

# **perfino Manual**

# Index

Introduction .....	3
Architecture .....	4
Installing .....	5
Monitoring JVMs .....	11
Basic concepts .....	16
UI .....	20
Transactions .....	27
Policies .....	33
Cross-VM monitoring .....	37
Probes .....	41
Method Sampling .....	45
Telemetries .....	51
Thresholds .....	55
Triggers .....	58
Alerts .....	63
End user experience monitoring .....	66
Memory .....	69
Historical comparisons .....	74
Cross-over to profiling .....	77
A Configuration .....	83
A.1 Server configuration .....	83
A.2 Server administration .....	86
A.3 Import/Export .....	89
A.4 Unattended installations .....	91
B Advanced topics .....	93
B.1 Annotation transactions .....	93
B.2 POJO transactions .....	95
B.3 DevOps transactions .....	98
B.4 Customizing net I/O methods .....	100

# Introduction To Perfino

## What is perfino?

perfino is a **monitoring tool for the JVM**. It is intended for **in-production** use and adds **extremely low overhead** to monitored applications. Its mode of operation is characterized as [APM](#), short for "application performance management". Rather than collecting performance data at a low level and with a broad scope, perfino presents selected operations at a high semantic level, called "business transactions". In addition, scalar data is monitored from a variety of sources. Based on that data, threshold violations can result in alerts that help you safeguard the quality of service of your applications.

perfino is intended to run with your application at all times. This enables it to focus on **historic data**, showing you how performance characteristics evolve over long periods of time. Data is automatically made less granular as time goes by, so you can look back years into the past with only a slow rate of storage space consumption.

perfino is designed to **monitor multiple VMs** and trace the interactions between them. Whether you have a number of fixed VM installations or a cloud deployment with hundreds of VMs, perfino can monitor and organize them at the same time.

The perfino UI is a **web interface** that can be used by multiple users to analyze the collected data at the same time. A system of access levels allows you to partition a single server for multiple groups.

## How do I continue?

This documentation is intended to be read in sequence, with later help topics building on the content of previous ones.

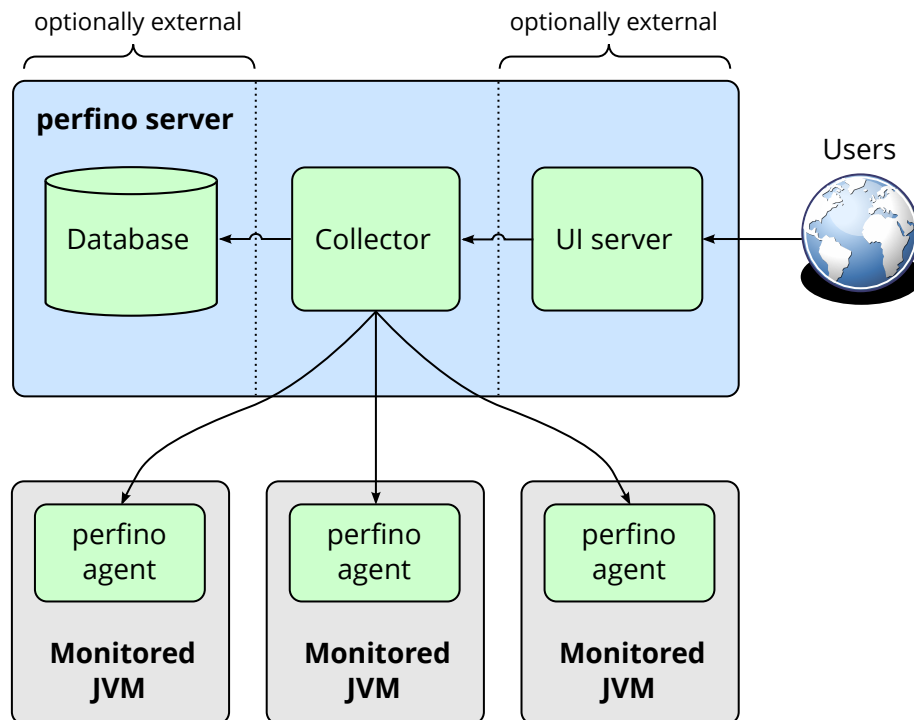
First, a broad overview over the architecture [\[p. 4\]](#) will help you to understand the components of perfino.

The help topics on installing perfino [\[p. 5\]](#) and monitoring your VMs [\[p. 11\]](#) will get you up and running.

Following that, the discussion of basic concepts [\[p. 16\]](#) and the overview of the UI [\[p. 20\]](#) take you to a level of understanding where you can explore perfino on your own.

Subsequent chapters build your expertise with respect to different functionality in perfino. The "Configurations" and "Advanced topics" sections are optional readings that should be consulted if you need certain features.

## Perfino Architecture



perfino consists of two main parts: the **server** and the **agent**. The agent is loaded in the monitored VM and records data. The agent connects to a perfino server. The server periodically queries all connected agents and processes their data. Historical information is written to a database. Users log in with their web browser to the perfino server to analyze the recorded data.

Internally, the perfino server consists of three components:

- The **collector** accepts TCP connections from perfino agents in monitored VMs. These connections can be encrypted and authenticated, so they are viable for wide area networks. The collector also consolidates data in the database, fires triggers and generates alerts.
- The **embedded H2 database** stores all persistent data. Optionally, you can use an **external mysql database** instead of the embedded H2 database. There are more tools available for the mysql database, such as hot backup and replication, so it is recommended for production.
- The **UI server** accepts HTTPS connections where users can log in, view and analyze the collected data and configure the server. Optionally, the UI server can be deployed as a WAR file into a separate JEE container like Apache Tomcat. It then uses RMI to talk to the collector. This is useful if the collector is located in an internal network and the perfino UI in a DMZ.

By default, the perfino server is the only process you need to run. In particular, you don't have to worry about installing and configuring an external database. This is ideal for evaluation and also appropriate for production with less than 10 VMs.

## Installing Perfino

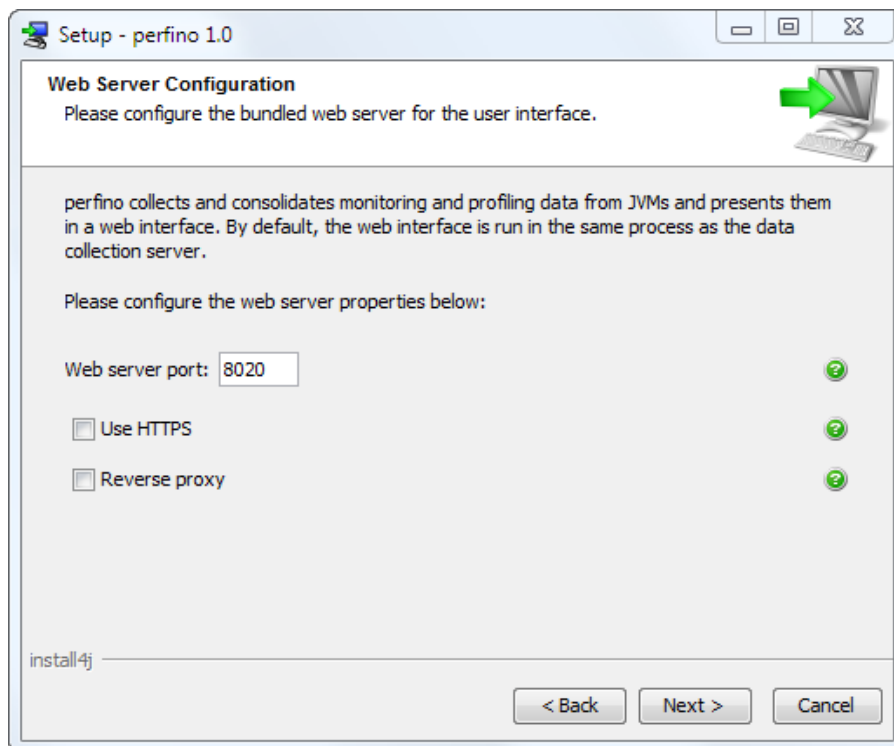
Installing perfino is done in two steps. First you run the installer. After perfino is running, you complete the post-installation setup in the browser.

### Installer

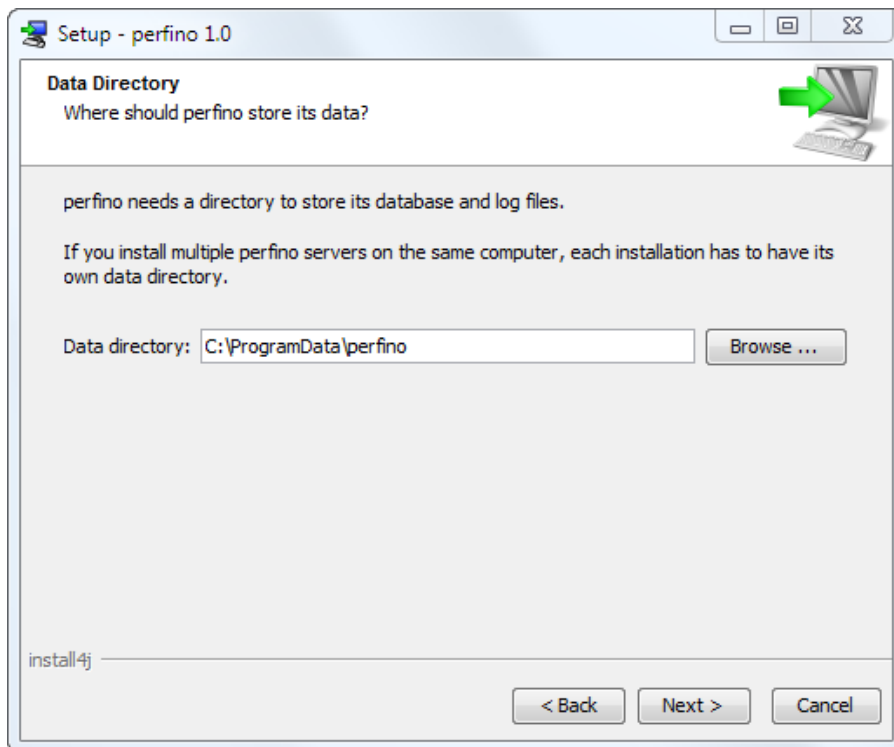
Installers are available for all major platforms. By default, the installer shows a GUI, but if you start the installer with the argument `-c`, it shows a **console interface**. This is required if you install perfino on a remote server with ssh.

For cloud-based deployments you may want to install and configure the perfino server without any user interaction. See the help topic on unattended installations [\[p. 91\]](#) for how to do that.

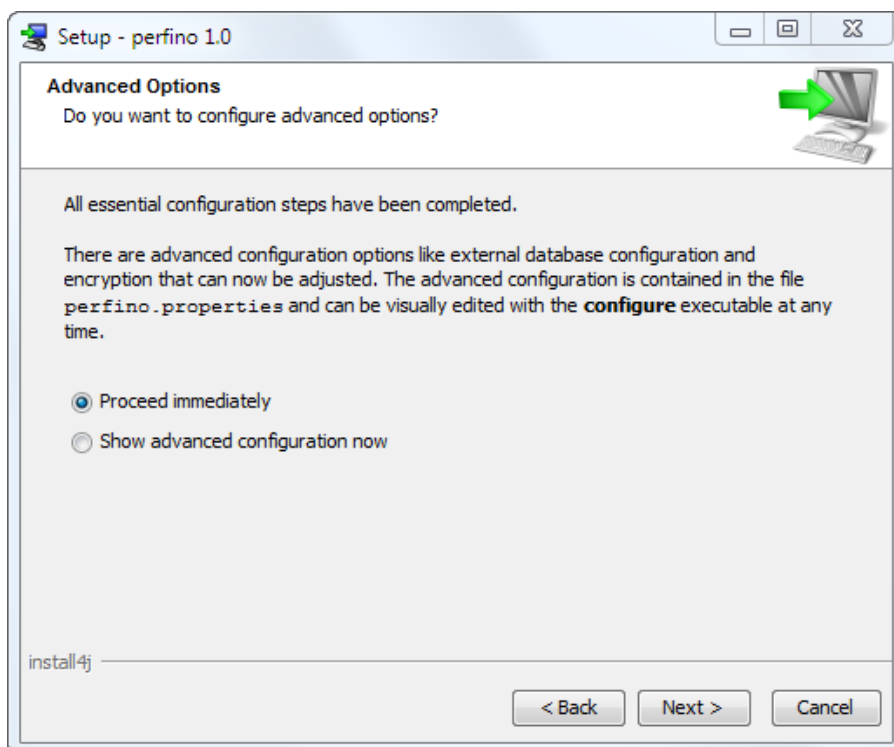
The installer asks you about all configuration options that are not configurable through the browser in the perfino UI. The most common setting that has to be adjusted is the the HTTPS/HTTP port for the web server.



Beside the installation directory, a perfino installation needs a **data directory** where all variable data is stored. This includes the database, log files and security certificates. By default, this directory is set to the program data directory on Windows (typically `C:\ProgramData`) and `/var/opt/perfino` on Unix and Mac OS X. In a scenario where you install multiple instances of perfino on the same machine, each perfino installation has to have its own data directory.

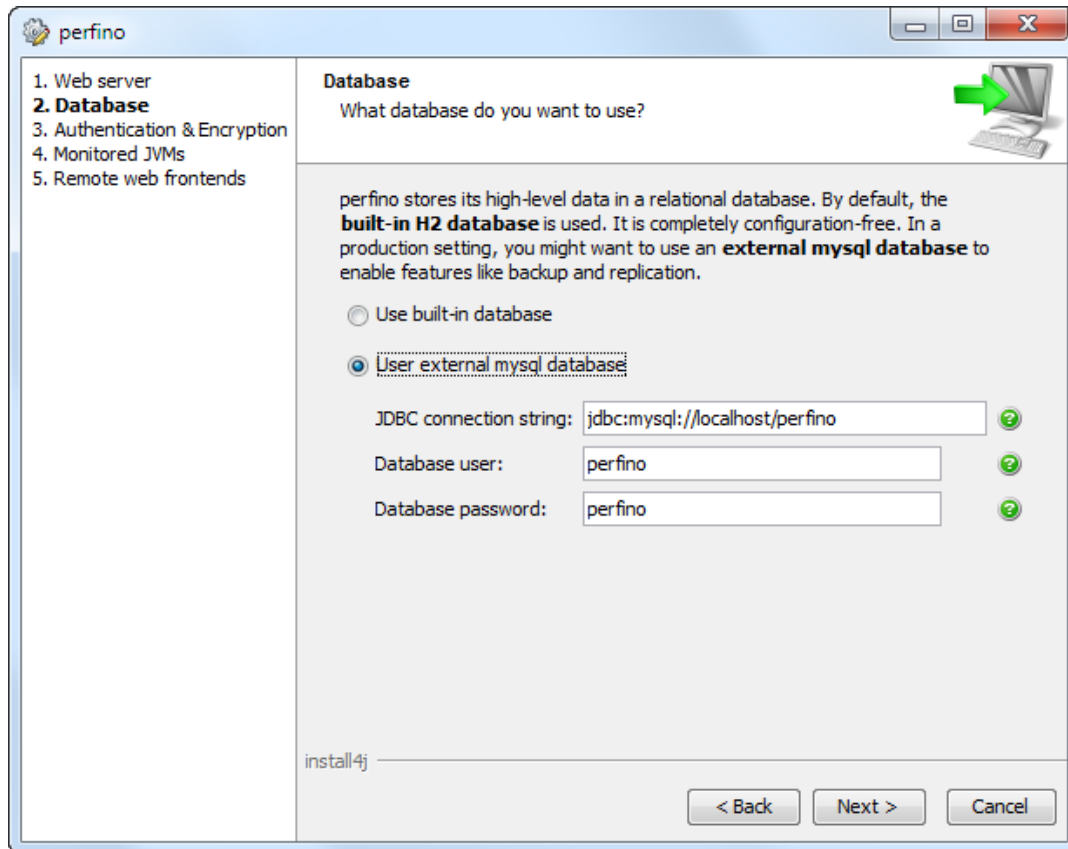


For evaluation purposes, this is all the installer needs to know. If you are deploying to a production environment, you might want to adjust other options, such as an external database.



## Changing the initial configuration later on

The installer saves the configuration to the file `perfino.properties` in the installation directory. To change the configuration at a later time, you can edit that file or run the `configure[.exe]` tool in the installation directory. The configure tool requests elevated privileges on Windows and Mac OS X, so that it can overwrite the configuration file and restart the server.



The GUI of the configure tool presents the configuration options in the same way as the installer if you choose to configure advanced settings. Just like the installer, you can run the configure tool in console mode by calling it with the `-c` argument:

```
configure -c
```

The various options are briefly documented in `perfino.properties` with the same explanations shown in the GUI when you hover the mouse over the question mark icons. For more information, see the help on server configuration [\[p. 83\]](#).

## Post-installation setup

After the installer has finished, the perfino server is running and you can open your browser at `http[s]://localhost:[port]`

to connect to the perfino UI. Note that if you have selected HTTPS for the web server, there is no port that serves HTTP requests.

Before you can use perfino, you have to complete the installation wizard in the perfino UI. The first step of the wizard asks you about the license key. Paste in the key that you have received by email.

With an **evaluation key**, a fixed number of VMs can be monitored for a limited period of time. If you want to evaluate perfino with more VMs, contact [sales@ej-technologies.com](mailto:sales@ej-technologies.com). After the expiration of the evaluation key, the perfino server will continue to run, but JVM monitoring is suspended.

For **permanent licenses**, there are two types of keys: A product key and a key for additional VMs. The product key allows one VM to be monitored. The additional VM keys allow more VMs to be monitored, but without a product key they will not work.



The second important step is the next one where you configure the **initial admin user**. You can configure other admin and non-admin users later on. The full name is displayed in the UI and the email can be entered so that an administrator can more easily identify or contact users.



perfino - Setup

Please complete the configuration of perfino.  
After the initial setup, this configuration is available in the "General Settings".

1. License Keys 2. User Account 3. Recording Options 4. SMTP Server 5. Finished

A user account with all privileges must be configured for working with perfino.

User Name \*

Password \*

Confirm Password \*

Full Name

Email

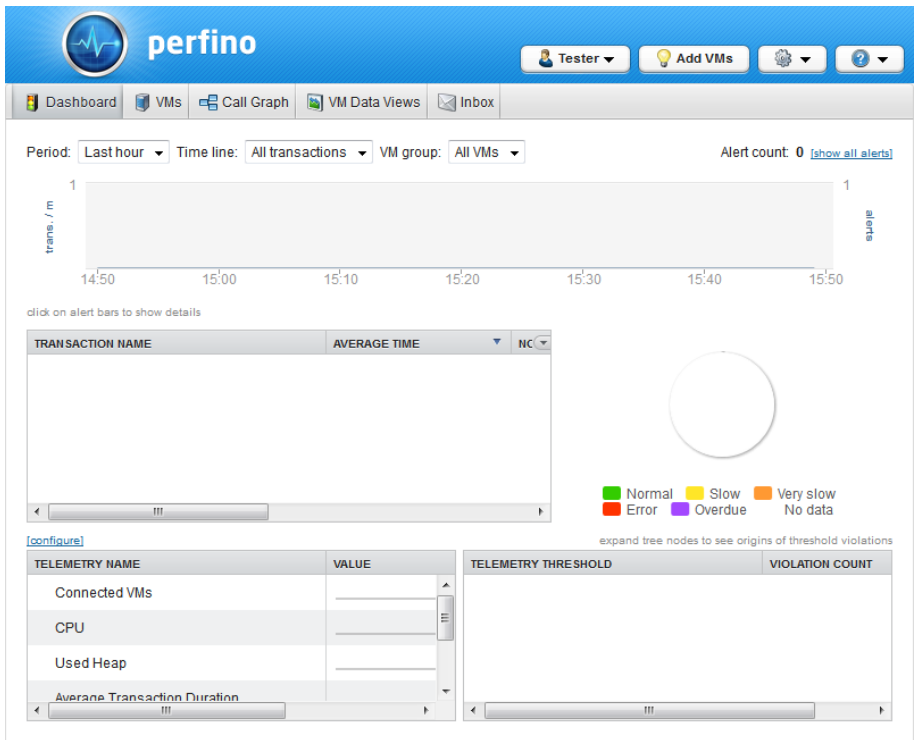
You can configure restricted users later on.

Back Next Finish

In the "Recording Options" step you already get a chance to configure **VM groups**. However, it is not necessary to do so at this point.

The server can only send emails if it has a valid **SMTP server configuration**, so in the next step it asks you to provide this information. If you do not intend to send emails, you can skip this step.

When you complete the wizard, the server is fully initialized and the installation wizard cannot be shown again. However, all settings in the installation wizard can also be adjusted in the perfino UI. You are now presented with the empty dashboard and can continue to set up monitored JVMs as explained in the next chapter.



## Monitoring JVMs

After perfino is installed, the agent can be added to JVMs on the local machine and on remote machines.

### Basic mechanism

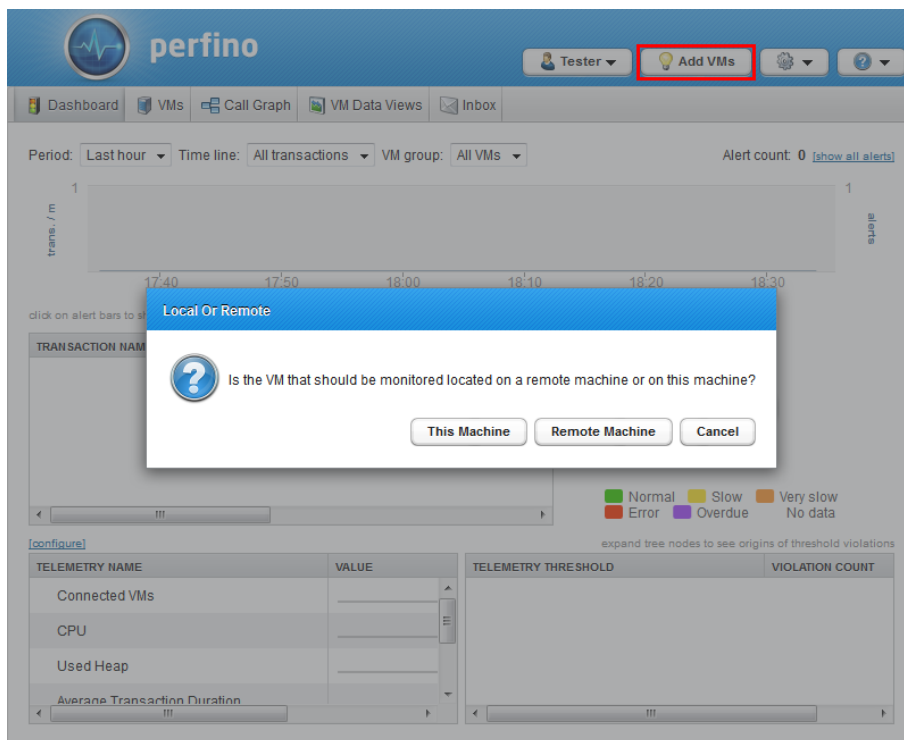
The perfino agent is a **Java agent**. It is loaded into a VM by specifying the `-javagent` VM parameter. Java agents are able to instrument classes as they are loaded and to retransform classes that have already been loaded.

To prepare a VM for monitoring, you need the perfino agent files and you have to know the actual VM parameter that needs to be inserted into the start script of your application server or your application. In a production deployment, the monitored VMs will usually be running on different machines. You do not have to install perfino on those machines, you just download the agent files from the perfino UI. Also, you will get instructions on how to construct the correct `-javaagent` VM parameter.

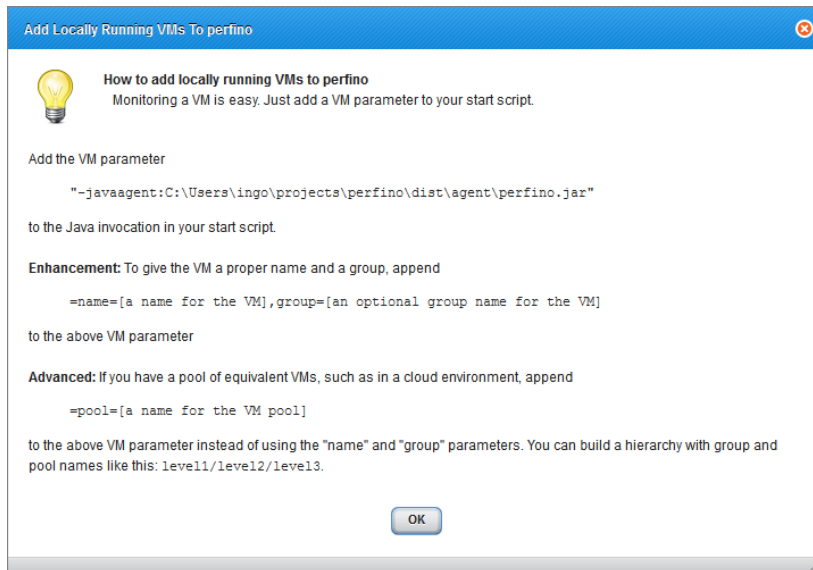
To get started, click on the "Add VMs" button in the header.

### perfino server and monitored VM running on the local machine

If the perfino UI server is running locally, the perfino agent is already available in the perfino installation directory and you can easily monitor a locally running VM. This is a likely scenario if you are evaluating perfino. In that case, you will be asked whether the VM is running on the local machine or on a remote machine.



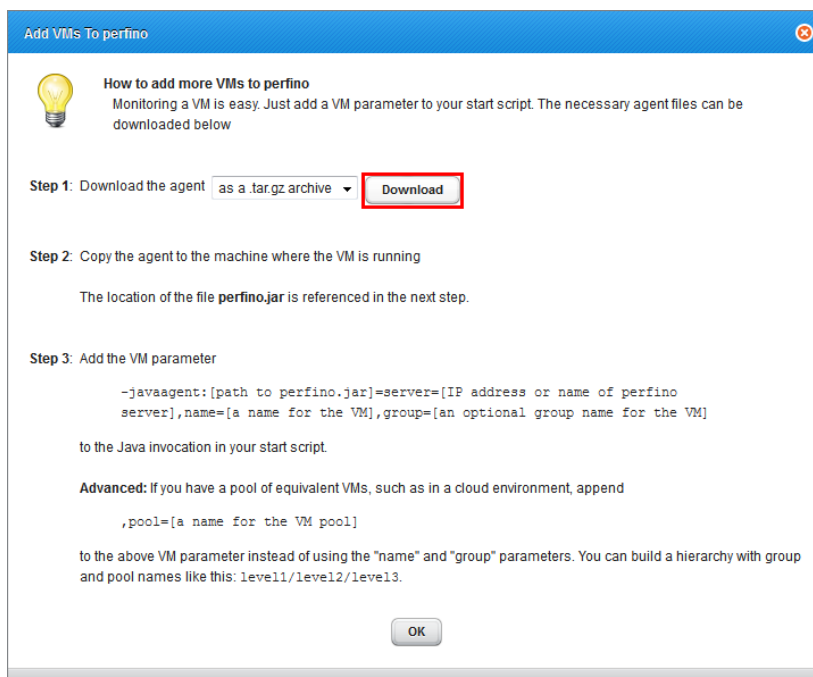
The "Add VMs" dialog displays the complete `-javaagent` VM parameter for the simplest case without any further configuration. It is suitable for monitoring a single VM.



Further instructions regarding group and pool names are explained below.

### perfino server and monitored VM running on the different machines

In the general case, where the perfino server and the monitored VMs are running on different machines, you can **download the agent files** from the dialog:



perfino has an integration with JProfiler [\[p. 77\]](#) that allows you to perform full sampling of a monitored VM and open the resulting snapshots in JProfiler. The libraries that implement full sampling are native libraries, and are included in the archive.

Depending on the selected option, a .zip or a .tar.gz file will be saved after clicking the **[Download]** button. Extract that archive anywhere on the computer where the monitored VM is running. In

the top-level directory of that archive there is a file named `perfino.jar`, whose path will be referenced in the `-javaagent` VM parameter.

The VM parameter is not specified completely in the "Add VMs" dialog, since it depends on where you extract the archive and the name or IP address of the perfino server:

```
-javaagent:[path to perfino.jar]=server=[IP or host name]
```


If you do not specify the server option, the agent will assume that the perfino server is running on localhost.

## Naming VMs

If you monitor multiple VMs, you have to give them names in order to be able to **identify them in the perfino UI**. To assign a VM name, pass the `name` option to the `-javaagent` VM parameter:

```
-javaagent:[path to perfino.jar]=server=[IP address or host name],name=[VM name]
```

Each name can only be used by one VM at the same time. A second VM that requests to be monitored with the same name will be rejected by the perfino server. If you have a pool of VMs that cannot be assigned with unique names, see the section on VM pools below.



In addition to the VM name, you can group VMs into a hierarchy. By default, a VM is inserted into the top-level group. To assign it to another group, set the `group` option in the `-javaagent` VM parameter:

```
-javaagent:[path to perfino.jar]=server=[IP or host name],name=[VM name],group=[group name]
```

In the group name, separate hierarchy levels with forward slashes, as in `Web/Workers/Gen3`. A complete `-javaagent` parameter looks like this:

```
-javaagent:/opt/perfino/perfino.jar=server=192.198.0.33,name=web,group=Web/Workers/Gen3
```

VM names must be unique within the same group.

## VM pools


Sometimes it is not possible to configure VMs with unique names, for example in a cloud environment where **instances are provisioned dynamically**. In that case, you can assign the VM to a VM pool:

```
-javaagent:/opt/perfino/perfino.jar=server=192.198.0.33,pool=Web/Workers
```

Like the `group` option, the `pool` option takes a hierarchical name with forward slashes as the hierarchy separator. If `pool` is specified, you cannot specify `name` or `group` and vice versa.

In a pool, a monitored VM is given a name with a unique identifier. When the VM detaches, for example because it is shut down, this identifier will never be used again. Unlike for named VMs,

the history of a single VM is limited to the connection time of that VM, so there is no associated long term history. However, you can follow the history of the entire pool to analyze trends and make historical comparisons.



### Server port

By default, a perfino server listens for VM connections on port **8847**. This is the port that needs to be **opened in firewalls** to allow monitored VMs to reach the perfino server.

A perfino server may be configured to use a different port by adjusting the "vmPort" property [\[p. 83\]](#) in `perfino.properties`. This is necessary if the port is already in use or if you install multiple instances of perfino on the same machine.

To tell the agent about such a non-default port, you have to add the `port` option to the `-javagent` VM parameter. For example:

```
-javagent:/opt/perfino/perfino.jar=server=192.198.0.33,port=8912,name=test
```

### Encryption and authentication

Mutual authentication and encryption are enabled by the agent keystore file `agent.ks` in the same directory as the `perfino.jar` file. The agent keystore is generated by the server and is located in the `ssl` directory below the perfino data directory.

When the "vmUseSsl" property in `perfino.properties` is set to `true`, the agent keystore file is automatically added to the agent files that you download from the perfino UI. In that case, encryption and authentication work out of the box.

If you switch encryption and authentication on or off after you have set up your monitored VMs, you have to make the following changes manually:

- **Switching encryption on**

Locate `[perfino data directory]/ssl/agent.ks` and copy it to the agent installations on all machines where VMs are monitored. The keystore file has to be copied to the same directory as the file `perfino.jar`. If the correct agent keystore is not present, the server will refuse the connection from the agent.

- **Switching encryption off**

Delete the file `agent.ks` next to the file `perfino.jar` in the agent installations on all machines where VMs are monitored. If that file is present and the server does not use encryption, the agent will refuse to connect to it.

If you would like to keep the agent keystore file in a different path, add the `keystore` option to the `-javagent` VM parameter. For example:

```
-javaagent:/opt/perfino/perfino.jar=server=192.198.0.33,keystore=/sec/agent.ks,name=test
```

## Basic Concepts

perfino collects data of two fundamentally different types: transactions and telemetries. Policies and thresholds are used to detect anomalous conditions while triggers take action if something is out of order.

### Transactions

In perfino, you analyze your business processes with transactions [\[p. 27\]](#). At a technical level, a transaction is simply a method invocation. To measure a transaction, perfino records its timing and constructs a transaction name that describes the business process.

The **transaction naming** has a significant impact on what you will see in the perfino UI.

- It enables you to understand what triggered the transaction.
- It groups all business processes with the same transaction name and so determines the granularity that is used to measure business processes.
- It can serve as a basis to filter out unwanted operations.

perfino cannot know what your business processes are, so configuring transactions is an important part in setting up an application for monitoring. Some frameworks are high-level by nature and so perfino can offer them as transaction types that can be configured with a minimal amount of work.


The most common example of a transaction is the invocation of a URL that is handled by your application server. In the default configuration, perfino intercepts the method that handles HTTP calls and constructs a transaction name including the first three segments of the URL. This is an arbitrary naming strategy that is just intended to get started. In your application, only the first segment of the URL may be relevant for the business process, or you may need a particular query parameter in the name.

Also, you will probably not want all URL invocations to become transactions. Many HTTP requests are for static resources, and those are not interesting in terms of business processes. In perfino, you can **discard transactions** based on the name that would be associated with a transaction.

The following figure shows how different URLs end up as the same transaction based on a transaction naming that

- adds the value of the query parameter "action"
- adds the fixed text "in shop"
- adds the second segment of the URL





## Policies

Transactions have associated policies. The policies determine

- the acceptable timing for a transaction
- the way errors are detected and handled
- when to perform method-level sampling

For each violated condition in the policies, you can see transaction details separately in the perfino UI. For example, you can inspect slow transactions or transactions that resulted in an error separately and not cumulated with other regular transactions of the same name.

perfino gets information from the monitored application by instrumenting methods. To keep the overhead low, very few methods are instrumented. In order to get more detailed information in the case of a very slow transaction, the policy can start method level sampling [\[p. 45\]](#) for a transaction once it is clear that it is taking too long.

With sampling, you get a cumulated call tree and hot spots on the method level that show you where the time is actually spent.

## Telemetries

The other fundamental type of data source in perfino is the periodic sampling of scalar values, like heap size or thread count. Each telemetry [\[p. 51\]](#) can be plotted as a time-resolved graph. In perfino, telemetries are often shown as sparklines, without defined axes and with a trailing current value.


There are many standard telemetries in perfino that collect their data from well-known subsystems of the JVM or popular databases and frameworks. In addition, you can use the `@Telemetry` annotation to define custom telemetries on static methods with a numeric return value.

## Thresholds

You will have different expectations with respect to different telemetries. For example, the heap usage often oscillates around a base line and where a steady increase is a sign of a bug in the application.

Or, the average duration of JDBC statements usually varies with server load and is an indicator of the health of the application.

To detect anomalous conditions, you define thresholds [\[p. 55\]](#) with an optional lower and an optional upper bound. Threshold violations are counted on a per-VM basis or for each VM group. They do not have actions associated with them. Often, you will not want to take any action for single threshold violations, but only for a cascade of such conditions.



## Triggers and alerts

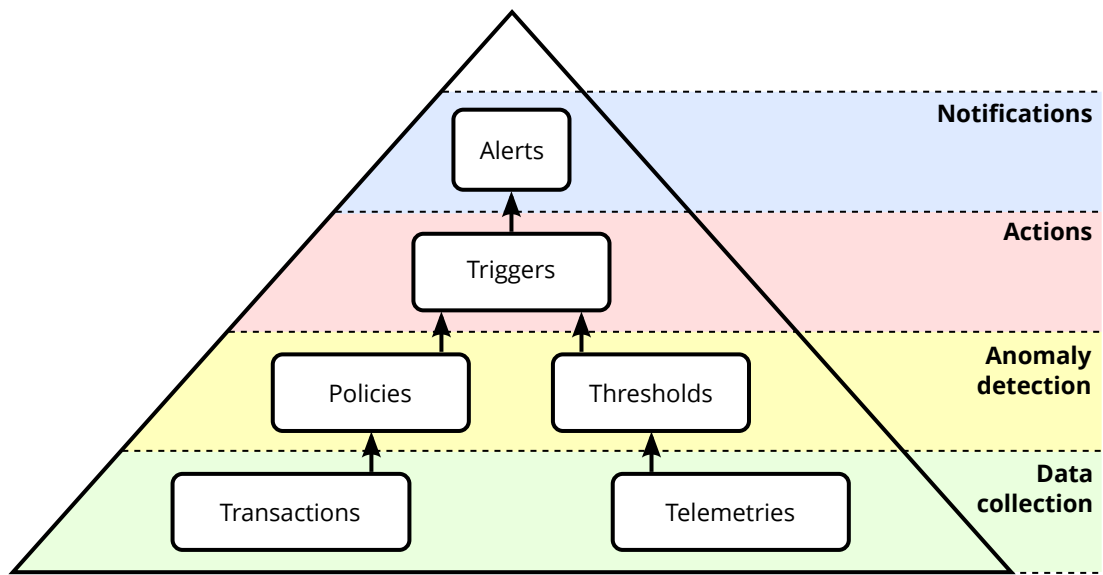
Both transactions and telemetries can lead to anomalous conditions: A transaction policy can identify a slow transaction and a telemetry threshold can be violated.

In order to take action on these conditions, you use triggers [\[p. 58\]](#). Triggers do not operate on a per-VM level, they process all recursively contained VMs in a VM group. Each VM group in a hierarchy of groups has its separate triggers.

For example, you could define a trigger for all VMs that fires when the number of connected VMs falls below 20. In the same VM hierarchy, you might have a group that only contains database VMs. In that group, you might want a separate trigger that fires when the number of connected database VMs falls below 3.

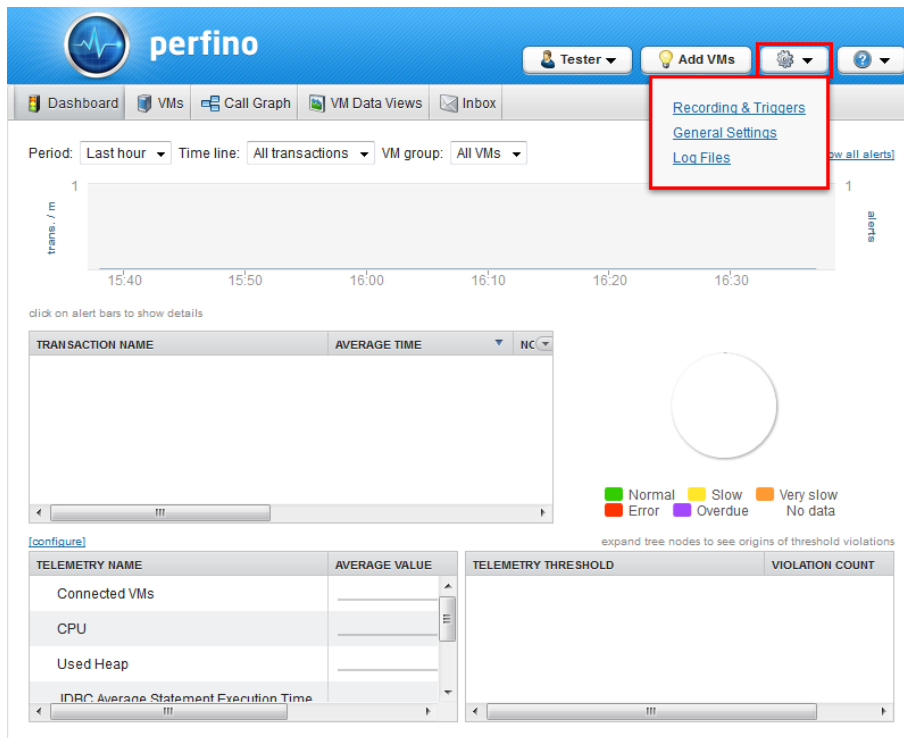
When a trigger fires, it executes its actions. Actions can start data collection, such as full VM sampling, send emails or create alerts.

Alerts [\[p. 63\]](#) are shown in the dashboard and are the highest level in perfino's pyramid of concepts:




# UI

The perfino UI separates data and configuration. When you log in, you are in the **data perspective**. The configuration drop-down in the top right corner gives you access to the **configuration**.



## The data perspective

The data perspective is divided into several tabs. The tabs are ordered from more general information on the left to more specific information on the right.

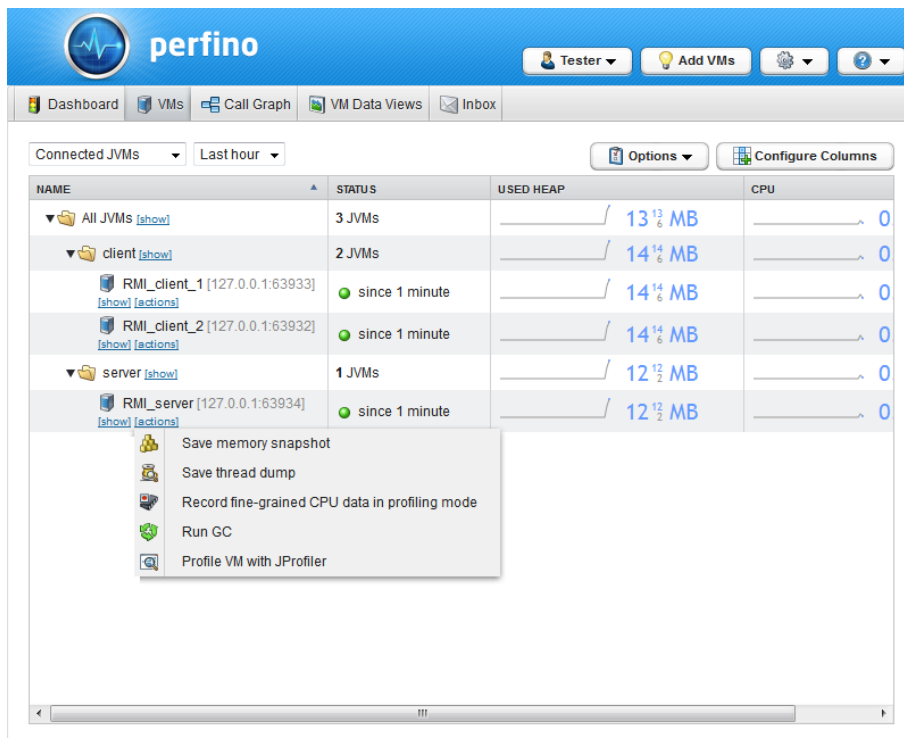


After logging in, you see the **dashboard** where important information about the activity and the health of the monitored VMs is presented on a single screen.

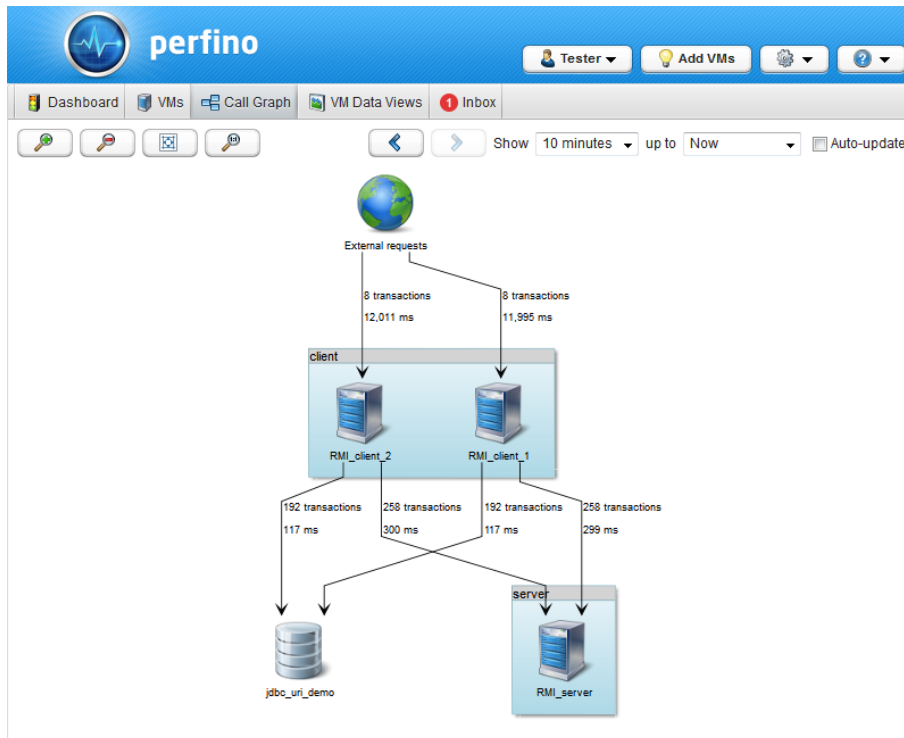


The **VMs** view shows you the hierarchy of all connected VMs, together with selected information in the form of sparklines.

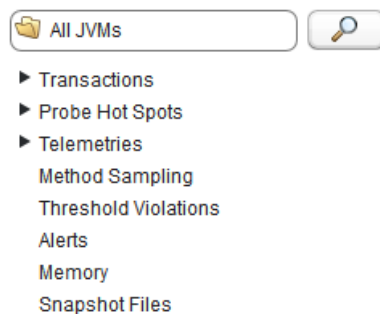
The "show" context action takes you to the "VM data views" for the selected group or VM and provides quick access to all recorded data that is available for your selection.



The **call graph** focuses on remote calls between monitored VMs. Each VM is a node and remote calls like EJB, RMI and web service invocations are edges in the graph. Clicking on nodes and edges shows further information.



The **VM data views** tab holds all views that show data for a single selected VM or a group of VMs. The actual VM selection and the desired view can be adjusted in the view selector.



The **inbox** view shows all messages that have been sent to you. For example, if you start full-VM sampling on a particular VM, a notice will be sent to your inbox once the snapshot is ready for downloading.

If there are unread messages in your inbox, you will see a red notification in the tab area.

perfino

Tester Add VMs

Dashboard VMs Call Graph VM Data Views **Inbox**

[Download selected file](#)

DATE	TYPE	NAME
08-May-2014 15:59:26	Memory snapshot	RMI_server


The inbox notifies you about new messages and snapshots. If you delete a snapshot from the inbox, it is still available in the data views. In the "Snapshot Files" view you can permanently delete snapshots.

[Show all snapshots in the "Data Views" section](#)

## Configuration

Configuration is divided into two parts: **general settings** and **recording options**. If you want to view and edit general settings, you have to be an administrator. Recording options can also be edited by the "profiler", but not by the "viewer" access level.

Both general settings, as well as recording options are modal panels. If you did not change anything, you have to click on the "Go Back" button to return to the data perspective.




If you made a change, the "Go Back" button is replaced by **[Apply Changes]** and **[Discard Changes]** buttons. No changes to the configuration have any effect before this confirmation.



## General Settings

In the general settings you administrate the perfino server. This includes

- the users and their access levels
- the license keys
- data consolidation options
- SMTP access



In addition, you can export and import the entire server configuration.

## Recording & triggers

Recording options and triggers are configured for each VM group. You can click the edit button or double-click on any column to jump directly to the desired step.

perfino

Tester Add VMs

Recording Options - Please click on "Go Back" to return to the data view

Go Back

GROUP	TRANSACTION	SAMPLING	OPTIONS	THRESHOLDS	TRIGGERS
All JVMs	✓	✓	✓	0	0
▼ Demo	✓	✓	✓	2	2
Web				0	0
▼ Workers				0	1
JDBC				0	0
JMS				0	0


Use drag and drop to reorder group configurations

Group configurations inherit recording settings from their parent group unless settings are overridden. Triggers and thresholds operate on all recursively contained VMs and are not overridden.

The All JVMs group is for JVMs at the top-level. New groups can be added here or are automatically added when a JVM connects with a specified group.

The configuration contains recording options as well as threshold and trigger settings. Recording options are sent to the monitoring agents in each connected VM. They are inherited to nested VM groups and nested VM groups can override settings in their parent groups.





**Create Group Configuration**

Please enter the details of the new group configuration.

1. Group Settings      2. ...

Group name \*

The group name can contain alphanumeric characters and spaces, but no special characters. The group will be added at the selected level.

Group type  
☒ VM Group  
☐ VM Pool

VM pools allow you to connect any number of exchangeable VMs with the exact same configuration.

Override settings for the following categories:

☐ Transactions  
☐ Method Sampling  
☐ Options

Thresholds and triggers operate on data from all recursively contained VMs in this group. These settings are not overridden.

Double click on a table column in order to go directly to the desired step.

Cancel Back Next Finish

Thresholds and triggers, on the other hand are handled in the perfino server and are processed for the VM group for which they have been defined. They operate on all VMs that are contained recursively in their VM group. If you need additional thresholds or triggers that operate only on a nested VM group, you can define them on that group.


## User settings

The user drop-down in the header shows selected information about your account and the current session and contains the **[Logout]** button that terminates your perfino session.



User settings, like your name and password can be adjusted by following the "Account Settings" link. If your access level is "profiler", this is where you can see the VM groups that you are allowed to modify.

Account Settings

 Edit your account settings.

Full Name

Email

☐ Change password

Old Password

Password

Confirm Password

# Transactions

## Configuration

Transactions are configured in the recording settings. When you edit a VM group or a VM pool, the first displayed step in the wizard shows the transactions settings.

**Edit Group Configuration**

Please modify the details of the existing group configuration.

1. Transactions   2. Method Sampling   3. Options   4. Thresholds   5. Triggers

BUSINESS TRANSACTION	NAMING	POLICIES
▼ Web requests	✓	✓
▼ EJB invocations	✓	✓
▼ Spring service invocations	✓	✓
▼ Annotated invocations		
▼ DevOps annotated invocations		✓
▼ Pojo invocations		
▼ RMI calls	✓	✓

Use drag and drop to reorder business transactions. Higher items are matched first. Policies and transaction naming fall through to the first defined matching item.

Cancel   Back   Next   Finish

If no transactions are defined for a nested VM group, VMs get their configuration from the nearest ancestor group where transactions are defined. The "All VMs" configuration always has associated transaction definitions and is always a successful last stop for this search in the hierarchy.

When a nested group overrides transaction settings, it overrides all of them and there is no merging with the settings of a parent group. If you want to reuse selected definitions from a parent group, save them to a set and include them in the configuration of the nested group.

**Save Business Transaction Set**

Please enter a name for the business transaction set.  
You can add saved business transaction sets to other JVM group configs.

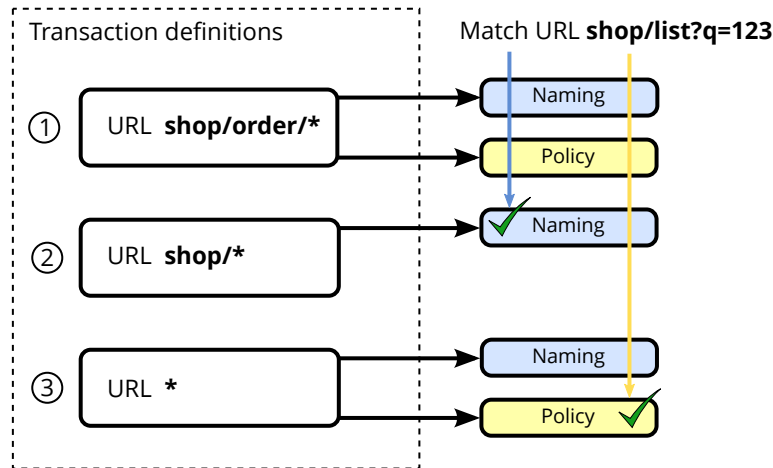
Name \*

Saved business transaction sets

BUSINESS TRANSACTION SET

Save   Cancel

A transaction definition selects a segment of all possible transactions in its subsystem. For example, web transactions are defined with a URL wildcard or regular expression. Each transaction definition has an optional associated naming scheme and optional associated policies. For both naming and policies, perfino keeps looking in the list of defined transactions until it finds a matching entry.



This allows you to define a series of transaction definitions that all have different naming schemes while keeping common policy definitions for a group of transaction definitions. Keep in mind that the transaction matching in perfino takes the first match, so more generic filters should be lower in the list of transaction definitions.

There are a number different transaction types in perfino. All types fall into two different categories: **predefined** transaction types and **custom** transaction types.

### Predefined transaction types

The most common mechanism that starts a business transaction is via a URL invocation. In perfino, this is called a **web transaction**. In the default configuration, perfino intercepts all URLs and creates transactions for them. For your application, you will want to limit the intercepted URLs to those that are meaningful from a business perspective. To do that, you can remove the transaction definition for the URL wildcard "\*" and add your own specific transaction definitions.

Another way to limit transaction matching is to use the "Discard" flag. In that case, any matching URLs will be excluded and no further matching will take place.

In the list of transaction definitions, the filter expression will be shown as the name of the transaction definition. Often, this is not descriptive enough. Use the "Custom description" check box shown above to enter a meaningful name. This is only used in the configuration and not in the displayed data where the transaction naming determines the name of the transaction.

The name of a transaction is composed of a chain of naming elements. The list of available naming elements depends on the transaction type. For web transactions you can use specific elements like "URL query parameter" and "URL segments".

An important consideration when defining complex transactions is what kind of nested transactions should be allowed. Perfino does not record a directly nested transaction that has the same name. More restrictive options are to prevent nested transactions

- **that match this entry**

No nested transactions will be created if the nested transaction would originate from this transaction definition.

- **with the same group name**

If you select this option, you have to configure a group name for this transaction definition. For example, you might have different entry points for a particular business transaction from a web transaction and from an EJB transaction and the URL handler always calls the EJB. If you do not want to see the EJB as a nested transaction, you can assign the same group name to the web transaction and the EJB transaction.

- **of the same transaction type**

This option is suitable if you only want to see entry points and not the internal structure. For example, EJBs will often call other EJBs. This reentry mode prevents a tree of nested transactions in that case.

- **all further entries**

Use this option to prevent any kind of nested transaction if a transaction is created from the current transaction definition.

The reentry inhibition only applies to the transactions that would be directly nested. If a nested transaction is allowed, its own reentry inhibition setting will be used for the next nesting level.

**EJB transactions** are created from calls into public methods of an EJB that is annotated with `@Stateful`, `@Stateless`, `@MessageDriven` or `@Singleton`. **Spring service invocations** are for calls into beans that are annotated with `@Component`, `@Controller`, `@Repository` or

@Service. You can limit transaction recording to a sub-set of those transactions for both EJB and spring transactions.

When creating transaction definitions, both these transaction types can be filtered with respect to class or package name. The discard mechanism is useful to exclude certain classes that are not generating significant data. By default, perfino shows all EJB and Spring service invocations.

Create EJB Invocation

Please enter the details of the new business transaction.  
Monitor EJBs of selected classes.

1. Filter 2. Naming 3. Policies

Class name  Choose Wildcard comparison

Custom description ☐

☐ Discard transactions

EJB types:

- ☒ Stateless EJBs
- ☒ Stateful EJBs
- ☒ Singleton EJBs
- ☒ Message-driven EJBs

**RMI transactions** handle incoming RMI calls into a VM. Like for EJB and Spring transactions, you can use a class filter to include or exclude implementation classes. The default configuration in perfino shows all RMI invocations.

### Custom transaction types

There are many frameworks where specific interceptions can capture all important business transactions. Also, your own code will often be structured in a such a way that a small set of methods will map to your business transactions. perfino offers three mechanism for monitoring these cases. Configuration of these transaction types is more elaborate so that each of them has its own chapter in the advanced topics.

For an **annotated invocation** transaction, you specify a particular annotation class name. In the further configuration [\[p. 93\]](#) you can decide whether the annotation should apply to classes or methods and how inherited classes should be handled.

**DevOps** transactions are defined in your code with annotations supplied by perfino. While the naming is completely specified with those annotations, the policy must be configured in the perfino UI. This is the most maintainable way to define complex transactions. See the advanced topic about DevOps transactions [\[p. 98\]](#) on how to add them to your project.

If you cannot modify the code that should be monitored, **POJO transactions** allow you to specify all the same interceptions as DevOps transactions directly in the perfino UI. See the detailed explanation of POJO transactions [\[p. 95\]](#) for more information.

### Call tree and hot spots

perfino builds a call tree from all recorded transactions. A call tree is a cumulated data structure that captures all the different sequences in which transactions are nested. If transaction C is called within transaction A and also within transaction B, it will occur twice in the call tree, once

as a child of A and once as a child of B. Each of these sequences can occur many times, increasing the invocation counts along the corresponding paths in the call tree.

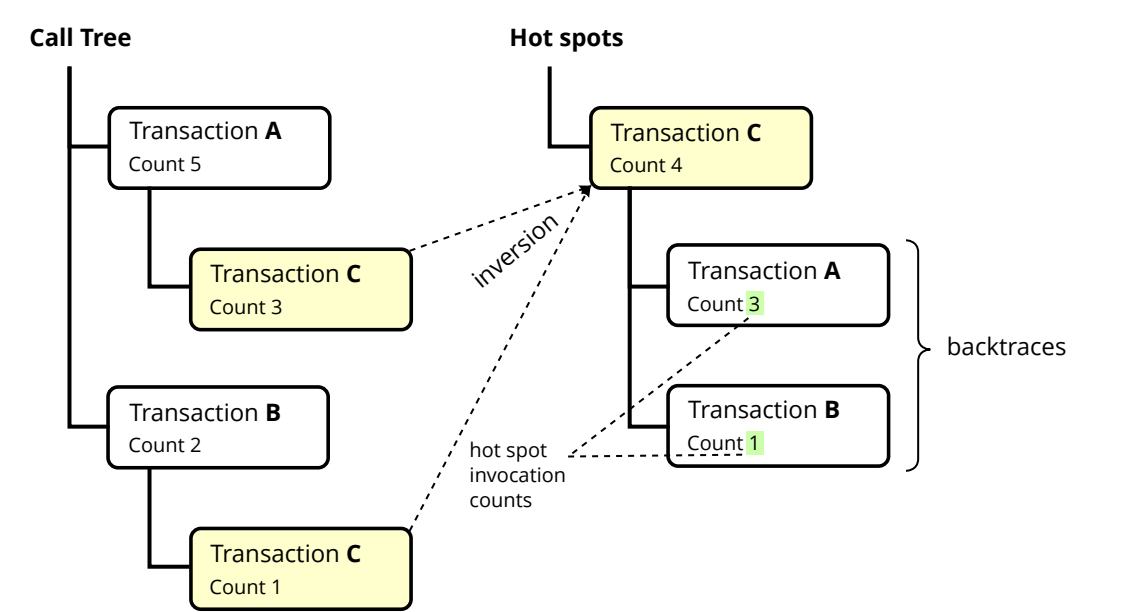
Policies: Split Remote origins: Merge Show 10 minutes up to Now Auto-update Compare

TRANSACTION	TOTAL TIME	INVOCATIONS	AVG. TIME
RmiHandler.remoteOperation [method samples] [time lines]	4,421 s	3,573	1,237 ms
demoView1 [method samples] [time lines]	1,149 s	221	5,200 ms
demoView5 [method samples] [time lines]	1,131 s	204	5,548 ms
Nested demo transaction [method samples]	1,129 s	204	5,539 ms
JPA/Hibernate	153 s	204	751 ms
JDBC	69,237 ms	850	81,456 µs
JNDI	14,957 ms	204	73,321 µs
Inventory checks via RMI [show remote call] [method samples]	769 s	204	3,772 ms
Exchange rate checks via web service [show remote call] [method samples]	70,731 ms	203	348 ms
[Very slow] Exchange rate checks via web service [show remote call] [method samples]	774 ms	1	774 ms
demoView3 [method samples] [time lines]	1,127 s	213	5,293 ms
demoView2 [method samples] [time lines]	1,101 s	205	5,370 ms

Filter: All Apply Clear Options

The numbers for each transaction - total time, invocation count and average time are displayed for the selected period. If the period is not fully measured, perfino will tell you the covered percentage at the bottom of the view. For example, when starting up the perfino server, the current hourly interval will not be fully measured until a full hour has passed. Similarly, if you stop the perfino server for some time, this will leave holes in the history with incomplete intervals at both sides.

Hot spots are the inversion of the call tree, where all occurrences of each transaction are summed and shown with their **backtraces**. The invocation counts and execution times in backtraces do not refer to the transactions, but rather to the number of times the hot spot was called along this path.



To emphasize the point about backtraces, all hot spots have a child node called "Cumulated backtraces". Searching for transactions is done with the filter bar at the bottom of the view. Many views in perfino have such a filter bar.

Policies: Split
Show 10 minutes up to Now
Auto-update
Compare

HOT SPOT	TOTAL TIME	INVOCATIONS	AVG. TIME
Inventory checks via RMI [method samples] [time lines]	3,928 s	1,061	3,702 ms
Remote demo transaction [method samples] [time lines]	2,792 s	3,188	876 ms
Exchange rate checks via web service [method samples] [time lines]	1,509 s	4,249	355 ms
Nested demo transaction [method samples] [time lines]	1,158 s	917	1,263 ms
Cumulated backtraces	1,158 s	917	1,263 ms
demo/view4 [method samples]	240 s	182	1,323 ms
demo/view3 [method samples]	240 s	193	1,247 ms
demo/view1 [method samples]	238 s	203	1,176 ms
demo/view5 [method samples]	224 s	161	1,392 ms
demo/view2 [method samples]	214 s	178	1,202 ms
Exchange rate to EUR [method samples] [time lines]	950 s	12,673	74,993 µs
Cumulated backtraces	950 s	12,673	74,993 µs
Called from group Demo/Workers [show remote origin]	709 s	9,488	74,799 µs

Filter:
Apply Clear Options



# Policies

## Configuration

Policies detect anomalous conditions with respect to transactions. Each transaction definition [p. 27] has associated policy definitions.

Create Web Request

Please enter the details of the new business transaction.  
Monitor web requests that come from outside the JVM

1. Filter 2. Naming 3. Policies

☒ Define policies

Slow events 500 \* % slower than average

Very slow events 1000 \* % slower than average

Overdue events 500000 \* ms and slower

☒ Split transaction tree for policy violations

Transaction errors ☒ Error throwables ☒ Runtime exceptions ☐ Checked exceptions

☒ Logged errors ☐ Logged warnings

Ignored HTTP error codes

Sampling for policy violations Very slow and more severe

Periodic sampling every 3 \* Hours

End user experience monitoring 20 \* % of all web requests

Cancel Back Next Finish

Policies create the following anomalous transaction states:

- **Slow**

The threshold after which a transaction is classified as slow can be specified as an absolute time or as a relative percentage value with respect to the average duration of transactions that are generated by this transaction definition.

- **Very slow**

Another time-based threshold like "Slow", only for a more severe category. You can partition "Slow" and "Very slow" so that "Slow" transactions can occur without notifications and "Very slow" transactions trigger notifications.

- **Overdue**

After a third time-based threshold which is usually an absolute time and much larger than the "Very slow" threshold, you can view a transaction as deadlocked or hanging.

- **Error**

Transactions that result in an error require special attention. There are three sources of errors:

1. **Exceptions**

Exceptions derived from `java.lang.Error` (errors), `java.lang.RuntimeException` (runtime exceptions) and `java.lang.Exception` (checked exceptions) can be activated

separately. Checked exceptions are usually handled by the application and do not necessarily result in a transaction error.

## 2. Logged errors

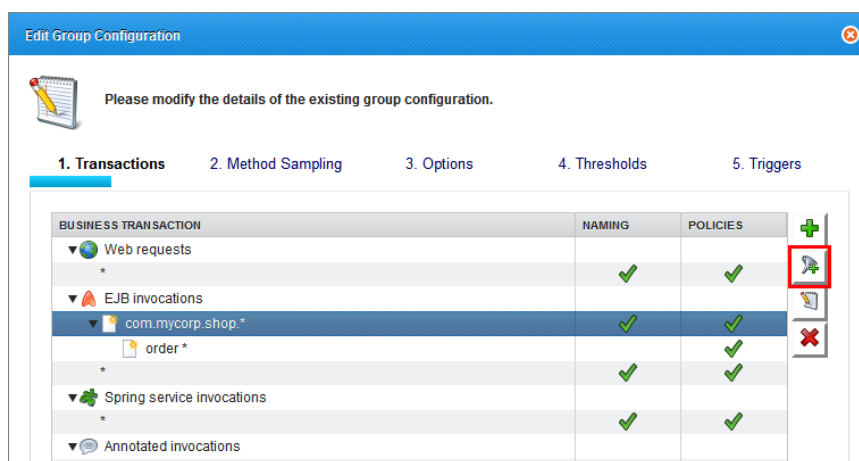
If your application logs an error with a logging framework, perfino can mark the transaction with an error. `java.util.logging`, `log4j`, `JBoss logging` and `Logback` are supported by perfino. Optionally, you can also configure that warnings result in a transaction error.

## 3. Specific conditions

Certain transaction types have their own built-in error conditions. For example, as shown in the screen shot above, web transactions can use the HTTP status code to detect an error. By default, all status codes  $\geq 400$  are handled as errors. If some of these status codes should not be viewed as errors, you can list them here.

If policies are not defined for a transaction definition, perfino continues in the list of transaction definitions until it finds a matching entry where policy definitions are enabled. This allows you to define the same policies for multiple transaction definitions, in the most extreme case with a "catch-all" entry that has no naming definition, but only policy definitions.

Sometimes, you want to go the other way and define more granular policies for certain transactions within a single transaction definition. For example, you might have a single EJB transaction definition which captures all your business transactions. If the acceptable durations depend on the transaction name, you can define **policy specializations** to avoid having to re-add the transaction definition and its naming scheme multiple times.



Unlike the parent transaction definition, which filters classes in this case, the policy specialization filters on the result of the transaction naming.

**Create Policy Specialization**

Please enter the details of the new policy specialization.  
Policy specializations allow you to discard selected transactions based on name patterns or apply different policy settings to them.

1. Filter 2. Override policies

Match transaction names  Wildcard comparison ☒ Comma-separated

Custom description \*

☐ Discard transactions

## Policies in the call tree

Both call tree and hot spots views have a policy mode drop-down at the top. By default, the call tree is **split for different policy violations**. For example, this means that you can see slow transactions separately from normal transactions. This is important, since the cause of a slow transaction will often be visible from nested transactions or a slow database statement.

Policies: **Split** Remote origins: Merge Show 10 minutes up to Now Auto-update Compare

TRANSACTION	TOTAL TIME	INVOCATIONS	AVG. TIME
RmiHandler.remoteOperation [method samples] [time lines]	4,531 s	3,660	1,238 ms
demo/View3 [method samples] [time lines]	1,251 s	234	5,349 ms
demo/View4 [method samples] [time lines]	1,154 s	215	5,368 ms
demo/View5 [method samples] [time lines]	1,148 s	211	5,441 ms
demo/View1 [method samples] [time lines]	1,141 s	216	5,285 ms
Exchange rate to EUR [method samples] [time lines]	1,092 s	14,555	75,045 µs
demo/View2 [method samples] [time lines]	980 s	186	5,273 ms
Demo JMS message [method samples] [time lines]	669 s	160	4,186 ms
<b>Very slow</b> Demo JDBC Job [method samples] [time lines]	433 s	6	72,333 ms
Demo JDBC Job [method samples] [time lines]	381 s	158	2,415 ms
processOrder [method samples] [time lines]	4,700 ms	176	538 ms
docs/invoiceService [method samples] [time lines]	93,469 ms	163	573 ms
docs/quoteService [method samples] [time lines]	7,412 ms	170	436 ms

Filter: **All** Apply Clear Options

When you set the policy mode to "merged", all policy splits are added so that there is only one transaction of the same name and type on the same call tree level.


In the screen shot above, you can also see the policy selector at the bottom the view. By default it is set to "All". If you choose a particular policy violation type, the tree will be filtered immediately. For example, selecting "Error" will show all transactions that have been marked as an error, regardless how deep they are in the call tree. These transactions will be expanded and all ancestor nodes are shown, even if they are not transactions with errors. When you also add a filter text, only transactions with errors will be searched for the filter expression.

## Policies in the dashboard

The dashboard is your first stop for checking the health of your application. That's why policy violations are featured prominently in the dashboard. The transaction table has a column for each policy violation type and you can sort the table by clicking on the column headers. Transactions are not shown in the form of a call tree, but completely flat. This means that it also shows transactions that are not entry points but only invoked as nested transactions. For detailed analysis, go to the data views.

TRANSACTION NAME	AVERAGE TIME	NORMAL	SLOW	VERY SLOW	OVERDUE	ERRORS
demo/view5	5,524 ms	1,243	0	0	2	0
demo/view4	5,434 ms	1,306	0	0	4	0
demo/view3	5,378 ms	1,283	0	0	2	0
Nested demo transaction	5,370 ms	6,269	0	0	0	0
demo/view1	5,288 ms	1,232	0	0	3	0
demo/view2	5,273 ms	1,205	0	0	2	0
Demo JDBC Job	5,039 ms	922	0	33	0	0
Demo JMS message	4,224 ms	971	0	0	0	0
Inventory checks via RMI	3,734 ms	7,240	0	0	0	0
RmiHandler.remoteOperation	1,243 ms	21,736	0	0	0	0
Remote demo transaction	1,243 ms	21,457	278	1	0	0
docs/invoiceService	786 ms	926	296	28	0	0

To visualize the proportion of policy violations with respect to the total amount of transactions, a policy pie chart is shown next to the transaction table. Numbers and percentages are shown as tool tips and by clicking on the legend entries, you can hide selected policy violation types.



The drop-downs at the top provide three ways to adjust the range of the displayed transactions in the dashboard:

- **Period**

By default, the dashboard shows the last running hour. Alternatively, you can switch to the last running day.

- **Request type**

By default, all transactions are shown. Alternatively, you can restrict the display to web transactions only.

- **VM group**

By default, all monitored VMs are shown. With the "VM group" drop-down, you can restrict the displayed data to a particular VM group and its nested groups. For example, if you select group "A/B", then all VMs that are contained in "A/B" are shown, but also VMs in "A/B/C".

## Cross-VM Monitoring

One of perfino's core features is the ability to track calls between monitored VMs. No configuration is necessary to enable cross-VM monitoring. As soon as an outgoing and incoming call of a supported subsystem can be matched between two monitored VMs, perfino shows it as a remote call in the call graph and in the call tree.

The following remote call mechanisms are supported:

- RMI
- Web services: JAX-WS-RI, Apache Axis2 and Apache CXF
- Remote EJB calls: JBoss 7.1+ and Weblogic 11+

### Call graph

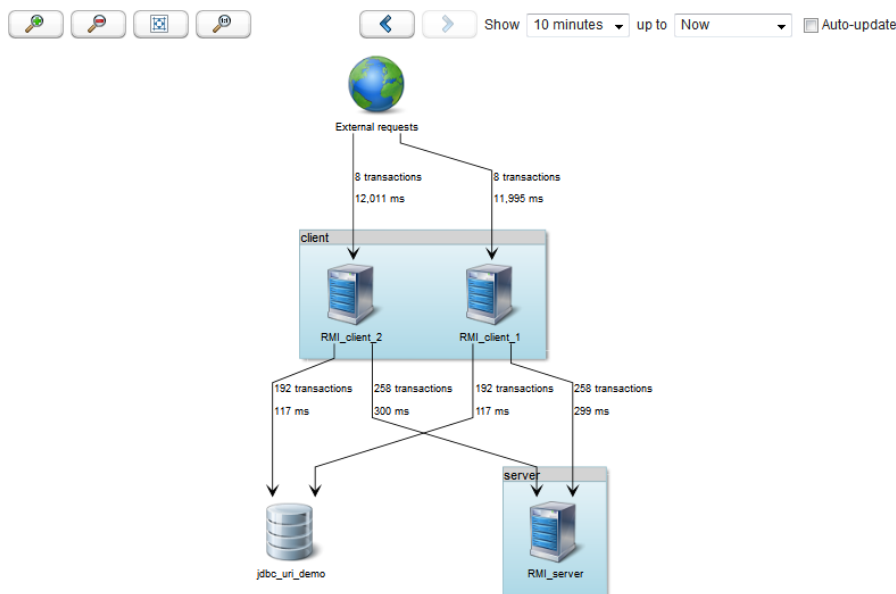
The natural representation for remote calls is the call graph, where each pair of VMs that participate in a remote call are connected by an edge. In addition, two other node types are present in the call graph that usually represent calls to and from external processes:

- **External calls**

The "world" node shows external calls that create transactions in monitored VMs, such as URL invocations.

- **Databases**

Database nodes show database operations. If perfino detects that two VMs are using the same physical database, it only shows a single node in the call graph.



When a node or an edge is selected in the call graph, a **detail panel** is shown that holds several tabs with information about the associated transactions. The number and names of the tabs depend on the selected object:

- **External calls node**

The only tab is the "Transactions" tab which shows the call tree of all transactions with external origin. Transactions are only shown in the first neighbor nodes, not across the whole graph.

- **Database node**

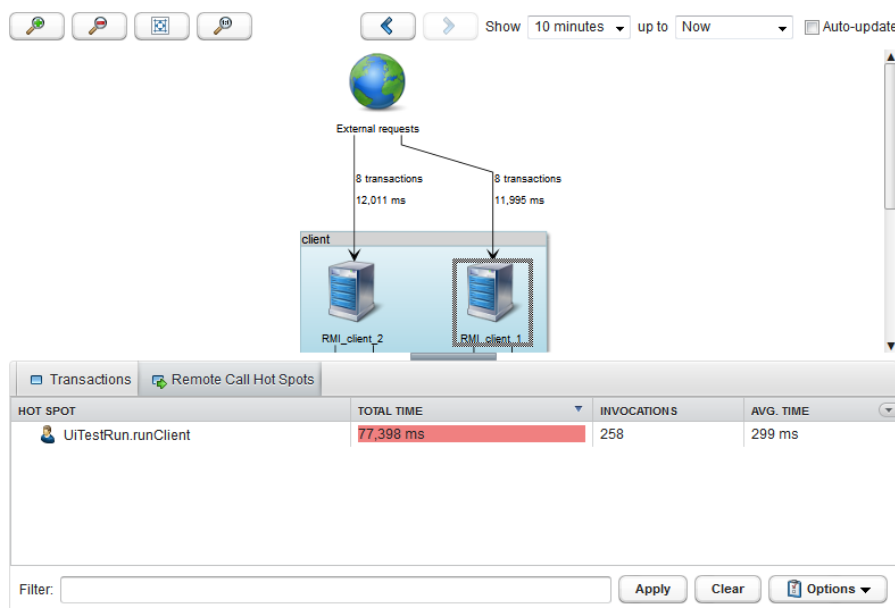
The only tab tab is the "Remote origin hot spots" tab which shows the database operations with cumulated backtraces. To restrict displayed data to a particular VM, select an incoming edge instead.

- **VM node**

The "Remote origin hot spots" tab shows the duration and count of handled remote calls mapped to the originating transactions from other monitored VMs. Backtraces are resolved up to the leafs that specify the VMs where the calls originated. The "Transactions" tab shows the call tree of all transactions executed in the selected VM. Finally, the "Remote call hot spots" tab shows the duration and count of remote calls originating in the selected VM. The top-level hot spots are the transactions that cause the jump into a different monitored VM. Check the outgoing edges to learn more about those calls.

- **Edge**

The "Call site hot spots" tab shows which transactions in the source node caused remote calls along the selected edge together with cumulated backtraces. For edges into "Database" nodes, this is the only tab and shows database operations instead of transactions. The "Execution sites" tab shows the call tree of the transactions that were caused by the selected edge. For edges starting at the "External calls" node, this is the only tab.



## Call tree

In the call tree, you can recognize transactions that cause a remote call by the "show remote call" link that is shown next to the transaction name. When following such a link, it is advisable to select the "Merge" option for the policy split [\[p. 33\]](#), otherwise the numbers on the remote side may not match the numbers in the source VM. This is because policy splitting is not tracked across VMs. In any case, there may be small discrepancies due to calls that are close to period boundaries which may be assigned differently by source and target VMs.

Policies: **Merge** Remote origins: Merge Show 10 minutes up to Now Auto-update Compare

TRANSACTION	TOTAL TIME	INVOCATIONS	AVG. TIME
RmiHandler.remoteOperation [method samples] [time lines]	4,531 s	3,660	1,238 ms
demoView3 [method samples] [time lines]	1,251 s	234	5,349 ms
demoView4 [method samples] [time lines]	1,154 s	215	5,368 ms
demoView5 [method samples] [time lines]	1,148 s	211	5,441 ms
▼ Nested demo transaction [method samples]	1,146 s	211	5,431 ms
JPA/Hibernate	156 s	211	743 ms
JDBC	71,074 ms	870	81,694 µs
JNDI	15,348 ms	211	72,742 µs
Inventory checks via RM [show remote call] [method samples]	777 s	211	3,686 ms
Exchange rate checks via web service [show remote call] [method samples]	74,193 ms	211	351 ms
demoView1 [method samples] [time lines]	1,141 s	216	5,285 ms
Exchange rate to EUR [method samples] [time lines]	1,092 s	14,555	75,045 µs

Filter: Apply Clear Options

Clicking on the "show remote call" link opens a new dialog that shows the call tree of transactions in the remote VMs that were started by the selected transaction. At the top-level, the call tree is partitioned into VM groups.

Remote Calls

**View Remote Calls**  
Below you see transactions in different VMs that were triggered by the selected transaction. Transactions close to the boundaries of the selected period may not be included.

TRANSACTION	TOTAL TIME	INVOCATIONS	AVG. TIME
▼ group DemoWorkers			
▼ RmiHandler.remoteOperation	773 s	630	1,228 ms
▼ Remote demo transaction	773 s	630	1,228 ms
JDBC	470 s	630	746 ms
JNDI	47,093 ms	630	74,751 µs
Exchange rate checks via web service [show remote call]	224 s	630	356 ms

OK

The call tree only shows the transaction in the target VMs. If there are further remote calls, you will see "show remote call" links, just like in the original call tree view. Clicking on those links replaces the call tree in the dialog and shows a **[Back In History]** button that can be used to return to the previous level. If the button disappears after clicking it, you are now looking at the first level.

It is also possible to move in the opposite direction and ask which remote calls in other VMs started transactions in the currently selected VMs. First, change the **remote origins** drop-down to "Split". By default, that drop down is set to "Merge" and you cannot see the remote origins. With the "Split" setting, "Called from ..." nodes appear at the top-level containing transactions that were started by remote calls from other monitored VMs. If you are looking at VM group or VM pool data, a split is performed for each VM group, but not for individual VMs. If the data is not granular enough, consider making your VM group structure more granular. If you are looking at the data for a single named VM, the split is performed for each remote VM.

Policies: Merge Remote origins: Split Show 10 minutes up to Now Auto-update Compare

TRANSACTION	TOTAL TIME	INVOCATIONS	AVG. TIME
Called from pool DemoWeb <a href="#">show remote origin</a>	3,608 s	5,500	656 ms
RmiHandler.remoteOperation <a href="#">[method samples]</a>	3,401 s	2,747	1,238 ms
Remote demo transaction <a href="#">[method samples]</a>	3,400 s	2,747	1,238 ms
JDBC	2,076 s	2,747	755 ms
Exchange rate checks via web service <a href="#">[show remote call]</a> <a href="#">[method samples]</a>	984 s	2,747	358 ms
Exchange rate to EUR <a href="#">[method samples]</a>	207 s	2,753	75,272 µs
demoView1 <a href="#">[method samples]</a> <a href="#">[time lines]</a>	1,036 s	199	5,208 ms
demoView3 <a href="#">[method samples]</a> <a href="#">[time lines]</a>	998 s	189	5,280 ms
demoView4 <a href="#">[method samples]</a> <a href="#">[time lines]</a>	987 s	180	5,485 ms
demoView2 <a href="#">[method samples]</a> <a href="#">[time lines]</a>	975 s	184	5,300 ms
demoView5 <a href="#">[method samples]</a> <a href="#">[time lines]</a>	897 s	165	5,440 ms
Demo JDBC Job <a href="#">[method samples]</a> <a href="#">[time lines]</a>	723 s	129	5,605 ms

Filter: Apply Clear Options

Similarly to remote calls, clicking on the "show remote origins" link brings up a modal dialog. In the dialog, the count and duration of the selected transaction is attributed to the remote transactions that started it. The data is presented as hotspots with the innermost transactions from the originating VM at the top level. Analogous to the "remote calls" dialog, you can jump further back in a chain of remote calls when you see a "show remote origins" link next to a transaction. The **[Back In History]** button brings you closer to the original transaction.

Remote Origins

**View Remote Origins**  
Below you see a list of remote origins that triggered the selected transaction. When you open the tree, you see further back traces in the same VM. Transactions close to the boundaries of the selected period may not be included.

HOT SPOT	TOTAL TIME	INVOCATIONS	AVG. TIME
Inventory checks via RMI	3,934 s	3,180	1,237 ms
Cumulated backtraces	3,934 s	3,180	1,237 ms
Nested demo transaction	3,934 s	3,180	1,237 ms
demoView3	873 s	702	1,243 ms
demoView1	811 s	648	1,251 ms
demoView4	783 s	642	1,221 ms
demoView5	773 s	630	1,228 ms
demoView2	691 s	558	1,239 ms

OK



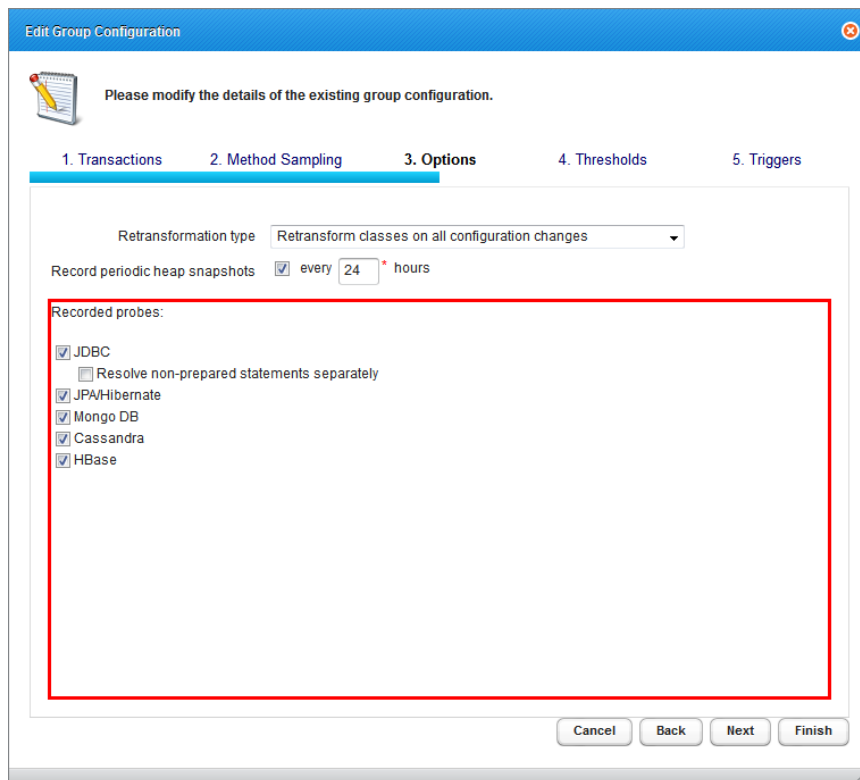
## Probes

Transactions typically use external resources that are often the cause for performance problems. The most prominent example is JDBC. Just looking at the name of a slow transaction may not generate any insight, but if you can see the SQL statement that is responsible for most of the time, you can often find the cause of the slowness right away.

perfino records this kind of data with probes. Probes generate payload data that is attached to the currently running transaction. If no transaction is running, the data is discarded.

### Configuration

To see the list of available probes, you can edit the recording settings of a VM group and go to the "Options" step. Most of the probes are for databases or messaging systems. You can switch off subsystems that you are not interested in. For example, if you use JPA, but do not want the overhead of measuring the duration of JPA calls, you can deselect the "JPA/Hibernate" check box.



### Probes in the call tree

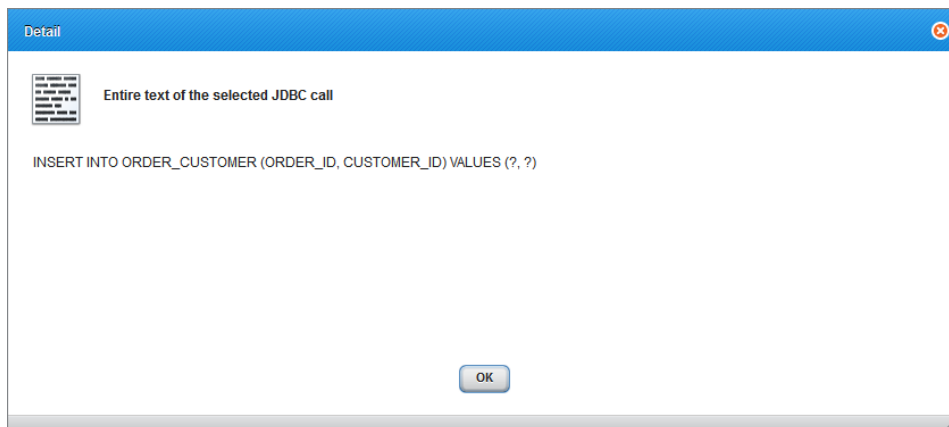
Every transaction that performs an operation measured by a probe shows this information as child nodes in the call tree. Each probe type has a separate grouping node with the name of the probe.

Policies: Split Remote origins: Merge Show 10 minutes up to Now Auto-update Compare

TRANSACTION	TOTAL TIME	INVOCATIONS	AVG. TIME
RmiHandler.remoteOperation [method samples] [time lines]	4,531 s	3,660	1,238 ms
demoView3 [method samples] [time lines]	1,251 s	234	5,349 ms
demoView4 [method samples] [time lines]	1,154 s	215	5,368 ms
demoView5 [method samples] [time lines]	1,148 s	211	5,441 ms
Nested demo transaction [method samples]	1,146 s	211	5,431 ms
JPA/Hibernate	156 s	211	743 ms
JDBC	71,074 ms	870	81,694 µs
SELECT * FROM ORDER O WHERE O.DATE >= ? [show]	31,436 ms	211	148 ms
INSERT INTO ORDER (ID, NAME, OPTIONS) VALUES (?, ?, ?) [show]	13,118 ms	211	62,172 µs
INSERT INTO CUSTOMER (ID, NAME, OPTIONS) VALUES (?, ?) [show]	12,937 ms	211	61,314 µs
INSERT INTO ORDER_CUSTOMER (ORDER_ID, CUSTOMER_ID) VALUES (?, ?) [show]	12,804 ms	211	60,685 µs

Filter: All Apply Clear Options

Often, the description of the operation is very long, like a complex SQL statement with many output fields, joins and conditions, and the finite length of the table column prevents further inspection. In this case, click the "show" link in the same table tell. A dialog will be shown, where you can scroll through the payload data and copy it to the clipboard.



## Probe hot spots

Sometimes your focus is not on transactions, but on the probe data. For example, when trying to speed up your database, the question is: What are my slowest database operations? For that purpose, perfino offers the probe hot spot views. Each probe has a separate view of this kind where a list of operations is shown together with cumulated backtraces.

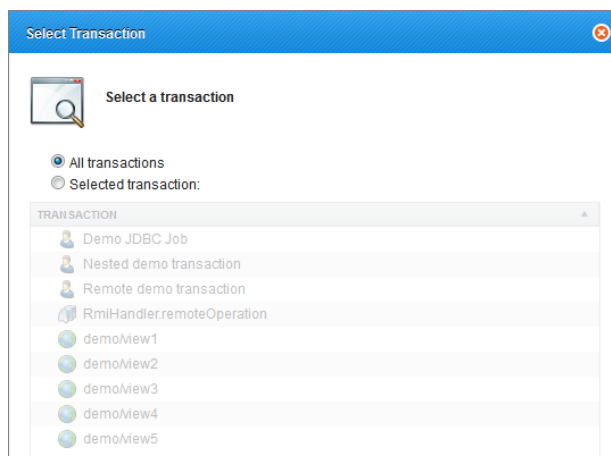
Policies: Split    Show 10 minutes up to Now    Auto-update    Compare

**Choose Transaction** [All transactions]

HOT SPOT	TOTAL TIME	INVOCATIONS	AVG. TIME
SELECT i.id, i.availability, i.name FROM inventory i WHERE i.delayed	2,397 s	3,174	755 ms
SELECT * FROM ORDER O WHERE O.DATE >= ?	137 s	917	150 ms
Cumulated backtraces	137 s	917	150 ms
Nested demo transaction	137 s	917	150 ms
demoView1	30,270 ms	199	152 ms
demoView3	28,657 ms	189	151 ms
demoView2	27,682 ms	184	150 ms
demoView4	26,898 ms	180	149 ms
demoView5	24,412 ms	165	147 ms
SELECT SUM(o.price * o.quantity) FROM customers c LEFT JOIN ord	97,231 ms	129	753 ms

Filter:    Apply    Clear    Options

The probe hot spots view can be filtered on a per-transaction basis. By default, cumulated data for all transactions in the selected time interval is shown. When clicking on the **[Choose transaction]** button, you can select one transaction from which the probe hot spots should be taken.



## Probe data granularity

Probe operations can be extremely frequent and they usually don't do the same thing over and over again. perfino makes an effort to modify the displayed name of a probe so that frequently varying information is replaced with placeholders. Imagine a loop of SQL queries where an ID varies in each statement. Showing the raw SQL in the UI would mean that perfino would potentially have to remember millions of SQL statements per hour.

This is why perfino by default does not show SQL for JDBC statements that have been created with the `createStatement(...)` method. Rather, perfino shows those operations with a description of "unprepared select statement". SQL is only shown for prepared statements, without resolving the single parameters. If you do not use prepared statements and need to see more detail, the "Resolve non-prepared statements separately" option in the probe configuration enables SQL recording for all unprepared statements.

These considerations only apply to the call tree where data is retained indefinitely at the highest consolidation level. For sampling, where data is discarded in the medium term, perfino always records the SQL of unprepared statements.


The same strategy with respect to data granularity is used for other databases. For example, in MongoDB operations, all values are replaced with question marks.

## Method Sampling

### What is method sampling?

Policies in perfino help you to identify slow transactions, but the cause of the slowness is often not clear, since it usually originates in a subsystem that is not monitored separately. In that case, you need **more information on the method level**.

The way perfino gets this information is by periodically inspecting the thread that executes a transaction and by recording its stack trace. By comparing subsequent "samples", perfino can build an approximate call tree. It has no information on invocation counts, since it is not possible to tell if a method is still being called since the last sample was taken - or if it has exited in the meantime and is being called again. The more time is spent in a method, the better the information in the sampling call tree is. As such, sampling is a great tool to **find hot spots** in the monitored VM.



### Automatic sampling


perfino can perform automatic sampling for transactions in two ways:

- **Periodic sampling**

With periodic sampling, perfino takes one transaction from a particular transaction definition every couple of minutes and samples it from start to finish. Alternatively, perfino can take every  $n$ -th transaction and sample it. The latter may be preferable for rarely executed transactions.

- **Sampling for slow transactions**

The most interesting samples are for slow transactions, because they let you investigate why a transaction is slow. perfino can be configured to sample slow transactions automatically. When a transaction passes the threshold of a policy violation like "slow" or "very slow", perfino then starts sampling until the end of the transaction. The period of time from the beginning of the transaction until the policy violation was reached is not sampled.



Both automatic sampling modes are configured in the "policies" step of a transaction definition.

Create Web Request

Please enter the details of the new business transaction.  
Monitor web requests that come from outside the JVM

1. Filter      2. Naming      3. Policies

☒ Define policies

Slow events: 500 \* % slower than average

Very slow events: 1000 \* % slower than average

Overdue events: 500000 \* ms and slower

☒ Split transaction tree for policy violations

Transaction errors: ☒ Error throwables ☒ Runtime exceptions ☐ Checked exceptions

☒ Logged errors ☐ Logged warnings

Ignored HTTP error codes:

Sampling for policy violations: Very slow and more severe

Periodic sampling every: 3 Hours

End user experience monitoring: 20 \* % of all web requests

Cancel Back Next Finish

## Filters

Method-level call stacks can become very deep. Too much detail is usually distracting, like the internal call structure of external libraries and the JRE. For each VM group, you can define the recording scope, either by specifying a list of packages that should be recorded (inclusive mode) or by listing packages that should not be recorded (exclusive mode). In both cases, the package filters are recursive.

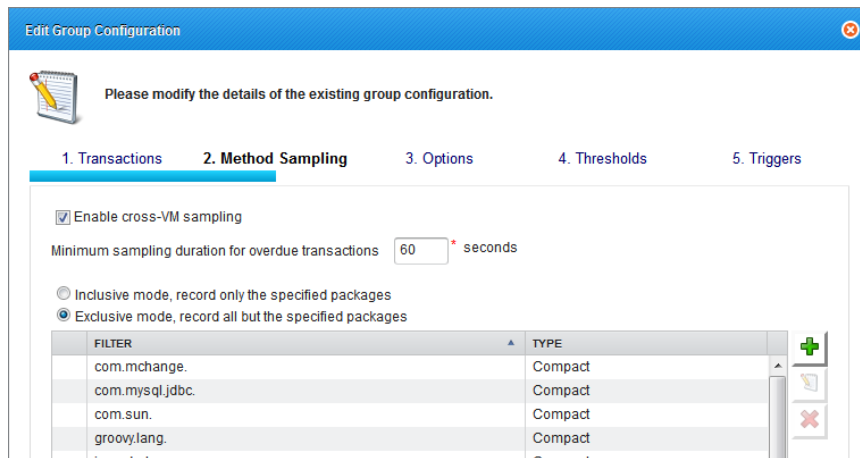
For the exclusive mode, there are two types of filters:

- **Compact**

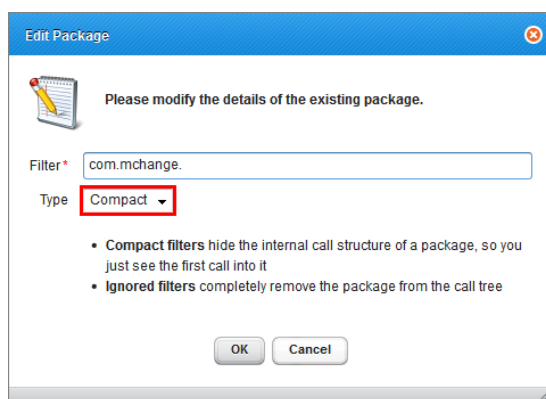
"Compact" means that internal calls will be excluded from the sampling call tree. If a filtered class is called as an entry point (like `java.lang.Thread.run()`) or from an unfiltered class, it will be displayed, but further calls within filtered classes will be hidden. This allows you to focus on the important methods that are in your code and where you are able to change something. However, you are still able to see calls into third-party code.

- **Ignored**

"Ignored" means that filtered classes will be completely removed from the tree.



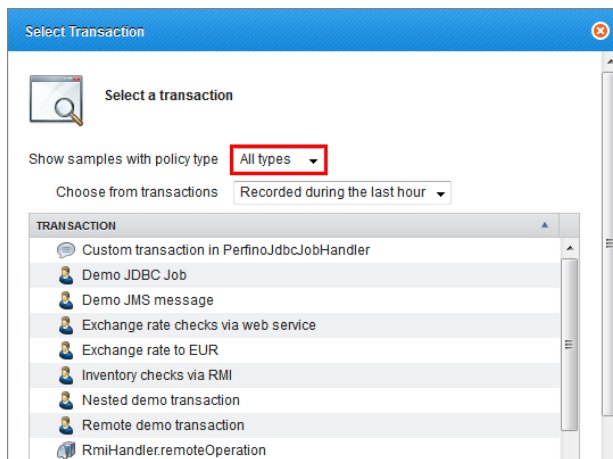
When you add a filter by browsing in archives or connected VMs, it is compact by default. To change its type to ignored, edit the filter and change the drop-down in the edit dialog.



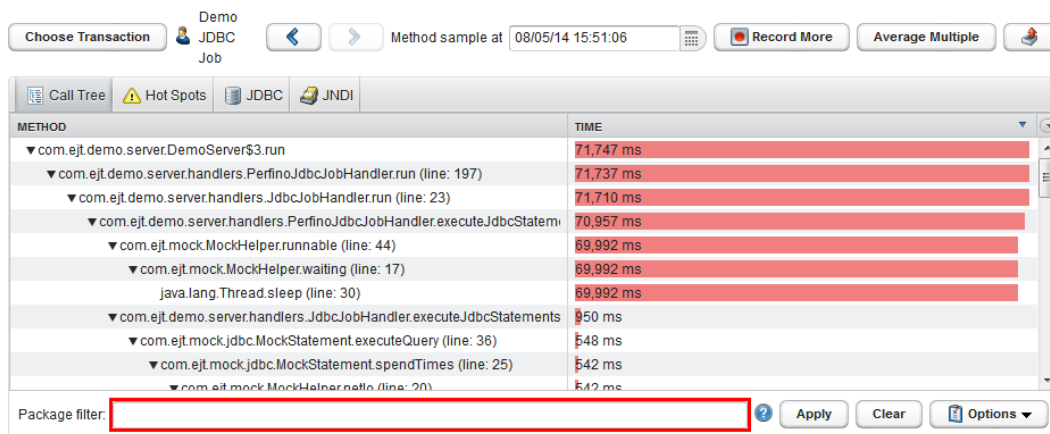
## Method sampling view


To see samples in the method perfino UI, you first have to select a transaction and a policy type. While you can do that with the **[Choose transaction]** button, it is usually easier to click the "show methods samples" link next to a transaction in the call tree or hot spots views.

When choosing a transaction manually, changing the selected policy type to something other than "All types" shows the transactions for which a corresponding policy violation has actually recorded at least one sample in the selected period. In the "All types" setting, all transactions in the selected period are shown regardless of whether a sample has been recorded. In that way you can select any transaction and tell perfino to record a sample.



The selected transaction and policy type are shown in the header and the navigation buttons move between different matching samples. You can use the date chooser to jump to the sample that is closest to a selected time and date.



The sampling call tree shows methods except when sampling crosses a VM boundary. In that case, a  VM node is inserted. If you want to limit sampling to the current VM, you can switch off cross-VM sampling in the sampling configuration of the VM group where the transaction is executed.

Where available, the sampling call tree shows **line numbers**. To interpret a line number, always look at the class of the parent item. Intuitively, you would expect the line number on a method item to show the line number of that method, but that would not be very useful information. The line number cannot be placed directly on the call site, because each call site can have multiple children with different associated line numbers.

After the initial inspection of a sampling call tree, you will often want to focus on a particular package. Modifying the sampling options to exclude all unwanted packages and re-recording the sample is not a good option, since the package of interest may vary or the performance problem may not be easily repeatable. To adjust your point of view in the sampling call tree, use the **package filter** at the bottom of the view. You can enter a single package, or a list of packages separated by commas.

Packages are included recursively, i.e. if you include `com.mycorp.`, then packages like `com.mycorp.project` are included as well. If you want to exclude a particular package, prefix



it with a minus sign. For example, if the classes in `com.mycorp.` are of interest, but `com.mycorp.algorithm` is distracting, enter

```
com.mycorp., -com.mycorp.algorithm.
```

into the filter bar. To better understand the call structure, the **first excluded node between included nodes is always shown**, so even after excluding a package you might see a few instances of it in the call tree, just not more than two consecutive method calls in such a package. This corresponds to the "Compact" mode for exclusive filters in the "Sampling" configuration step.

If the first package in the filter bar is an exclusion (with a prefixed minus sign), then all other packages will remain visible. If the first package in the filter bar is an inclusion, then all other packages will be hidden.

The **hot spot tab** shows the methods where most time is spent together with cumulated back traces. Note that the set of hot spots strongly depends on the granularity of the measurement which is governed by the defined filters. Defining different filters in the sampling options or applying different view filters may change the list of hot spots completely.

The screenshot shows the 'Hot Spots' tab in a performance monitoring tool. At the top, there is a 'Choose Transaction' button, a 'Demo' label, and a 'JDBC Job' label. Below these, there are navigation arrows, a 'Method sample at' field showing '08/05/14 15:51:06', a 'Record More' button, and a highlighted 'Average Multiple' button. The main area displays a table with two columns: 'METHOD' and 'TIME'. The table lists several methods and their execution times. The first method, 'java.lang.Thread.sleep', has a time of '71,695 ms' and is highlighted in red. The other methods have times in microseconds. At the bottom, there is a 'Package filter:' field with an 'Apply' button, a 'Clear' button, and an 'Options' dropdown menu.

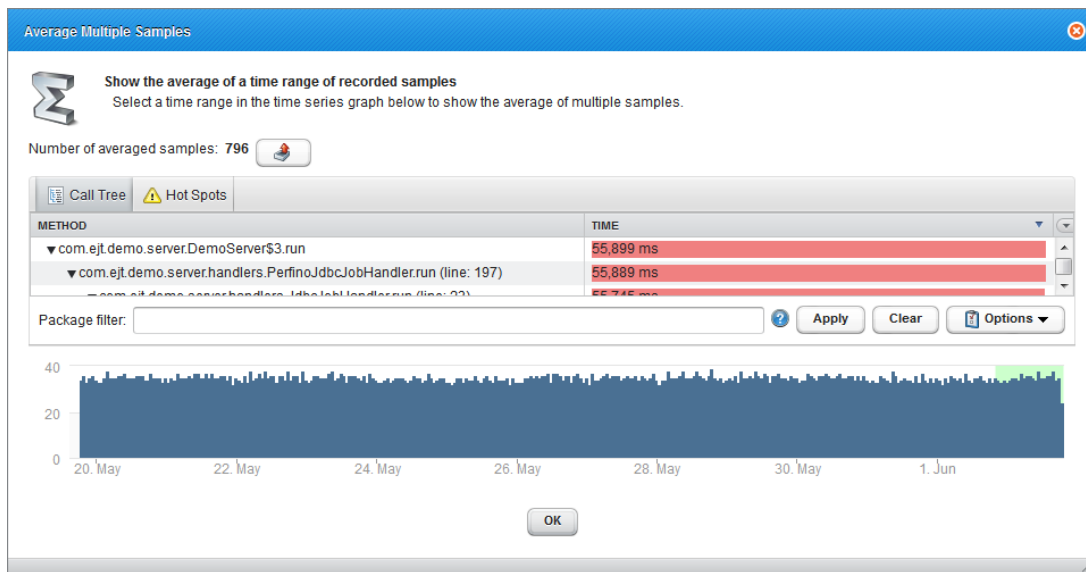
METHOD	TIME
▶ java.lang.Thread.sleep	71,695 ms
▶ com.ejt.demo.server.handlers.PerfinoJdbcJobHandler.executeJdbcStatements	15,682 µs
com.ejt.demo.server.DemoServer\$3.run	10,461 µs
▶ com.ejt.demo.server.handlers.JdbcJobHandler.run	10,436 µs
▶ com.ejt.mock.jdbc.MockStatement.executeQuery	5,226 µs
▶ com.ejt.demo.server.handlers.PerfinoJdbcJobHandler.run	5,215 µs
▶ com.ejt.mock.jdbc.MockPreparedStatement.executeQuery	5,201 µs

For each probe that recorded data during the sampling period, a separate tab is added that shows the probe hot spots [\[p. 41\]](#).



### Averaging multiple samples


The quality of sampling increases, the longer you sample. Especially for periodic sampling you collect a large number of comparable samples over time. perfino can average these samples to give you better statistics that comes closer to the actual execution times. When you click on the **[Average Multiple]** button highlighted in the screen shot above, a modal dialog will be shown where you can select a range of samples to be averaged.




The time line at the bottom shows all available samples. Note that samples are not kept indefinitely due to their high storage space requirements. The retention threshold can be configured in the general settings. By clicking and dragging a range in the time line, you chose the samples to be averaged. The total selected number of samples is shown at the top.

Just like for the single samples, you can switch between call tree, hot spots and available probe hotspots.

### Manual recording of samples

There may be cases where you require a current sample of a particular transaction and you cannot wait for periodic sampling or a policy violation that would trigger sampling. For that purpose, you can use the  **[Record More]** button in the method sampling view.

**Record More Samples**

 **Record more samples**  
perfino will record more samples for the selected transaction as they become available.

Number of samples to be recorded\*

You can even request more than one sample which will give you the ability to average the samples after they have been recorded. For rarely executed transactions, it may take a long time to fulfil the recording request.


## Telemetries

perfino observes scalar value from three different types of sources:


- Counts of entities that are managed by the perfino collector, like the number of connected VMs or the number of transactions. This is not a direct measurement in the monitored VM, but an evaluation in the perfino server.
- Telemetries that are measured inside the monitored VM, like the heap size or the average duration of a JDBC statement. Many of these telemetries are produced by probes [\[p. 41\]](#).
- Custom telemetries where you have used the perfino API in your own code.

### Telemetry data view

All telemetries are available from the "VM Data Views". Each telemetry has its own entry under the "Telemetries" node. Many telemetries are stacked area plots where the single lines add up to a total value. For example the "Transactions" telemetry shows the total number of transactions over time, split into single lines for the various policy violations [\[p. 33\]](#).




Initially, the telemetries show a 10 minute interval up to the current time. The data that is displayed here has been recorded with a resolution of one minute. Other telemetries that do not get its data from the observation of transactions will have a resolution of 10 seconds. You can move back in time with the navigation buttons, but at some point the one-per-minute resolution will end. perfino consolidates telemetries to progressively more coarse-grained resolutions and keeps them for progressively longer periods of time. Switching to the next interval shows a total extent of 3 hours, with a resolution of 2 minutes.



The full table with all intervals, resolutions and retention times is given below:

Display interval	Resolution	Retention time
10 minutes	10 seconds (1 minute for transaction-based data)	6 hours (48 hours for transaction-based data)
3 hours	2 minutes	10 days
3 days	1 hour	1 year
30 days	12 hours	unlimited

Some telemetries have multiple data lines that do not add up to a total value. In that case, there is a drop-down box above the graph and the telemetries are shown as line plots.



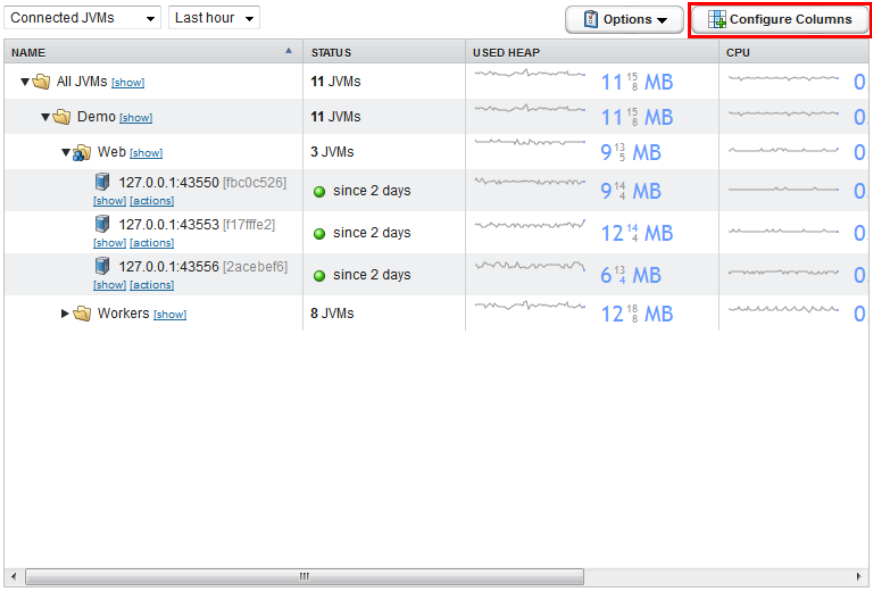
Detailed numbers can be obtained by hovering the mouse over the graph. in stacked area plots, you can toggle single data lines by clicking on the legend items.

## Sparklines

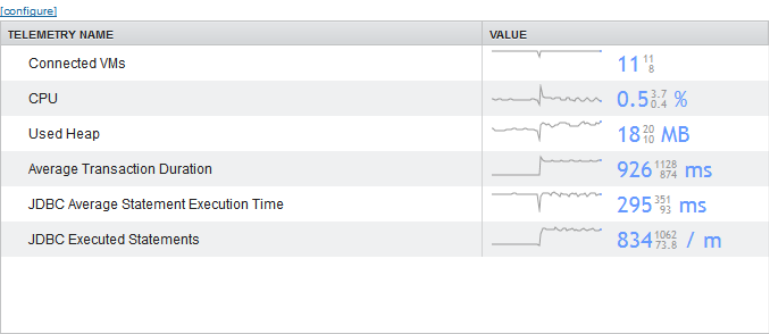
The VMs view and the dashboard show small versions of the current telemetry data as so-called "sparklines". Sparklines do not have labeled axes and are intended to give a visual impression

of the recent development of an observed scalar value. The graph is followed by the current numeric value. The superscript indicates the observed maximum value, the subscript the minimum value in the displayed time range.

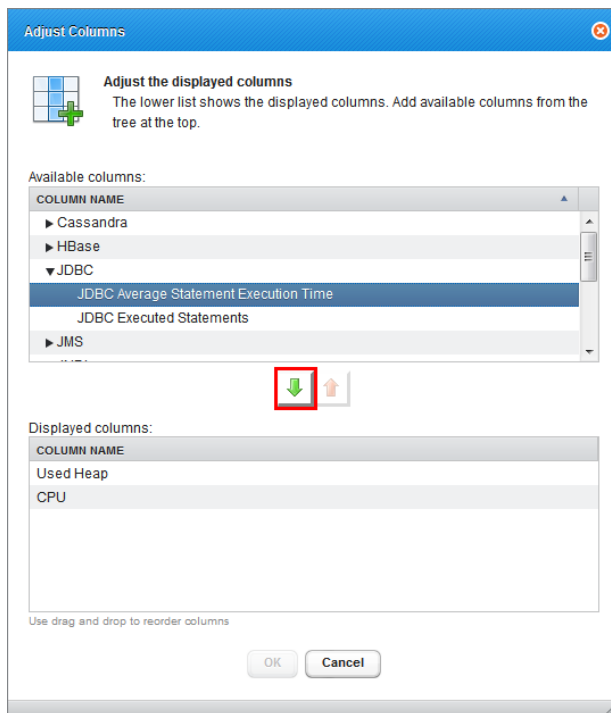
In the "VMs" view, you can configure a set of sparkline columns. This allows you to make a relative assessment of the different VMs and VM groups with respect to the monitored value. Sparklines can be scaled separately, with a common scale for each group or with a common scale for all VMs. This is configured in the options popup.



In the dashboard, sparklines are displayed in a table rather than as columns. The useful number of sparkline columns that you can add in the VMs view is limited due to the finite width of the table. In the dashboard, you can add a lot of sparkline rows without any such restrictions. Rather than showing data for separate VM groups or VMs, the dashboard only shows data for the selected VM group.



As in the VMs view, there is a "Configure" action that takes you to the list of all available telemetries. The telemetries are grouped into categories. The lower list shows the currently configured sparklines.



## Custom telemetries

To monitor any scalar value in your application, add a static method that returns a number anywhere in your code. Then, annotate the method with the `@Telemetry` annotation. You have to make sure that the containing class is actually loaded, otherwise perfino will not detect the annotated method.

You can display the custom telemetry by going to the "VM Data Views", selecting Telemetries->Custom telemetries and choosing your telemetry from the drop-down list at the top.

For more information, please see the Javadoc of the `com.perfino.annotation.Telemetry` class in the `api/doc` directory of your perfino installation.

## Time zones

All displayed times are expressed in the time zone of the server. If you are in a different time zone, there will be an offset. perfino detects this condition by inspecting the time zone provided by the request headers from your browser. In that case, the current time with an explicit time zone is shown in the header.



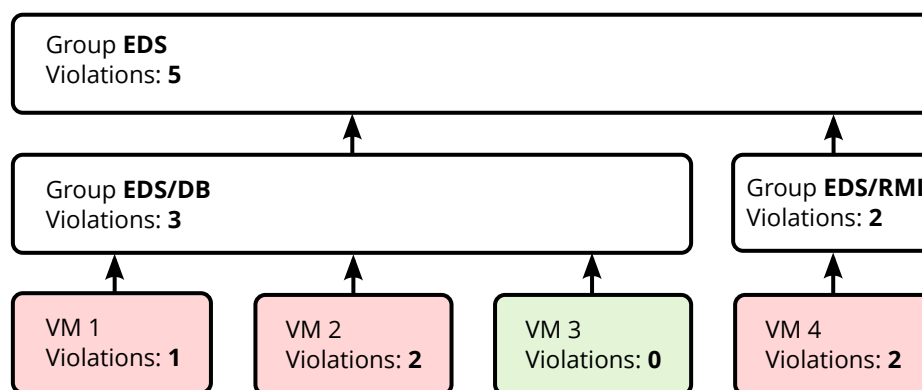
## Thresholds

Thresholds detect anomalous conditions for telemetries. Threshold violations are not directly coupled to alerts or other actions, they just **increase an associated counter**. Thresholds can be configured for single VMs or for VM groups. This is different from triggers [\[p. 58\]](#) that always operate on a VM group level rather than for single VMs.

For **single VM thresholds**, the telemetry value of each VM is checked and for each offending VM a threshold violation is created. For example, you may have an upper bound on the used heap and each VM that uses more than that gets its own threshold violation.

For **VM group thresholds**, the telemetry value of the VM group is checked and only one threshold violation is created no matter how many VMs violate the threshold individually. Imagine a database that serves 10 VMs in a VM group. If that database becomes very slow and you have defined a threshold for the average JDBC execution time, all 10 VMs will report a threshold violation. In the end, you just want one alert and not 10. In addition, the averaged group value is smoother and fewer spurious threshold violations will occur than for single VM thresholds.

All telemetries that show frequency data are summed for VM groups. If you have 10 VMs that perform 10 JDBC statements per second, the VM group will show 100 JDBC statements per second. If your acceptable range is defined for the total values rather than the loads on the single VMs, then you have to configure your threshold as a VM group threshold.



The counter is maintained on a per-VM level for single VM thresholds as well as on a VM group level. When a threshold violation occurs, it bubbles up through the parent hierarchy, increasing all the associated counters by one. At each VM group in that hierarchy, you can define triggers that react to the corresponding number of threshold violations.

### Configuration

You can define thresholds for each telemetry in the recording settings of a VM group. A threshold definition has an associated telemetry and optional **lower and upper bounds**. The available units depend on the selection of the telemetry.

**Create Threshold**

Please enter the details of the new threshold.

Telemetry: JDBC Average Statement Execution Time Choose

Custom name:

Threshold target: Single VMs

Lower bound:  s

Upper bound: ☒ 300 ms

Minimum time: 10 seconds


Inhibit duplicate time: 1 minutes

☒ No duplicates if threshold remains violated continuously

OK Cancel

In some cases, you need multiple threshold definitions for the same telemetry, for example to designate different severities. In that case, you have to give the thresholds **different custom names**, like "High average JDBC execution time" and "Very high average JDBC execution time". These definitions can be used as the basis for triggers with different escalation strategies.

To avoid spurious firing of threshold violations, a minimum time can be configured for which the bounds have to be exceeded before a threshold violation is detected. After a threshold has fired, there is an inhibition time, during which the threshold is muted and cannot fire again. This serves to prevent frequent firing in the case of oscillating conditions.



If the threshold is continuously violated, perfino will not fire any more threshold violations during that time. There is a check box to disable this constraint. If disabled, new threshold violations will be fired at a constant rate with a period of the configured inhibition time.

Like for transactions and method sampling filters, you can save and load sets of thresholds. This makes it possible to copy and paste threshold definitions between VM group configurations.






## Triggers

While policies [\[p. 33\]](#) and thresholds [\[p. 55\]](#) detect anomalous conditions and display them in the dashboard and the VM data views, they cannot take any actions by themselves. With triggers, you can react to policy and threshold violations and execute a list of configurable actions.

### Mechanism

Triggers operate on a different level than policies and thresholds. The latter are configuration options that are applied to single VMs or VM groups. Triggers, on the other hand are not directly coupled to single policy or threshold violations. Rather, they **react to sequences of such events** that originate from all monitored VMs in a VM group.


This mechanism is intended to give you greater flexibility for deciding what constitutes a condition that requires a particular action. For example, you might expect up to 1 slow URL invocation when a cache is rebuilt. However, if there are 5 slow URL invocations per hour or more, then something is wrong. The definition of what is acceptable and what is not, strongly depends on the type and the implementation of your application.



### Configuration

In the recording settings, you can edit triggers for each VM group. Triggers operate on all recursively contained VMs. It is possible to define different triggers for a VM group and an ancestor VM group, both sets of triggers are handled separately.

In some cases, triggers are fired too often. For that case, perfino allows you to disable a trigger until you have time to figure out how to change the underlying configuration.



Like for transaction definitions and other entities in perfino, triggers can be saved to and loaded from trigger sets. This enables you to copy and paste trigger definitions as starting points to multiple VM groups.

### Trigger types

There are three types of triggers:

- **Policy triggers**

Policy triggers are fired when the policy violations of a set of transactions exceed a defined number during a specified period of time. The filter text field takes a transaction name pattern, either a wildcard expression or a regular expression. Also, policy triggers are configured for particular policy violation types.

You can have multiple policy triggers, each matching different transaction names. Unlike for transaction definitions, there is no name matching for policy triggers where only the first matching entry is used. If you add more than one trigger for the same transaction name, you will probably want to set different event rates, otherwise both triggers will be fired at the same time.

The trigger condition is decoupled from the actual condition of the policy violation. Various transactions can define different times after which a transaction is characterized as "very slow", the policy trigger then counts these events.

- **Threshold triggers**

Threshold triggers are fired when the rate of threshold violations for a selected telemetry exceeds a configured value. This requires that you have configured at least one threshold [\[p. 55\]](#) for the same VM group.

Like for the policy trigger, the trigger condition does not define an actual threshold. Thresholds can be defined differently in different descendant VM groups and the threshold trigger then counts threshold violations.

As an example, imagine you have two groups of machines, powerful machines and legacy machines. On a powerful machine, the number of threads may not exceed 1000, on a legacy machine that threshold is just 500. You would create VM groups named "Powerful" and "Legacy" and define the corresponding thresholds in the recording settings of each group as well as a default threshold in the "All VMs" group. Then, in the "All VMs" group, you would define a threshold trigger for the thread count telemetry. That trigger would handle both VM groups at the same time.

- **Connection trigger**

The number of connected VMs is a scalar value that originates in the perfino collector and not in the monitored VMs. The connection trigger is intended to take action in the case that too

few VMs are running. The configured minimum number is not reached immediately when the perfino server is started, so you have the option of arming the trigger only after the minimum number has been reached for the first time, or after a fixed amount of time has passed.

Create Connection Count Trigger

Please enter the details of the new trigger.  
Execute a list of actions when the number of connected VMs drops below a specified threshold.

Minimum number of VMs \* 10

Trigger is armed  
☐ After the configured threshold has been reached  
☒ Immediately

Minimum time 1 \* minutes

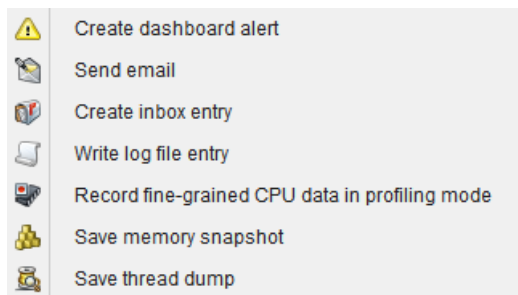
Inhibition time 12 \* hours

Trigger actions

ACTION
--------

## Trigger actions

Each trigger can have an arbitrary list of actions.



The types of actions that can be added to a trigger can be ordered into two categories:

- **Notification actions**

Notifications can be created for consumption inside perfino. Apart from the alerts mechanism [p. 63], you can write to the log file or to the inbox. Alerts are visible in the dashboard and are saved as historical data. The inbox is maintained for each user and has an "unread status" on a per-user basis. Externally, you can send emails. For the latter, the SMTP configuration in the general settings has to be valid.

- **Data collection actions**

As a reaction to anomalous conditions, you can decide to record more data for a detailed analysis. These data recording options are more intrusive than the regular perfino recording.

For example, saving an **HPROF heap snapshot** should not be done on a regular basis, but if memory is low, it will help you to find a memory leak. HPROF snapshots are written directly by the JVM and do not require a native JVMTI profiling agent to be loaded. As such, this is a low-risk activity, although the VM is halted until the snapshot is saved. JProfiler and other Java profilers can open HPROF snapshots.

**Thread dumps** are a basic low-overhead way to inspect what is currently happening in a JVM. A policy trigger with event type "overdue" can save a thread dump, so you can immediately see in which method a transaction is hanging.

**Recording data in profiling mode** is an escalation in CPU data gathering that requires a native JVMTI agent to be loaded. The native perfino agent operates in a restricted mode that is optimized for minimum overhead and minimum risk. More information on this topic is available in the chapter on cross-over to profiling [\[p. 77\]](#).

The data collection actions are also available in the "VMs view" when clicking on the "actions" link next to a connected VM.

The list of actions is executed in order. If one action fails, perfino jumps to next action and does not terminate the execution of the trigger actions.




## Alerts

Alerts communicate anomalous conditions to you and other users of the perfino UI server. By default, perfino does not create any alerts since the conditions that are noteworthy on such a level are highly individual.

### Configuration

To create alerts, you have to add a trigger [\[p. 58\]](#) that fires for a specified policy or threshold violation. In the list of actions, add a **"Create dashboard alert" action**.




The screenshot shows the 'Create Policy Trigger' dialog box. It has a blue header bar with the title 'Create Policy Trigger' and a close button. Below the header, there is a document icon with a green plus sign and the text: 'Please enter the details of the new trigger. Execute a list of actions when a policy is violated for a specified number of times.' The dialog contains several fields: 'Policy types' with checkboxes for 'Normal', 'Slow', 'Very slow' (checked), 'Overdue', and 'Error'; a 'Filter' field with 'Order \*' and a 'Wildcard comparison' dropdown; 'Fire after' set to '10' events in one 'hour'; and 'Inhibition time' set to '12' hours. At the bottom, there is a 'Trigger actions' section with a table header 'ACTION' and one row containing a yellow warning icon and the text 'Create dashboard alert'. To the right of the table are icons for adding (+), deleting (-), and refreshing (circular arrow).

When configuring the "Create dashboard alert" action you can specify a category. This selection determines the color of the alert bars in the dashboard:

- **Error**
  - Red alert bars
- **Warning**
  - Orange alert bars
- **Info**
  - Green alert bars

With the configured text message, you will be able to identify the origin of the alert in the data views.




The screenshot shows the 'Create Create Dashboard Alert' dialog box. It has a blue header bar with the title 'Create Create Dashboard Alert' and a close button. Below the header, there is a document icon with a green plus sign and the text: 'Please enter the details of the new action. Add a dashboard alert. The alert will be shown in the time line of the dashboard, and you can click on it to see details.' The dialog contains a 'Category' dropdown menu set to 'Error' and a large text area labeled 'Text \*' containing the text 'Order handling SLA violation'. At the bottom, there are 'OK' and 'Cancel' buttons.

## Data

Alerts are displayed in the dashboard. They are drawn as **alert bars** over the transaction timeline. In this way, you can quickly correlate an alert with an increase in the general activity of your application.

If there is more than one alert in the same time slot, just one alert bar is shown with a height proportional to the number of contained alerts. If those alerts are from different categories, the associated colors are shown stacked on top of each other. The tool tip shows the number of alerts in each category.



Clicking on the alert bar brings up the **alert detail dialog** where you can see exact times, categories and the text messages that were configured in the trigger action. To view all alerts in the time range that is currently displayed by the transaction time line, click on the "show all alerts" link in the top-right corner.

The Alert Details dialog shows a table of alerts with the following columns: SEVERITY, TIME, DESCRIPTION, and LAST VM. The table contains 6 rows of data.

SEVERITY	TIME	DESCRIPTION	LAST VM
Error	5/8/14 3:50:00 PM CEST	Demo alert caused by	Demo/Workers/RMI ha
Warning	5/8/14 3:23:00 PM CEST	JDBC performance del	Demo/Workers/JDBC/
Warning	5/8/14 3:20:00 PM CEST	Multiple very slow trans	Demo/Workers/Web s
Error	5/8/14 3:20:00 PM CEST	Demo alert caused by	Demo/Workers/RMI ha
Error	5/8/14 3:12:00 PM CEST	Demo alert caused by	Demo/Workers/RMI ha
Warning	5/8/14 3:03:00 PM CEST	JDBC performance del	Demo/Workers/JDBC/

At the bottom of the dialog, there is a Filter dropdown set to 'All', an input field, and buttons for Apply, Clear, Options, and OK.

The "Last VM" column in the detail dialog shows the VM that was responsible for triggering the alert. Note that triggers always react to policy and threshold violations of entire VM groups, so multiple VMs may contribute to the condition that creates a single alert.

The dashboard only shows alerts from the selected period and from VMs that are recursively contained in the selected VM group. By changing the period or the VM group selection, you can limit the displayed alerts to your focus of interest.



To analyze alerts in a historical context, go to the alerts view in the "VM data views". Here you can see alerts from previous days, weeks or months. If the "Auto-update" check box is selected, alerts are displayed as soon as they are generated by perfino. Also, you can inspect alerts from a single VM only, which is not possible in the dashboard.

⏪ ⏩ Show 1 day up to Now ☐ Auto-update

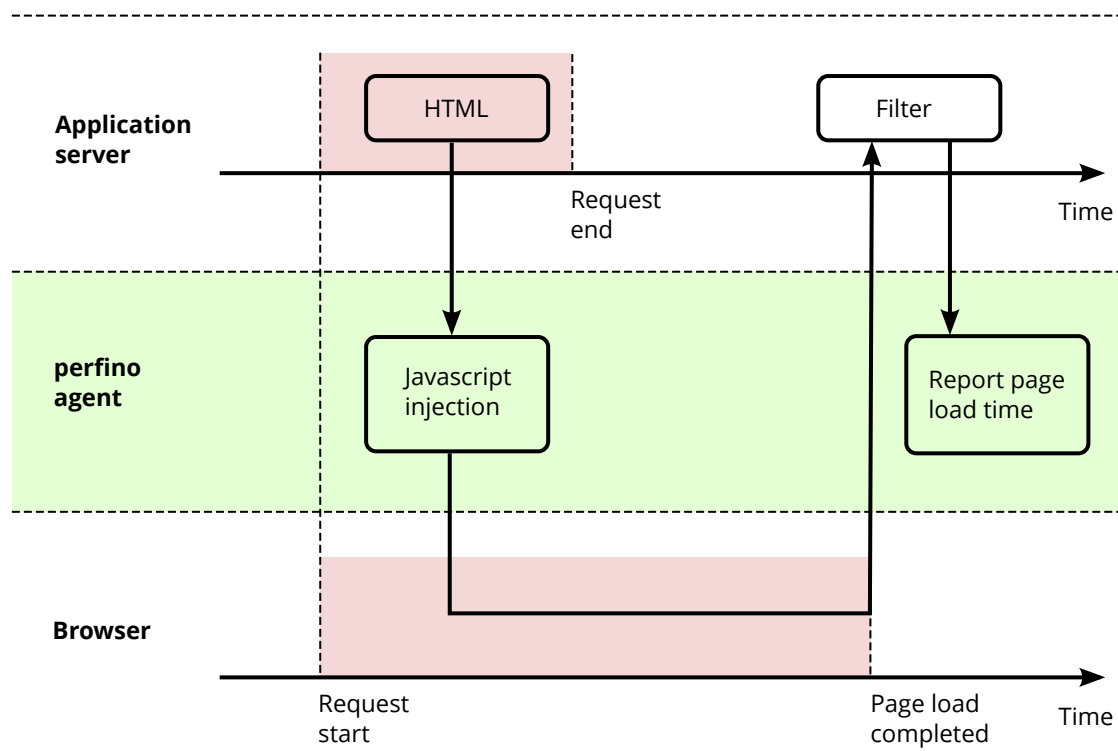
SEVERITY	TIME	DESCRIPTION	LAST VM
Error	5/8/14 3:50:00 PM CEST	Demo alert caused by a threshold violation of the d	Demo/Workers/RMI handler
Warning	5/8/14 3:23:00 PM CEST	JDBC performance deterioration	Demo/Workers/JDBC/127.0.0.1:49514 [837d7310]
Warning	5/8/14 3:20:00 PM CEST	Multiple very slow transactions were detected	Demo/Workers/Web service handler
Error	5/8/14 3:20:00 PM CEST	Demo alert caused by a threshold violation of the d	Demo/Workers/RMI handler
Error	5/8/14 3:12:00 PM CEST	Demo alert caused by a threshold violation of the d	Demo/Workers/RMI handler
Warning	5/8/14 3:03:00 PM CEST	JDBC performance deterioration	Demo/Workers/JDBC/127.0.0.1:44092 [69ef739a]
Warning	5/8/14 2:56:00 PM CEST	Multiple very slow transactions were detected	Demo/Workers/JDBC/127.0.0.1:44093 [dabb662]
Error	5/8/14 2:42:00 PM CEST	Demo alert caused by a threshold violation of the d	Demo/Workers/RMI handler
Warning	5/8/14 2:15:00 PM CEST	Multiple very slow transactions were detected	Demo/Workers/JDBC/127.0.0.1:44093 [dabb662]
Warning	5/8/14 2:14:00 PM CEST	JDBC performance deterioration	Demo/Workers/JDBC/127.0.0.1:44092 [69ef739a]
Error	5/8/14 2:03:00 PM CEST	Demo alert caused by a threshold violation of the d	Demo/Workers/RMI handler
Warning	5/8/14 1:34:00 PM CEST	Multiple very slow transactions were detected	Demo/Web/127.0.0.1:44089 [1dbfaf0d]
Error	5/8/14 1:32:00 PM CEST	Demo alert caused by a threshold violation of the d	Demo/Workers/RMI handler

Filter: All Apply Clear Options

## End User Experience Monitoring

For web requests, the server transaction time is shorter than the page load time in the browser. Secondary requests that are triggered by the page can lead to an unacceptable end user experience even though your server monitoring indicates that everything is fine.

perfino includes a servlet filter that injects a Javascript snippet into HTML pages and reports the page load time back to the server. In addition, a transaction ID is sent back so that perfino can correlate execution times of web transactions and the associated page load times. The snippet is very small and does not require any external libraries. The overhead is also extremely small, since the snippet only runs after the page has loaded and the Web Performance API is used to get the timing directly from the browser.



There are no cross origin request problems with this scheme, because the browser sends the data back to the application server where the page was loaded from. A servlet filter provided by perfino reads the timing data and reports it to the perfino monitoring agent. From there it is transmitted to the perfino collector.

Typically, you do not need page load times from all HTTP requests that create HTML pages, so you just sample a particular percentage. This percentage value can be configured in the policies configuration of web transactions. To disable this feature, set the percentage value to 0%.

Create Web Request

Please enter the details of the new business transaction.  
Monitor web requests that come from outside the JVM

1. Filter 2. Naming 3. Policies

☒ Define policies

Slow events  \* % slower than average ▾

Very slow events  \* % slower than average ▾

Overdue events  \* ms and slower ▾

☒ Split transaction tree for policy violations

Transaction errors ☒ Error throwables ☒ Runtime exceptions ☐ Checked exceptions

☒ Logged errors ☐ Logged warnings

Ignored HTTP error codes

Sampling for policy violations  ▾

Periodic sampling every  \* Hours ▾

End user experience monitoring  \* % of all web requests

Cancel Back Next Finish

In the VM data views, the "Transactions" view category holds the "End user monitoring" view. It is structured like the hot spots view with different time ranges and navigation buttons to show previous intervals. The table lists all web transactions for which end user monitoring data has been received by the perfino collector.

◀ ▶ Show 10 minutes ▾ up to Now ▾ ☐ Auto-update

WEB TRANSACTION NAME	TRANSACTIONS	Ø TRANSACTION DURATION	SAMPLES	Ø PAGE LOAD DURATION	Ø OVERHEAD
demo/view3	223	5,336 ms	45	6,850 ms	28 % <a href="#">[actions]</a>
demo/view5	212	5,485 ms	43	6,731 ms	22 % <a href="#">[actions]</a>
demo/view4	223	5,375 ms	46	6,569 ms	22 % <a href="#">[actions]</a>
demo/view1	221	5,229 ms	44	6,497 ms	24 % <a href="#">[actions]</a>
demo/view2	181	5,255 ms	36	6,451 ms	22 % <a href="#">[actions]</a>

Filter:  Apply Clear Options ▾

From the ratio of the page load duration and the transaction execution time, an **overhead** column is calculated. The "Samples" column shows you how many measurements were taken, while the "Transactions" column contains the total number of transactions. Their ratio should correspond to the sampling percentage above. It can be slightly lower due to incomplete page loads and browsers that do not support the Web Performance API. To get better statistics, switch the time range to a longer interval.

## Using the servlet filter

The servlet filter that injects the monitoring Javascript snippet into HTML pages is contained in the perfino API that is located in the JAR file `api/perfino_api.jar`. See the javadoc overview for how to download this JAR file with Maven, Ivy or Gradle.

To enable end user experience monitoring, you have to add a filter definition to your `web.xml` file:

```
<filter>
  <filter-name>euem</filter-name>
  <filter-class>com.perfino.filter.EndUserFilter</filter-class>
</filter>
```

In addition, you need at least one filter mapping that passes your HTML pages through this filter:

```
<filter-mapping>
  <filter-name>euem</filter-name>
  <url-pattern>/*</url-pattern>
</filter-mapping>
```

Only HTML files with a head tag are processed, other files are left unmodified. HTML pages that are not generated as part of a business transaction are not injected.

If you have filters that compress your pages, you have to make sure to insert the `EndUserFilter` after any such filters, so that it can detect and modify the HTML content.

The Javascript snippet is loaded from and reports back to URLs below `[context path]/__perfino`. You have to make sure that URLs matching `/__perfino/**` are passed to the filter. In the above example configuration that would be the case. If the filter mapping only handles selected URLs, the following additional filter mapping has to be added as the first mapping:

```
<filter-mapping>
  <filter-name>euem</filter-name>
  <url-pattern>/__perfino/*</url-pattern>
</filter-mapping>
```

## Cached pages

If you use a filter to cache certain HTML pages, you might want to disable end user experience monitoring for those pages by adjusting the URL pattern for the `EndUserFilter`. If you really want to monitor cached pages, make sure that the sampling percentage for end user experience monitoring is set to 100%, otherwise not all cached pages will include the Javascript snippet.

Since the load time of the page will not be correlated to the transaction time on the server in that case, the average overhead may become negative. In that case, the perfino UI shows `[Cached]` instead of the negative percentage number. In the time lines for the overheads of a single transaction, negative overhead percentages due to caching are set to zero.

For cached pages, a dynamic change of the transaction naming configuration will not propagate to the browser. perfino uses a hashing technique to detect pages that were generated with a different configuration and ignores the reported load times.

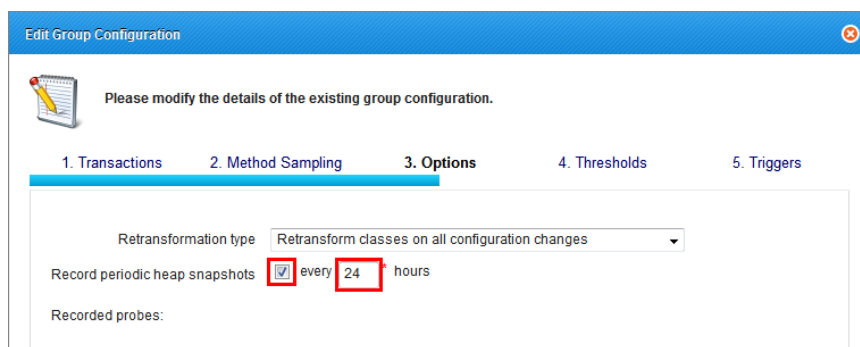
# Memory


Memory analysis in perfino is always on a per-VM level. Two kinds of memory snapshots can be taken and analyzed: Low-overhead memory snapshots and HPROF snapshots.

## Memory snapshots


Low overhead memory monitoring is only possible at a "shallow" level. The JVM knows about loaded classes and instance counts and exposes this information. There is no kind of reference information available, which is important for solving memory leaks. Nevertheless, classes and instance counts give you a good initial overview of the memory consumption inside a JVM. For historical comparisons [p. 74], plotting the instance count of a particular class against time often shows important trends.


To get a basis for historical comparisons, memory snapshots are recorded once per day. This default can be changed on the "Recording" step of the VM group configuration. You can change the period for memory snapshot recording or disable it altogether.





When you open the "Memory" view in the "VM data views", you are required to **select a single VM or a VM pool** with the VM selector in the top-left corner. After you make this selection, the navigation controls in the top right corner are enabled and you can move to the most recent snapshot with the  Previous button.

[no snapshot selected]

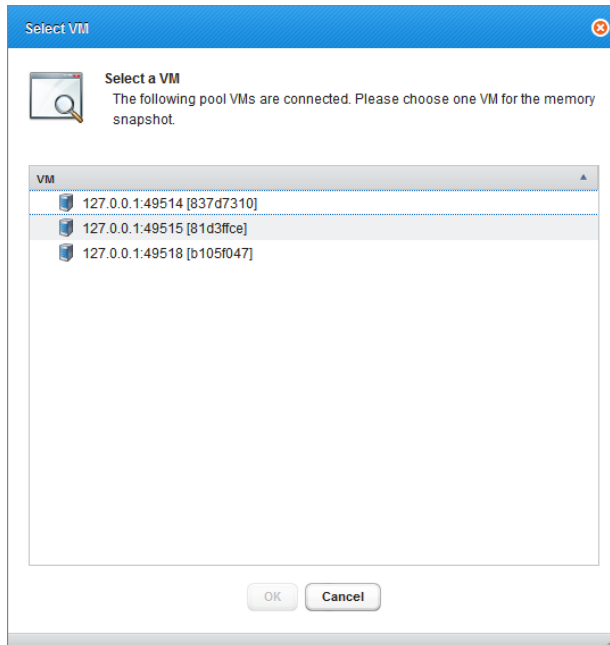


If you select a **VM pool**, you must  select a specific snapshot. In this way, you also narrow your focus to a particular VM from the VM pool and the VM selection will be set to that VM for all VM data views. To move back to the entire pool, simply click on the "Move one level up" link below the VM selector.



The  record button takes a new snapshot and displays it immediately. Since periodic snapshots are taken quite rarely, this strategy may be necessary if you want to analyze a current situation.

If you have selected a **VM pool**, perfino asks you for which VM the snapshot should be taken. Just as when selecting an existing snapshot for a VM pool, the VM selection is then changed to that particular VM.




A memory snapshot is a list of loaded classes together with their instance counts and shallow sizes. "Shallow" size means that only the direct size of the object is counted. This includes pointer sizes for referenced objects, but not the referenced objects themselves. Any other size calculations like deep size or retained size would require information about references which are not available with this low-overhead technique.

Aggregation: Classes Snapshot at: 08-May-2014 00:00:00

CLASS NAME	INSTANCE COUNT	SHALLOW SIZE
char[]	14399	1,470 kB <a href="#">[actions]</a>
java.lang.String	13786	330 kB <a href="#">[actions]</a>
java.util.HashMap\$Entry	6627	212 kB <a href="#">[actions]</a>
byte[]	5017	1,536 kB <a href="#">[actions]</a>
int[]	4983	251 kB <a href="#">[actions]</a>
short[]	4865	287 kB <a href="#">[actions]</a>
java.lang.Class	3467	417 kB <a href="#">[actions]</a>
com.perfino.agent.e.c.a	2032	81,280 bytes <a href="#">[actions]</a>
java.util.LinkedHashMap\$Entry	1806	72,240 bytes <a href="#">[actions]</a>
java.util.HashMap	1684	80,832 bytes <a href="#">[actions]</a>
java.util.HashMap\$Entry[]	1644	166 kB <a href="#">[actions]</a>
java.lang.Object[]	1497	87,392 bytes <a href="#">[actions]</a>
java.lang.Object	1182	18,912 bytes <a href="#">[actions]</a>
int[]	1076	95,256 bytes <a href="#">[actions]</a>

Filter:  Apply Clear Options

The snapshot selector shows you all memory snapshots that are stored in perfino for the current VM selection. This is also the place to delete snapshots.



## HPROF snapshots


Advanced memory analysis that can give answers to questions like

- What is the cause for the memory leak in my application?
- My application needs a lot of memory, how can I reduce memory consumption?
- Where does this class loader leak come from?

require a full heap dump.

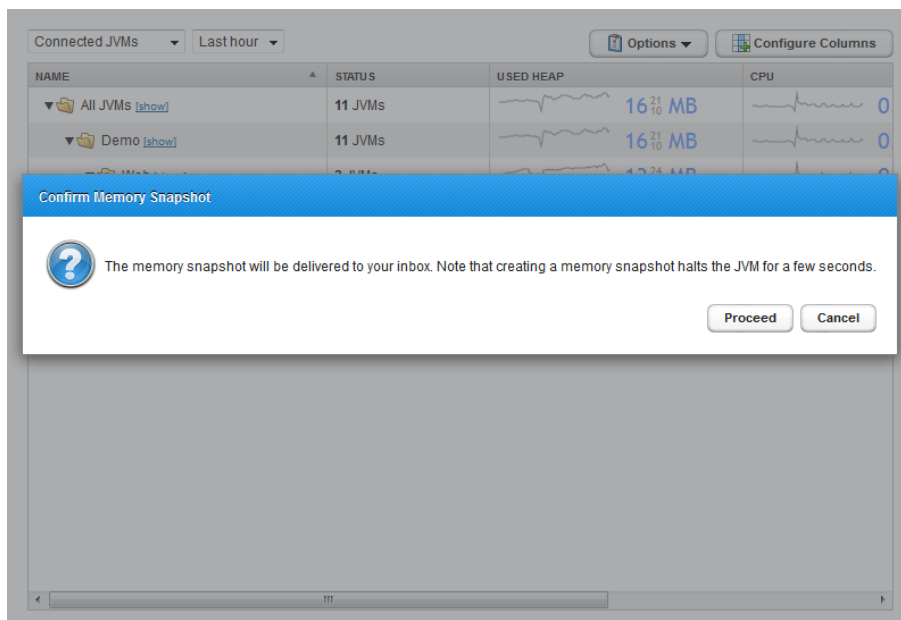
The JVM can save memory snapshots directly in a format called HPROF without the help of native JMVTI profiling agents. Saving HPROF snapshots is a low-risk operation, but during the time that such a snapshot is saved, the JVM is slowed down or temporarily halted.

You can save HPROF snapshots with triggers [\[p. 58\]](#) or manually in the VMs view, by clicking on the "actions" link next to a VM and selecting "Save a memory snapshot".



The snapshot will be delivered to your inbox when it is ready. This can take a few seconds. If the snapshot is large and the network connection between the perfino collector and the monitored VM has low bandwidth, it can also take several minutes. From the inbox, you can download the snapshot and load it with JProfiler or other Java profilers.

Note that if a sampling operation started with the "Record fine-grained CPU data in profiling mode" is currently in progress, the heap dump will only be taken after that operation has completed.






Snapshots are stored indefinitely until you delete them from the "Snapshots" view. The inbox item just notifies you about a new snapshot, deleting the inbox item will not delete the snapshot itself. In the "Snapshots" view you can select the VM group or VM to show only snapshots that have been taken for the selected VMs.

## Heap telemetry

Information about free heap and used heap is always available with a fine time resolution. As such, they are an ideal basis for trigger conditions that save HPROF snapshots in case of low memory.



To do that, define a **threshold** for the "Free heap" telemetry with a specified lower bound in the recording settings. In the trigger settings, add a **threshold trigger** for that threshold and add a "Save a memory snapshot" action to the trigger.

perfino only takes a memory snapshot for the last VM where a threshold violation was detected, not for all such VMs.

## Historical Comparisons

perfino stores data for long period of times, in some highly aggregated forms even indefinitely. Apart from telemetries where the main content of the view already is a historical comparison by itself, the VM data views show data from a certain time interval and let you move back and forth to adjacent intervals or select arbitrary intervals in the past.

In addition to looking at this historical data to analyze it by itself, many VM data views offer facilities to make historical comparisons. There are two kinds of comparisons: Comparing all the content from two selected points in time and comparing one scalar value for many points in time.

### Content comparisons

Content comparisons are available in the views that show transactions as well as the probe hot spots views. To start a comparison, click on the "Compare" link in the top right corner. This immediately shows you a comparison between the following intervals:

- **Second interval**

The second interval is the interval you were looking at before clicking on the "Compare" link. You can adjust the second interval with the same navigation controls as before.

- **First interval**

The first interval is set to the interval just preceding the second interval. New navigation buttons are shown that let you adjust the first interval as required. With the date and time chooser you can jump back to arbitrary points in time.

Policies: Split

Show 1 hour up to 10/04/14 16:00

Compare with: 10/04/14 15:00

HOT SPOT	TOTAL TIME	TOT. TIME DIFF.	INVOCATIONS	INV. DIFF.	AVG. TIME
Inventory checks via RMI	231 m	+231 m	7799	+7,799	1,783 ms
Remote demo transaction	183 m	+1,463 s	12471	+1,611	882 ms
Nested demo transaction	164 m	-2,235 s	7727	-1,782	1,275 ms
Exchange rate checks via web service	5,697 s	+5,697 s	19213	+19,213	296 ms
Demo JDBC Job	4,242 s	-396 s	836	-218	5,074 ms
Exchange rate to EUR	2,688 s	+1,055 s	35853	+14,134	74,994 μs
Inventory check via RMI	1,065 s	-163 m	1057	-9,801	1,008 ms
Exchange rate to GBP	899 s	+899 s	11965	+11,965	75,144 μs
docs/invoiceService	835 s	-116 s	1065	-130	784 ms
processOrder	797 s	-10,855 ms	1098	-57	726 ms
docs/quoteService	692 s	-39,579 ms	1154	-46	608 ms

Filter: [ ] Apply Clear Options

Transaction data is aggregated, so as you move the first interval back in time, you will hit the limit of stored data. At that point you have to switch to a larger resolution for which data is retained longer. The available display intervals with their retention times are:

Display interval	Retention time
1 minute	24 hours
10 minutes	24 hours
1 hour	90 days
1 day	unlimited

In a comparison view, all measurements that are shown in the regular view get an additional column that shows the difference between the second and the first interval. In addition, the measurement columns are shown with a **difference bar** in the background:

- If the value has **increased**, the total length of the bar corresponds to the value in the second interval, with the increase shown in **green**.
- If the value has **decreased**, the total length of the bar corresponds to the value in the first interval, with the decrease shown in **red**.
- The **unchanged fraction** is shown in **gray**.

## Timelines

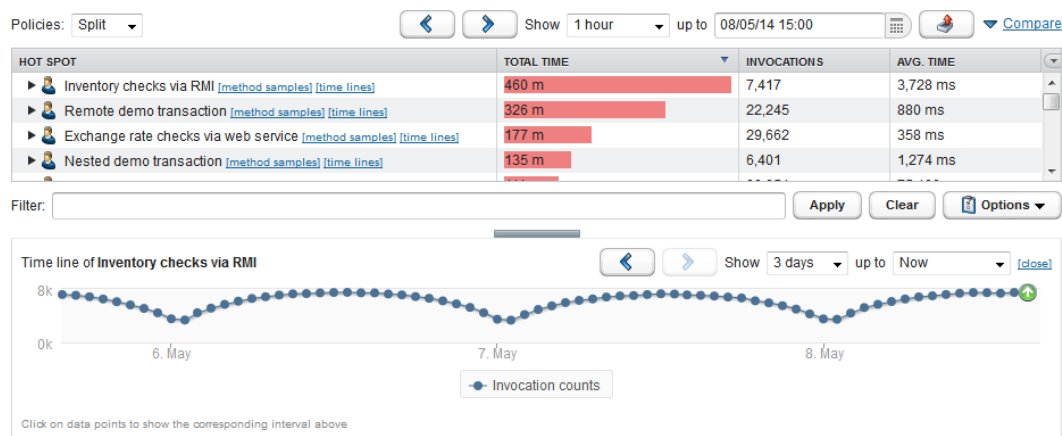
If you click on the "timelines" link next to an element in any of the transaction views, you can choose which column should be plotted over time.

Policies: Split ⏪ ⏩ Show 1 hour up to 08/05/14 15:00 🔍 🔥 [Compare](#)

HOT SPOT	TOTAL TIME	INVOCAIONS	AVG. TIME
▶ <a href="#">Inventory checks via RMI [method samples] [time lines]</a>	460 m	7,417	3,728 ms
▶ <a href="#">Remote demo transaction [method sample]</a>	Show a time line of invocation counts	22,245	880 ms
▶ <a href="#">Exchange rate checks via web service [m]</a>	Show a time line of average times	29,662	358 ms
▶ <a href="#">Nested demo transaction [method samples]</a>	Show a time line of total times	6,401	1,274 ms
▶ <a href="#">Exchange rate to EUR [method samples] [time lines]</a>	111 m	89,051	75,100 µs
▶ <a href="#">Demo JDBC Job [method samples] [time lines]</a>	4,564 s	959	4,759 ms
▶ <a href="#">docs/invoiceService [method samples] [time lines]</a>	1,008 s	1,268	795 ms
▶ <a href="#">processOrder [method samples] [time lines]</a>	861 s	1,237	696 ms
▶ <a href="#">docs/quoteService [method samples] [time lines]</a>	769 s	1,279	602 ms
▶ <a href="#">esb/sendMessage [method samples] [time lines]</a>	530 s	1,315	403 ms
▶ <a href="#">data/load [method samples] [time lines]</a>	372 s	1,232	302 ms
▶ <a href="#">rest/delete [method samples] [time lines]</a>	369 s	1,248	296 ms
▶ <a href="#">rest/put [method samples] [time lines]</a>	251 s	1,254	200 ms

Filter:  Apply Clear Options ▼

The timeline will be opened in a split panel below the view.

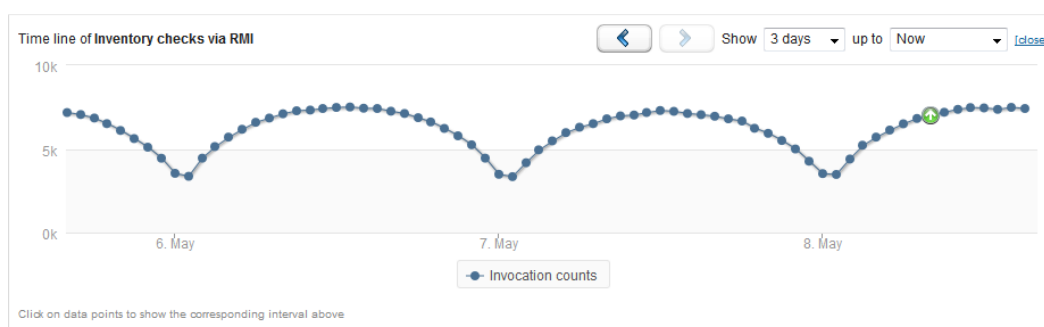


Time line data is not available for the one minute resolution, so the time line links are not shown for that display interval. For other display intervals, the data points in the time line initially match the display interval in the data view, so that the numbers are comparable. The timeline has navigation buttons itself which take you to different intervals of the entire timeline, just like for telemetry views.

The following display intervals in the data view and the time line are compatible with respect to the measured intervals:

Data view	Time line
10 minutes	3 hours
1 hour	3 days
1 day	30 days

To get the chronological context for the data displayed above, look for the up-arrow icon in the timeline. You can navigate to different points in time by clicking on them.



If you change the display interval of the time line and then click on a data point in the time line, the resolution of the data view will be adjusted to be compatible with the time line. Then, the current time can be marked.

If you select a new timeline in the data view, the old timeline will be replaced. Any change in the parameters of the data view will automatically close the timeline.

## Cross-over To Profiling

There are several reasons why you should not have a profiler running in production at all times. For one, there is the overhead that may be unacceptably high depending on the profiling settings. Profilers are geared towards maximizing the extraction of useful information with no explicit guarantees as for the incurred overhead. Also, the use of the native profiling interface of the JVM (JVMTI) is something that is an additional risk in a production environment. Depending on whether you use non-standard garbage collectors or other JVM tuning options, there may be stability concerns, since the JVM is not tested as extensively with JVMTI as it is without.

However, the best defense against performance problems is defence in depth. Sometimes there are situations that require more information than what can be obtained from the low overhead monitoring and sampling techniques that are available in perfino.


### Recording CPU snapshots

For these cases, perfino offers a **full sampling mode that loads a native JVMTI library** that is optimized for in-production use. The result of this operation is a JProfiler snapshot file that can be downloaded from perfino and opened in JProfiler - similarly to the HPROF memory snapshot files [p. 69] that are saved directly by the JVM.

This profiling mode is only available if the correct native libraries are present in the `lib/[platform]` directory next to the `perfino.jar` file that was used in the `-javaagent` VM parameter. If you download the agent [p. 11] from the perfino UI, native libraries for all supported platforms are already included. If you copy the agent yourself from the perfino installation directory, copy the `agent/lib` directory along with it.


Once the JVMTI has been turned on, it cannot be turned off again. As long as the JVM is running, it will remain in this state. Typically, the overhead of the JVMTI without any data recording is less than 1%, though.

To take a CPU snapshot of a particular VM, go to the VMs view, click on the "actions" link next to the VM and select "Record fine-grained CPU data in profiling mode".




The duration of CPU recording can be configured. perfino monitoring is not impacted by CPU sampling. Also, you can specify if the call tree should be split for each transaction or not. This is

a top-level split and the sampling call tree is appended to the perfino transaction tree (see below). Whether this split is beneficial or not depends on your transaction definitions and what kind of problem you want to find. The overhead for sampling with transaction splitting is higher than without.



The "Record Data" dialog box has a blue title bar and a red close button. It contains a "Configure recording options" section with a red dot icon and a note: "A JProfiler snapshot is delivered to your inbox after data recording has completed." Below this are two progress indicators: "1. Recording" (active) and "2. Finished". The "Recording" section includes a "Recording time" field set to "1" with a "Minutes" dropdown. Under "During recording:", there is a bulleted list: "there will be no progress information", "the profiled VM will be specially marked as being profiled", and "no data will be shown in perfino for the profiled VM". A checkbox "Split the call tree for each transaction" is present. A yellow warning box states: "Warning: This action will switch on the Jvmti profiling interface in one of the VMs that caused the trigger to fire. This carries an inherent stability risk and is only intended for solving immediate problems." At the bottom are "Cancel", "Back", "Next", and "Finish" buttons.

To take snapshots automatically if the CPU load is too high, you can set up a threshold for the CPU telemetry and configure a trigger that includes the "Record fine-grained CPU data in profiling mode" trigger action.



The "Create Threshold Violation Trigger" dialog box has a blue title bar and a red close button. It includes a document icon and a plus sign. The text says: "Please enter the details of the new trigger. Execute a list of actions when a configured threshold is violated for a specified number of times." The "Threshold" section shows "CPU" selected, with a "Choose" button. The "Fire after" field is set to "10" with "events in one" and a "hour" dropdown. The "Inhibition time" field is set to "12" with "hours" dropdown. The "Trigger actions" section has a table with one row: "Record fine-grained CPU data in profiling mode [1 minute]". To the right of the table are icons for adding (+), deleting (X), and reordering (arrows). At the bottom are "OK" and "Cancel" buttons.

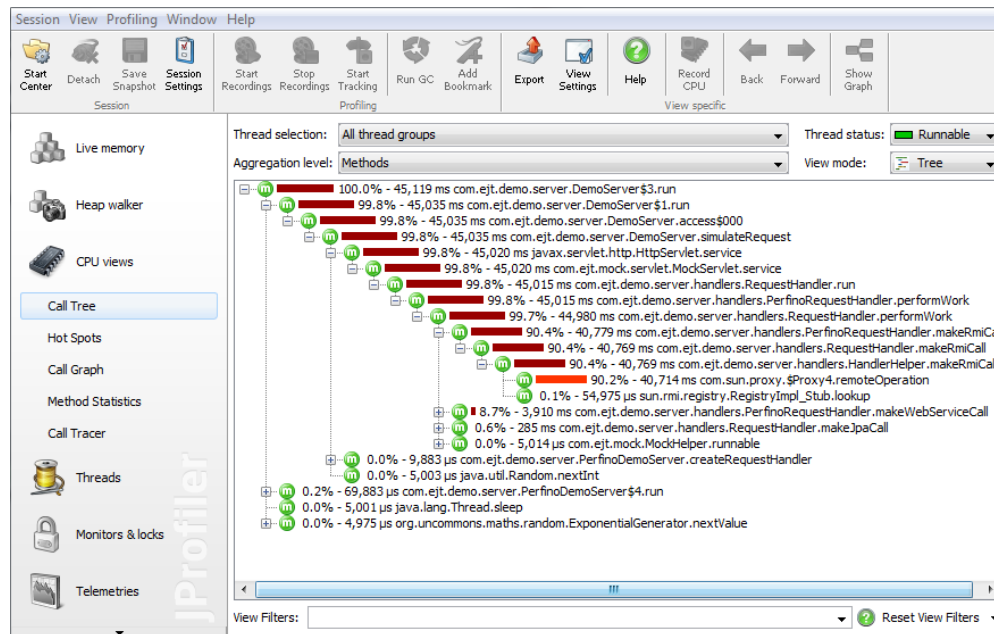
Note that only the last offending VM will be profiled, not all of them.

If you use a custom library for network I/O, you can add selected methods to the net I/O thread state [\[p. 100\]](#), so that you get useful hot spots for your analysis.

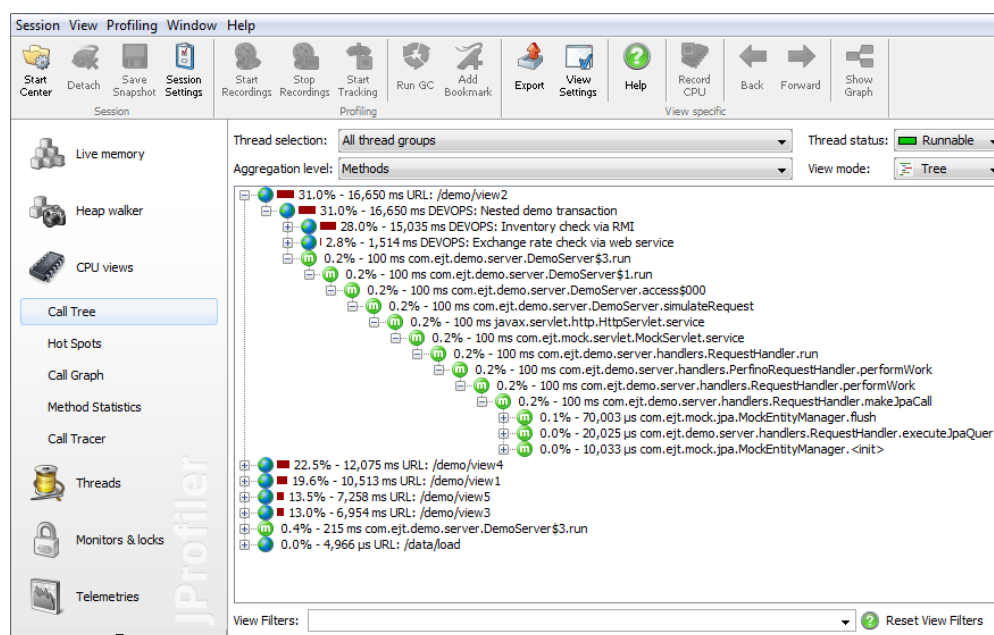
## Viewing CPU snapshots in JProfiler

CPU snapshots are delivered to your perfino inbox from where you can download them. The file extension is ".jps" which stands for "JProfiler snapshot". If you have JProfiler installed, you can double-click on the snapshot file to open it, or choose **Session -> Open Snapshot** from JProfiler's main menu.

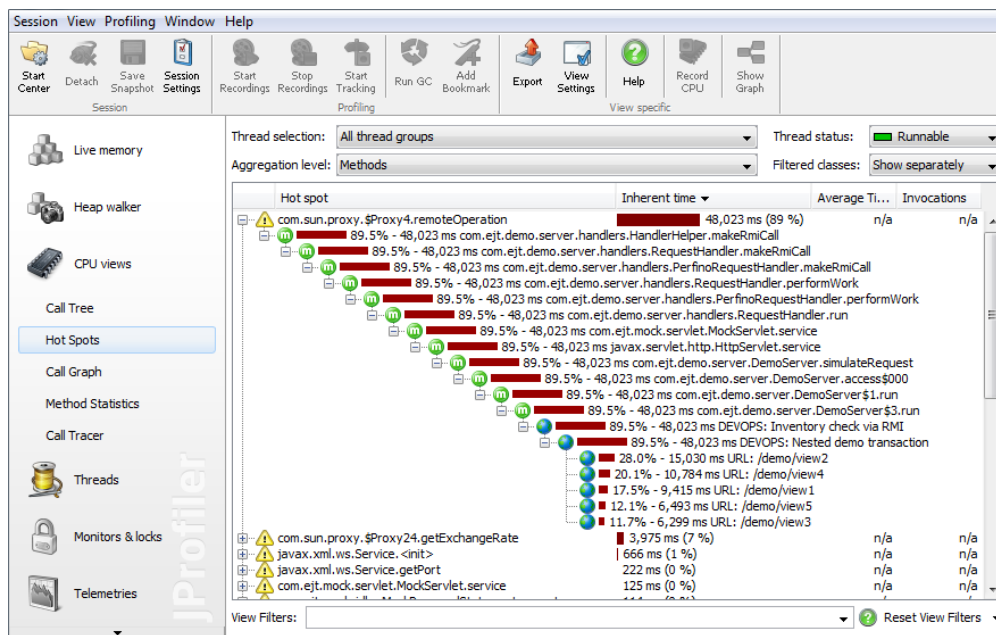
CPU snapshots taken by perfino only contain data for the call tree, hot spots and call graph views as well as the thread history and thread monitor views. Other views are disabled when you open a perfino CPU snapshot.



If transaction splitting was enabled for CPU recording, the perfino call tree is shown at the top and the sampling data is attached as appropriate.

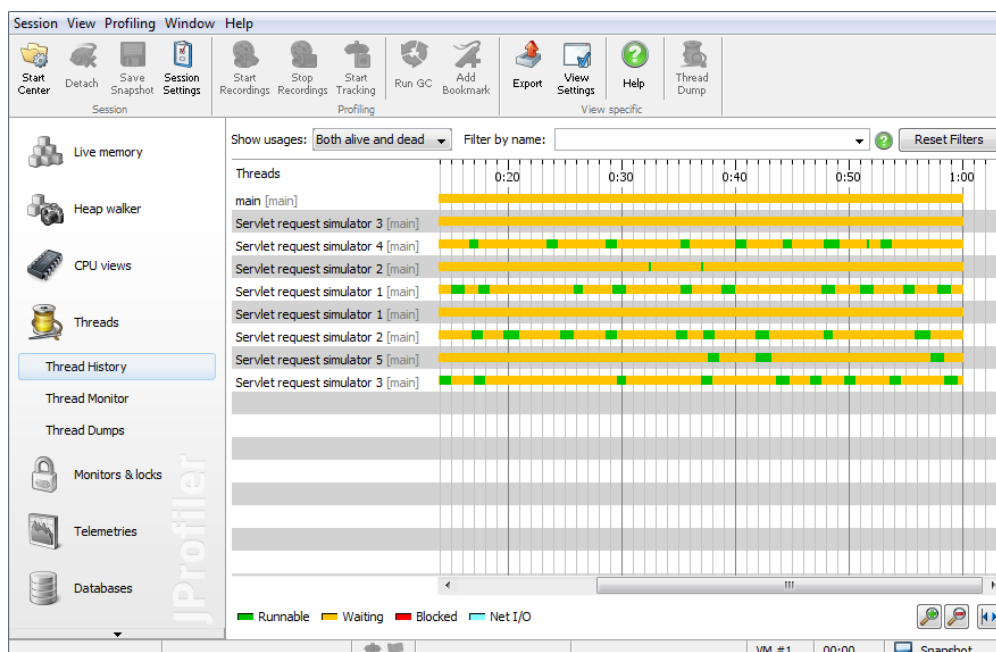


The back traces in the hot spots view also show the perfino transactions in this case.



The attribution of sampling data to transactions is not exact, but only an estimation based on the analysis of subsequent sampling call stacks. For very short running transactions (< 5ms), sampling data may be attributed to another transaction or even be placed outside any transaction.

In the thread history view, you can check when threads were created or terminated and what kind of activity they performed over their recorded lifetime. Only green areas indicate times when the thread was actually eligible to execute methods. Other thread statuses designate idle time while waiting or blocking on monitors or reading or writing to network sockets.






## Full profiling

Sometimes performance or memory problems are so tricky that you need the entire arsenal of recording techniques that a profiler has to offer. For example, the capabilities of the heap walker for live heap snapshots can be instrumental for solving a memory leak or the CPU and probe profiling features may be required for understanding a performance problem. Questions involving threading issues, monitors and locks can only be solved in JProfiler and not in perfino.

perfino helps you to cross over to full profiling with the minimum amount of intrusion into your production environment. If a JProfiler installation can be found, perfino can load its agent into a monitored VM and prepare it for a connection from the JProfiler GUI. Use the "Detach VM for JProfiler" action in the VMs view to initiate this process. The VM will be detached from perfino and has to be restarted after profiling to reconnect to perfino.



To take advantage of this integration, you have to [download the JProfiler archive](#) and extract it in one of the following directories:

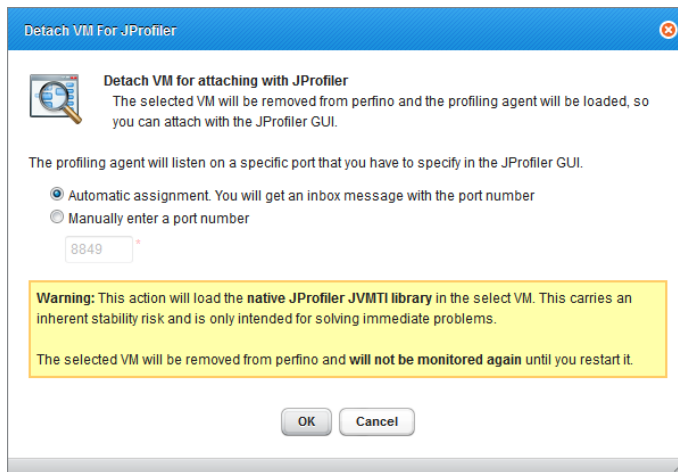
- The directory where the `perfino.jar` file has been installed
- The directory `$HOME/.perfino` on Linux/Unix or `%USERPROFILE%\.perfino` on Windows, for the user who is running the monitored VM.

The JProfiler archives from the download page contain a single top-level directory named `jprofiler[major version]`. The perfino agent sequentially looks for such directories in the above locations and uses the directory with the highest version number. Any version of JProfiler starting from 8.0.6 can be used for this integration.

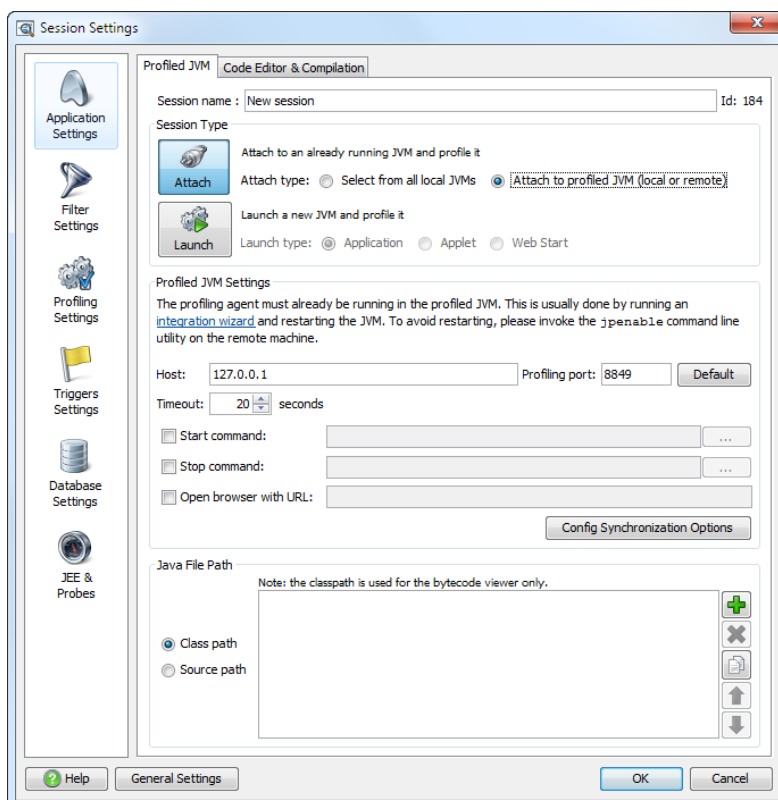
If none of the standard directories above are suitable or if you already have a JProfiler installation elsewhere, you can append the option `,jprofiler=[JProfiler installation directory]` to the VM parameter for monitoring [p. 11]. If the appropriate native library file can be found in that JProfiler installation, it will be used preferentially, otherwise the perfino agent will continue to look in the standard directories.

If no JProfiler installation can be found, specific instructions for integrating JProfiler are given. You can download and extract JProfiler on the machine where the monitored VM is running and execute the "Detach VM for JProfiler" action again, there is no need to restart the monitored VM.

If JProfiler was found, a confirmation dialog is shown, where you get the chance to configure the profiling port to which the JProfiler GUI will attach. You can either enter the desired port number yourself or let perfino find a free port for you. In the latter case, an inbox message will be sent to you with the actual port number.



In JProfiler, create a new session of type "Attach to profiled JVM (local or remote)" and enter the IP address or host name where the VM is running as well as the port on which the profiling agent is listening.



perfino will stop monitoring the VM when you prepare it for profiling. To monitor it again, you have to restart the JVM.

# A Configuration

## A.1 Server Configuration

perfino configuration options that cannot be changed in the perfino UI are contained in the text file `perfino.properties` in the perfino installation directory. You can either edit that text file in a text editor or use the `configure` executable in the perfino installation directory. The latter presents an organized view of all properties, saves the file even if elevated privileges are required and can restart the perfino server to apply your changes.

The properties themselves are documented with comments in `perfino.properties`. Here, a couple of scenarios are discussed where it is necessary to adjust the default parameters.

### Data directory

The **dataDirectory** property points to the directory where all variable data is located. The following subdirectories are created by perfino:

- **db**  
Contains the embedded H2 database. If an external mysql database is configured, this directory is not used.
- **log**  
Contains all log files. By default, log files are rotated. The rotation settings can be changed in the `log4j.properties` file in the perfino installation directory. Different settings can be applied to the three different log files, "server", "connection" and "event". All log files can be viewed in the perfino UI.
- **snapshots**  
Memory snapshots and profiling snapshots are saved in this directory. These snapshots can be downloaded or deleted in the "Snapshots" view in perfino.
- **ssl**  
This directory contains the file pair `agent.keystore` / `server.keystore` for authentication and encryption as well as the SSL certificate for the web server.

If you run two perfino servers on the same machine, they have to have different data directories. By default the installer always suggests the same location for the data directory, so in the case of multiple installations you have to adjust it in the installer or after the installation in the `perfino.properties` file.

### Web server

perfino comes with a built-in web server that listens on port 8020 by default. You can adjust that port with the **httpPort** property and switch to HTTPS by setting **useHttps** to `true`.

When you use HTTPS, perfino will generate a self-signed certificate `ssl/self_signed.keystore` in the perfino data directory. Browsers will display warning messages with this certificate. If you have a certificate that is signed by a recognized certificate authority, you can copy it in PKCS12 format to `ssl/web.pkcs12`. If the certificate file has a different name, you can specify the **keystoreName** property.

If the certificate is protected with a password, you can specify it in the **keystorePassword** property. While the password cannot be encrypted, it can at least be obfuscated with the command line tool `perfino_obfuscate`:

```
perfino_obfuscate [password]
```

If you put perfino behind a reverse proxy, you have to set the **reverseProxy** property to `true`. The web server will then analyze the proxy headers to create correct URLs. If this should not work due to a problem with the reverse proxy, set the **reverseProxyHost** to the host name of the proxy.

If you have infrastructure that can check the health of a web server by making an HTTP GET request, you can set the **healthCheckPort** to a non-zero value in order to create such an HTTP port in perfino. Any HTTP request to that port will return a document with HTML mime type and the text "Alive". For example, Amazon Web Services provides a health check service that is used by Route 53 to determine if an IP address can be routed to or not.

### Communication with monitored VMs

Monitored VMs create a TCP connection on the port that is configured with the **vmPort** property.

By default, the communication between monitored VMs and the perfino collector is unencrypted and unauthenticated. This means that every VM can connect to the perfino server and the perfino agent has no way of knowing if the perfino server on the other side can be trusted. This can be acceptable in certain local area networks but it is not suitable for connecting over WANs or even over the internet.

To enable authentication and encryption set the property **vmUseSsl** to `true`. In that case, the file pair `ssl/agent.ks` and `ssl/server.ks` will be created. Now, the server will only allow connections from agents who possess the `agent.ks` file and agents will only connect to servers who have the `server.ks` file. In addition the communication protocol will be encrypted.

For more information on this topic, see the chapter on monitoring JVMs [\[p. 11\]](#).

### Remote perfino UI

Running the perfino UI on a different server than the collector can have two purposes. First, it allows you to split the server load of the UI and the collector to separate machines which is a good idea if you have many users. Second, some network topologies require that the collector runs in one and the perfino UI in another network. For example, if the collector runs in an internal network that is protected by a dual firewall, and the perfino UI should be available to the outside, the perfino UI has to run on a machine in the DMZ.

The first change you have to make is to set the **startRemoteServer** property to `true`. Then, the perfino server will listen for remote perfino UIs on the port configured with the **remoteServerPort** property.

The perfino UI can be deployed as a WAR file to a servlet container like Tomcat or Jetty. The WAR file has to be generated with the `deploy` tool in the `deploy` directory of your perfino installation. It will create the file `perfino-standalone.war` in the same directory.

In the application server, define the JNDI property **perfino/server** with a value of `"server name"[:port]` where "server name" is the name or IP address of the server where the perfino collector server is running. The port is optional and defaults to 1099 unless you have configured the **remoteServerPort** property differently. Then, deploy the generated WAR file into your application server. The perfino web application will make an RMI connection to the configured data collection server automatically.

### External MySQL database

For your convenience, perfino ships with an embedded H2 database. No separate process has to be started for the database and no manual database configuration is necessary. In production, you might want to use the external MySQL database option instead. MySQL has an extensive tool set that allows hot backup and replication as well as a multitude of database viewers.

To switch perfino to a MySQL database, set the **dbType** property to `mysql`. The **externalDbUrl** determines the name of the database and where the MySQL database is located. By default, it is set to `jdbc:mysql://localhost/perfino`, so it connects to a MySQL installation on the local machine and selects a database with the name "perfino".

If you do not want to enter the password in plain text, you can obfuscate it with the command line tool `perfino_obfuscate`:

```
perfino_obfuscate [password]
```


When you configure the mysql database in the installer or the `configure` tool, the password will always be obfuscated.

In MySQL, you have to create that database first and grant all privileges to the user that is specified with the properties **externalDbUser** and **externalDbPassword**. The following statements, executed by the root user, create a database named "perfino" and a user named "perfino\_user" with a password of "apmrocks":

```
CREATE DATABASE perfino;  
GRANT ALL ON perfino.* TO 'perfino_user'@'localhost' IDENTIFIED BY 'apmrocks';  
FLUSH PRIVILEGES;
```

## A.2 Server Administration

Many server options are contained in the file `perfino.properties` in the installation directory and can only be changed when the server is restarted [p. 83]. The server administration settings that can be changed in the perfino UI are contained in the "General Settings". These settings include user configuration, license keys and data consolidation options.



### Access control

There are three different access levels in perfino: admin, profiler and viewer. With these access levels you can implement a safe usage policy for the perfino UI in a larger organization.

- **Admin**

When you install perfino, you have to create an admin user. Only an admin user can open the "General Settings" and make modifications to it.


- **Profiler**

With the "Profiler" access level, you can assign full control for a particular VM group, including the ability to put the JVM in profiling mode.

First, you create the VM group in the recording settings, if it does not already exist. When creating the user, set its access level to "Profiler" and add the VM group to the list of access rights, removing the entry for "All VM groups".

This user will now be able to change the recording settings for this VM group and all VM groups that are contained within it. In addition, heap dumps and CPU snapshots in profiling mode can be taken for the associated VMs.

Viewing rights are not restricted in perfino.



**Create User**

Please enter the details of the new user.

User Name\*

Full Name

Email

Password\*

Confirm Password\*

☐ Must change on next login

Configuration access level: **Profiler** ▼

Access rights for JVM groups:

GROUP
DemoWeb

OK Cancel

- **Viewer**


To give somebody access to monitoring data without the ability to change any monitoring settings, use this access level.

## License keys

As an administrator, you can add or remove license keys from perfino.

If more VMs are monitored than licenses are available, then some VMs will be marked as "unlicensed" and no monitoring data will be available for those VMs. When you add a new license key, perfino will start monitoring additional VMs right away, without the need for a restart.

The license key entry dialog can take multiple license keys, with each license key on a separate line. One of the license keys in perfino has to be a product key, otherwise no VMs can be monitored. License keys for additional VMs increase the maximum number of monitored VMs.



**Add License Keys**

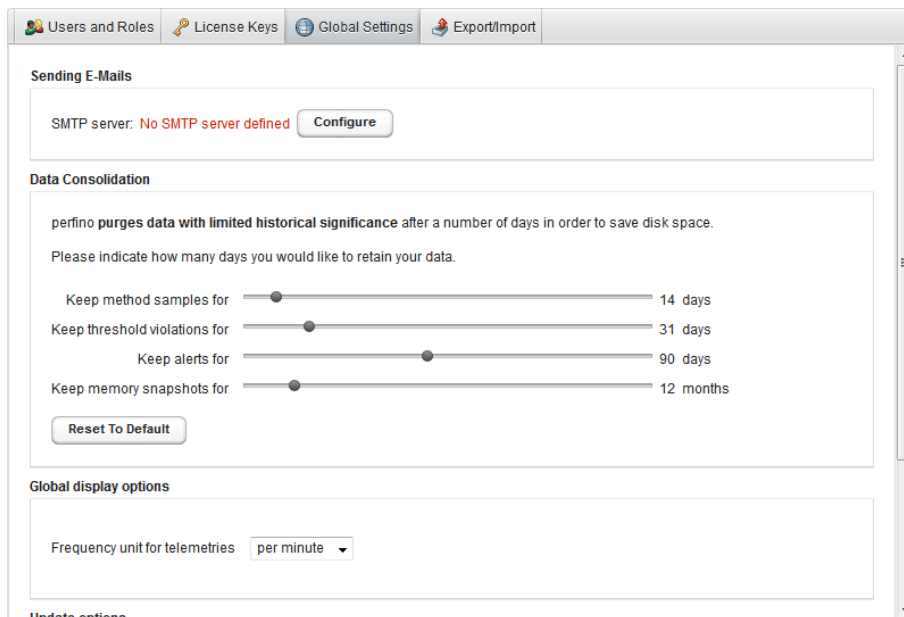
Please enter one or more license keys.  
License keys should be entered on separate lines (press SHIFT-RETURN to start a new line) or separated by spaces.

License keys: \*

OK Cancel

## Global settings

In the global settings, you can configure an SMTP server that is used for sending emails. Emails are only sent by "Send email" trigger actions [\[p. 58\]](#) . If no SMTP server is defined, these actions do not have any effect.



The data consolidation options determine how long certain types of data are kept in the database.

Transaction data [p. 74] and telemetries [p. 51] are consolidated automatically, the highest aggregation levels remain in the database indefinitely. Due to the large interval size, the linear increase in required storage space is small.

Detailed data like method sampling takes a lot of space and is of decreasing interest the farther back you go. After some time, it can be deleted. The higher the retention times are set, the more disk space will be used by the database.

With the "frequency unit for telemetries" option you can make the numbers in telemetries easier to interpret. It should be set to a value that fits with the typical throughput in your applications. If your monitored VMs handle many transactions per second, you can set it to "per second", if the transaction frequency is more on the order of minutes, the setting "per hour" will be appropriate.

The global settings also show the currently installed version of perfino and let you configure automatic update checking. If selected, perfino will check for updates within in the current major version once a day. If an update is detected, an inbox message will be sent to all admin users. The inbox message contains the new version and hyperlinks for change log and download. An update notification will only be shown once even if perfino is restarted. When another version is released, a new update notification will be sent, even if you have not updated perfino in the meantime.

Whether you have enabled or disabled automatic update checking, the manual **[Check For Updates Now]** button performs the same check live and shows you its results in a dialog right away.



## A.3 Import And Export

There are several kinds of motivations for exporting and importing your perfino configuration:

- **Backup**

Backing up your perfino server configuration from time to time is a good idea. The configuration contains a lot of knowledge about your business transactions and their expected behavior. It is a valuable piece of data that could be stored in a version control system.

- **Staging**

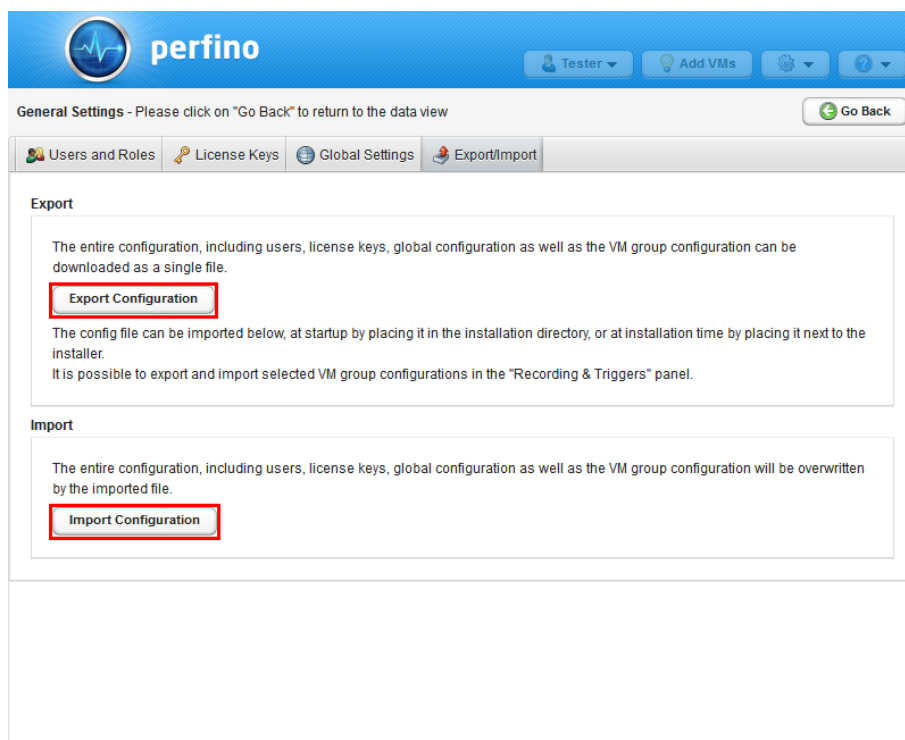
If you have a staging environment where you try out changes in the perfino configuration before deploying them to a production server, you need a way to transport the configuration from one perfino installation to another.

- **Unattended deployment**

If you deploy your perfino server to the cloud, you probably use the unattended installation mode [p. 91] of perfino and install it on a server instance where no perfino server has been installed yet. In that case you have to put the exported configuration next to the installer.

### Server configuration

If you are an administrator, you can go to the general settings and select the "Export/Import" tab. The entire configuration is exported to an XML file. When you import that file, the entire server configuration is replaced. The change does not happen until you click on **[Apply Settings]**. It is not possible to delete the current user that way.



The XML file format is always backwards compatible, a more recent perfino server can read the exported file of an older version.

Naming that file `perfino_server_config.xml` and copying it to the perfino data directory overwrites the entire configuration at startup. This mechanism is used by the unattended installation [p. 91]. After the import, the file is deleted automatically.

## Recording configuration

The exported XML file with the entire configuration contains license keys and user data. This may not be what you want for staging or also for backup purposes. In the recording settings, you can also export selected VM group configurations. Here, you can select "All VMs" or a particular VM group and click on the export button on the right.



Recording Options - Please click on "Go Back" to return to the data view

GROUP	TRANSACTION	SAMPLING	OPTIONS	THRESHOLDS	TRIGGERS
All JVMs	✓	✓	✓	0	0
▼ Demo	✓	✓	✓	2	2
Web				0	0
▼ Workers				0	1
JDBC				0	0
JMS				0	0

Use drag and drop to reorder group configurations

Group configurations inherit recording settings from their parent group unless settings are overridden. Triggers and thresholds operate on all recursively contained VMs and are not overridden.

The All JVMs group is for JVMs at the top-level. New groups can be added here or are automatically added when a JVM connects with a specified group.

This export is not only available for administrators, but also for users with "profiler" access rights.

Naming this file `perfino_recording_config.xml` and copying it to the perfino data directory overwrites the entire recording configuration at startup. The file is deleted automatically after the import has been completed.

## A.4 Unattended Installations

### Unattended mode and response files


In case you want to fully automate the installation of perfino, you can pass the argument `-q` to the installer. This makes the installer run in **unattended mode**. To set the installation directory, pass the argument `-dir [installation directory]`:

```
./perfino_unix.sh -q -dir /opt/perfino
```

To apply user input from a previous GUI or console installation, locate the response file `.install4j/response.varfile` in the installation directory and pass it to the installer with the argument `-varfile response.varfile`. The response file is a plain text file and the contained property definitions can be changed in a text editor.

### Automatic server configuration

In a cloud environment, you might want to recycle an instance or set up scripts that initialize a new instance with a completely configured perfino installation. To help you with that task, the perfino installer tries to read a number of **optional configuration files** with special names from the same directory.



Each of those files is explained in the following sections.

### Server configuration

While you could use the installer response file to adjust values in `perfino.properties`, it is much easier and more maintainable to copy a `perfino.properties` file from the data directory of a configured installation into the same directory as the installer.

The installer will use the values in that file with the highest priority. You can trim the contents of that file to the properties that deviate from the defaults, the installer will supply the default values for all other options. The basis for the structure of the actual `perfino.properties` file is the default template in the installer, so it does not matter if you delete comments or change the order of properties. Extra properties that are not present in the default template are merged in at the end.

### Monitoring configuration

The monitoring configuration includes everything you can adjust in the perfino UI. Open the general settings in a configured installation, select the "Export/Import" tab and click on "Export configuration". This will save the entire server configuration to a file [\[p. 89\]](#). If you rename that

file to `perfino_server_config.xml` and put it into the same directory as the installer, the installer will automatically apply this configuration in the new installation.

### Agent configuration

The monitoring configuration in the previous section includes the configuration for the agent. However, there is one case where you might want to supply a separate agent configuration: When the perfino agent connects to a perfino server for the first time, it receives its configuration from the server and, as a consequence, some classes have to be reinstrumented for monitoring.

If your policy is to avoid all class retransformations, you can specify that on the "Options" step of the VM group configuration. In that case, any configuration change will only be applied when the monitored VM is restarted. To avoid the need for a restart in an unattended deployment, the configuration for the agent can be imported in advance.

First, you have to open the recording settings and export the VM group configurations [\[p. 89\]](#) of interest. Then, rename the exported file to `perfino_recording_config.xml` and place it next to the installer. The installer will perform the import for agents that are running on the local machine. For other machines, you have to perform this import yourself by calling

```
java -jar perfino.jar import perfino_recording_config.xml
```

The agent extracts its config from that file and writes it in binary form to the directory

`$HOME/.perfino/config`.

This directory is read by all agents on the local machine.

Note that these steps are only necessary if you want to avoid class retransformations. Otherwise all configuration changes are applied on the fly.

### Agent and server keystores

You can encrypt and authenticate [\[p. 83\]](#) the communication between monitored VMs and the perfino server. The files `agent.ke` and `server.ke` constitute a key pair that enables both encryption as well as mutual authentication.

For an unattended deployment, you will probably already be using a particular key pair with your monitored VMs. In a configured perfino installation you can find these files in the `ssl` directory below the data directory. Placing them next to the installer ensures that they are copied to the same location in the new installation and that the server does not generate a new key pair.

### Server SSL certificate

It is recommended to use SSL [\[p. 83\]](#) to encrypt the communication between the perfino UI server and browsers. If you enable SSL during the installation, a self-signed SSL certificate is generated and saved to `ssl/web.pkcs12` in the perfino data directory. You can replace that file with a certificate that is signed by a well-known certificate authority.

If you have such a certificate, you can put it next to the installer with the name `web.pkcs12`. No self-signed certificate will be generated in that case.

## B Advanced Topics

### B.1 Annotation Transactions

Many frameworks use annotations for designating important classes and entry points that could be turned into transactions in perfino. With the "Annotated invocations" transaction type you can let perfino do this work for you.

#### Method and class annotations

An annotation transaction definition takes the fully qualified class name of the annotation that you are interested in. Annotations can be either used on **classes** or on **methods**. Before proceeding with the configuration, you have to tell perfino for which target the selected annotation is used.

For an annotated method, each method invocation becomes a transaction in perfino. Annotations on classes create transactions for all calls into **public instance methods**.

The second choice is whether derived classes should be considered as well. For methods, it can be important to capture the time in an overridden method. For example, if a framework creates an implementing proxy class and overrides the annotated method in order to add database transactions, you want to intercept the method in the proxy class, and not only the annotated method.

For annotated classes, the "Method selection" options determine which methods are selected to generate transactions:

- **Implementing methods only**

If the annotation is placed on an interface that already defines all operations of interest with its methods, you should select this option.

- **All public methods**


If the derived classes have their own methods of interest that do not implement or override the methods in the annotated class, use this option to create transactions from all public methods.

## Naming

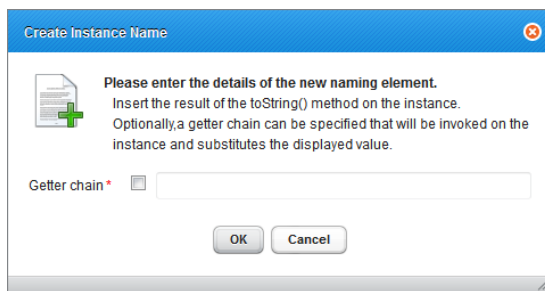
When you use annotations on classes with the setting for intercepting public methods in derived classes, the question remains what class name should be used in the transaction naming.

By default, a naming element of type "Class name" adds the name of the class where the intercepted method is defined. If you want to add the name of the annotated super class instead, select the "Use annotated class name for filter and naming" check box in the "Annotation" step of the wizard.

If you want to go the other way and add the class name of the actual instance on which the method was called, use the "Instance class name" naming element instead.



Even more specific, the "Instance name" naming element adds the `toString()` invocation on the object where the transaction method was called.



This naming element can also apply a getter chain to this object and append the `toString()` invocation on the result to the transaction name instead. For example, if the instance class has a `getVerbose()` method that returns the desired text, set the getter chain to `getVerbose()`. You can mix public field accesses and parameter-less invocations of public methods like this:

```
getParent().descriptor.getVerbose()
```

If you use the `getClass()` method to append a class name, there are two special fields that are provided by perfino to simulate the abbreviated and simple class name modes that are available for the "Class name" and "Instance class name" naming elements:

- With `getClass().simpleName`, the name of the class without its package is added. For example, `com.mycorp.MyClass` becomes `MyClass`.
- With `getClass().abbrevName`, the abbreviated package names are added. For example, `com.mycorp.MyClass` becomes `c.m.MyClass`.

Note that the "Instance name" naming element generates a far higher overhead than the "Instance class name" or "Class name" naming elements, since it always involves actual method calls.

## B.2 POJO Transactions

With POJO transactions you can take any method call in the JVM and turn it into a transaction. The functionality for POJO transactions mirrors that of the DevOps transactions [\[p. 98\]](#) which are specified directly in your code by using the perfino annotations.

POJO transactions are necessary if you cannot add annotations into your code or if the classes of interest are in an external library. Otherwise, DevOps transactions are recommended for easier maintainability.

### Classes or methods

Similar to annotation transactions [\[p. 93\]](#), the first question is whether you want to choose a single method or all public methods from a particular class.

POJO transactions have the same configuration options regarding inheritance as annotation transactions. You can switch on transactions from derived classes with the "intercept subclasses" check box.

For a single method transaction, this means that overridden methods will create transactions as well. If such a method makes a super call, that super call will create a separate transaction. If the names of both transactions are equal, the super call will not be recorded separately. However, with an "Instance class name" naming element it is easy to create two different names. In that case, you will see the super method as a nested transaction. If this is not what you want, you can set the "no nested transaction" option on the "Naming" step of the wizard to the value "That match this entry".

The method of interest is conveniently chosen with the method chooser. It can show you classes from all connected VMs or from JAR, WAR and EAR files that can be uploaded on the fly. You can also edit the method manually in which case you have to take care to enter the method signature in [bytecode format](#).

Create Pojo Invocation

Please enter the details of the new business transaction.  
Monitor a selected method of an arbitrary class

1. Interception Type    2. Pojo Method    3. Filter    4. Naming    5. Policies

Class name:  Choose Method

Method name:

Method signature:

When editing directly, please note that the signature of the selected method is in bytecode format.

Custom description: ☐

For classes, all public methods of the selected class will create transactions. When the inheritance option is turned on, you can choose between two different strategies for method selection in the derived classes.

- **Implementing methods only**

This makes sense if you have selected an interface that defines all operations of interest with its methods. Additional public methods in implementing classes will be ignored.

- **All public methods**

The selected class may be a base class or a marker interface and the methods of interest are in the derived classes. In this case, all public methods in derived classes will create transactions.

As an additional way to collect transaction methods you have the option to select static methods. If the inheritance option is selected and all public methods are used, this will collect public static methods in derived classes as well.

Create Pojo Invocation

Please enter the details of the new business transaction.  
Monitor a selected method of an arbitrary class

1. Interception Type    2. Pojo Class    3. Filter    4. Naming    5. Policies

Class name:  Choose

Custom description: ☐

Method selection

☒ Implementing or overriding public methods

☐ All public methods of implementing or derived classes


If you select a marker interface or an abstract base class, select "All public methods of implementing or derived classes".

☐ Include static methods

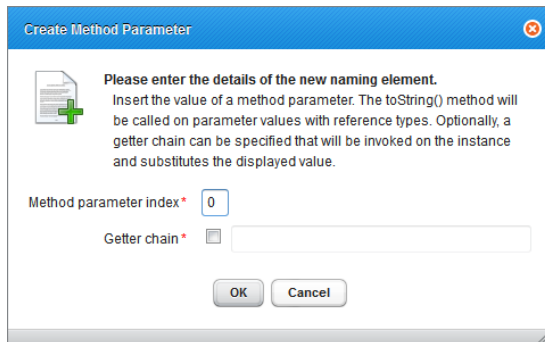


## Naming

In addition to the naming elements that are available for annotation transactions, POJO method transactions have a "method parameter" naming element. Method parameter interception only makes sense if a particular method is selected, so it is not available for POJO class transactions.



You can select a parameter by its zero-based index. In addition, a getter chain can be applied to the parameter just like for the "Instance name" element that is described for annotation transactions [\[p. 93\]](#).



## B.3 DevOps Transactions

Maintaining POJO transactions in the perfino VM group configurations is an extra step in the development workflow and can get easily out of sync with the actual code. While POJO transactions are ideally suited for external classes, it is much more maintainable to directly annotate the methods and classes of interest in your own code.

The naming of this transaction type comes from the [DevOps](#) software development method that stresses the importance of increased collaboration between the development and the operations departments. Thinking about and implementing monitoring aspects at development time falls into this category.

The perfino annotations are located in the JAR file `api/perfino_api.jar`, see the javadoc overview for how to download this JAR file with Maven, Ivy or Gradle.

All annotations have a class retention policy. This means that they are present in the class file, but the JVM does not load them into memory. Code that queries runtime annotations on a particular method cannot be confused by the additional perfino annotations because they are only detected by perfino at the class loading stage and do not appear in the loaded class objects.

The usage of DevOps transactions is **described in detail in the Javadoc** that is present in the `api/doc` directory of the perfino installation.


### Policies

DevOps annotations only define the transaction naming, but not the policies for transactions. Policies are closer to the operations side and often need to be adjusted in production.

It is possible to select different DevOps annotations in the perfino configuration by way of the **group name**. Each DevOps annotations can have a "group" parameter that is set to the empty string by default.

If the configured groups are not sufficiently granular or if you want to single out a particular class, go to the next step to add a class filter. Here, it is also possible to discard all DevOps transactions that originate from the selected group and classes.

Create DevOps Annotated Invocation



Please enter the details of the new business transaction.


Monitor methods or classes with perfino annotations. This transaction type allows developers to define transactions in code.

1. Annotation

2. Filter

3. Policies

Package filter

 Choose

Wildcard comparison ▾

Custom description \*

☐

☐ Discard transactions

On the last step of the wizard, you define the policies for the selected DevOps transactions. As with other transaction types, it often makes sense to keep a catch-all transaction definition at the bottom and add more specific transaction definitions at the top.

## B.4 Customizing Net I/O Thread States For CPU Recording

perfino offers you the possibility to record natively sampled CPU data [\[p. 77\]](#) with minimum overhead and stability risk for the entire VM. The resulting snapshot files can be viewed in JProfiler.

One pre-condition for useful sampling data is that all thread states where the thread is waiting are separate from the default "Runnable" thread state. If that is not the case, the top hotspots usually consist of methods that wait, block or perform network input and output (net I/O).

While waiting and blocking are comprehensively handled by the perfino sampling library, network I/O is often performed via native libraries that perfino does not know about. To mitigate this problem, perfino offers a mechanism to specify a list of additional methods that will be considered as net I/O.

To do that, you can specify the system property `-Dperfino.netioMethods=[path to text file]` in the Java invocation of the monitored VM. The referenced text file must be located on the machine where the monitored VM is running. Alternatively, perfino looks for the file `$HOME/.perfino/netio.txt` on Linux/Unix or `%USERPROFILE%\perfino\netio.txt` on Windows.

In the `netio.txt` file, add your net I/O method definitions, one definition per line. A method definition can have one of three forms:

- **Class wildcard**

Add a class name with a trailing asterisk, like

```
com.mycorp.MyClass.*
```

If you reference an inner class, the inner class separators must be written as dollar signs:

```
com.mycorp.MyClass$InnerClass.*
```

In this case, all methods of the selected class will be recorded in the net I/O state.

- **Signature wildcard**

Add a class name and a method name, separated by a dot, like

```
com.mycorp.MyClass.myMethod
```

In this case, all methods of the selected class with the specified name but with an arbitrary signature will be recorded in the net I/O state.

- **Specific method**

Like "Signature wildcard", but with the signature in bytecode format appended at the end. For example,

```
java.net.AbstractPlainSocketImpl.doConnect(Ljava/net/InetAddress;II)V
```

This is an actual net I/O method that is handled by default, the signature is for a method that returns void and takes a `java.net.InetAddress` and two `int` parameters. This is the same format that is used by [JNI type signatures](#).

The single specified method will be recorded in the net I/O state as a result of this definition.