



VOSS



**VOSS Insights
Windows Forwarder Install Guide**

Release 23.3

Dec 13, 2023

Legal Information

- Copyright © 2023 VisionOSS Limited.
All rights reserved.
- This information is confidential. If received in error, it must be returned to VisionOSS ("VOSS"). Copyright in all documents originated by VOSS rests in VOSS. No portion may be reproduced by any process without prior written permission. VOSS does not guarantee that this document is technically correct or complete. VOSS accepts no liability for any loss (however caused) sustained as a result of any error or omission in the document.

DOCUMENT ID: 20231213175955

Contents

- 1 Introduction** **1**
- 2 Install Windows Forwarder** **2**
- 3 Configure Windows Forwarder** **3**
 - 3.1 Overview 3
 - 3.2 Event Logs 4
 - 3.3 Files 5
 - 3.4 Disk Space 8
 - 3.5 Performance Monitoring 9
 - 3.6 Commands 10
 - 3.7 Database Queries 12
 - 3.8 Registry 23
 - 3.9 Network 24
 - 3.10 Windows Services 25
 - 3.11 Forwarder Configuration 26
- 4 Save and Retrieve Settings** **28**
- 5 Reporter Integration** **31**
- 6 Centralized Management** **33**
- 7 Setting up a Data Source Name (DNS)** **34**

1. Introduction

VOSS Insights Windows Forwarder (the Forwarder) is an application that runs on a Windows server or workstation and monitors the system activity.

Information is passed along to the primary and/or secondary Forwarder indexing engine.

The Forwarder is an NT service and is supposed to start automatically each time the operating system is started.

The Forwarder is capable of monitoring any activity, either directly by watching for Registry changes, SNMP traps, event logs, Performance counters, and application logs, or indirectly by executing commands and SQL queries for any ODBC compliant database.

2. Install Windows Forwarder

To install VOSS Insights Windows Forwarder, run the following file (which is included in the installation package): `setup.exe`

Setup installs the Visual C++ runtime environment, then proceeds to run `setup.msi`, which will display all of the installation dialogs, prompting for the target directory as well as the installation locale parameters, IP address of Arbitrator (or Dashboard) and its port number.

Note: If VC++ runtime environment is already installed, there is no need to run `setup.exe`. In this case, you can execute `setup.msi` directly.

You can also execute the setup in “silent mode”, without displaying the GUI, by running:

```
msiexec /i <path to setup.msi> /qn /l* <path to setup.log> PORT=62009 EDEVICESITE=<site>␣  
↵EENTITY=<entity> IPADDR="10.13.37.185"
```

The installation includes the Forwarder, the configuration application, and all system components required to make both applications run.

The installation configures the Forwarder as an NT service. All the shortcuts required to launch the configuration application are also created.

The installation collects the minimum amount of information required (i.e. IP address of local monitoring appliance). All the additional information can and should be added later using the configuration program.

The Forwarder is fully functional after the installation and should be up and running when the system is rebooted. It can also be started from the configuration application.

3. Configure Windows Forwarder

3.1. Overview

The VOSS Insights Windows Forwarder configuration application provides a graphical user interface (GUI) for configuring the Forwarder. All settings are saved and retrieved from the Windows registry.

The current version of the Forwarder has multiple sources of data, including Windows event logs, files and disk space, Windows Registry, Event logs, Performance counters, Windows Services, Database queries, as well as the output of various commands.

The configuration application design reflects these data sources. For example, parameters for configuring the Forwarder are found in the **Forwarder** tab.

The main pane of the application displays the VOSS logo.




You can click an icon in the menu tree to edit any of the settings.



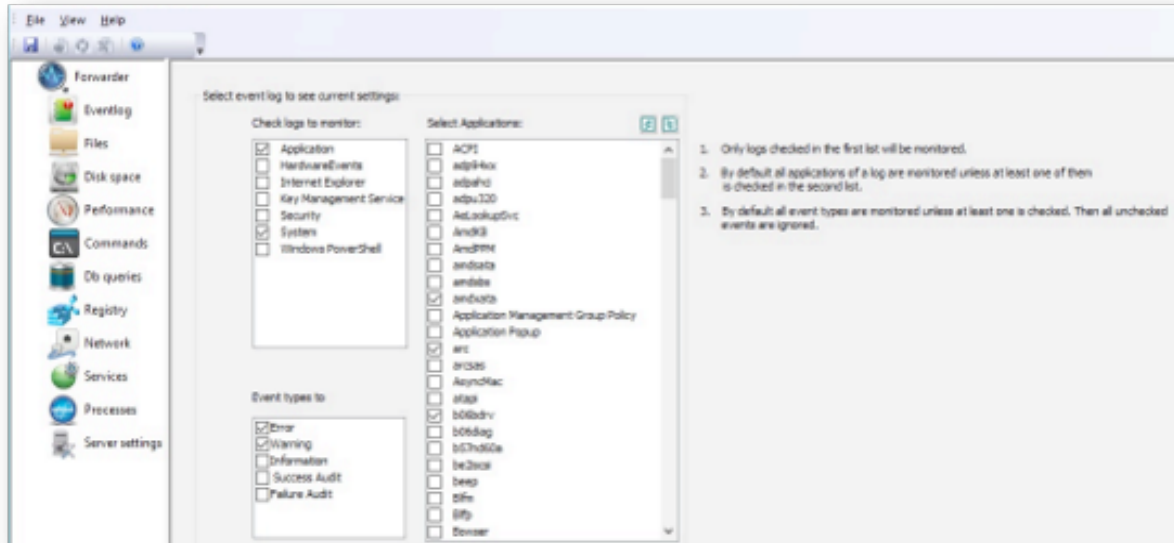
Note: The GUI toolbar is activated when any screen displays a list of items (files, commands, queries). You can either use the toolbar buttons or right-click in the list to display a drop-down menu.

3.2. Event Logs

The Event logs tab displays all the event logs available on a current system. Those logs that are selected for monitoring are checked.



You will need to click on a specific event log in the **Check logs to monitor** list to view the selection for that log:



For each log, one or more of the five event types can be selected by checking the appropriate box in the **Event types to monitor** list.

By default, events of checked types are being monitored for all the applications in the system. You can select only those applications that have to be monitored by checking appropriate boxes in the **Select applications** list.

To summarize:

The Forwarder has three levels of configuration for event logs:

1. Select only those event logs that are of interest.
2. Select only specific event types for each log.
3. Select applications for each log.


This granularity is intended to reduce traffic between the Forwarder and the local Arbitrator/Dashboard appliance, and make troubleshooting easier.

3.3. Files


The Forwarder can monitor both flat files and named pipes. Each monitored file has to be added using the configuration program by clicking on an **Add** button.

Files already in the system can be reconfigured by right-clicking on a specific line. The pop-up menu, which displays after right-clicking on a line, allows you to remove the file altogether or change its settings.

To stop monitoring a specific file, you don't need to remove it; just uncheck the relevant checkbox.




The same dialog box displays for both *new* and *updated* files.



- **Logical name** - Should describe the origin of the file. Its contents are completely up to the person configuring the application.
- **Monitor file with static name** - click the Browse button adjacent to the field to select the file (even though file name can be typed in). The Browse button provides the opportunity to automatically populate the field with the file name and the directory name.
- **Monitor files with name pattern** - Applies to a situation where the program creates a series of traces (files) in a specific directory, with a specific naming pattern. To configure monitoring of these files, click the Browse button adjacent to the field and select one of the files (it doesn't matter which one). The system tries to find all or the instances of the files with a similar name. The resulting string should be fine tuned by the user. The permanent part of the name should be left intact and the variable part replaced by the *.

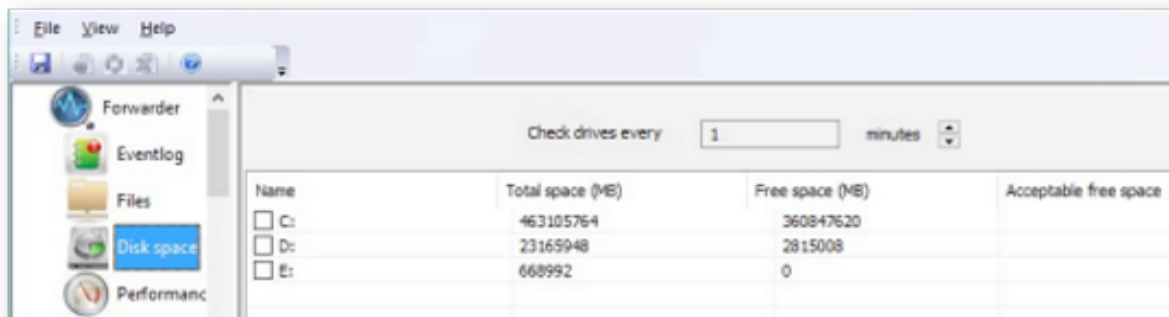
On this screen you can also choose which lines from the file will be processed. The selection can be applied based on sequential line numbers (all lines, odd lines, even lines), and based on specific keywords found in the line. The first option will handle traces where each line's text description is followed by the line with a hexadecimal representation of data. The line is selected if any of the keywords in the list is found.

To update the File settings, right-click on a file and select **Update File**:



3.4. Disk Space

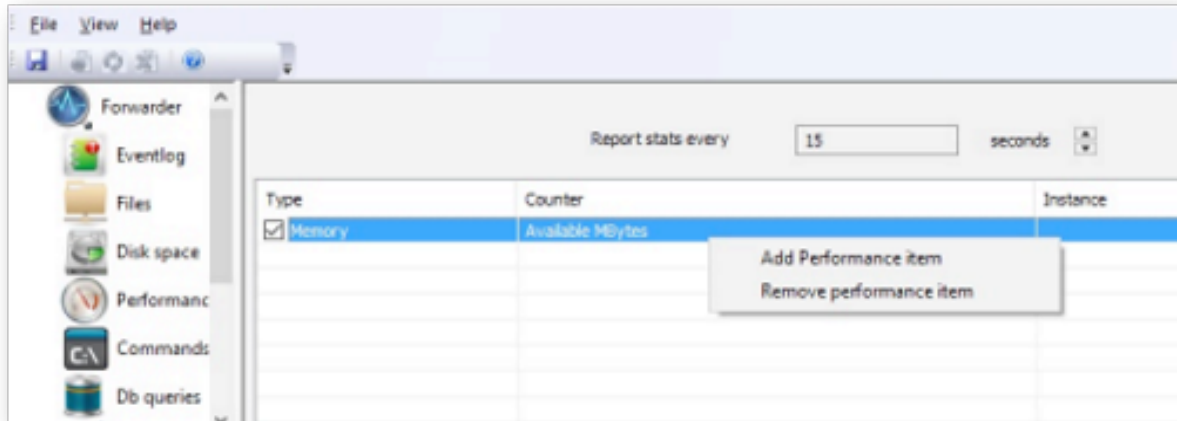
The Disk Space tab displays information about all the logical drives, along with their maximum capacity and current utilization. It allows you to set an acceptable free space limit for each device. The Forwarder raises an alarm once the limit is reached. The free space may be set as either a percent of total capacity or an explicit amount of space in megabytes.



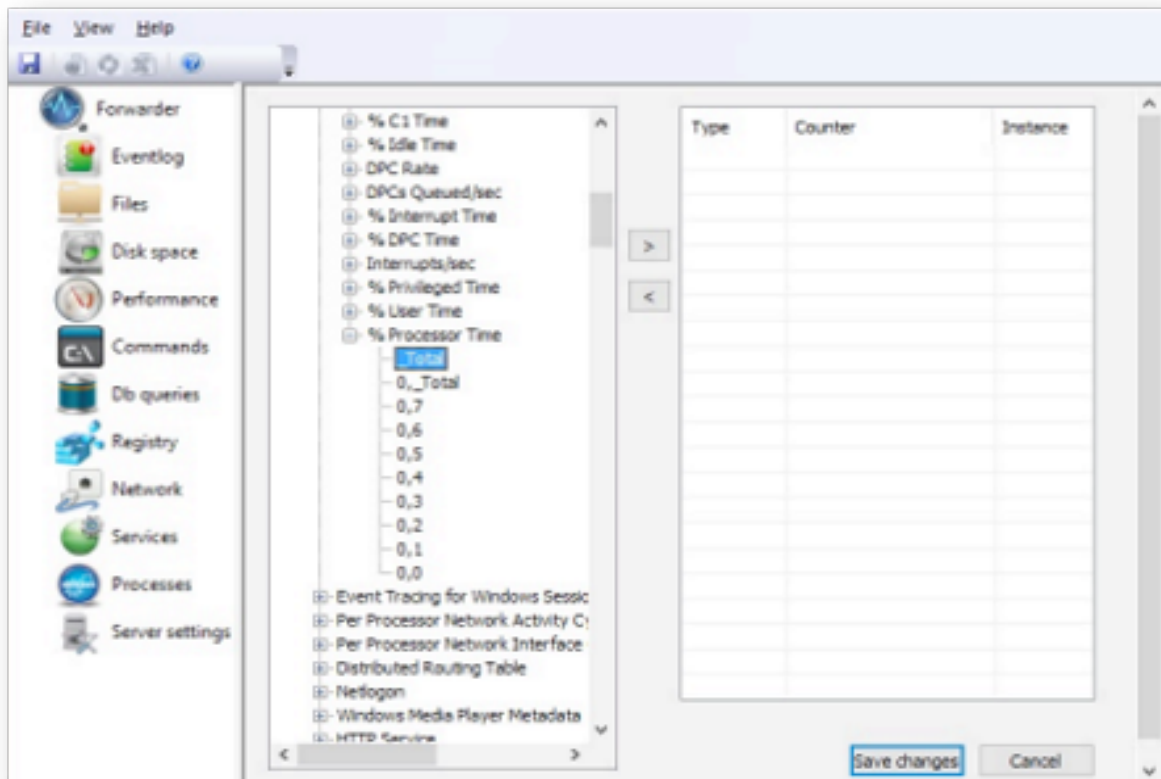
The last option on this screen relates to reporting absence of activity in the current file to the local monitoring appliance. There are situations when we can presume that a specific application is malfunctioning if it is not writing anything to the trace. This option would allow us to raise a red flag under these circumstances.

3.5. Performance Monitoring

The Performance Monitoring tab displays all performance counters currently monitored by the Forwarder. When the configuration application is first launched, the list includes all predefined counters.




The Performance Monitoring program collects the list of all the counters available and saves them in a list that serves as a basis for adding new counters:




Click the right-pointing arrow (>) to add a selected counter to the list. Click the left-pointing arrow (<) to

remove it from the list.

3.6. Commands



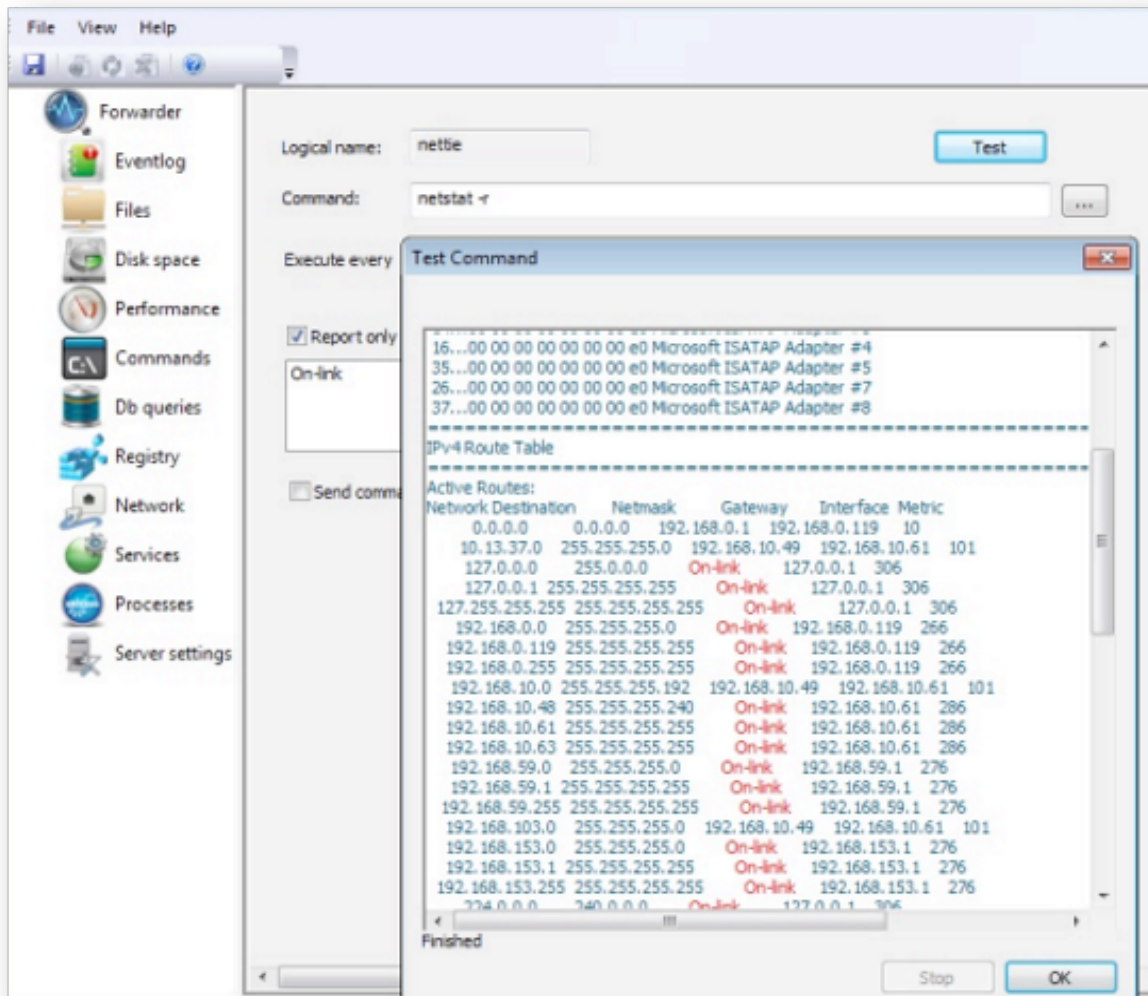
The image shows the dialog that displays when adding a new command or when modifying an existing command:




Special features have been added to handle the output from commands that produce CSV file types. It is strongly recommended that you select **Send command output one line at a time** for CSV formatted files

since data sent to the Arbitrator server will closely resemble the result of a SQL query.

Each command must be tested before it is added to the configuration. To do this, click the **Test** button:



3.7. Database Queries

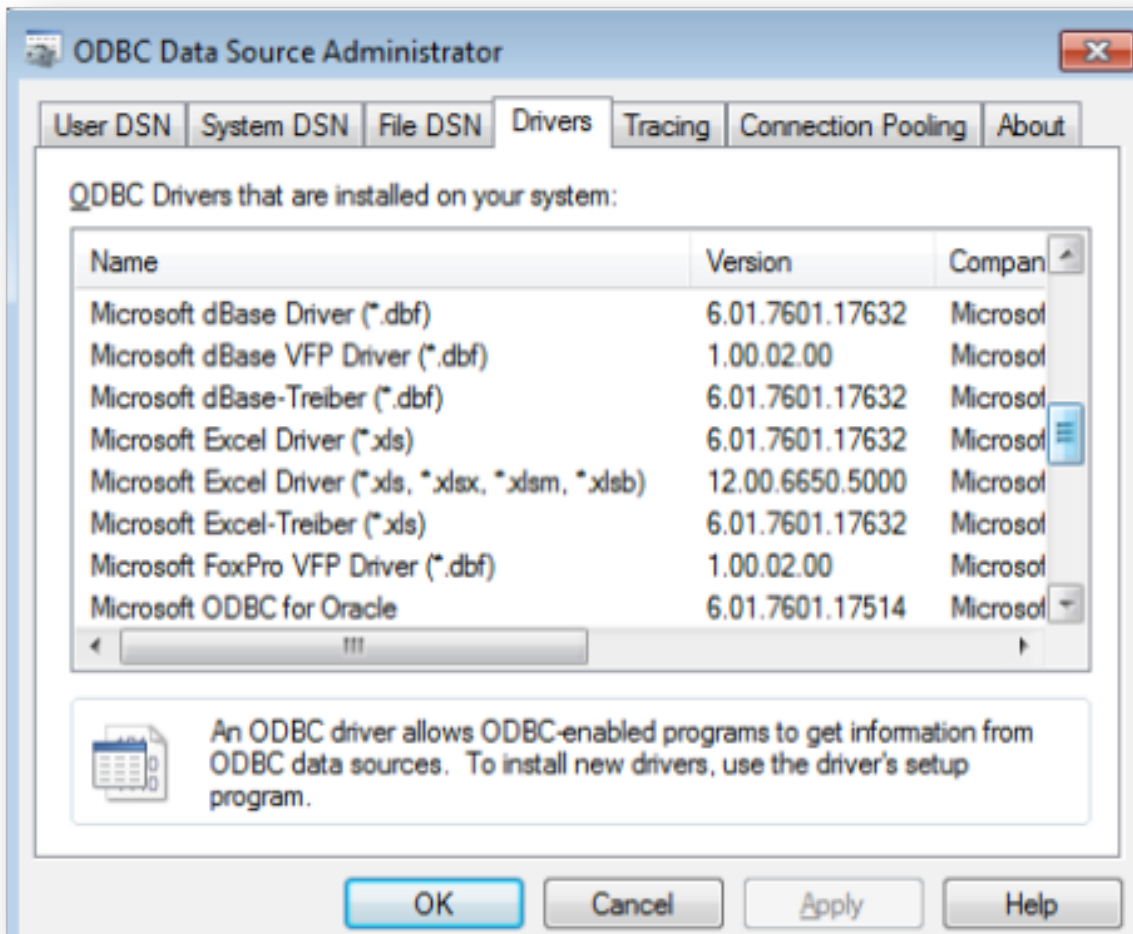


The Forwarder is able to run queries against any ODBC compliant database at defined intervals, and to stream the results of the query to the Arbitrator server.

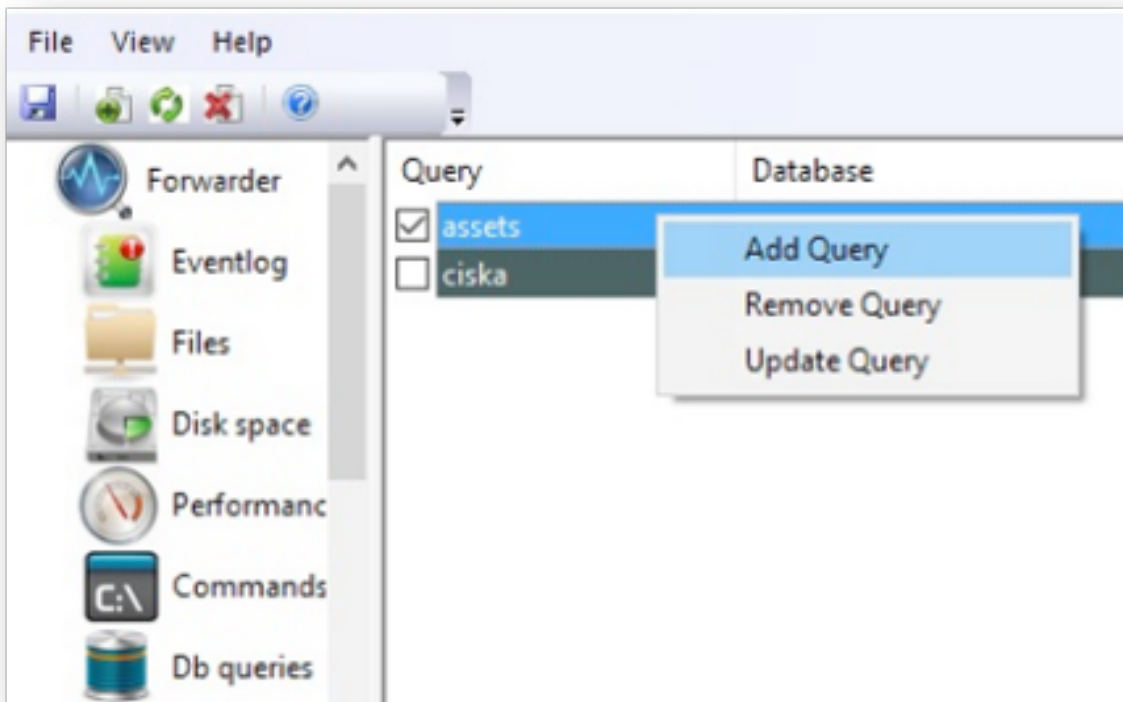
To configure a query, the local system must have an ODBC driver installed and a system DSN configured. The Forwarder has been tested against multiple databases, including Microsoft SQL server, Postgres, and Intersystems Cache.

Since the Forwarder is a 32 bit application it will use 32 bit drivers and related DSN-s that can be checked by executing:


```
C:\Windows\SysWOW64\odbcad32.exe:
```




To configure a new query, select **Add Query**, which displays a list of all existing ODBC data sources:




Choose a data source and click **OK**.



The system attempts to connect to the database to retrieve information about its tables and views that will be displayed on the following screen in the leftmost window.



You can extend tables and views to see the database architecture. The tree structure will depend on the database layout and schemas (when applicable). Clicking on a table or view displays all the fields and their types:

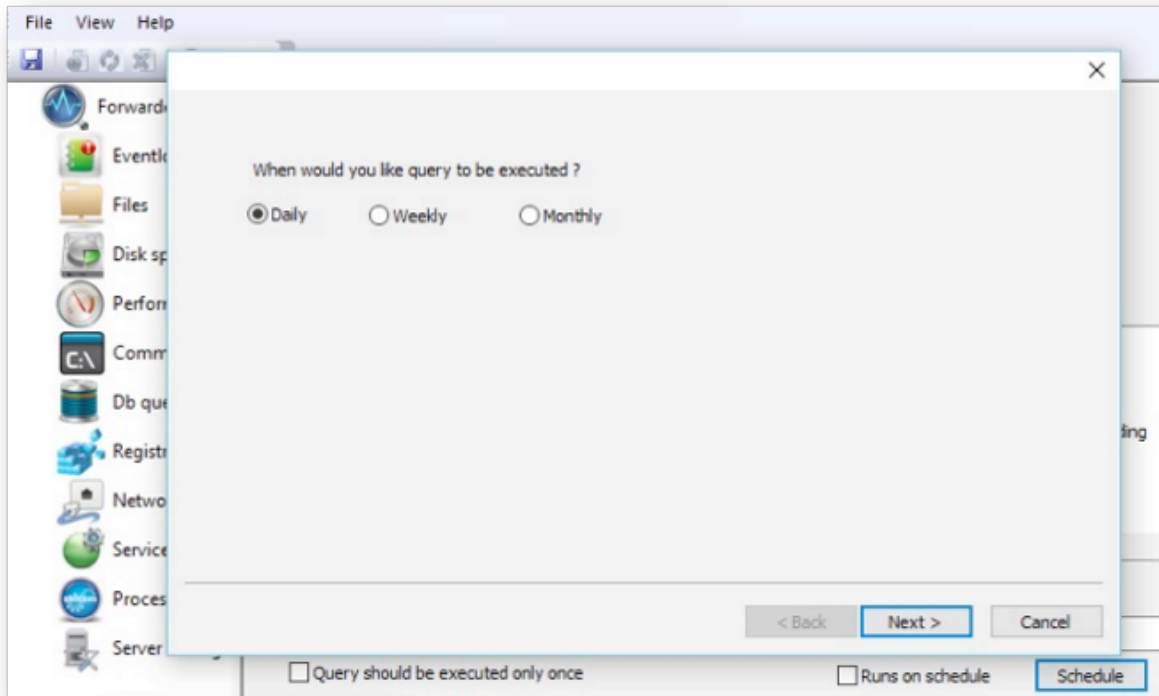


The SQL statement should be typed in the third window. Clicking **Run** triggers an execution of the statement and displays the result of the query.


There are three options for executing the SQL query:

- 1) every “xxx” seconds as defined on configuration screen above
- 2) only once (which may be used to import large amount of data)
- 3) based on a schedule that can be configured by pressing the “Schedule” button

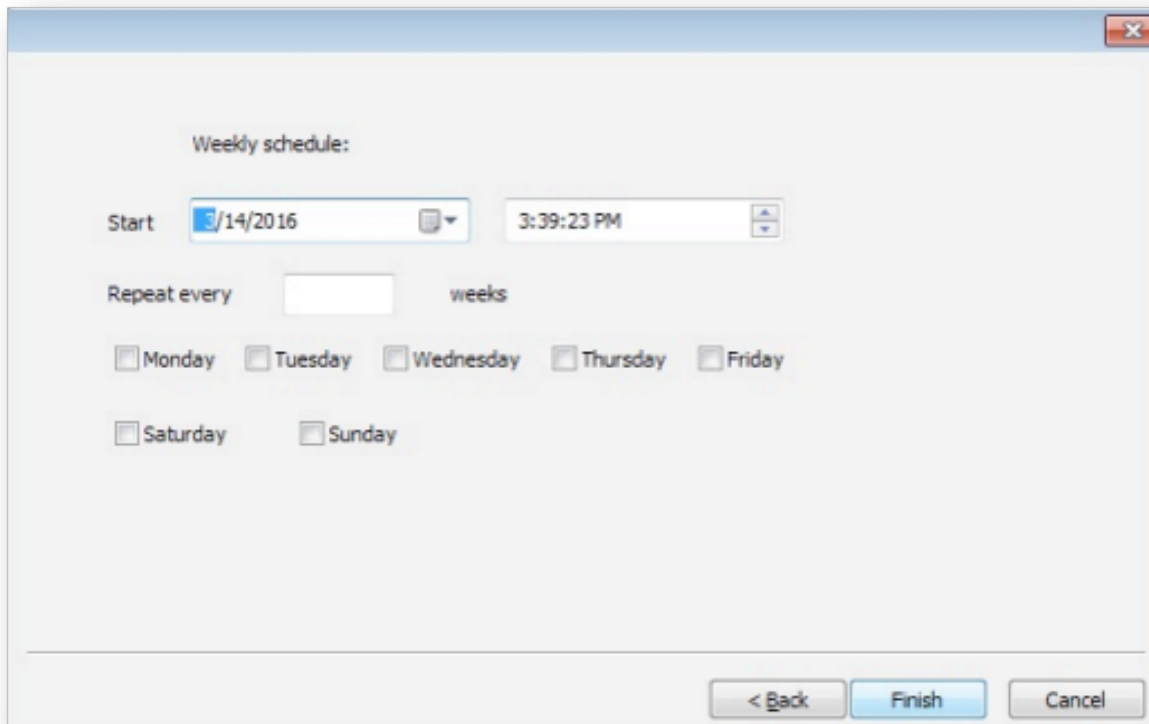
The first screen allows you to choose the schedule mode:



When using a daily schedule, the query is executed starting at the date and time selected, and repeated the same time every n days, based on the configuration:



When using a weekly schedule, the query executes on the days of the week selected:



Weekly schedule:

Start

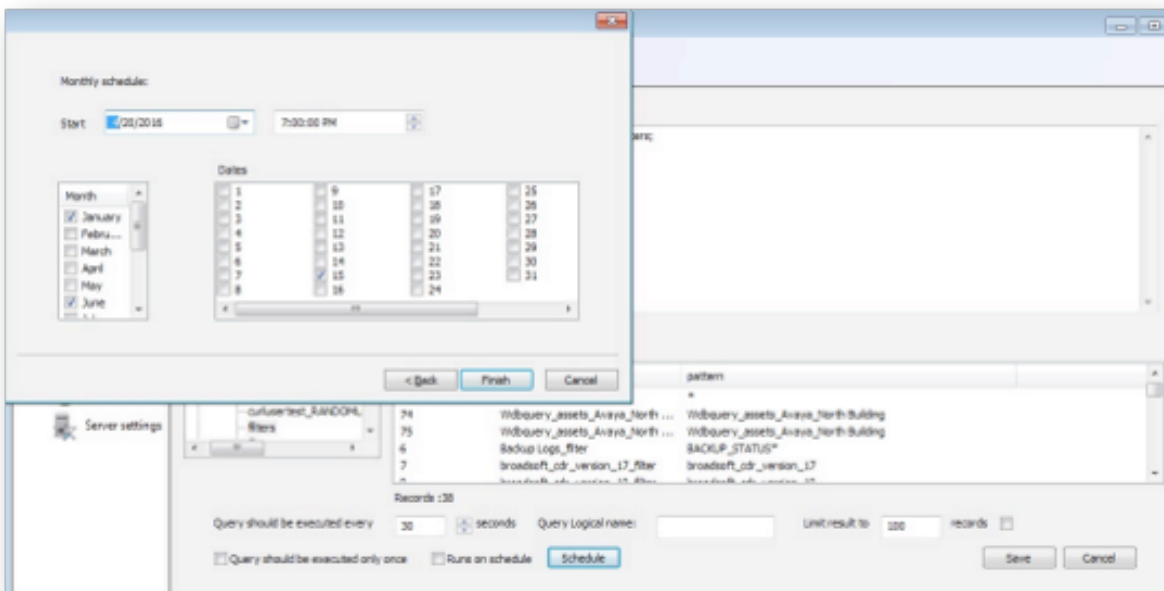
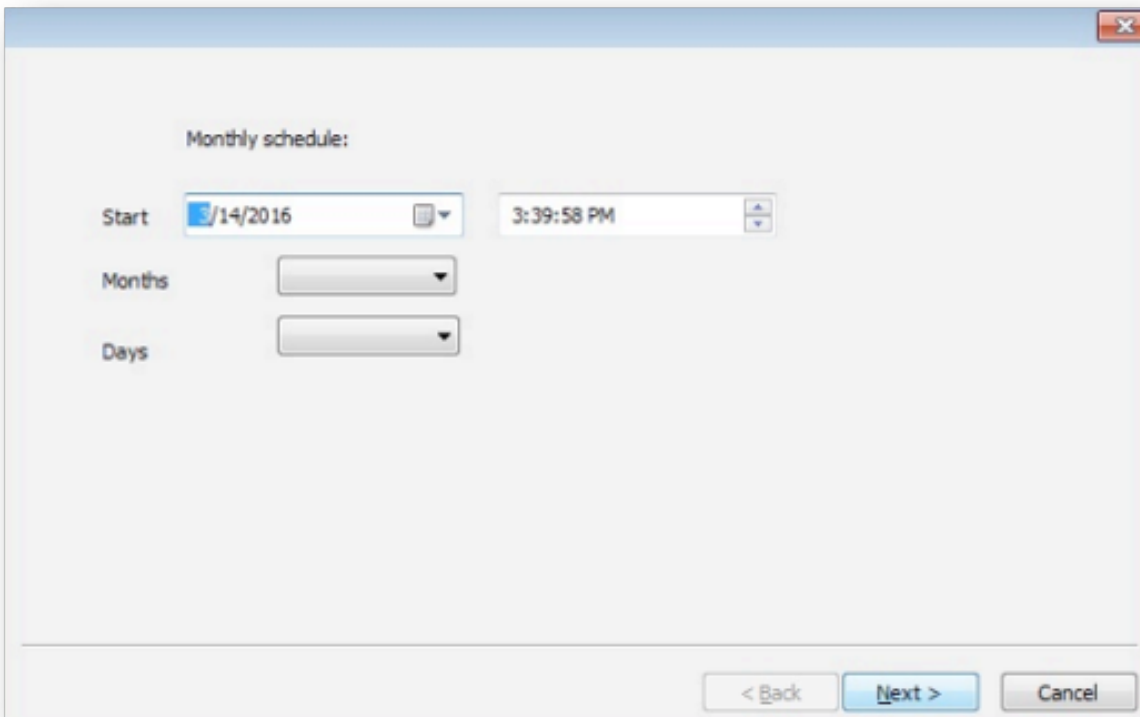
Repeat every weeks

Monday Tuesday Wednesday Thursday Friday

Saturday Sunday

< Back Finish Cancel


When using a monthly schedule, you can select months and dates of the month:




Modifying the query with a new database connection

You may need to modify the database connection for an existing query after importing the resulting data on a machine different from the one where the query was originally configured.

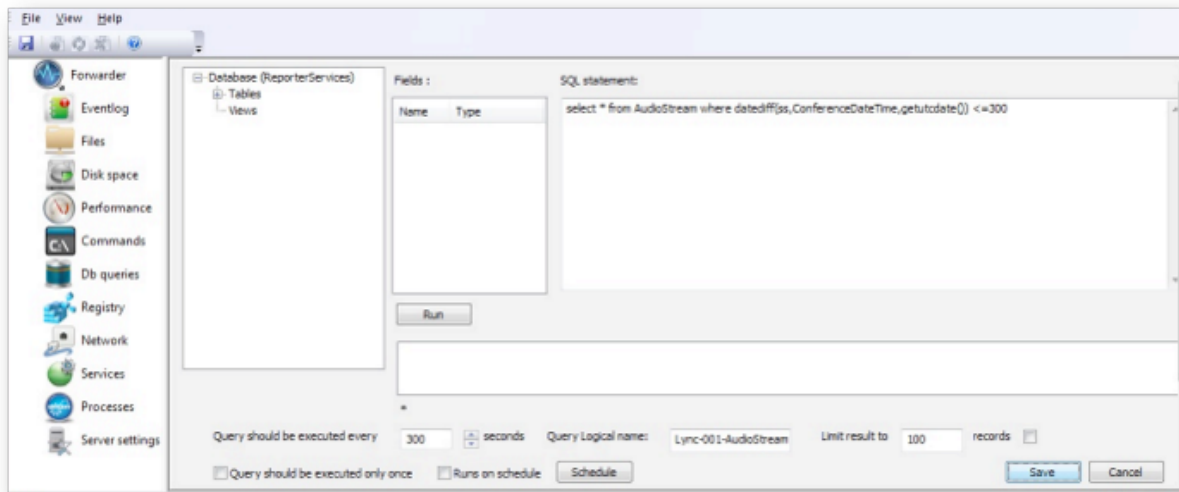
Since the database connection is embedded in the overall query information, a special mode was added to handle this situation:




Selecting **Modify database connection** displays a list of available DSN-s. These are presented in the same way as if the query is being created from scratch:



Once a DSN is selected, the application shifts to a single query view that combines SQL and scheduling data from previously existing query and new database connection:




When saving the modified query, a list of queries will now reflect the new database information:



3.8. Registry


A separate service (which runs in tandem with the Forwarder), performs registry monitoring.

Registry monitoring can be enabled or disabled and you can specify what hives should be included.



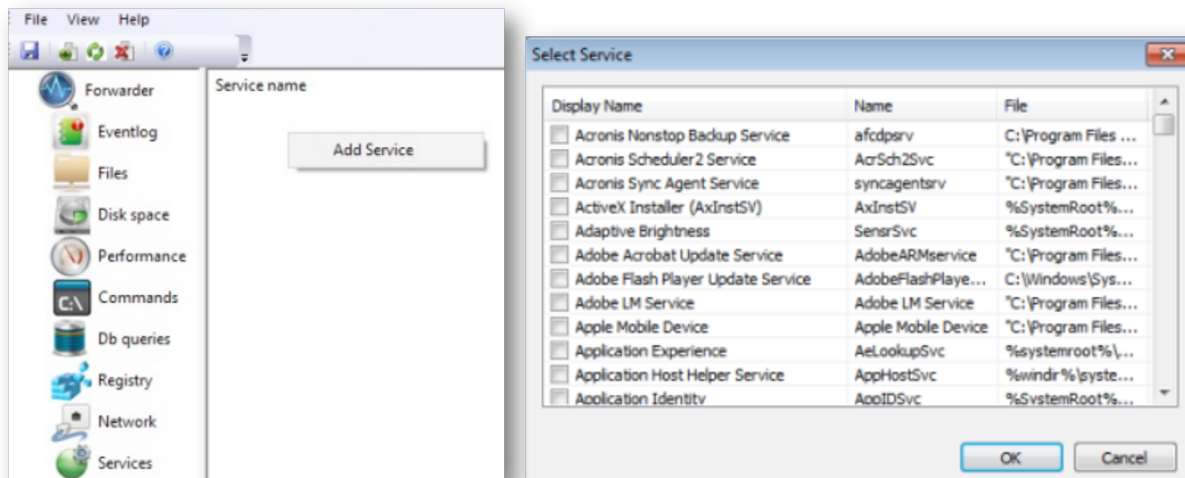
3.9. Network

The Forwarder can monitor all network connections providing the same information as the netstat command, including process information. Monitoring is performed based on the configuration and will include TCP and UDP connections depending on the options you select:

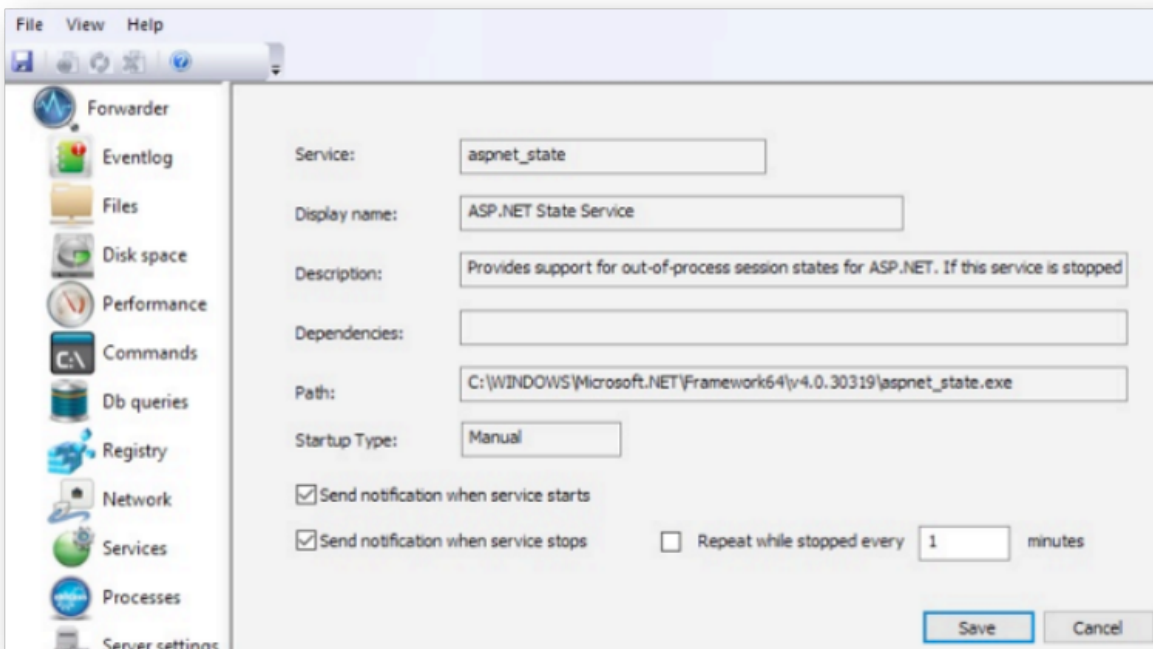
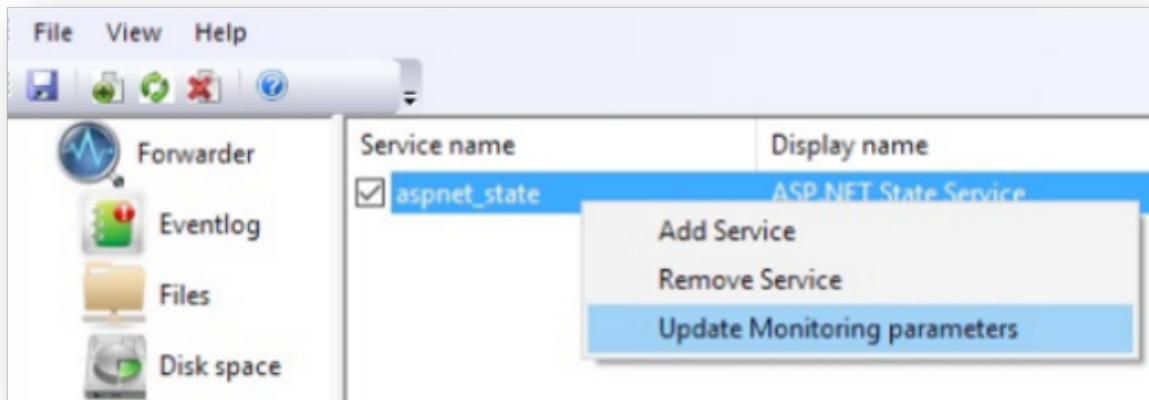


3.10. Windows Services

The Forwarder can monitor the status of any Windows service and report it to the Arbitrator server.

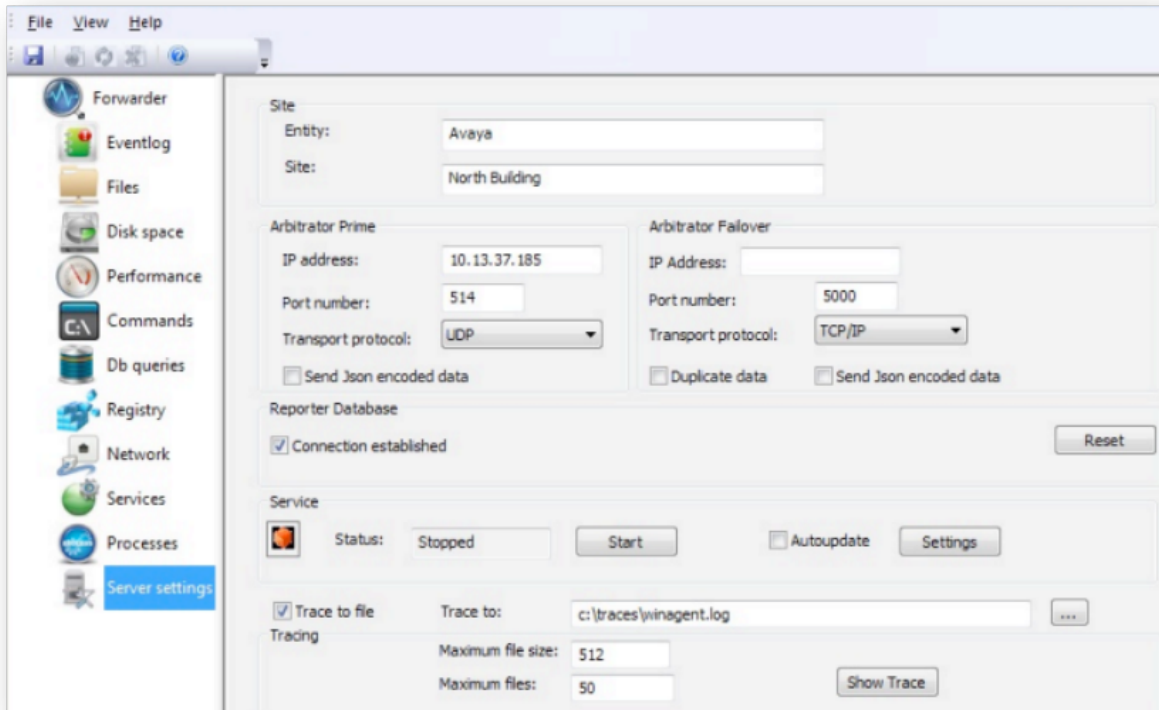


For each monitored service, you can specify the following options:



3.11. Forwarder Configuration

This page allows you to define connection information between the Forwarder and the local Arbitrator server.



Although the configuration screen explicitly mentions Arbitrator (Correlation Server), the Forwarder can work with both the Arbitrator and Dashboard server. The ports will depend on which system the Forwarder will be connected to. In the case of Dashboard, it is preferable to send the data in JSON format.

The latest version of the Forwarder can save data directly into the Dashboard server database, if it is configured to communicate directly with the Dashboard server. Even though the configuration dialog mentions Arbitrator (Correlation Server), you can use the Dashboard IP address.

The Forwarder service will send a request to the server upon startup and upon receiving the type of the server, will change its behavior.

To be able to write data into the Dashboard/Reporting database, the Forwarder needs a locally configured system DSN pointing at the Dashboard/Reporting database. This DSN will be created by a configuration program when it's first launched after the installation.

If the DSN has been configured and tested, a **Connection established** checkbox is selected, and the service will be set to create several "static" tables that will hold performance, eventlog, network, and process data. These tables are created by calling the REST APIs executed by the Dashboard/Reporting server. Each of these tables will be named following the same pattern, <data type>_<site>_<entity>, for example, evtlog_avaya_north building.

Tables collecting the results of database queries will be named: w̄bquery_.

This screen also provides a way to configure traces of the Forwarder service.

4. Save and Retrieve Settings

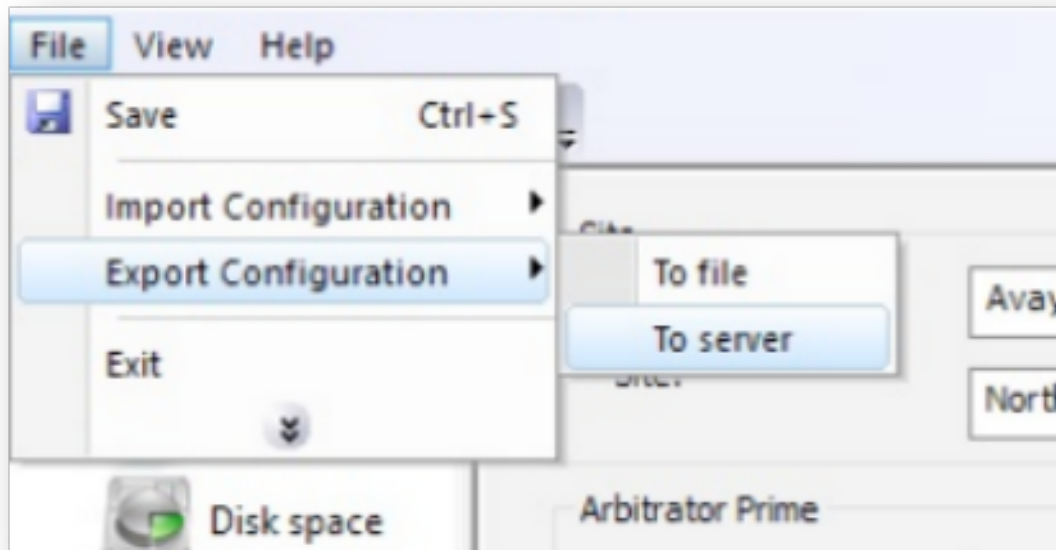
All settings for the VOSS Insights Windows Forwarder are kept in the Windows registry. They are loaded by both the Forwarder itself and the configuration program on startup.

All changes made in the configuration program are kept locally until you click the **Save Setting** button. Only then are they saved in the registry.

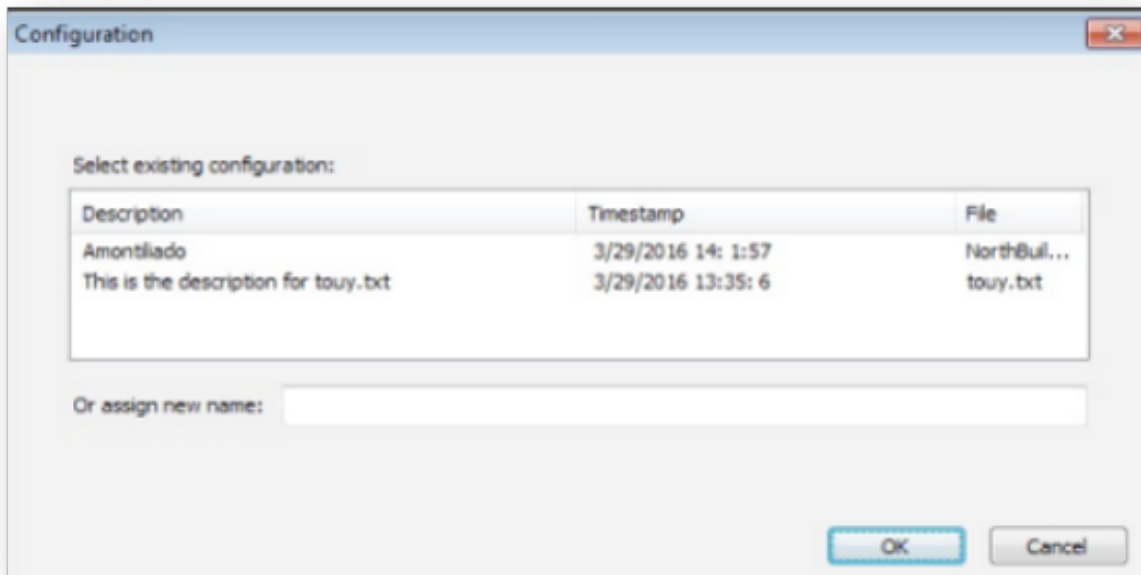
The Forwarder was designed to pick up the changes “on the fly” (i.e. service does not have to be restarted to make changes effective).

There are two ways to propagate the Forwarder settings:


1. Select **File > Export Configuration > To file** to save the settings as a file. This file can be later used to restore settings or configure different instance of a Forwarder.
2. Select **File > Export Configuration > To server** to upload settings to the Arbitrator server and use it as a repository for the configuration files.



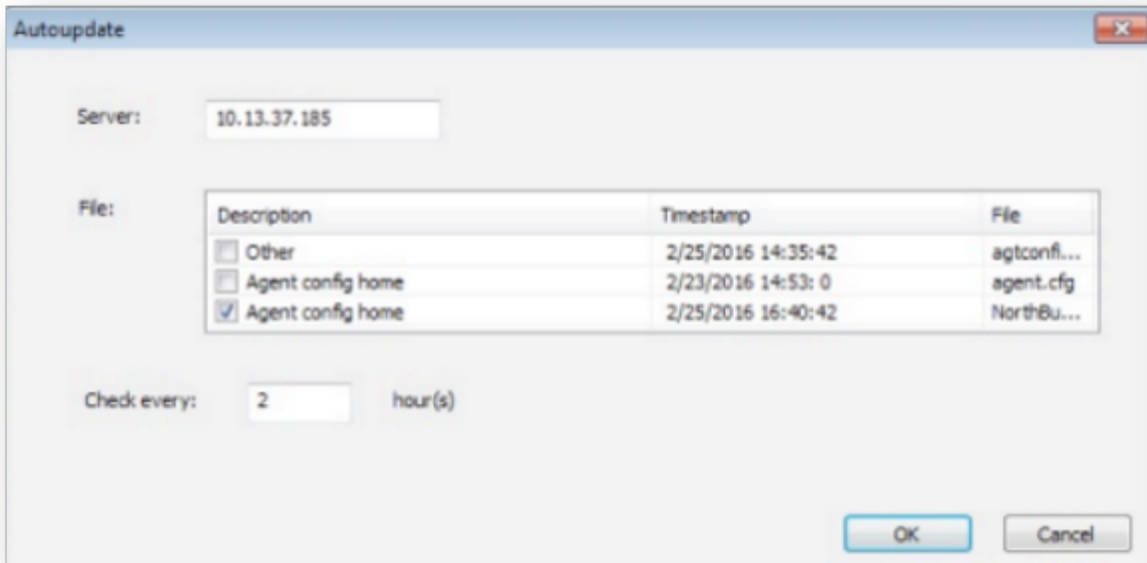
The system displays the list of currently saved configuration files. You will have an option to either replace the existing file or to create a new one:



The server-based repository may be used in conjunction with the Forwarder's *Autoupdate* feature, and can be configured from the same screen:




The system displays a list of existing files. You can select the one it will use to update itself.



5. Reporter Integration

As described earlier, the Forwarder can send data to both Arbitrator and Dashboard/Reporting server. In the case of the Dashboard/Reporting server, the Forwarder can add certain data elements as a custom resource to be used for the widget:

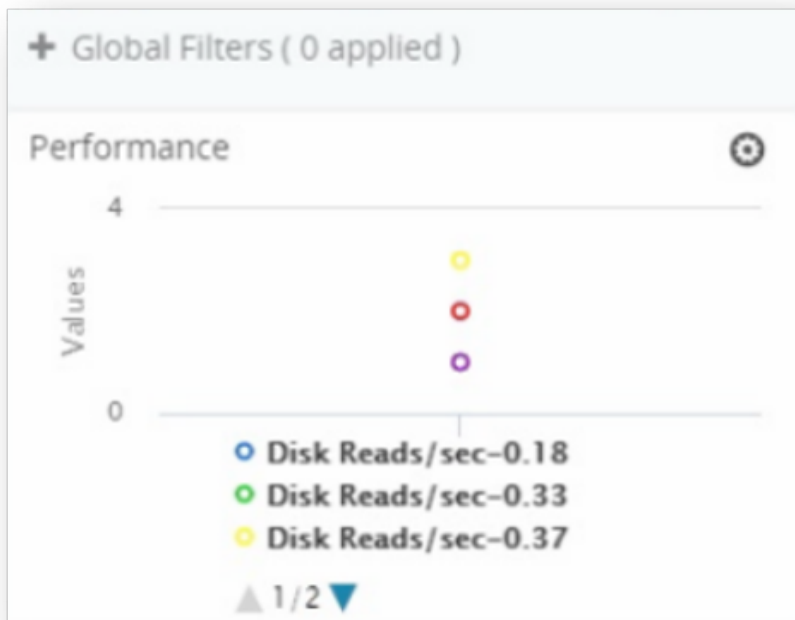


The current version creates three types of custom resources for each instance of the Forwarder, where *Locale* is the "Entity" and "Site" of each Forwarder:

- Evtlog_<locale>
- Perfmon_<locale>
- Wbquery_<locale>

For each instance of a Forwarder there will be one Evtlog and Perfmon source and as many Wbquery sources as the database queries configured on the Forwarder.

This feature is designed to simplify the process of creating reports and Widgets based on the data that can be easily tabulated:



6. Centralized Management

Each instance of the Forwarder registers with the Dashboard server or the Arbitrator server configured as the main server.


The registration creates a record in the central database and enables the server to display information about all of the installed Forwarder instances.

Once a minute, each Forwarder sends a *keep alive* message to the server, thus maintaining its status.

It will also periodically check to see if the profile assigned to it has changed. If so then and if it will download and replace itself with the updated version.

Note: The AutoUpdate feature needs to be enabled.

Centralized management allows you to configure one or multiple profiles, upload it to a central repository, and distribute it to multiple instances of the Forwarder without touching the actual server it's running on.



7. Setting up a Data Source Name (DNS)

This section describes how to set up a data source name (DSN) to execute database queries via the VOSS Insights Windows Forwarder.

The Forwarder uses a DSN connection mechanism in order to connect to an external database and collect information periodically.


Requirements for the environment:

The following assumptions are made for the environment:


- The Forwarder is installed on the same host machine as the database.
- The target database type is MS-SQL 2XXX.
- The Host machine's operating system is Windows 64-bit.
- The Host machine is part of a windows domain.
- The domain (admin/local/system) account that is used to run the Forwarder as a service has at least read-only access to the database that is supposed to be queried periodically.

Perform these steps to set up a DSN:

1. Use Windows Explorer to go to the following directory: C:\Windows\SysWOW64
2. Locate and then execute the following file: odbcad32.exe
3. Select the **System DSN** tab.

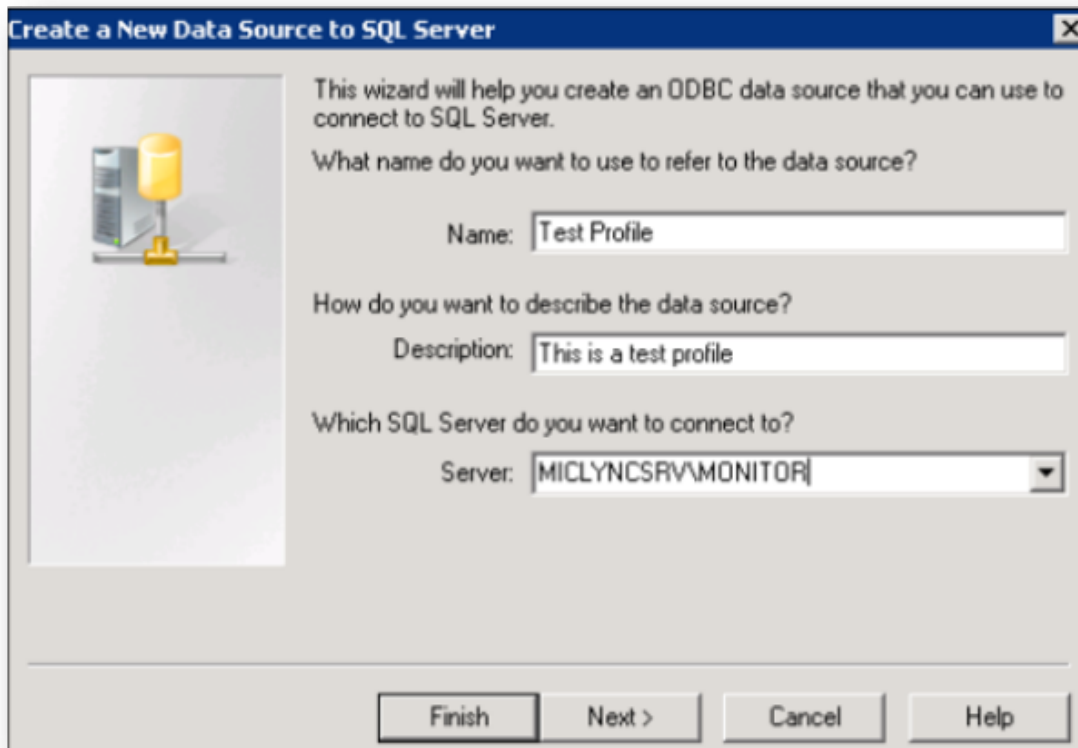


4. Click **Add**, then select **SQL Server** from the list, and click **Finish**.

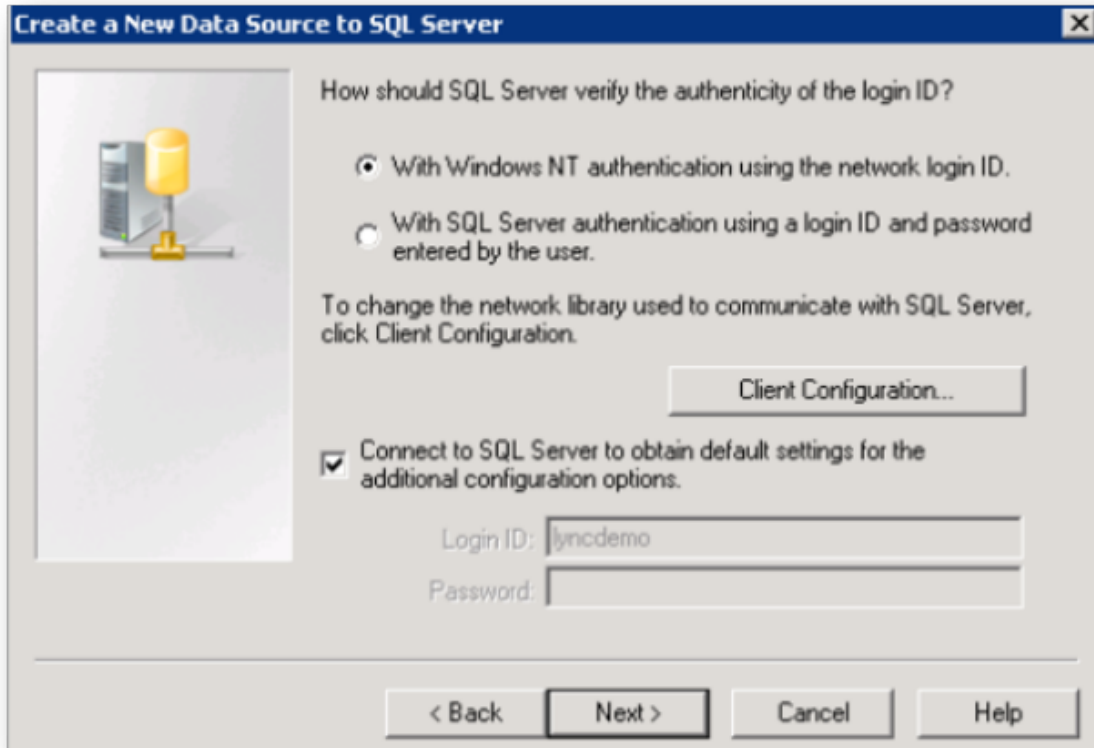


5. On the **Create a New Data Source to SQL Server** dialog:

- Fill out the **Name** and **Description** fields.
- Select the down-arrow at the **Server** drop-down, then let the system discover all the database instances available in your sub-network. Choose the appropriate “database hostnameinstance name” to connect to.
- Click **Next**.

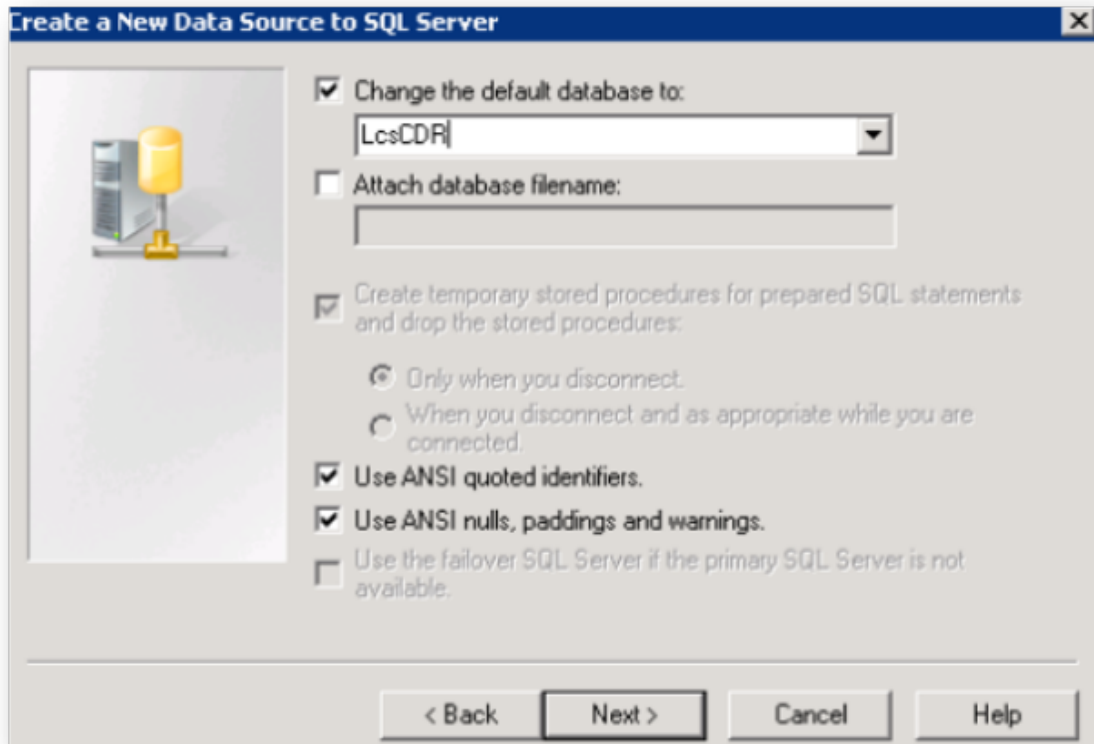


6. Do not change any settings on the authentication page. Click **Next**.

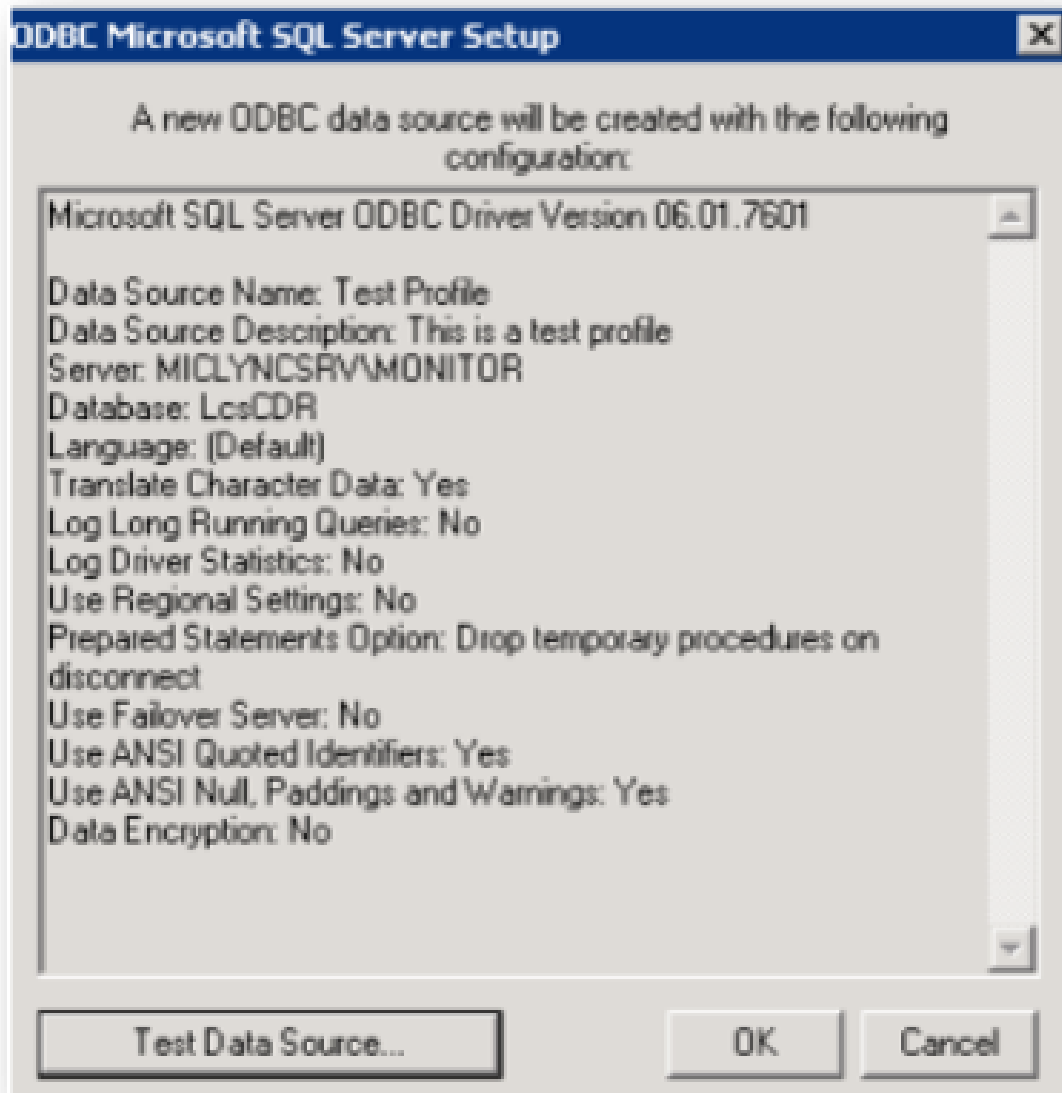


Note: Before proceeding, review the prerequisites for this procedure, which describe the assumptions around your environment setup.


7. If your credentials are correct, you will be able to connect to the database instance and set a default database name on the next screen. Choose the appropriate database name that you would like to query and click **Next**, then click **Finish**.



8. Click **Test Data Source...**



9. If the test passes, you will see the following message:



10. Click **OK** to exit setup.

The next time you try to add a “Db query” on Forwarder (see section Database Queries) where you will see your DSN name listed in the *ODBC Data Source List*.