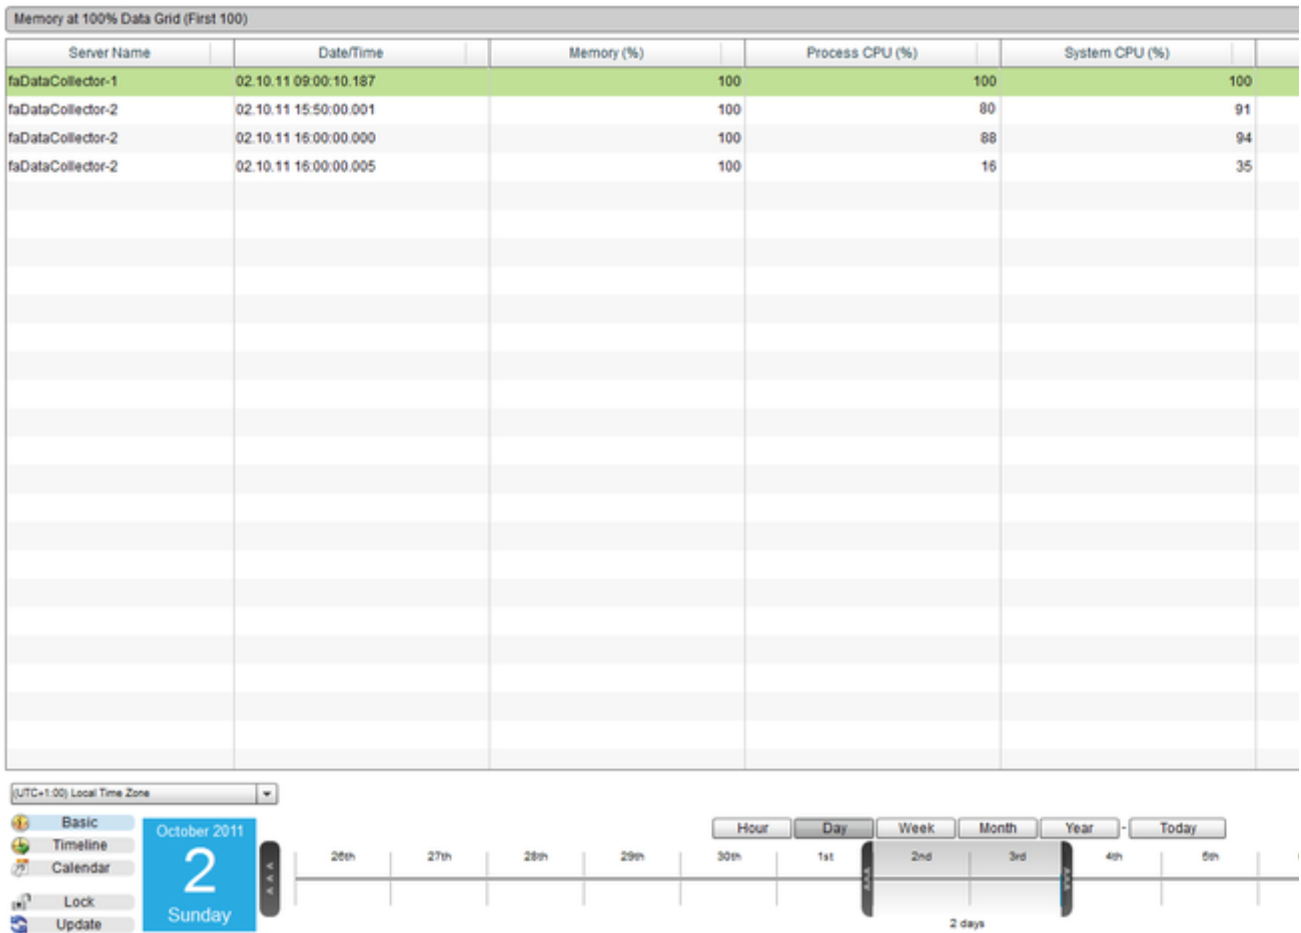


Figure 2: Cluster version of the Memory at 100% perspective

[Back to the top](#)

Memory Highest Usage


Memory Highest Usage

Description

This perspective shows a table of the top 100 times that had the 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 covers.

Perspective View



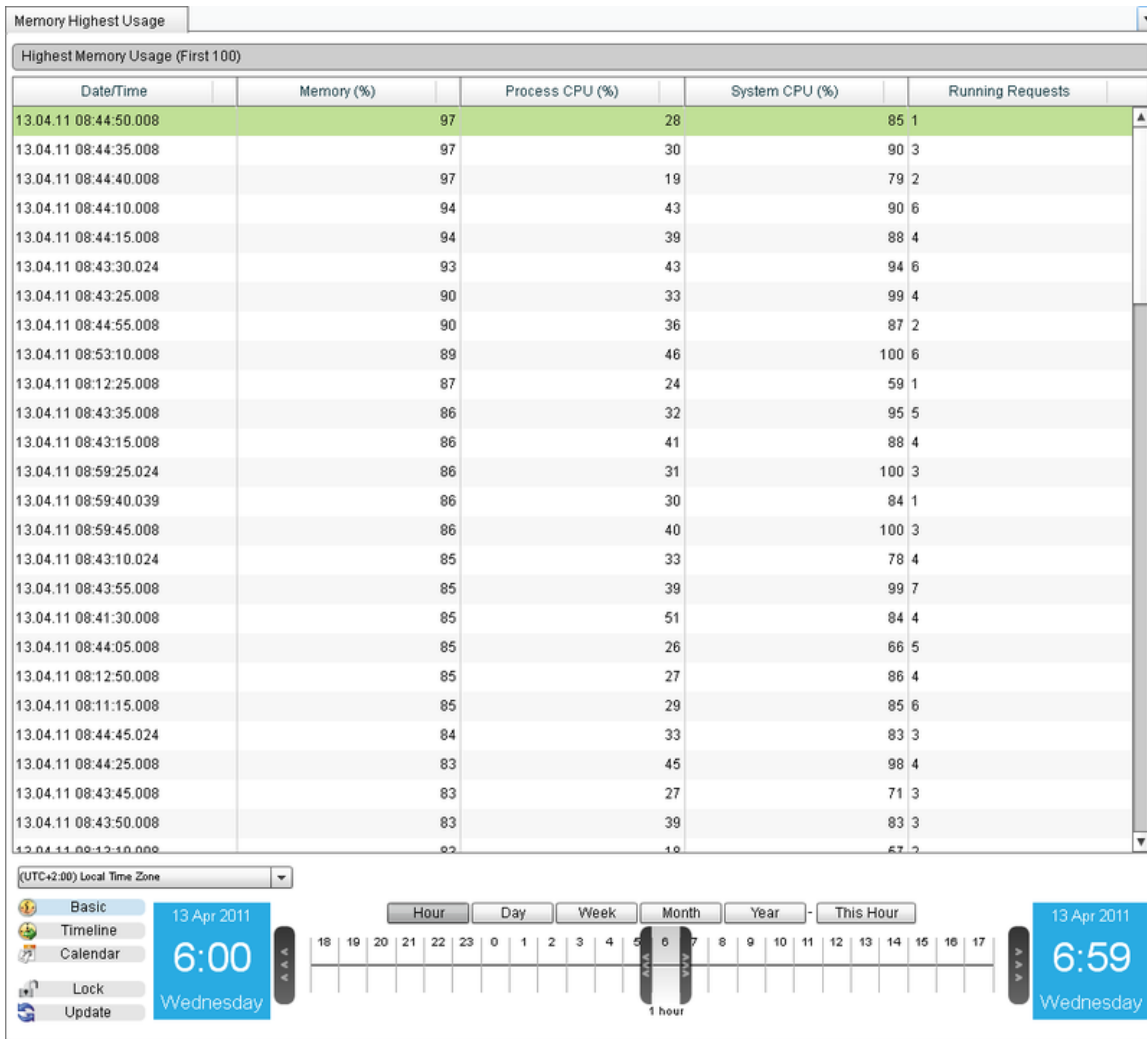


Figure 1: Memory Highest Usage Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Hot-Spots > Memory Hot-Spots > Comparison Grid.

It shows the top 100 instances of high memory usage from any of the servers in the cluster. You can see which server the data came from Server Name.

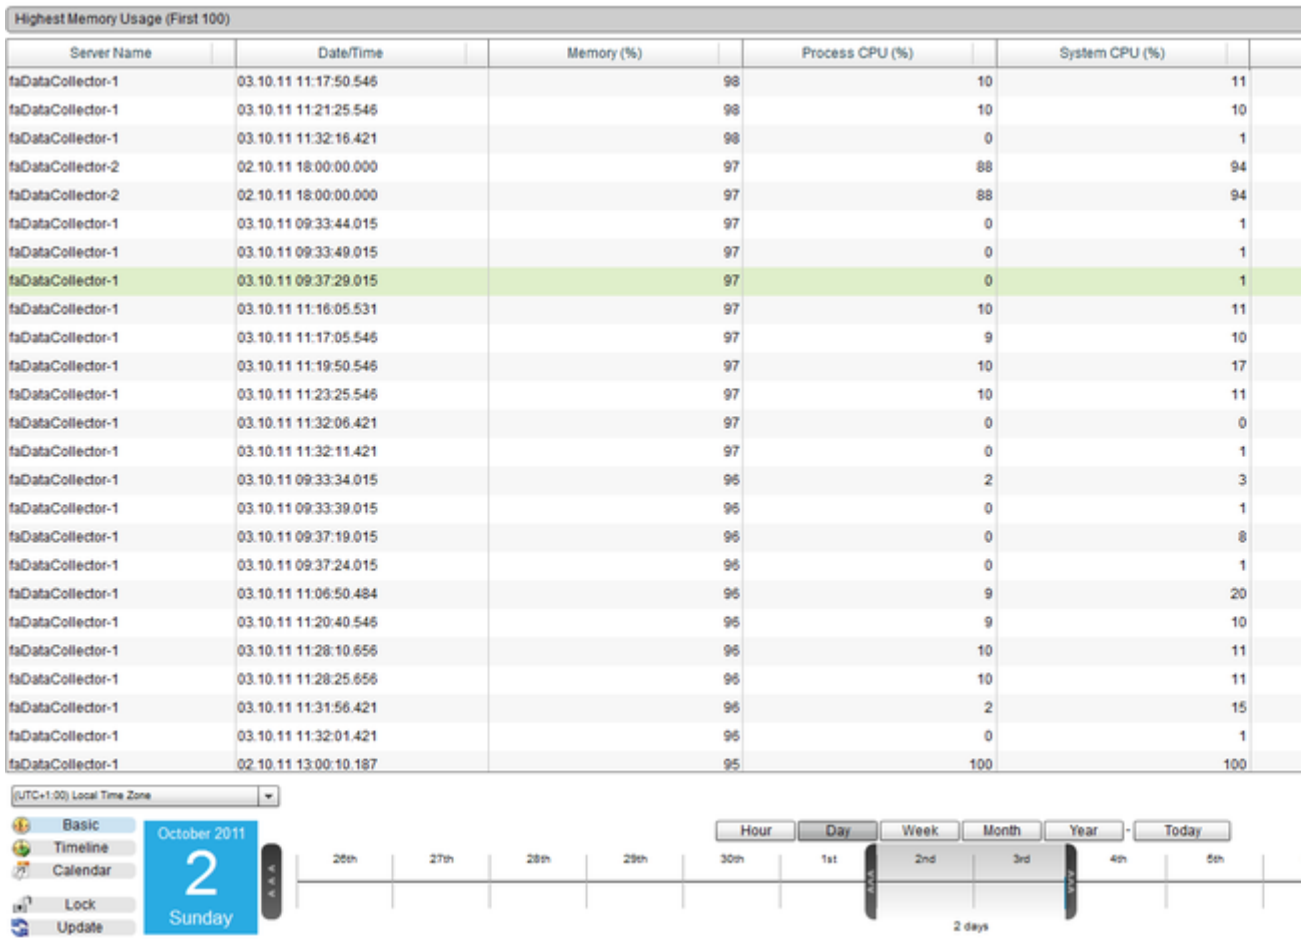


Figure 2: Cluster version of the Memory Highest Usage perspective

Back to the top

Memory Events

Memory Events

Description

This perspective shows a table of Memory Annotations that were created within the specified time period and a Time Block Chart of when Usage

To change the order in which the annotations are shown in the table, click on one of the table column headings. Use the Date Navigation i perspective covers.

By hovering the mouse over the Annotation on the timeline, its details will be displayed in a popout box.

Perspective View

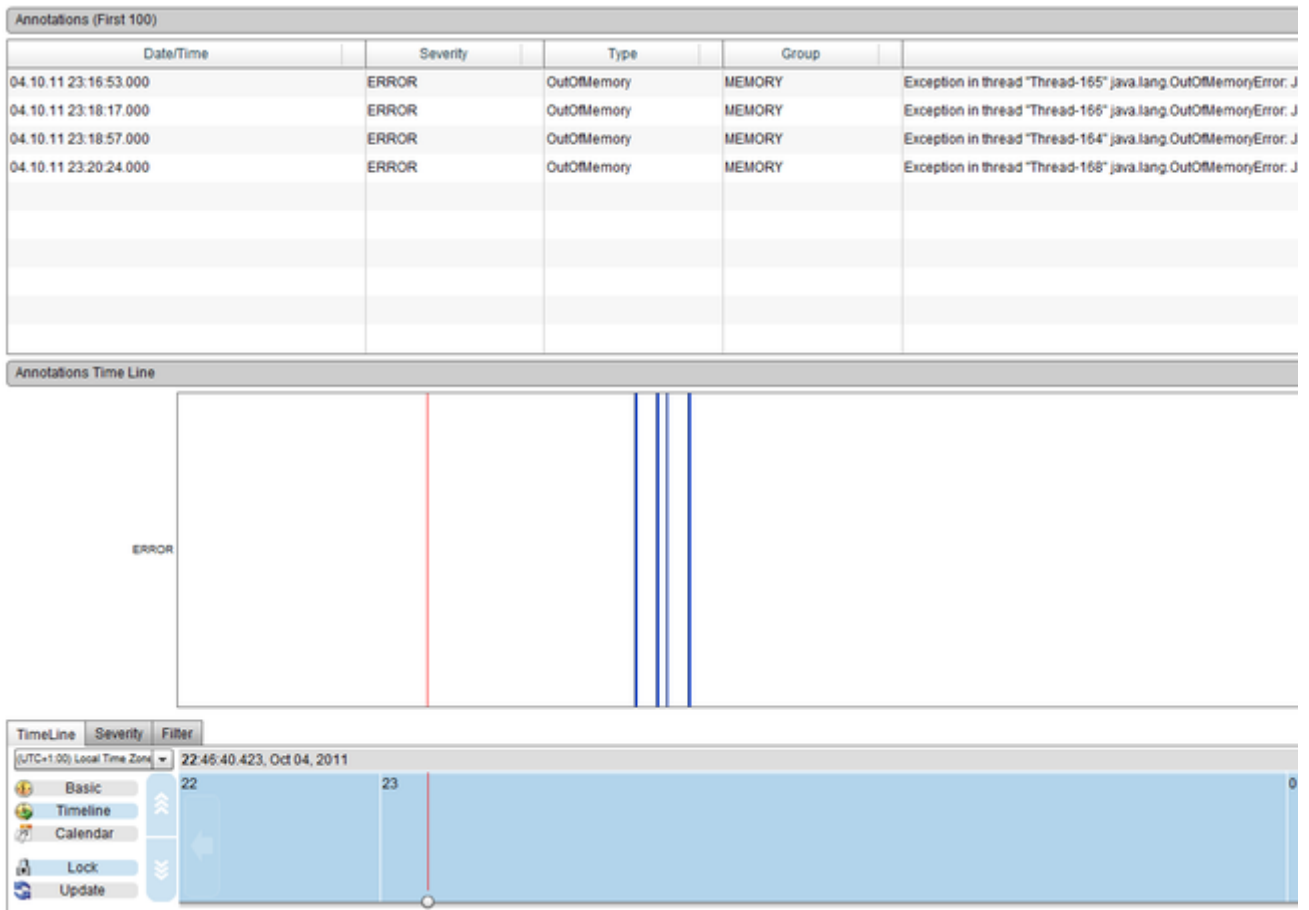


Figure 1: Memory Events Perspective

Cluster Perspective

The cluster version of this perspective shows a table of the Memory Events that were created within the specified time period, across the v events occurred.

You can see which of the servers the Event relates to, by viewing the leftmost column in the table (named Server Name) or by hovering over displayed.

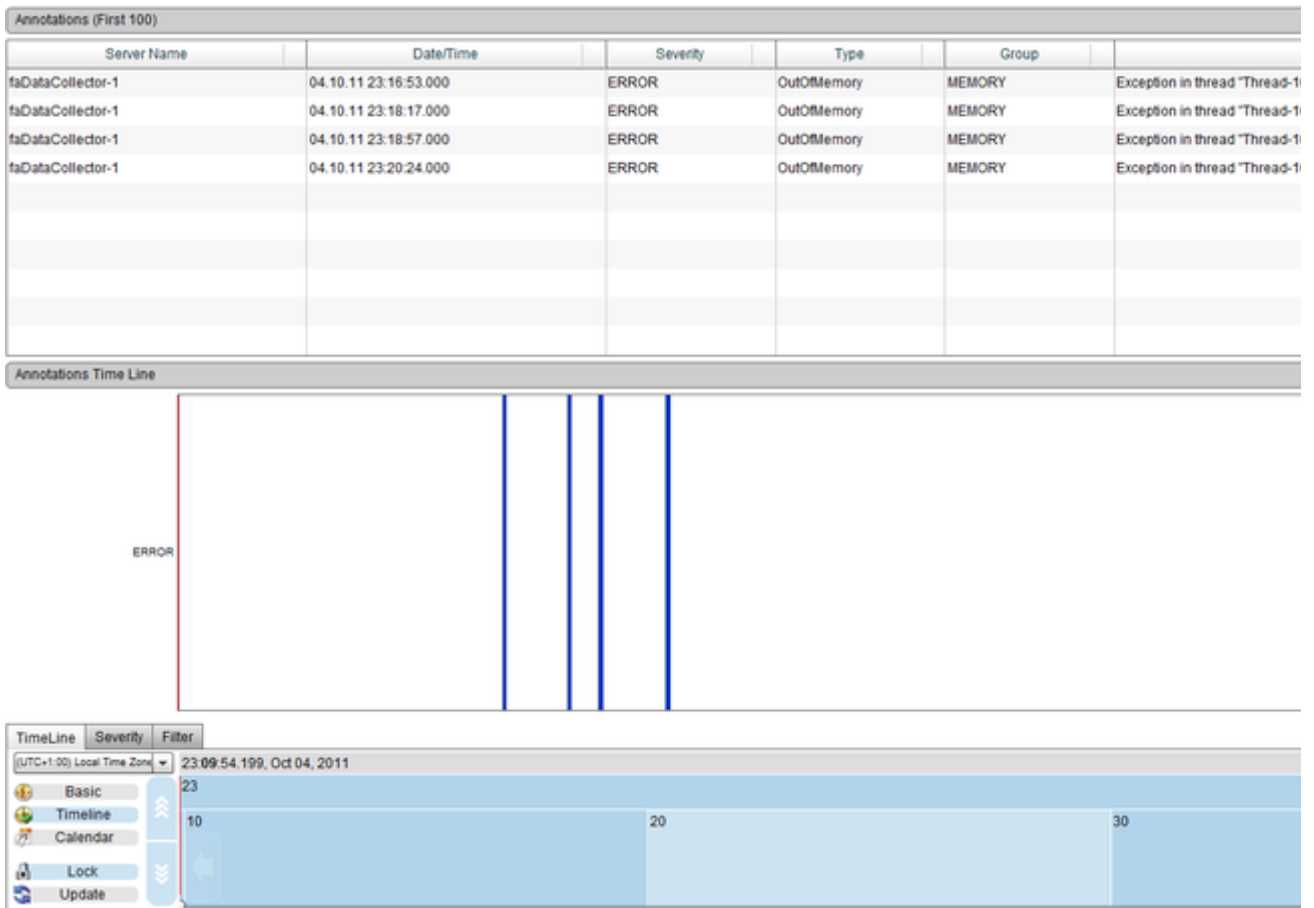


Figure 2: Cluster version of the Memory Events perspective

[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

This perspective shows the first 100 times that the Process CPU of the server reached 100%, within the specified time period. Usage

To change the order in which the data is shown, click on one of the table column headings. Use the Date Navigation in order to change the You can right click on any of the rows in the table to bring up a context menu. This will allow you to navigate to other useful perspectives for and System Metrics.

Perspective View

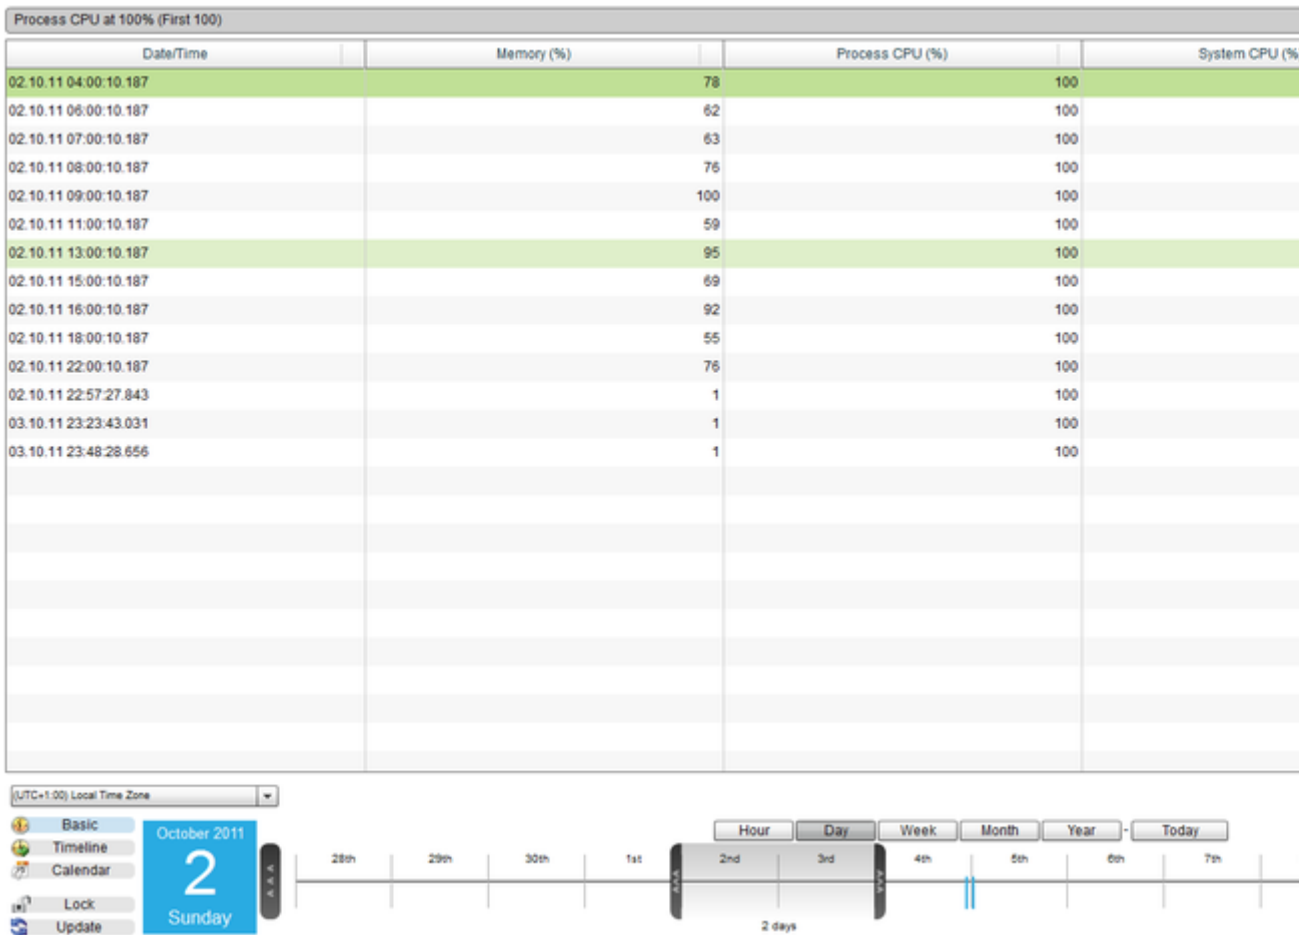


Figure 1: Process CPU at 100% Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Hot-Spots > CPU Hot-Spots > Comparison Grid. It shows the first 100 times that the Process CPU of any of the servers in the cluster reached 100% usage. You can see which server it is in the table, which displays Server Name.

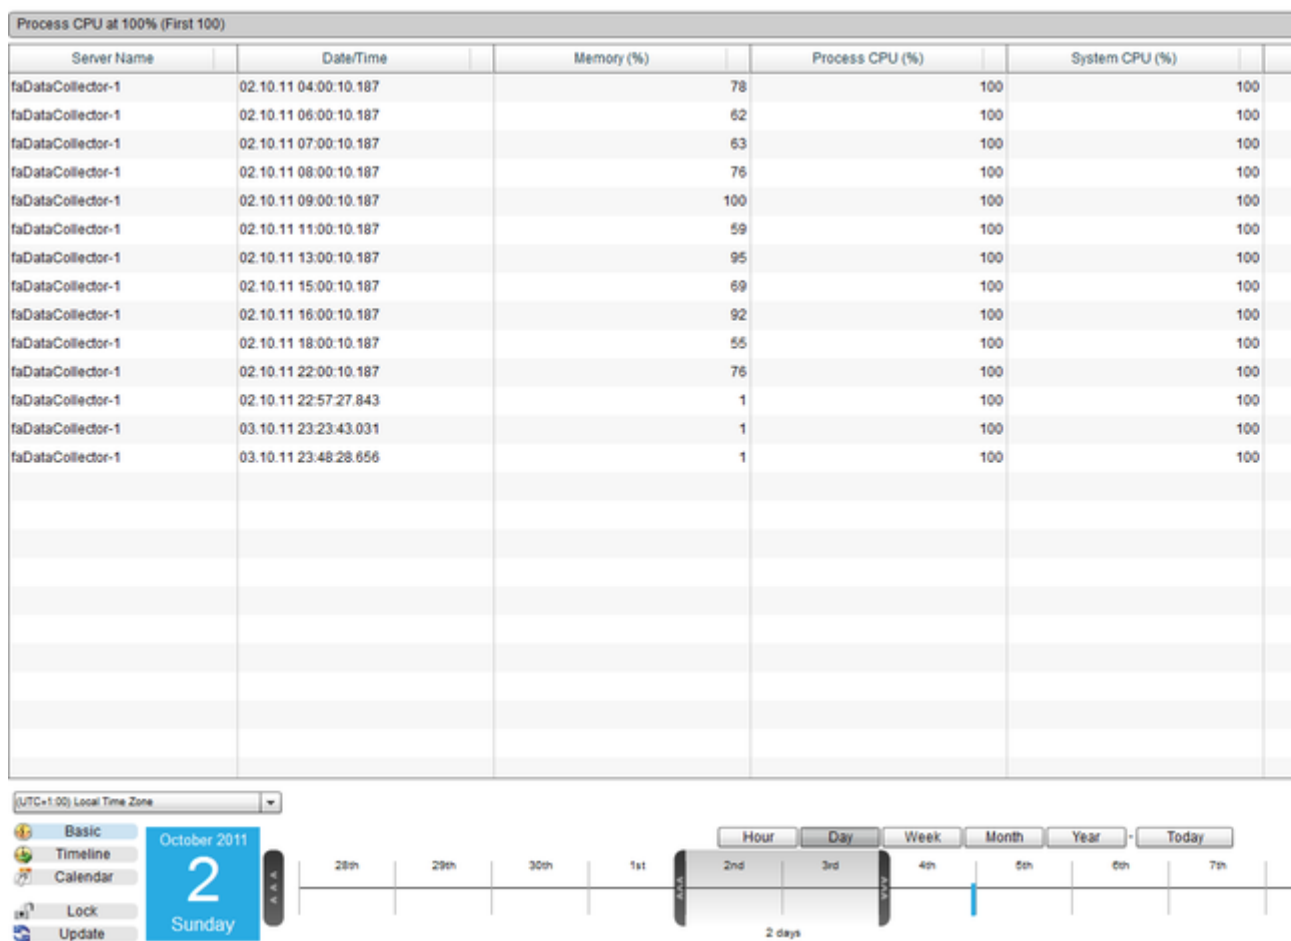


Figure 2: Cluster version of the Process CPU at 100% perspective

[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 shown, click on one of the table column headings. Use the Date Navigation in order to change covers.

Perspective View

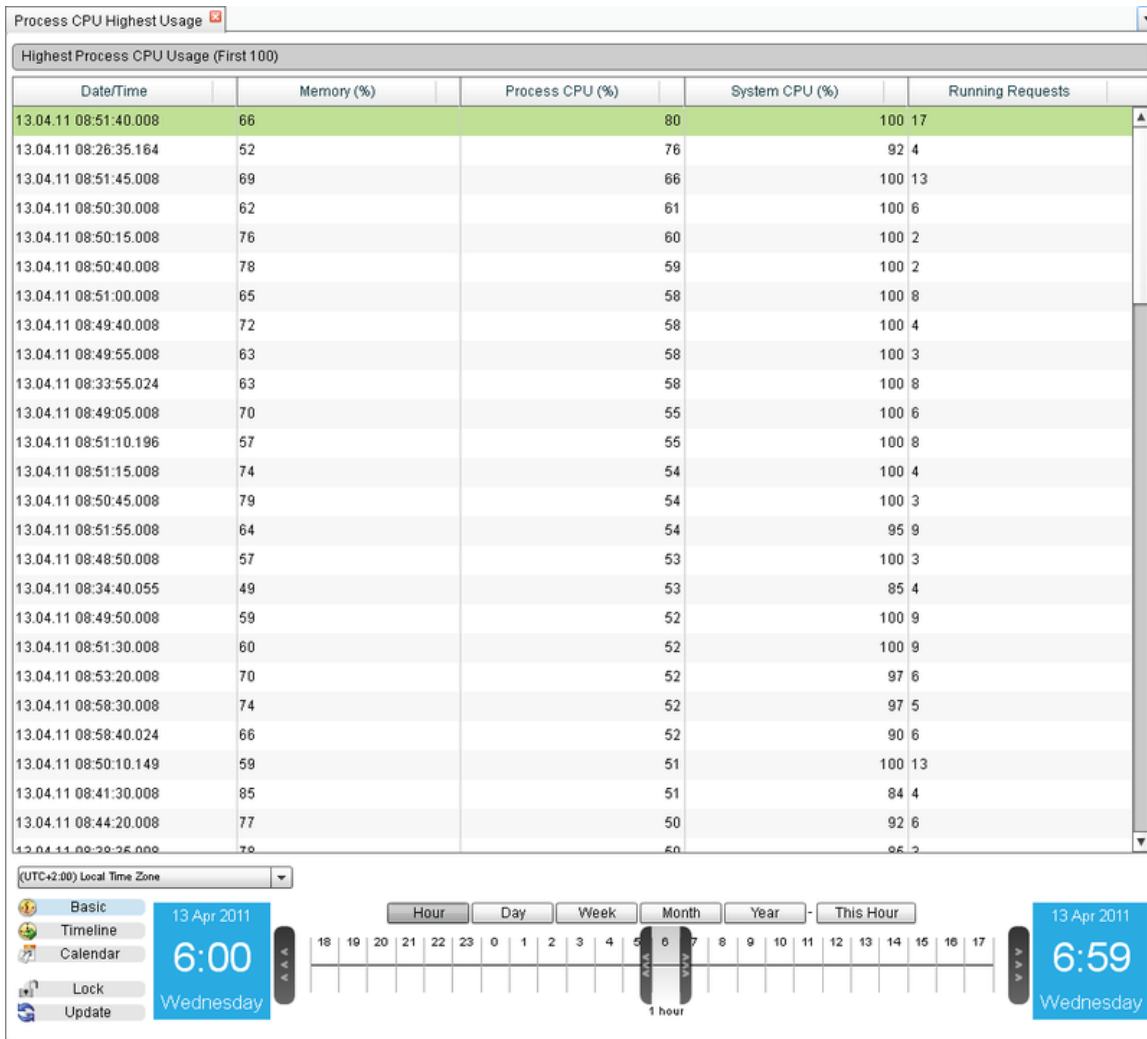


Figure 1: Process CPU Highest Usage Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Hot-Spots > CPU Hot-Spots > Pr Comparison Grid.

It shows the first 100 Highest Process CPU Usage times of the any of the servers in the cluster. You can see which server the data came from by displaying Server Name.

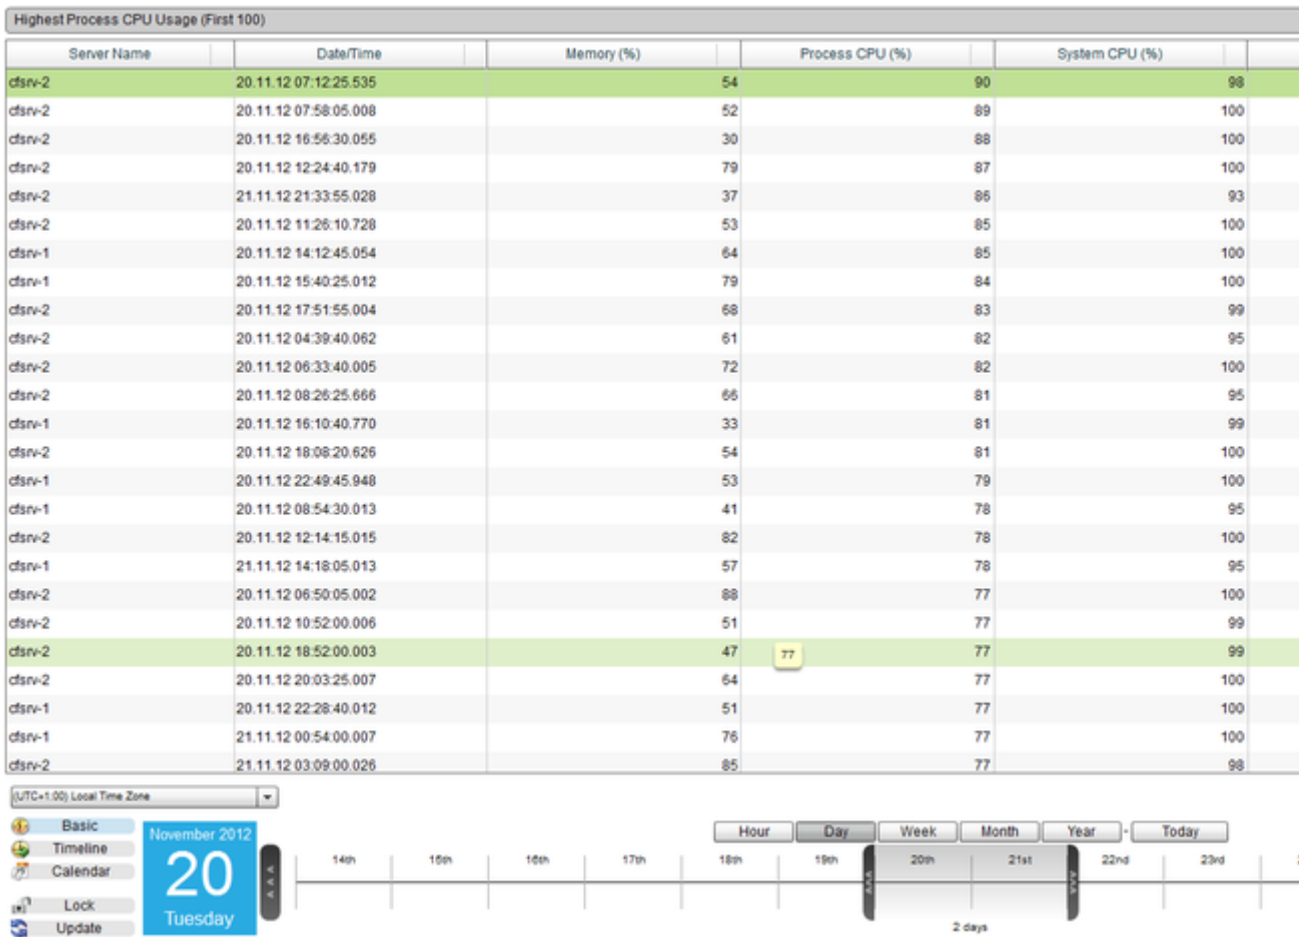


Figure 2: Cluster version of the Process CPU Highest Usage perspective

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 shown, click on one of the table column headings. Use the Date Navigation in order to change covers.

Perspective View

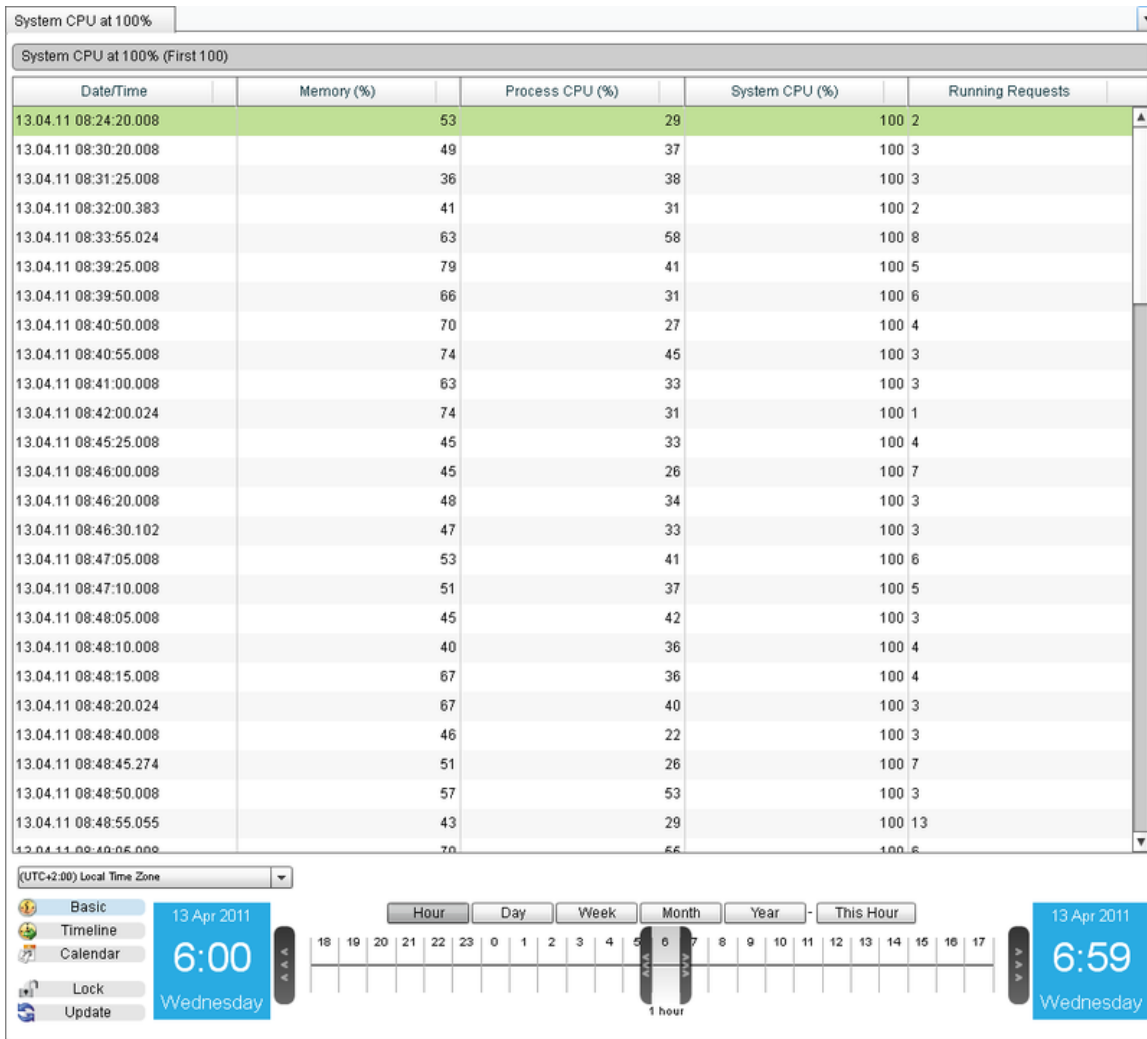


Figure 1: System CPU at 100% Perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Hot-Spots > CPU Hot-Spots > System Comparison Grid.

It shows the first 100 times that the System CPU of any of the servers in the cluster reached 100% usage. You can see which server triggered the event, which displays Server Name.

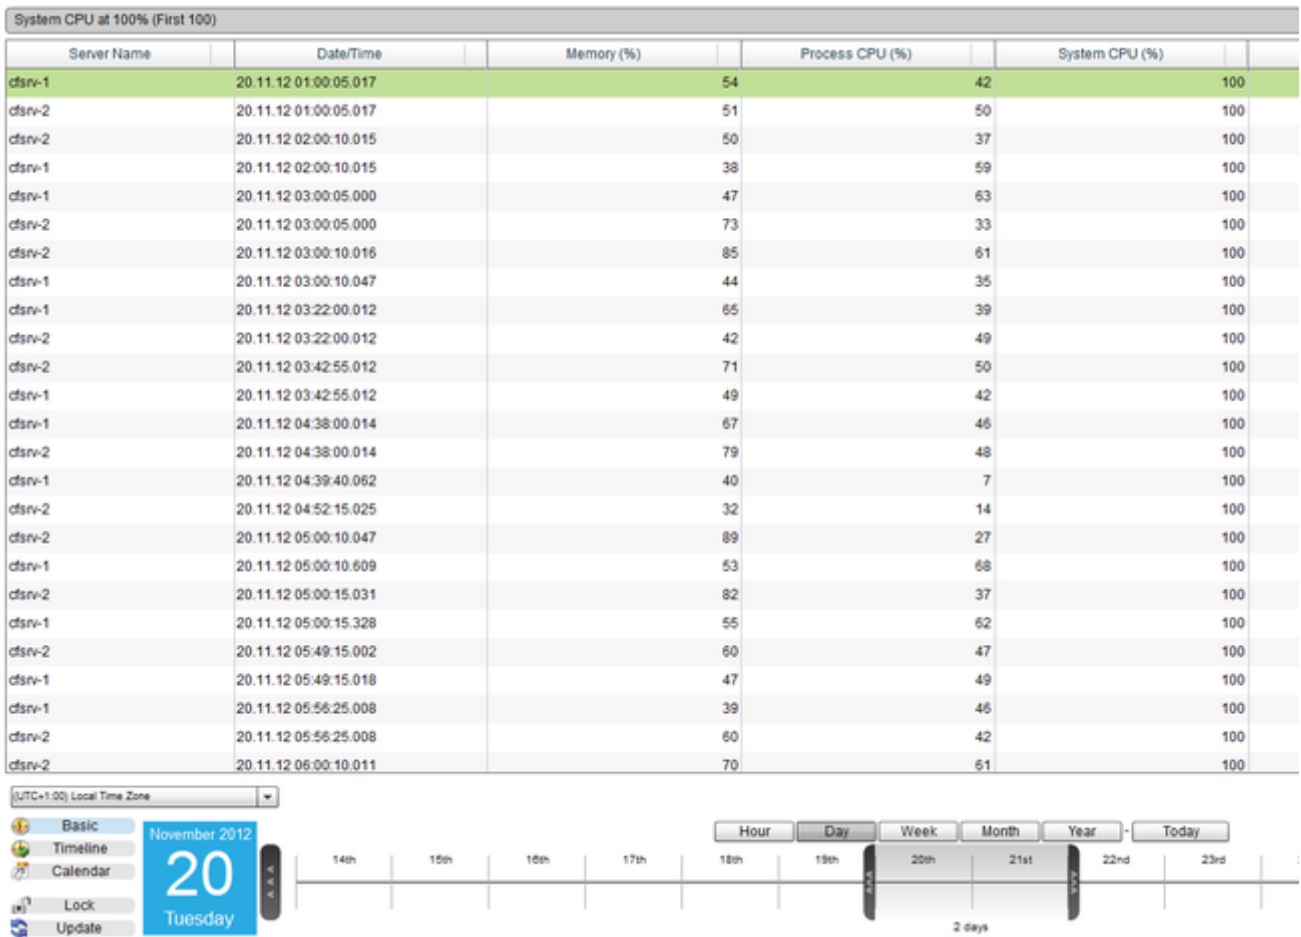


Figure 2: Cluster version of the System CPU at 100% perspective

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 covers.

Perspective View

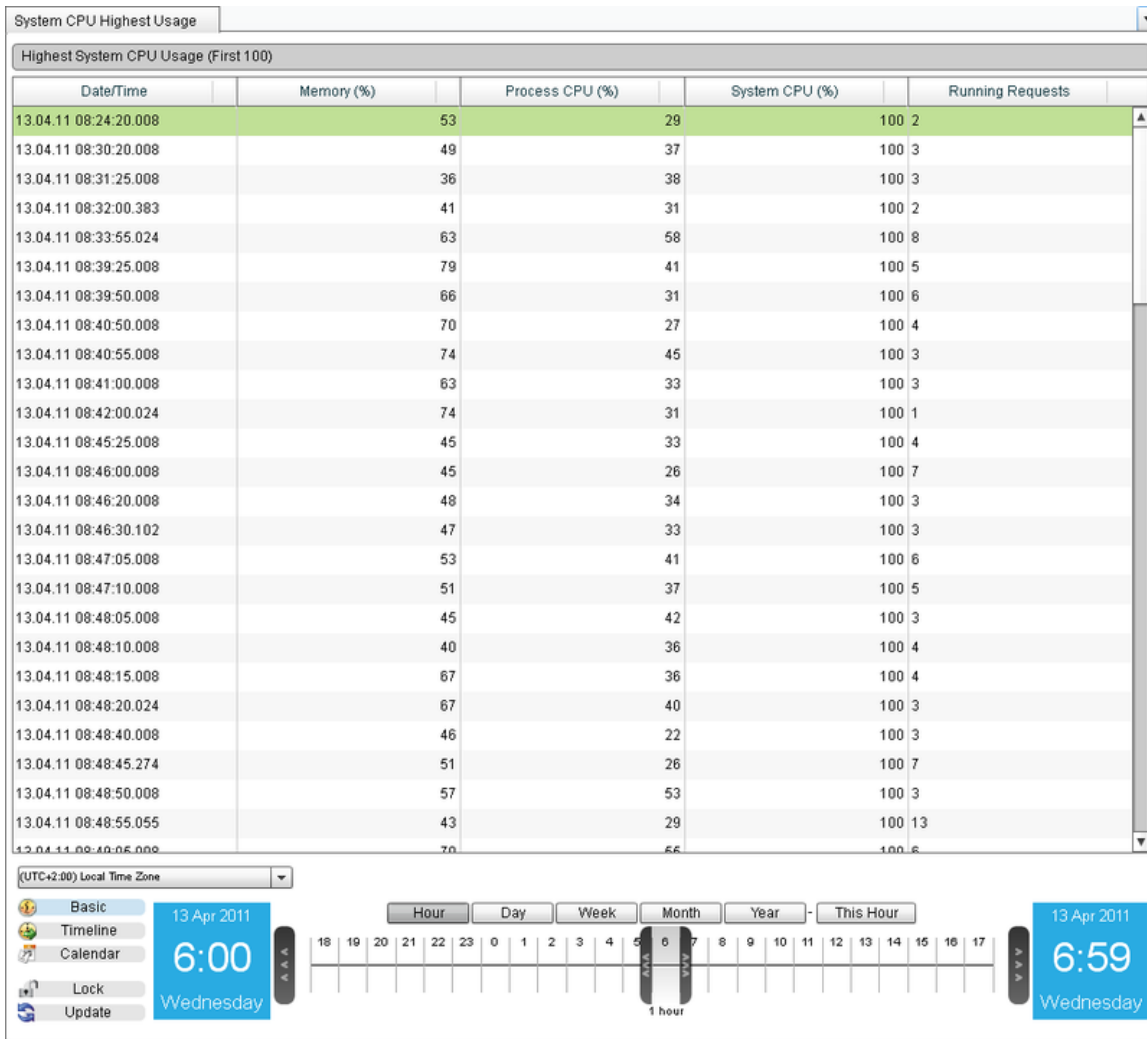


Figure 1: System CPU Highest Usage

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Hot-Spots > CPU Hot-Spots > Server Comparison Grid.

It shows the top 100 instances of high System CPU Usage from any of the servers in the cluster. You can see which server the data came from by displaying the Server Name.

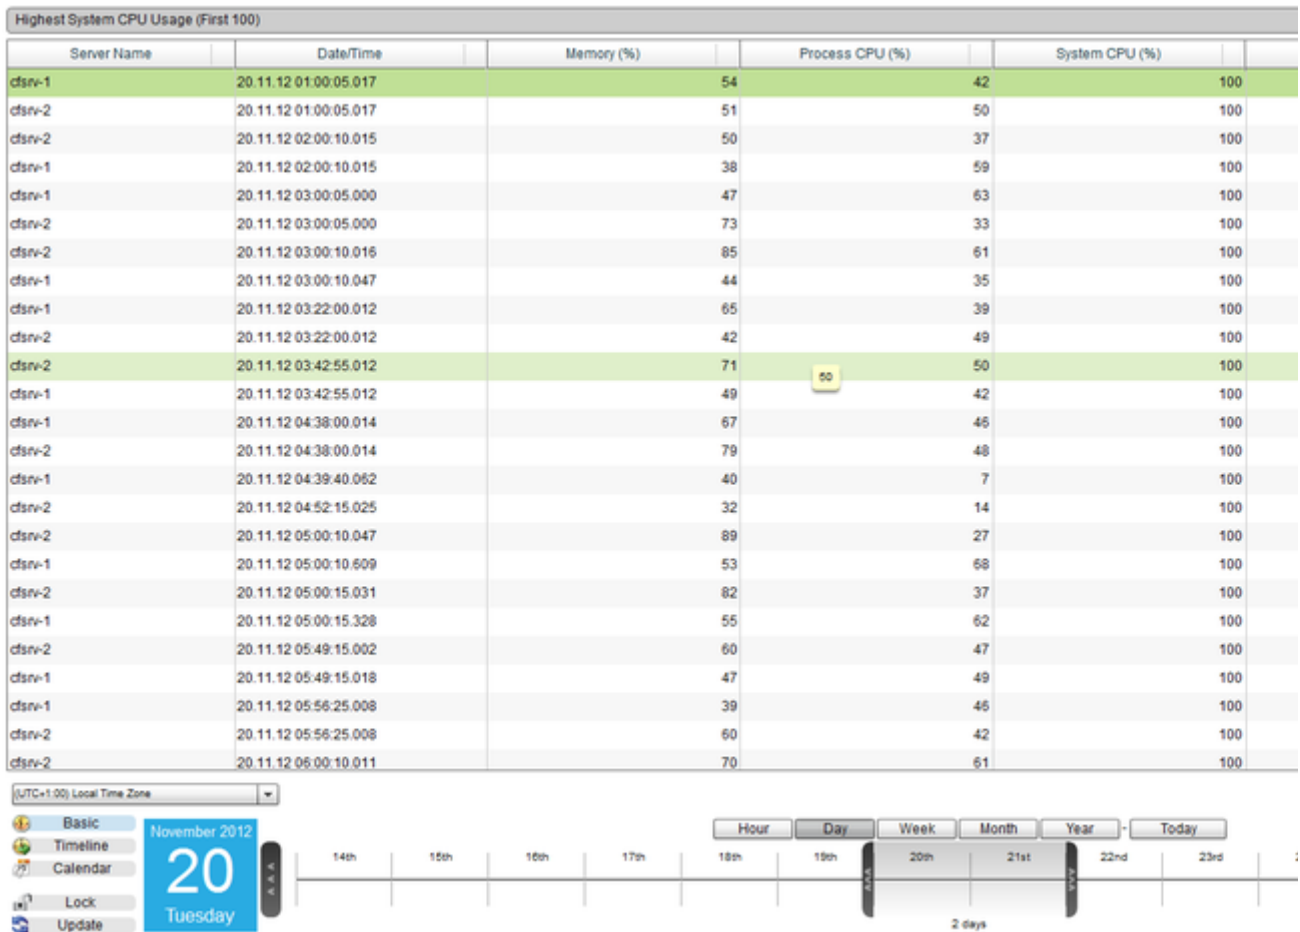


Figure 2: Cluster version of the System CPU Highest Usage perspective

[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 Restart Coverage

Requests Not Completed

Requests Not Completed

Description

This perspective shows a Table of the first 100 requests that failed to complete within the specified time period, and a Time Block Chart showing CPU Usage

To change the order in which the requests are shown in the table, click on one of the table column headings. Use the Date Navigation in this perspective covers.

By hovering the mouse over the requests on the timeline, its details will be displayed in a popout box.

Perspective View

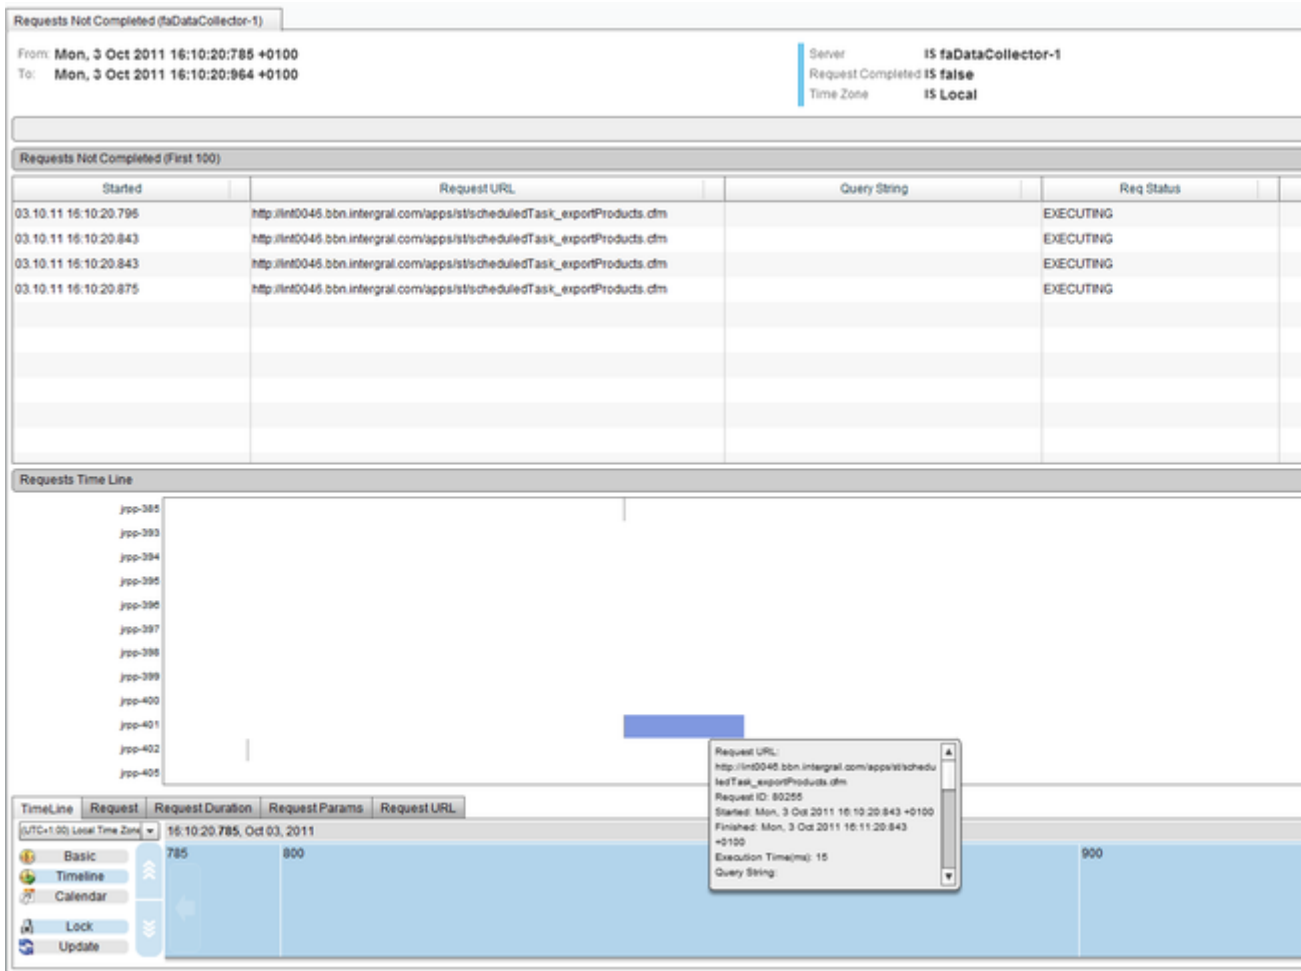


Figure 1: Requests Not Completed perspective

Cluster Perspective

The Cluster version of this perspective, found under Clusters > "Cluster Name" > Deep Server Analysis > Hot-Spots > Request Hot-Spots : Comparison Grid.

It shows the first 100 requests that were not completed during the time frame, from any of the servers in the cluster. You can see which server the table, which displays Server Name.

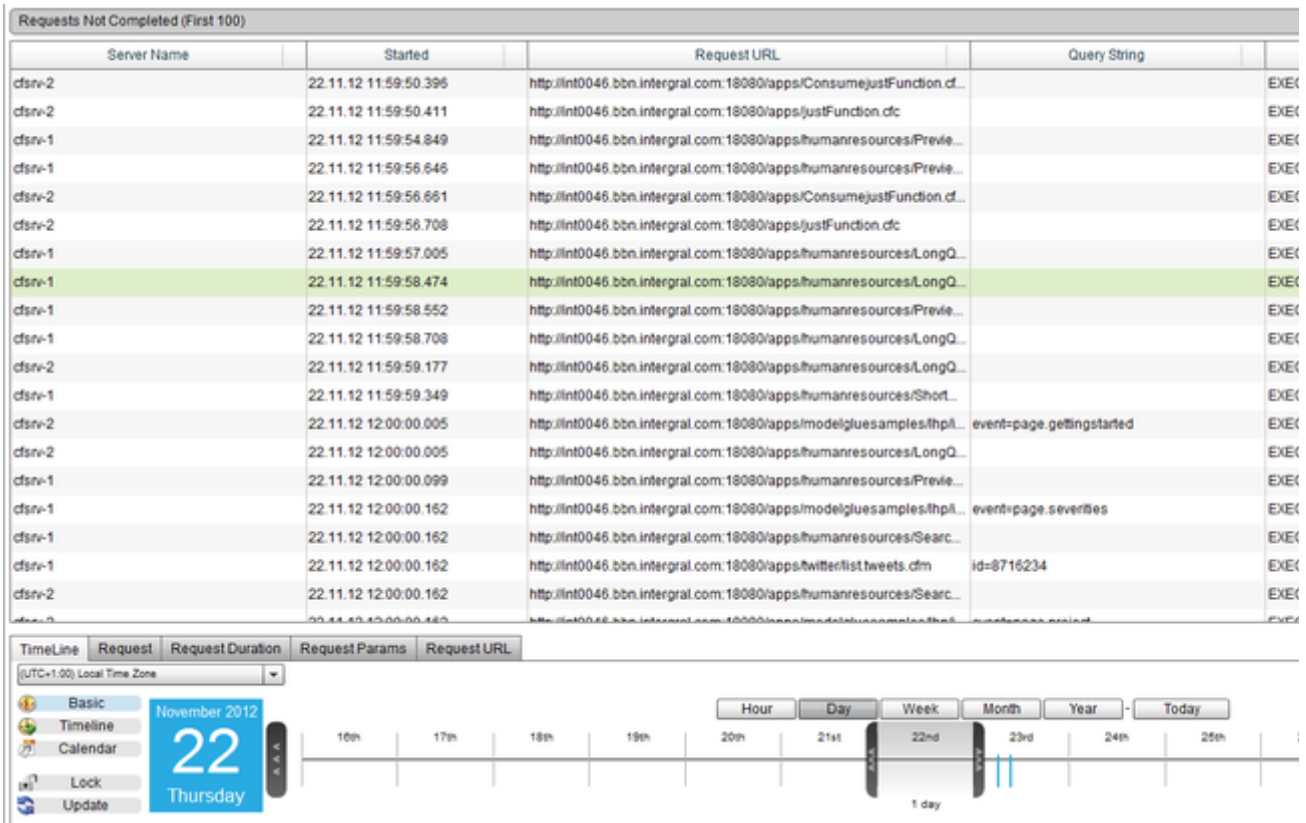


Figure 2: Cluster version of the Requests Not Completed perspective

[Back to the top](#)

Slowest Requests (Hot-Spots)

Slowest Requests (Hot-Spots)

See Slowest Requests

[Back to the top](#)

Slowest Requests On Average (Hot-Spots)

Slowest Requests On Average (Hot-Spots)

See Slowest Requests On Average

[Back to the top](#)

Slowest DB Requests On Average (Hot-Spots)

Slowest DB Requests On Average (Hot-Spots)

See Slowest DB Requests On Average

[Back to the top](#)

Requests Restart Coverage

Requests Restart Coverage

Description

This perspective makes it easy to identify commonly executing requests at the time of a server failure. The report offers the number of times a request failed and the percentage coverage over all server failures in a given time-frame. It shows the top 100 requests based on the number of times a request failed, for any given request URL within the specified time period. The number of restarts that the request failed on and the number of covered restarts as a percentage of the total number of restarts that occurred (Restart Coverage%) are also shown.

Usage

To change the order in which the requests are shown in the table, click on one of the table column headings. Use the Date Navigation in the perspective covers.

From easily identifying the failure URL, you can then drill into individual failures identifying the root-cause of each - often finding a simple fix for the development team.

Perspective View

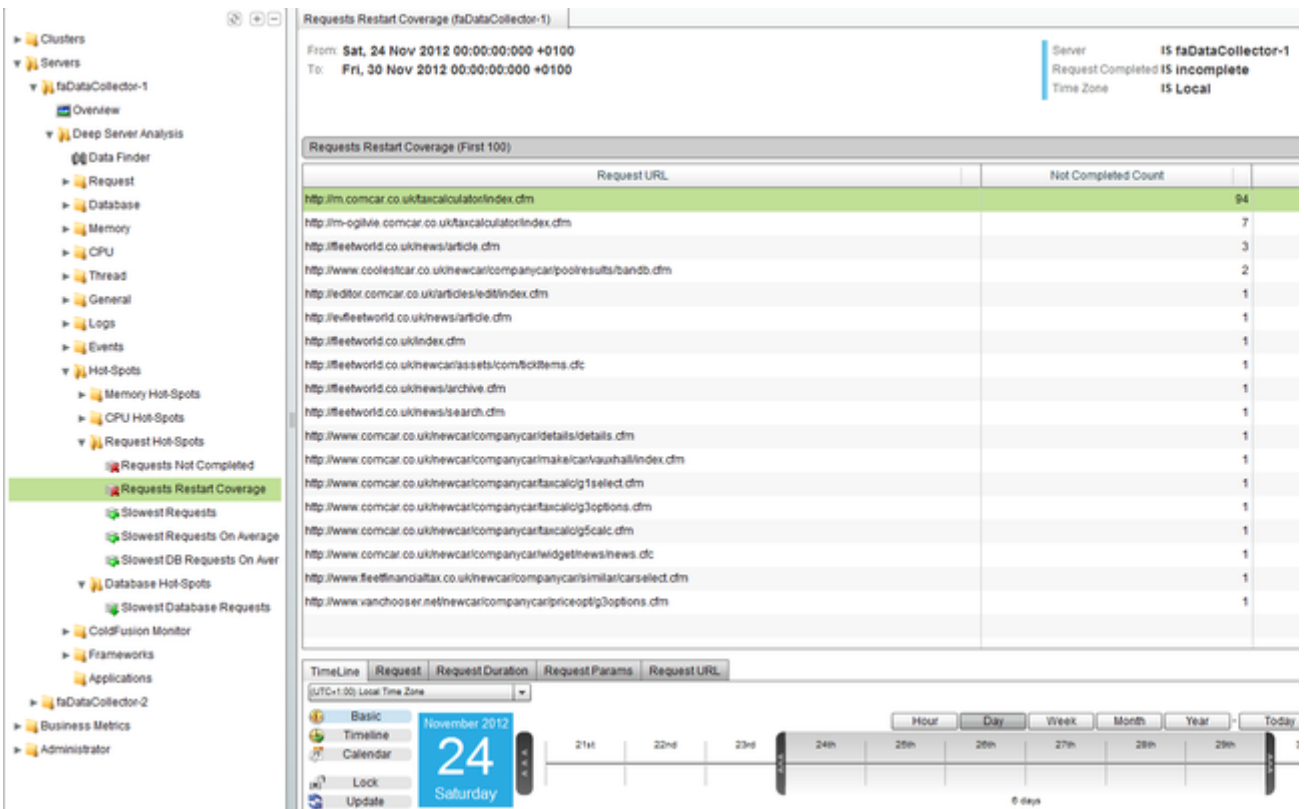


Figure 1: Requests Restart Coverage perspective

DATABASE HOT-SPOTS

Database Hot-Spots

Overview

Next Steps

Slowest Database Requests (Hot-Spots)

Slowest Database Requests (Hot-Spots)

Slowest Database Requests (Hot-Spots)

See Slowest Database Requests

Back to the top

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

The CF General Statistics perspective has several different views for lots of different ColdFusion Statistics, that you can switch between using the corner of the perspective within the specific period of time.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers. You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to see the remaining series.

Perspective View



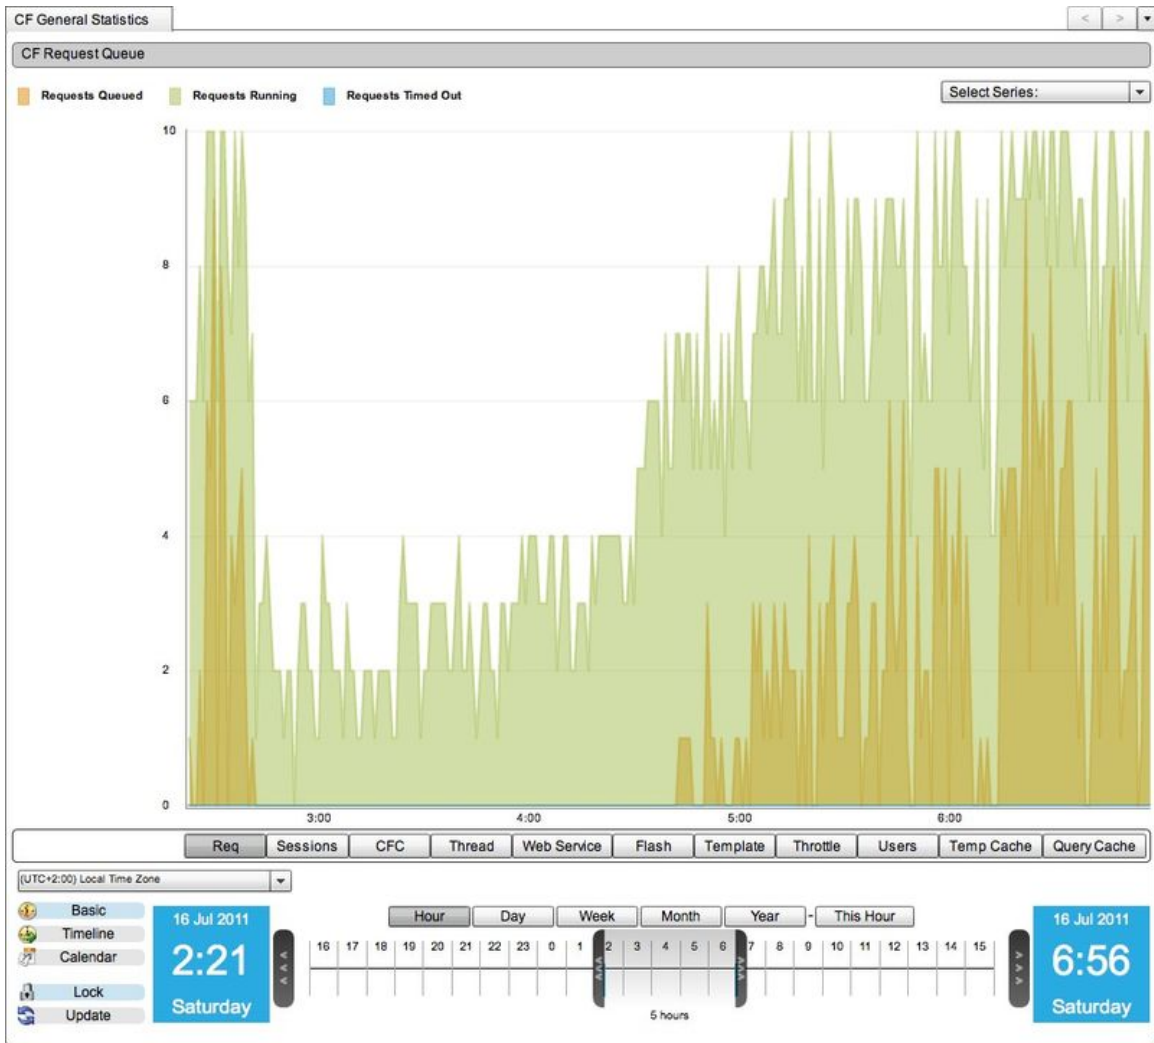


Figure 1: CF General Statistics Perspective

Cluster Perspective

All of the views that are available in the server version of this perspective are present in the cluster version, in the form of one Server Com. You can switch between the views using the drop down box in the uppermost right hand corner of the perspective.

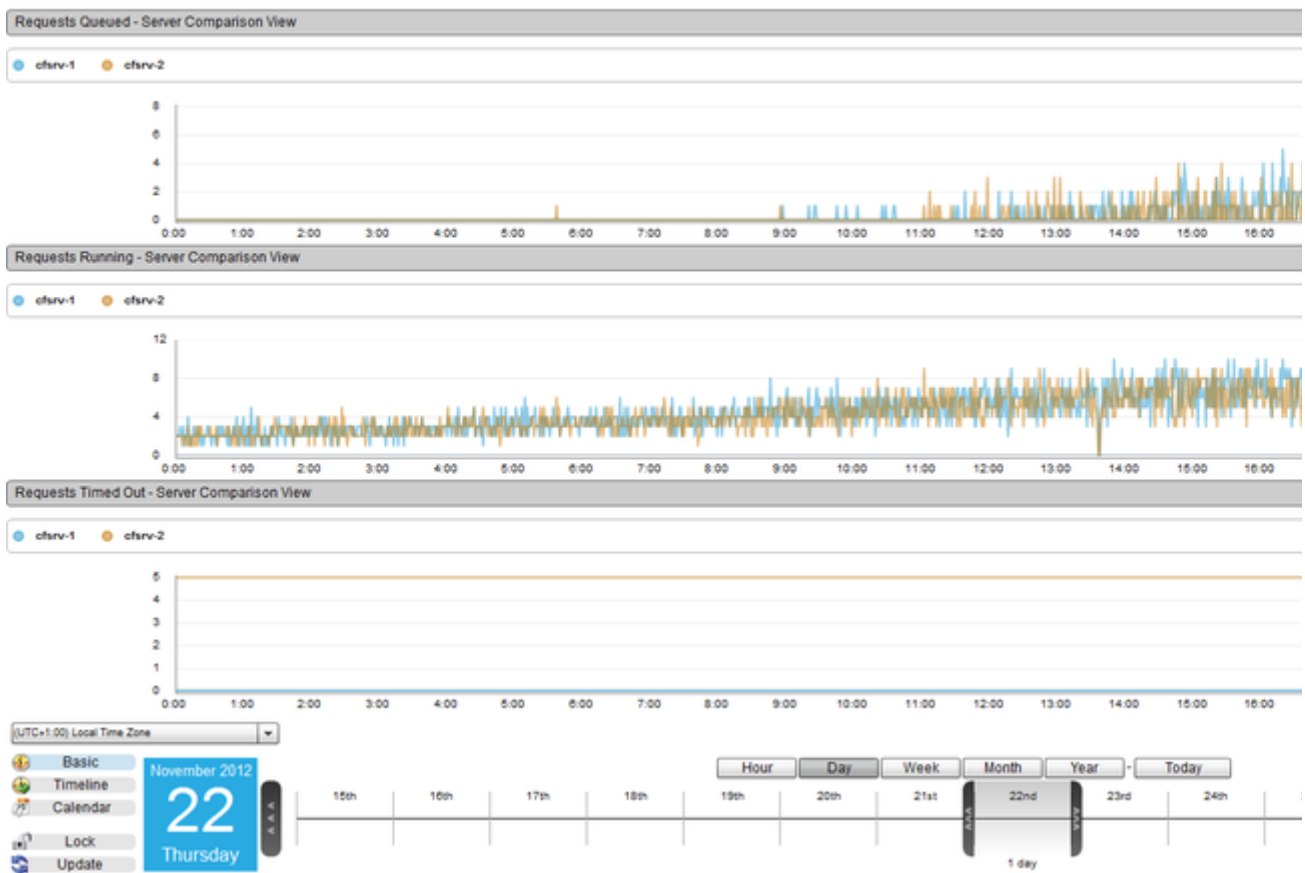


Figure 2: Cluster version of the CF General Statistics perspective

Back to the top

CF SCOPE SIZES

CF Scope Sizes

Description

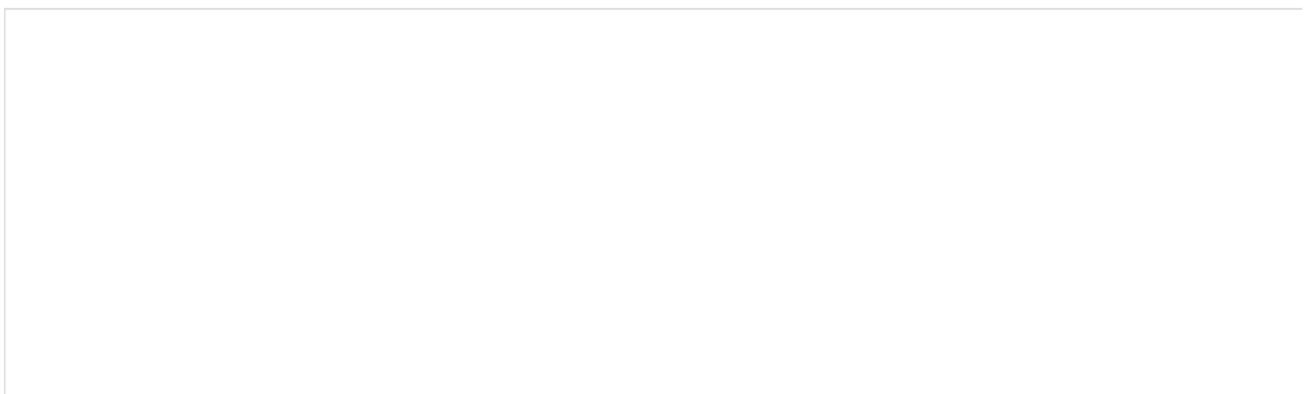
This perspective shows a chart of the CF Scope Sizes within the specific period of time.

Usage

Use the Date Navigation in order to change the period of time over which this perspective covers.

You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining :

Perspective View






Figure 1: CF Scope Sizes Perspective

Cluster Perspective

The cluster version of this perspective consists of four Server Comparison views: Memory Used(bytes), Application Scope Size(bytes), Se (bytes). Each graph has it's own "Select Series" drop down box in the right hand corner of the view.

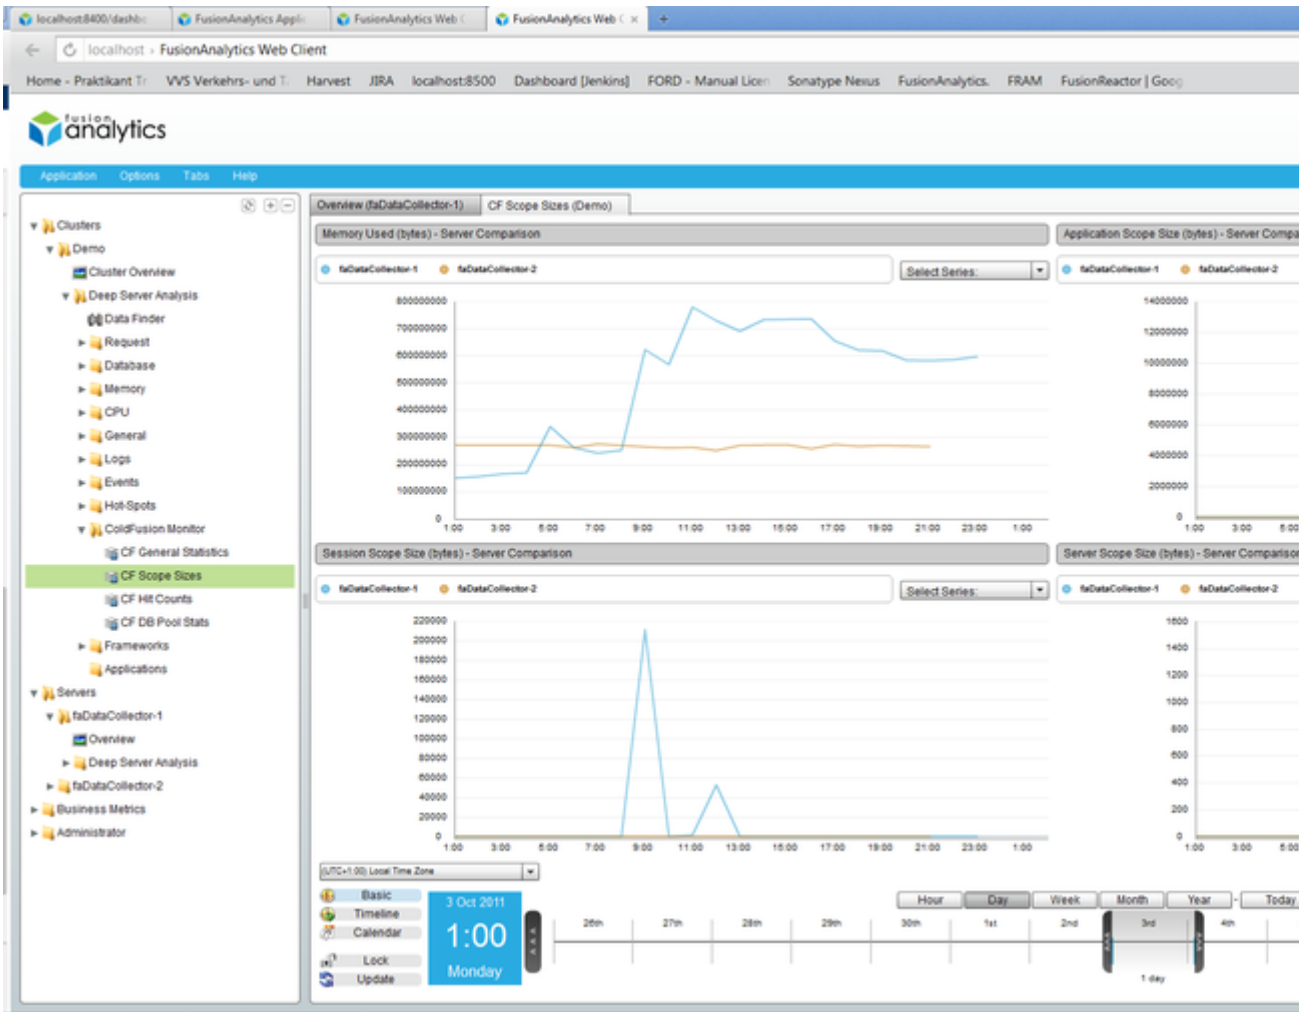


Figure 2: Cluster version of the CF Scope Sizes perspective

[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

Use the Date Navigation in order to change the period of time over which this perspective covers.

You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining :

Perspective View






Figure 1: CF Hit Counts Perspective

Cluster Perspective

The cluster version of this perspective consists of a Server Comparison and a Cluster View for Hit Counts and a Server Comparison view. You can switch between the views using the drop down box in the uppermost right hand corner of the perspective.



Figure 2: Cluster version of the 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

Use the Date Navigation in order to change the period of time over which this perspective covers.

You can temporarily remove a series from the graph by clicking the "Select Series" drop down box in the right hand corner of the view, and simply uncheck the series that you want to remove by clicking it. This will remove the series from the graph allowing you to the remaining :

Perspective View






Figure 1: CF DB Pool Stats Perspective

Cluster Perspective

The cluster version of this perspective a Server Comparison view of the DBPoolStats for each server in the cluster and one Cluster View v servers in the cluster.

You can switch between the views using the drop down box in the uppermost right hand corner of the perspective.




Figure 2: Cluster version of the CF DB Pool Stats perspective

[Back to the top](#)

Applications

Applications

Applications are an easy way to filter your request data. It works in the same way as the filtering options in each of the Request perspectives, but it will be applied to each of the perspectives in the Application straight away.

This saves you from having to configure and apply the filters that you most commonly use, every time you open a Request perspective.

Application filters will be applied to all the Request perspectives that are Grids and the Database perspectives that are grids, including the perspectives for Clusters. When applied, you can find the perspectives under the ToC item labeled "Applications".

Creating an Application

To create an application, you must go to the "Manage Applications" page from the ToC.

You will be taken to the following page. To add an Application, simply enter you desired name and click the "Add Application Filter" button.

The name cannot match an already existing name and cannot contain any spaces or the following characters: ! # \$ % ^ & * () + = ~ ` [] \ ' ; , { } | " < > ? .

Configuring an Application

Once you have added an Application, the following options page will appear:

Figure 1: Example Application Filter

Your new Application will now appear under the ToC menu item "Applications", with the name you have given it.

You can add as many Applications as you like, and each of the existing Applications will be displayed under "Current Application Filters", with there configuration options hidden.

| Icon | Name | Description |
|----------------------|--------|---|
| ? Unknown Attachment | Edit | Opens the configuration options for this application to allow you to make changes. |
| ? Unknown Attachment | Rename | Allows you to rename the application. Names cannot already exist, contain spaces or the characters ! # \$ % ^ & * () + = ~ ` [] \ ' ; , { } " < > ? . |
| ? Unknown Attachment | Copy | Creates a copy of this application with a new name of your choosing. |
| ? Unknown Attachment | Delete | Deletes this application. |

Request Params

You can filter the requests by parameters that are passed down through the query string, such as the event parameter or the fuseaction parameter.

To set this filter option, click the "Request Params" tab in the Application Configuration Options (this tab is open by default).

Enter the parameter you want to filter by (e.g. event) into the first text box. Choose a filtering action from the drop down box (e.g. IS, IS NOT) and finally enter the filter condition that you want the param to meet (in Fig1, page.user is used).

Once you have done this, simply click the button to add the filter. You can chain together as many filters as you want to further narrow down your result sets, and you can remove a filter at any time by clicking the small next to the filter option in the applications configuration options.

Figure 2: Slowest Requests Perspective with 'Request Param' Application Filters

Figure 2 shows the "Slowest Requests" perspective, after the Application Filter from Figure1 has been applied. You can see that the data in the table is limited to show only the requests with a Query String containing an "event" parameter that matches the condition "page.user".

Request URL

You can filter the data by the URLs of the requests. Each request has its URL broken down into segments. Each / in the URL denotes the start of a new segment.

For instance, the following URL: `http://int0046.bbn.intergral.com:18080/apps/modelgluesamples/lhp/index.cfm` can be broken down into the following segments;

| Segment | Value |
|---------|-------|
|---------|-------|

| | |
|--|---------------------------|
| URL Protocol (protocol://*...) | http |
| URL Hostname (*://hostname/*...) | int0046.bbn.intergral.com |
| URL Port (*://*:port/*...) | 18080 |
| Path Element 1 (*://*[1]/*...) | apps |
| Path Element 2 (*://*/[2]/*...) | modelgluesamples |
| Path Element 3 (*://*/*/[3]/*...) | lhp |
| Path Element 4 (*://*/*/*/[4]/*...) | index.cfm |

To set this filter option, click the "Request URL" tab in the Application Configuration Options.

Enter the segment you want to filter by (e.g. URL Port) from the first drop down box. Choose a filtering action from the second drop down box (e.g. IS, IS NOT) and finally enter the filter condition that you want the url segment to meet into the text box.

Once you have done this, simply click the button to add the filter. You can chain together as many filters as you want to further narrow down your result sets, and you can remove a filter at any time by clicking the small next to the filter option in the applications configuration options.

Figure 3: The Slowest Request perspective with 'Request URL' Application filters.

Figure 3 shows the "Slowest Requests" perspective, after a Request URL Application Filter has been applied. In this case, the Application uses two filter. It first filters by "Port is 18080", and second filters by "Path Element 3 is lhp". You can see that the data in the table is limited to show only the requests that match these two filters.

[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 quickly 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.](#)



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](#)

SLA Reports:

[Accessing the SLA Report Configuration Page](#)

? Unknown Attachment

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.](#)

? Unknown Attachment

Figure 2: Part of a Daily Status Report

SLA Reports

The SLA Reports are a designed to give an overview of your server performance at a glance using business-level statistics.

[Click here to learn more about the SLA Reports.](#)

Next Steps

TAP Reports

- [What is TAP Score - and how is it calculated?](#)
- [Running a TAP Report](#)
- [TAP Report Provider Arguments](#)
- [Understanding the TAP Report](#)
- [The TAP Report configuration page](#)

Daily Status Reports

- [Running a Daily Status Report](#)
- [Daily Status Report Provider Arguments](#)
- [Understanding the Daily Status Report](#)
- [The Daily Status Report configuration page](#)

SLA Reports

- [Accessing the SLA Report Configuration Page](#)
- [Adding a new SLA Report Configuration](#)
- [The SLA Report Configuration](#)
 - [Report Configuration Tab](#)
 - [Email Recipients Tab](#)
 - [Domain Filters Tab](#)
 - [Thresholds Tab](#)
- [Running the SLA Reports](#)

TAP Reports

TAP REPORTS

OVERVIEW

? Unknown macro: 'tooltip'

The 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.

Adding a new SLA Report Configuration

Running the SLA Reports

The SLA Report Configuration

- [Report Configuration Tab](#)
- [Email Recipients Tab](#)
- [Domain Filters Tab](#)
- [Thresholds Tab](#)

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](#) or [Setting up the SLA Reports](#) quick start guides.

This score gives you a standardized comparison of applications regardless of function or hardware capability.

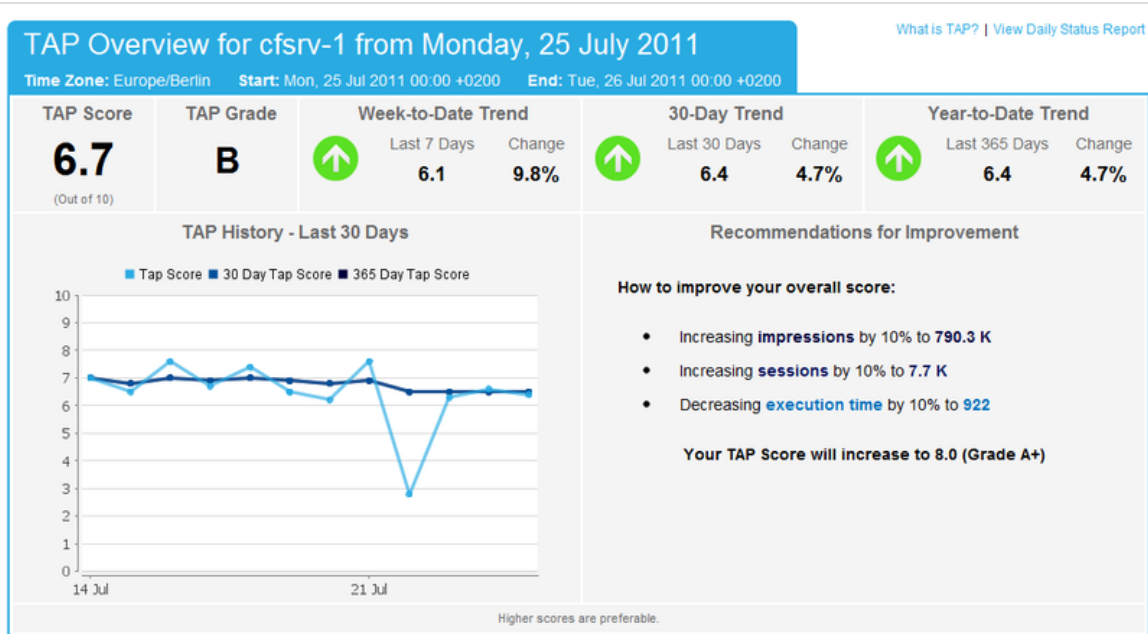


Figure 1: TAP Report Overview

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 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 Score - and how is it calculated?

Running a TAP Report

TAP Report Provider Arguments

Understanding the TAP Report

The TAP Report configuration page

What is TAP Score - and how is it calculated?

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.

On This Page:

- TAP Score
- Impressions
- Sessions

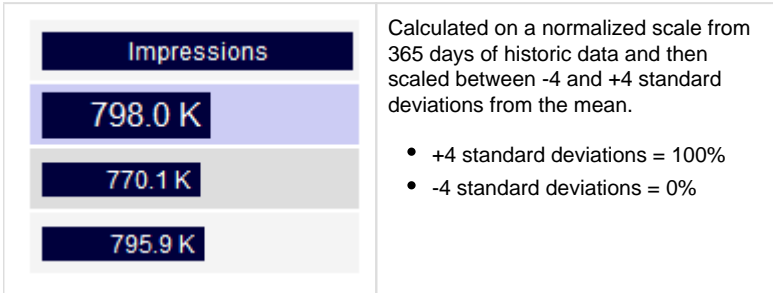
TAP Score

The TAP Score is measured from 0 (worst score) to 10 (best score).

There are 5 sections, awarding up to 2 points each.

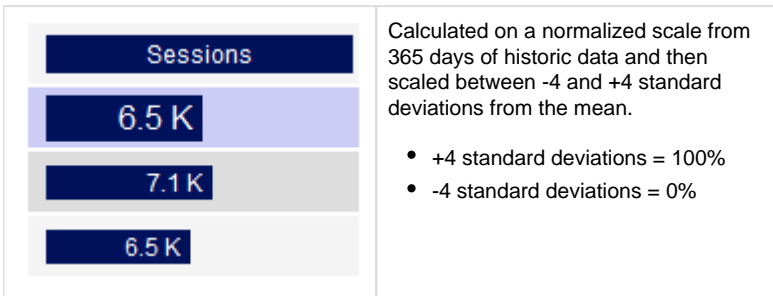
[Back to the top](#)

Impressions



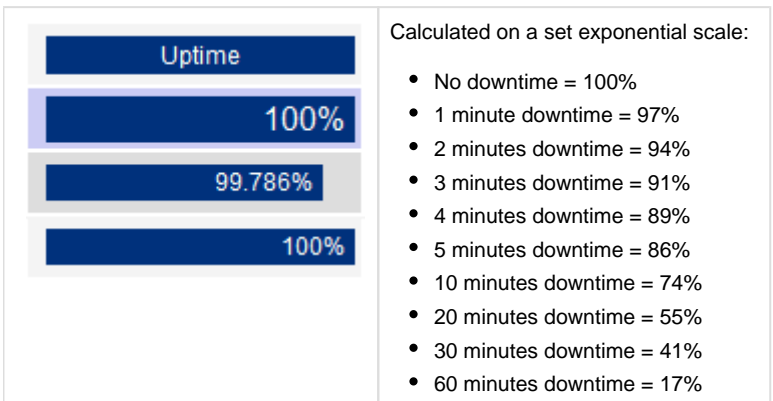
[Back to the top](#)

Sessions




[Back to the top](#)

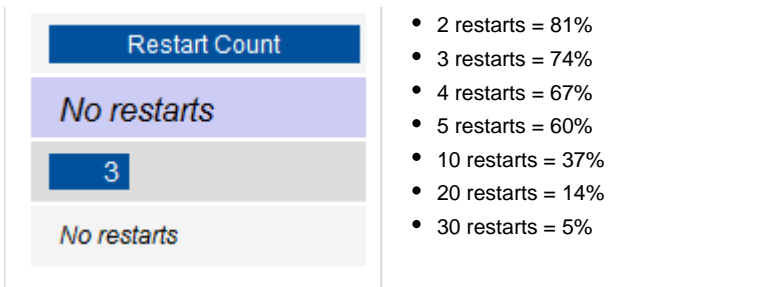
Uptime



[Back to the top](#)

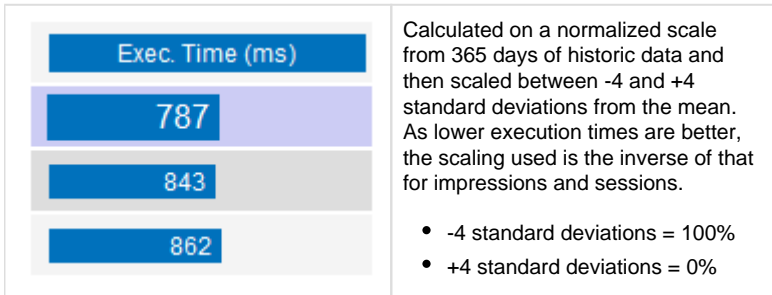
Restart Count





[Back to the top](#)

Execution Time



[Back to the top](#)

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 (

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).



Application Scheduled Tasks

| Task Name | Provider Name | Duration | Schedule | Last Run | Status |
|---|---------------------------|----------------------------|---------------------|-----------------|-----------|
| Daily Report | DailyStatus (coldfusion9) | Jan 1, 2011 - INDEFINITELY | Daily at 4:30:00 AM | 7/25/11 4:30 AM | Scheduled |
| Edit Run Now Deactivate | | | | | |
| T.A.P Report | TAPReport (coldfusion9) | Jan 1, 2011 - INDEFINITELY | Daily at 4:15:00 AM | 7/25/11 4:15 AM | Scheduled |
| Edit Run Now Deactivate | | | | | |

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 you receive y

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.

Task Name

T.A.P Report

Active

Duration

Start Date End Date (optional)

Frequency

One-Time at
 Recurring at
 Daily every Hours Minutes Seconds
Start Time End Time

Provider Name

TAPReport

Provider Arguments

Name Value -
+ Add New Argument

[Update Task](#)

[Cancel](#)

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**.

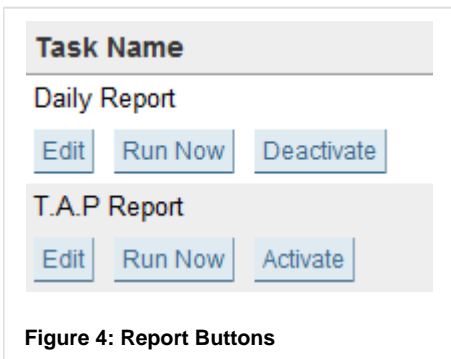


Figure 4: Report Buttons

[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.

| From | To | Report Name | Report Description | Server |
|-----------------------|-----------------------|---------------------|-------------------------------------|---------|
| 10.08.11 00:00:00.000 | 11.08.11 00:00:00.000 | Daily TAP Report | Daily TAP Report (Europe/Berlin) | cfsrv-1 |
| 10.08.11 00:00:00.000 | 11.08.11 00:00:00.000 | Daily TAP Report | Daily TAP Report (Europe/Berlin) | cfsrv-2 |
| 10.08.11 00:00:00.000 | 11.08.11 00:00:00.000 | Daily Status Report | Daily Status Report (Europe/Berlin) | cfsrv-1 |
| 10.08.11 00:00:00.000 | 11.08.11 00:00:00.000 | Daily Status Report | Daily Status Report (Europe/Berlin) | cfsrv-2 |
| 09.08.11 00:00:00.000 | 10.08.11 00:00:00.000 | Daily TAP Report | Daily TAP Report (Europe/Berlin) | cfsrv-1 |
| 09.08.11 00:00:00.000 | 10.08.11 00:00:00.000 | Daily TAP Report | Daily TAP Report (Europe/Berlin) | cfsrv-2 |
| 09.08.11 00:00:00.000 | 10.08.11 00:00:00.000 | Daily Status Report | Daily Status Report (Europe/Berlin) | cfsrv-1 |

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.

| Filter | Action | Argument |
|-------------|--------|---------------------|
| SERVER | IS | cfsrv-1 |
| Report Name | IS | Daily Status Report |

Figure 6: Report Filters

i You can also configure your TAP Reports so that you can view them as emails. See the **users** argument in TAP Report Provider Arg

Back to the top

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 | <p>The date that the report will run for.</p> <ul style="list-style-type: none"> 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. | <p>1312243200000 (Report Date: 1st-2nd August GMT)</p> <p>1310281200000 (Report Date: 9th-10th August GMT)</p> |
| https | False | <p>Defines if HTTP or HTTPS will be used.</p> <ul style="list-style-type: none"> 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 | None | <p>Collects any data that is available for previous days.</p> <ul style="list-style-type: none"> The value must be a UTC time in milliseconds. Make sure the time is the start of the day i.e. midnight. 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 produce one report for each day onwards from the specified time.  Warning: You should not use this every time you run a TAP Report: <ul style="list-style-type: none"> 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. | 1361404800000 |
| timezone | DataCollector's timezone | <p>Used to change the default timezone of the report.</p> <ul style="list-style-type: none"> 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 | <p>Used to email reports to the specified users.</p> <ul style="list-style-type: none"> 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

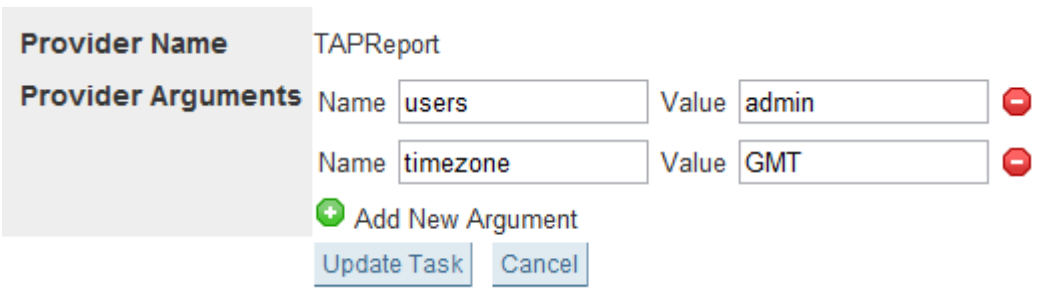
To add your Provider Arguments, you will need to go to the **Application Tasks Summary** in FADS and click the **Edit** button on the TAP Report task.

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.



Provider Name TAPReport

Provider Arguments

Name users Value admin 

Name timezone Value GMT 


 Add New Argument

Figure 1: Adding Provider Arguments

Next Steps

What is TAP Score - and how is it calculated?

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
- Last 7 Days
- Weekly Report
- Day of Week

Overview

The Overview header displays the following information (Figure 1):





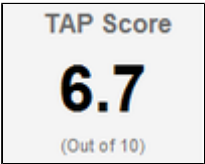
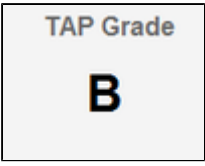
Figure 1: TAP Overview section header example

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: <ul style="list-style-type: none"> • 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 |

| | | |
|---------------------------------|---------------------------|---|
| | | <ul style="list-style-type: none"> • 3.0 to 3.9 = E • 2.9 and below = F |
| Week-to-Date Trend | | <p>The Last 7 Days score is an average of the TAP scores from today and the 6 previous days.</p> <p>The arrow indicates if today's score has improved, decreased or stayed the same compared to the week's average.</p> <p>Change gives the percentage change from the Last 7 Days average and today's score.</p> |
| 30-Day Trend | | <p>The Last 30 Days score is an average of the TAP scores from today and the 29 previous days.</p> <p>The arrow indicates if today's score has improved, decreased or stayed the same compared to the month's average.</p> <p>Change gives the percentage change from the Last 30 Days average and today's score.</p> |
| Year-to-Date Trend | | <p>The Last 365 Days score is an average of the TAP scores from today and the 364 previous days.</p> <p>The arrow indicates if today's score has improved, decreased or stayed the same compared to the year's average.</p> <p>Change gives the percentage change from the Last 365 Days average and today's score.</p> |
| TAP History | <p>(click to enlarge)</p> | <p>The TAP History graph shows a maximum of 30 days. There are 3 lines, showing:</p> <ul style="list-style-type: none"> • The TAP score for each day • The 30 day average TAP scores • The 365 day average TAP scores <p>⚠ Until you have more than 30 days of TAP data, the 30 and 365 day averages will be the same.</p> |
| Recommendations for Improvement | <p>(click to enlarge)</p> | <p>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).</p> <p>You can also see the score you will achieve by making all of the improvements.</p> |

[Back to the top](#)

Last 7 Days

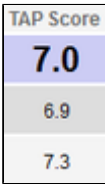

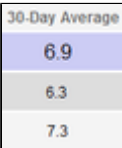
The top row of the Last 7 Days table is highlighted; this is the data that has been generated for the day the report has been run (i.e. the previous day).

Last 7 Days contains the following sections:

| Name | Example | Description |
|-----------|-------------------------|---|
| Date | Friday, 1 July 2011 | The top row, highlighted in purple, is the date that the report |
| g. | Thursday, 30 June 2011 | |
| Friday, 1 | Wednesday, 29 June 2011 | |


Example Last 7 Days

(click to enlarge)

| | | |
|-----------------------|---|---|
| <p>July 2011</p> | | <p>has been generated for.</p> <p>The next rows show the previous 7 days, going back to the same day of the previous week (e.g. Friday 24 July 2011).</p> |
| <p>TAP Score</p> |  | <p>The TAP Score that your server received for each day.</p> |
| <p>30-Day Trend</p> |  | <p>This compares the TAP Score from each day with the 30-day Average score from each day.</p> <p>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.</p> |
| <p>30-Day Average</p> |  | <p>This is the average of TAP Scores from the ? past 30</p> |

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 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.

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

| | | |
|--------------------------|--|--|
| | | <p>this score; the shorter the bar, the better the score.</p> |
| Exec. Time (Performance) | <p>A bar chart titled 'Exec. Time (ms)' showing three horizontal bars. The top bar is blue with the value '787'. The middle bar is grey with the value '843'. The bottom bar is blue with the value '862'.</p> | <p>Exec. Time is the average execution time of requests between the start and end date of the report.</p> <p>For Exec. Time, lower values are preferable.</p> <p>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.</p> |

[Back to the top](#)

Weekly Report

| Name | Example | Description |
|------|----------------------------|-------------|
| Mon | 27 June 2011 - 3 July 2011 | Each |



day to Sunday

20 June 2011 - 26 June 2011

13 June 2011 - 19 June 2011

row is one week; the dates are from Monday to Sunday.

The top row includes the date that the report has been generated for.

[\(click to enlarge\)](#)

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

t
we
ek
has
not
yet
fini
she
d,
the
top
row
will
sho
w
the
ave
rag
e
fro
m
Mo
nd
ay
to
the
cur
ren
t
day
.

30
-
D
ay
Tr
en
d



This
co
mp
are
s
the
TA
P
Sc
ore
fro
m
eac
h
day
wit
h
the
30-
day
Av
era
ge
sco
re
fro
m
eac
h
day
.

A
gre
en
arr
ow
indi
cat
es
you
r
ser
ver
has
imp
rov
ed
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 the 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



?(click to enlarge)

This shows the 7-Day


| | |
|--------------------------------------|--|
| T A P S c o r e | y a v e r a g e a s a b a r g r a p h, a l l o w i n g y o u t o e a s i l y c o m p a r e w h i c h w e e k s h a d b e t t e r T A P s c o r e s o n a v e r a g e. |
|--------------------------------------|--|

[Back to the top](#)

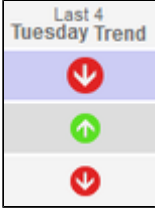
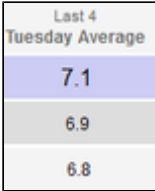
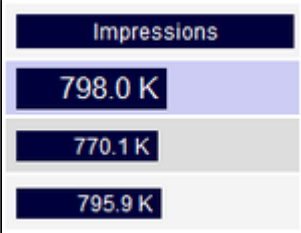
Day of Week

| N a m e | Example | Descri ption |
|---|---|---|
| D a t e . T u e s d a y, 2 8 J u n e 2 0 1 1 | <div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;">Tuesday, 28 June 2011</div> <div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;">Tuesday, 21 June 2011</div> <div style="border: 1px solid black; padding: 5px;">Tuesday, 14 June 2011</div> | <p>The top row, highlighted in purple, is the date that the report has been generated for.</p> <p>The next rows show the same day of the week, going back over 5 weeks.</p> |
| T A P S c o r e | <div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;">TAP Score</div> <div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;">7.0</div> <div style="border: 1px solid black; padding: 5px;">6.9</div> | The TAP Score that your server |

Example Day of Week



(click to enlarge)


| | | |
|---|---|--|
| | <div data-bbox="316 145 419 174" style="border: 1px solid black; padding: 2px; text-align: center;">7.3</div> | received for each day. |
| <p>Last 4 Day Trend e.g. Last 4 Tuesday Trend</p> |  | <p>This compares the TAP Score from each day with the Last 4 Day Trend score from each day.</p> <p>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.</p> |
| <p>Last 4 Day Average e.g. Last 4 Tuesday Average</p> |  | <p>This is an average of the scores from the last 4 rows (the same day of each week).</p> |
| <p>Impressions (Traffic)</p> |  | <p>Impressions are the total number of hits received during the start and end date of the</p> |

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 receive a score between 0.0 - 2.0 (see ? What is TAP? for more information).

tion). 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

| | | |
|----------------------|--|--|
| or ma nc e) | <div style="border: 1px solid black; background-color: #0070c0; color: white; padding: 2px; text-align: center; width: fit-content; margin-bottom: 5px;">862</div> | <p>of requests between the start and end date of the report.</p> <p>For Exec. Time, lower values are preferable.</p> <p>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.</p> |
|----------------------|--|--|

[Back to the top](#)

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.

The TAP Report configuration page

THE TAP REPORT CONFIGURATION

The TAP Reports come with a new report configuration feature will allows you to add email recipients to the report. This feature can be found within the FusionAnalytics Client.

There are two ways to access the client, either through a web browser (see The FusionAnalytics Web Client) or using our Air Dashboard (see The FusionAnalytics AIR Client).

On this page:

Adding a recipient from FADS users

Adding a custom email address

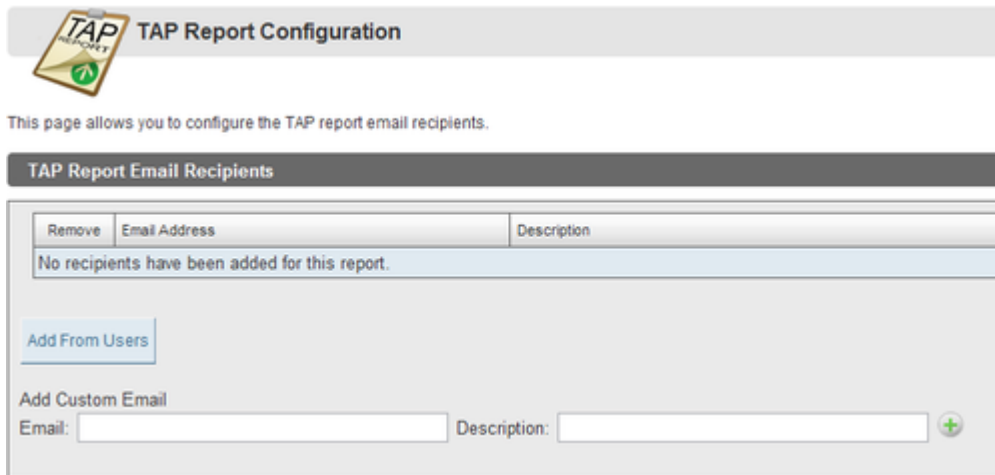
Where to go from the client

Session Configuration

Once you've accessed the client, you'll find a new ToC item under Administrator > TAP Report Configuration.

Clicking this will display the new configuration screen which may contain existing report configurations if some have previously been created.

Adding email recipients.



The screenshot shows the 'TAP Report Configuration' page. At the top left is a 'TAP REPORT' logo. Below it is a header 'TAP Report Configuration'. A sub-header reads 'This page allows you to configure the TAP report email recipients.' Below this is a section titled 'TAP Report Email Recipients'. It contains a table with columns 'Remove', 'Email Address', and 'Description'. The table is empty, with a message 'No recipients have been added for this report.' Below the table is an 'Add From Users' button. At the bottom, there is an 'Add Custom Email' section with 'Email:' and 'Description:' input fields and a green plus button.

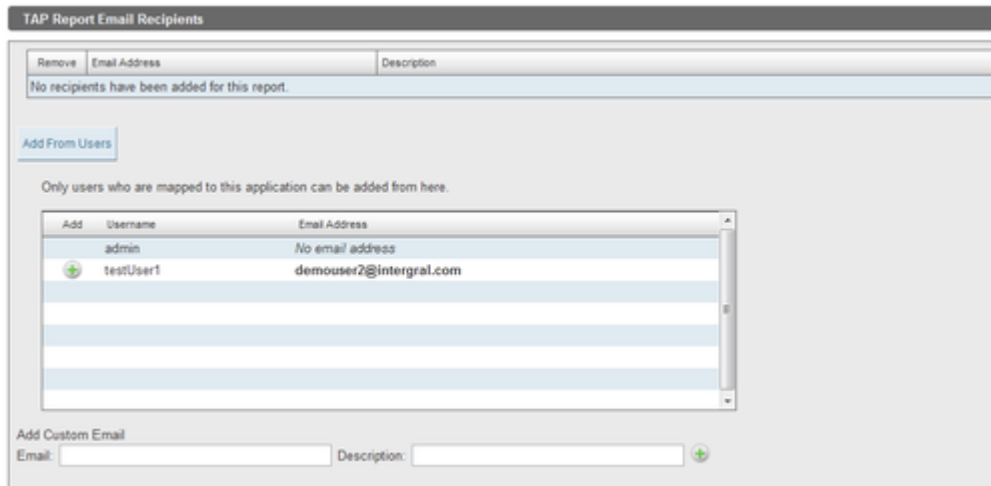
This page allows you to configure who will receive the report via email.

There are two methods of adding recipients:

- Adding a recipient from FADS users. Only users who are mapped are available to add with this method.
- Adding a custom email address.

Adding a recipient from FADS users.

To add from FADS users, simply click the Add From Users button to view a list of viable FADS users and click the button to add them to the recipients list.



This screenshot shows the 'TAP Report Email Recipients' section with the 'Add From Users' button clicked. A dialog box is open, displaying a list of users. The dialog has an 'Add' button and a table with columns 'Username' and 'Email Address'. The table contains two entries: 'admin' with 'No email address' and 'testUser1' with 'demouser2@integral.com'. Below the dialog is the 'Add Custom Email' section with 'Email:' and 'Description:' input fields and a green plus button.

Adding a custom email address

To add a custom email address, simply enter the email address and optional description into the corresponding fields and click the button.

Session Configuration

Click on the Report Configuration tab. You will see a page that looks like this:

This page will allow you to tell the TAP report which of your servers use sessions.

For those that use sessions, the box MUST be checked for the TAP report to calculate the data.

If a server does not use sessions, then the check box should be left clear. In these cases, the TAP report will not take session data into account when calculating the final score.

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 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.

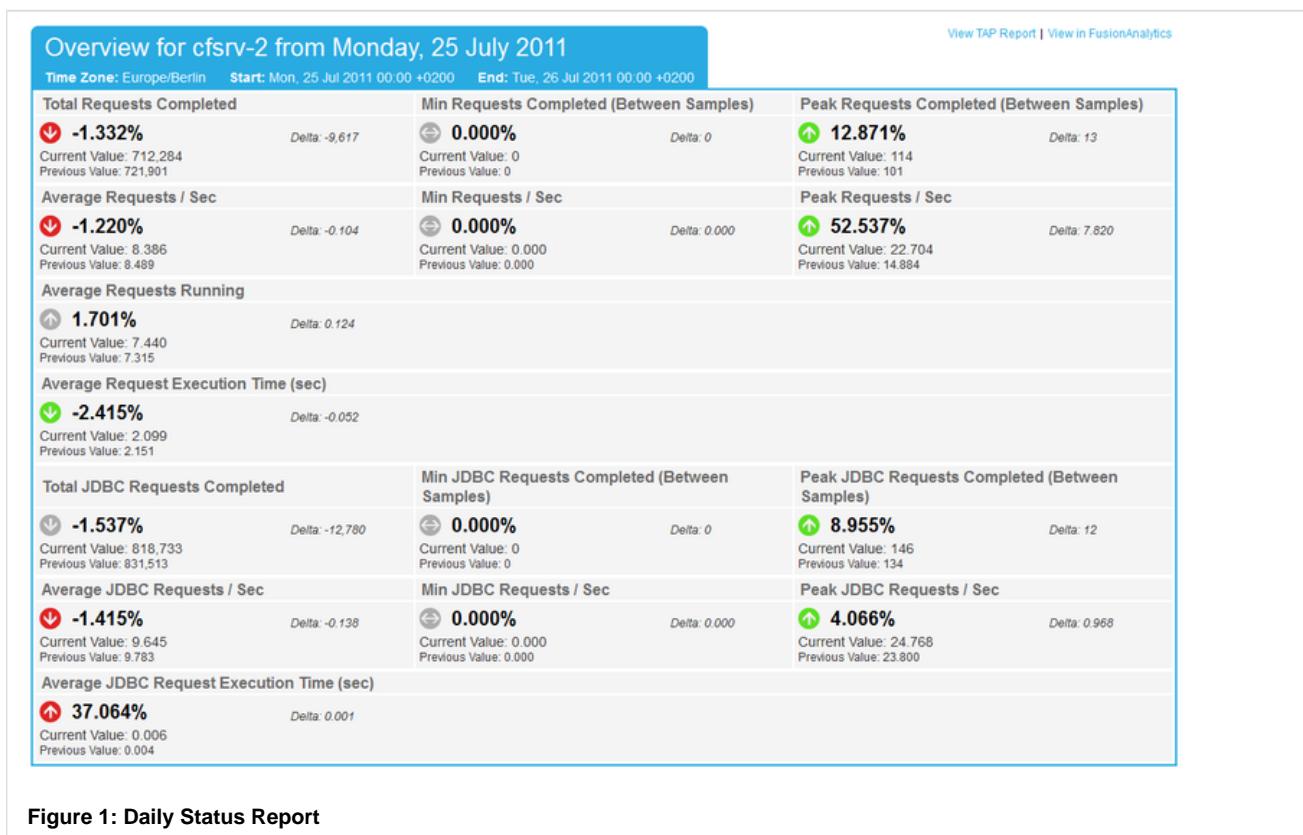


Figure 1: Daily Status Report

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 Report](#) 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

The Daily Status Report configuration page

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 Sta

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).

| Application Scheduled Tasks | | | | | |
|-----------------------------|---------------------------|----------------------------|---------------------|-----------------|-----------|
| Task Name | Provider Name | Duration | Schedule | Last Run | Status |
| Daily Report | DailyStatus (coldfusion9) | Jan 1, 2011 - INDEFINITELY | Daily at 4:30:00 AM | 7/25/11 4:30 AM | Scheduled |
| Edit | Run Now | Deactivate | | | |
| T.A.P Report | TAPReport (coldfusion9) | Jan 1, 2011 - INDEFINITELY | Daily at 4:15:00 AM | 7/25/11 4:15 AM | Scheduled |
| Edit | Run Now | Deactivate | | | |

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 re

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.

| | |
|---------------------------|---|
| Task Name | Daily Report |
| Active | <input checked="" type="checkbox"/> |
| Duration | Start Date <input type="text" value="01/01/2011"/> End Date (optional) <input type="text"/> |
| Frequency | <input type="radio"/> One-Time at <input type="text"/> <input checked="" type="radio"/> Recurring <input type="text" value="Daily"/> at <input type="text" value="04:30 PM"/> <input type="radio"/> Daily every Hours <input type="text"/> Minutes <input type="text"/> Seconds <input type="text"/> Start Time <input type="text"/> End Time <input type="text"/> |
| Provider Name | DailyStatus |
| Provider Arguments | Name <input type="text" value="users"/> Value <input type="text" value="admin"/> <input type="button" value="−"/> |
| | <input type="button" value="+"/> Add New Argument |
| | <input type="button" value="Update Task"/> <input type="button" value="Cancel"/> |

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.

| |
|--|
| Daily Report |
| <input type="button" value="Edit"/> <input type="button" value="Run Now"/> <input type="button" value="Deactivate"/> |

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**.

| |
|--|
| Task Name |
| Daily Report |
| <input type="button" value="Edit"/> <input type="button" value="Run Now"/> <input type="button" value="Deactivate"/> |
| T.A.P Report |
| <input type="button" value="Edit"/> <input type="button" value="Run Now"/> <input type="button" value="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

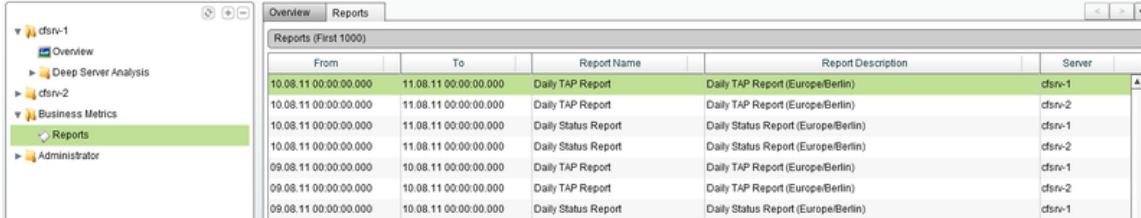
See Daily Status Report Provider Arguments for more information.

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.

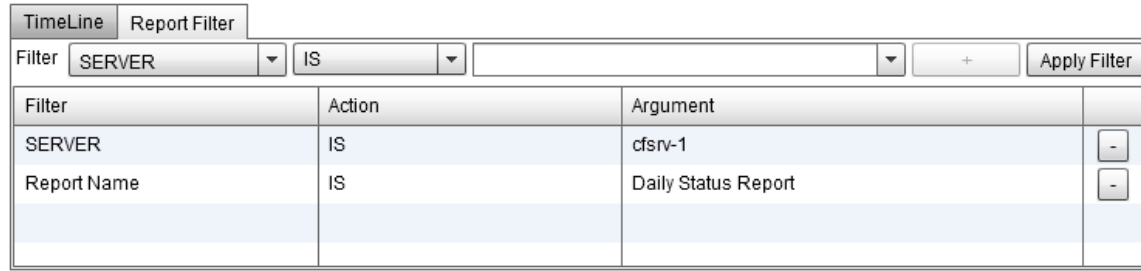


| From | To | Report Name | Report Description | Server |
|-----------------------|-----------------------|---------------------|-------------------------------------|---------|
| 10.08.11 00:00:00.000 | 11.08.11 00:00:00.000 | Daily TAP Report | Daily TAP Report (Europe/Berlin) | dfsrv-1 |
| 10.08.11 00:00:00.000 | 11.08.11 00:00:00.000 | Daily TAP Report | Daily TAP Report (Europe/Berlin) | dfsrv-2 |
| 10.08.11 00:00:00.000 | 11.08.11 00:00:00.000 | Daily Status Report | Daily Status Report (Europe/Berlin) | dfsrv-1 |
| 10.08.11 00:00:00.000 | 11.08.11 00:00:00.000 | Daily Status Report | Daily Status Report (Europe/Berlin) | dfsrv-2 |
| 09.08.11 00:00:00.000 | 10.08.11 00:00:00.000 | Daily TAP Report | Daily TAP Report (Europe/Berlin) | dfsrv-1 |
| 09.08.11 00:00:00.000 | 10.08.11 00:00:00.000 | Daily TAP Report | Daily TAP Report (Europe/Berlin) | dfsrv-2 |
| 09.08.11 00:00:00.000 | 10.08.11 00:00:00.000 | Daily Status Report | Daily Status Report (Europe/Berlin) | dfsrv-1 |

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.



| Filter | Action | Argument |
|-------------|--------|---------------------|
| SERVER | IS | dfsrv-1 |
| Report Name | IS | Daily Status Report |

Figure 6: Report Filters

i You can also configure your Daily Status Reports so that you can view them as emails. See the **users** argument in ?Daily Status Rep

Back to the top

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. <ul style="list-style-type: none">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. <ul style="list-style-type: none">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. <ul style="list-style-type: none">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. <ul style="list-style-type: none">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

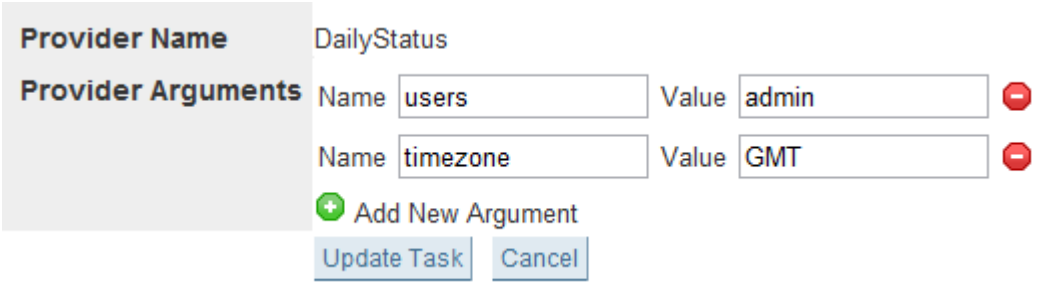
To add your Provider Arguments, you will need to go to the **Application Tasks Summary** in FADS and click the **Edit** button on the Daily Status Report task.

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.



Provider Name DailyStatus

Provider Arguments

Name users Value admin 

Name timezone Value GMT 


 Add New Argument

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 |
|--|--|
|   | <p>A green arrow indicates a good change.</p> <p>In some cases a lower value is preferable, e.g. Average Request Execution Time.</p> |
|   | <p>A red arrow indicates a bad change.</p> <p>In some cases a higher value is not preferable, e.g. Average Memory Used.</p> |
|   | <p>A grey arrow indicates that the change is neither good nor bad.</p> <p>For example, the Average Requests Running uses this type of arrow.</p> |
|  | <p>This indicates no change between the current and previous values</p> |

On this page:

- [Overview](#)
- [Performance Charts](#)
- [System Load Statistics](#)
- [Session Statistics](#)
- [Slowest Requests](#)
- [Slowest Requests on Average](#)
- [Application Data Overview](#)
- [Application JDBC Overview](#)
- [Applications](#)

Overview


The Overview header displays the following information (Figure 1):














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 | <div style="border: 1px solid black; padding: 5px;"> <p>Total Requests Completed</p> <p> -1.332% Delta: -9,617</p> <p>Current Value: 712,284 Previous Value: 721,901</p> </div> | The total number of requests that have completed during the 24 hours. |
| Min Requests Completed (Between Samples) | <div style="border: 1px solid black; padding: 5px;"> <p>Min Requests Completed (Between Samples)</p> <p> 0.000% Delta: 0</p> <p>Current Value: 0 Previous Value: 0</p> </div> | The smallest number of requests completed at sampled points throughout the 24 hours. |
| Peak Requests Completed (Between Samples) | <div style="border: 1px solid black; padding: 5px;"> <p>Peak Requests Completed (Between Samples)</p> <p> 12.871% Delta: 13</p> <p>Current Value: 114 Previous Value: 101</p> </div> | The largest number of requests completed at sampled points throughout the 24 hours. |
| Average Requests / Sec | <div style="border: 1px solid black; padding: 5px;"> <p>Average Requests / Sec</p> <p> -1.220% Delta: -0.104</p> <p>Current Value: 8.386 Previous Value: 8.489</p> </div> | The average number of requests that are completed per second. |
| Min Requests / Sec | <div style="border: 1px solid black; padding: 5px;"> <p>Min Requests / Sec</p> <p> 0.000% Delta: 0.000</p> <p>Current Value: 0.000 Previous Value: 0.000</p> </div> | The smallest number of requests that are completed per second. |
| Peak Requests / Sec | <div style="border: 1px solid black; padding: 5px;"> <p>Peak Requests / Sec</p> <p> 52.537% Delta: 7.820</p> <p>Current Value: 22.704 Previous Value: 14.884</p> </div> | The largest number of requests that are completed per second. |
| Average Requests Running | <div style="border: 1px solid black; padding: 5px;"> <p>Average Requests Running</p> <p> 1.701% Delta: 0.124</p> <p>Current Value: 7.440 Previous Value: 7.315</p> </div> | The average number of requests running at any time. |
| Average Request Execution Time (sec) | <div style="border: 1px solid black; padding: 5px;"> <p>Average Request Execution Time (sec)</p> <p> -2.415% Delta: -0.052</p> <p>Current Value: 2.099 Previous Value: 2.151</p> </div> | The average execution time, in seconds, of all requests in the 24 hours. |
| Total JDBC Requests Completed | <div style="border: 1px solid black; padding: 5px;"> <p>Total JDBC Requests Completed</p> <p> -1.537% Delta: -12,780</p> <p>Current Value: 818,733 Previous Value: 831,513</p> </div> | The total number of JDBC requests that have completed during the 24 hours. |
| Min JDBC Requests Completed (Between Samples) | <div style="border: 1px solid black; padding: 5px;"> <p>Min JDBC Requests Completed (Between Samples)</p> <p> 0.000% Delta: 0</p> <p>Current Value: 0 Previous Value: 0</p> </div> | The smallest number of JDBC requests completed at sampled points throughout the 24 hours. |
| Peak JDBC Requests Completed (Between Samples) | <div style="border: 1px solid black; padding: 5px;"> <p>Peak JDBC Requests Completed (Between Samples)</p> <p> 8.955% Delta: 12</p> <p>Current Value: 146 Previous Value: 134</p> </div> | The largest number of JDBC requests completed at sampled points throughout the 24 hours. |
| Average JDBC Requests / Sec | <div style="border: 1px solid black; padding: 5px;"> <p>Average JDBC Requests / Sec</p> </div> | |


| | | |
|---|--|---|
| | <div style="border: 1px solid black; padding: 5px;"> ▼ -1.415% Delta: -0.138 Current Value: 9.645 Previous Value: 9.783 </div> | The average number of JDBC requests that are completed per second. |
| Min JDBC Requests / Sec | <div style="border: 1px solid black; padding: 5px;"> ↔ 0.000% Delta: 0.000 Current Value: 0.000 Previous Value: 0.000 </div> | The smallest number of JDBC requests that are completed per second. |
| Peak JDBC Requests / Sec | <div style="border: 1px solid black; padding: 5px;"> ▲ 4.066% Delta: 0.968 Current Value: 24.768 Previous Value: 23.800 </div> | The largest number of requests that are completed per second. |
| Average JDBC Request Execution Time (sec) | <div style="border: 1px solid black; padding: 5px;"> Average JDBC Request Execution Time (sec) ▲ 37.064% Delta: 0.001 Current Value: 0.006 Previous Value: 0.004 </div> | The average JDBC execution time, in seconds, of all requests in the 24 hours. |

[Back to the top](#)

Performance Charts

The Performance Charts section contains the following:


| Graph | Example | Description |
|------------------------|---|--|
| Request Execution Time |  | A graph showing the average execution time, in milliseconds, for |



r
t
h
e
r
e
q
u
e
s
t
s
t
h
r
o
u
g
h
o
u
t
t
h
e
d
a
y.


T
h
e
r
e
i
s
o
n
e
p
o
i
n
t
o
n
t
h
e
g

R
e
q
u
e
s
t
E
x
e
c
u
t
i
o
n
T
i
m
e
:
?
P
e
a
k



A
g
r
a
p
h
s
h
o
w
i
n
g
t
h
e
m
a
x
i
m
u
m
a
v
e
r
a
g
e
e
x
e
c
u
t
i
o
n


| | |
|---------------|--|
| | |
| Daily Traffic | time, in milliseconds, for the requests throughout the day. There is one point on the graph for each hour. A graph showing the average |



g
e
n
u
m
b
e
r
o
f
r
e
q
u
e
s
t
s
p
e
r
h
o
u
r
f
o
r
t
h
e
r
e
q
u
e
s
t
s
t
h
r
o
u
g
h
o
u
t
t
h
e
d
a
y.

T
h
e
r
e
i
s
o
n
e
p
o
i
n
t
o
n
t
h
e
g
r
a
p
h
f
o
r
e
a
c
h
h
o
u
r
.

D
a
t
a
b
a
s
e
R
e
q
u
e
s
t



A
g
r
a
p
h
s
h
o
w
i
n
g
t
h


Execution Time: Average

average JDB execution time, in

milliseconds, for the requested throughput today.

There is one point on the graph for each hour


Data base Request Execution Time: Peak



ur . A graph showing the maximum average JDB C 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 day.

a
y.

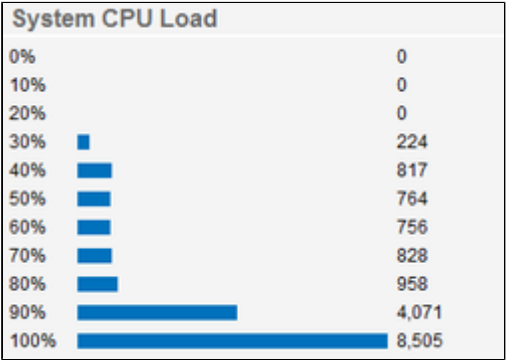
T
h
e
r
e
i
s
o
n
e
p
o
i
n
t
o
n
t
h
e
g
r
a
p
h
f
o
r
e
a
c
h
h
o
u
r
.

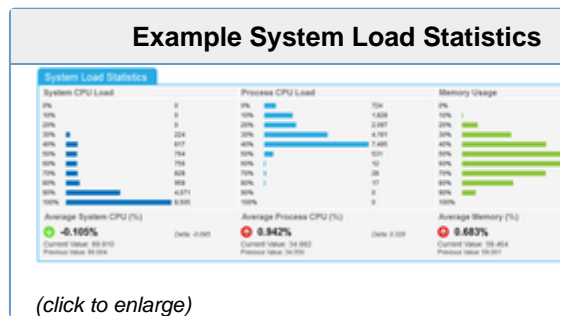
i Clicking on the Performance Charts header will take you to the Overview perspective in the Web Client.

[Back to the top](#)

System Load Statistics

The System Load Statistics section contains the following:


| Name | Example | Description | | | | | | | | | | | | | | | | | | | | | | | | |
|---------------------|--|---------------------|-------|----|---|-----|---|-----|---|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-------|------|-------|---|
| System CPU Load |  <table border="1" data-bbox="209 1458 716 1816"> <thead> <tr> <th>System CPU Load (%)</th> <th>Value</th> </tr> </thead> <tbody> <tr><td>0%</td><td>0</td></tr> <tr><td>10%</td><td>0</td></tr> <tr><td>20%</td><td>0</td></tr> <tr><td>30%</td><td>224</td></tr> <tr><td>40%</td><td>817</td></tr> <tr><td>50%</td><td>764</td></tr> <tr><td>60%</td><td>756</td></tr> <tr><td>70%</td><td>828</td></tr> <tr><td>80%</td><td>958</td></tr> <tr><td>90%</td><td>4,071</td></tr> <tr><td>100%</td><td>8,505</td></tr> </tbody> </table> | System CPU Load (%) | Value | 0% | 0 | 10% | 0 | 20% | 0 | 30% | 224 | 40% | 817 | 50% | 764 | 60% | 756 | 70% | 828 | 80% | 958 | 90% | 4,071 | 100% | 8,505 | A bar graph showing the number of samples |
| System CPU Load (%) | Value | | | | | | | | | | | | | | | | | | | | | | | | | |
| 0% | 0 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 10% | 0 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 20% | 0 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 30% | 224 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 40% | 817 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 50% | 764 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 60% | 756 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 70% | 828 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 80% | 958 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 90% | 4,071 | | | | | | | | | | | | | | | | | | | | | | | | | |
| 100% | 8,505 | | | | | | | | | | | | | | | | | | | | | | | | | |



[\(click to enlarge\)](#)

taken at each percentage load for the system CPU.


Process CPU Load



A bargraph showing the number of samples taken at each percentage load for the pro


cess CPU.

Memory Usage



Abargraphs showing the number of samples taken at each percentage for memory.


Average System CPU (%)



The average percentage for the


a
m
o
u
n
t
o
f
S
y
s
t
e
m
C
P
U
u
s
e
d.

A
v
e
r
a
g
e
P
r
o
c
e
s
s
C
P
U
(
%
)

Average Process CPU (%)
 **0.942%** *Delta: 0.326*
Current Value: 34.882
Previous Value: 34.556

T
h
e
a
v
e
r
a
g
e
p
e
r
c
e
n
t
a
g
e
f
o
r
t
h
e
a
m
o
u
n
t
o
f
P
r
o
c
e
s
s
C
P
U
u
s
e
d.

A
v
e
r
a
g
e
M
e
m
o
r
y
(
%
)

Average Memory (%)
 **0.683%** *Delta: 0.403*
Current Value: 59.464
Previous Value: 59.061

T
h
e
a
v
e
r
a
g
e
p
e
r
c
e
n
t
a
g
e
f
o
r
t
h
e
a
m
o
u
n
t
o
f

i 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](#)

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. |


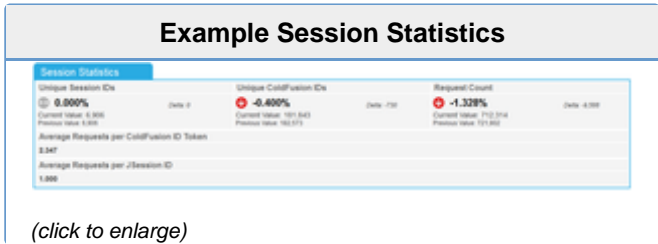
i Clicking on the Session Statistics header will take you to the Table of Requests perspective in the Web Client.

[Back to the top](#)


Slowest Requests

Slowest Requests contains the following sections:

| Name | Example | Description |
|------|---------|-------------|
| | | |




| | | p t i o n | | | | |
|--|---|----------------------------------|--|--|---|--|
| Start ed | <table border="1"> <tr><td>Started</td></tr> <tr><td>15:52:01:640</td></tr> <tr><td>01:48:55:609</td></tr> <tr><td>18:54:14:406</td></tr> </table> | Started | 15:52:01:640 | 01:48:55:609 | 18:54:14:406 | The time when the re q u e s t s t a r t e d. |
| Started | | | | | | |
| 15:52:01:640 | | | | | | |
| 01:48:55:609 | | | | | | |
| 18:54:14:406 | | | | | | |
| R e q u e s t U R L | <table border="1"> <tr><td>Request URL</td></tr> <tr><td>http://int0046.bbn.intergral.com:18080/apps/fads/test.cpuspike.cfm</td></tr> <tr><td>http://int0046.bbn.intergral.com:18080/apps/fads/test.cpuspike.cfm</td></tr> <tr><td>http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm</td></tr> </table> | 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 re q u e s t U R L. |
| 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 | | | | | | |
| Ex e c u t i o n T i m e (s e c) | <table border="1"> <tr><td>Execution Time (sec)</td></tr> <tr><td>182.344</td></tr> <tr><td>182.172</td></tr> <tr><td>182.016</td></tr> </table> | Execution Time (sec) | 182.344 | 182.172 | 182.016 | The le n g t h o f t i m e, in s e c o n d s, th at the re q u e s t to ok to ru n. |
| Execution Time (sec) | | | | | | |
| 182.344 | | | | | | |
| 182.172 | | | | | | |
| 182.016 | | | | | | |

 Clicking on the Slowest Requests header will take you to the Slowest Requests perspective in the Web Client.

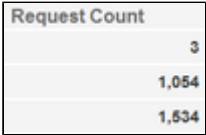

You can also click **Expand / Collapse** to see more slow requests.

[Back to the top](#)



Slowest Requests on Average

Slowest Requests on Average contains the following section:

| Name | Example | Description | | | | |
|---|--|---------------|--|---|---|---|
| Request Count |  <table border="1" data-bbox="389 479 595 613"> <thead> <tr> <th>Request Count</th> </tr> </thead> <tbody> <tr> <td>3</td> </tr> <tr> <td>1,054</td> </tr> <tr> <td>1,534</td> </tr> </tbody> </table> | Request Count | 3 | 1,054 | 1,534 | The total number of requests that have occurred for this URL. |
| Request Count | | | | | | |
| 3 | | | | | | |
| 1,054 | | | | | | |
| 1,534 | | | | | | |
| Request URL |  <table border="1" data-bbox="209 1323 775 1451"> <thead> <tr> <th>Request URL</th> </tr> </thead> <tbody> <tr> <td>http://int0046.bbn.intergral.com:18080/apps/fads/test.cpuspike.cfm</td> </tr> <tr> <td>http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm</td> </tr> <tr> <td>http://int0046.bbn.intergral.com:18080/apps/ConsumejustFunction.cfm</td> </tr> </tbody> </table> | 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 URLs of the requests that take the longest time to c |
| 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 | | | | | | |

Example Slowest Requests on A



| Request Count | Request URL |
|---------------|---|
| 3 | http://int0046.bbn.intergral.com:18080/apps/fads/test.cpuspike.cfm |
| 1,054 | http://int0046.bbn.intergral.com:18080/apps/test.cfthread.cfm |
| 1,534 | http://int0046.bbn.intergral.com:18080/apps/ConsumejustFunction.cfm |
| 1,482 | http://int0046.bbn.intergral.com:18080/apps/justFunction.cfm |
| 1,882 | http://int0046.bbn.intergral.com:18080/apps/testTime.cfm |

(click to enlarge)

| | | o m p l e t e o n a v e r a g e. | | | | |
|--|--|---|---------|--------|--------|---|
| Ex ec u t i o n T i m e (s e c) | <table border="1"> <tr><th>Execution Time (sec)</th></tr> <tr><td>182.177</td></tr> <tr><td>72.706</td></tr> <tr><td>21.945</td></tr> </table> | Execution Time (sec) | 182.177 | 72.706 | 21.945 | T h e a v e r a g e e x e c u t i o n t i m e i n s e c o n d s f o r e a c h U R L. |
| Execution Time (sec) | | | | | | |
| 182.177 | | | | | | |
| 72.706 | | | | | | |
| 21.945 | | | | | | |

i 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](#)

Application Data Overview

Application Data Overview contains the following sections:

| Name | Example | Description | | | | |
|-------------------|--|---|-------------------|---------|--------------|---|
| Application Key | <table border="1"> <tr><td>Application Key</td></tr> <tr><td>test.cfthread.cfm</td></tr> <tr><td>twitter</td></tr> <tr><td>flex2gateway</td></tr> </table> | Application Key | test.cfthread.cfm | twitter | flex2gateway | Each application filter that you define will be listed. |
| Application Key | | | | | | |
| test.cfthread.cfm | | | | | | |
| twitter | | | | | | |
| flex2gateway | | | | | | |
| | | To find out how to configure your own application filters, see ? Manage Applications. | | | | |

Example Application Data Overview

| Application Key | Amount of Views | Average Execution Time (sec) | Total Execution Time | Average Bytes Sent | Total Bytes Sent | Unique Session IDs | Unique ColdFusion IDs |
|-------------------|-----------------|------------------------------|----------------------|--------------------|------------------|--------------------|-----------------------|
| test/cfthread.cfm | 1,088 | 72.706 | 219.176.124 | 78 | 81,237.438 | 1 | 5 |
| twitter | 191,477 | 0.419 | 219.86.274 | 0 | 0.8064 | 1 | 191,478 |
| flex2gateway | 4,808 | 0.248 | 486.154 | 802 | 2,229.848 | 4,808 | 5 |
| all-data | 712,214 | 0.888 | 712.76.884 | 9,287 | 4,171.048 | 4,808 | 191,842 |
| apps | 705,406 | 0.888 | 712.61.284.884 | 9,287 | 4,167.048 | 1 | 191,842 |
| testTime.cfm | 1,892 | 19.888 | 61.486.84 | 0 | 0.8064 | 1 | 1 |
| mongodbexamples | 248,242 | 0.288 | 142.19.86.224 | 4,764 | 1,088.048 | 1 | 288 |

[\(click to enlarge\)](#)

| | | | | | | |
|-------------------------------------|--|-------------------------------------|-------------|--------------|----------|--|
| Amount of Views | <table border="1"> <tr><td>Amount of Views</td></tr> <tr><td>1,054</td></tr> <tr><td>181,477</td></tr> <tr><td>6,908</td></tr> </table> | Amount of Views | 1,054 | 181,477 | 6,908 | The total amount of times the pages within this filter have been viewed. |
| Amount of Views | | | | | | |
| 1,054 | | | | | | |
| 181,477 | | | | | | |
| 6,908 | | | | | | |
| Average Execution Time (sec) | <table border="1"> <tr><td>Average Execution Time (sec)</td></tr> <tr><td>72.706</td></tr> <tr><td>0.419</td></tr> <tr><td>0.348</td></tr> </table> | Average Execution Time (sec) | 72.706 | 0.419 | 0.348 | The average execution time in seconds for the requests within each application filter. |
| Average Execution Time (sec) | | | | | | |
| 72.706 | | | | | | |
| 0.419 | | | | | | |
| 0.348 | | | | | | |
| Total Execution Time | <table border="1"> <tr><td>Total Execution Time</td></tr> <tr><td>21h 17m 12s</td></tr> <tr><td>7d 7h 9m 56s</td></tr> <tr><td>40m 10s</td></tr> </table> | Total Execution Time | 21h 17m 12s | 7d 7h 9m 56s | 40m 10s | The total execution time for the requests within each application filter. |
| Total Execution Time | | | | | | |
| 21h 17m 12s | | | | | | |
| 7d 7h 9m 56s | | | | | | |
| 40m 10s | | | | | | |
| Average Bytes Sent | <table border="1"> <tr><td>Average Bytes Sent</td></tr> <tr><td>78</td></tr> <tr><td>0</td></tr> <tr><td>502</td></tr> </table> | Average Bytes Sent | 78 | 0 | 502 | The average number of bytes for each application filter. |
| Average Bytes Sent | | | | | | |
| 78 | | | | | | |
| 0 | | | | | | |
| 502 | | | | | | |
| Total Bytes Sent | <table border="1"> <tr><td>Total Bytes Sent</td></tr> <tr><td>81.237 KB</td></tr> <tr><td>0 bytes</td></tr> <tr><td>3.309 MB</td></tr> </table> | Total Bytes Sent | 81.237 KB | 0 bytes | 3.309 MB | The total number of bytes sent for each application filter. |
| Total Bytes Sent | | | | | | |
| 81.237 KB | | | | | | |
| 0 bytes | | | | | | |
| 3.309 MB | | | | | | |
| Unique Session IDs | <table border="1"> <tr><td>Unique Session IDs</td></tr> <tr><td>1</td></tr> <tr><td>1</td></tr> <tr><td>6,906</td></tr> </table> | Unique Session IDs | 1 | 1 | 6,906 | The number of unique session IDs per application. |
| Unique Session IDs | | | | | | |
| 1 | | | | | | |
| 1 | | | | | | |
| 6,906 | | | | | | |
| Unique ColdFusion IDs | <table border="1"> <tr><td>Unique ColdFusion IDs</td></tr> <tr><td>1</td></tr> <tr><td>181,475</td></tr> <tr><td>1</td></tr> </table> | Unique ColdFusion IDs | 1 | 181,475 | 1 | The number of unique ColdFusion IDs per application. |
| Unique ColdFusion IDs | | | | | | |
| 1 | | | | | | |
| 181,475 | | | | | | |
| 1 | | | | | | |

[Back to the top](#)

Application JDBC Overview

Application JDBC Overview contains the following sections:

| Name | Example | Description | | | | |
|------------------------|---|------------------------|-------------------|---------|--------------|---|
| Application Key | <table border="1"> <tr><td>Application Key</td></tr> <tr><td>test.cfthread.cfm</td></tr> <tr><td>twitter</td></tr> <tr><td>flex2gateway</td></tr> </table> | Application Key | test.cfthread.cfm | twitter | flex2gateway | Each application filter that you define will be listed. |
| Application Key | | | | | | |
| test.cfthread.cfm | | | | | | |
| twitter | | | | | | |
| flex2gateway | | | | | | |
| | | To find out how | | | | |

Example Application JDBC Overview

| Application Key | Average JDBC Queries | Total JDBC Queries | Average JDBC Time (sec) | Total JDBC Time | Average JDBC Rows | Total JDBC Rows |
|-------------------|----------------------|--------------------|-------------------------|-----------------|-------------------|-----------------|
| test.cfthread.cfm | 0.000 | 0 | 0.000 | 0s | 0.000 | 0 |
| twitter | 0.000 | 0 | 0.000 | 0s | 0.000 | 0 |
| flex2gateway | 0.000 | 0 | 0.000 | 0s | 0.000 | 0 |
| all-data | 1.148 | 818,190 | 0.004 | 10 3m 00s | 37,150 | 26,408,892 |
| apps | 1.161 | 818,190 | 0.004 | 10 3m 00s | 37,484 | 26,408,892 |
| testTime.cfm | 0.000 | 0 | 0.000 | 0s | 0.000 | 0 |
| mongodbsamples | 0.007 | 666,797 | 0.009 | 40m 00s | 1,630 | 387,076 |

[\(click to enlarge\)](#)

| | | | | | | |
|--------------------------------|--|--|------------|-------|---------|---|
| | | to configure your own application filters, see ?Manage Applications. | | | | |
| Average JDBC Queries | <table border="1"> <tr><td>Average JDBC Queries</td></tr> <tr><td>1.161</td></tr> <tr><td>0.000</td></tr> <tr><td>2.287</td></tr> </table> | Average JDBC Queries | 1.161 | 0.000 | 2.287 | The average number of database queries specifically from the requests within each application filter. |
| Average JDBC Queries | | | | | | |
| 1.161 | | | | | | |
| 0.000 | | | | | | |
| 2.287 | | | | | | |
| Total JDBC Queries | <table border="1"> <tr><td>Total JDBC Queries</td></tr> <tr><td>818,750</td></tr> <tr><td>0</td></tr> <tr><td>560,797</td></tr> </table> | Total JDBC Queries | 818,750 | 0 | 560,797 | The total number of database queries specifically from the requests within each application filter. |
| Total JDBC Queries | | | | | | |
| 818,750 | | | | | | |
| 0 | | | | | | |
| 560,797 | | | | | | |
| Average JDBC Time (sec) | <table border="1"> <tr><td>Average JDBC Time (sec)</td></tr> <tr><td>0.005</td></tr> <tr><td>0.000</td></tr> <tr><td>0.009</td></tr> </table> | Average JDBC Time (sec) | 0.005 | 0.000 | 0.009 | The average database execution time in seconds for the requests within each application filter. |
| Average JDBC Time (sec) | | | | | | |
| 0.005 | | | | | | |
| 0.000 | | | | | | |
| 0.009 | | | | | | |
| Total JDBC Time | <table border="1"> <tr><td>Total JDBC Time</td></tr> <tr><td>1h 3m 25s</td></tr> <tr><td>0s</td></tr> <tr><td>40m 22s</td></tr> </table> | 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. |
| Total JDBC Time | | | | | | |
| 1h 3m 25s | | | | | | |
| 0s | | | | | | |
| 40m 22s | | | | | | |
| Average JDBC Rows | <table border="1"> <tr><td>Average JDBC Rows</td></tr> <tr><td>37.466</td></tr> <tr><td>0.000</td></tr> <tr><td>1.620</td></tr> </table> | Average JDBC Rows | 37.466 | 0.000 | 1.620 | The average number of rows returned from the database for each application. |
| Average JDBC Rows | | | | | | |
| 37.466 | | | | | | |
| 0.000 | | | | | | |
| 1.620 | | | | | | |
| Total JDBC Rows | <table border="1"> <tr><td>Total JDBC Rows</td></tr> <tr><td>26,428,852</td></tr> <tr><td>0</td></tr> <tr><td>397,178</td></tr> </table> | Total JDBC Rows | 26,428,852 | 0 | 397,178 | The total number of rows returned from the database for each application. |
| Total JDBC Rows | | | | | | |
| 26,428,852 | | | | | | |
| 0 | | | | | | |
| 397,178 | | | | | | |


[Back to the top](#)

Applications

Applications contains the following sections:

| Name | Example | Description | | |
|--------------|---|--------------|--|----|
| App | <table border="1"> <tr> <td style="background-color: #0070C0; color: white;">Applications</td> <td></td> </tr> </table> | Applications | | Ea |
| Applications | | | | |

Example Applications



[\(click to enlarge\)](#)

Application Name

Application Name: test.cfthread.cfm

ch
a
p
p
l
i
c
a
t
i
o
n
f
i
l
t
e
r
t
h
a
t
y
o
u
h
a
v
e
c
o
n
f
i
g
u
r
e
d
w
i
l
l
h
a
v
e
i
t
s
o
w
n
s
u
b
s
e
c
t
i
o
n
w
i
t
h
i
n
t
h
i
s
s
e
c
t
i
o
n.

T
o
f
i
n
d
o
u
t
h
o
w
t
o
c
o
n
f
i
g
u
r
e
y
o
u
r
o
w
n
a
p
p
l
i

cat
ion
fil
ters
,
see
?
Ma
na
ge
Ap
pli
c
ati
ons
.

Sl
ow
est
Re
que
sts

| Request ID | Request URL | Execution Time (ms) |
|--------------|---|---------------------|
| 10.10.10.101 | http://localhost:8080/intergral.com/10000/applytest/afterwait.php | 175.210 |
| 10.10.10.102 | http://localhost:8080/intergral.com/10000/applytest/afterwait.php | 175.210 |
| 10.10.10.103 | http://localhost:8080/intergral.com/10000/applytest/afterwait.php | 175.210 |
| 10.10.10.104 | http://localhost:8080/intergral.com/10000/applytest/afterwait.php | 175.210 |
| 10.10.10.105 | http://localhost:8080/intergral.com/10000/applytest/afterwait.php | 175.210 |

(click to enlarge)

This section will only show you the slowest requests for this application. It is displayed

a
m
e
a
s
t
h
e
S
l
o
w
e
s
t
R
e
q
u
e
s
t
s
s
e
c
t
i
o
n.

S
l
o
w
e
s
t
R
e
q
u
e
s
t
s
o
n
A
v
e
r
a
g
e



(click to enlarge)

T
h
i
s
s
e
c
t
i
o
n
w
i
l
l
o
n
l
y
s
h
o
w
y
o
u
r
e
q
u
e
s
t
s
t
h
a
t
m
a
t
c
h
t
h
e
a
p
p
l
i
c
a
t
i
o
n
f
i
l
t
e
r,
g
r
o
u
p
e
d
b
y
U
R
L.

I
t
i
s
d
i
s

pl
a
y
e
d
th
e
s
a
m
e
a
s
th
e
Sl
o
w
e
st
R
e
q
u
e
st
s
o
n
A
v
er
a
ge
s
e
cti
o
n.

[Back to the top](#)

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**.

The Daily Status Report configuration page

THE DAILY STATUS REPORT CONFIGURATION

The Daily Status Reports come with a new report configuration feature will allows you to add email recipients to the report. This feature can be found within the FusionAnalytics Client.

There are two ways to access the client, either through a web browser (see The FusionAnalytics Web Client) or using our Air Dashboard (see The FusionAnalytics AIR Client).

Where to go from the client

Once you've accessed the client, you'll find a new ToC item under Administrator > Daily Status Report Configuration.

Clicking this will display the new configuration screen which may contain existing report configurations if some have previously been created.

Adding email recipients.

On this page:

Adding a recipient from FADS users

Adding a custom email address



Daily Status Report Configuration

This page allows you to configure the Daily Status report email recipients.

Daily Status Report Email Recipients

| Remove | Email Address | Description |
|--|---------------|-------------|
| No recipients have been added for this report. | | |

[Add From Users](#)

Add Custom Email

Email: Description: [+](#)

This page allows you to configure who will receive the report via email.

There are two methods of adding recipients:

- Adding a recipient from FADS users. Only users who are mapped are available to add with this method.
- Adding a custom email address.

Adding a recipient from FADS users.

To add from FADS users, simply click the Add From Users button to view a list of viable FADS users and click the [+](#) button to add them to the recipients list.

Daily Status Report Email Recipients

| Remove | Email Address | Description |
|--|---------------|-------------|
| No recipients have been added for this report. | | |

[Add From Users](#)

Only users who are mapped to this application can be added from here.

| Add | Username | Email Address |
|-------------------|-----------|------------------------|
| | admin | No email address |
| + | testUser1 | demouser2@integral.com |
| | | |
| | | |
| | | |

Adding a custom email address

To add a custom email address, simply enter the email address and optional description into the corresponding fields and click the [+](#) button.

SLA Reports

SLA REPORTS

OVERVIEW

The SLA Report provides business-level information on your server's performance.

You can choose to receive daily, weekly and monthly reports, compare to previous reports and even change the way the reports are displayed based on whether your server met your targets.

If you are looking for a quick start guide for getting your SLA reports up and running as soon as possible, see [Setting up the SLA Reports in the Quick Start Guide](#).

Find out more

Accessing the SLA Report Configuration Page

Where to find the new report configuration area.

Adding a new SLA Report Configuration

How to add new reports and configurations to your application.

The SLA Report Configuration

Understand what each of the configuration elements mean and how they should be used.

NEXT STEPS

Accessing the SLA Report Configuration Page

Adding a new SLA Report Configuration

The SLA Report Configuration

- Report Configuration Tab
- Email Recipients Tab
- Domain Filters Tab
- Thresholds Tab

Running the SLA Reports

Accessing the SLA Report Configuration Page

Accessing the SLA Report configuration page

The SLA Reports come with a new report configuration feature will allows you to change the way reports will behave. This feature can be found within the FusionAnalytics Client.

There are two ways to access the client, either through a web browser (see [The FusionAnalytics Web Client](#)) or using our Air Dashboard (see [The FusionAnalytics AIR Client](#)).

Where to go from the client

Once you've accessed the client, you'll find a new ToC item under Administrator > SLA Report Configuration.

Clicking this will display the new configuration screen which may contain existing report configurations if some have previously been created.

Next Steps

Adding a new SLA Report Configuration

Discover how to add new SLA report

The SLA Report Configuration

Learn what each of the configuration options mean and how to use them.

Take me back to [FusionAnalytics Reporting](#).

Adding a new SLA Report Configuration

ADDING A NEW SLA REPORT CONFIGURATION

There are a two ways you can add a report configuration:

- Create a new configuration
- Copy a configuration

Create a new configuration

Whether configurations already exist or not you will see the Add New Report section. Enter a suitable report name and press the "Add Report" button.

Add new report

Report Name:

Report names cannot contain the following characters: ! # \$ % ^ & * () + = ~ ` [] \ ' ; , { } | " < > ?

Copy a configuration

If you already have an existing configuration, you can create a copy using the button.

You will be prompted to enter a name for the new configuration. Again, names cannot use the characters shown above.

Once you have added a new report, the configuration options for it will be added to the page and automatically expanded for editing.

Current SLA Reports

Name

DemoReport

Report Configuration | Email Recipients | Domain Filters | Thresholds

Run report:

Daily
 Weekly
 Monthly

Email Prefixes

Custom email title:

Successful Report:

Failed Report:

NEXT STEPS

The SLA Report Configuration

Learn what each of the configuration options mean and how to use them.

Report Configuration Tab

Choose from daily, weekly and monthly reports or any combination of the three, and customise how the reports will appear in your inbox.

Email Recipients Tab

Learn how to select who should receive the report via email.

Domain Filters Tab

Understand the different options that can be found on the Domain Filters tab.

Thresholds Tab

Configure your report to show you whether your server is performing correctly at a glance.

Take me back to **FusionAnalytics Reporting**.

The SLA Report Configuration

THE SLA REPORT CONFIGURATION

The report configuration allows you to tweak the settings for each report. You can add as many reports as you like.

Your reports will be listed on the page with their configuration options hidden:

| Icon | Name | Description |
|----------------------|--------|--|
| ? Unknown Attachment | Edit | Opens the configuration options for this report to allow you to make changes. |
| ? Unknown Attachment | Rename | Allows you to rename the report. Names cannot already exist, contain spaces or the characters ! # \$ % ^ & * () + = ~ ` [] \ ' ; , { } " < > ? . |
| ? Unknown Attachment | Copy | Creates a copy of this report with a new name of your choosing. |
| ? Unknown Attachment | Delete | Deletes this report. |

NEXT STEPS

Start configuring your SLA Reports: Report Configuration Tab

REPORT CONFIGURATION TAB

Report Configuration tab

This tab allows you to choose when you would like the report to run as well as some additional settings.

Schedule

| Time period | Description |
|-------------|---|
| Daily | Your report will run every day showing data from the previous day.. |
| Weekly | Checking this box will allow you to choose from 3 different week periods. <ul style="list-style-type: none">Weekdays (M, T, W, T, F)Week starting Monday (M, T, W, T, F, S, S)Week starting Sunday (S, M, T, W, T, F, S) Your report will run every week showing data for the previous week based on the type you selected. |
| Monthly | Your report will run on the 1st of each month showing data from the entire previous month. |

Email Prefixes

These settings allow you to change the email subject that will be displayed in your recipients inbox.

The default subject is: "[fusionanalytics] [timeframe] SLA Report for [Report_Name] ([DataCollector_Name]) from [Date].

For this example, our default title would look like:

Other steps in this tutorial:

[Accessing the SLA Report Configuration Page](#)

[Adding a new SLA Report Configuration](#)

The SLA Report configuration

- [Report Configuration Tab](#)
- [Email Recipients Tab](#)
- [Domain Filters Tab](#)
- [Thresholds Tab](#)

[Running the SLA Reports](#)

On this page:

[Schedule](#)

[Email Prefixes](#)

Other steps in this tutorial:

[Accessing the SLA Report Configuration Page](#)

[Adding a new SLA Report Configuration](#)

The SLA Report configuration

- [Report Configuration Tab](#)
- [Email Recipients Tab](#)
- [Domain Filters Tab](#)
- [Thresholds Tab](#)

[Running the SLA Reports](#)

[fusionanalytics] Daily SLA Report for DemoReport (cfsrv-2) from Monday, 12 March 2012

Custom email title

This setting allows you to replace the text "SLA Report for DemoReport" in the email subject.

Successful Report

A report is considered successful if all targets were met. If the report is successful, this value will directly follow [fusionanalytics] in the email subject.

Failed Report

A report is considered failed if one or more targets were not met. If the report has failed, this value will directly follow [fusionanalytics] in the email subject.

This is an example of how these settings will effect the title we were getting before on a failed report:

[fusionanalytics] (Fail) Daily Demonstration report for documentation (faDataCollector-1) from Monday, 12 March 2012

Next Steps

Configure who will receive the reports via email: Email Recipients Tab

EMAIL RECIPIENTS TAB

Email Recipients tab

This tab allows you to configure who will receive the report via email.

There are two methods of adding recipients:

- Adding a recipient from FADS users. Only users who are mapped are available to add with this method.
- Adding a custom email address.

Adding a recipient from FADS users.

To add from FADS users, simply click the Add From Users button to view a list of viable FADS users and click the button to add them to the recipients list.

Adding a custom email address

To add a custom email address, simply enter the email address and optional description into the corresponding fields and click the button.

On this page:

Adding a recipient from FADS users

Adding a custom email address

Other steps in this tutorial:

Accessing the SLA Report Configuration Page

Adding a new SLA Report Configuration

The SLA Report Configuration

- **Report Configuration Tab**
- **Email Recipients Tab**
- **Domain Filters Tab**
- **Thresholds Tab**

Running the SLA Reports

Next Steps

Select which domains should be included or excluded on your report: Domain Filters Tab

[DOMAIN FILTERS TAB](#)

Domain Filters Tab

The SLA report shows request information for all domains by default. This tab allows you to only show information for specific domains, or show information for all domains except the ones you choose.

Filter Type

| Method | Description |
|---------|--|
| Include | Selecting this option will only show information on the report for domains you add to the list. Any others will be ignored. If you add incorrect domains or domains that receive no requests to the list they will also be displayed on the report. |
| Exclude | Selecting this option will show all domains on the report except the domains you added to the list. |

Example

Assume our server takes requests on the following domains:

- domain1.intergral.com
- domain2.intergral.com
- domain3.intergral.com

We might want to configure one report to display all of this information, and another that only contains information on "domain1.intergral.com" and "domain2.intergral.com". For this we would choose the "Include" option and add our domains to the list.

Now only information for those two domains would be visible on the report.

Next Steps

Set threshold targets to define successful and failed reports: Thresholds Tab

[THRESHOLDS TAB](#)

Thresholds Tab

Thresholds tab

This tab allows you to set target thresholds for your SLA report. These targets will be used to determine whether your report was successful or not and will alter the appearance of the report.

| Name | Default Value | Enabled By Default | Description |
|--------|---------------|--------------------|---|
| Uptime | 10 | Yes | This allows you to set the minimum uptime allowed before your |

On this page:

[Filter Type](#)

[Example](#)

Other steps in this tutorial:

[Accessing the SLA Report Configuration Page](#)

[Adding a new SLA Report Configuration](#)

[The SLA Report Configuration](#)

- [Report Configuration Tab](#)
- [Email Recipients Tab](#)
- [Domain Filters Tab](#)
- [Thresholds Tab](#)

[Running the SLA Reports](#)

Other steps in this tutorial:

[Accessing the SLA Report Configuration Page](#)

[Adding a new SLA Report Configuration](#)

[The SLA Report Configuration](#)

- [Report Configuration Tab](#)

| | | | |
|--------------------------|--------|-----|---|
| | 0% | | report is flagged as failed. |
| Outages | 0 | Yes | This is the number of times your server can go offline before your report is flagged as failed. This value is per-day, so it will be multiplied for the number of days the report runs. Example: If your outages target is 0 and you run a 7-day report, your target for that report will be 7 outages. |
| Total Requests | 0 | No | The minimum number of requests on your server before this report is flagged as failed. This is not enabled by default. This value is per-day, so it will be multiplied for the number of days the report runs. Example: If your outages target is 0 and you run a 7-day report, your target for that report will be 7 outages. |
| Average Page Performance | 1000ms | Yes | If the average request execution time is longer than the specified value your report will be flagged as failed. |

- **Email Recipients Tab**
 - **Domain Filters Tab**
 - **Thresholds Tab**
- Running the SLA Reports**

Running the SLA Reports

RUNNING THE SLA REPORTS

By default, the SLA Report will run at 4:00am every morning. There are two steps involved with this:

1. FusionAnalytics will generate all the data required for the reports that are going to run.
2. The SLA reports will be generated and the emails will be sent.

Depending on the size of your database and the type of report that will run (daily, weekly, monthly) the data generation process can take a long time. Once this has run however, reports should run instantly if you choose to re-run them.

Re-running a report manually

Running the SLA Report Task manually will cause all configured SLA reports for that application to be emailed again. This can be useful when you are creating new report configurations and would like to see how they will appear in your inbox.

To run the report task manually, log into FusionAnalytics DataServices (FADS) and navigate to Scheduled Tasks > Application Tasks Summary. Here you'll see a list of tasks for your applications including the "SLA Report" task.

Clicking the "Run Now" button will force the task to run. If the task has already run that day then the process should be almost instant. You can tell the task is finished when the status changes from "Running" to "Scheduled".

If you have configured your reports correctly they should now be waiting in your inbox.

Changing when the report automatically runs

You can change the default scheduled time of 4:00am from the same page. Simply click the Edit button for the "SLA Report" task for your application to be presented with the task scheduler:

Do not change any other settings!

The report task must run every day to ensure the correct data is generated. This schedule page was designed for non-configurable reports. If you wish to disable the SLA Report emails simply uncheck the schedule options in the client configuration. See Report Configuration Tab.

- On this page:**
- Re-running a report manually**
 - Changing when the report automatically runs**

NEXT STEPS

Having trouble receiving SLA Reports via email? See SLA Report FAQ

Using FusionAnalytics to analyze a cluster of servers / instances

Clusters

The user now has the ability to group their servers together into a 'Cluster'. This will allow you to easily compare the performance of each independent cluster member, and identify which member is causing problems.

You can also see the average performance of the entire cluster itself to get a general idea of how the cluster is performing.

Any number of servers, and any combination of servers, can be added to a cluster at any time. There is no limit on how many clusters can be configured.

Configuring a Cluster

Once you have accessed the client, you'll find a new ToC item under Administrator > Manage Clusters.

Clicking this will display the new configuration screen which may contain existing cluster configurations if some have previously been created.

The screenshot shows a web interface for managing clusters. At the top, there's a 'Manage Clusters' tab. Below it, a 'Cluster Configuration' section contains a descriptive sentence: 'FusionAnalytics allows you to represent a cluster by grouping together servers.' Underneath is an 'Add New Cluster' section with a dark header, followed by two input fields: 'Cluster Name' and 'Description', and an 'Add Cluster' button. Below that is an 'Existing Clusters' section with a table. The table has two columns: 'Actions' and 'Cluster Name'. The 'Cluster Name' column contains the text 'test cluster (test cluster)'. The 'Actions' column contains four icons: a gear, a pencil, a document with a plus sign, and a red 'X'.

Your existing clusters will be listed with their configuration options hidden.

| Icon | Name | Description |
|----------------------|--------|---|
| ? Unknown Attachment | Edit | Opens the configuration options for this cluster to allow you to make changes. |
| ? Unknown Attachment | Rename | Allows you to rename the cluster. Names cannot already exist, contain spaces or the characters ! # \$ % ^ & * () + = ~ ` [] \ ' ; , { } " < > ? . |
| ? Unknown Attachment | Copy | Creates a copy of this cluster with a new name of your choosing. |
| ? Unknown Attachment | Delete | Deletes this cluster. |

ADDING A NEW CLUSTER

You can add as many clusters as you like.

There are a two ways you can add a report configuration:

- Create a new configuration
- Copy a configuration

Create a new configuration

On this page:

Configuring a Cluster

Adding a new Cluster

Configuration Options

Whether configurations already exist or not you will see the Add New Cluster section. Enter a suitable cluster name and description(optional) and then press the "Add Cluster" button.

Cluster names cannot contain the following characters: ! # \$ % ^ & * () + = ~ ` [] \ ' ; , { } | " < > ?
You cannot use a name if a cluster already exists with that name.

Copy a configuration

If you already have an existing cluster configured, you can create a copy using the button.

You will be prompted to enter a name for the new configuration. Again, names cannot use the characters shown above, and cannot match the name of an existing cluster.

CONFIGURATION OPTIONS

The following page is displayed when adding a new cluster or when clicking the icon.

The Servers tab allows you to choose which servers you want to be included in the cluster. Any existing servers will appear in the drop down box.

To add a server to the cluster simply select it from this drop down box and add it using the button.
Adding a description is optional.

The server will then appear in the table underneath and will be saved to your cluster. You can add any number of existing servers to a cluster at any time.

To remove a server from the cluster, simply click the button in the table, next to the server you wish to remove. You can remove a server from the cluster at any time.

Next Steps

Cluster Perspectives

Learn about the new types of perspectives and where to find them.

Cluster Perspectives

Cluster Perspectives

Before you are able to use functionality of clusters, you must set up the cluster correctly. See [Setting Up A Cluster](#).

CHANGES TO THE TOC

There have been some minor changes to the ToC to accommodate clusters. You will now see that there is a cluster folder at the top.

The folders for the individual servers(e.g. fa-DataCollector-1) and all their sub-folders have now been moved into a folder named "Servers" . This is a simplified, easier to understand, layout for the ToC.

If a cluster has been added, but it contains no servers, the message "No Servers" will be displayed next to the folder, and you will be prompted to the "Manage Clusters" page so that you can add servers.

When a cluster has been set up, you can access all the usual perspectives that are included in the Server folders, by expanding the Cluster folder from the ToC.

The perspectives are ordered in the same way, with the same folder structure as in Servers.

NEW VIEWS

The cluster perspectives are based off the same perspectives in the Server folders. I.e. Slowest Requests for Clusters, works in the same way as Slowest Requests for Servers, but it displays data for all the servers in the Cluster.

There are two basic types of view within the cluster perspectives. These are the **Server Comparison View** and the **Cluster View**

Server Comparison View

In a server comparison view, the data from one of the servers in the cluster is taken and displayed directly alongside the data from the other servers in the cluster.

This allows the user to directly compare the performance of each of the servers that are included in the cluster.

In data grids, such as 'Slowest Requests'(figure 1), the Server Comparison Views have an extra column which displays the **Server Name**. This is so that you can see which server any particular data entry comes from.

Figure 1: Server Comparison Example. Shows the Slowest Request perspective for a cluster. You can see which server had the slowest request of the day.

In Graphs, such as 'Memory Usage', the Server Comparison Views are displayed as a line graph where each line (or 'series') represents the data from one particular server.

For example, in Memory Usage, the 'Used Memory %' for all servers is displayed on the same line graph. In Fig.2 you can see that the data for server 1 is displayed as the blue line, and the data for server 2 is displayed as the orange line. In this example, both servers are performing equally.

Figure 2: Server Comparison Example. Shows the Used Memory(%) graph from the Memory Usage perspective for a cluster.

Cluster Views

The cluster view is used when the data for the cluster is an aggregate rather than specific to a server. For some perspectives the data needs to be averaged, such as "Slowest Requests On Average", where execution times of the requests need to be averaged across the entire cluster.

We call these cases '**Cluster Views**'. An example is given in Fig.3, in which you can see the cluster data for Slowest Requests On Average

Figure 3: Cluster View example. Showing the Slowest Requests On Average perspective for a cluster.

In Fig.3, you can see that the perspective looks the same as the perspective in Servers. The same columns are displayed, however, this is a cluster view, so in this case the 'Request Count' for a particular request is summed up for all servers in the cluster. The 'Exec Time' is taken as an average across all the servers in the cluster.

Below is an example of a cluster view for graphs. Again this appears the same as the perspective for Servers, however, the data for each series is taken as either an average or a sum across all the servers in the cluster.

Figure 4: Cluster View example. Shows the Request and JDBC Execution Time perspective for a cluster.

FAQs

Frequently Asked Questions (FAQ) for FusionAnalytics

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.

FAQ Sections

GENERAL FAQ

SERVER FAQ

DATACOLLECTOR FAQ

FUSIONANALYTICS APPLICATION FAQ

DESKTOP CLIENT FAQ

LICENSING FAQ

SLA REPORT FAQ

General FAQ

General FAQ

- What's the difference between FusionReactor and FusionAnalytics?
- Do I need FusionReactor to be able to use FusionAnalytics?
- Do I need a separate database for FusionAnalytics?
- Which Version and Edition of FusionReactor does FusionAnalytics require?
- Should I have a separate server for FusionAnalytics?
- How can I evaluate FusionAnalytics?
- Does the edition of Adobe ColdFusion I have effect whether or not I can use FusionAnalytics?
- How can I find more help on FusionAnalytics?
- Which levels of support services do you offer for FusionAnalytics?
- How do I find the size of the database?

WHAT'S THE DIFFERENCE BETWEEN FUSIONREACTOR AND FUSIONANALYTICS?

FusionReactor is a professional server monitor, it captures metrics (instrumentation) as the server is running and it is able to store many requests online for an immediate diagnostic view about the health of your server, but it is not able to provide information about long-term data - this keeps FusionReactor compact and ensures that overhead is kept to an absolute minimum.

FusionAnalytics is a visual tool for analyzing the data which has been provided by FusionReactor. FusionAnalytics loads the extensive log file data captured by FusionReactor into the FusionAnalytics DB. It then provides numerous graphs, charts, tables and reports, to visualize the data. FusionAnalytics is all about "Making IT Better" and giving you the information and data to help you make better business decisions, improve application performance and quality of your applications as well as measure exactly how your applications are performing on a continuous day to day basis.

DO I NEED FUSIONREACTOR TO BE ABLE TO USE FUSIONANALYTICS?

Yes – FusionAnalytics relies on the metric (log) files which have been generated from FusionReactor.

DO I NEED A SEPARATE DATABASE FOR FUSIONANALYTICS?

Yes – FusionAnalytics relies on using a separate database. For information on supported databases - please refer to our system requirements page.

WHICH VERSION AND EDITION OF FUSIONREACTOR DOES FUSIONANALYTICS REQUIRE?

FusionAnalytics works with both the Standard and Enterprise versions of **Fusion Reactor 4** and above.

SHOULD I HAVE A SEPARATE SERVER FOR FUSIONANALYTICS?

Yes – FusionAnalytics analyzes your historical monitoring data. The amount of data you collect and the processing required to 'crunch' this information can be high (depending on how busy your target platform is). For this reason, we recommend installing FusionAnalytics on a separate machine. View more information about hardware requirements on our system requirements page.

HOW CAN I EVALUATE FUSIONANALYTICS?

We offer a **30-DAY** free trial of FusionAnalytics. After 30 days, the Trial simply disables itself, without the need to uninstall it. Evaluators can turn the Trial into any full version by entering a valid license key without reinstalling the software.

DOES THE EDITION OF ADOBE COLD FUSION I HAVE EFFECT WHETHER OR NOT I CAN USE FUSIONANALYTICS?

Indirectly. FusionAnalytics relies on FusionReactor 4 or above – FusionReactor works with ColdFusion 6, 7, 8, 9 and 10 You can purchase either the FusionReactor Enterprise or Standard Edition, regardless of which version of ColdFusion you currently wish to monitor.

HOW CAN I FIND MORE HELP ON FUSIONANALYTICS?

We offer a wide range of resources and documentation on the FusionAnalytics product. As well as this online documentation, you can also refer to our website support area which contains our Knowledge Base and links to Google groups.

WHICH LEVELS OF SUPPORT SERVICES DO YOU OFFER FOR FUSIONANALYTICS?

Documentation, tutorials, knowledge base and our Google groups are available for free to help you get up and running. Additionally, we offer a Quick Start Installation package - in this case one of our trained engineers will install & configure FusionAnalytics for you. This is the fastest, hassle-free way to get FusionAnalytics installed. We also offer Fusion Premium Support which is a paid support option and can be purchased on an incident by incident basis. We also provide professional onsite and remote ColdFusion Consulting Services if you have serious CF problems which you need detail assistance with. Take a look at our Support Area for a complete overview.

HOW DO I FIND THE SIZE OF THE DATABASE?

In FusionAnalytics the size of the database is displayed on the Application Detail page under the Database Details section.

Alternatively you can display the size of the database in the Microsoft SQL Server Management Studio by selecting the FA database and choosing `Properties` from the context menu.

[Back to the top](#)

Server FAQ

Server FAQ

CAN I RUN MY OWN WEBAPPS IN FUSIONANALYTICS?

FusionAnalytics consists of two webapps: **Data Services (FADS)** and **Data Collector (FADC)**. These run within a Tomcat application server as separate webapps. Although Tomcat doesn't restrict how many webapps can be run within one container, we do not support adding any further webapps to this Tomcat instance.

WE ALREADY HAVE A STANDARD J2EE CONTAINER, CAN I RUN FUSIONANALYTICS IN THAT?

No; the only supported container for FusionAnalytics is the one that we ship with the installer. All other containers and versions are unsupported.

HOW DO I INCREASE THE MEMORY USED BY FUSIONANALYTICS?

To change the VM memory settings for the FusionAnalytics Server do the following:

1. Stop the FusionAnalytics Service
2. Backup `\FusionAnalytics\Server\bin\FusionAnalytics Server.vmoptions`
3. Edit `\FusionAnalytics\Server\bin\FusionAnalytics Server.vmoptions`

Edit the line:

```
-Xmx768m
```

and increase the memory for example to 1GB (if the server has the memory available).

```
-Xmx1024m
```

4. Start the FusionAnalytics Service

When would I do this?

Typically you do not have to alter the default settings. However, if you manually increase the number of import threads, the JVM could run out of memory. We have tested bulk imports of data consisting of more than 250000 files with more than 60GB in size without any problems.

See also: [Windows Service Configuration](#).

WHAT SPECIFICATION SERVER DOES FA REQUIRE?

DataServices: FADS operates an asynchronous message service to deal with requests from clients. The faster the system, the faster responses will be processed and delivered to the client systems. Memory is required to perform analytics on data sets (although as much of this as possible is offloaded to the database) and to prepare and transfer analytics to the client.

DataCollector: FADC benefits most of fast disk I/O as it's main task is to read files from the file system, parse and import them to the database. The database should be best located on a dedicated computer and not on the same machine as FADC.

For more information on FusionAnalytics System Requirements, [click here](#).

CAN I RUN FUSIONANALYTICS ON AMAZON EC2?

Yes, you will be able to run FusionAnalytics Server on Amazon EC2, however, you should be aware of matching the Amazon EC2 instance to your application on FusionAnalytics, as well as Amazon's pricing policies for EC2 before attempting to set up any instances.

Please refer to our [Amazon EC2 test system guidelines](#) page for more information.

[Back to the top](#)

DataCollector FAQ

DataCollector FAQ

- What is the DataCollector ?
- What is DCML?

WHAT IS THE DATACOLLECTOR ?

The 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 Client. 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).

WHAT IS DCML?

DCML is the abbreviation of DataCollector Markup Language. A DCML file is used to define how a FADC application works. See DCML (FADC) for more information.

[Back to the top](#)

FusionAnalytics Application FAQ

FusionAnalytics Application FAQ

- Why can't I reinstall a FusionAnalytics DataServices application with the same name?
- In FADS, What is the "Controllers" page for? What about "Asynchronous Jobs"?

WHY CAN'T I REINSTALL A FUSIONANALYTICS DATASERVICES APPLICATION WITH THE SAME NAME?

Due to the way Java locks files, when you uninstall a FADS application, the actual files may not be immediately deleted. On some operating systems, the files are locked until the server stops. In this case, FADS will delete them when the server restarts. Until this deletion occurs, FADS won't allow you to install another application with the same name. If you require the application to have the same name as a deleted application, restart the service to ensure the files are actually deleted. You should then be able to add the application under the required name.

IN FADS, WHAT IS THE "CONTROLLERS" PAGE FOR? WHAT ABOUT "ASYNCHRONOUS JOBS"?

The Controllers page details all the modules which make up DataServices. Some modules have optional operations (e.g. Dump, Clear etc.). These are intended for debugging purposes, should you need to get in touch with support. You'll be directed to use these operations as required for troubleshooting.

The Asynchronous Jobs page shows the last few requests that clients made for analytical datasets. Again, this is used for troubleshooting. A non-zero value for "cancelled" is normal, but a non-zero value for "Failed" might indicate an issue processing requests from clients.

[Back to the top](#)

Desktop Client FAQ

Desktop Client FAQ

- What is the Desktop Client ?
- Do I need to install the AIR Client to use FusionAnalytics ?
- Do I need Adobe AIR to run the Desktop Client ?
- How do I get the Desktop Application ?
- How do I add an application to the FA Desktop Application ?
- How do I login to an FA application from the AIR Desktop ?
- What is an alternative to using the Desktop Client ?

WHAT IS THE DESKTOP CLIENT ?

The Desktop Client enables the FusionAnalytics Web / Browser Client to be taken, installed and ran on the users Desktop. The Desktop application that is built using (and ran inside of) Adobe AIR, therefore it is required that the AIR runtime is installed.

Provided that the AIR runtime has been successfully installed, the Desktop Client can be installed trouble free as all files go into the app

The main advantage of using the Desktop Client is that you can quickly open and access FusionAnalytics applications and analyze data the need of the browser. The Desktop client has access to the local computer, e.g. access to local storage allowing multiple applications client.

[Back to the top](#)

DO I NEED TO INSTALL THE AIR CLIENT TO USE FUSIONANALYTICS ?

No - the AIR Client is not a mandatory install component. You can always access FusionAnalytics from the browser interface (Web Clier have the AIR client installed or not. All the same, we do recommend that you download and install the FusionAnalytics AIR Client.


NoteL In this case, you will also need to download the Adobe AIR Runtime if you've not already got it installed.

[Back to the top](#)

DO I NEED ADOBE AIR TO RUN THE DESKTOP CLIENT ?

Yes. The Desktop Client is built with and runs inside of Adobe AIR, therefore it is required that the Adobe AIR runtime has been installed can be used.

The Adobe AIR runtime can be downloaded from the Adobe Air website. Once AIR has been downloaded, simply run the installer to inst



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

Download for free: Latest version of Adobe AIR

[Back to the top](#)

HOW DO I GET THE DESKTOP APPLICATION ?

If you have not already got a copy of the FusionAnalytics Desktop Client, It is available to download from our website.

Note: Before the Desktop Application can be installed the Adobe Air runtime must also be installed.

[Back to the top](#)

HOW DO I ADD AN APPLICATION TO THE FA DESKTOP APPLICATION ?

When the Desktop Client is first loaded, an address bar (URL input box) is positioned at the top of the screen (similar to the layout of a v can be used to add applications to the client as URLs are used to identify and connect to applications.

The FusionAnalytics Application URL's use the following 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 a |
| 8400 | The machine port. If you changed the default port during the installation, use that instead; if not, use the d |

| | |
|------------------------|--|
| /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 be <i>coldfusion9</i> . |

Once the application URL has been entered into the address bar click Add Application button. The application will be added and a short

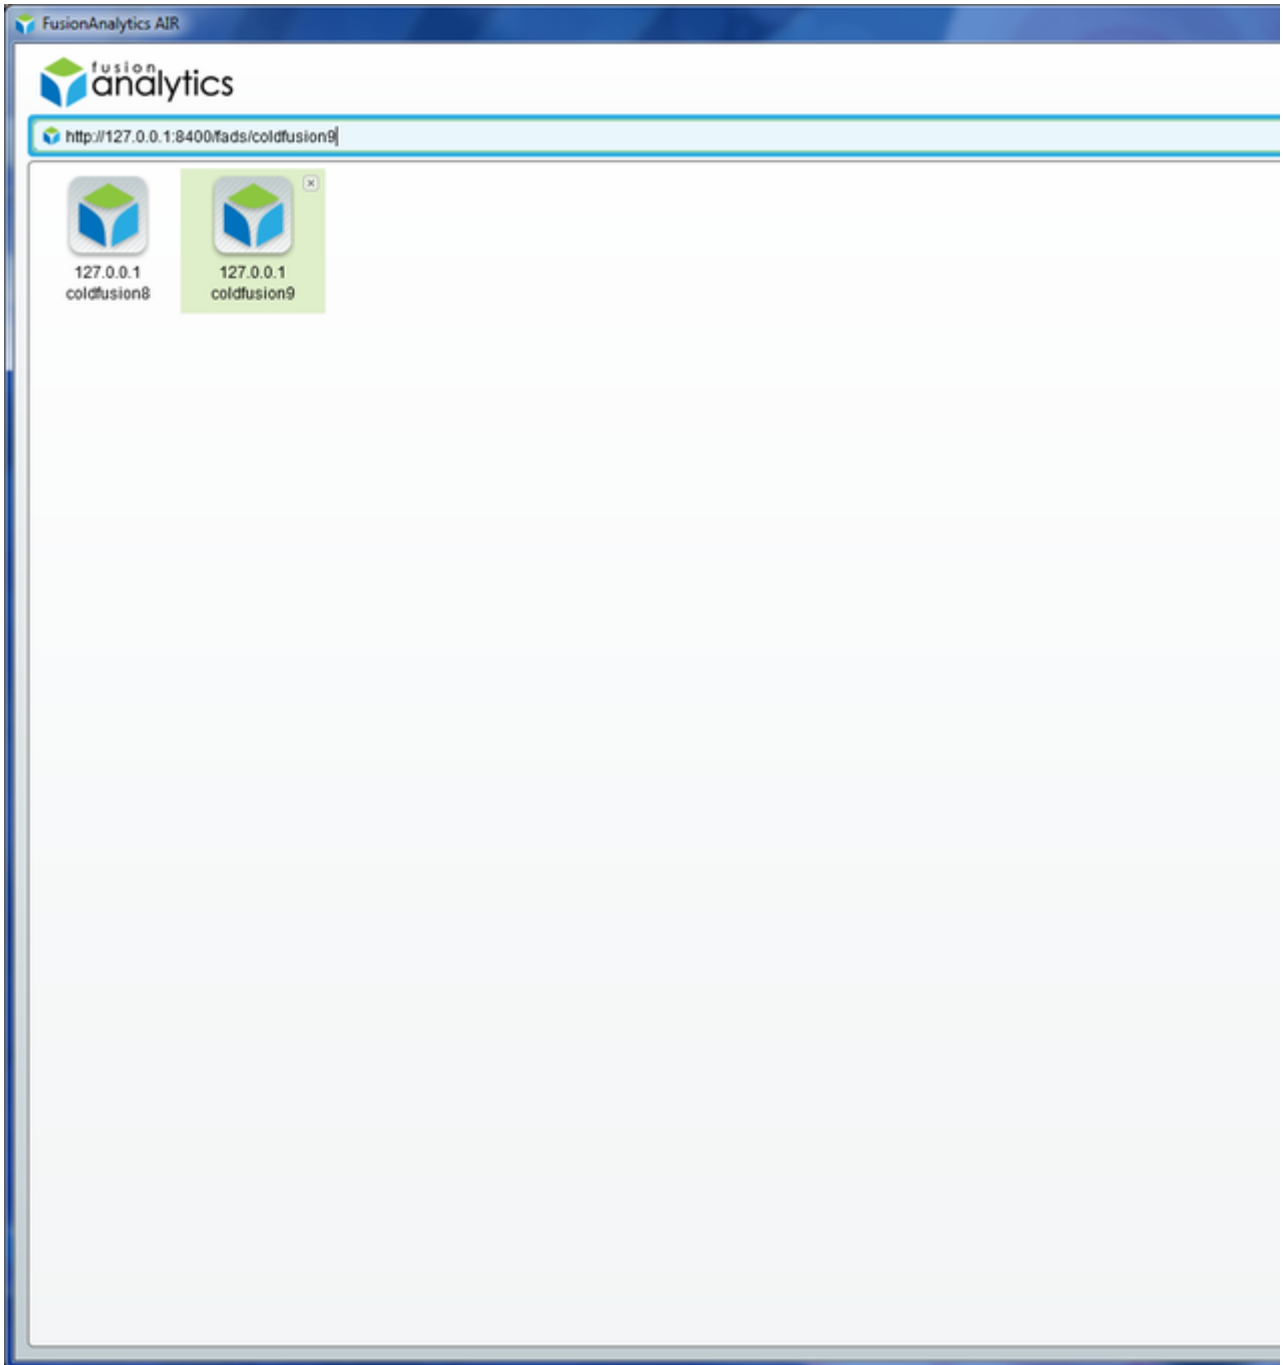


Figure 1: AIR Client Shortcuts

For more information about adding applications to the Desktop Client see: [Setting Up the AIR Client.](#)

[Back to the top](#)

HOW DO I LOGIN TO AN FA APPLICATION FROM THE AIR DESKTOP ?

After an application(s) has been added to the Desktop Client the application shortcuts (application icons) can be selected. Simply double open up the login screen for the selected application (Figure 2):

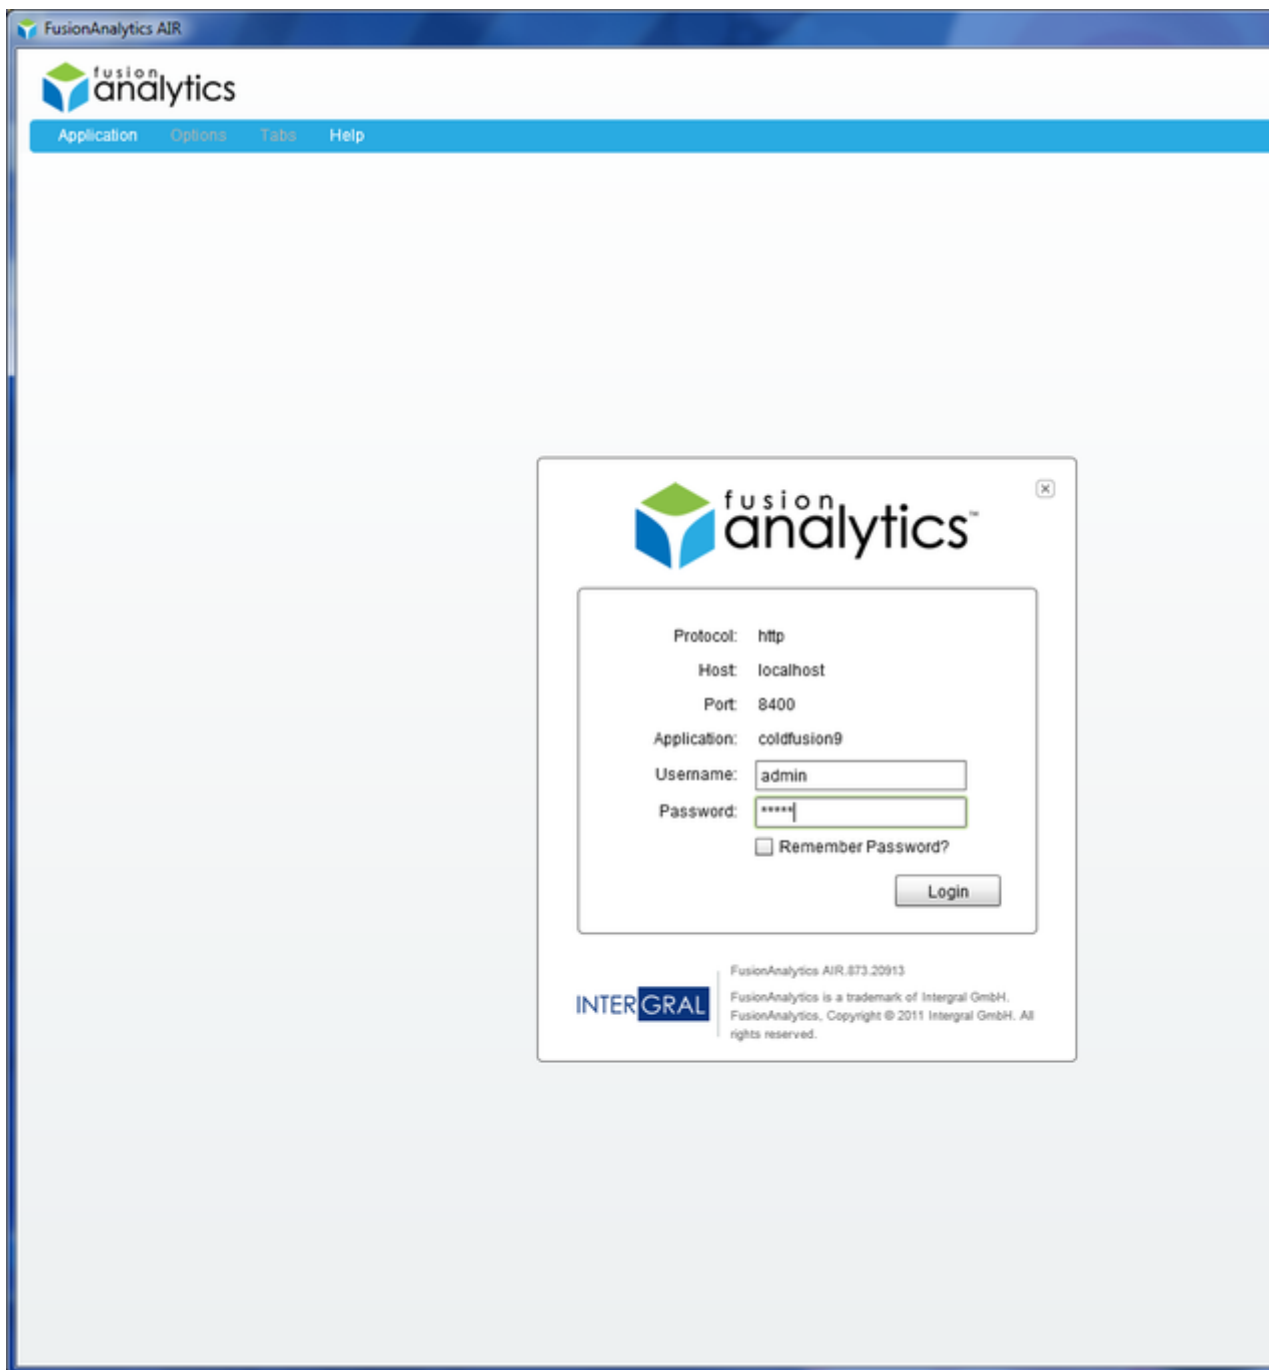


Figure 2: Login Screen

At the login page, you can enter the relevant username and password for the selected application and click the login button. If the correct username and password are entered and the application has been started the application will be logged into successfully. For more information about logging into applications from the Desktop Client see: [Setting Up the AIR Client](#).

[Back to the top](#)

WHAT IS AN ALTERNATIVE TO USING THE DESKTOP 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 FusionAnalytics DataServices (FADS) and click Launch on the application that you want to view.

[Back to the top](#)

Licensing FAQ

- What is a product family?
- What are "constraints" in a FusionAnalytics license?
- What is a user (as a license constraint)?
- What is a DataCollector (as a license constraint)?
- What is an Application (as a license constraint)?
- I tried to activate my license and was told I have no activations remaining
- I didn't activate my license in time and the activation period has expired. Have I lost my configurations?
- My license expires soon. Can I purchase another activation for it?
- I need a bigger license. Do I have to purchase a new one?
- Can I install two licenses for the same product family?

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.

In order to check the number of activations left for that license key, navigate to the FADC or the FADS application and go to **System > License Info**. There you will be able to see the section indicating the remaining DataCollector, Applications and Users. If you want a bigger license key, more information about pricing can be found here or if you want to use the same license key, then you must **uninstall the license from another machine before you can use it again**.

More information about the license activation and license deactivation can be found in the links below.

- [Installing a valid License Key](#)
- [Uninstalling a FusionAnalytics License Key](#)

[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](#)

SLA Report FAQ

SLA Report FAQ

- [Why am I not receiving SLA reports via email?](#)
- [Why can't I see my FADS users in the report configuration?](#)
- [Why can't I see my SLA Reports in the reports perspective of the client?](#)

[WHY AM I NOT RECEIVING SLA REPORTS VIA EMAIL?](#)

If you are not receiving your reports via email as expected, please make sure you have configured the following settings:

- Added email server details to FADS on the Server Settings page.
- Added a new report configuration. See [Adding a new SLA Report Configuration](#).
- Selected an appropriate time period to run the report. Daily reports are a good way to test your settings as these will always run. See [Report Configuration Tab](#).
- Configured the recipients for this report. See [Email Recipients Tab](#).

Once you have configured the steps above, try manually running the report task to see if you receive the emails.

[WHY CAN'T I SEE MY FADS USERS IN THE REPORT CONFIGURATION?](#)

For security reasons, the report configuration will only display users who are mapped to the application being edited. You can find out how to map your users here.

WHY CAN'T I SEE MY SLA REPORTS IN THE REPORTS PERSPECTIVE OF THE CLIENT?

SLA Reports are not currently supported in the reports perspective due to the introduction of the new report configuration feature. This will be added at a future date.

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](#)

Known Issues

Introduction

This section lists **answers** to some of the common known problems encountered using the latest version of FusionAnalytics. You can also find support information, updates and articles at the FusionAnalytics Website or go straight to the FAQ located there.

Common Issues

[NOT A VALID LICENSE FILE](#)

One of the most common issues with the latest version of FusionAnalytics is the license key activation. If you are not able to activate a valid license key and you get the error message "**Not a valid license file**", the section below will provide you with a couple of possible fixes.



Possible Fixes

- Make sure that a **Firewall configuration** or a **policy** is not blocking the license activation process.
- Make sure that an anti-virus software running in your system is allowing you to completely **exclude** the FusionAnalytics directory and **exclude** peeking into the web requests of the JVM process associated to FusionAnalytics.
- Make sure that the hash value of the license key is not changed. There are a number of open source tools that can be used for that purpose.

If you are still unable to activate a valid license key, please contact the FusionReactor Support Team in the following email address: support@fusion-reactor.com for more information.

MANUAL LICENSE ACTIVATION IS FAILING

Another common issue is related to the manual license activation of FusionAnalytics. If the server running FusionAnalytics is not connected to the Internet, you will be taken to the manual activation page.

Fusion Analytics is offline

Manual License Activation

Activation Input String

On a computer that is connected to the Internet go to the FusionAnalytics Online Activation page:

<http://www.fusion-analytics.com/fa/activate/activate.php>

and enter the Activation Input String shown below:

FA1Y1XXL-FE-0003041-1-T5Atr-F4F6-8960-86DD-984F-D853-4FC3-7468-F919

You will receive an Activation Key. Enter that Activation Key in the text input field below:

Activation Key

Activation Key: Enter the key as returned from the web page shown above.

Please Note: Once you have obtained a license activation key by visiting the link above, the license will be immediately marked as active in our records, and it will be difficult to reactivate it using the automatic or manual process later. For this reason, once you have obtained an Activation key, you must immediately enter it in this form, and click the "Activate License" button.

[Activate License](#)

Possible Fixes

- Make sure that a **Firewall configuration** or a **policy** is not blocking the license activation process.
- Make sure that an anti-virus software running in your system is not blocking the license activation process.
- Make sure that you have copied the whole License activation key string from the Manual License Activation page of FusionAnalytics.
- Make sure that you have press the "**Activate License**" button.

More information about the Manual License activation can be found here, [Manual License activation](#).

NOT RECEIVING THE TAP/SAP/DAILY REPORT

There are many reason why you are not able to receive any reports from FusionAnalytics. The most common is that the mail server configuration in FusionAnalytics is not correct or the configured email is not correct.

Unfortunately, there is no way to test the mail server in FusionAnalytics however, there a a number of things that you can check in order to make sure that your configuration is correct.

1. More information about the things you can check can be found in the documentation below and in section "Why am I not receiving SLA reports via email?"
 - a. <http://docs.intergral.com/display/FA207/SLA+Report+FAQ>
2. Additionally, you can try to re-run the reports manually. For more information please check out the link below.
 - a. <http://docs.intergral.com/display/FA207/Running+the+SLA+Reports>
3. Another way to test the mail server is by "abusing" the Disk Usage check in the FADC application. For example, this can be achieved by temporary changing a value that will always create/send an email alert. More information about the Disk Usage Check can be found in the following link; <http://docs.intergral.com/pages/viewpage.action?pageId=34898117>.

CANNOT SEE ANY DATA IN THE REPORTS

Another issue that often occurs is that the TAP/SAP/Daily report might contain no data and there are many reasons for that behavior. In order to solve the issue, please follow the steps mentioned below.

1. Make sure that FusionReactor communicates successfully with FusionAnalytics and the archive files are send to the defined target. If you have issues accessing the FusionAnalytics server, please make sure that a Firewall policy/configuration is not blocking that traffic.
2. Check the status of the Datacollector. This can be done by accessing the FADC application and then by navigating to **Applications > Import History**.
3. If you are able to see data in the "Import History" area, check the **fusionanalytics-datacollector-errors.log** file for error messages. This can be found by navigating to C:\FusionAnalytics\Server\logs.

If you are unsure about the error messages and you require further assistance, please feel free to get in touch with the FusionReactor Support Team on support@fusion-reactor.com.

[FNS-1] Creating Daily Status Reports in FusionAnalytics

Creating Daily Status Reports in FusionAnalytics

INTRODUCTION

A major new feature in the release candidate is the addition of a Daily Status Report. This report contains a breakdown of the server's performance on the previous day including detailed information about the requests, database activity, resource usage and even a breakdown of individual applications running. The report can be scheduled to be emailed every day to a number of users. Details on how to setup and configure the Daily Reports can be found in the article below.

1.0 SETUP

1.1 USERS

During the installation of FusionAnalytics, you will have created an administrator account. To create other accounts, you need to go to FusionAnalytics DataServices > Users tab > Users Summary. Here you will be able to see any existing user accounts and create new ones.

To add an email address to an existing user: Click the **Edit** button under the username of the account that you want to add an address to. Type an email address into the email field and click the **Modify User Details** button.

Editing 'admin'

| | |
|--------------------|---|
| First Name | <input type="text" value="System"/> |
| Last Name | <input type="text" value="User"/> |
| Password | <input type="password" value="*****"/> |
| Password (Confirm) | <input type="password" value="*****"/> |
| Email | <input type="text" value="username@company.com"/> |
| System User | <input checked="" type="checkbox"/> |

[Modify User Details](#)

To create a new user: Click the **Create A New User** button. You must fill in the username and password fields; the first name, last name and email fields are optional. You can always go back and edit these details later.

There is the option to make the account a System User, which means they will have access to all applications (this is the same settings as the administrator account). Once you have clicked **Add New User**, you can add any application mappings that the user will need. An account with an administrator role can make changes to the application, but an account with a user role can only view the data in a client.

Current Application Role Mappings

You may only modify roles for applications for which you are an administrator.

| Application | Role | Actions |
|-------------|-----------------------------------|---|
| fads-app2 | <input type="text" value="User"/> | Modify Delete |

Available Application Role Mappings

You may only add roles for applications for which you are an administrator.

| Application | Role | Actions |
|-------------------|---|--------------------------|
| fads-application1 | <input type="text" value="None (No Access)"/> | Add Role |

1.2 SERVER SETTINGS

During installation, you will have set the host address and port number of the server for FusionAnalytics. If needed, these settings can be changed under FusionAnalytics DataServices > System tab > Server Settings. This needs to be correct for the reports (and any other background tasks) to run.

The server host field is the server or IP address that FusionAnalytics is running on; the server port is the port that FusionAnalytics is configured to use.

| | |
|--------------|--|
| Server Host: | <input type="text" value="127.0.0.1"/> |
| Server Port: | <input type="text" value="8400"/> |

[Update Server Settings](#)

1.3 EMAIL SETTINGS

You will also need to make sure that the email server settings are correct before you can receive any emails from FusionAnalytics. To do this, go to FusionAnalytics DataServices > System tab > Server Settings. These fields need to be set, as any emails that are sent to and from FusionAnalytics use these as the default. This includes the sending of all notification emails.

The mail server field needs to be set to the server 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 need to put this in front of the address: **username:password@mail.company.com:25**.

1.4 DAILY STATUS REPORTS

The reports are configured in FusionAnalytics DataServices > Scheduled Tasks tab > Create A New Task. You can set the name, how long and how often the task will run. The task will run indefinitely if no end date is specified. There are several options for how often you want the task to run; as these are Daily Status Reports, you probably want to have them recurring daily at a suitable time, e.g. early in the morning when there are no other tasks running.

The **Provider Name** must be set to **DailyStatus**. You can then provide particular arguments to the report if needed (see 2.2 receiving reports as emails and 2.3 running reports for a specific date).

2.0 CUSTOMISING YOUR REPORTS

2.1 APPLICATION FILTERS

By making use of the application filters in the FusionAnalytics client, you can make sure you get the most out of the information you receive in your Daily Reports. They allow you to see data specifically for particular parts of your application.

Under the Administrator tab in your FusionAnalytics client, you need to select **Manage Applications**. Click the **Add New Application** button; specify an appropriate name and set the filter either by URL or by writing your own specific parameters. The filters allow you to sort the requests from each application by protocol, username, password, hostname, port or by a particular path element.

| Filter | Action | Argument |
|----------------------------------|--------|----------|
| URL Username (*://username:*...) | IS | admin |
| URL Port (*://*:port*...) | IS | 8501 |
| Path Element 2 (*://*/[2]*...) | IS | web |
| Path Element 3 (*://*/[3]*...) | IS NOT | apps |

Once you have added application filters, you must log out of the client and log in again to see the applications listed. You can add as many application filters as you need; these allow you to sort the data to only see requests from particular parts of the application.

2.2 RECEIVING REPORTS AS EMAILS

To send the report straight to your inbox, you must first have an email address set on your user account. To do this, see 1.1 users. The advantage of setting this up is that every time a report is scheduled to run, it will also email a copy to any users that are specified.

Once you have ensured that the required users have email addresses, you will need to create a scheduled task (see 1.4 daily status reports). In the provider arguments section, you will need to put **users** in the name field and insert **your username** in the value field.

To send the report to multiple users, in the value field insert the usernames separated by a comma and no spaces; e.g. **admin,user1, user2**. All these user accounts must have email addresses set, otherwise they will not receive the report.

2.3 RUNNING REPORTS FOR A SPECIFIC DATE

You can also set the reports to run on data from a specific day. When you are creating your scheduled task, add the following provider arguments: **date** (for the name field) and a **UTC time format in milliseconds** (for the value field). There are many free UTC calculators available on the Internet to help you calculate this value. From the date entered, the report will show the previous 24 hours.

3.0 RUNNING THE REPORTS

3.1 RUNNING DAILY STATUS REPORTS INSTANTLY

The report will run at the time you have specified, but you can also force it to run at any time by clicking **Run Now** in the FusionAnalytics DataServices > Scheduled Tasks tab > Scheduled Tasks Summary.

Once the report has run, you will see it in your FusionAnalytics client > Business Metrics > Reports. You can use the timeline to navigate to the reports from previous dates. To open a report, right click the row and select **Open Report**.

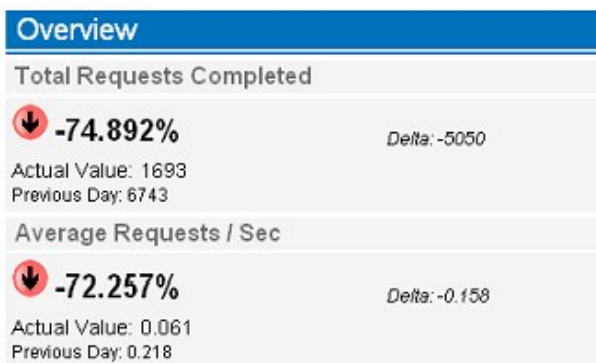
| Date/Time | Date/Time | ReportName | |
|-----------------------|-----------------------|---------------------|---------------------|
| 09.11.10 01:00:00.000 | 10.11.10 01:00:00.000 | Daily Status Report | Daily Status Report |
| 09.11.10 01:00:00.000 | 10.11.10 01:00:00.000 | Daily Status Report | Daily Status Report |

3.2 READING A DAILY STATUS REPORT


By comparing data with the previous day, the Daily Status Reports allow you to see the changes in your server from day to day. This gives you information on each section and what it all means.

Overview: As the name suggests, this section allows you to quickly see vital statistics about the requests that have run in the last 24 hours. The arrows easily show if the change is positive, negative or neutral. For example, a red arrow means a negative change, but this may point up, as an increase may indicate requests are taking longer to complete.

Each attribute also shows the amount of change as a percentage, the delta value, the actual value for the last 24 hours and the value for the previous day (for the 24 prior to the current report). Several rows also list the minimum and maximum requests that occurred between each sample.



System Load Statistics: These graphs show the system and process CPU load and memory usage. If the graphs are showing that the loads or usage is high, this can indicate that the application needs to be looked at. Underneath the graphs, there is also a percentage change from the previous day, with actual and previous day values.



Session Statistics: This gives the recorded number of session and ColdFusion IDs, including the average number of requests that run per ID.



Slowest Requests: This section allows you to easily see which requests are taking the longest to complete. Only the longest 5 requests are shown, but if viewed in a web browser, this can be expanded to show the 20 slowest requests. If the same URLs constantly give slow requests, it may need to be investigated further in the application.

Slowest Requests on Average: Similar to the above section, this allows you to see the slowest requests grouped by URL. It also shows how many times this URL was used in the request count. This section can also be expanded in a web browser to see more requests.

Application Data Overview: If you have application filters configured (see 2.1 application filters), you will see them listed in this section. This gives similar information shown in the overview and session statistics sections, but only counts data that is relevant to that part of the application.

Application JDBC Overview: Similar to the above section, it lists JDBC data that is relevant to each application filter that has been configured (see 2.1 application filters).

Applications: This section divides each configured application filter into a separate sub-section, which lists the overall slowest requests and the slowest requests on average for each one. This is the same as the previous sections on slowest requests and slowest requests on average, but with data relevant to a specific part of the application.

Copyright

Copyright

All documentation is (c) Intergral Information Solutions GmbH 2011. All Rights Reserved.

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

Responsible entity for website content may be found here.