Configuring ELMA to Work under High Load

User Manual

Business Process Management Software

Contents

Introduction

This book is a manual on how to configure ELMA to operate under high load. This book describes how to install and configure the failover cluster for MS SQL Server databases and MS Web Farm Framework. The book also details on how to maintain the system components and solve various issues. It is aimed at specialists, who intend to implement the system on the professional level.

It is assumed, that the user already knows the architecture of ELMA and has the basic skills, which are described in the ELMA BPM Platform user manual. The user is also supposed to be familiar with the basics of ELMA administration, described in the ELMA Administration user manual.

Below is the list of user manuals:

- ELMA BPM Platform user manual
- ELMA Web Portal user manual
- ELMA ECM+ user manual
- ELMA Projects+ user manual
- ELMA KPI user manual
- ELMA Administration user manual

You can find information about all the ELMA functions in ELMA Help. It is available in the ELMA knowledge base: https://kb.elma-bpm.com/help.

Solutions for many technical issues are provided in ELMA knowledge base: https://kb.elma-bpm.com. The knowledge base is constantly updated by the company's specialists.

Chapter 1. General Information

Corporate systems are often designed to support numerous end users – up to several thousands and more. For a system to work under significant load and maintain a sufficient performance, it is necessary to configure the system itself as well as the company's IT infrastructure: match server capacities, account for the network capacity and so on. Scaling up is expensive and not always possible. In this case, it is more reasonable to **scale out**, i.e. deploy an **application farm** to distribute the load.

For example, if you need 1000 users to work at the same time (a back office of 4000 employees), it is more advantageous to have four regular servers than one high power server. The four regular servers will handle the load just as well.

In this manual, the term **high load** is used. The term implies operation of ELMA with numerous concurrent users (more than 300-1000, depending on the intensity of their work), when it becomes more beneficial to increase the number of servers, instead of increasing the capacity of one server.

ELMA web farm is a group of ELMA system servers, linked into a single network that has two-way data transfer and works as a single unit. All the servers work with the same database and use the same cache (**cache cluster**). Certain settings allow connecting several servers to one database and synchronize them.

Roughly, it works like this: a user opens ELMA Web Application and performs a certain action; after that, the system sends a request to ELMA server. The request can be processed not only by the first server, but also by any other server, which is a part of the farm. This way, the server load is distributed among several servers.

In addition, the load is distributed evenly among the servers due to the **load balancer**. It receives requests from users and redirects it to the servers, thus distributing the load. Load balancer can be easily and quickly configured in the IIS manager and does not require resources.

High load on the system is a result of numerous operations carried out in the system. It also often leads to strict requirements to fail-safety: a thousand users cannot wait in case of a crush. **Failover clusters** can help with that. Without clustering, server malfunction causes its network services to become unavailable until the server is functional again. Failover clusters allow writing data to several

databases at the same time. It means that if a server malfunctions, users will be able to continue their work.

When you deploy the system to work under high load, you are supposed to use different server machines for different roles. Among those roles are:

- application servers, that will run the ELMA system;
- DBMS servers, where the database will be deployed;
- controller-server for distributing the load on ELMA servers and user requests;
- server for storing documents, attachments, temporary files and configurations.

Fig. 1 shows the architecture of the web farm, which is used as an example in this manual.

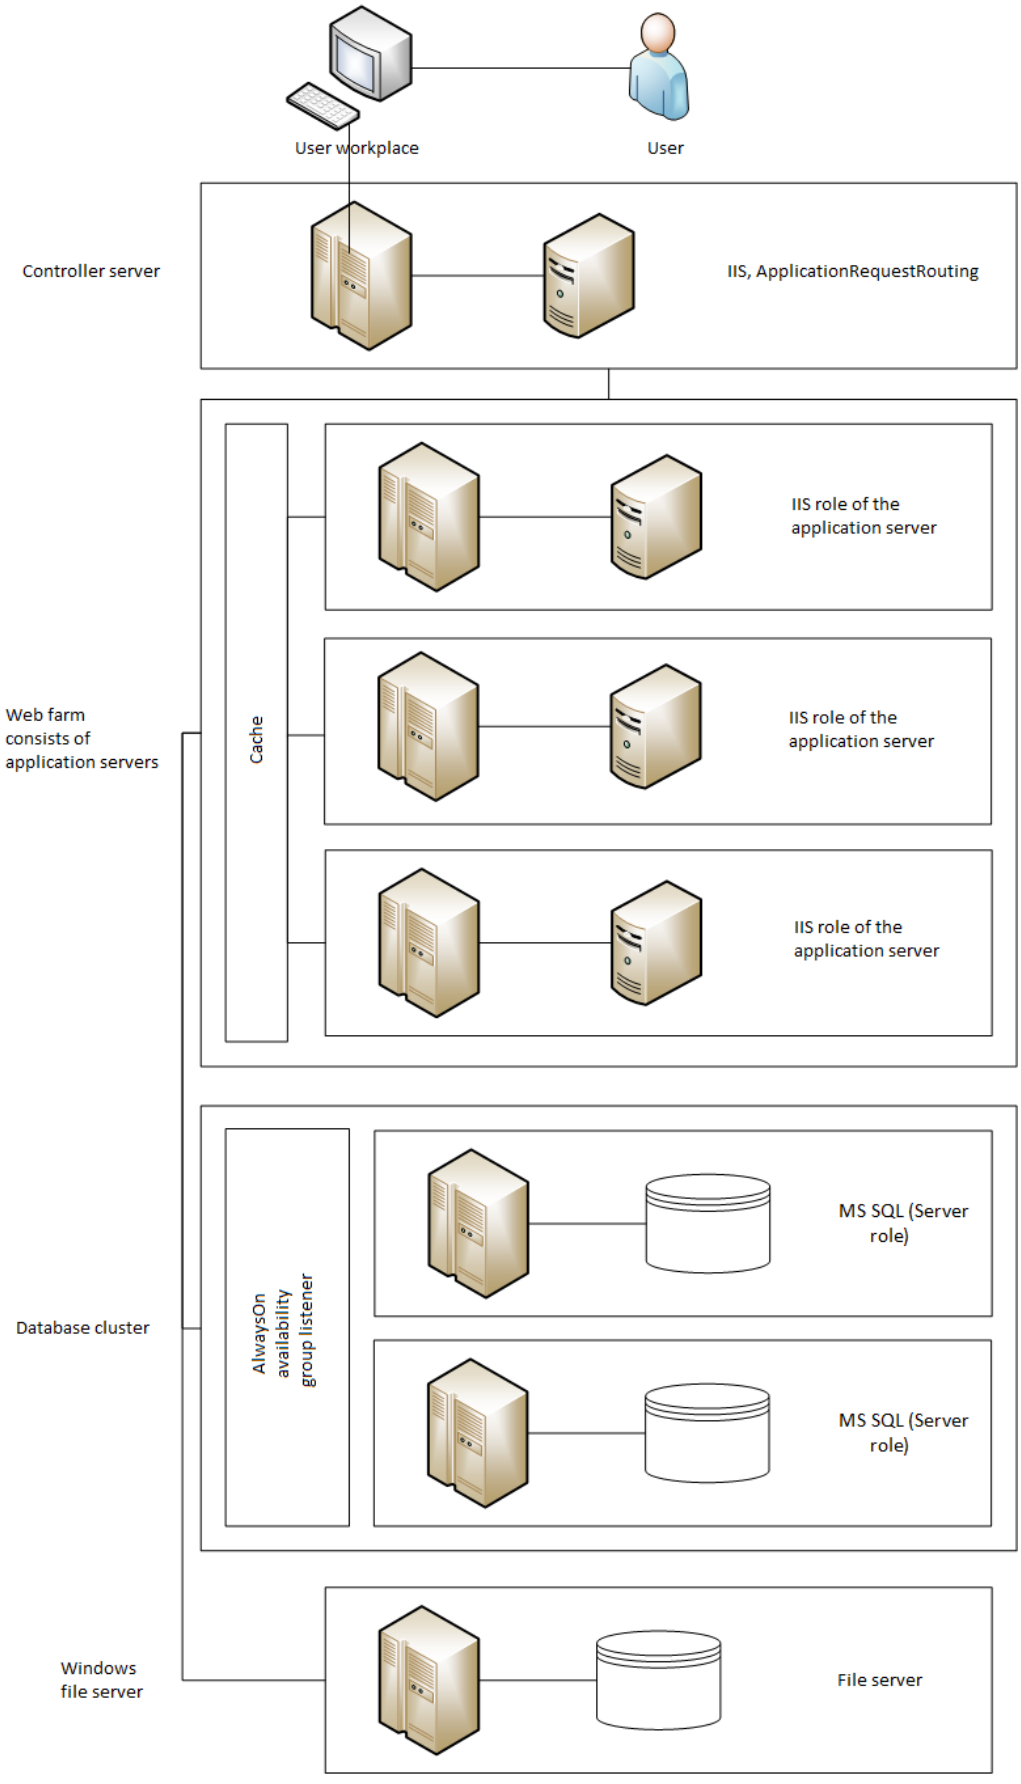

Fig. 1 Web farm architecture

Chapter 2. Servers Used

When you deploy the system to work under high load, you are supposed to use different server machines for different roles. Among those roles are:

- application servers, that will run the ELMA system;
- DBMS servers, where the database will be deployed;
- controller-server for distributing the load on ELMA servers and user requests;
- server for storing documents, attachments, temporary files and configurations.

Servers can be virtual, under the condition that enough resources will be allocated and CPU and memory time reservation will be provided.

Attention! Do not try to save on reservation for virtual servers, use a burdened or an old host-server. If necessary, you will be able to reduce resources later, when the system operates at full capacity. It is much simpler, than gather statistics and diagnose the reasons behind performance drops.

If you expect high load, then you must take into account the peculiarities of using shared resources of the host machine in case of virtual servers. You can find some advice in the section 6.1.

All the servers, listed below, must be in the same domain (ELEWISE, or elewise.local) and you must use a domain account with the local administrator permissions on each of the servers.

It is strongly discouraged to change network addresses and domain names of these servers after deploying the services. Normally, the server domain names will be specified. However, to avoid issues and a long debugging procedure, it is discouraged to change server network addresses after you start configuring the system.

Application servers:

- 93.158.134.3- SRV12-1 – SRV12-1.elewise.local
- 213.180.204.3- SRV12-2 – SRV12-2.elewise.local

Farm controller server:

- 93.158.134.30- WFCONTROLLER – WFCONTROLLER.elewise.local

Server for files:

- 93.158.134.35- ELMAConfig – ELMAConfig.elewise.local

Database servers:

- 192.168.18.230 – DBCLUSTERTEST01 – DBCLUSTERTEST01.elewise.local – primary replica
- 192.168.18.23 – DBCLUSTERTEST02 – DBCLUSTERTEST02.elewise.local – secondary replica

Database cluster network address:

- 192.168.18.25- SRV12DB – SRV12DB.elewise.local

Domain name for the availability group listener in the network:

- 192.168.18.40- SRV12DBLst – SRV12DBLst.elewise.local

Account for all the services:

- EleWise\ELMAadmin

Note! Do not change the password of this account unless it is necessary!

Since many services work under the same account, changing the password will require stopping all the services, and therefore change account information. Note that restoring all the services may take several hours.

Chapter 3. Installing and Configuring Failover Cluster for MS SQL Databases

This section contains step-by-step instructions for installing and configuring MS SQL Server and for combining two database servers into a single failover cluster.

It is important to understand, that at least one MS SQL server must be installed for the system to work, while the failover cluster is an additional feature for preventing fails and data losses.

Attention! A domain account with permissions to create computers in the domain will be required at one of the steps of installing and configuring the failover cluster.

3.1. Installing and configuring MS SQL Server

Before creating and configuring a failover cluster, you must install and configure MS SQL Server. You should install and configure ELMA web farm only after performing all the settings, described in Chapter 3. Later on, to switch ELMA from using a single MS SQL server to using a database cluster you will only need to change the address in the configuration file.

You need to install MS SQL Management Studio, which is included in most MS SQL Server packages. You can also download it from the official Microsoft website. It is recommended that you use MS SQL Server 2008 R2 or higher.

3.1.1. Installing MS SQL Server

Attention! You need to apply the procedures described below to all database servers you intend to use.

Step 1. Start SQL Server installation center.

After starting the installation software, select **Install** in the left menu and click **New installation or add features to an existing installation** (Fig. 2).

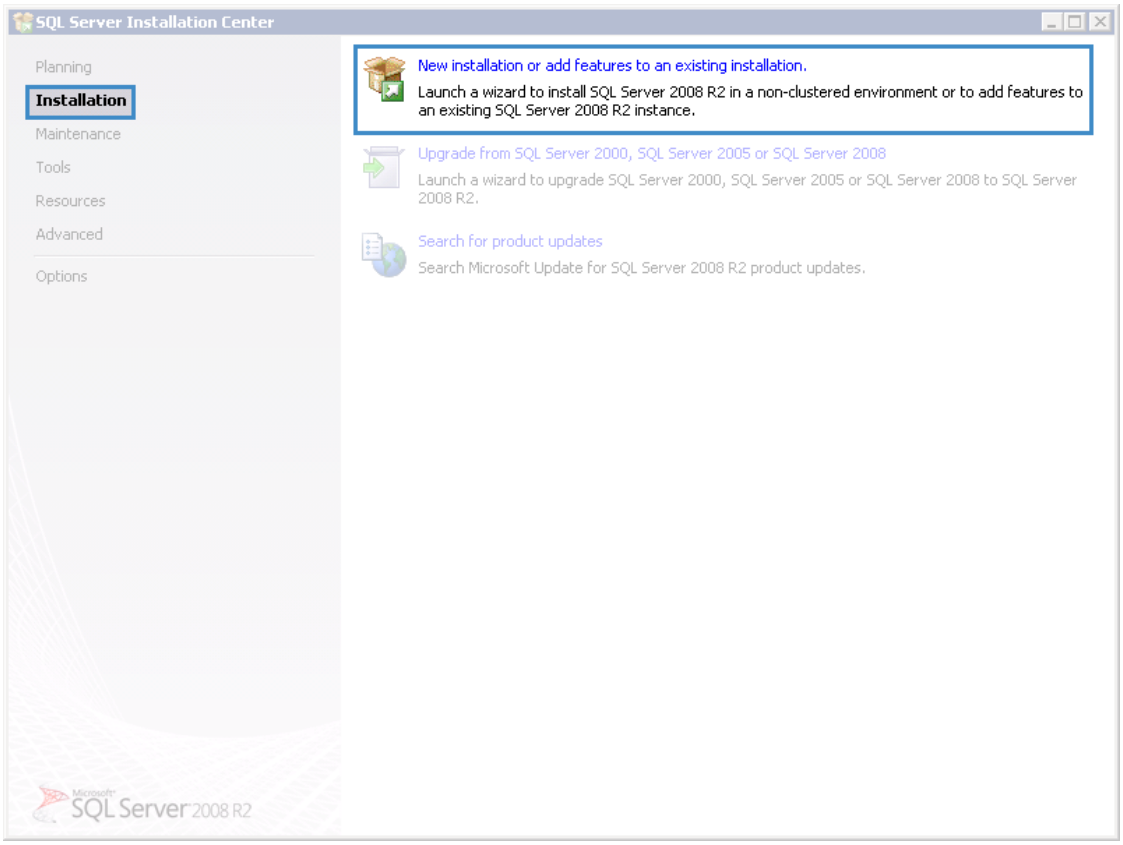

Fig. 2 SQL Server installation center

Step 2. Wait until **Installation wizard** (Fig. 3) initializes the installation files and completes checking global rules. To continue installation, click **Next**.

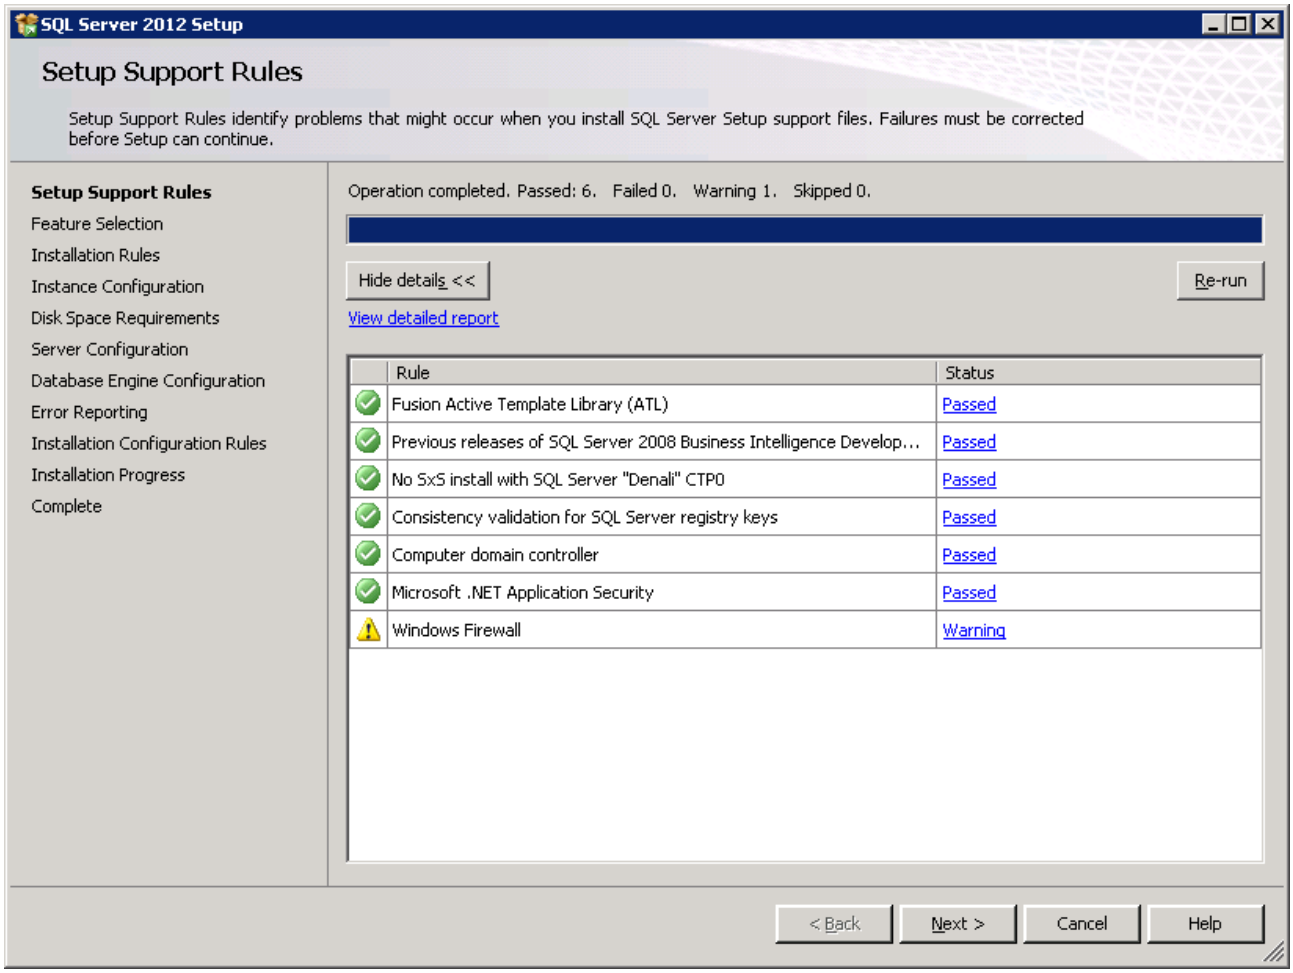

Fig. 3 Installation wizard. Setup Support Rules

If there are warnings and/or errors, you must eliminate them before continuing setup. After that, restart the check by clicking **Re-run**.

Step 3. Select a licensing option and click **Next**.

Step 4. Read the licensing agreement.

After reading the agreement, if you agree with it, select **I accept the license terms** (Fig. 4). To continue, click **Next**.

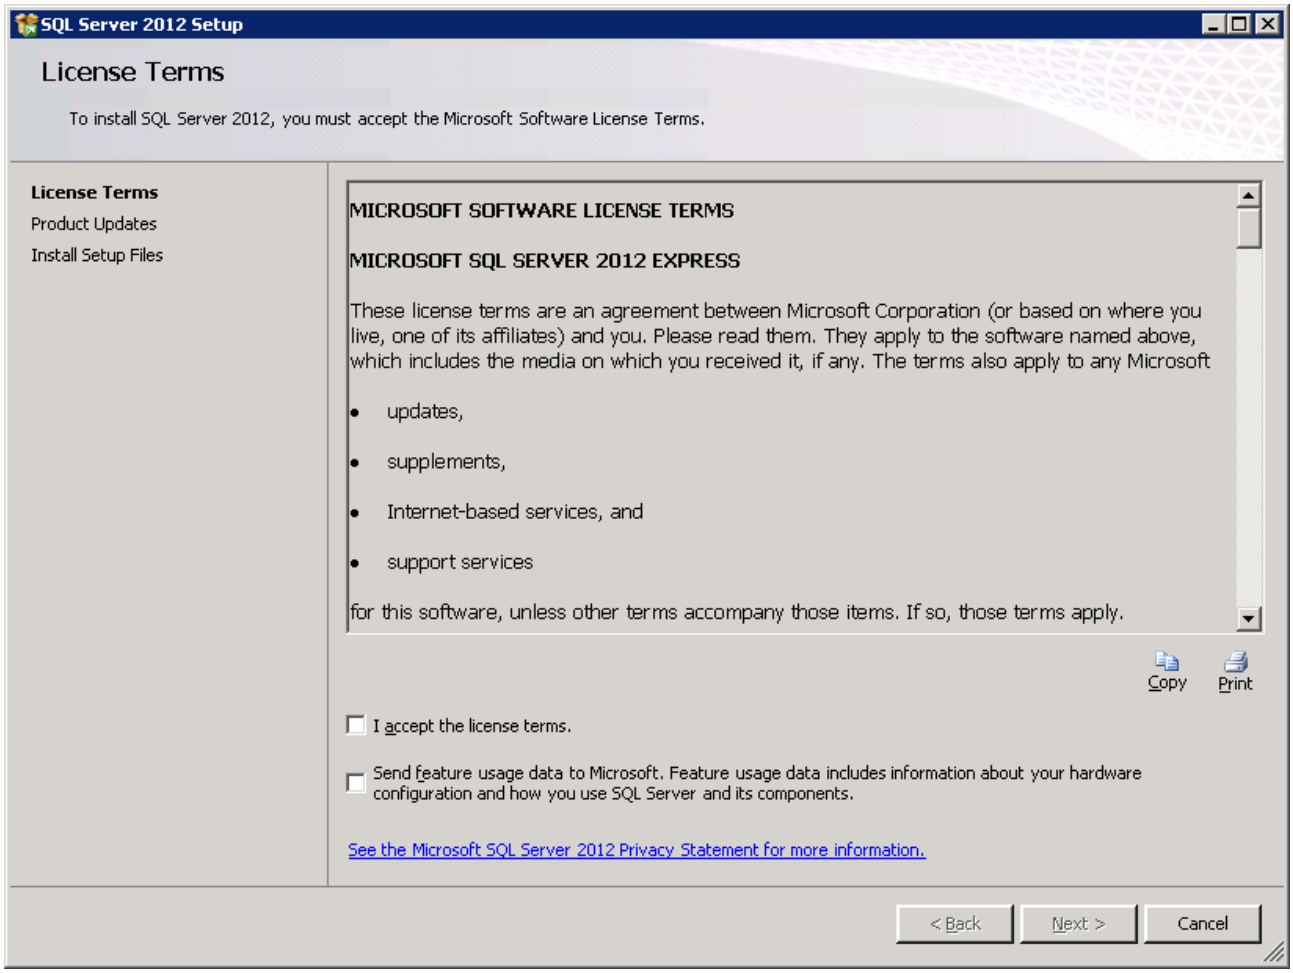

Fig. 4 Installation wizard. Licensing agreement

Step 5. Select components to install.

Check the boxes to left of the components you need. Fig. 5 shows the list of components that must be installed for the system to operate correctly. Other components are optional.

The list of components depends on the version of the SQL Server.

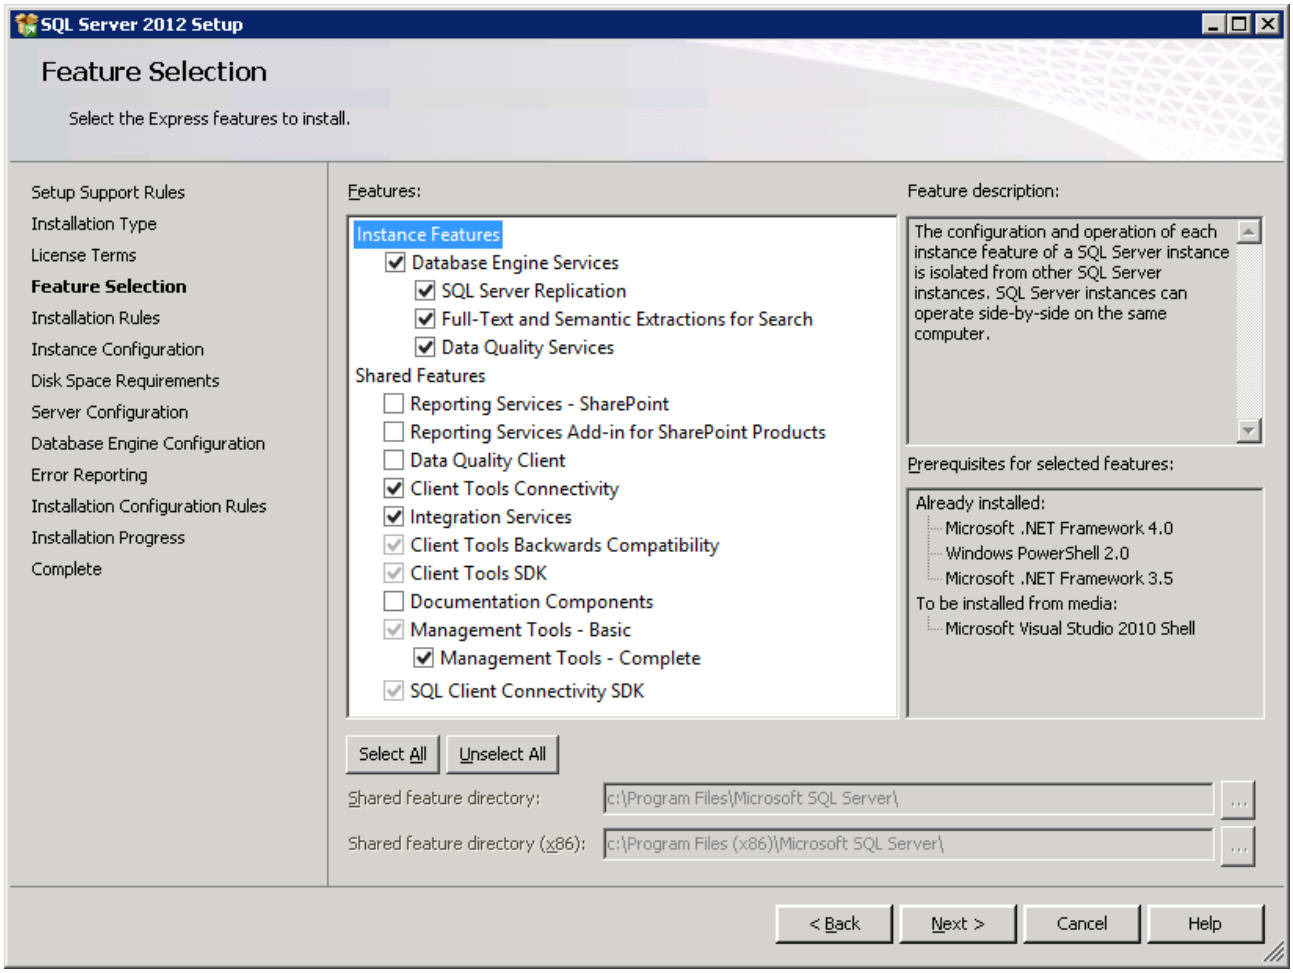

Fig. 5 Installation wizard. Selecting components

It is recommended that you use the default directory for installation. To change the default directory, click "**...**", and select another directory. To continue installation, click **Next**.

Pay attention to the Management Tools. You must install basic management tools and MS SQL Management Studio, which is included in complete management tools.

Step 6. Select an instance to install.

You can install several MS SQL servers on one machine, but we recommend that you install and use the default instance (Fig. 6).

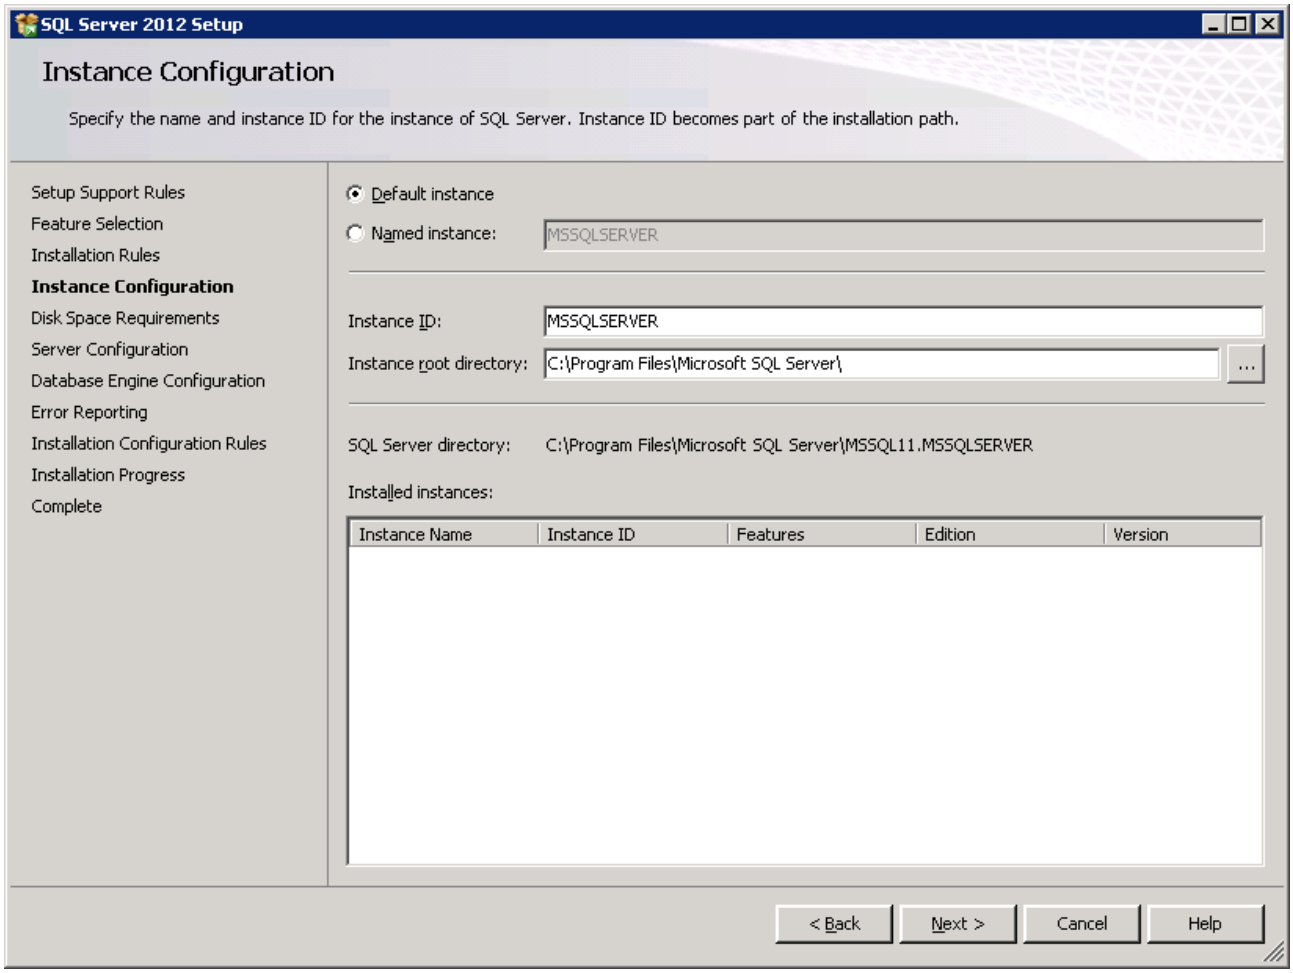

Fig. 6 Installation wizard. Instance configuration

You can select Named instance if it is necessary, but it will make the database connection string more complex. You will have to enter the connection information correctly every time it is required by installation of other services and applications.

If an error occurred at this step, you can find possible solutions in the section 6.2.

To continue installation, click **Next**.

Step 7. Make sure that there is enough disk space for the installation of the components.

Step 8. Leave the default settings at the **Server Configuration** step (you can change the account later). To continue installation, click **Next**.

Step 9. Switch authentication mode.

Select **Mixed Mode (SQL Server authentication and Windows authentication)**, and specify a password for logging in as the system administrator (sa), and add the current Active Directory user (Fig. 7).

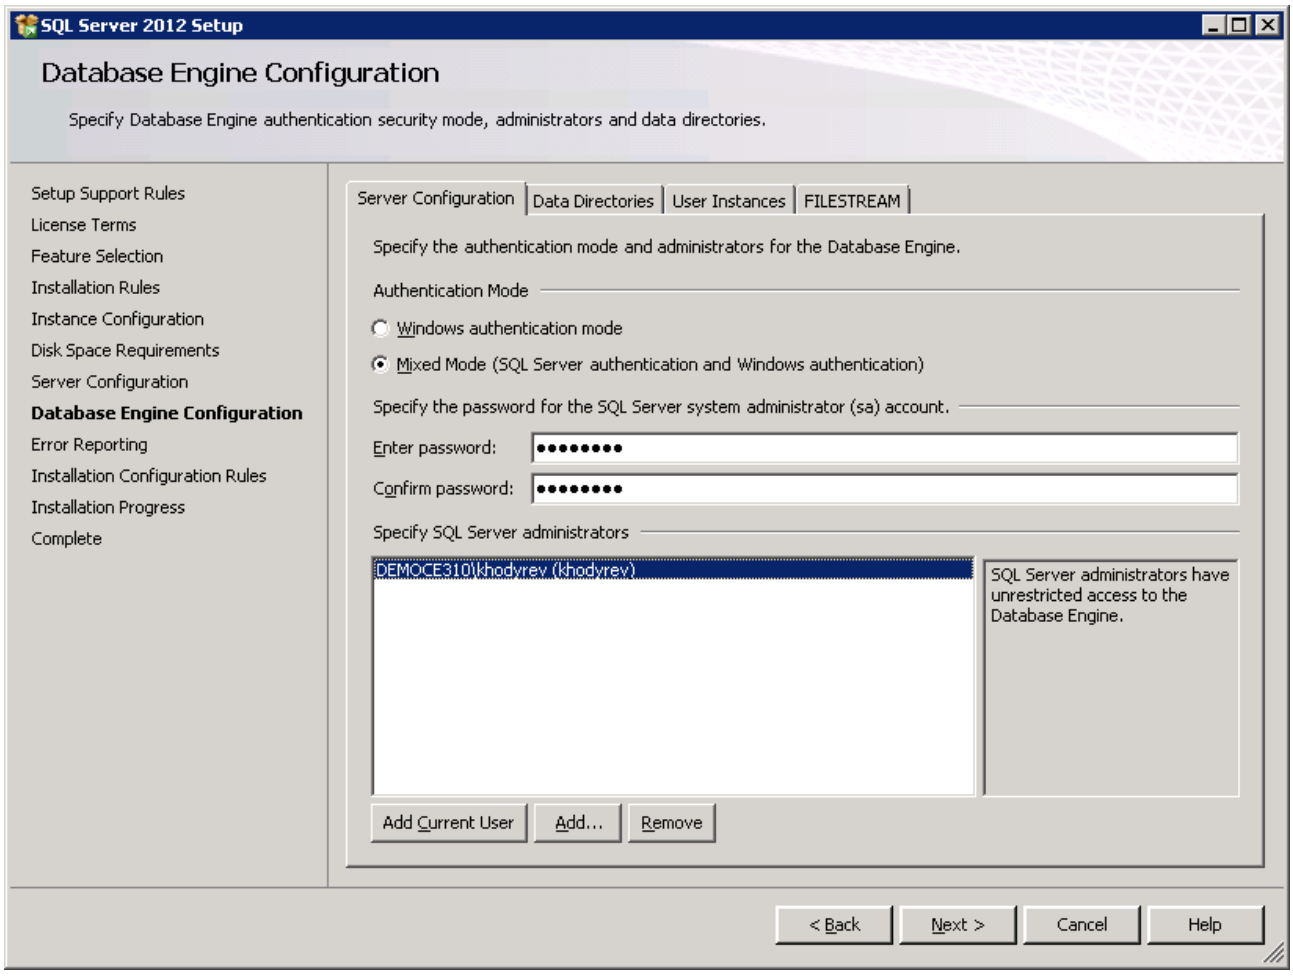

Fig. 7 Installation wizard. Database Engine configuration

To continue installation, click **Next**.

Step 10. Skip the following two steps by clicking **Next**.

Step 11. Check the selected components.

At this step (Fig. 8), you can review the list of components selected for installation at the previous steps. If necessary, you can change this list. To do so, return to the required step of the wizard by clicking **Back**.

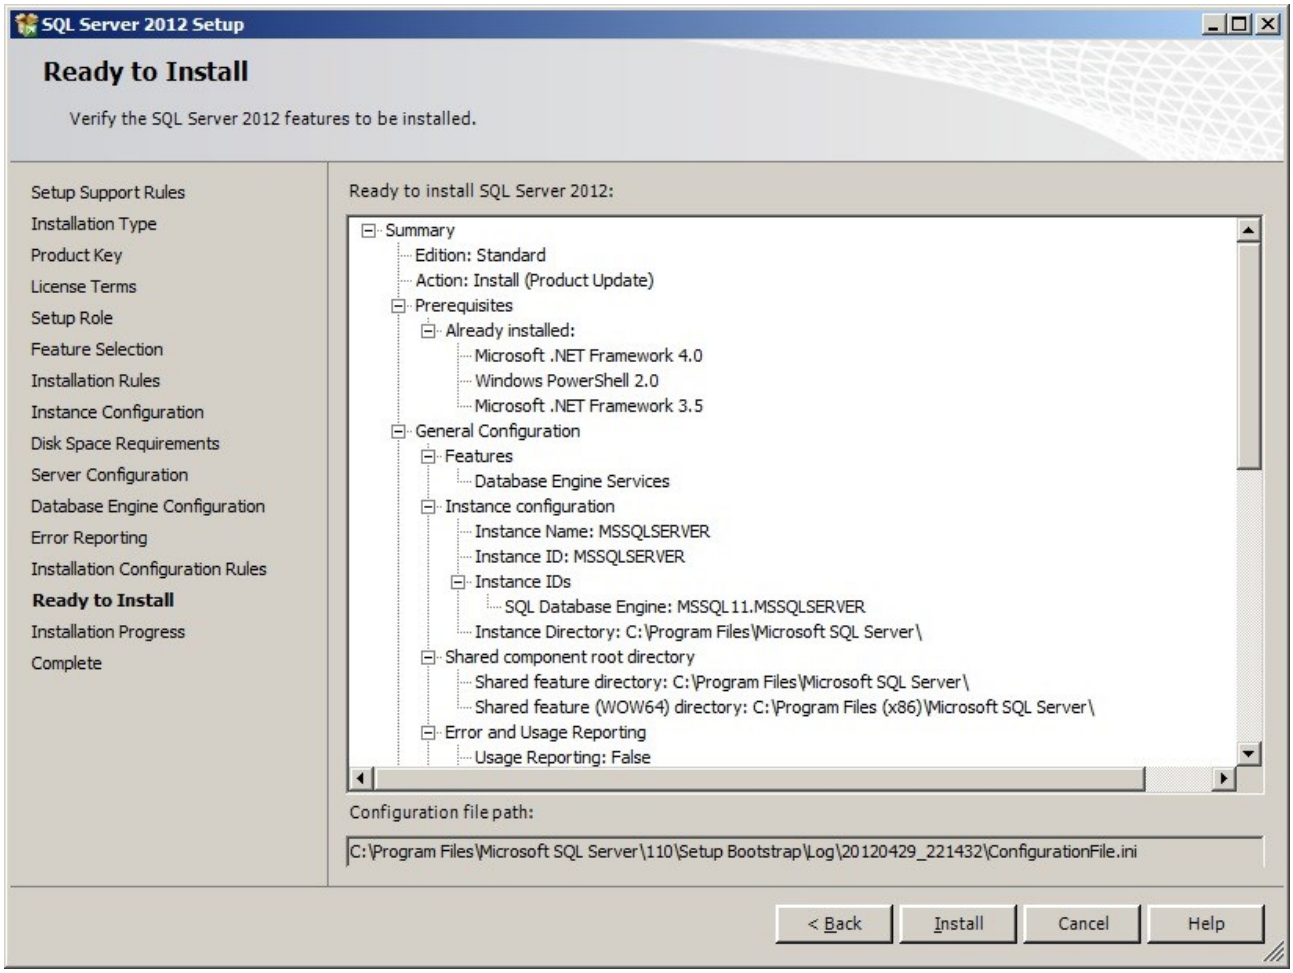

Fig. 8 Installation wizard. Ready to install

To start installation, click **Install**.

Step 12. Installation.

At this step, the components are being installed on the computer. This process can take several minutes.

Once the installation has completed, click **Next** to go to the next step.

Step 13. Complete the installation.

At this step, the installation wizard displays the results of installing MS SQL Server. To complete the installation process, click **Close**.

3.1.2. Configuring MS SQL Server authentication

Attention! You need to apply the procedures described below to all the database servers you intend to use.

This procedure is necessary to provide the possibility to log in to the ELMA server and cache services database using different accounts (single sign on for the cache and local account for the ELMA server).

This difference is due to the cache peculiarities. It is recommended using a domain account for cache. At the same time, you should use a local DBMS account for the ELMA server in order to avoid the Double-Hop Problem.

In practice, there is only one possibility when the Double-Hop problem may occur for ELMA – when using single sign on. Quite likely, your module configuration excludes the problem; however, we recommended taking safety measures beforehand.

Step 1. Start MS SQL Server Management Studio.

Start MS SQL Server Management Studio (**Start -> SQLServerManagementStudio**) and log in using the **sa** account (Fig. 9), specified during the installation.

In the **Authentication** field, select **SQL Server Authentication** (Fig. 9) from the drop down list. The login and password fields will become available.

Fig. 9 Connect to Server dialog box

Check the **Remember Password** box.

The first start may take several minutes.

Step 2. Configure server authentication.

In MS SQL Server Management Studio, right click on the **(local)(SQLServer)** pool and select **Properties** in the context menu (Fig. 10).

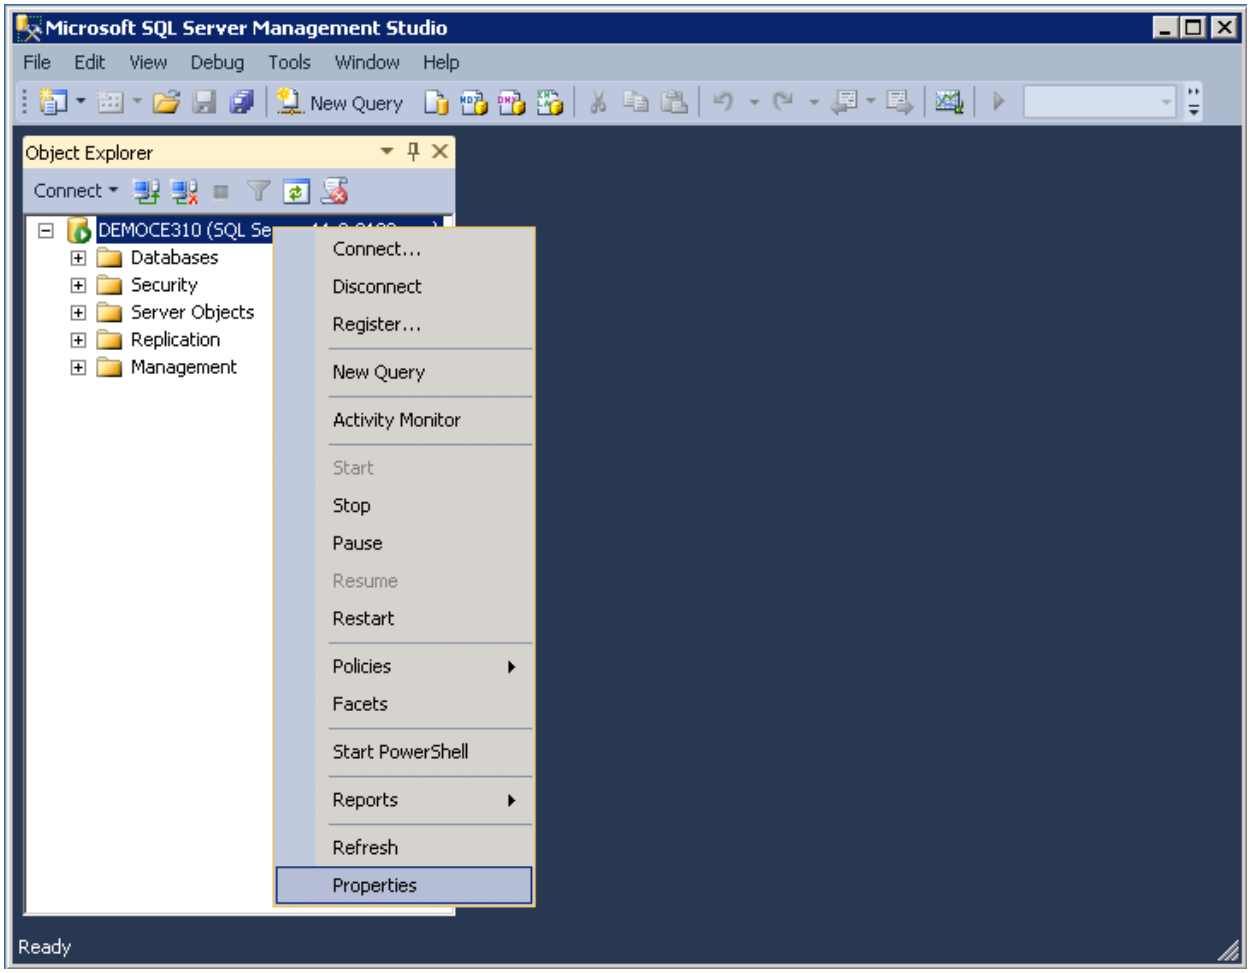

Fig. 10 Context menu of the "(local)(SQLServer)" pool

In the opened dialog box, select the **Security** page and make sure, that **SQL Server and Windows authentication mode** is selected in the **Server Authentication** unit (Fig. 11).

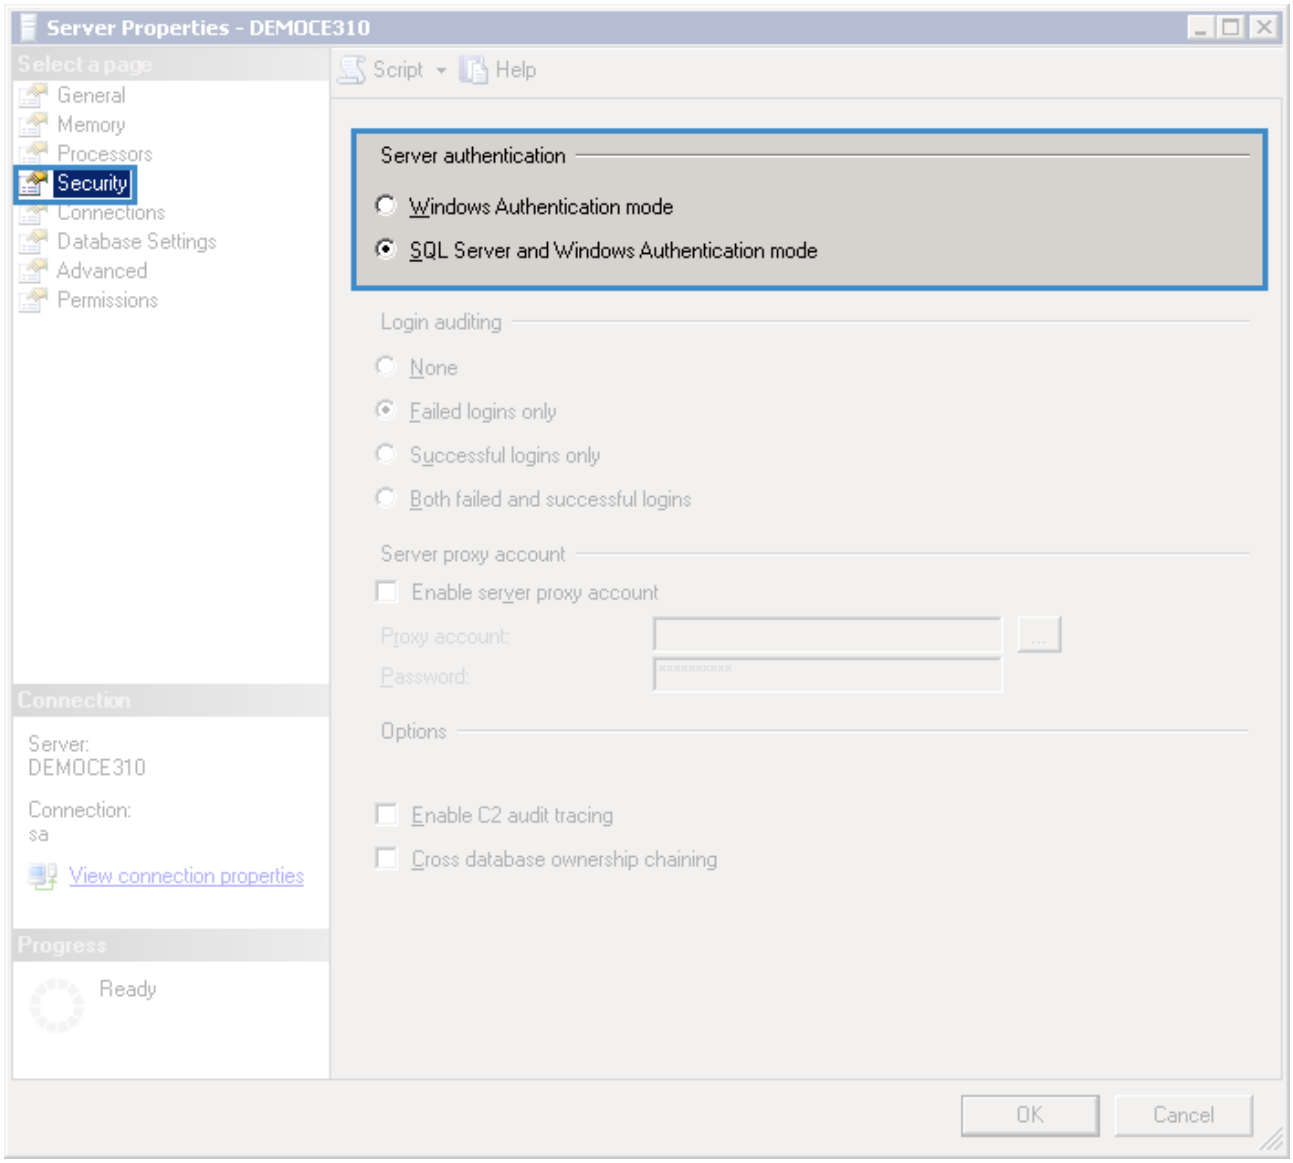

Fig. 11 Security page. Server authentication unit

To confirm the settings and close the dialog box, click **OK**.

Step 3. Create a login.

In the object explorer of MS SQL Server Management Studio, open the context menu of the **Security -> Logins** and select **New Login**... (Fig. 12).

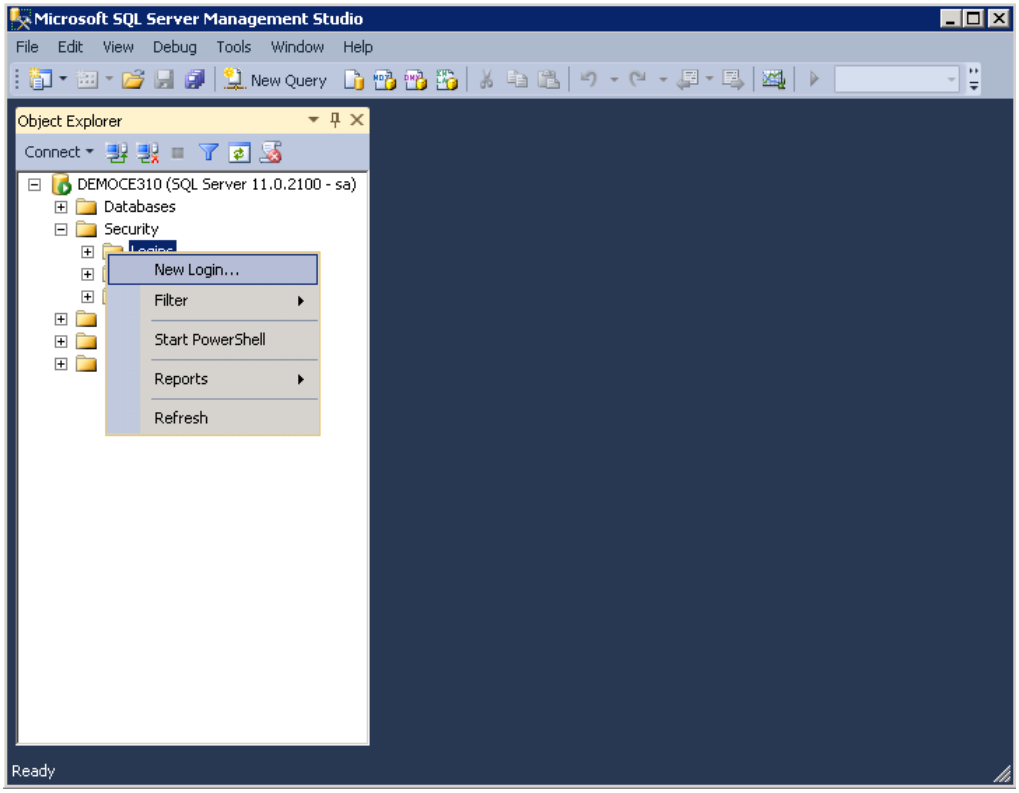

Fig. 12 Context menu. New Login… item

In the opened dialog box, select the **General** page and specify the domain and name of the required user (Fig. 13).

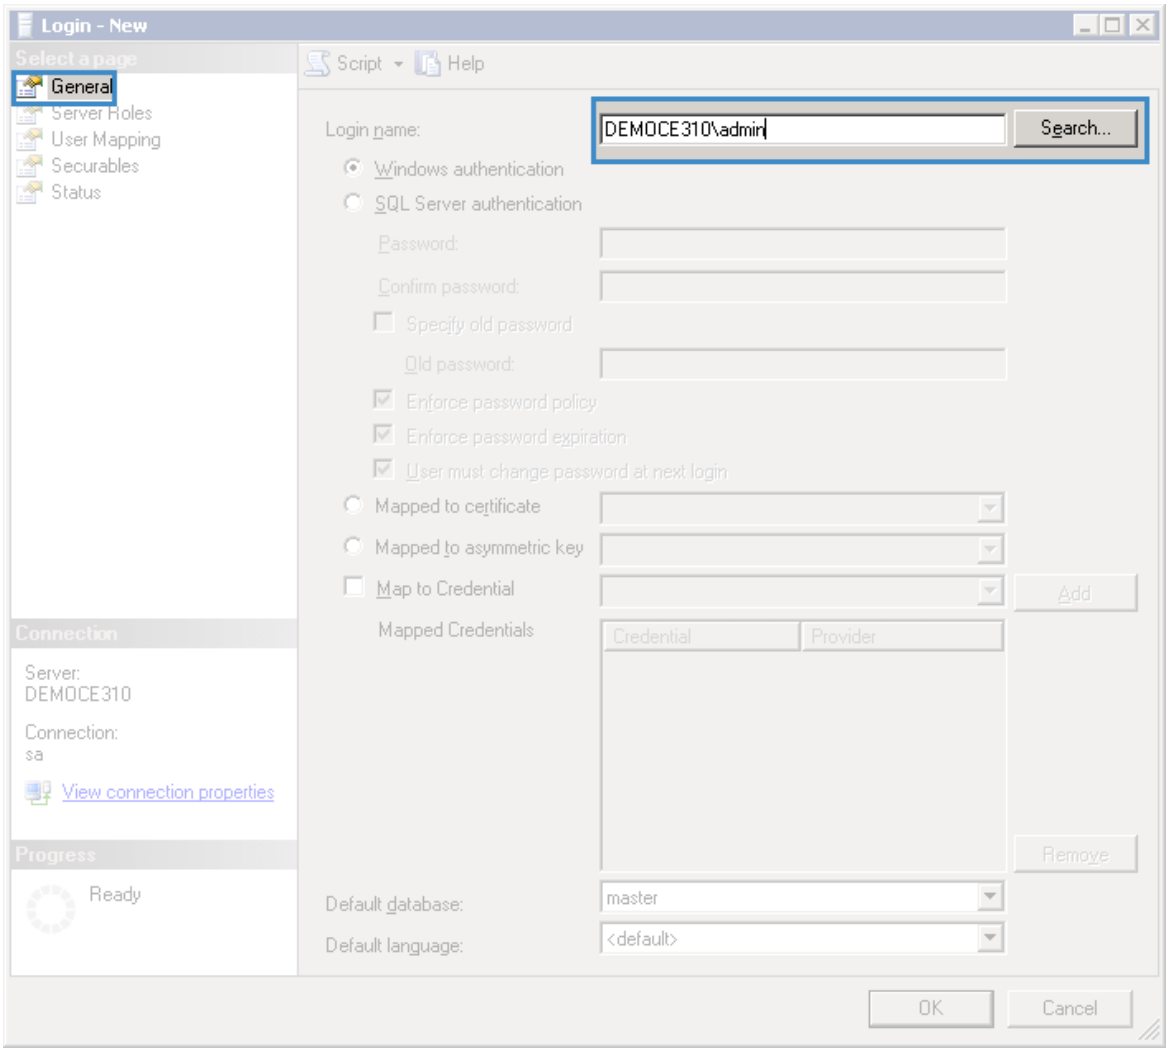

Fig. 13 Creating a login. General page

Next, open the **Server Roles** page and check all the available boxes (Fig. 14).

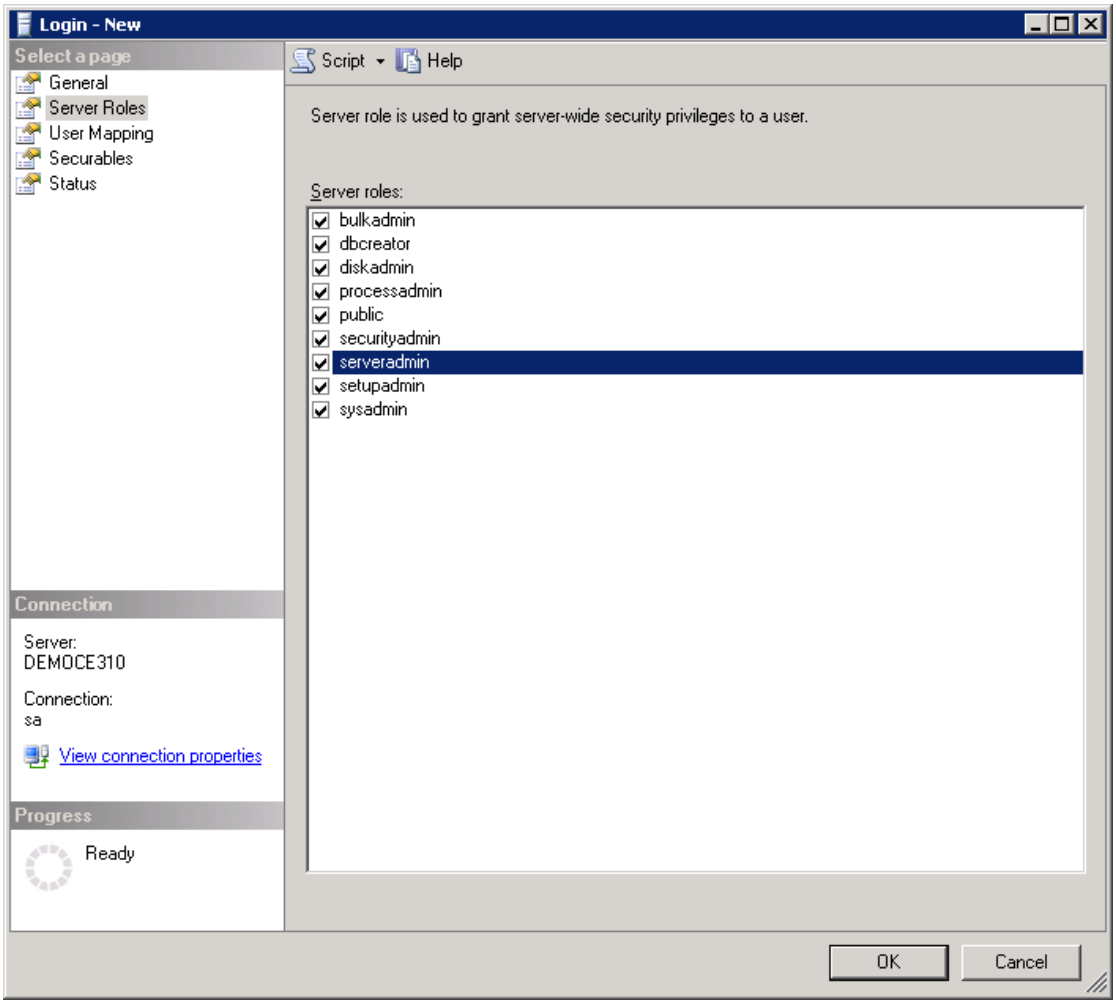

Fig. 14 Creating a login. Server Roles page

To save the changes and close the dialog box, click **OK**. After that, you can log in to the system using the created login.

3.1.3. Change the MS SQL Server services account

For the MS SQL server to work correctly in the failover cluster, you need MS SQL services to be started under one domain account.

Attention! You need to apply the procedures described below to all database servers you intend to use.

In this example, 192.168.18.23 and 192.168.18.230.

Step 1. Add the computer to the domain.

The computer must be included in the domain. The domain account, which will be used to log in, must be included in the administrators group of this server.

In this example, the account **EleWise\ELMAadmin** is used.

Step 2. Close MS SQL Server Management Studio and open SQL Server Configuration Manager. Go to **Start -> SQL Server Configuration Manager** (Fig. 15).

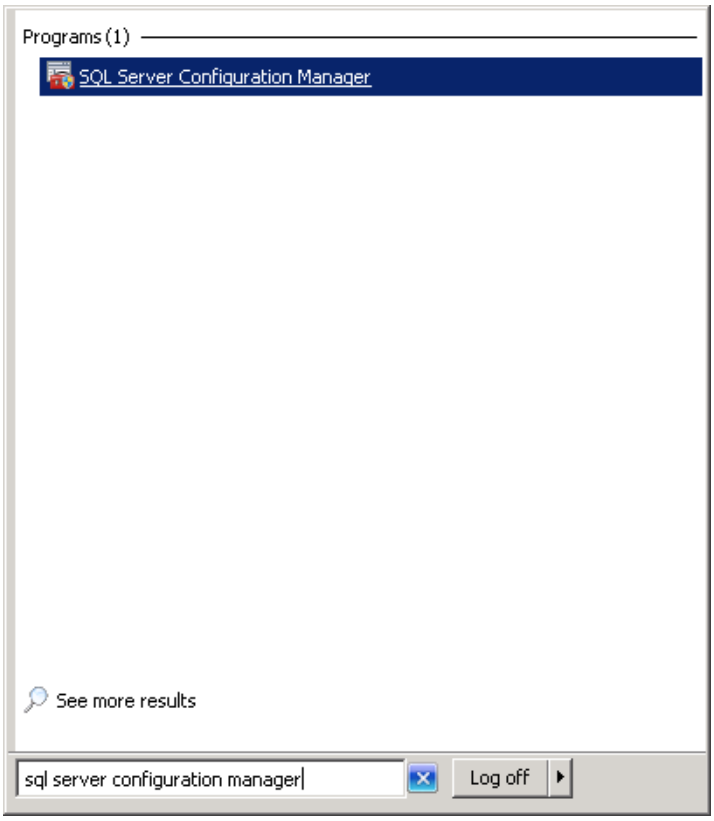

Fig. 15 Start -> SQL Server Configuration Manager

Step 3. Stop the SQL server.

In the opened window, select **SQL Server Services** and right click **SQL Server (MSSQLSERVER)** to open its context menu. In the context menu, click **Stop** (Fig. 16).

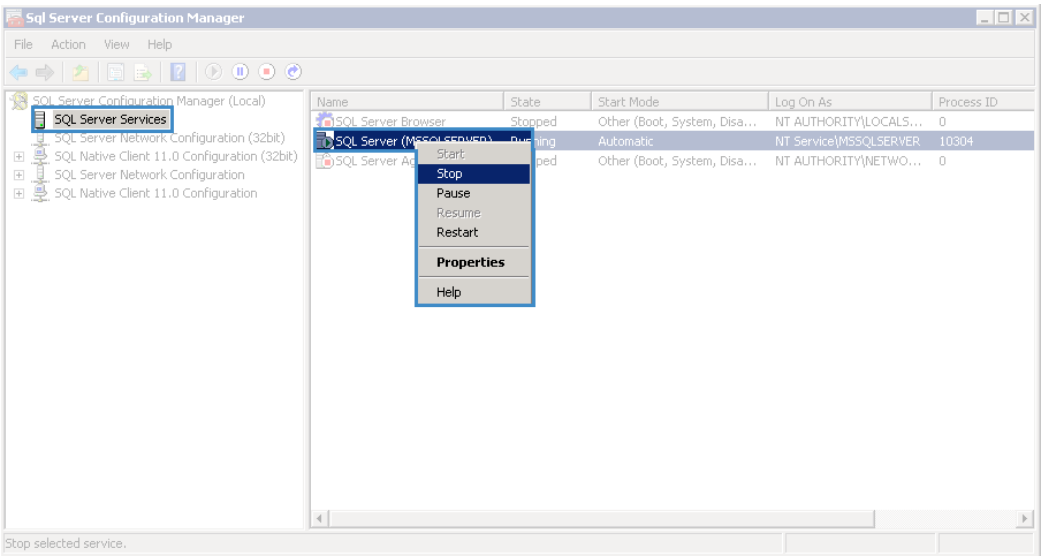

Fig. 16 SQL Server Configuration Manager. Stopping the SQL server

When you stop the SQL server, the SQL Server Agent will also stop.

Step 4. Configure SQL server properties.

Open the properties window of the server, by selecting **Properties** in its context menu (Fig. 17).

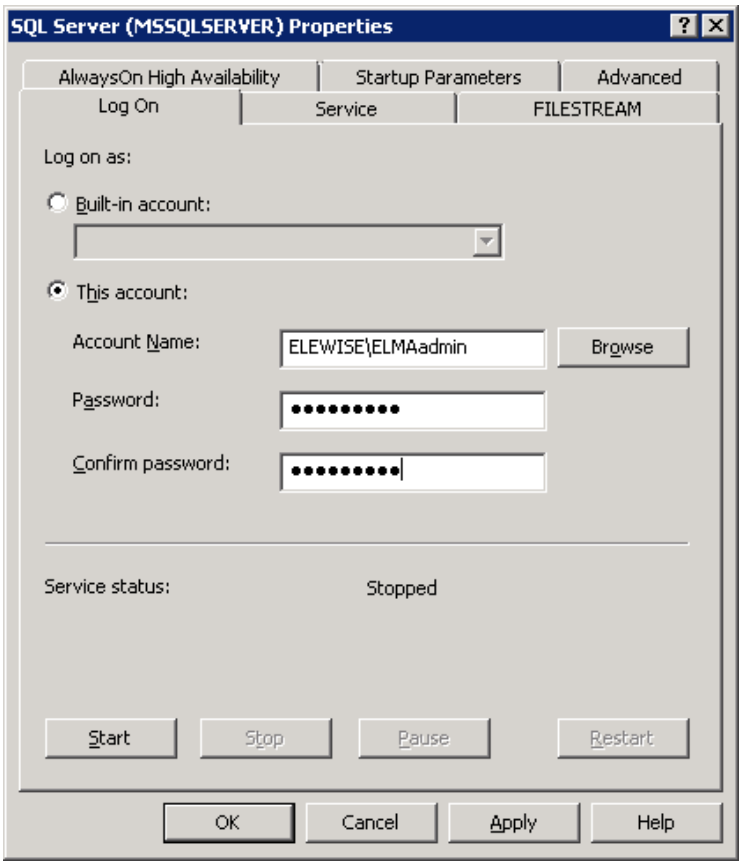

Fig. 17 SQL server properties

In this dialog box, open the **Log On** tab, select **This account** and enter the account information for logging in to the system. In this example, it is **EleWise\ELMAadmin**.

To save the changes and close the dialog box, click **OK**.

Step 5. Start the SQL server.

Start the SQL server by clicking **Start** in the context menu of **SQL Server (MSSQLSERVER)** (Fig. 16).

Repeat steps 3-5 for the **SQL Server Agent (MSSQLSERVER)**.

3.1.4. Create a database for ELMA

Attention! You need to apply the procedures described below to all database servers you intend to use.

In this example, it is 192.168.18.230.

To use ELMA, you must create a database. In this example, it will be named **ELMAPROD38**, but you can name it whatever you like. This name will be used in the ELMA configuration file.

To create a new database, in the object explorer of MS SQL Server Management Studio right click on **Databases** and select **New Database...** in the context menu (Fig. 18).

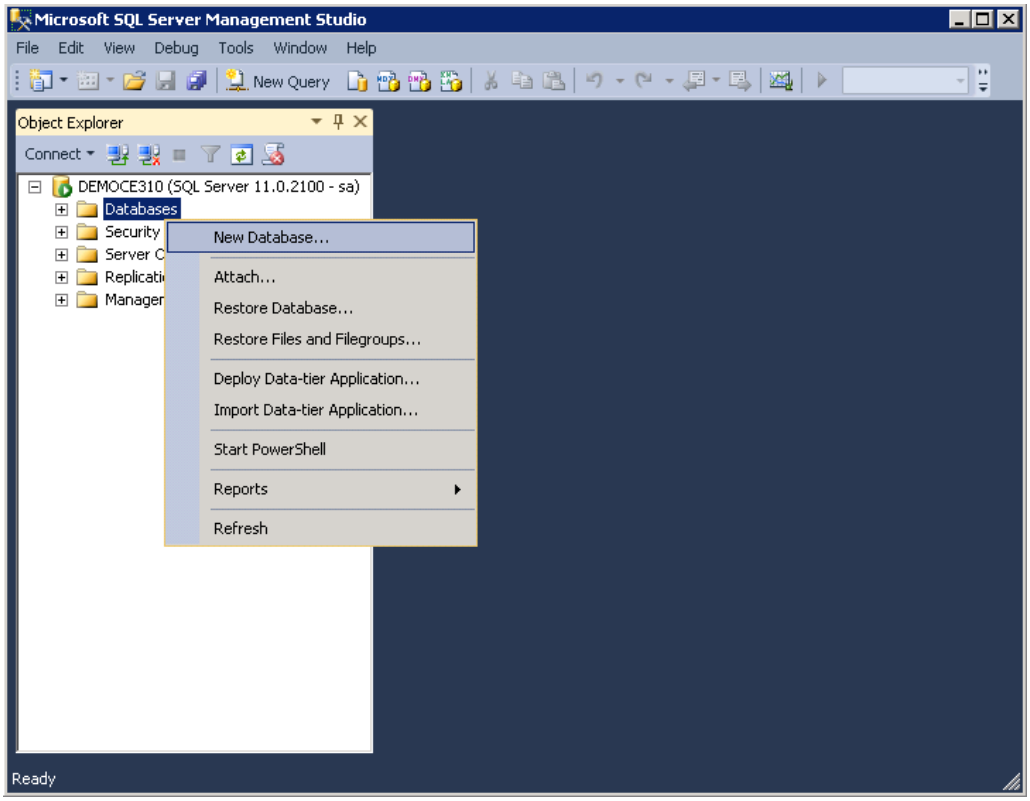

Fig. 18 Context menu. New Database… item

In the opened dialog box, select the **General** page and enter the database name in the **Database name** field (Fig. 19). In this example, it is **ELMAPROD38**.

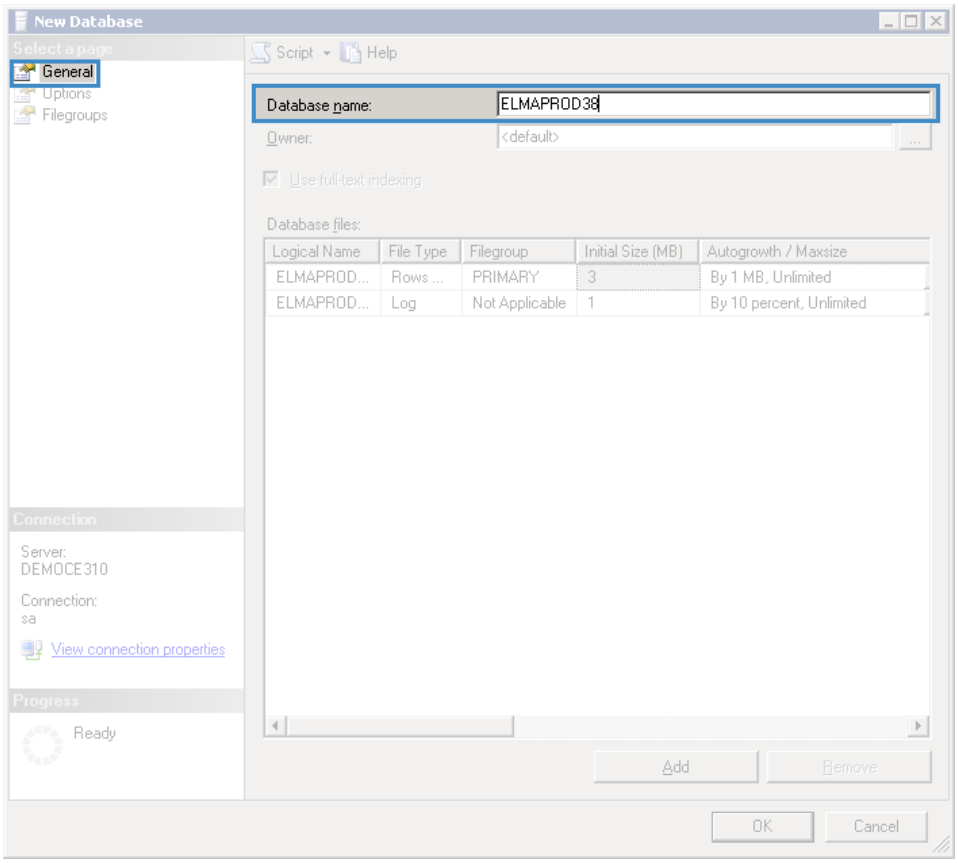

Fig. 19 Creating a database. General page

Next, open the **Options** page and perform the following settings (Fig. 20):

- in the **Recovery model** field, select "Full". The full recovery model is required for participating in the AlwaysOn availability group of the failover cluster. It also means that you will need to make backup copies of the transaction log regularly (to learn more, see the section 5.1).
- in the **Compatibility level** field, select **SQL Server 2012 (110)**. You need to set the compatibility level due to the fact, that core changes were made in the MS SQL Server 110 and higher, which caused some select queries to work longer than usual.

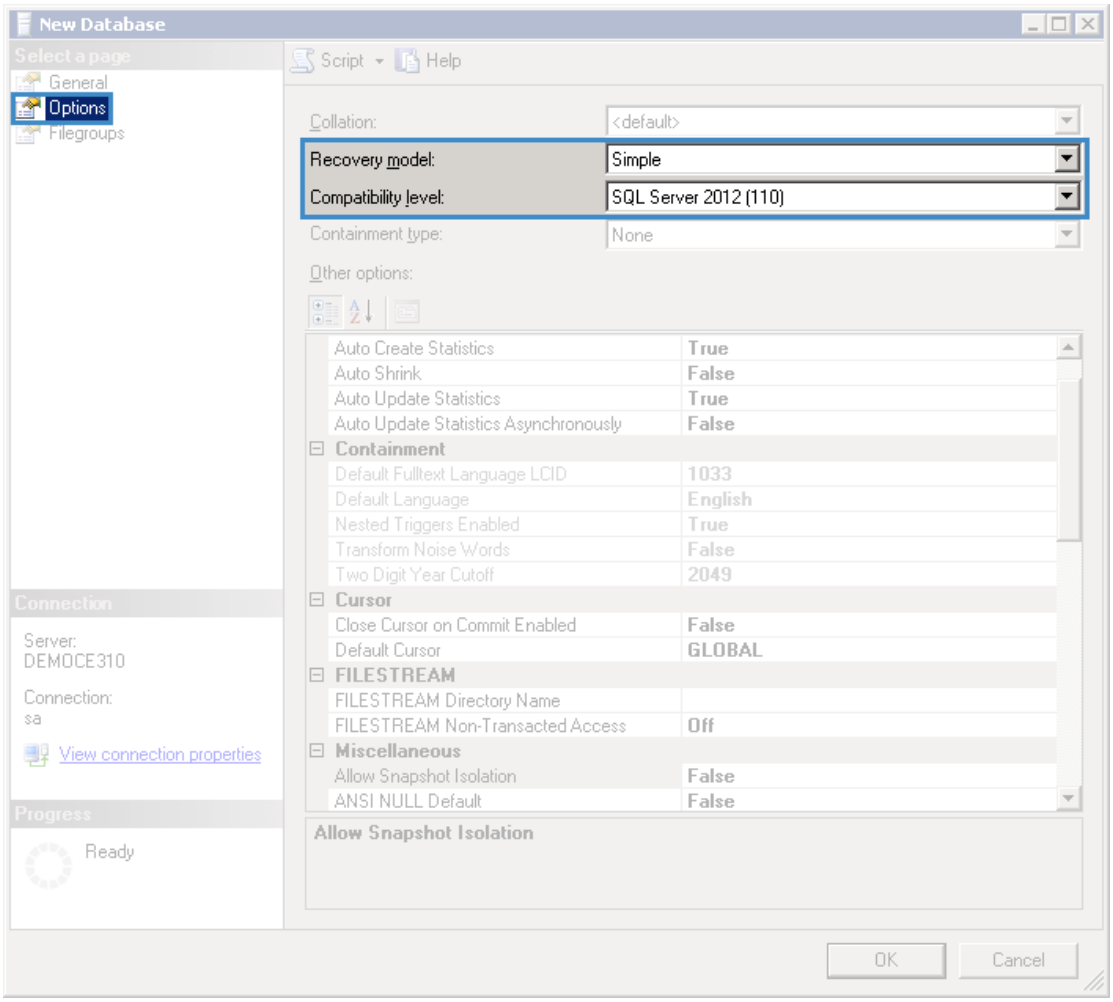

Fig. 20 Creating a database. Options page

Note, that for ELMA to work with MS SQL SERVER 2014, you need to select the **MSSQL SERVER 2012** compatibility level.

To save the database and close the dialog box, click **OK**.

3.1.5. Creating a database for ASPState

Attention! You need to apply the procedures described below to all database servers you intend to use.

In this case, it is 192.168.18.230.

An ASPState database is required for ELMA web farm. There are several implementation options for ASPState.

We highly recommend using an ASPState database with permanent table structure storage, since it is the only option that can be successfully used in a failover cluster.

Step 1. Log in to the main database server using the ELMAadmin account. This account must already have been added to sysadmin in MS SQL Server.

Step 2. Open the command prompt **Start -> cmd.exe**

Step 3. An ASPState database is created with a special **.NET Framework** utility. To run it in an x64 Windows OS use the following command:

C:\Windows\Microsoft.NET\Framework64\v4.0.30319\aspnet_regsql.exe -S <server name>\<server instance name> -E -ssadd -sstype p

where: **aspnet_regsql.exe** – utility call. By default, the ASPState database creation wizard opens, but it is unnecessary, since the parameters are known.

S – server on which the database should be installed;

E – use the account of the current user;

ssadd – key for creating a new ASPState database;

sstype p – database type with permanent structure storage.

By default, the database is created in RAM and in case of a restart, it has no structure. Due to this, you will have to manually initialize it again.

3.2. Creating Failover Cluster

3.2.1. Installing the failover component

Attention! You need to apply the procedures described below to all database servers that will be included in the failover cluster.

In this example, these are the 192.168.18.23 and 192.168.18.230 servers.

Step 1. Open the Server Manager (**Start -> Server Manager**), go to **Dashboard**, click **Manage** and select **Add Roles and Components** (Fig. 21).

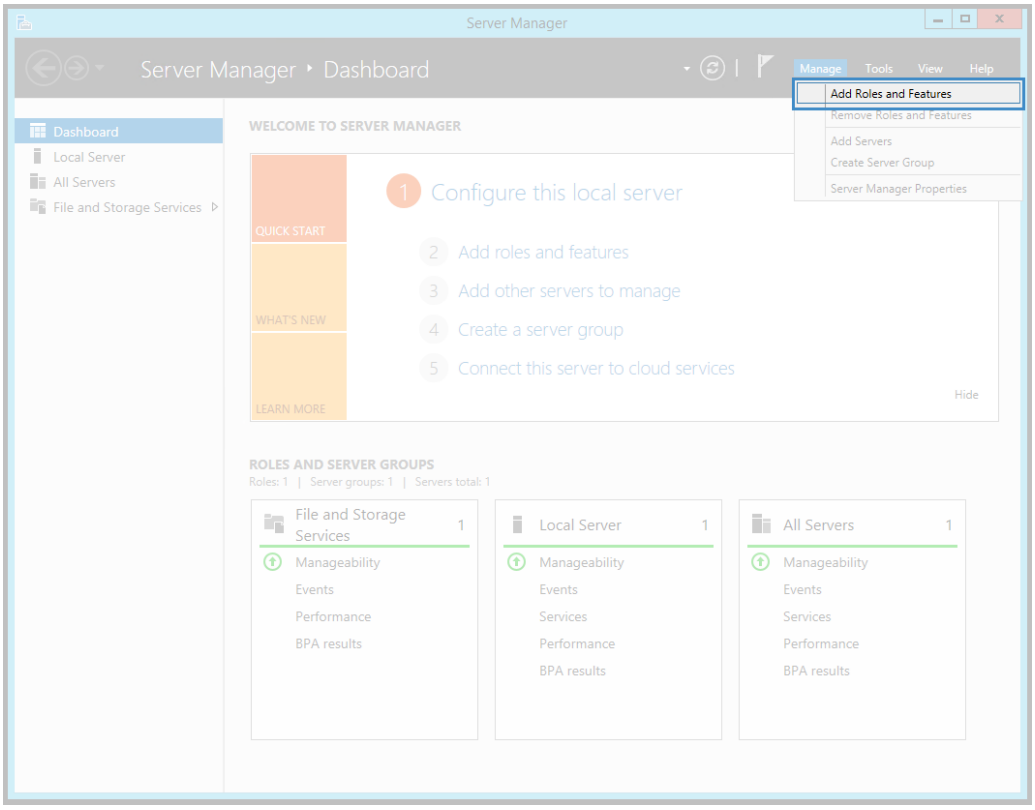

Fig. 21 Server Manager. Dashboard. Manage – Add Roles and Features button

In the wizard for adding roles and features, skip all the steps until the **Features** step, by clicking **Next**. At the **Features** step, check the **Failover Clustering** box (Fig. 22) and click **Next.**

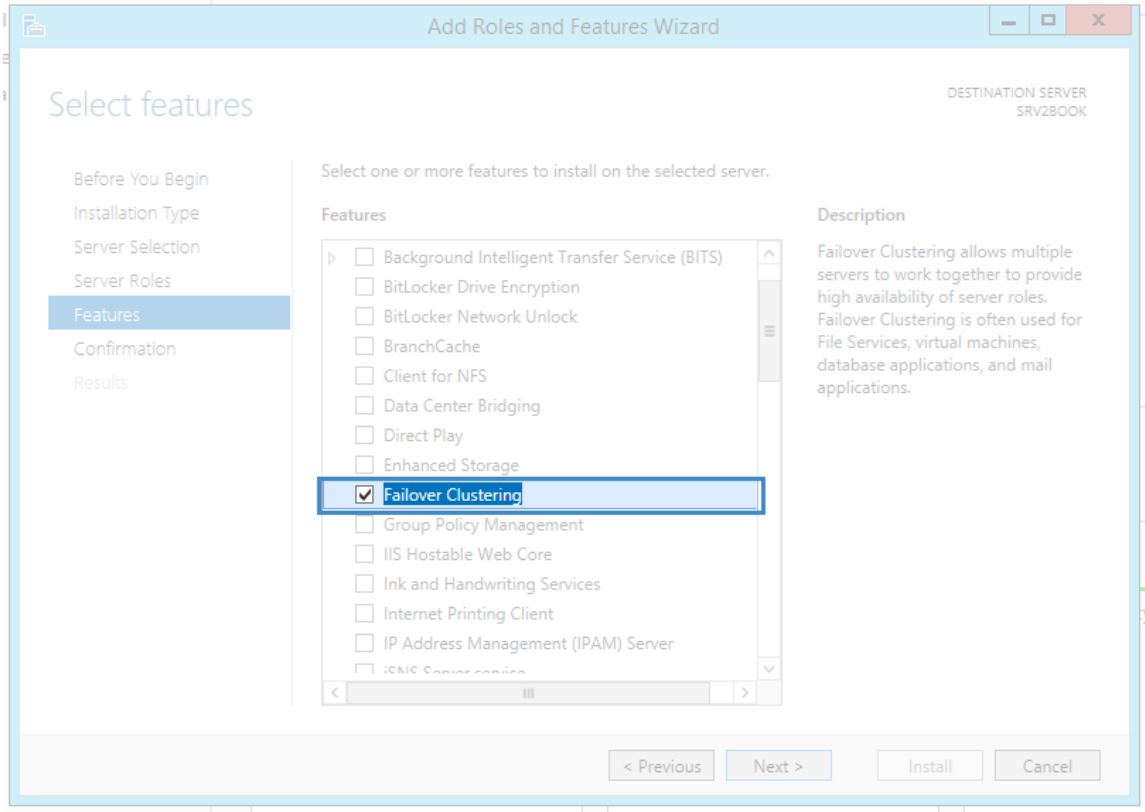

Fig. 22 Add Roles and Features Wizard. Features step

To start installation of the selected feature, click **Install** at the **Confirmation** step. The installation process may take a long time, wait until it is complete.

Step 2. In the Server Manager, open **Dashboard**, click **Tools** and select **Failover Cluster Manager** (Fig. 23).

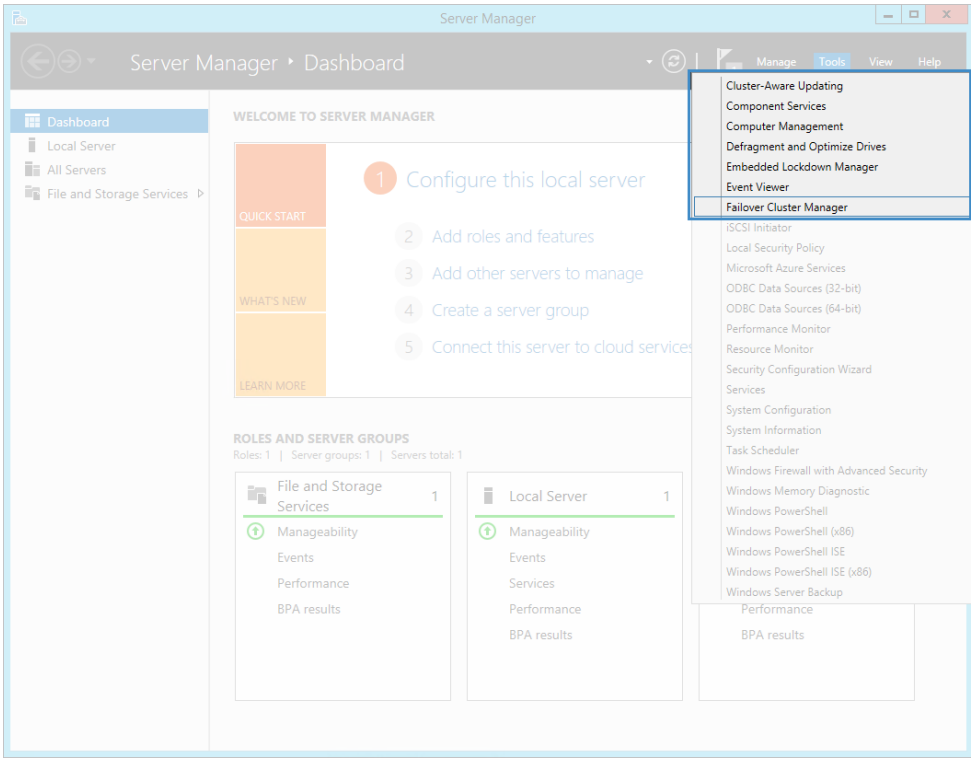

Fig. 23 Server Manager. Dashboard. Tools – Failover Cluster Manager button

If you click on this item and an error occurs (Fig. 24), log in to the system using the domain user account (the server must be added to the domain).

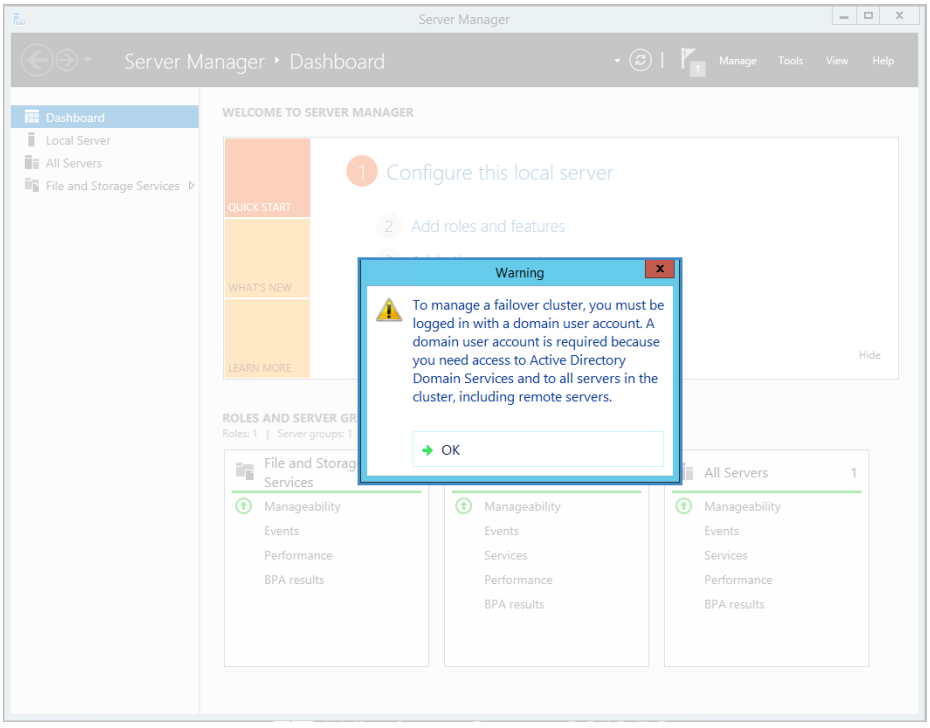

Fig. 24 Server Manager. Access error

The following steps of creating a failover cluster will require an account with permissions to create computers in the domain.

3.2.2. Creating a failover cluster

Now you can start creating and configuring a failover cluster.

Attention! You need to apply the procedures, described below, only to one database server that will be included in the failover cluster.

In this example, it is the server 192.168.18.230.

Step 1. Log in to the system using the domain account with permissions to create computers in the domain. You can grant these permissions to the EleWise\ELMAadmin account (if there aren't any) or use the system administrator's account to carry out these steps.

Step 2. Open the Failover Cluster Manager (Fig. 25).

To do so, in the Server Manager open **Dashboard**, click **Tools** and select **Failover Cluster Manager** (Fig. 23). You can also open this manager via **Start -> Failover Cluster Manager.**

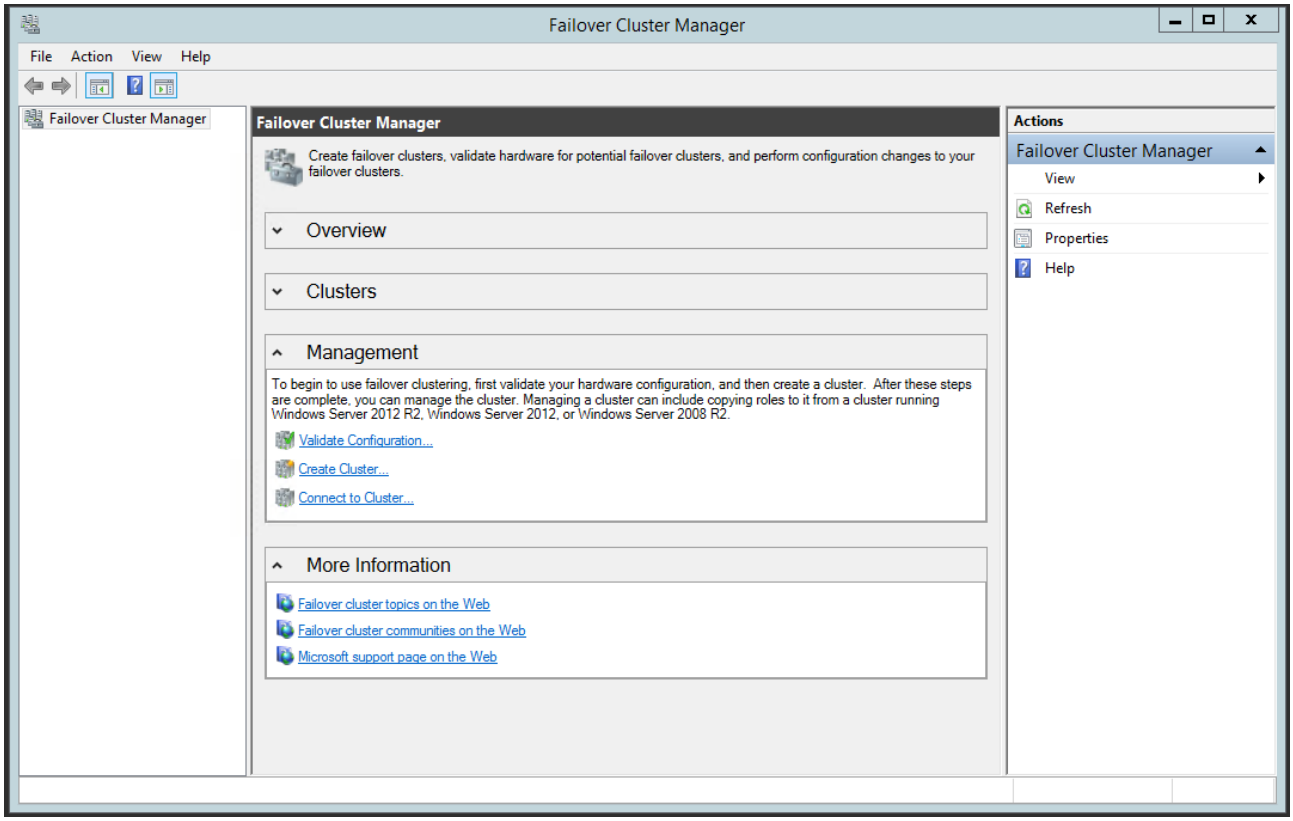

Fig. 25 Cluster Failover Manager

Fig 3. Create a failover cluster.

Attention! The failover cluster creation wizard also starts the configuration validation wizard, while it stays at the initial steps.

To create a failover cluster, click **Create a Cluster...** in the work pane of the **Failover Cluster Management** panel. The cluster creation wizard will open. The first step **Before you begin** is a welcome screen, you can skip it by clicking **Next**.

At the second step, **Select Servers** (Fig. 26), add the names of the required servers and click **Next**. Make sure that the failover feature is installed on all of the selected servers.

In this example, these are the servers 192.168.18.23 and 192.168.18.230. Order is irrelevant.

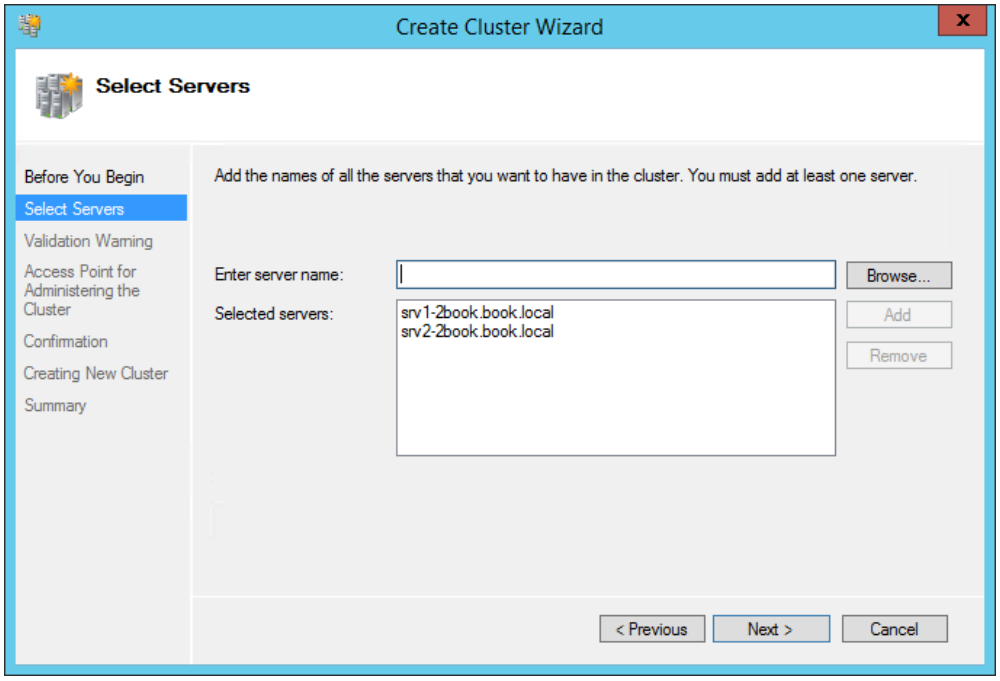

Fig. 26 Cluster creation wizard. Select Servers

At the **Validation Warning** step (Fig. 27) select **Yes** and click **Next**.

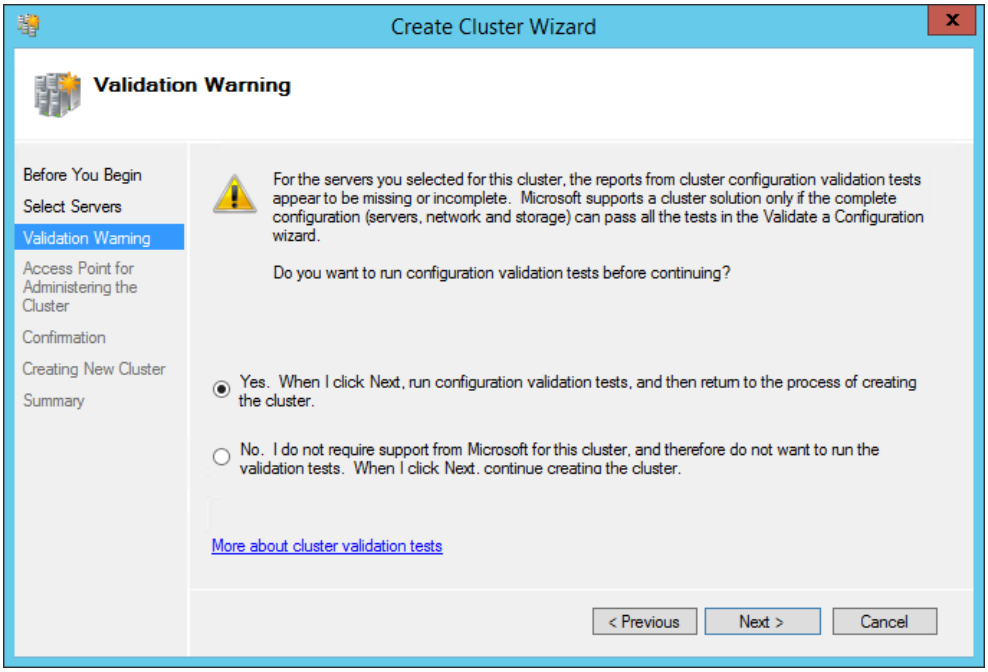

Fig. 27 Cluster creation wizard. Validation Warning

This switch starts the configuration validation wizard (Fig. 28).

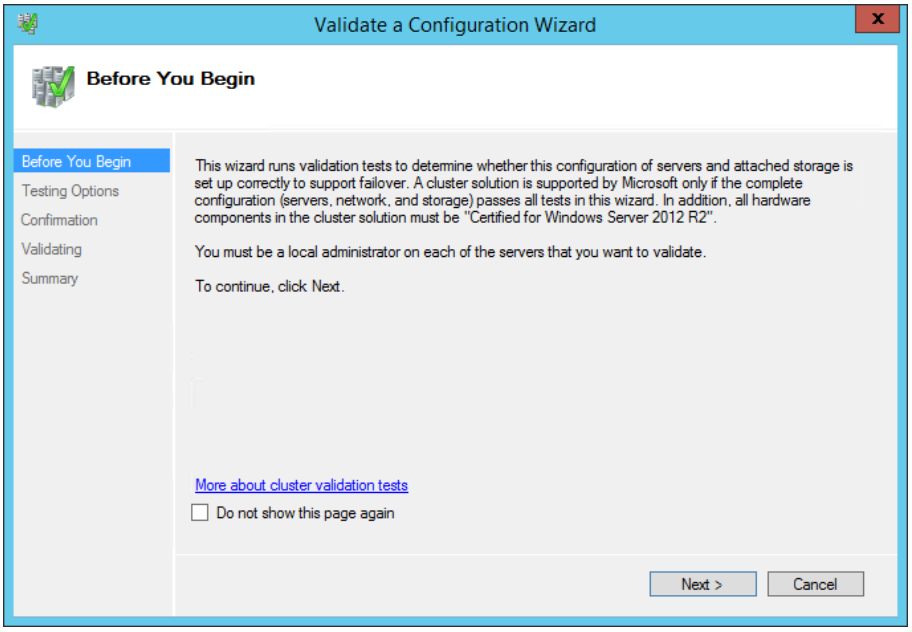

Fig. 28 Configuration validation wizard. Before you begin

Leave the settings at the following steps unchanged and skip them by clicking **Next**.

After that the configuration will be validated and a report will be generated (Fig. 29). To view the report, click **View Report**. The report contains recommendations on improving the failover cluster.

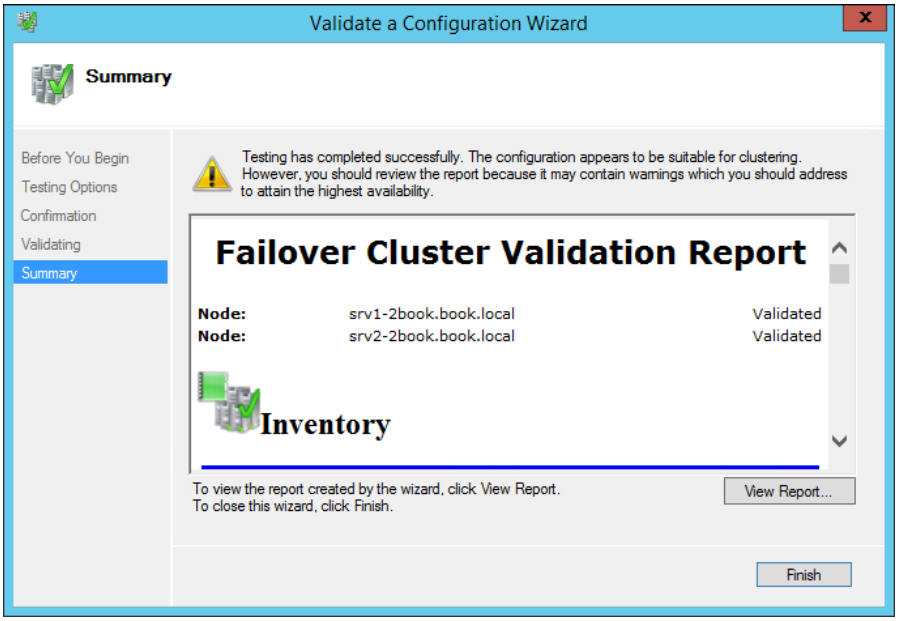

Fig. 29 Configuration validation wizard. Summary

Take into account the warnings and make the changes according to the recommendations. If there are any errors, eliminate them.

To close the configuration validation wizard and return to the cluster creation wizard (Fig. 30) click **Finish**.

At the **Access Point for Administering the Cluster** step (Fig. 30), enter the cluster name that will be used for administering it, and click **Next**.

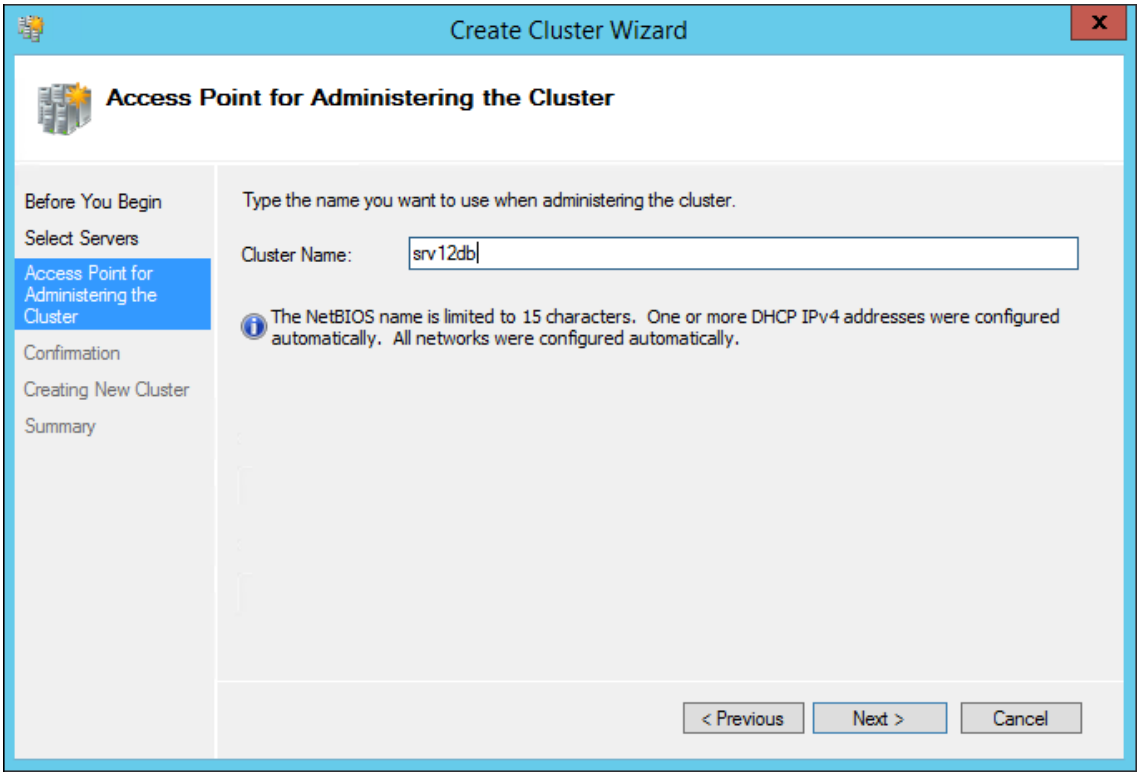

Fig. 30 Cluster creation wizard. Access Point for Administering the Cluster

In this example, it is **srv12db**. You are going to specify this server name when connecting to the cluster. The name of the availability group listener will be used for connecting to the database. The listener will be created later.

At the **Confirmation** step (Fig. 31), you can review the settings of the new cluster, specified at the previous steps. To continue creating the cluster, click **Next**.

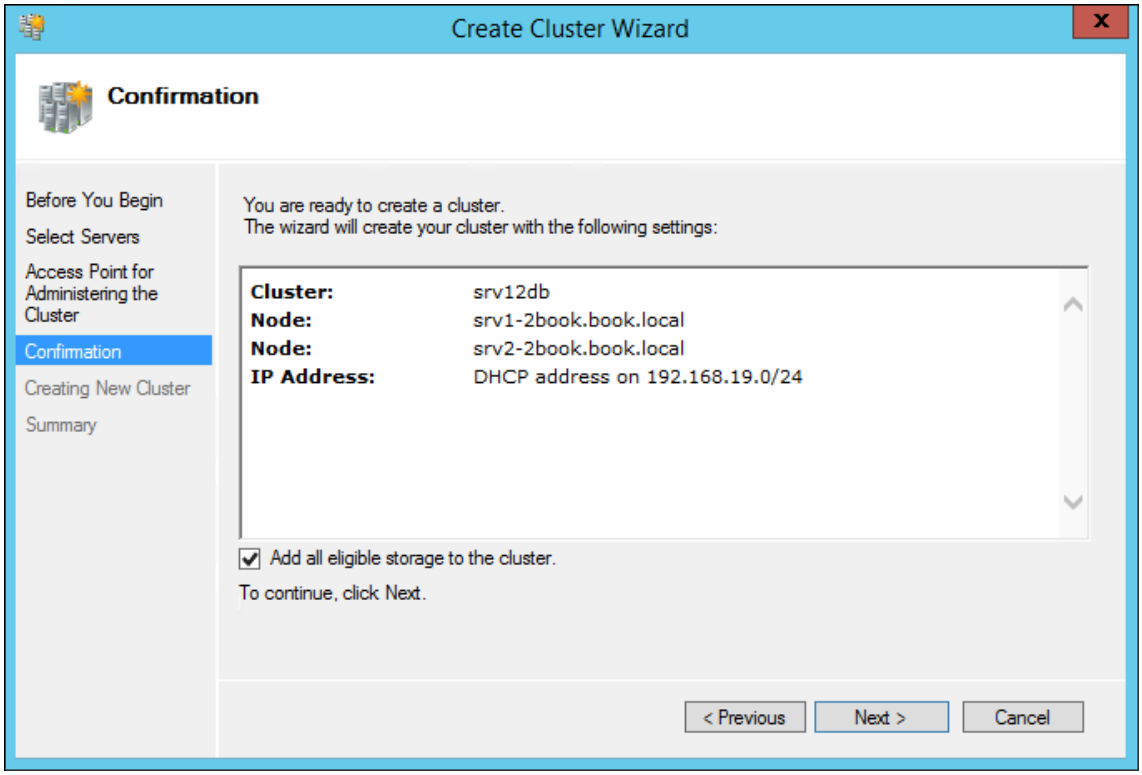

Fig. 31 Cluster creation wizard. Confirmation

Once the installation is complete, the cluster will be created and a report will be generated. The report provides recommendations on improving the cluster's reliability. To view the report, click **View Report**. If there are any errors, you should eliminate them according to the provided recommendations.

To close the cluster creation wizard, click **Finish**.

The created cluster will be added to the Failover Cluster Manager (Fig. 32).

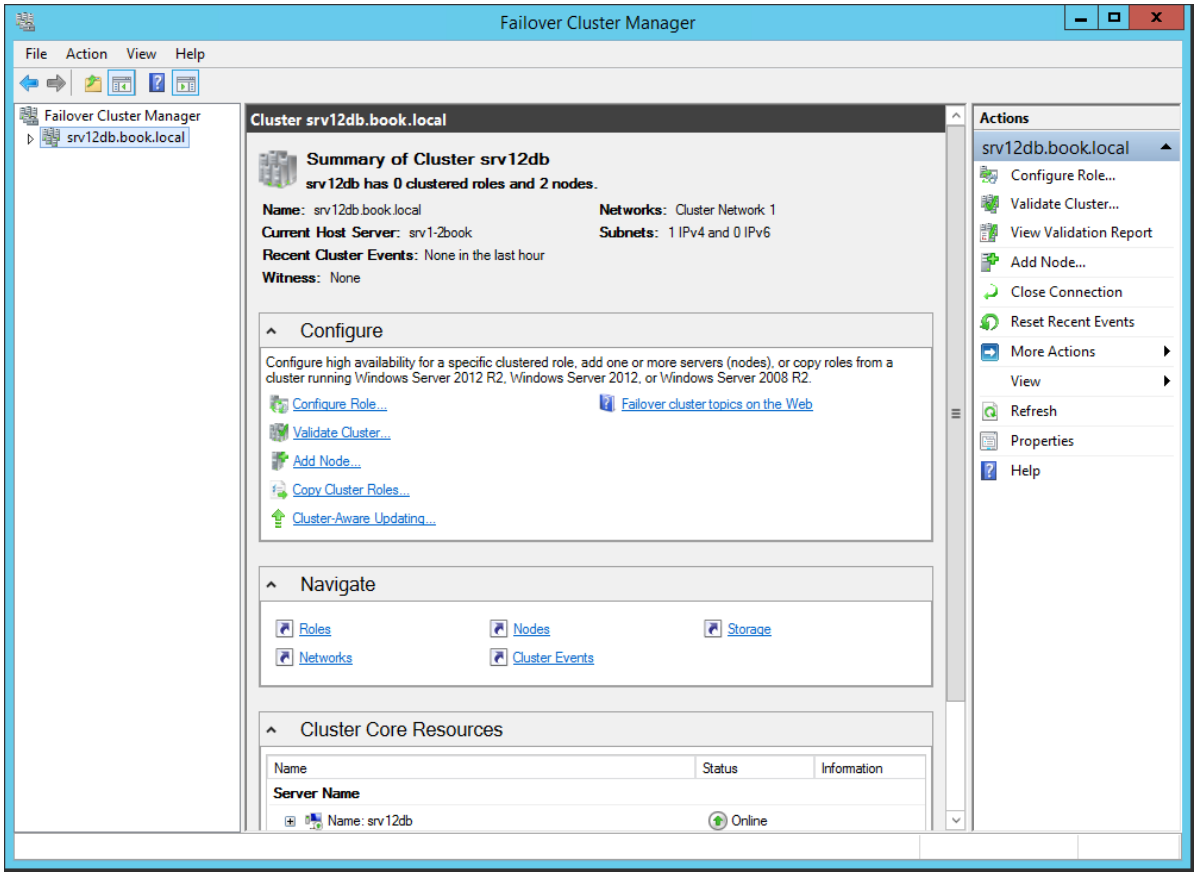

Fig. 32 Failover Cluster Manager

Step 5. Next, you can configure "Services or applications" in the failover cluster. In this case, it is the MS SQL AlwaysOn availability group. It is configured in MS SQL Management Studio.

3.2.3. Creating a shared folder for storing backup copies

For MS SQL AlwaysOn to function correctly, you need to provide all the servers in the cluster with access to backup copies of the databases. Shared access is required because initial synchronization in the MS SQL AlwaysOn availability group in the failover cluster is carried out via a backup copy, available on all the database servers.

We recommend allocating a separate hard drive on the main database server (in this example – 192.168.18.230) for backup copies, create a folder on the hard drive and grant this folder shared access to read and write (in this example, it is the folder **C:\SRV12**). Next, specify this folder as the default folder for creating backup copies.

Step 1. Start MS SQL Server Management Studio (to learn more, see paragraph 3.1.2).

After starting MS SQL Server Management Studio, open the server context menu in the tree and select **Properties** (Fig. 10).

In the opened dialog box, go to the **Database Settings** page and change the location for backup copies in the **Backup** field of the **Database default locations** unit (Fig. 33).

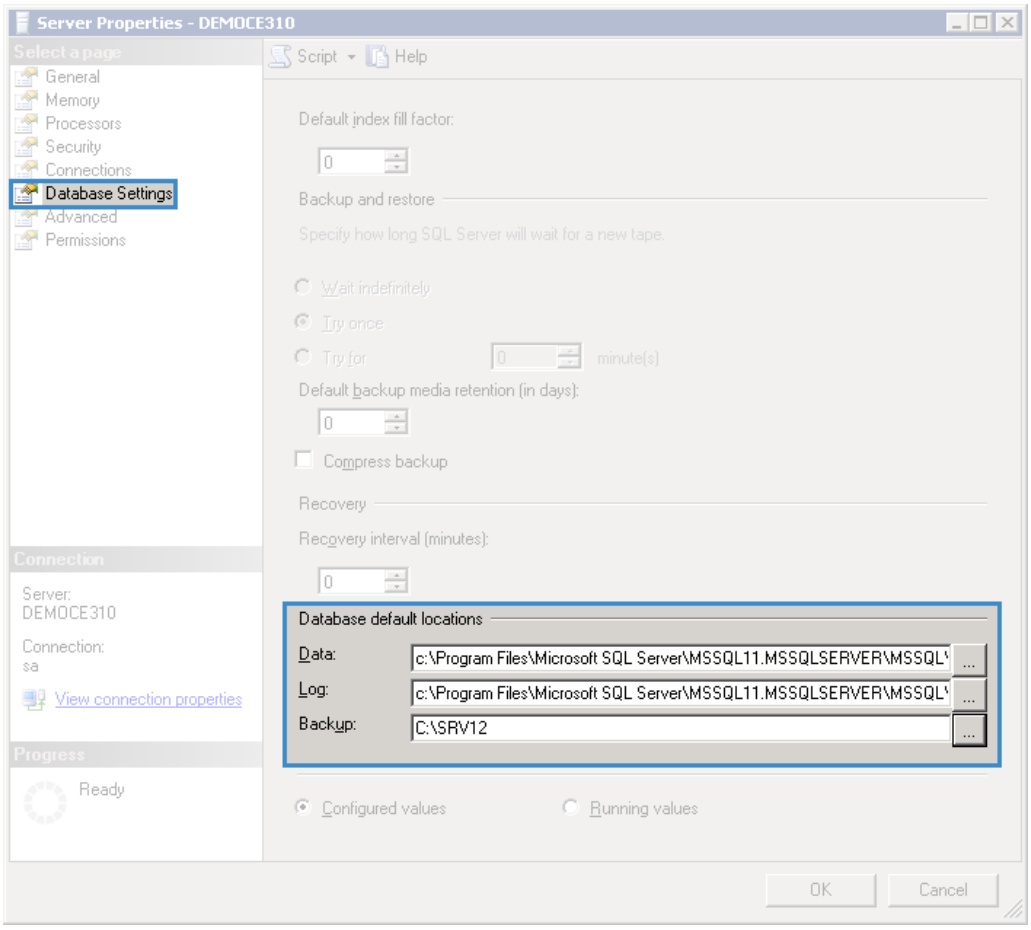

Fig. 33 Database Settings. Database default locations

To confirm the settings and close the dialog box, click **OK**.

Attention! It is not recommended to specify a shared folder located on the primary server for a secondary replica, since according to the further settings, a backup copy will be taken from the secondary replica only if the primary server and, therefore, the folder is unavailable.

3.2.4. Creating an MS SQL AlwaysOn availability group

The MS SQL AlwaysOn availability group provides availability of the database in case one or several database servers crush (depends on the total number), and prevents data losses. The availability group is based on the Windows failover cluster component.

Attention! You should apply the procedures described below only to the primary database server.

In this example, it is the server 192.168.18.230.

Place the databases that directly affect the ELMA farm to the availability group:

- **ELMA38PROD** – main ELMA database;
- **ASPState** – database for saving sessions;
- **ELMACache** – cache cluster configuration database.

An availability group can include only the databases with the full recovery model and only after creating a backup copy (even if a database is new and has no data). That is why first you need to make sure that you have selected the appropriate recovery model and manually create backup copies.

Step 1. Configure each of the abovementioned databases.

In MS SQL Server Management Studio, open the database settings by clicking **Properties** in the context menu (Fig. 34).

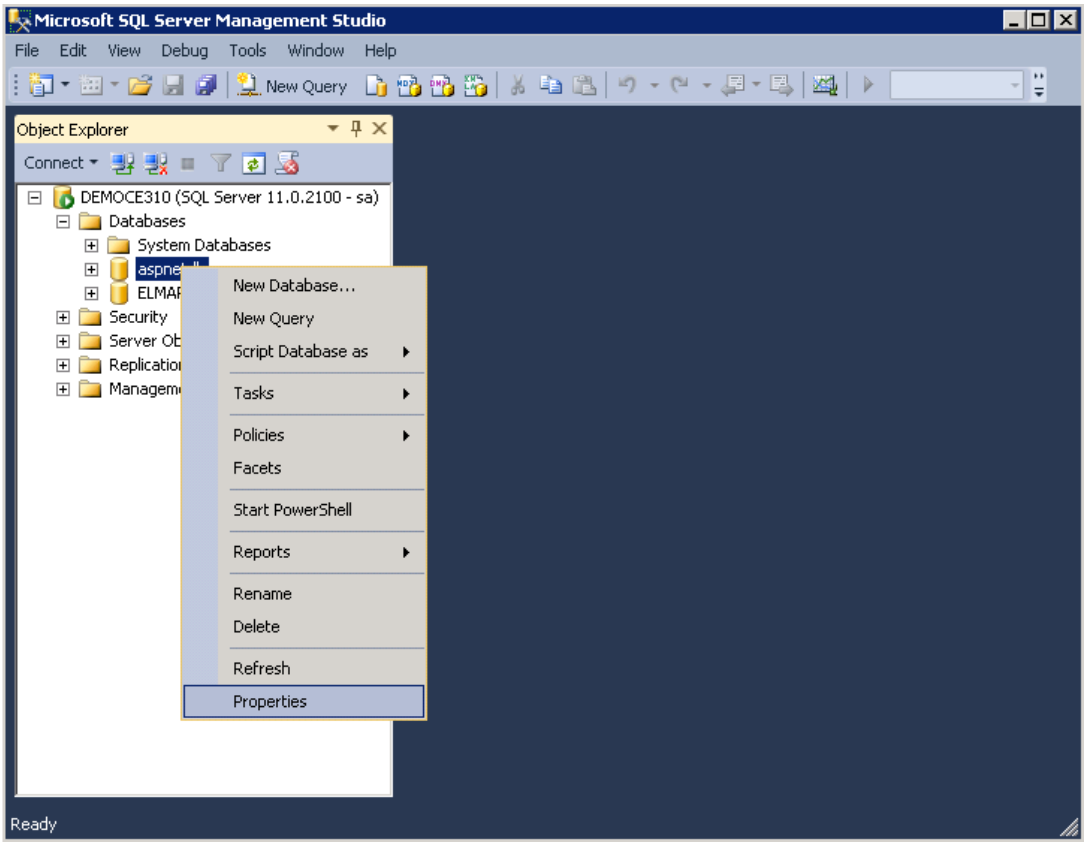

Fig. 34 MS SQL Server Management Studio. Database context menu. Properties item

In the opened dialog box, go to the **Options** page and make sure that **Full** is selected in the **Recovery model** field (Fig. 35).

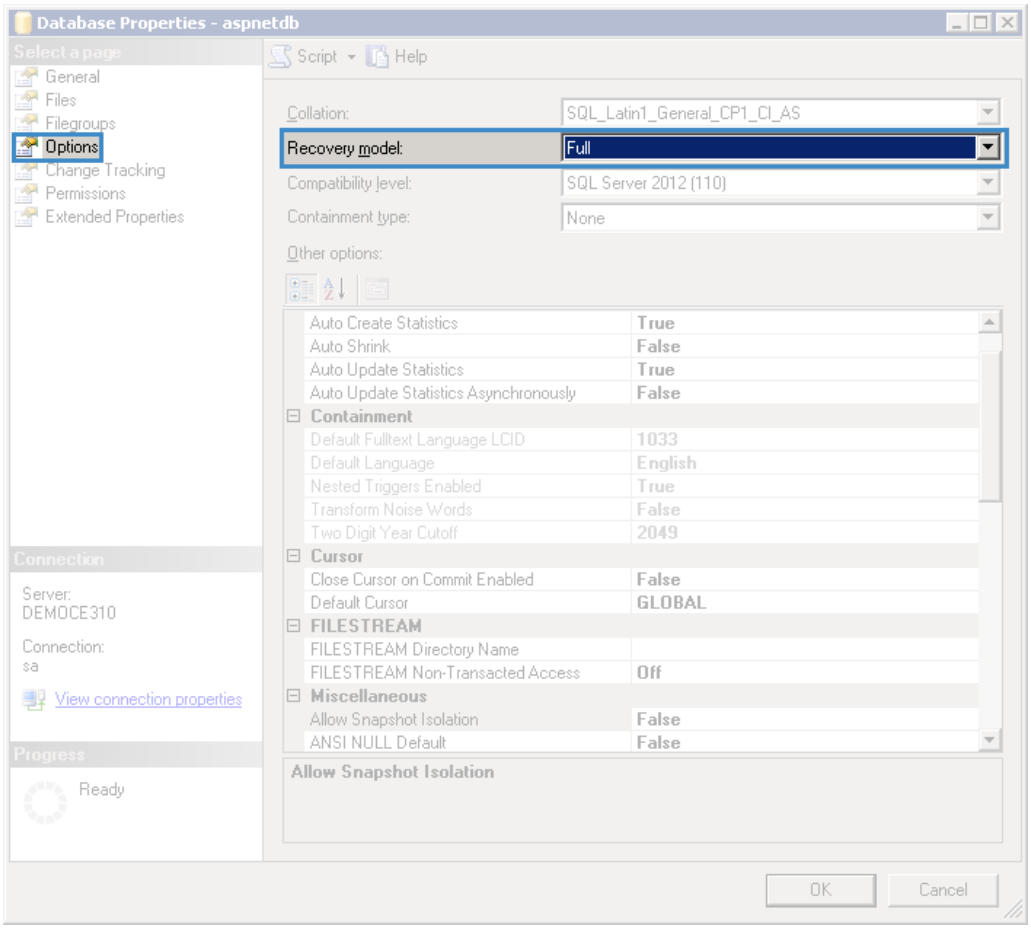

Fig. 35 Database properties. Options page. Recovery model field

Step 2. Create a backup copy of the database.

In the database context menu, select **Tasks – Back Up...** (Fig. 36).

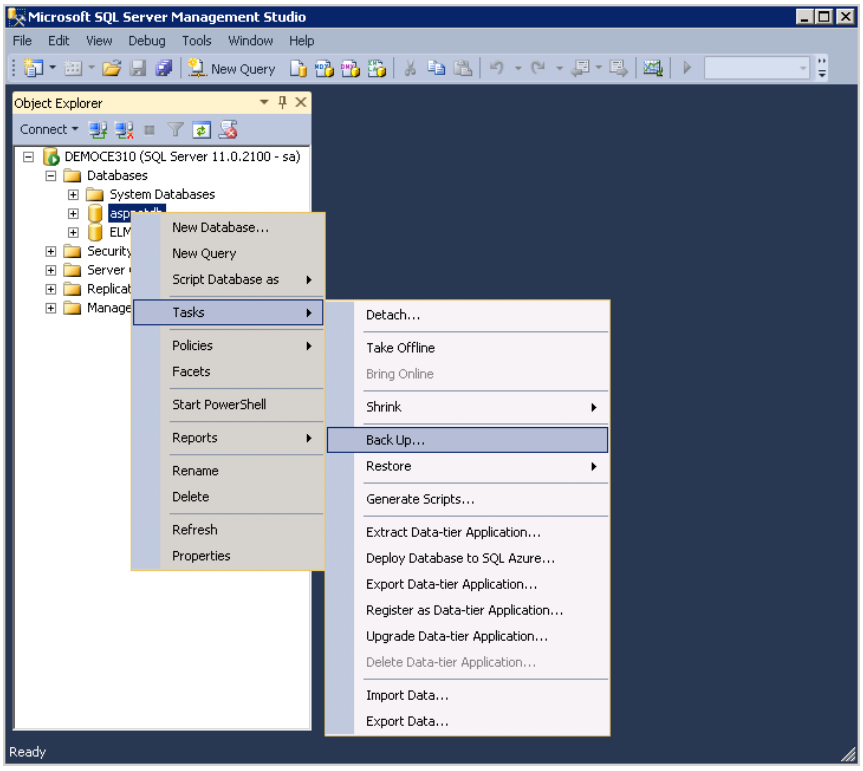

Fig. 36 MS SQL Server Management Studio. Database context menu. Tasks – Back Up…

In the opened dialog box, go to the **General** page, make sure, that **Full** is selected in the **Backup type** field and specify a location in **Back up to** field (Fig. 37). The copy must be created in the shared folder for backup copies that was created earlier.

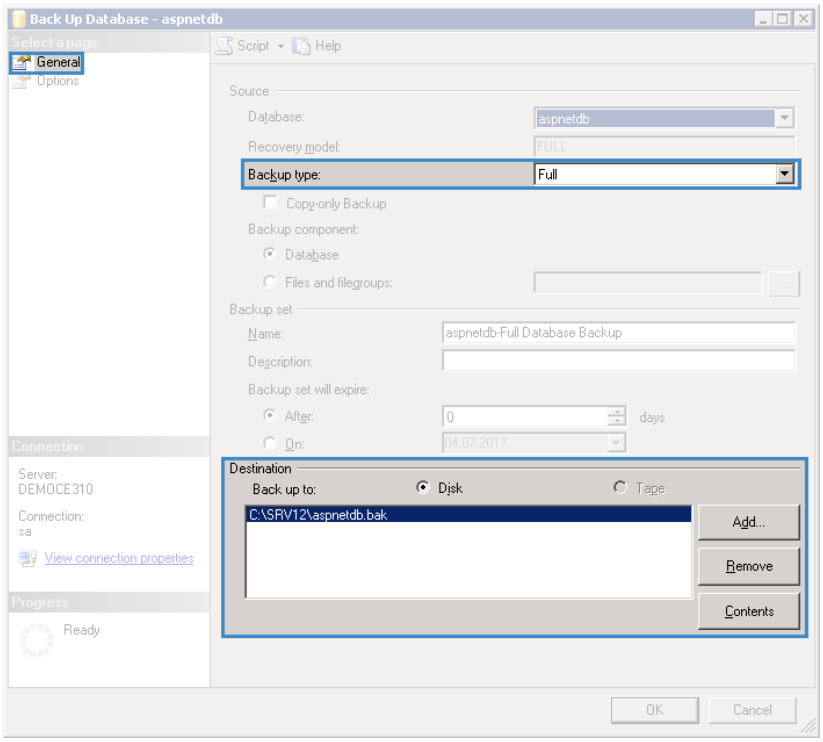

Fig. 37 Back Up Database dialog box

You should repeat these operations (steps 1 and 2) for all the databases.

Step 3. Create an availability group.

In MS SQL Server Management Studio, open the context menu of **AlwaysOn High Availability** and select **New Availability Group Wizard**... (Fig. 38).

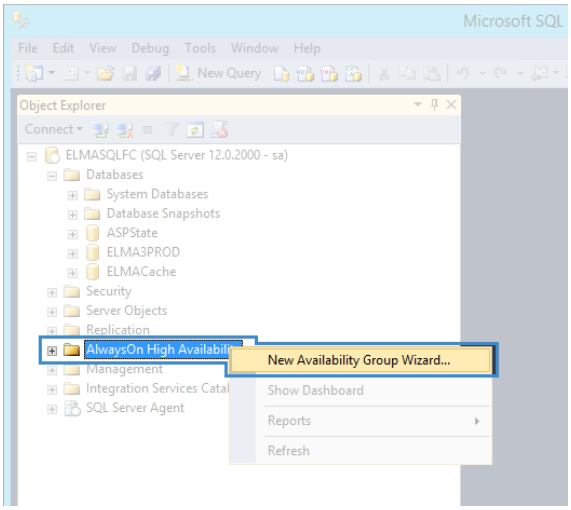

Fig. 38 MS SQL Server Management Studio. AlwaysOn High Availability context menu

Skip the first step by clicking **Next**.

At the **Specify Name** step, enter a name for the new availability group (Fig. 39) and click **Next**. The availability group name is not used for connecting to the database and is not specified in any ELMA configuration files; therefore, you can enter any name you like.

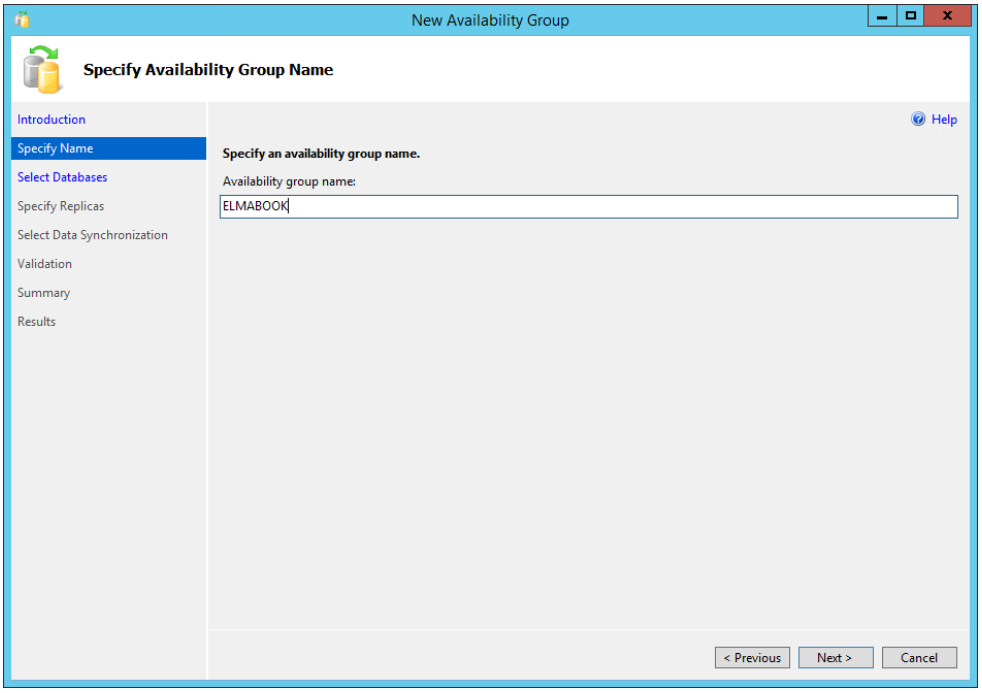

Fig. 39 New Availability Group wizard. Specify Name step

To go to the next step, click **Next**.

At the next step, **Select Databases** (Fig. 40), select the databases created earlier:

- ASPState
- ElmaCache
- ELMAPROD38

If you cannot select all the necessary databases now, you can do it later.

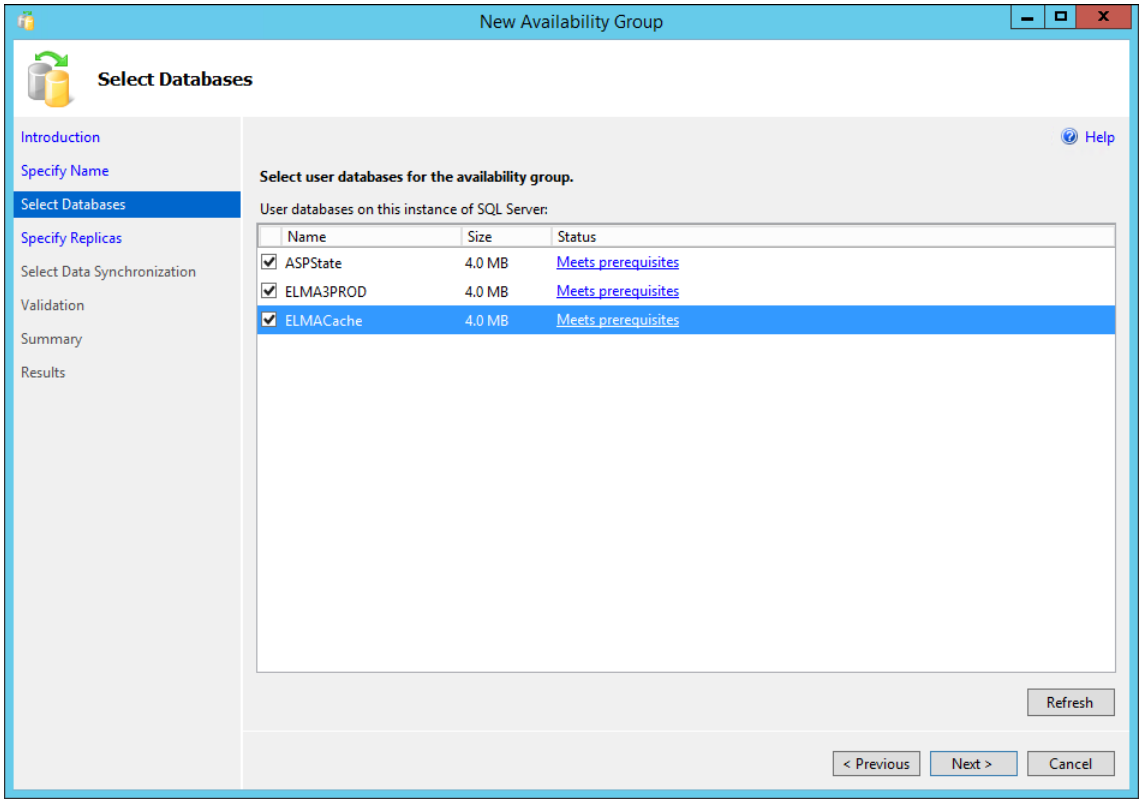

Fig. 40 New Availability Group wizard. Select Databases

You can add databases to the availability group only if the **State** column displays "Meets prerequisites". If the State column displays Full backup is required, you need to create a backup copy of this database.

To go to the next step, click **Next**.

At the **Specify Replicas** step, on the **Replicas** tab, click **Add Replica**... and specify the information for connecting to the second server in the opened window. Click **Connect**. The selected server will be added to the list of availability replicas on the **Replicas** tab (Fig. 41).

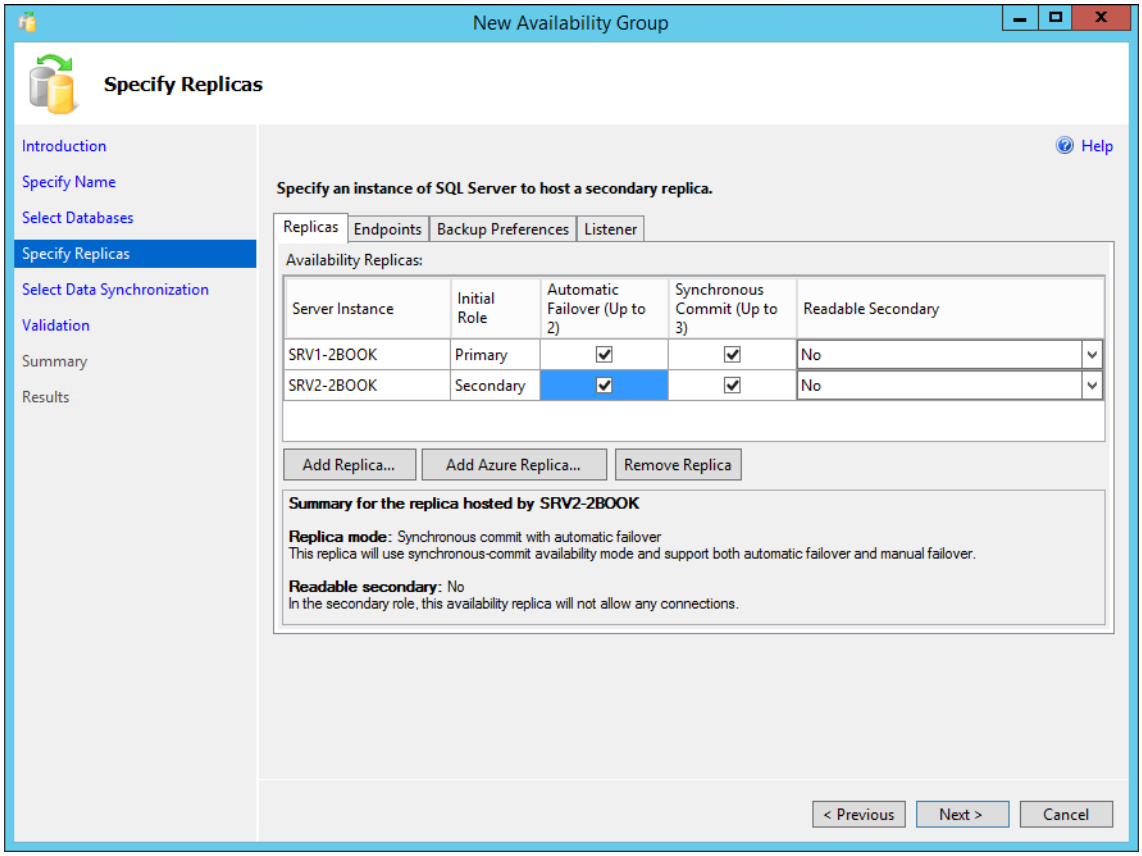

Fig. 41 New Availability Group wizard. Specify Replicas. Replicas tab

On this tab, you also should check the boxes in the columns **Automatic Failover (Up to 2)** and **Synchronous Commit (Up to 3)**.

Attention! Sometimes the synchronous commit mode slows down the system. It is caused by the network connection settings. Because of this, asynchronous mode was selected for the secondary replica (see paragraph 6.10).

At the **Specify Replicas** step, open the **Backup Preferences** tab and select "Primary" (Fig. 42).

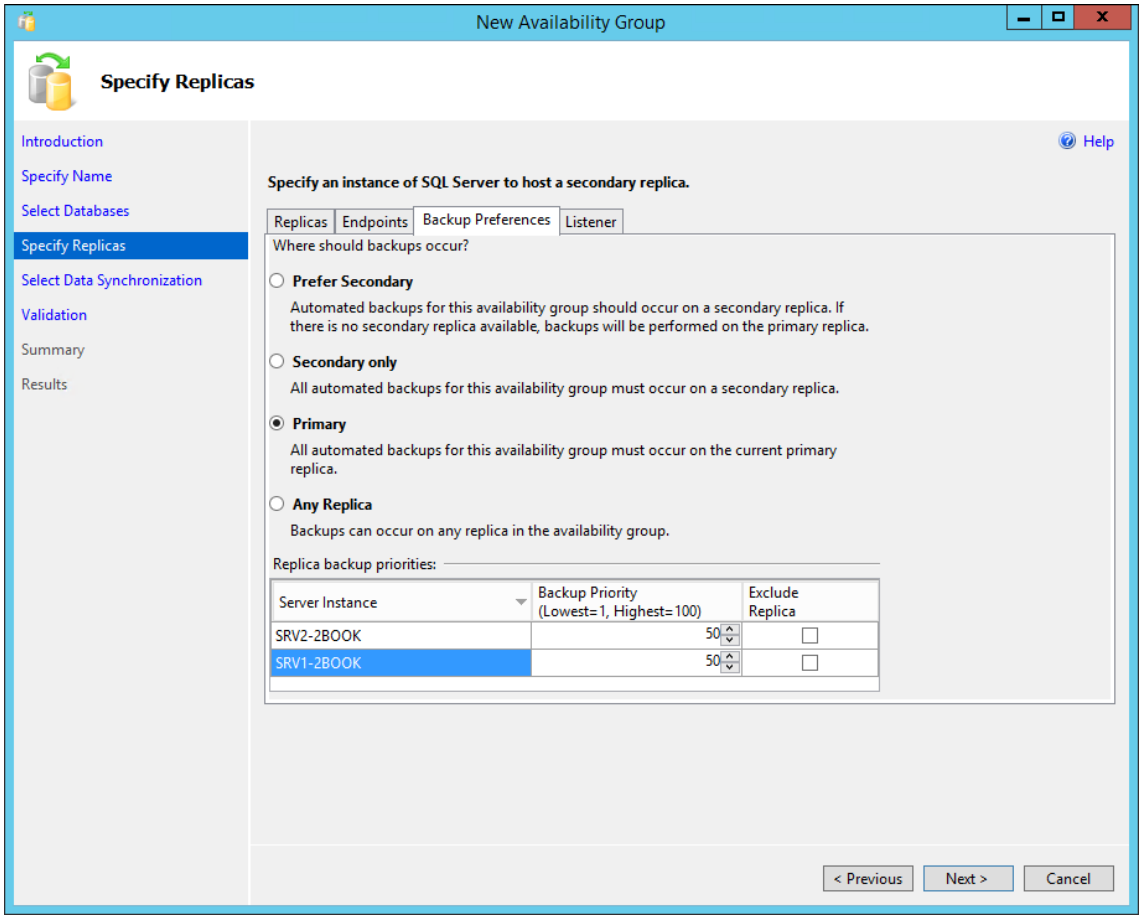

Fig. 42 New Availability Group wizard. Specify Replicas. Backup Preferences

The reason behind this choice is that depending on the reading availability settings of the secondary replica, it may be impossible to take backup copies from the secondary replica.

This option does not affect the general performance of the system since it is recommended to make backup copy and perform other maintenance outside the working hours (see the paragraph 5.1).

To go to the next step, click **Next**.

At the **Select Data Synchronization** step, select **Full** and select a network folder, accessible to all the replicas (Fig. 43). By default, this field contains the available network folder for backup copies on the 192.168.18.230 server, which was configured earlier.

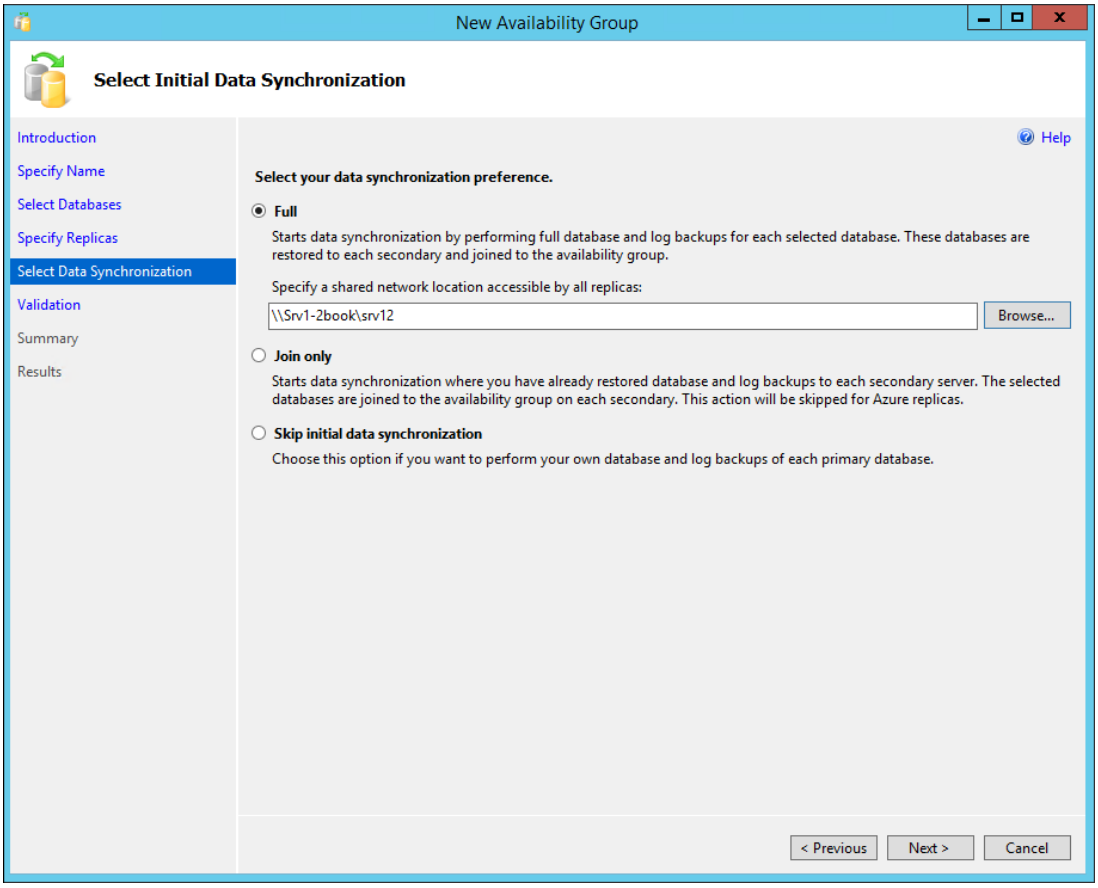

Fig. 43 New Availability Group wizard. Select Data Synchronization

To go to the next step, click **Next**.

At the **Validation** step (Fig. 44), the availability group will be validated.

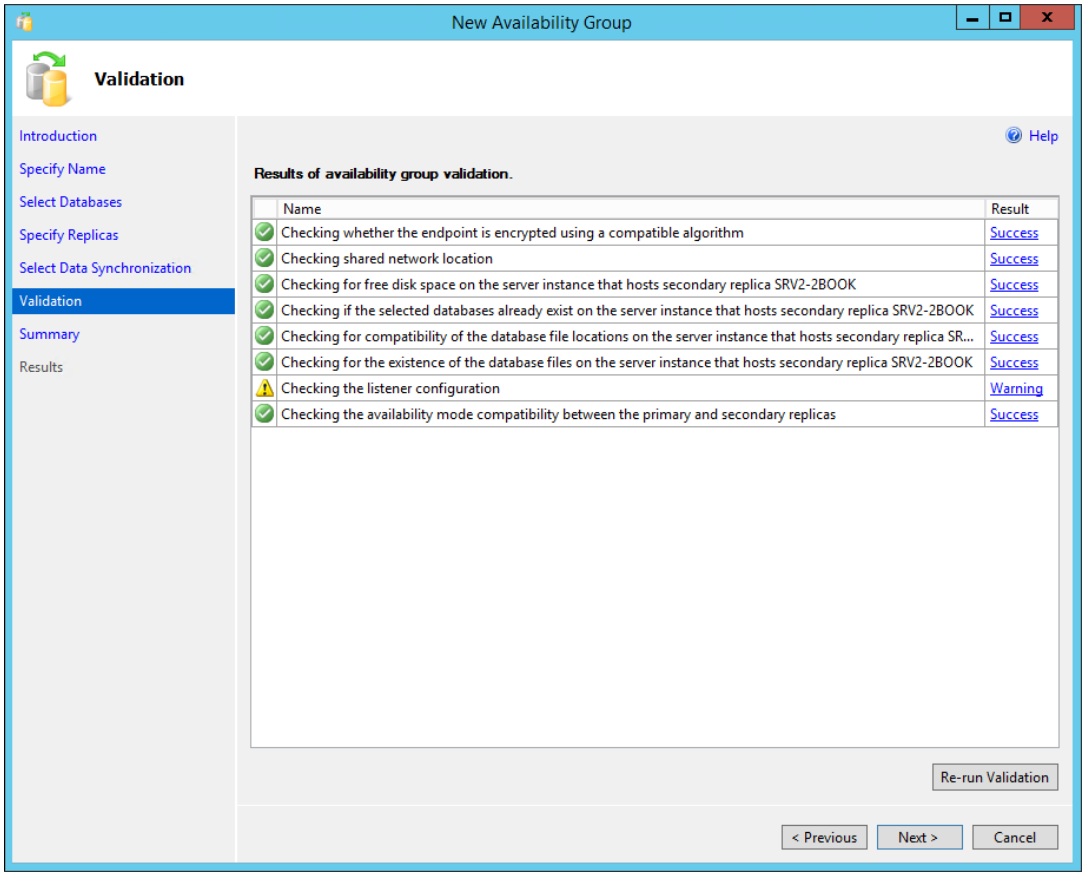

Fig. 44 New Availability Group wizard. Validation

To go to the next step, click **Next**.

At the **Summary** step, you can review the full list of parameters, selected at the previous steps of the wizard. To continue creating the availability group, click **Finish**.

The creation process may take a while; wait until it is complete. If any errors occur, follow the recommendations provided in the paragraph 6.8.

3.2.5. Adding databases to the availability group

Attention! You need to follow through this paragraph only if databases were not added to the availability group upon its creation.

If you have failed to add a database to the availability group when creating the MS SQL AlwaysOn group (see paragraph 3.2.4), you can do it at any moment later.

This paragraph repeats some of the steps mentioned above, since adding a database when creating the MS SQL AlwaysOn availability group is not significantly different from adding a database to an existing availability group.

In MS SQL Server Management Studio, open the context menu of **AlwaysOn High Availability – Availability Groups – <…> (Primary) – Availability Databases** click **Add Database**... (Fig. 45). The wizard for adding a database to the availability group will open.

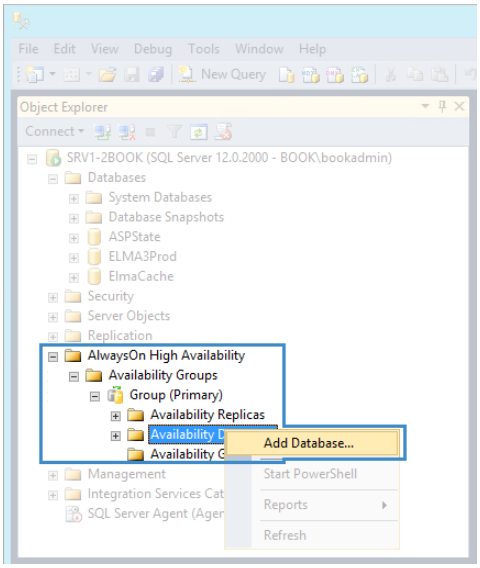

Fig. 45 MS SQL Server Management Studio. Availability Database context menu

Skip the first step of the wizard by clicking **Next**.

At the **Select Databases** step, select the database created earlier and click **Next**.

At the **Select Data Synchronization** step, select **Full** and specify the network folder, accessible for all replicas. To go to the next step, click **Next**.

At the **Connect to Replicas** step, click **Connect** (next to the required replica) and in the opened dialog box select **Windows authentication** in the **Authentication** field. Click **Connect**. The server instance will be connected to the replica. To go to the next step click **Next**.

At the next step, the availability group will be validated. To go to the next step, click **Next**.

At the **Summary** step, all the parameters, selected at the previous steps of the wizard will be displayed. To continue adding databases to the availability group, click **Finish**. The process of creating the availability group may take a while; wait until it is complete.

3.2.6. Creating a database listener in the availability group

Now, server fail-safety is already implemented via MS FailoverCluster, however, such a cluster is not resistant to MS SQL Server failures. To provide higher failover protection and reduce the time for switching between replicas in case of a service or server failure, it is recommended that you create an availability group listener.

Attention! The account you intend to use to perform these actions must have permissions to create computers in the domain.

In MS SQL Management Studio, select your availability group – **Availability Group Listeners – Add Listener...** (Fig. 46).

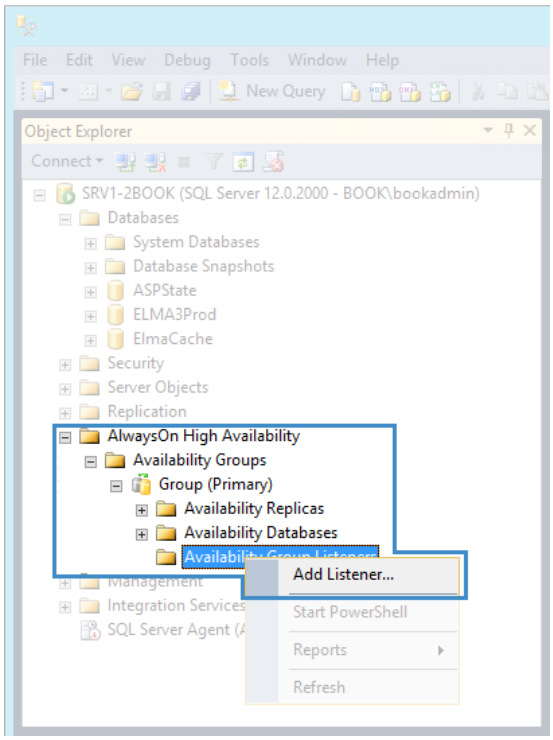

Fig. 46 MS SQL Management Studio. Availability Group Listeners context menu

In the opened dialog box (Fig. 47), fill in the following fields:

- **Listener DNS Name** – name, used for connecting to databases. In this case, it is **SRV**.
- **Port** – connection port. By default – **1433**.
- **Network Mode and Subnet** – the default value is **DHCP**; upon creation, an attempt will be made to automatically obtain an IP address and link the domain name to it.

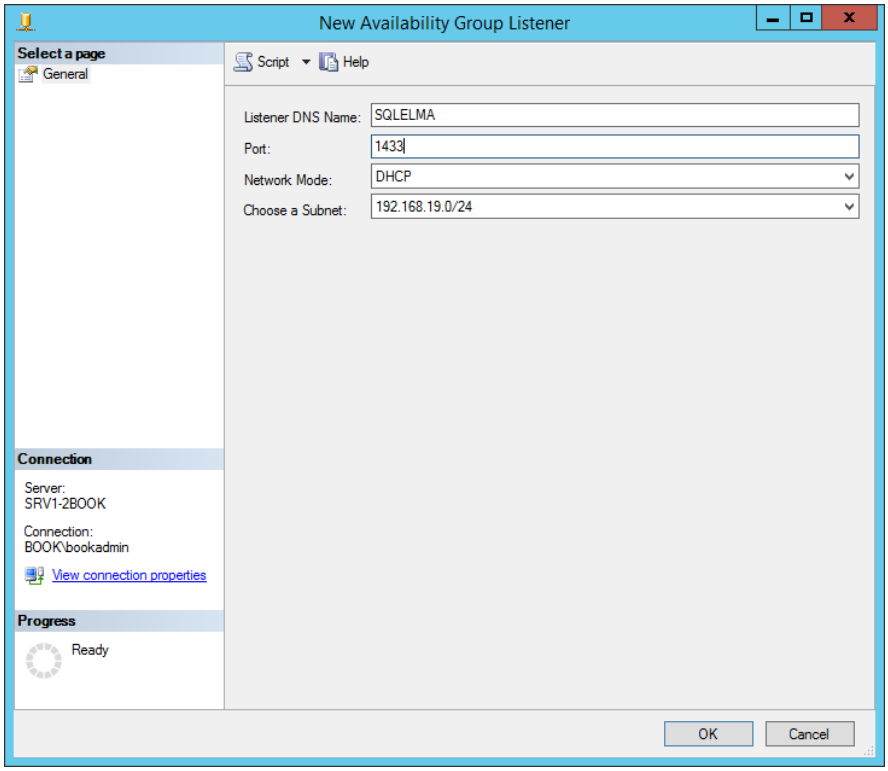

Fig. 47 New availability group listener dialog box

To finish registering the listener, click **OK**. If no errors occurred in the process, you can proceed to the next step (see below).

If an error occurred when registering the listener, try changing the **Network Mode** from **DHCP** to **Static** and manually specify an unoccupied IP address.

Chapter 4. Installing and Configuring MS Web Farm Framework

This step provides step-by-step instructions for installing and configuring an ELMA servers web farm.

A web farm allows distributing the load (activity of ELMA users) between several servers (virtual and physical), that work as a single whole. This way, you can significantly increase the number of users working in the system simultaneously.

4.1. Prerequisites

This section describes the requirements for deploying a web farm.

A web farm consists of two application servers and a controller server (load-balancer), which distributes the load between them:

- 93.158.134.3 – SRV12-1
- 213.180.204.3 – SRV12-2

Farm controller server:

- 93.158.134.30 – WFCONTROLLER

Prerequisites:

- disable Windows Firewall;
- disable monitors/real-time of anti-virus software;
- install and configure the IIS component as described in paragraph 4.2;
- install and configure the database server.

Some of the installation steps will temporary require:

- Internet access – direct or via a proxy server.
- enabled Windows automatic update at least to "Check for update, but let me choose whether to download and install them". Actual Internet access and installation of updates is not required.

Install the following applications and components, which are not included in the standard Windows OS:

- App Fabric 1.1 http://www.microsoft.com/en-us/download/details.aspx?id=27115

- Application Request Routing (ARR) 2.5 http://www.microsoft.com/web/gallery/install.aspx?appid=ARRv2_5
- WPI launcher (requires Internet connection) http://www.microsoft.com/WEB/downloads/platform.aspx
- Web Farm Framework (installed from WPI launcher)
- URL Redirect (installed from WPI launcher)

4.2. Installing and configuring IIS

Attention! You need to apply the steps described below to each application server and to the controller.

You need to install the IIS component before installing ELMA. If IIS is configured correctly, ELMA will function properly.

4.2.1. Installing IIS

1. Open the Server Manager (**Start -> Server Manager**), go to **Dashboard**, click **Manage** and select **Add Roles and Features** (Fig. 21). The Add Roles and Features wizard will open (Fig. 48).

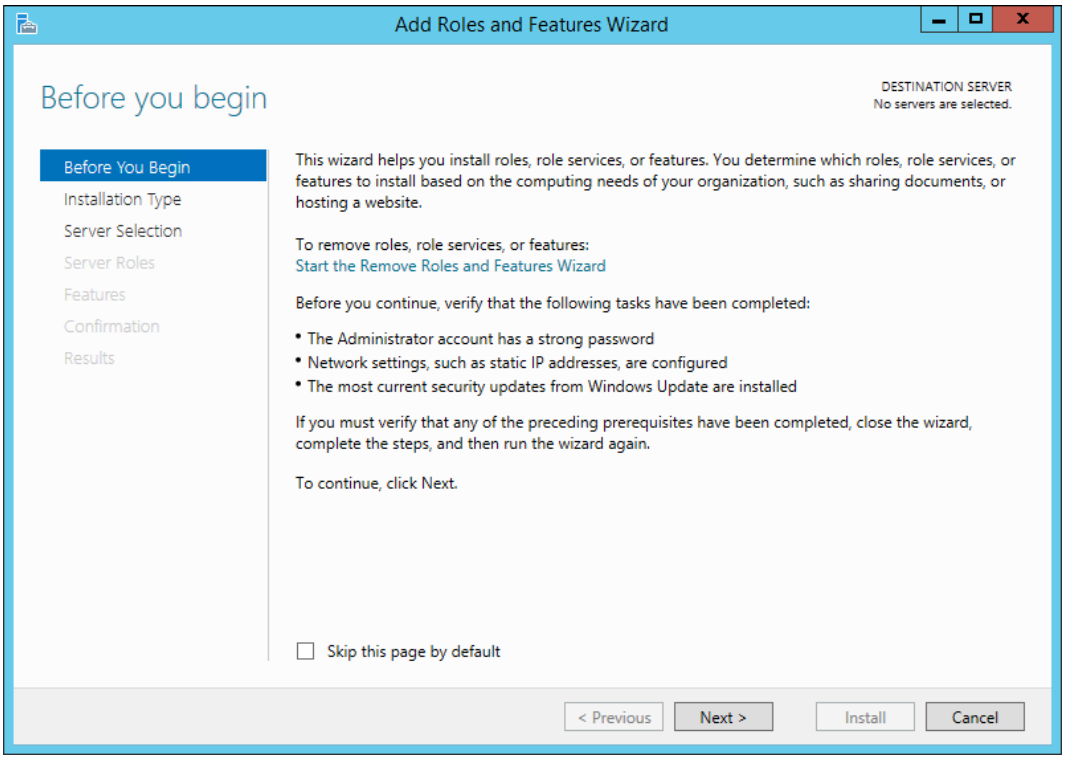

Fig. 48 Add Roles and Features wizard. Before you Begin

Skip the first step (Fig. 48) by clicking **Next**.

2. At the **Installation Type** step, select **Role-based or feature-based installation** (Fig. 49) and click **Next**.

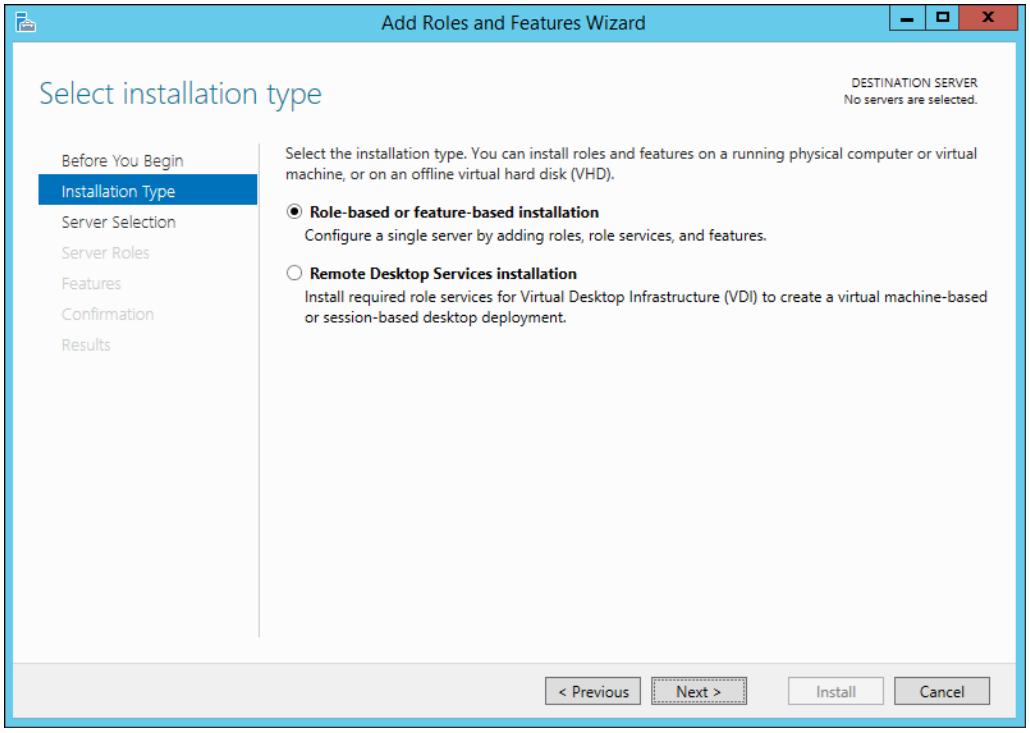

Fig. 49 Add Roles and Features wizard. Installation type

3. At the **Server Selection** step, choose **Select a server from the server pool** (Fig. 50) and select the required server. To continue installation, click **Next**.

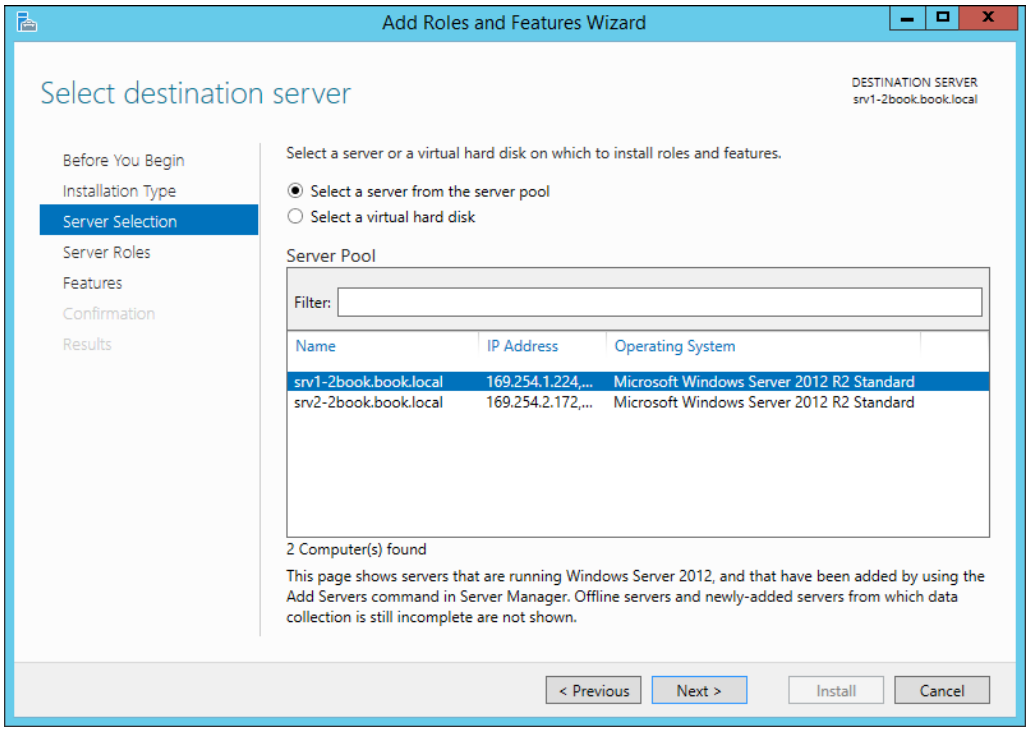

Fig. 50 Add Roles and Features wizard. Server Selection

4. At the **Server Roles** step, check **Web Server (IIS)** (Fig. 51) and click **Next**.

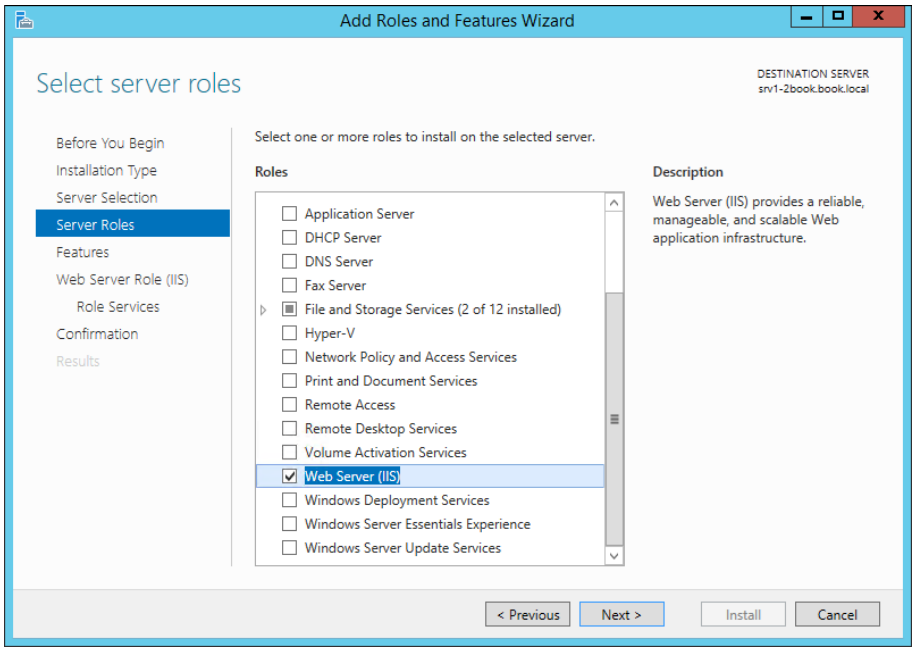

Fig. 51 Add Roles and Features wizard. Server Roles

5. Skip the **Features** and **Web Server Role (IIS)** steps by clicking **Next**.
6. At the **Role Services** step (Fig. 52), select the following role services, installed for IIS.

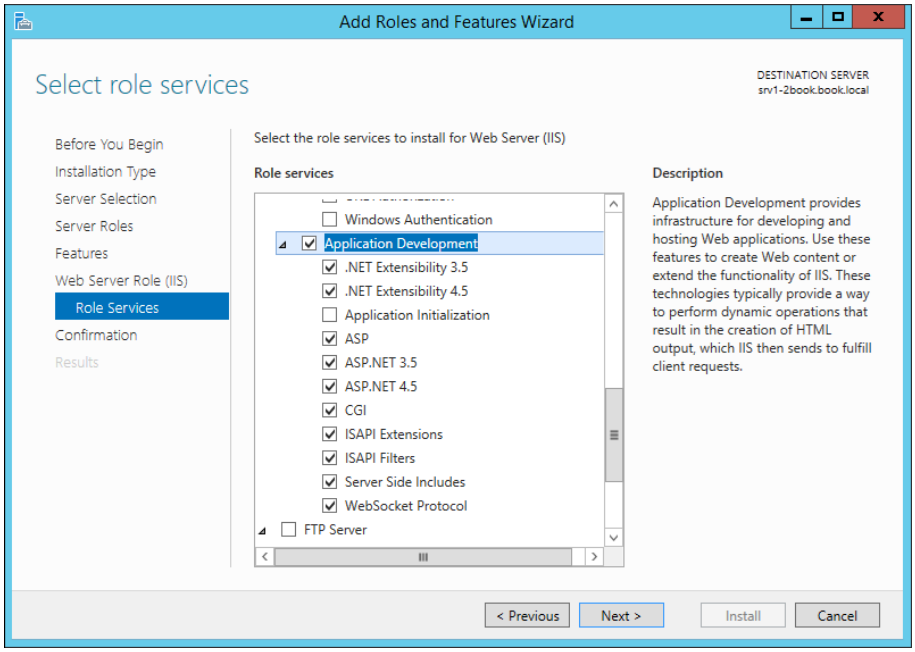

Fig. 52 Add Roles wizard. Role Services

Select the following components (Fig. 53):

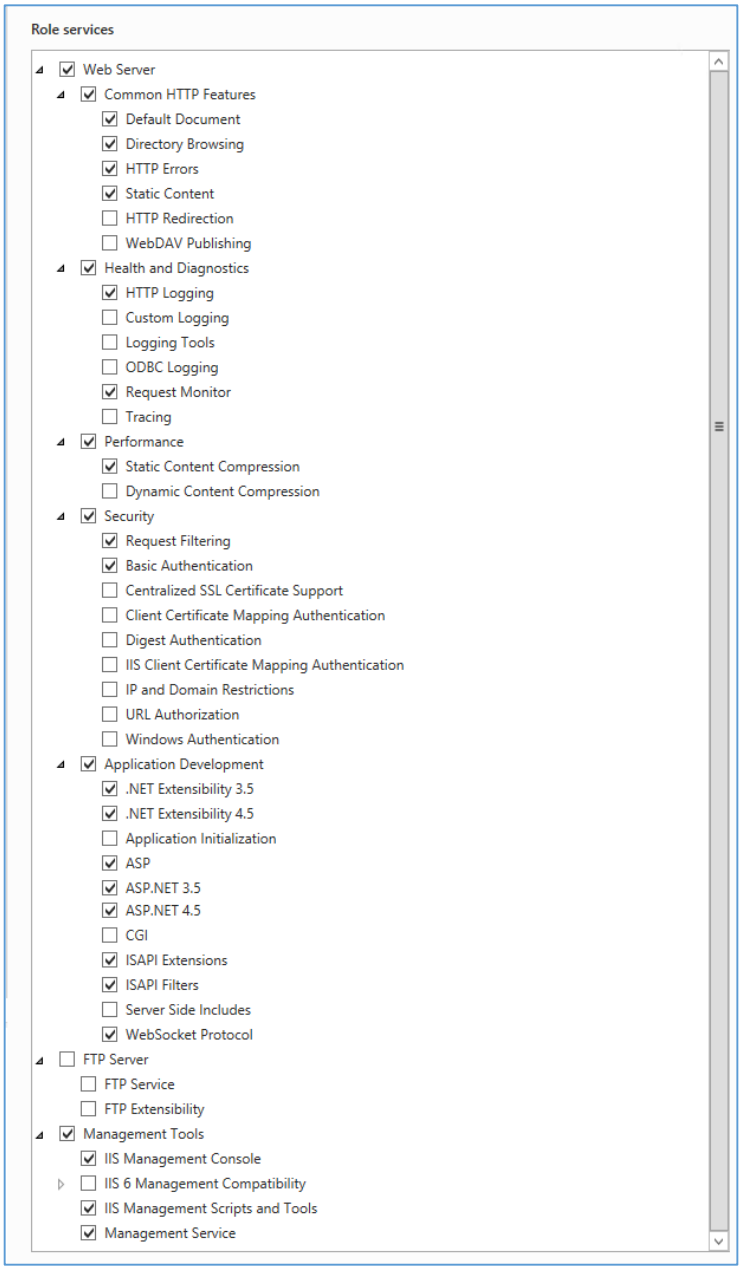

Fig. 53 Components to install

After selecting all the required components, click **Next.**

7. At the **Confirmation** step (Fig. 54), you can review role settings, role services and features, selected at the previous steps. To start the installation process, click **Install**.

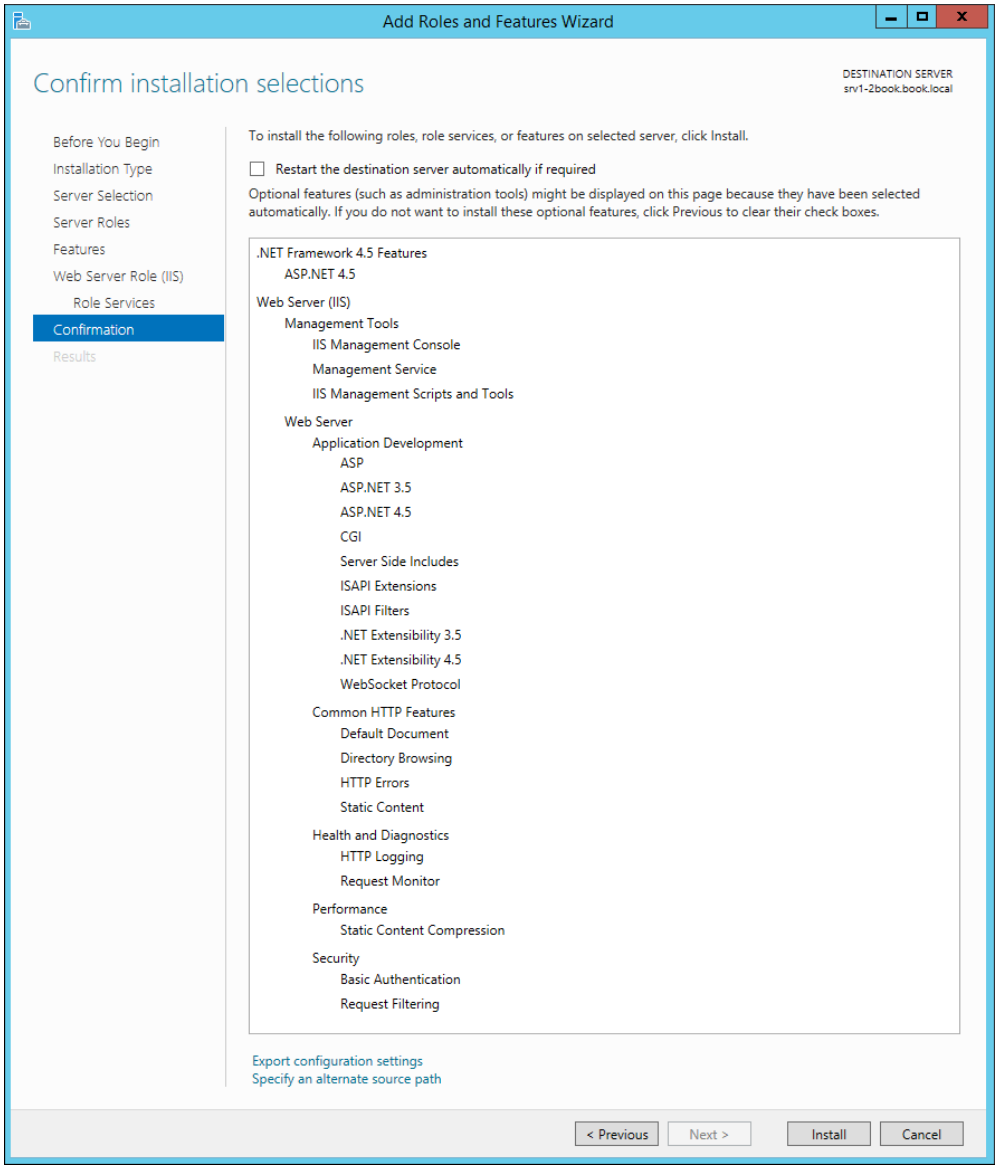

Fig. 54 Add roles wizard. Confirmation

Close the wizard after the installation is completed.

4.2.2. Configuring IIS

To configure IIS:

1. Start the Internet Information Services (IIS) Manager (**Start –> IIS Manager**) and click on the **Authentication** icon in the **IIS** block (Fig. 55).

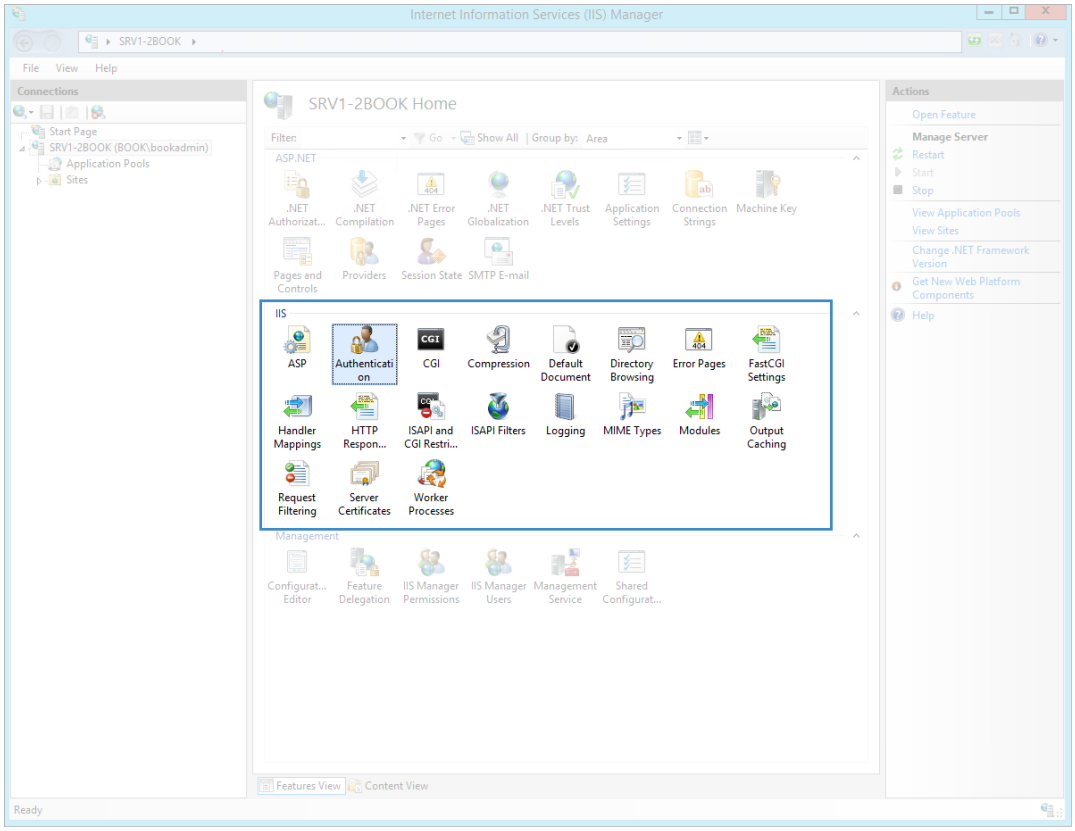

Fig. 55 IIS Manager. Authentication icon

2. In the opened window, select **Anonymous Authentication** and click **Edit...** in the actions menu on the right (Fig. 56).

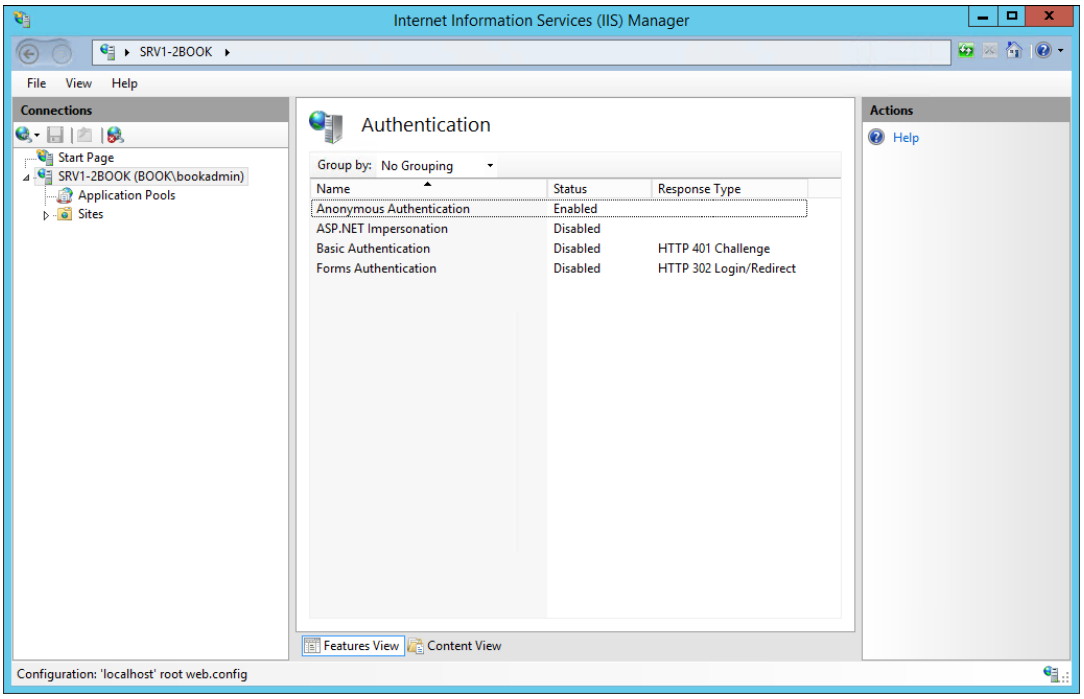

Fig. 56 IIS Manager. Authentication section

In the opened dialog box, select **Application Pool Identity** (Fig. 57).

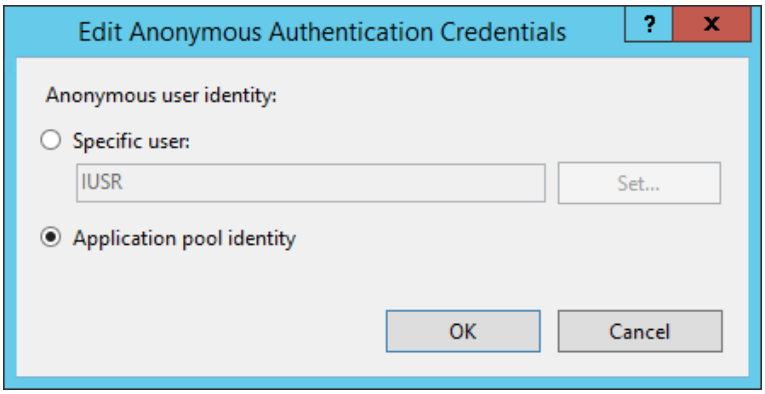

Fig. 57 IIS Manager. Edit Anonymous Authentication Credentials

If you skip this step, ELMA may open pages without static content (images, styles).

Attention! Steps 3 – 6 (marked with [*]) are optional, but if you intend to employ ELMA web farm, it is recommended that you perform them immediately.

3. **[*]** Additionally, to establish remote management from one server machine, go to the **Management Service** in the IIS Manager (Fig. 58).

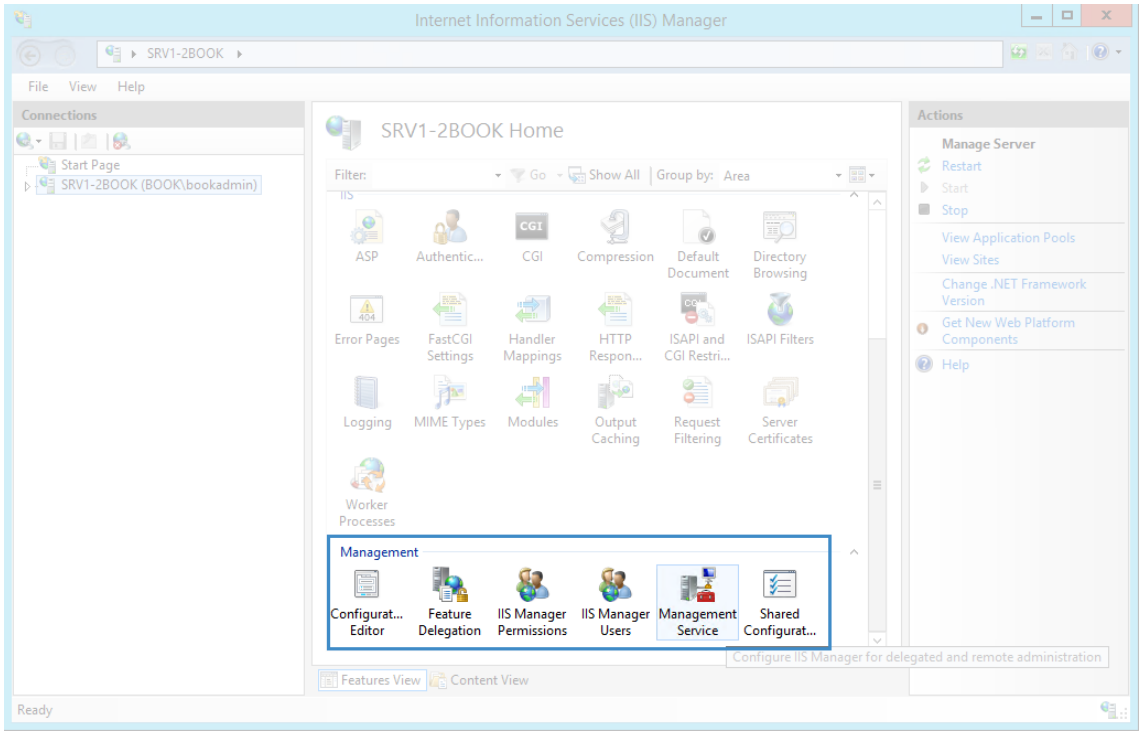

Fig. 58 IIS Manager. Management Service icon

If you do not have this menu item, then the **Management Service** component is not installed. Install this component and restart the IIS Manager.

4. **[*]** Ignore the warnings in the upper right corner (Fig. 59); they specify the correct order of actions and the mentioned stopping and restarting refers to this service, not the ELMA system/server/pool.

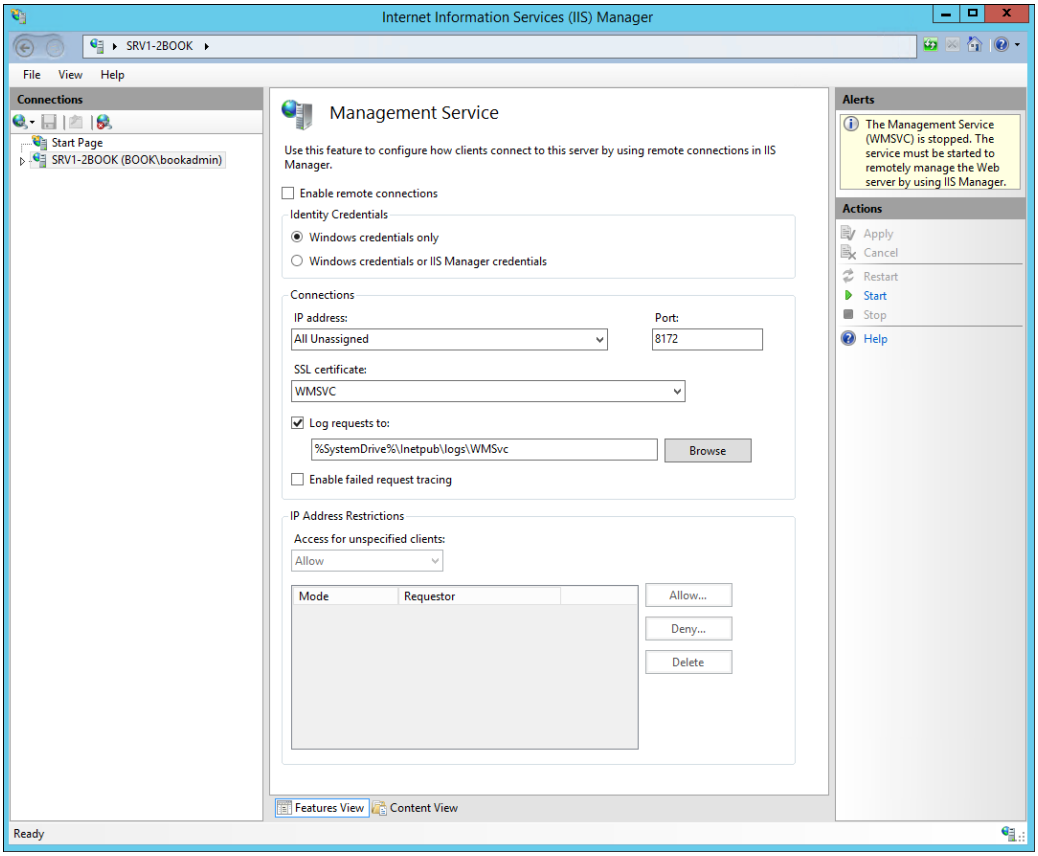

Fig. 59 IIS Manager. Management Service section

5. **[*]** Stop the service using the **Stop** button in the right part of the window.
6. **[*]** Check the **Enable Remote Connectors** box and start the server by clicking **Start** in the right part of the window (Fig. 60). The other settings are required for fine-tuning; you can skip them.

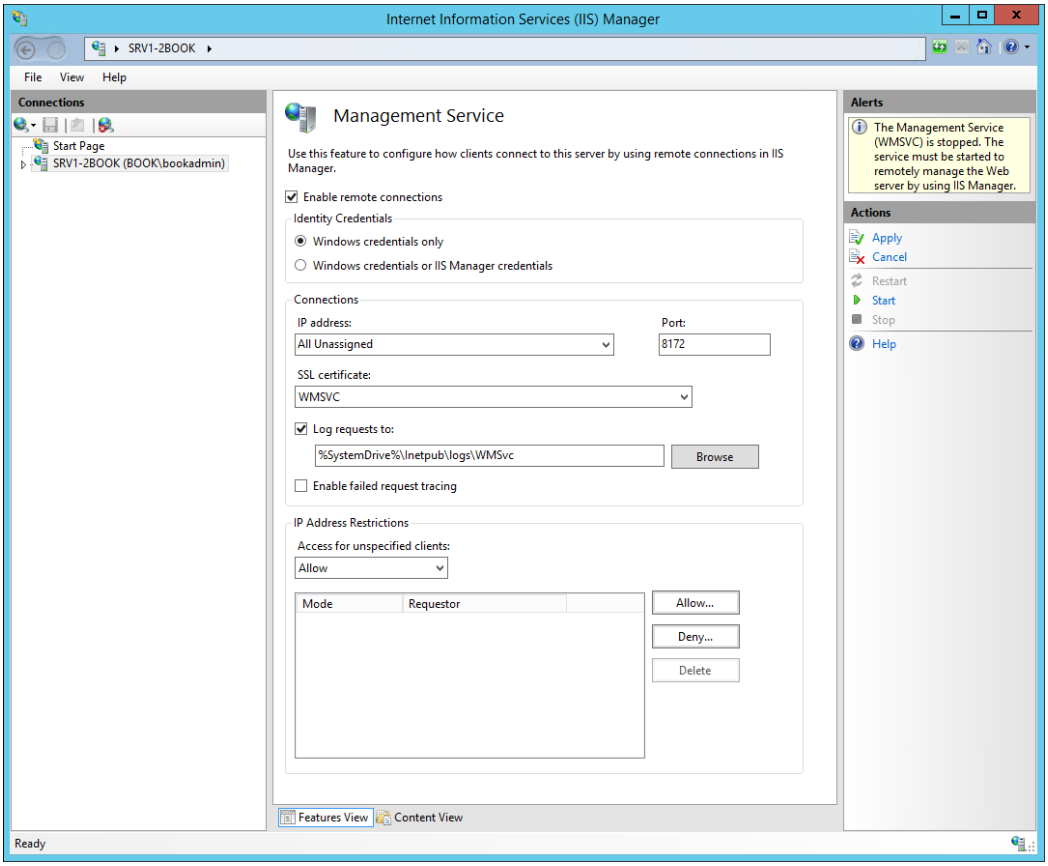

Fig. 60 IIS Manager. Management Service section

Now, everything is ready for installing ELMA.

In addition, it is recommended that you disable logging in IIS (see section 4.6).

Configuring IIS to work with ELMA is similar in other editions of Windows Server.

4.3. Installing and Configuring Farm Servers.

Attention! You should apply the steps described below to all the database servers.

4.3.1. Installing ELMA

Before you start using ELMA, you must install and register it. To install the system, use the installation software that you can get from an ELMA Company representative. You can find more information about installing and registering ELMA in ELMA BPM Platform user manual, and in Help.

This subsection partially describes the ELMA installation process. The description includes only the steps you need to pay attention to when installing ELMA Enterprise.

Step 1. ELMA Enterprise must be installed with a new configuration, using the MS SQL Server database (Fig. 61).

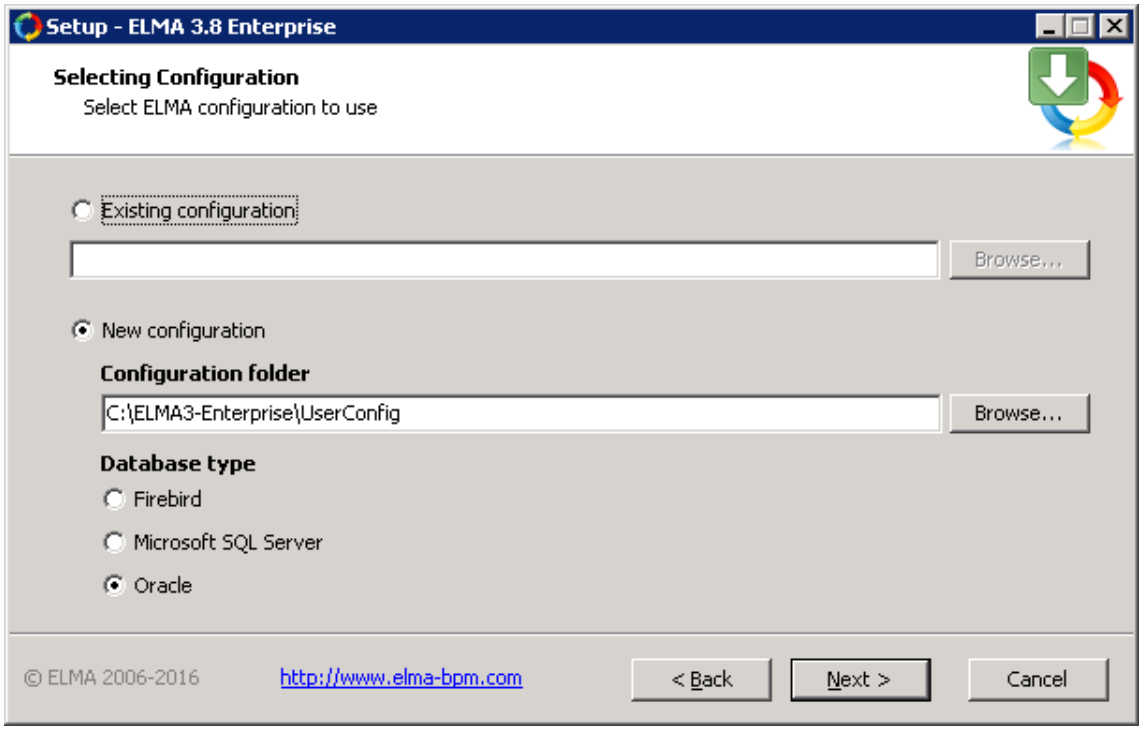

Fig. 61 Installing ELMA

Step 2. ELMA must be installed on the IIS Server. You need to set the following parameters for the system to work correctly (Fig. 62):

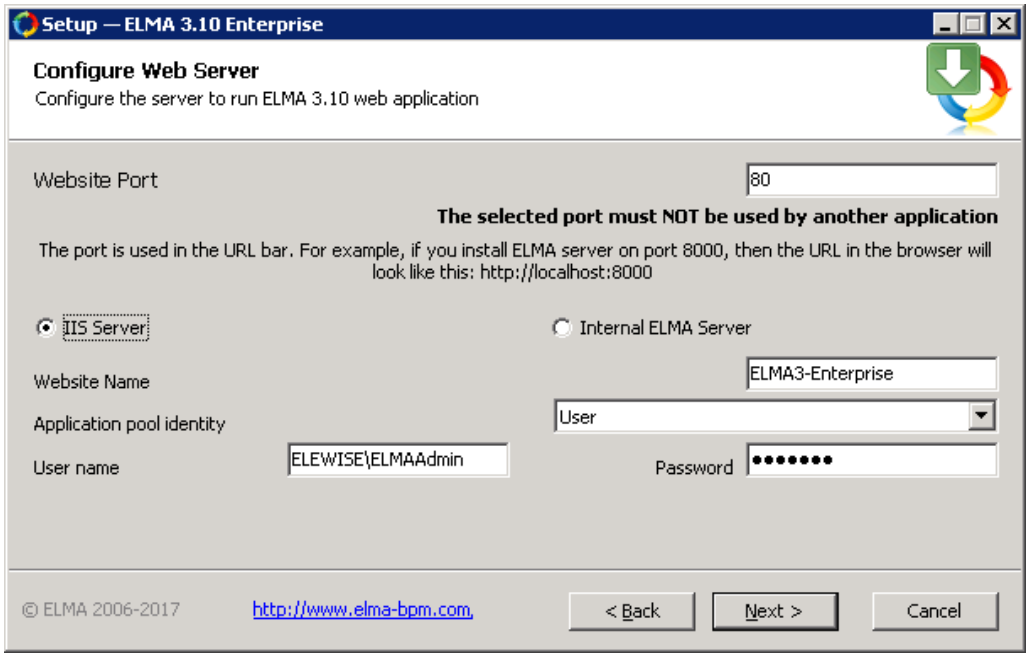

Fig. 62 Installing ELMA

Website Port must be set to 80. If the 80 port is occupied, go to the IIS Manager (**Start –>IIS Manager**) and in the **Sites** branch stop the Default website (by default, it occupies the 80 port). ELMA can use any unoccupied port, however, it is more convenient for users to open ELMA via the 80 port, since it is the default website port and users do not need to type it in the URL bar manually.

Website Name – ELMA3-Enterprise.

Web applications pool – ELMA3-Enterprise.

Application pool identity – from the drop-down list select **User** and enter their login – EleWise\ELMAAdmin – and password.

This account must have administrator permissions on the current server.

Step 3. After installing ELMA on the webserver, it is better not to start servers. On each virtual machine, stop the ELMA3-Enterprise pool and site, since you need to configure server pools and configuration files.

To do so, start the IIS manager, open **Application Pools**, right click on **ELMA3-Enterprise** (matches the name, specified during the installation) and select **Advanced Settings**... in the context menu (Fig. 63).

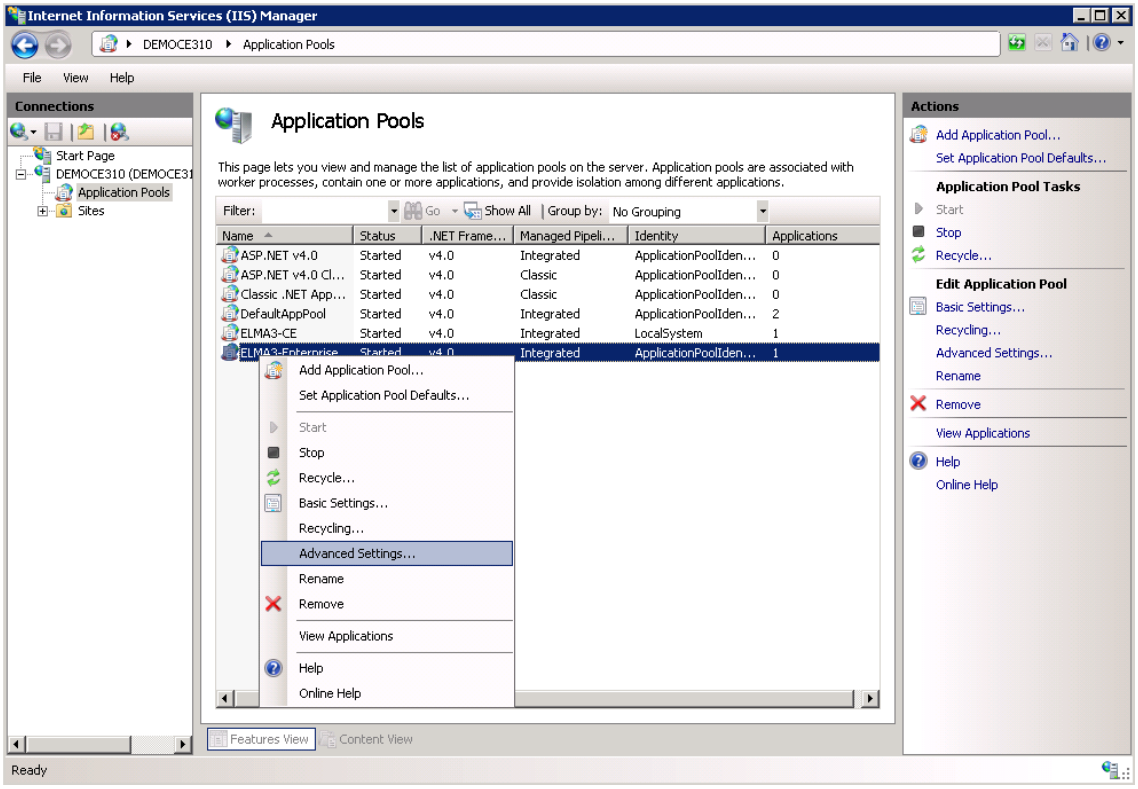

Fig. 63 IIS Manager. Pool context menu

The **Advanced Settings** dialog box will open (Fig. 64).

In **Idle Time-out (minutes)**, specify “0”. The default value is “20” (minutes). It means that the application pool will sleep if it does not receive web requests for 20 minutes. The “0” value disables this feature.

In **Regular Time Interval (minutes)**, specify “0”.

In **Disable Recycling for Configuration Changes**, select “True”. It is necessary to avoid delays and errors that occur in case of a time-out or unexpected restart, since all the HTTP requests and responses go through the application requests routing system.

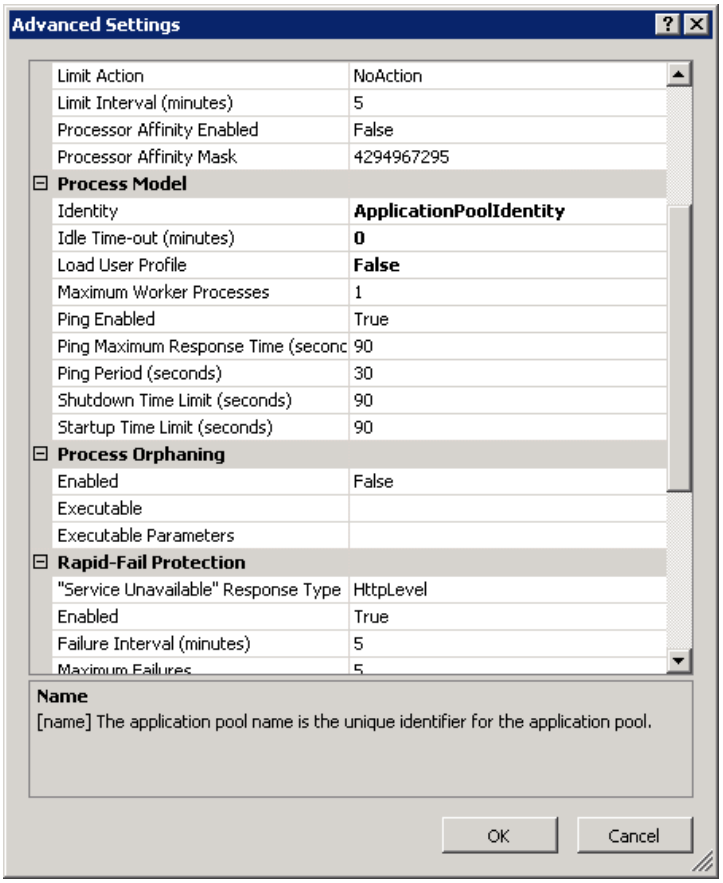

Fig. 64 IIS Manager. Advanced Settings dialog box

To apply the changes you have made, click **OK**.

Step 4. Next, you need to create a shared network folder for storing ELMA configuration. In this example, the folder will be created on the server 93.158.134.35 – ELMAConfig.

On ELMAConfig, allocate a separate hard drive with at least 500 GB for storing files. In this example, it will be disk E. Create a folder and name it **ELMAShared**.

Configure sharing for this folder. Open the context menu of the **ELMAShared** folder and select **Properties**. A dialog box will open, where you need to select the **Share** tab and click **Share**... (Fig. 65).

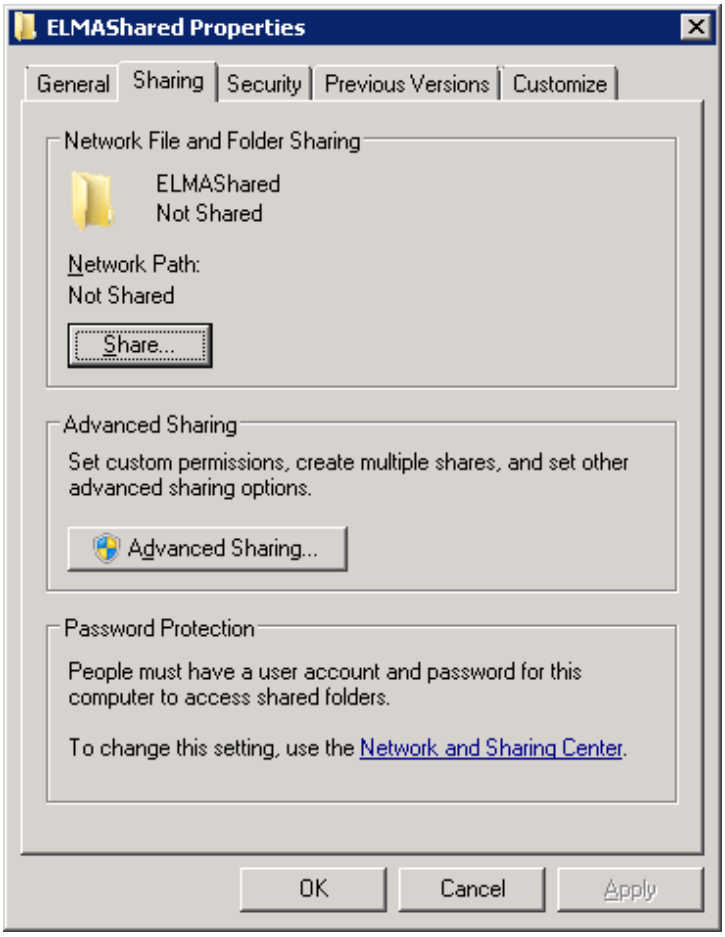

Fig. 65 Folder properties dialog box

In the opened dialog box (Fig. 66), add the account, which will run the services (ELMAAdmin), and assign permissions to **Read/Write**.

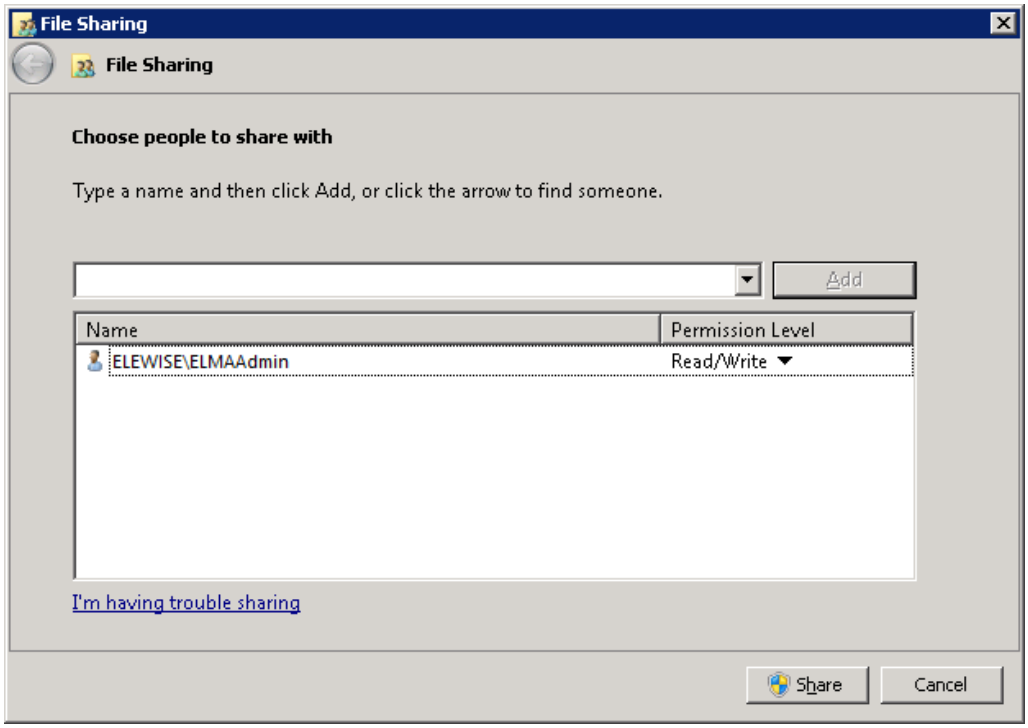

Fig. 66 Configuring shared access

To save the changes, click **Share –> Done**.

Step 5. Proceed to creating configuration files. To do so, open the ELMA configuration file (**configuration.config**) on any of the application servers. In this example, it is 93.158.134.3. You can find the configuration.config file in **C:\ELMA3-Enterprise\UserConfig**.

To the **configuration.config** file, add the following:

```
<section name="cacheService" type="EleWise.ELMA.Configuration.GenericProviderFeatureSection`1[[EleWise.ELMA.Cache.CacheServiceManager, EleWise.ELMA.SDK]], EleWise.ELMA.SDK" />
<section name="AppFabricProviderSettings" type="EleWise.ELMA.Cache.AppFabric.AppFabricProviderSettings, EleWise.ELMA.Cache.AppFabric"/></configSections>
```

Also, add the database location:

```
<connectionStrings>
 <add name="MainDB" connectionString="Data Source=SRV12DBLst;Initial Catalog= ELMAProd38;User ID=sa;Password=123456;"/>
</connectionStrings>
```

Note that **<connectionString>** contains the string for connecting to the database on the database server. In this case, it is:

- server **srv12** (database cluster connection point name);
- database **ELMAProd38**;
- log in using the **sa** account with the password **123456**.

In the string:

```
<fileStore defaultProvider="FSProvider">
    <providers>
            <clear/>
            <add                                        name="FSProvider"
type="EleWise.ELMA.Runtime.Providers.Impl.FileSystemFileStoreProvider,
EleWise.ELMA.SDK"    filesPath="\\    ELMAConfig\ELMAShared\Config\Files"
tempFilesPath="\\ ELMAConfig\ELMAShared\Config\TempFiles"/>
    </providers>
</fileStore>
```

Specify the location of the network folder for storing files and temporary files.

In this example, it is the shared folder **\\ELMAConfig\ELMAShared**

Next, copy and paste this code to the file:

```
<cacheService defaultProvider="AppFabric">
            <providers>
               <clear/>
               <add                                      name="AppFabric"
type="EleWise.ELMA.Cache.AppFabric.AppFabricCacheService,
EleWise.ELMA.Cache.AppFabric"/>
           </providers>
          </cacheService>
     <AppFabricProviderSettings                channelOpenTimeout="20000"
requestTimeout="10000"        dataCacheServiceAccountType="DomainAccount"
MainCacheName="ELMACache"  NHCacheName="ELMACache_NH"  LockTimeout="30000"
LocksRegionName="ELMACache_NH_Locks">
            <hosts>
              <host name="127.0.0.1" cachePort="22233"/>
            </hosts>
    </AppFabricProviderSettings>
```

Note that such values as ELMACache; ELMACache _NH; ELMACache _NH_Locks; 22233 are taken from the **AppFabricCacheService** settings, which are described below.

Attention! If you specified different names or port when configuring the cache, you should use those in the file.

Do not change **host name="127.0.0.1"** in any case. Since one configuration file is used for several servers, the host address must refer to the local cache instance, i.e. 127.0.0.1 or localhost. Changing this address in the configuration file to a specific network address or a domain name will cause each server to store its cache in that address, instead of storing it locally.

Step 6. Save and move all the content of the folder **C:\ELMA3-Enterprise\UserConfig** to the folder **\\ELMAConfig\ELMAShared\Config**. Thus, you will create a copy of all the configurations in the folder, accessible to all the application servers.

The resulting configuration file should look like this:

```
<?xml version="1.0"?>
<configuration>
  <configSections>

        <section name="cacheService"
type="EleWise.ELMA.Configuration.GenericProviderFeatureSection`1[[EleWis
e.ELMA.Cache.CacheServiceManager, EleWise.ELMA.SDK]], EleWise.ELMA.SDK"
/>
        <section name="AppFabricProviderSettings"
type="EleWise.ELMA.Cache.AppFabric.AppFabricProviderSettings,
EleWise.ELMA.Cache.AppFabric"/>

        <section name="main"
type="EleWise.ELMA.Configuration.MainBaseSettingsSection,
EleWise.ELMA.SDK"/>
        <section name="blobStore"
type="EleWise.ELMA.Configuration.GenericProviderFeatureSection`1[[EleWis
e.ELMA.Runtime.Providers.BLOBStoreProviderManager, EleWise.ELMA.SDK]],
EleWise.ELMA.SDK"/>
        <section name="settingsStore"
type="EleWise.ELMA.Configuration.GenericProviderFeatureSection`1[[EleWis
e.ELMA.Runtime.Providers.SettingsStoreProviderManager,
EleWise.ELMA.SDK]], EleWise.ELMA.SDK"/>
        <section name="fileStore"
type="EleWise.ELMA.Configuration.GenericProviderFeatureSection`1[[EleWis
e.ELMA.Runtime.Providers.FileStoreProviderManager, EleWise.ELMA.SDK]],
EleWise.ELMA.SDK"/>
  </configSections>
  <connectionStrings>
        <add name="MainDB" connectionString="Data
Source=SRV12DBLst;Initial Catalog=ELMAPROD38;Integrated Security=false;
User ID = sa; Password=123456"/>
```

```
    </connectionStrings>
    <main                                    connectionStringName="MainDB"
type="EleWise.ELMA.Extensions.MSSQL.MSSQLProvider,
EleWise.ELMA.Extensions.MSSQL"  backupEnabled="false"/>
    <blobStore defaultProvider="MemoryBLOBStoreProvider">
        <providers>
            <clear/>
            <add                          name="MemoryBLOBStoreProvider"
type="EleWise.ELMA.Runtime.Providers.MemoryBLOBStoreProvider,
EleWise.ELMA.SDK"/>
        </providers>
    </blobStore>
    <settingsStore defaultProvider="NHSettingsProvider">
        <providers>
            <clear/>
            <add                                name="NHSettingsProvider"
type="EleWise.ELMA.Runtime.Providers.Impl.NHSettingsStoreProvider,
EleWise.ELMA.SDK"/>
        </providers>
    </settingsStore>
    <fileStore defaultProvider="FSProvider">
        <providers>
            <clear/>
            <add                                          name="FSProvider"
type="EleWise.ELMA.Runtime.Providers.Impl.FileSystemFileStoreProvider,
EleWise.ELMA.SDK"       filesPath="\\ELMAConfig\ELMAShared\Config\Files"
tempFilesPath="\\ELMAConfig\ELMAShared\Config\TempFiles"/>
        </providers>
    </fileStore>
    <cacheService defaultProvider="AppFabric">
            <providers>
                <clear/>
                <add                                        name="AppFabric"
type="EleWise.ELMA.Cache.AppFabric.AppFabricCacheService,
EleWise.ELMA.Cache.AppFabric"/>
            </providers>
        </cacheService>
    <AppFabricProviderSettings                channelOpenTimeout="20000"
requestTimeout="10000"       dataCacheServiceAccountType="DomainAccount"
MainCacheName="ELMACache"  NHCacheName="ELMACache_NH"  LockTimeout="30000"
LocksRegionName="ELMACache_NH_Locks">
            <hosts>
```

```
                <host name="127.0.0.1" cachePort="22233"/>
            </hosts>
        </AppFabricProviderSettings>
    </configuration>
```

4.3.2. Switching servers to using a single configuration file

To switch servers to using the same configuration file on each application server, specify the configuration file location in the **connection.config** file in the **C:\ELMA3-Enterprise\Web** folder:

```
<?xml version="1.0"?>
<connectionStrings>
    <add name="ConfigurationFile" connectionString="\\
ELMAConfig\ELMAShared\Config\configuration.config"/>
</connectionStrings>
```

In this file you specify, that the configuration located in the shared folder on the server 93.158.134.35 ELMAConfig should be used.

Next, you need to share the **Web** folder, located in **C:\ELMA3-Enterprise\Web**. Later on, it will not be necessary to remotely use or change the content of this folder. However, it will help to gather log files and diagnostics results.

Attention! ELMA3-Enterprise application pools must be disabled, since you have specified that AppFabric cache cluster and ASPState should be used, but have not completed their configuration. If for some reason the application pool was started, then you will most likely see an ELMA server start error. Stop the ELMA3-Enterprise pool and continue configuring.

4.4. Installing and Creating Cache Cluster

This chapter discusses two cache cluster options for a web farm – Redis and AppFabric 1.1. In this guide, we recommend using Redis cache cluster. Note that this cluster can be replaced with its analog AppFabric 1.1.

4.4.1. Installing and creating the Redis cache cluster

This section deals with installing, configuring and creating the Redis cache cluster.

4.4.1.1 Before you start

To work with the Redis cache cluster, you must have the following:

- ELMA 3.9.22, 3.10.12, 3.11.2 or higher;
- At least three Redis cache cluster servers with Unix OS (Linux, OSX, OpenBSD, NetBSD, FreeBSD and derived OS).

After that you can download the latest **Stable** version of Redis from the official website.

4.4.1.2 Building and installing cache cluster

An example of the Redis cache cluster installation on Ubuntu 14/16.

The installation process consists of several steps and is described in the table below. Each step requires a series of commands.

sudo apt-get update **sudo apt-get install build-essential** **sudo apt-get install tcl8.5**	Install additional packages, required for assembling a Redis package. In case of Ubuntu, those are **build-essentials** and **tcl 8.5** (or higher).
wget http://download.redis.io/releases/redis-stable.tar.gz	Download the Redis package to any folder (e.g. /home/<user>). (get a URL on the page).
tar xzf redis-stable.tar.gz	Unzip the package (the name of the downloaded package may be different).
cd redis-stable	Open the folder with the unzipped package (the name may also be different).
make	Assemble the package.

make test	Check the assembled package.
make install **cd utils** **./install_server.sh**	Install the package and register it for running. **./install_server.sh** – this command is interactive; when requesting parameters, you can use default values. After running these commands, the binary files **redis-cli** and **redis-server** will be installed to **…/usr/local/bin**. Additionally, the service with the parameters specified in **install_server.sh** will be registered.
sudo service redis_6379 start	Start the service. **Note that** the server name may be different, depending on the values, specified in the **install_server.sh** command. To stop the service, type **'stop'** instead of **'start'**.

4.4.1.3 Configuring the Redis cache cluster

You can find links to official Redis documentation below:

- configuring a cluster;
- configuring a high-availability solution.

This guide describes how to configure the simplified variant. The first server is Master, the second and consequent are Slave servers.

Step 1. To configure the Redis cache cluster configuration, you must edit the **/etc/redis/redis.conf** file on each of the servers. To do this, you can use the command line text editor **Nano**. In order to open the configuration file, open the terminal and execute the **nano/<absolute file path>** command (e.g., **nano/etc/redis/redis.conf**). To perform a quick search in **Nano**, use the keyboard shortcut **Ctrl+W.**

In addition, the following commands may be helpful while working in the terminal:

- **ls** – display the contents of the current directory;
- **cd** – go to your home directory.

#bind 127.0.0.1	Make the server available for all the IP addresses of this server. **#** – comment line, i.e. In this case, it makes the Redis service available for external addresses.

To ensure maximum performance, disable data backup copying to the disk. To do so:

#save	comment out all the lines starting with 'save' (e.g. #save 900 1).
appendonly no	Set the value.

Step 2. To configure Master server, execute the following commands:

tcp-keepalive 60	Set the parameter.
requirepass your_redis_password	Set the Master access password.

If necessary, set the used memory limit (upon reaching the threshold, some values will be removed depending on the **maxmemory-policy** strategy, defined in the parameter):

maxmemory <bytes>	Specify the maximum volume in bytes or leave the line commented.
maxmemory-policy volatile-lru	Set **volatile-lru** or another value, specified in the comments in this file **redis.conf**.
maxmemory-policy noeviction	Set **noeviction** if using a limit is not required.

Step 3. To configure the Slave server, you must additionally execute the following commands:

requirepass your_redis_password	Specify the Slave access password (for the correct functioning of ELMA, specify the same password, as for the Master).
slaveof your_redis_master_ip 6379	Specify the Master address and port.
masterauth your_redis_password	Specify the Master password.

Step 4. To complete the configuration, restart all the servers by running the following command:

sudo service redis_6379 restart	Restart all the servers (first Master, then Slaves).

4.4.1.4 Configuring Redis servers for failover operation

Attention! Cache Redis.Cluster is not used in ELMA! Use Redis Sentinel instead.

First, Configure Redis Master and several Redis Slave servers (see 4.4 above). Then configure Redis sentinel architecture according to the documentation.

4.4.1.4.1 Configuring Redis Sentinel

If Redis.sentinel is configured and the master server is down, then one of the slave servers will be reconfigured as the master. This decision is made by all the started Redis.sentinel servers, having reached a quorum (the **sentinel monitor** parameter in the configuration). Once the master server is up again, it will be reconfigured as a slave.

Step 1. To configure, create a file **.../etc/redis/redis.sentinel.conf** on each server with the help of the **nano/etc/redis/redis.sentinel.con command:**

```
# *** IMPORTANT ***
#
# By default Sentinel will not be reachable from interfaces different than
# localhost, either use the 'bind' directive to bind to a list of network
# interfaces, or disable protected mode with "protected-mode no" by
# adding it to this configuration file.
#
# Before doing that MAKE SURE the instance is protected from the outside
# world via firewalling or other means.
#
# For example you may use one of the following:
#
# bind 127.0.0.1 192.168.1.1
protected-mode no
```

```
# port <sentinel-port>
# The port that this sentinel instance will run on
port 16379

# By default Redis does not run as a daemon. Use 'yes' if you need it.
# Note that Redis will write a pid file in /var/run/redis.pid when daemonized.
daemonize yes

# Specify the log file name. Also the empty string can be used to force
# Redis to log on the standard output. Note that if you use standard
# output for logging but daemonize, logs will be sent to /dev/null
logfile /var/log/redis/redis-sentinel.log

# sentinel monitor <master-name> <ip> <redis-port> <quorum>
#
# Tells Sentinel to monitor this master, and to consider it in O_DOWN
# (Objectively Down) state only if at least <quorum> sentinels agree.
#
# Note that whatever is the ODOWN quorum, a Sentinel will require to
# be elected by the majority of the known Sentinels in order to
# start a failover, so no failover can be performed in minority.
#
# Slaves are auto-discovered, so you don't need to specify slaves in
# any way. Sentinel itself will rewrite this configuration file adding
# the slaves using additional configuration options.
# Also note that the configuration file is rewritten when a
# slave is promoted to master.
#
# Note: master name should not include special characters or spaces.
# The valid charset is A-z 0-9 and the three characters ".-_".
sentinel monitor elma-redis your_redis_master_ip  6379 2

# sentinel auth-pass <master-name> <password>
#
# Set the password to use to authenticate with the master and slaves.
```

```
# Useful if there is a password set in the Redis instances to monitor.
#
# Note that the master password is also used for slaves, so it is not
# possible to set a different password in masters and slaves instances
# if you want to be able to monitor these instances with Sentinel.
# However you can have Redis instances without the authentication enabled
# mixed with Redis instances requiring the authentication (as long as the
# password set is the same for all the instances requiring the password) as
# the AUTH command will have no effect in Redis instances with authentication
# switched off.
sentinel auth-pass elma-redis your_redis_password


# sentinel down-after-milliseconds <master-name> <milliseconds>
#
# Number of milliseconds the master (or any attached slave or sentinel) should
# be unreachable (as in, not acceptable reply to PING, continuously, for the
# specified period) in order to consider it in S_DOWN state (Subjectively
# Down).
#
# Default is 30 seconds.
sentinel down-after-milliseconds elma-redis 30000


# sentinel failover-timeout <master-name> <milliseconds>
# sentinel failover-timeout <master-name> <milliseconds>
#
# Specifies the failover timeout in milliseconds. It is used in many ways:
#
# - The time needed to re-start a failover after a previous failover was
#   already tried against the same master by a given Sentinel, is two
#   times the failover timeout.
#
```

```
# – The time needed for a slave replicating to a wrong master according
#   to a Sentinel current configuration, to be forced to replicate
#   with the right master, is exactly the failover timeout (counting
since
#   the moment a Sentinel detected the misconfiguration).
#
# – The time needed to cancel a failover that is already in progress
but
#   did not produced any configuration change (SLAVEOF NO ONE yet not
#   acknowledged by the promoted slave).
#
# – The maximum time a failover in progress waits for all the slaves
to be
#   reconfigured as slaves of the new master. However even after this
time
#   the slaves will be reconfigured by the Sentinels anyway, but not
with
#   the exact parallel-syncs progression as specified.
#
# Default is 3 minutes.
sentinel failover-timeout elma-redis 180000
```

Step 2. Fill the Master-server data:

sentinel monitor elma-redis your_redis_master_ip 6379 2	Specify the Master address and port and the value for reaching a quorum.
sentinel auth-pass elma-redis your_redis_password	Specify a password for accessing the Master.

Step 3. Create a folder **…/var/log/redis/** with the help of the **sudo mkdir/var/log/redis/** command and configure access to it.

Step 4. To complete the configuration of **Redis Sentinel**, run the following command:

protected-mode no	Configure bindings to network interfaces.

Step 5. Attention! When a sentinel server is operational, configuration files of the server and of each Redis server change, therefore you must grant access to overwriting them. You can use the **sudo chmod a+rwx/etc/redis/** command for it.

To configure a sentinel server as a service, create a file **/etc/init.d/redis-sentinel** on each server and configure access permissions.

```
#!/bin/bash
# Start/Stop/restart script for Redis Sentinel

NAME=`basename ${0}`
EXEC=/usr/bin/redis-server
PIDFILE="/var/run/redis/${NAME}.pid"
CONF="/etc/redis/redis.sentinel.conf"

PID=`cat $PIDFILE 2> /dev/null`
case "$1" in
     start)
          echo "Starting $NAME ..."
          touch $PIDFILE
          exec $EXEC $CONF --sentinel --pidfile $PIDFILE
          ;;
     stop)
          echo "Stopping $NAME ..."
          kill -9 $PID
          ;;
     restart)
          echo "Restarting $NAME ..."
          $0 stop
          sleep 2
          $0 start
          ;;
     *)
          echo "Usage $0 {start|stop|restart}"
          ;;
esac
```

EXEC=/usr/bin/redis-server or **/usr/local/bin/redis-server** (executable redis-server application; you can take it from the file /etc/init.d/redis-server)

Step 6. Create a folder **.../var/run/redis**, configure access to it.

Step 7. To complete the configuration of **Redis Sentinel,** run the following command:

sudo systemctl unmask redis-sentinel.service	Register the service.

Step 8. Start the Sentinel-server on each server:

sudo service redis-sentinel start or **sudo redis-server /etc/redis/ redis.sentinel.conf –sentinel**	Start server. **/etc/redis/ redis.sentinel.conf** is the path to the required sentinel configuration.

4.4.2. Installing and creating the AppFabric cache cluster

This section deals with installing, configuring and creating the AppFabric 1.1. cache cluster.

4.4.2.1 Before you start

First, download or copy the AppFabric 1.1 installation file from the official site. In case the downloaded file is a web installer, you will need an active internet connection.

Second, enable automatic Windows update at least in the "Check for updates but let me choose whether to download and install them" mode. In this case, the actual access to the Internet and the installation of updates is not required.

4.4.2.2 Installing and creating a cache cluster

Initial installation of AppFabric includes creating a cache cluster and automatic creation of a database for saving a configuration – these operations are performed once on the first application server (in this case, it is 93.158.134.3 – SRV12-1). After that, the other servers join the cluster as described below.

Step 1. Start installation of AppFabric and select all the suggested features.

Attention! If automatic update is not enabled, AppFabric installation will not start. If automatic update is enabled, but the AppFabric installation still displays an error, disable and re-enable the automatic update.

Step 2. After installing AppFabric on the web server, **AppFabric for Windows Server** will appear in **Start -> All Programs**. Select **Configure AppFabric** (Fig. 67), to start the AppFabric Server configuration wizard:

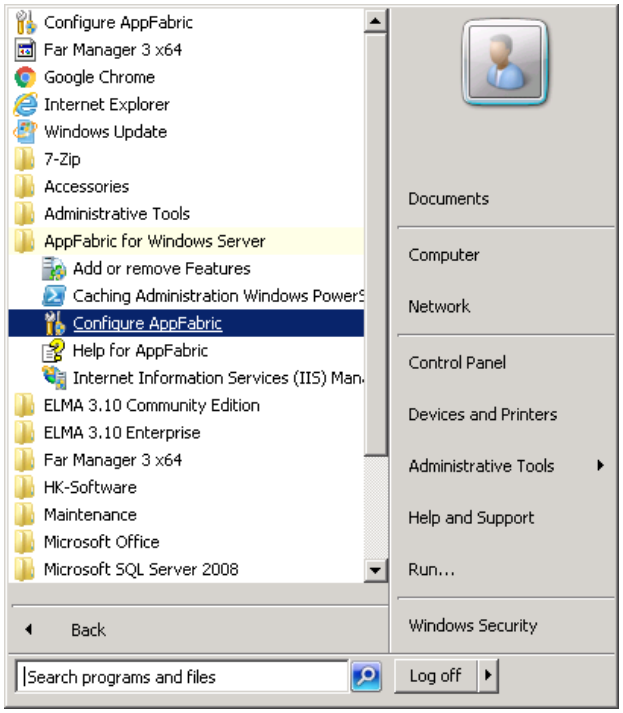

Fig. 67 Start -> All Programs -> AppFabric for Windows Server -> Configure AppFabric

1. At the first step (Fig. 68), you will be offered to participate in the customer experience improvement program. You can refuse to send Microsoft information about any issues and to participate in the program.

   To go to the next step, click **Next**.

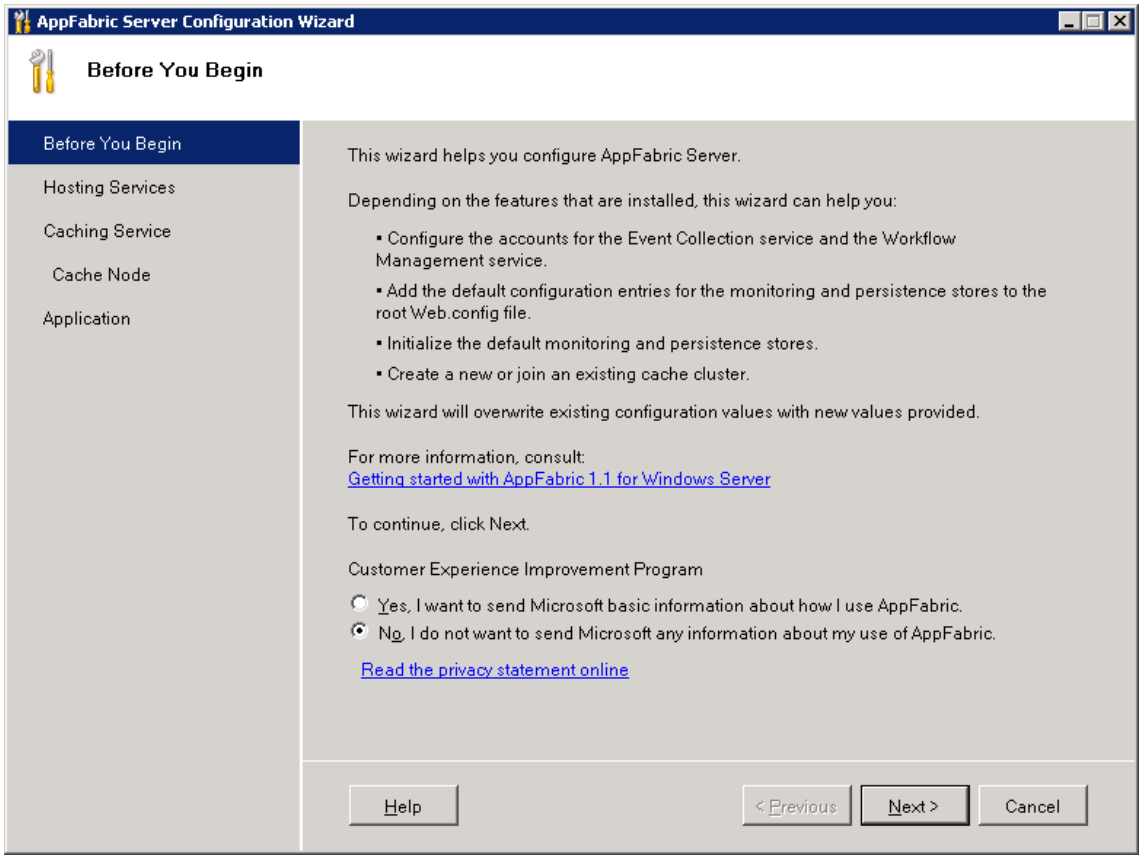

Fig. 68 AppFabric Server configuration wizard. Before Your Begin

2. At the second step of the wizard (Fig. 69), you need to configure hosting services. Leave all the settings unchanged.

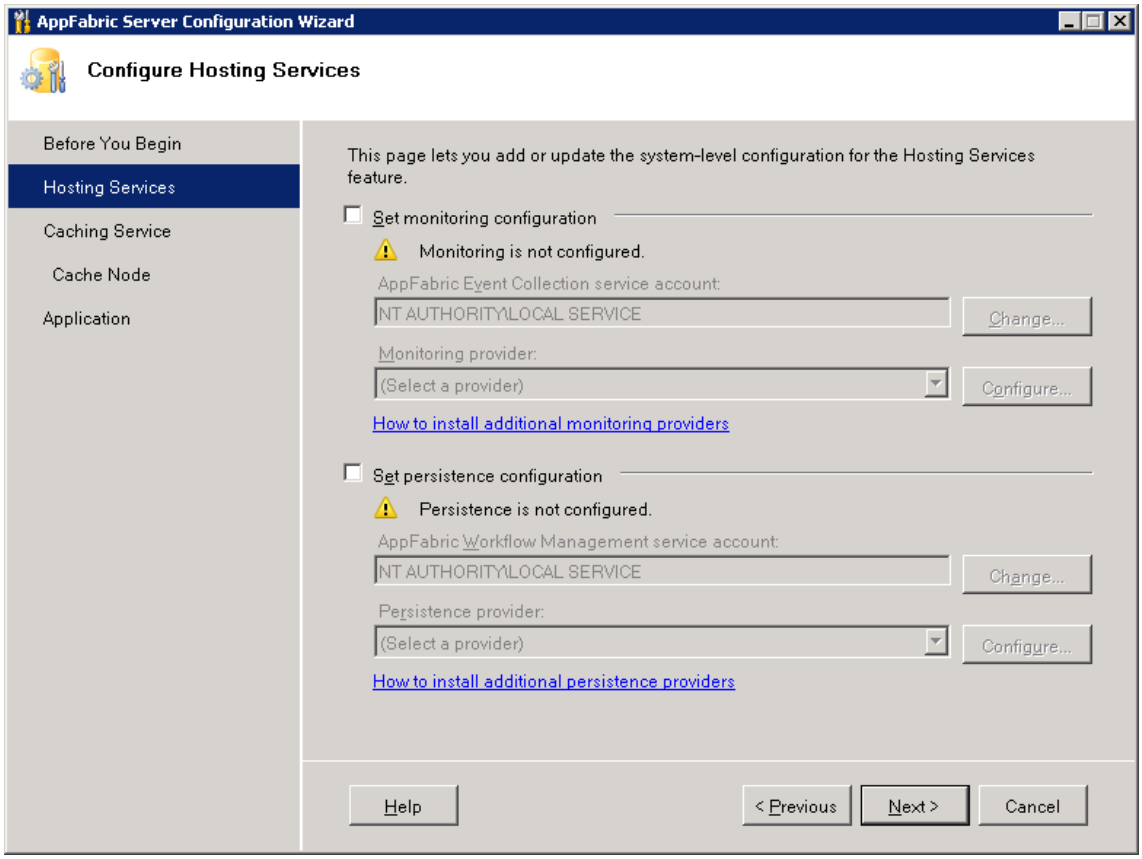

Fig. 69 AppFabric Server configuration wizard. Hosting Services

3. At the third step of the wizard (Fig. 70), you need to configure caching services. Check the **Set Caching Service configuration** box and specify the parameters.

Specify the **Caching Service account** by clicking **Change**... You must select a domain account with administrator permissions for all the machines. Moreover, this account has to have access to the SQL-server. In this example, it is the EleWise\ELMAadmin account.

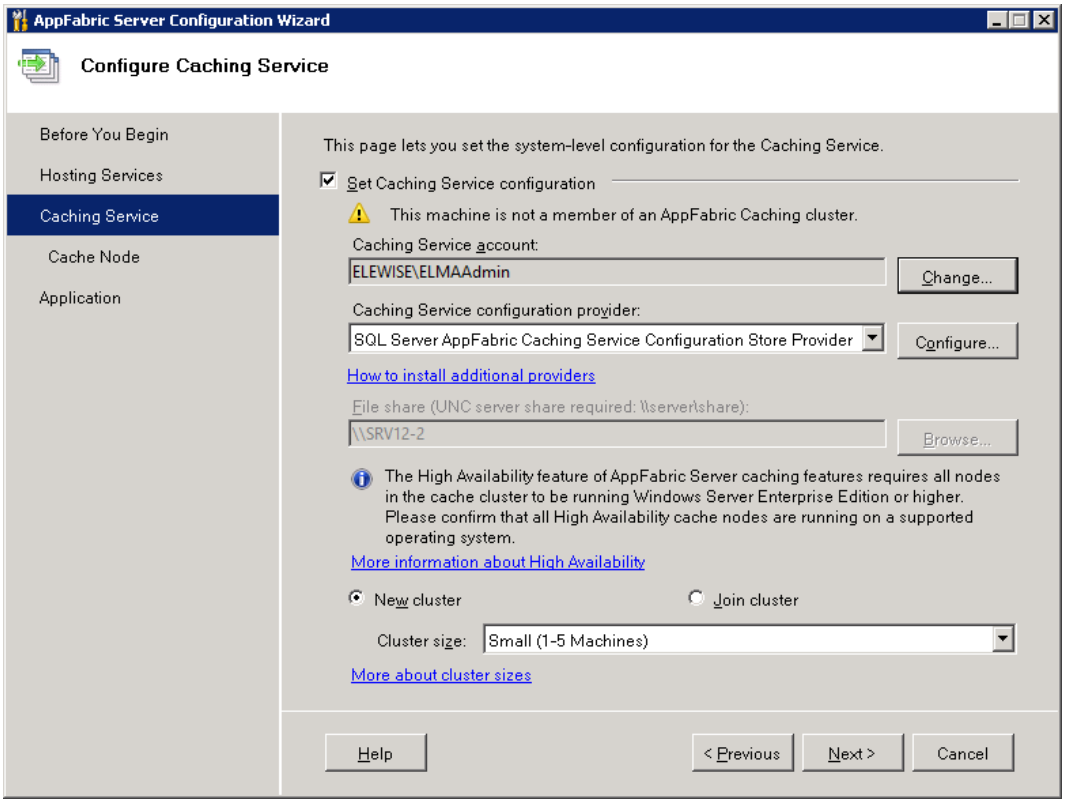

Fig. 70 AppFabric Server configuration wizard. Cache Service

In the drop down list of **Caching Service configuration provider** (Fig. 70), select **SQL Server AppFabric Cashing Service Configuration Store Provider** and click **Configure...** next to the provider name.

The dialog box **AppFabric Server Caching Service configuration Store** will open (Fig. 71).

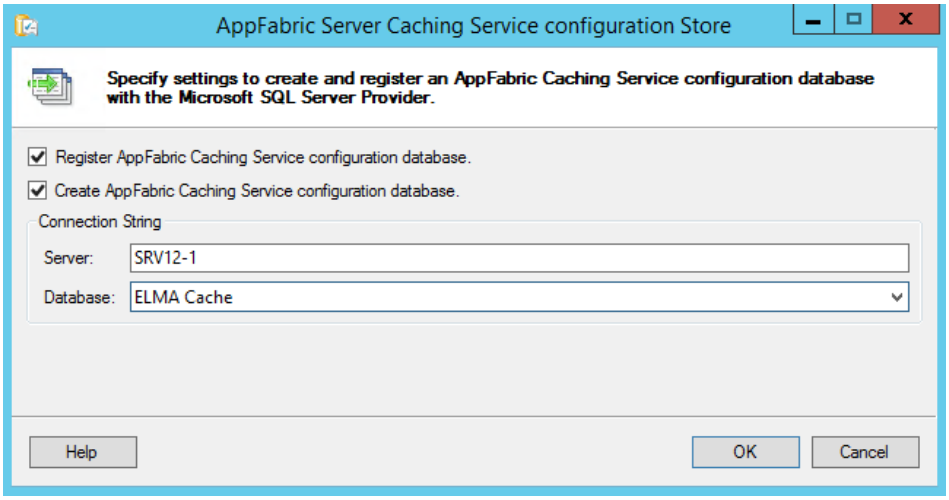

Fig. 71 AppFabric Server Caching Service configuration Store

In this dialog box (Fig. 71), you need to:

- check the **Create AppFabric Caching Service configuration database** box. It will allow creating a database with a configuration and connect other servers to the cache;
- specify the name of your SQL server (VM name\SQL login) in the **Server** field. In this example, the VM name is SRV12-1, login – local (it is not specified);
- specify the name of the database for storing cache in the **Database** field. It is not the same database that was created for ELMA3-Enterprise; it is a new database for storing cache settings. In this example, the name ELMACache has also been specified in the configuration file.

  To save the changes, click **OK**.

You also need to select **New cluster** at the **Caching Service** step (Fig. 70) and select **Small (1-5 Machines)** in the **Cluster size** field.

4. At the fourth step of the wizard (Fig. 72), you can specify the cache node ports and firewall parameters. Leave all the settings unchanged.

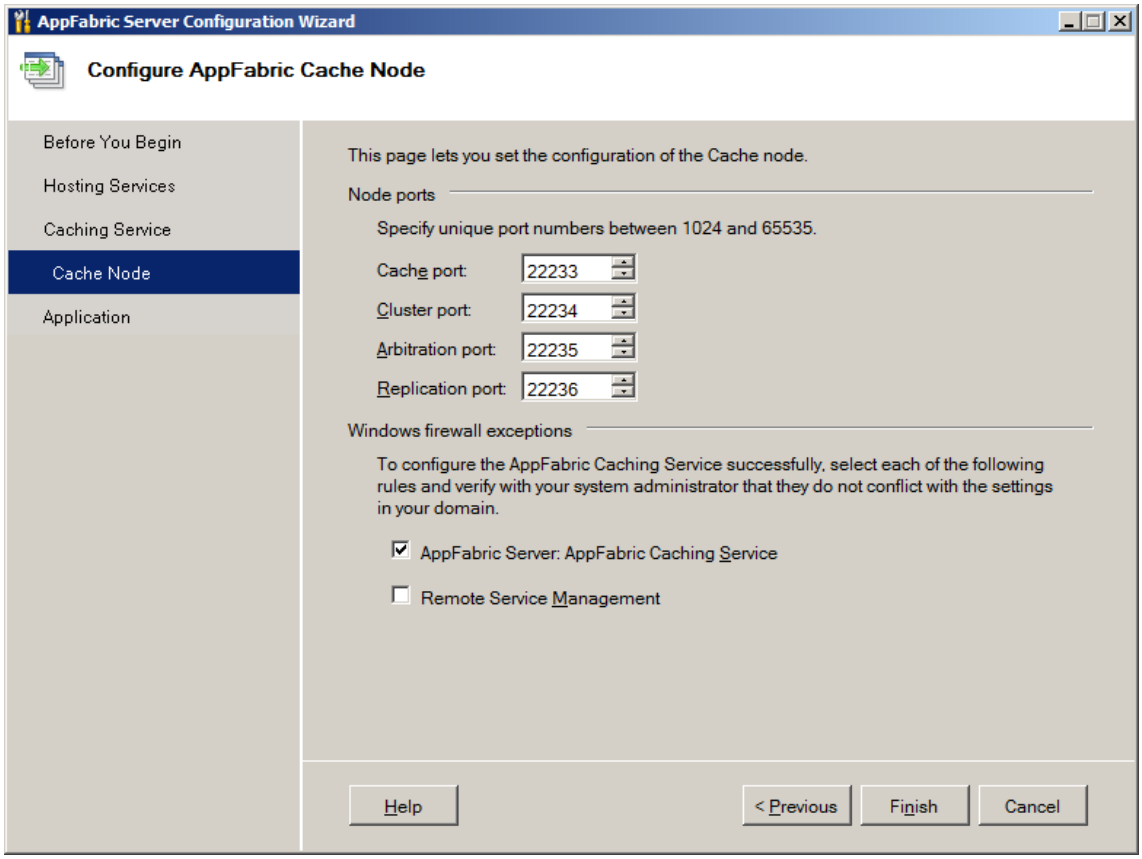

Fig. 72 AppFabric Server configuration wizard. Cache Node

5. At the last step of the wizard (Fig. 73), check the **Start Internet Information Services (IIS) Manager** box and click **Finish**. Now you can proceed to connecting the other servers to the cache cluster.

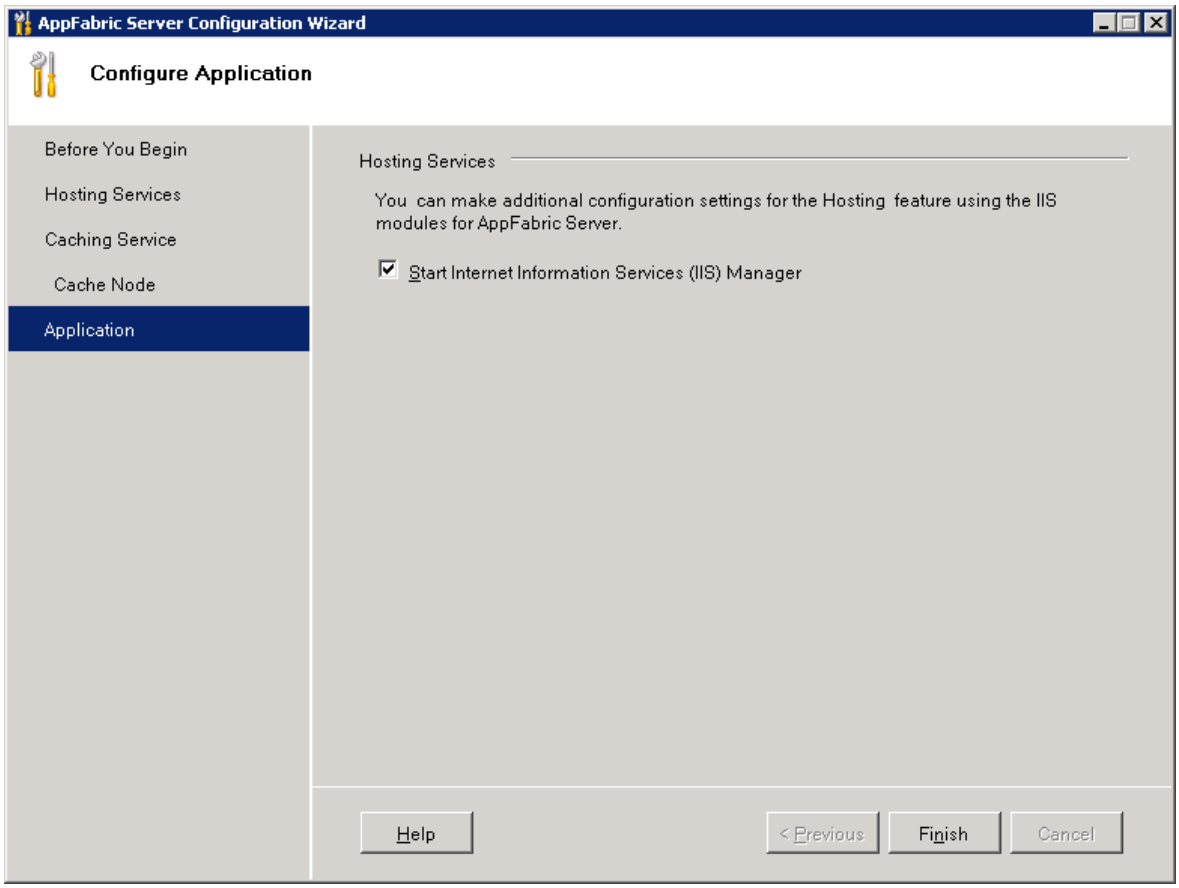

Fig. 73 AppFabric Server configuration wizard. Results

4.4.2.3 Connecting the other servers to the AppFabric cache cluster

For you to be able to connect servers, a cluster must already be created (described выше). Cache must be installed and connected on each application server where ELMA will function.

On the server 93.158.134.3 – SRV12-1, cache has been connected during the installation. For the other server (213.180.204.3 – SRV12-2), you need to apply all the steps, described in this section.

Step 1. Start installation of AppFabric using the same installation file, as in section 4.4.1.3, and select all the features.

Attention! If automatic update is not enabled, AppFabric installation will not start. If automatic update is enabled, but the AppFabric

installation still displays an error, disable and re-enable the automatic update.

Step 2. After installing AppFabric on the web server, **AppFabric for Windows Server** will appear in **Start -> All Programs**. Open it, and select **Configure AppFabric** (Fig. 67), to start the AppFabric Server configuration wizard:

The first two steps are described in the section 4.4.1.3.

At the third step of the wizard (Fig. 74), you need to configure the caching service. Check the **Set Caching Service configuration** box and specify the following parameters:

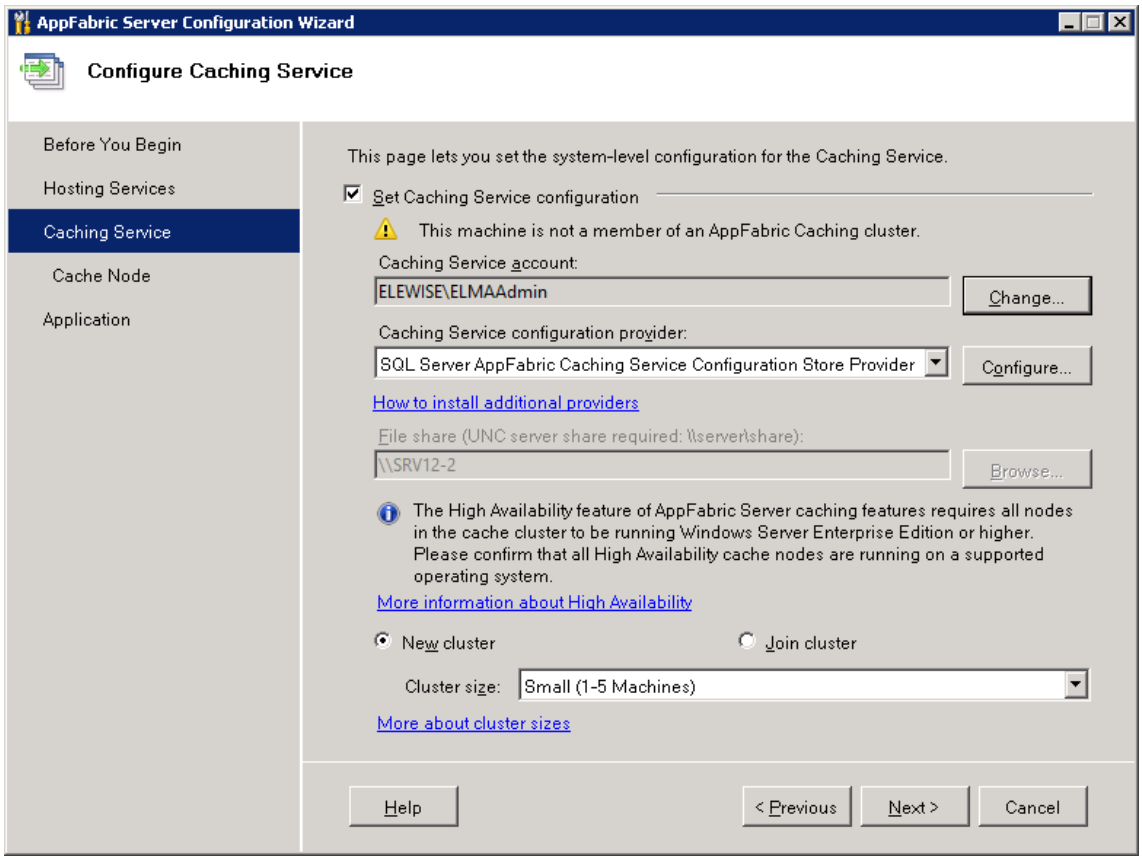

Fig. 74 AppFabric Server configuration wizard. Cache Services

Specify a **Caching service account** by clicking **Change...** You must select a domain account with administrator permissions for all the machines. Moreover, this account has to have access to the SQL-server. In this example, it is the EleWise\ELMAadmin account.

In the drop down list of **Caching Service configuration provider** (Fig. 74), select **SQL Server AppFabric Cashing Service Configuration Store Provider** and click **Configure...** next to the provider name.

The dialog box **AppFabric Server Caching Service configuration Store** will open (Fig. 75). In this dialog box, you need to:

- uncheck the **Create AppFabric Caching Service configuration database** box, since the database was created at the previous step (see выше).
- specify the name of your SQL server (VM name\SQL login) in the **Server** field. In this example, the VM name is SRV12DBLst, login – local (it is not specified).
- specify the name of the database for storing cache in the **Database** field. It is not the same database that was created for ELMA3-Enterprise; it is a new database for storing cache settings. In this example, the name ELMACache has also been specified in the configuration file.

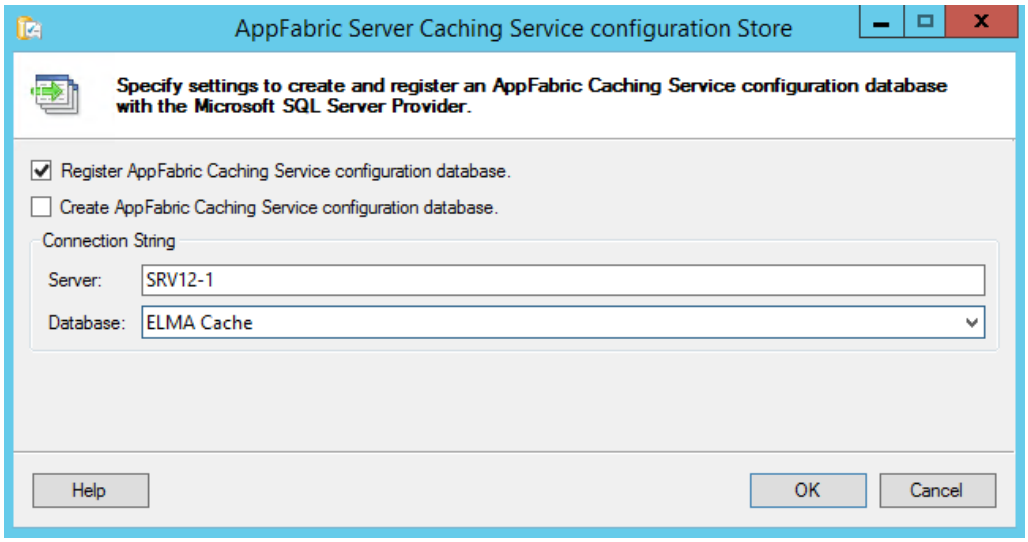

Fig. 75 AppFabric Server Caching Service configuration Store

To save the changes, click **OK**.

You also need to select **Join cluster** at the **Caching Service** step (Fig. 74).

At the fourth step of the wizard (Fig. 72), you can specify the cache node ports and firewall parameters. Leave all the settings unchanged.

At the last step (Fig. 73), leave everything unchanged. Repeat these steps on each of the application servers.

4.4.2.4 Creating a cache and starting a cluster

After you have installed the AppFabric caching service on all the application servers, you need to execute several commands for creating a cache and starting a cluster.

Run the **PowerShell** console as the administrator on any application server (for example, 93.158.134.3 – SRV12-1). Go to **Start -> All Programs -> AppFabric for Windows Server -> Caching Administration Windows PowerShell** (Fig. 76).

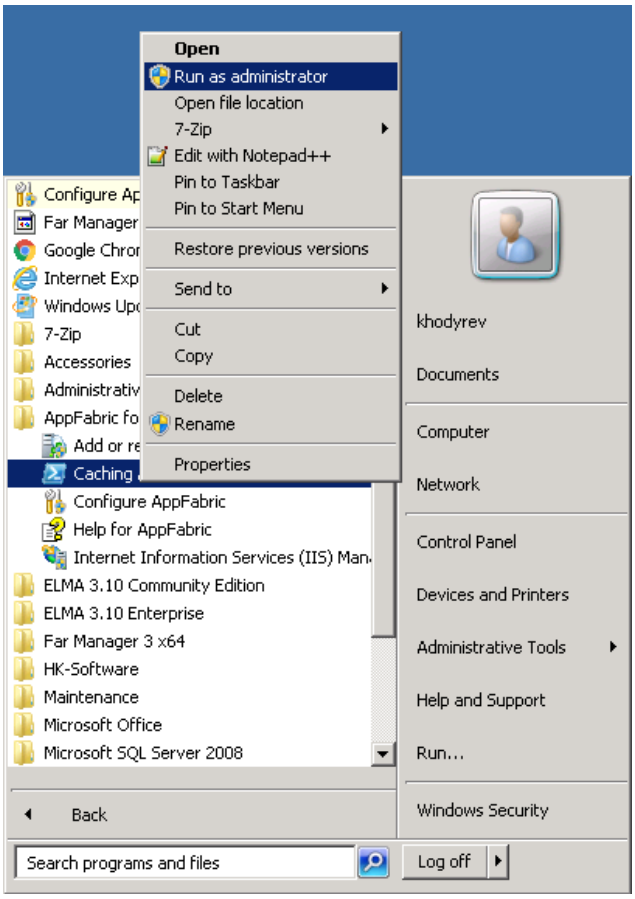

Fig. 76 Start -> All Programs -> AppFabric for Windows Server -> Caching Administration Windows PowerShell

Consequently execute the following commands in the console:

Start-CacheCluster	Start the cluster. It was disabled before that.
Grant-CacheAllowedClientAccount EleWise\ELMAAdmin	Grant permissions to the user, who starts the cache.

Grant-CacheAllowedClientAccount ELMAAdmin	Grant permissions to the user, who starts the cache.
New-Cache ELMACache	Create a cache. **Attention!** The name must match the one specified in the configuration file in the folder \\ELMAConfig\ELMAShared\Config
New-Cache ELMACache_NH	Create a cache. **Attention!** The name must match the one specified in the configuration file in the folder \\ELMAConfig\ELMAShared\Config
New-Cache ELMACache_NH_Locks	Create a cache. **Attention!** The name must match the one specified in the configuration file in the folder \\ELMAConfig\ELMAShared\Config

After that, you can start ELMA3-Enterprise and make sure that each server works without errors.

4.5. Starting the Farm and Creating Extended Configuration Files

At this stage, everything must be ready to start ELMA on each of the application servers. Two additional configuration extension files will be created.

Step 1. Start the ELMA3-Enterprise pool and site in the IIS manager on each application server.

Step 2. Check http://localhost on each application server – ELMA must start consequently on each server. If errors occur, check the section 6.9.

Step 3. Stop the pool on each application server.

Go to the server with the common configuration. In this example, it is 93.158.134.35 – ELMAConfig.

Step 4. Open the folder with the common configuration E:\ELMAShared\Config\WebApplication (Fig. 77). The **WebApplication** folder is created after the first start of the servers in the farm.

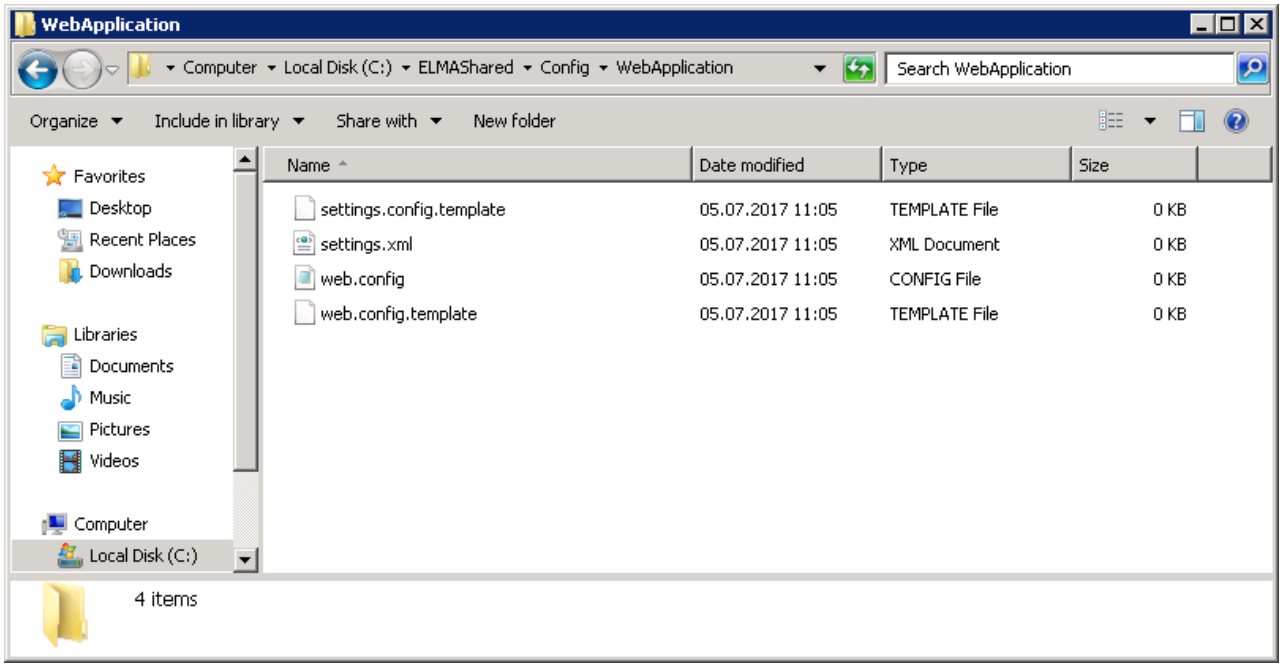

Fig. 77 Folder with the common configuration

Now you will work with the configuration extension files: **web.config.xslt.template** and **settings.config.template**.

Step 5. Rename the file **web.config.xslt.template** to **web.config.xslt**. This file will be used for making changes to ELMA configuration on each server from this folder.

All the lines from these two files are automatically copied when starting ELMA. This way, if you write a new line in the configuration file **web.config.xslt**, it will be copied to all the application servers to C:\\ELMA3-Enterprise\Web\.

Generate a machineKey. You can use online services for generating such keys. A key may look like this:

```
   <machineKey
validationKey="C9ABD1186BBB9C9129586CD47F504395D26E7612344C29F895AB23E4B
D23D9E24B71DFF3B7BC73A3F81F9C0DC5C8A0A85
4EE1F00EFD010381C60659165557875"
decryptionKey="A5E13EF6BD6F1825B904B60DB9B42A9C2E113DE3124C80055F6FA6B7B
2DBE69B" validation="SHA1" decryption="AES" />
```

Add the generated key to the **web.config.xslt** file like this:

```
   <xsl:template
match="/configuration/system.web[not(machineKey)]/trust">
       <xsl:copy>
         <xsl:apply-templates select="node()|@*"/>
       </xsl:copy>
       <machineKey
validationKey="C9ABD1186BBB9C9129586CD47F504395D26E7612344C29F895AB23E4B
D23D9E24B71DFF3B7BC73A3F81F9C0DC5C8A0A854EE1F00EFD010381C60659165557875"
decryptionKey="A5E13EF6BD6F1825B904B60DB9B42A9C2E113DE3124C80055F6FA6B7B
2DBE69B" validation="SHA1" decryption="AES" />
     </xsl:template>
```

In this file, `xsl:template match` defines the computer key, which must be the same on all the servers.

For the farm to work correctly, enable the web session storing provider. To do so, add the following lines to the file:

```
   <xsl:template match="sessionState">
       <!-- Web session storing provider -->
       <sessionState      mode="SQLServer"      sqlConnectionString="data
source=SRV12DBLst;User     ID=sa;Password=123456;"     cookieless="false"
timeout="30"/>
     </xsl:template>
```

DataSource – DBMS connection string. The system will use the ASPState database, created earlier.

The resulting file will look as follows:

```
   <?xml version="1.0" ?>
   <!-- Copy this file to the same folder under the name web.config.xslt,
if it is necessary to configure additional application settings -->
```

```
   <xsl:stylesheet       xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
version="1.0">
     <xsl:output indent="yes" />
     <xsl:template match="node()|@*">
       <xsl:copy>
         <xsl:apply-templates select="node()|@*"/>
       </xsl:copy>
     </xsl:template>
     <xsl:template match="sessionState">
       <!-- Web session storing provider -->
       <sessionState      mode="SQLServer"      sqlConnectionString="data
source=SRV12DBLst;User     ID=sa;Password=1234567;"     cookieless="false"
timeout="30"/>
     </xsl:template>
     <xsl:template
match="/configuration/system.web[not(machineKey)]/trust">
       <xsl:copy>
         <xsl:apply-templates select="node()|@*"/>
       </xsl:copy>
       <machineKey
validationKey="C9ABD1186BBB9C9129586CD47F504395D26E7612344C29F895AB23E4B
D23D9E24B71DFF3B7BC73A3F81F9C0DC5C8A0A854EE1F00EFD010381C60659165557875"
decryptionKey="A5E13EF6BD6F1825B904B60DB9B42A9C2E113DE3124C80055F6FA6B7B
2DBE69B" validation="SHA1" decryption="AES" />
     </xsl:template>
   </xsl:stylesheet>
```

Step 6. Rename the second configuration extension file **settings.config.template** to **settings.config**.

Add the following lines to this file:

<add key="Workflow.QueueLockTimeout" value="1"/> – message queue timeout lock.

<add key="NHibernate.QueryCacheEnabled" value="false"/> – enabling this parameter on the farm will cause cache operations to be executed consequently, which will significantly reduce performance.

<add key="Workflow.StateLockTimeout" value="30"/> – timeout for releasing locked resources in case of parallel tasks in a process.

<add key="Workflow.WaitNextTaskTime" value="3"/> – timeout for opening the next process task. If a task/script/gateway is executed longer, a user will see the loading icon of the next task.

These parameters are required.

Also, add this line to the file:

```
<add key="Notifications.TemplatesEnabled" value="false"/>
```

This command disables system notifications about new tasks, completed tasks, etc. in the messages section.

Since users don't always need to use messages to receive information about new tasks (if ELMA is used as front-end or only for business processes), then notifications can be disabled.

The resulting file will look like this:

```
<?xml version="1.0" ?>
<!-- Copy this file to the same folder under the name settings.config, if it is necessary to configure additional application settings -->
<appSettings>
 <!—Location of the shared folder, where the information about generated reports will be stored --> <add key="FastReportStoragePath" value="\\BM2\ELMAShared\Reports"/>
   <!-- Time, during which the generated report is available (minutes) -->
  <add key="FastReportStorageTimeout" value="10080"/>
  <add key="Workflow.QueueLockTimeout" value="1"/>
  <add key="NHibernate.QueryCacheEnabled" value="false"/>
  <add key="Workflow.StateLockTimeout" value="30"/>
  <add key="Workflow.WaitNextTaskTime" value="3"/>
  <add key="Notifications.TemplatesEnabled" value="false"/>
</appSettings>
```

Step 7. Now start the ELMA3-Enterprise pool on each server.

Make sure that each server is started. Take into account the fact, that only one server starts at a time, while others show "0% Waiting for another server to start".

Step 8. Now you need to enable data reading from the database snapshot. It will greatly lower the probability of locking in case of numerous requests.

To do so, open MS SQL Management Studio, connect to the main DBMS server and execute the following command in the exclusive mode:

```
ALTER DATABASE [DATABASE_NAME] SET READ_COMMITTED_SNAPSHOT ON;
```

While there are no other connections, or if this step has not been fulfilled when creating a database, force-switch the database to the exclusive mode.

Attention! If you intend to restore the database using a backup copy from the server, where this command has not been executed, in order to transfer data or a configuration, you need to repeat it for the relevant database after restoring.

Next, proceed to configuring the farm controller. Users are going to access ELMA via the controller, instead of directly via each application server.

4.6. Configuring the farm controller

At this moment, ELMA must be running on all the application servers.

Allocate a separate machine for the farm controller. In this example, it is 93.158.134.30 – WFCONTROLLER. Users will log in to ELMA via this server.

Step 1. Install and configure IIS (see paragraph 4.2).

Step 2. Install Microsoft Application Request Routing Version 3 for IIS. Three additional components will be installed, including WebFarm framework and URL Rewrite. The easiest way to install it is by using WPI launcher (web platform installer). For the launcher to work, active Internet connection is required (direct or via a proxy server).

To open the web platform installer, go to the IIS Manager and select **Web Platform Installer** in the **Management** unit (Fig. 78).

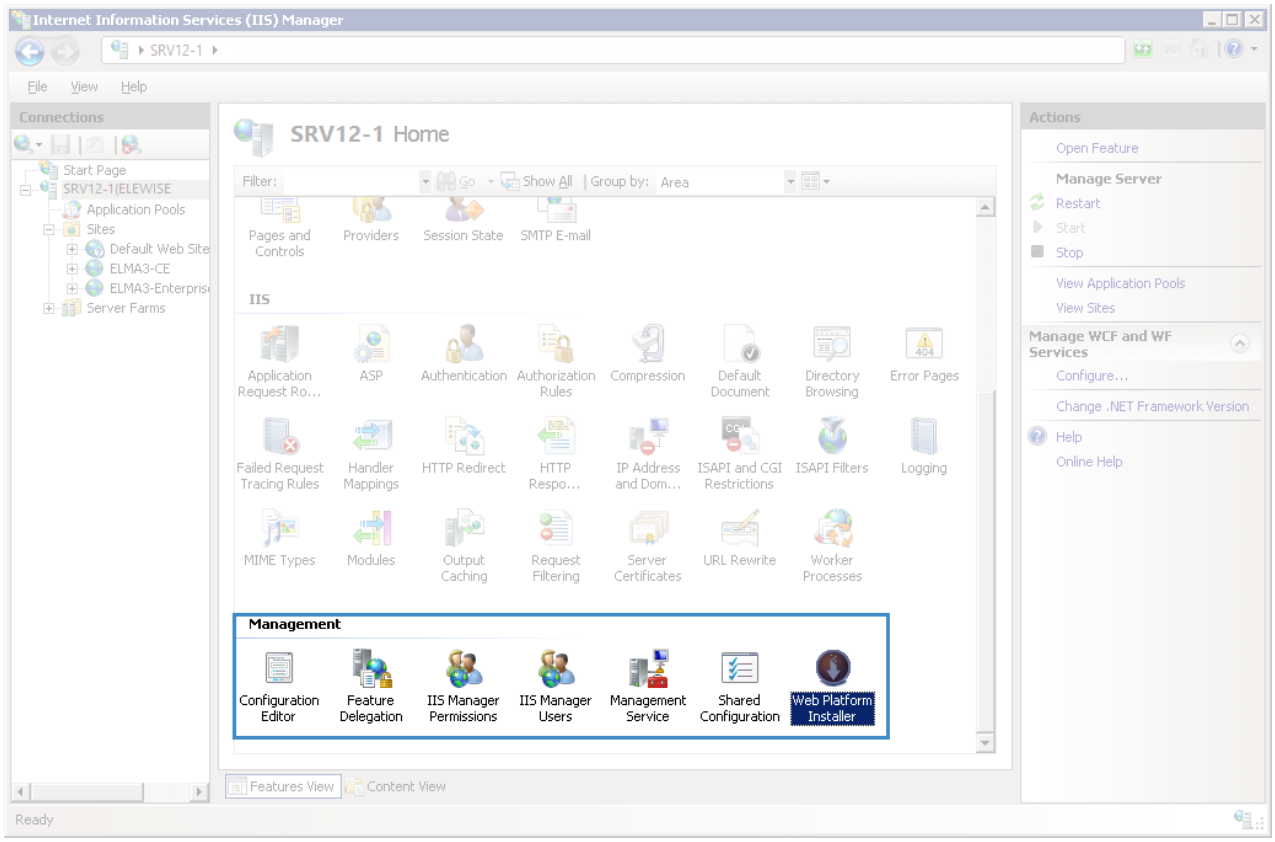

Fig. 78 IIS Manager. Web Platform Installer icon

In WPI launcher, find and install:

- Application Request Routing (ARR)
- URL Rewrite

- WebFarm framework

If all the components are installed successfully, a respective notification will be shown (Fig. 79).

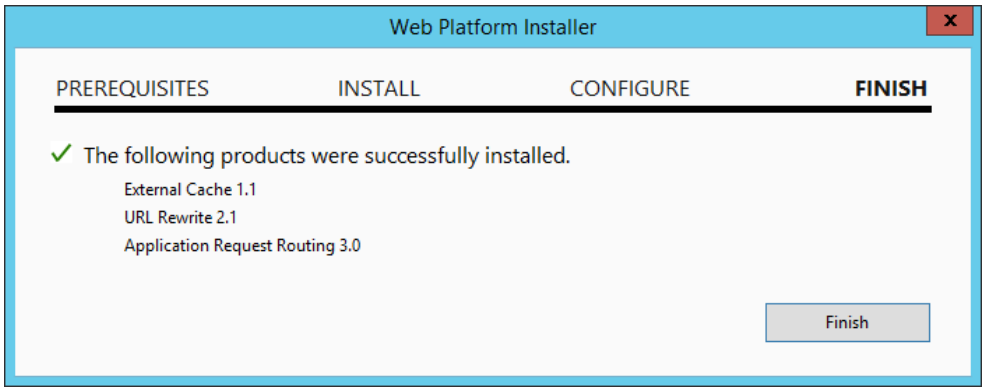

Fig. 79 Web Platform Installer dialogue box

If these packages were not found, you can download and install them manually. Note, that Windows Server 2012 does not support WebFarm Framework version 2, i.e. you need to install version 1 or 3.

Step 3. Open the IIS Manager. In the application pool, select **DefaultAppPool**, click **Advanced Settings**... in the right menu and configure the following settings (Fig. 80):

- In **Idle Time-out (minutes)**, set "0";
- In **Regular Time Interval (minutes)**, set "0".

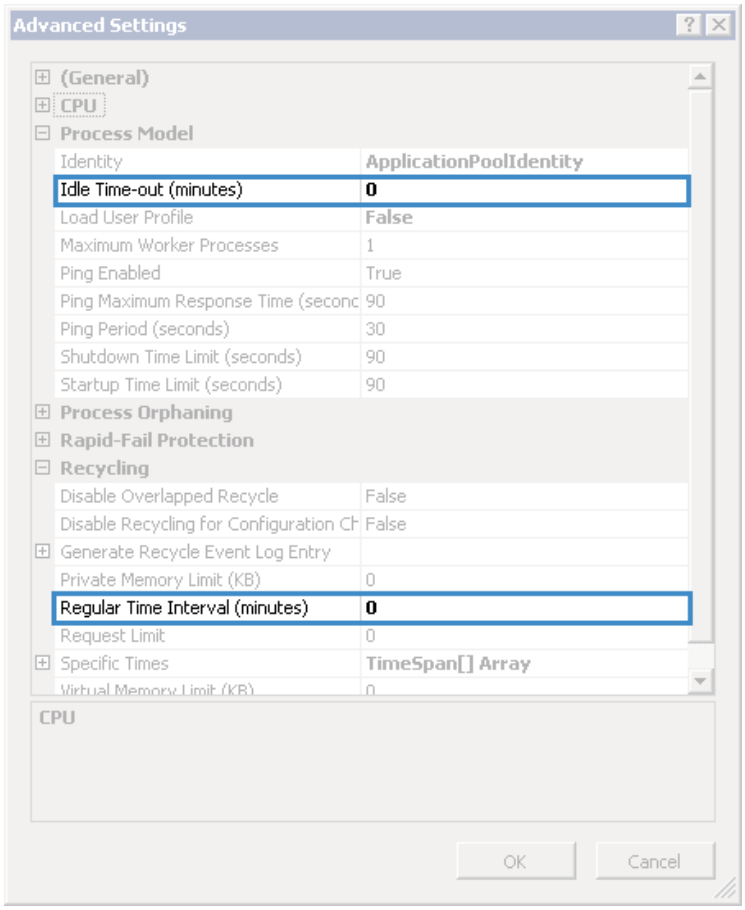

Fig. 80 Pool advanced settings

It is necessary to avoid delays or errors that occur in case of a time-out or unexpected restart, since all the HTTP requests and responses go through the application request routing system.

To save the changes, click **OK**.

Step 4. Since requests will go through the controller, it is important to remove certain default restrictions to avoid receiving errors due to the request length (404 not found) and stand-by period (502).

Attention! If you skip this step, the Tasks tab in the process monitor will return the 404 error.

A controller with a configured URL Rewrite uses the settings of the site that is being requested and the request is redirected to the farm from this site. By default, it is the Default Web Site; therefore, you need to configure all the settings mentioned in this section in the IIS Manager for **Default Web Site.**

1. In the IIS Manager, select **Sites – Default Web site** and click **Configuration Editor** in the **Management** unit (Fig. 81).

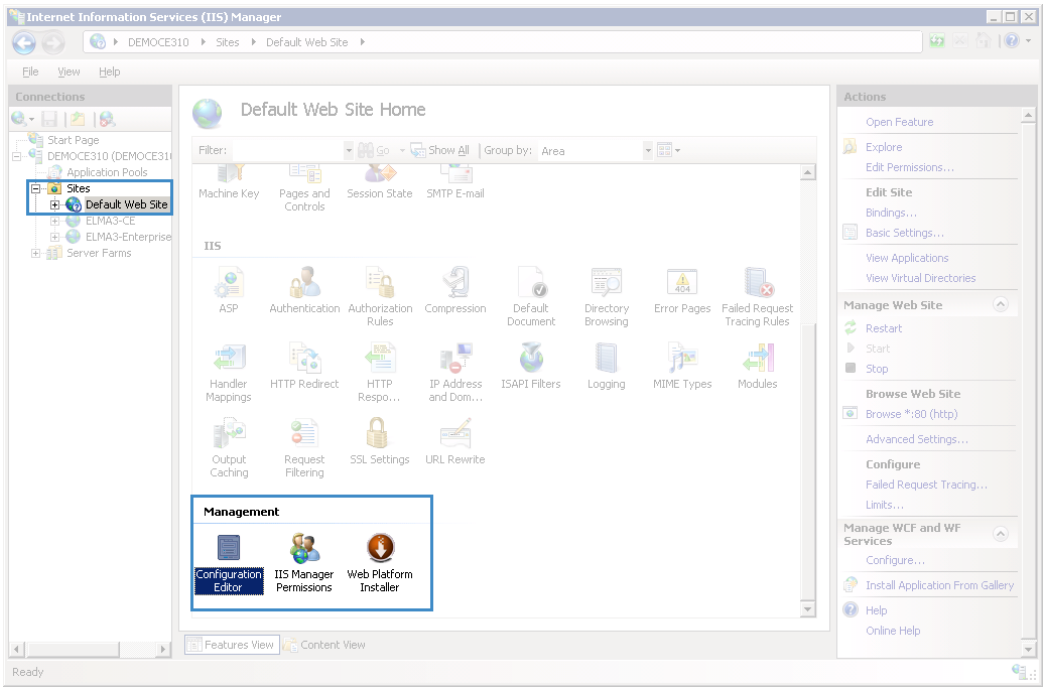

Fig. 81 IIS Manager. Default Web Site. Configuration Editor icon

2. In the configuration editor, open the **Section** drop-down list and go to **system.webServer/security/requestFiltering** (Fig. 82).

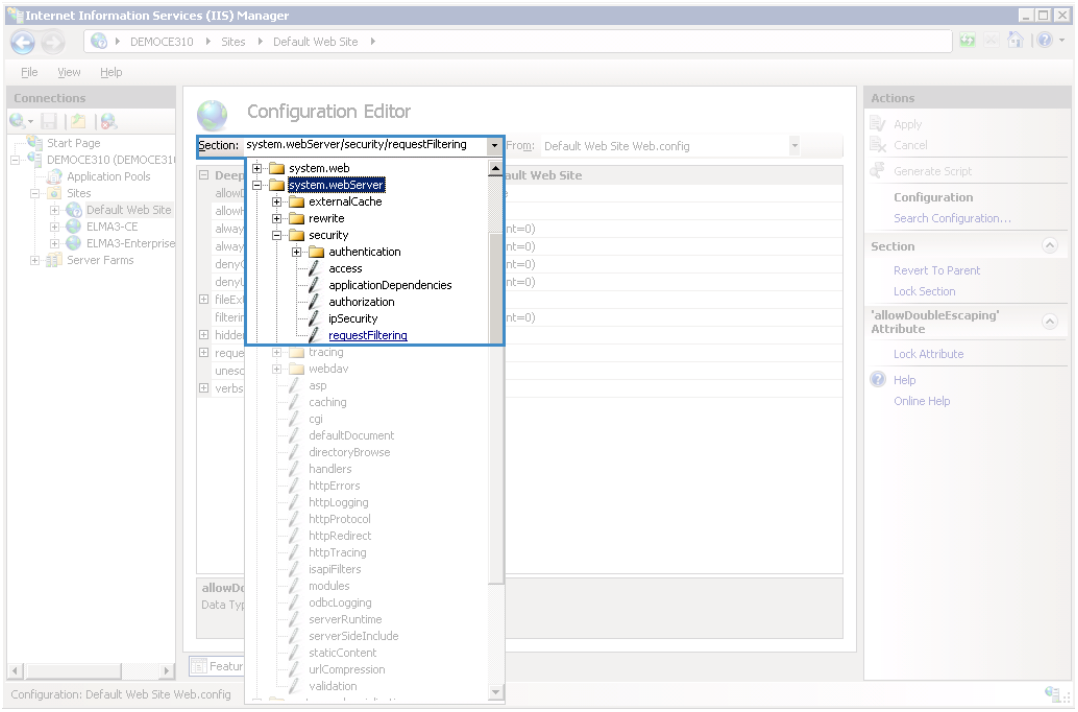

Fig. 82 IIS Manager. Section drop-down list

3. In the opened section, set the **maxQueryString** parameter to "2000000" (Fig. 83) and click **Apply** in the right menu.

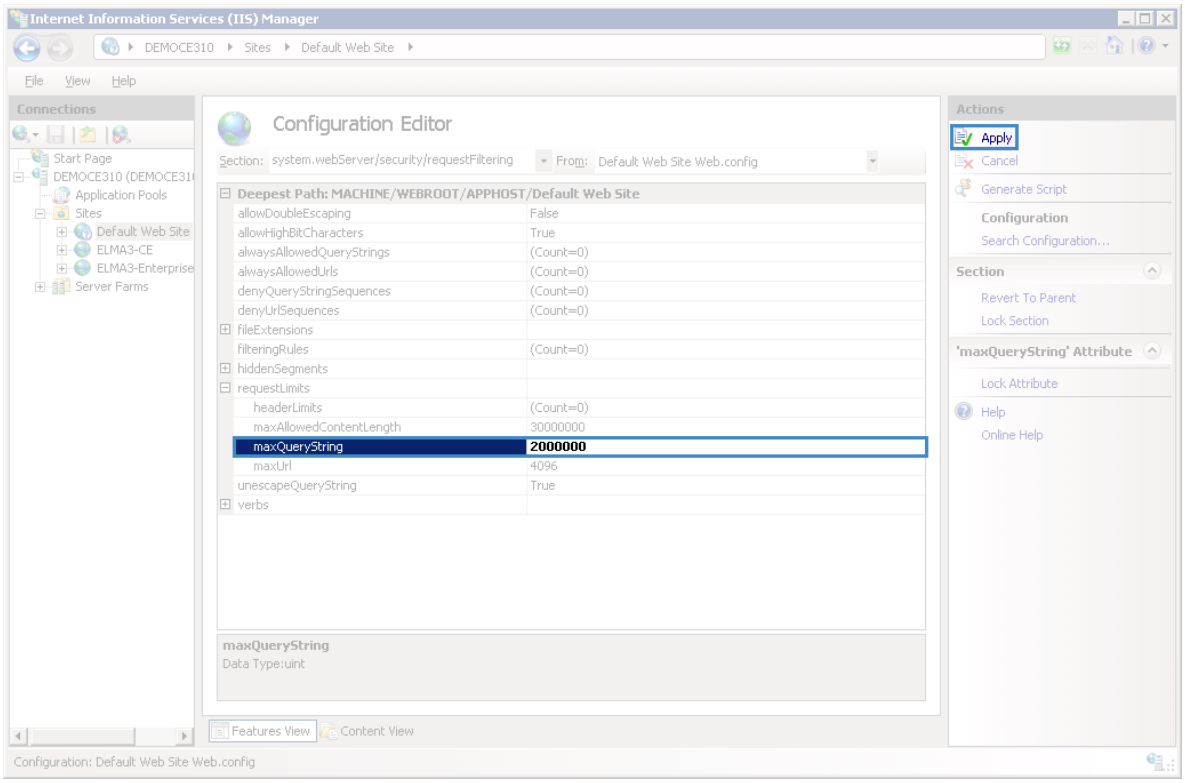

Fig. 83 IIS Manager. "system.webServer/security/requestFiltering" section

This parameter filters requests by the maximum length and this value has to be large enough. That is why it is recommended to use the value "2000000".

Step 5. After completing the **Steps 2-4**, the **Server Farms** item should appear in the IIS Manager (Fig. 84). If this item is absent, try restarting the IIS Manager. If it did not work, then WebFarm Framework was not installed. Install it as described above (see **Step 2**).

Right click on this item to open the context menu and select **Create Server Farm**... (Fig. 84).

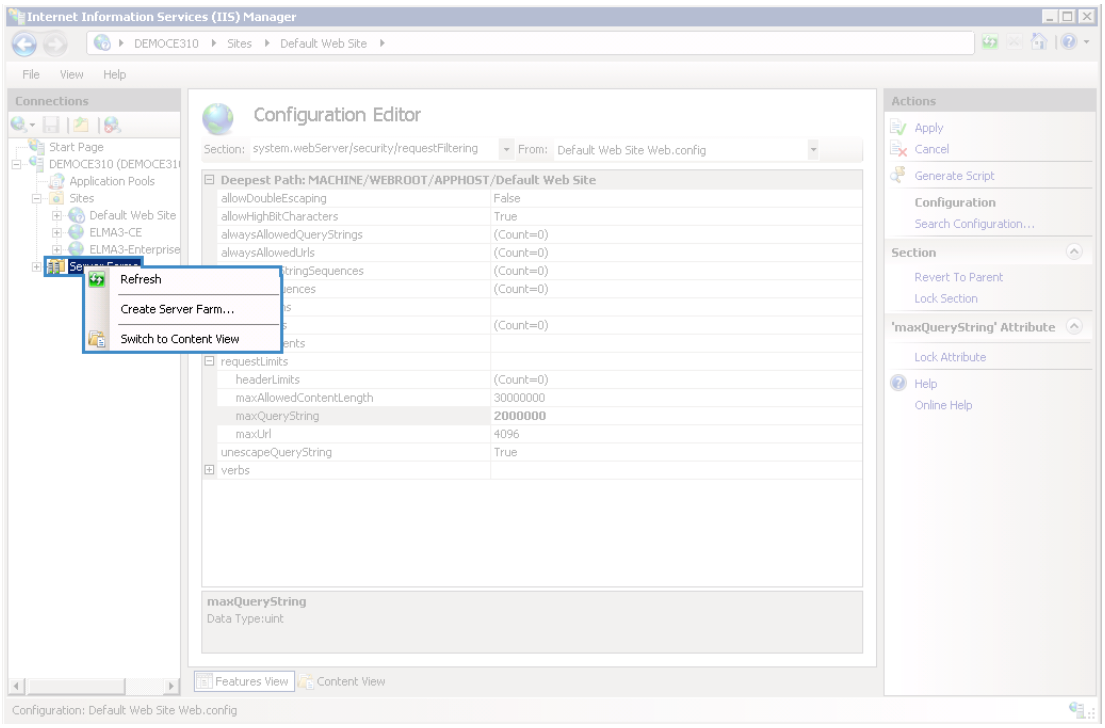

Fig. 84 IIS Manager. Server Farms context menu

The **Create Server Farm** dialog box will open (Fig. 85). In the **Server farm name** field, enter the name of the created farm and click **Next**. In this example, the farm name is **ELMA3**.

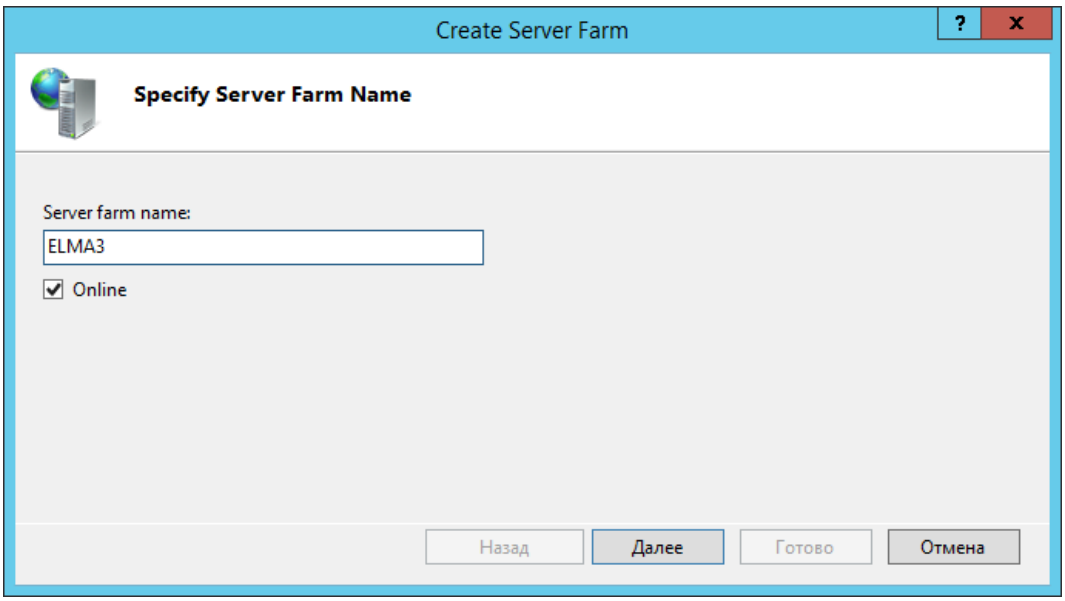

Fig. 85 "Create Server Farm" dialog box

At the next step (Fig. 86), in the **Server address** field enter the IP-address of the application server or its name (in this case, **SRV12-1**) and click **Add**.

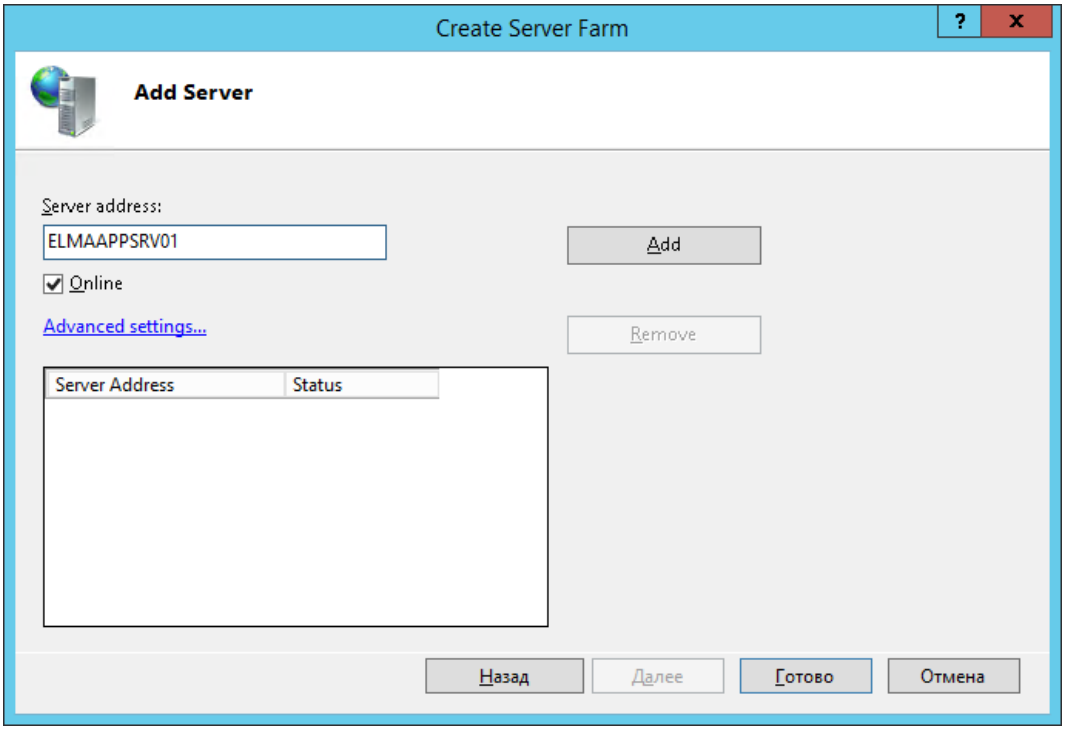

Fig. 86 "Create Server Farm" dialog box

In a similar way, add the IP-addresses or names of all the required application servers (Fig. 87):

- 93.158.134.3 – SRV12-1
- 213.180.204.3 – SRV12-2

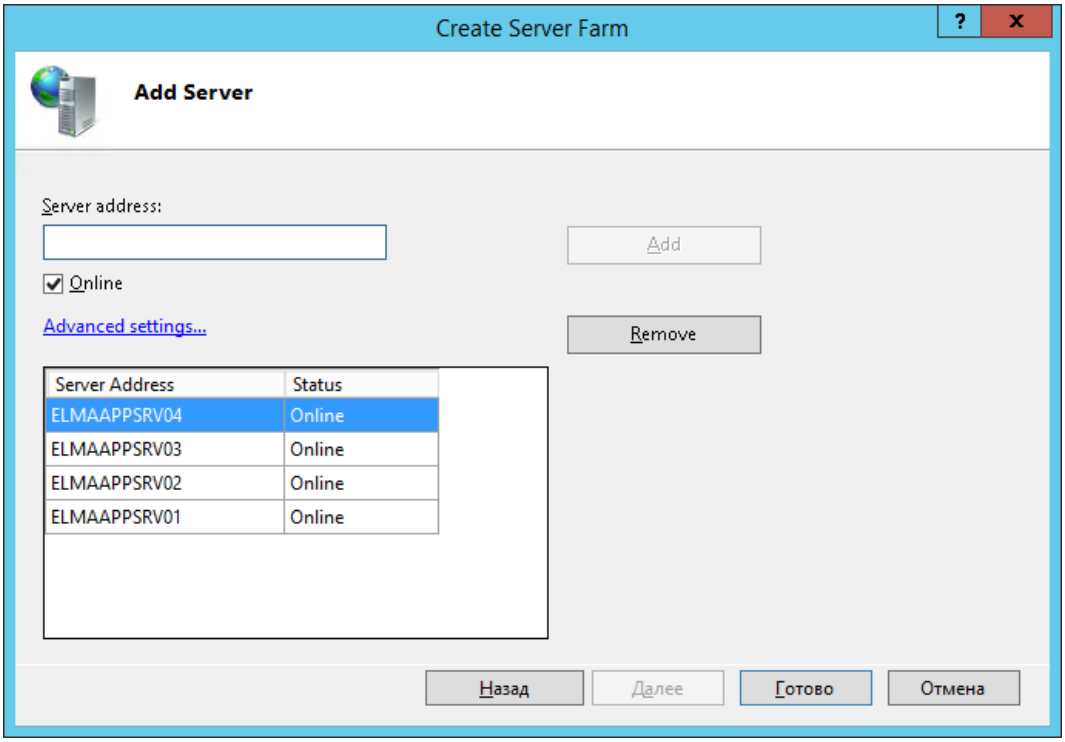

Fig. 87 "Create Server Farm" dialog box

To save the changes, click **Finish**.

If all the components were installed successfully, a dialog box will open (Fig. 88) notifying, that the **URL Rewrite** rule will be created automatically for routing incoming requests to this server farm. Click **Yes** in this dialog box.

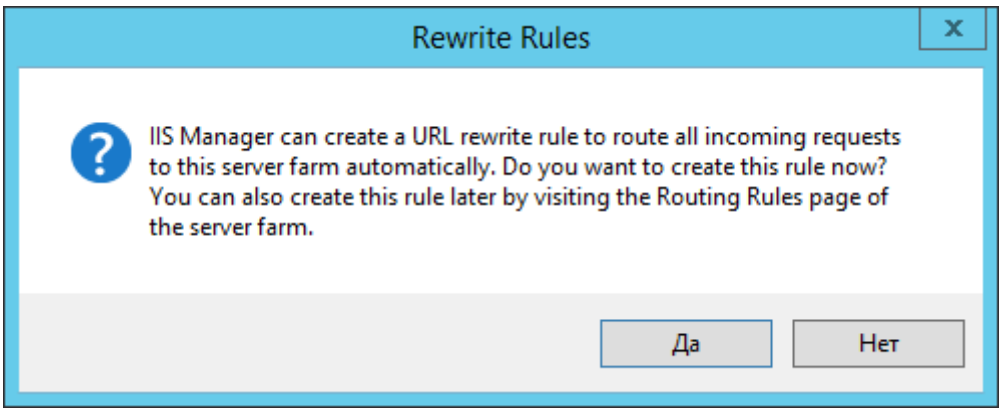

Fig. 88 Notification about creating a URL Rewrite rule

If the dialog box does not open, make sure that the **URL Rewrite** is installed: go back and return the **Step 5**. If necessary, you can create the rule manually (see section 6.6).

After that, the **ELMA3** nested item will be added to the **Server Farm** item in the IIS Manager. Next, you need to configure certain parameters of the farm.

Step 6. Configure the load balancing using the application requests routing. This procedure defines the application requests routing with an algorithm that evenly distributes incoming requests among the web servers.

In the IIS Manager click on the name of the created farm (**ELMA3**) and go to the **Load Balance** unit on the **Server Farm** panel (Fig. 89).

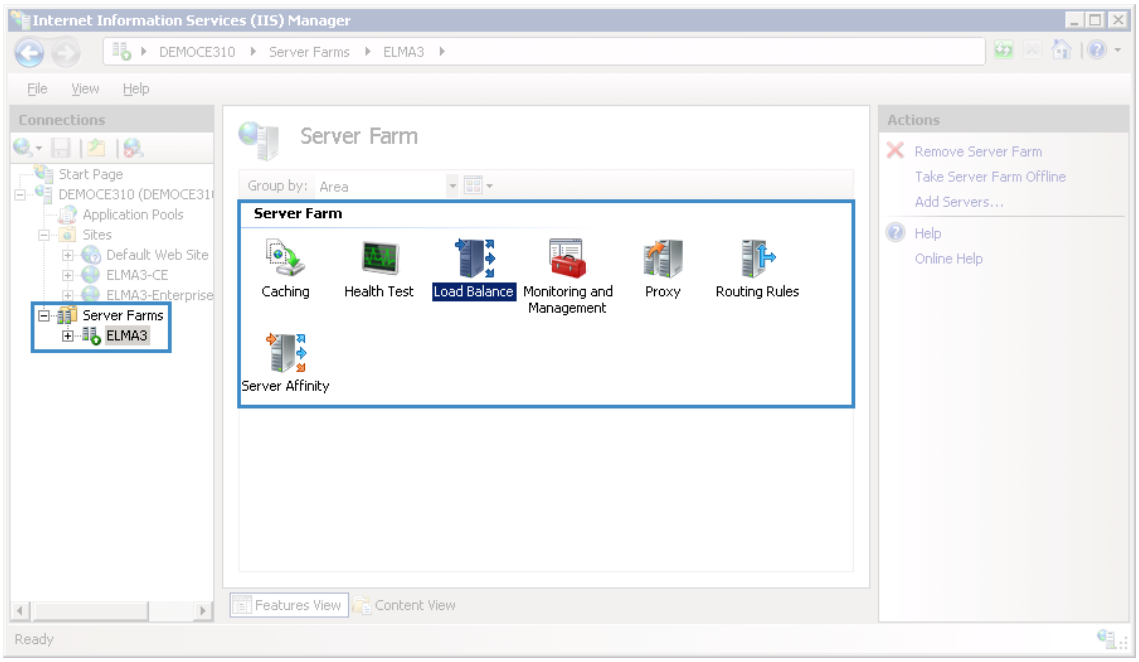

Fig. 89 IIS Manager. Server Farm. Load Balance icon

In the opened section (Fig. 90), select **Weighted round robin** in the **Load balance algorithm** field; select **Even distribution** in the **Load distribution** field. After that, click **Apply**.

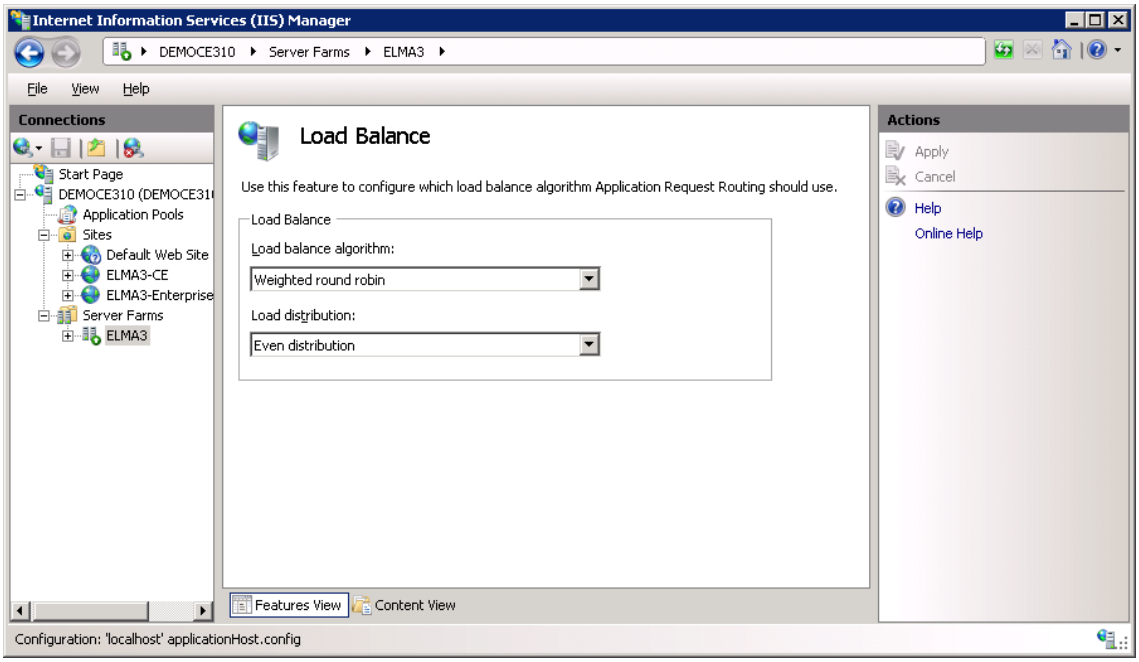

Fig. 90 IIS Manager. "Load Balance" section

Next, open the ELMA3 farm management panel and go to the **Routing Rules** section on the **Server Farm** panel (Fig. 91).

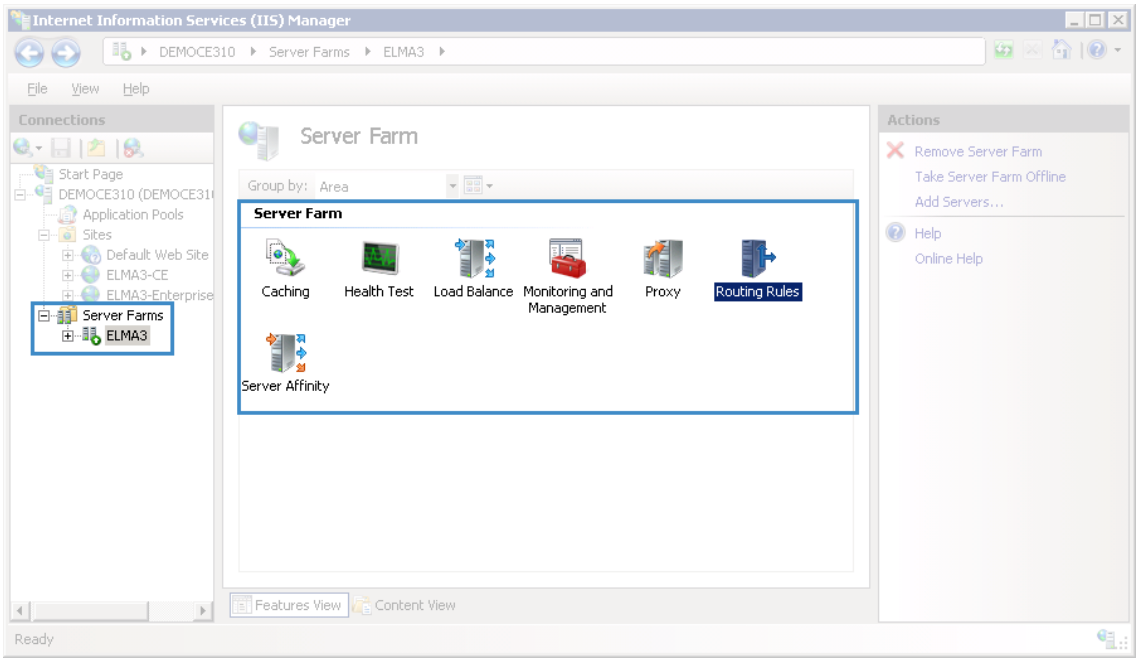

Fig. 91 IIS Manager. Server farm. Routing Rules icon

In the opened section (Fig. 92), check the boxes **Use URL Rewrite to inspect incoming requests** and **Enable SSL offloading**. To save the changes click **Apply**.

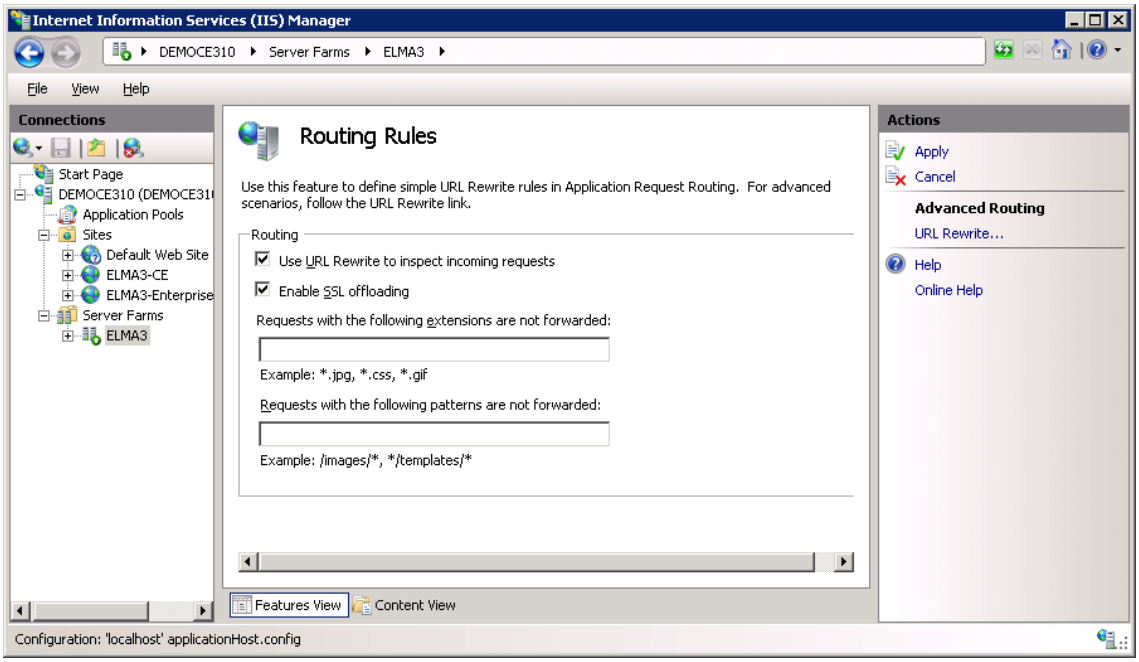

Fig. 92 IIS Manager. "Routing Rules" section

Next, open the **ELMA3** farm management panel and go to the **Health Test** section on the **Server Farm** panel (Fig. 93).

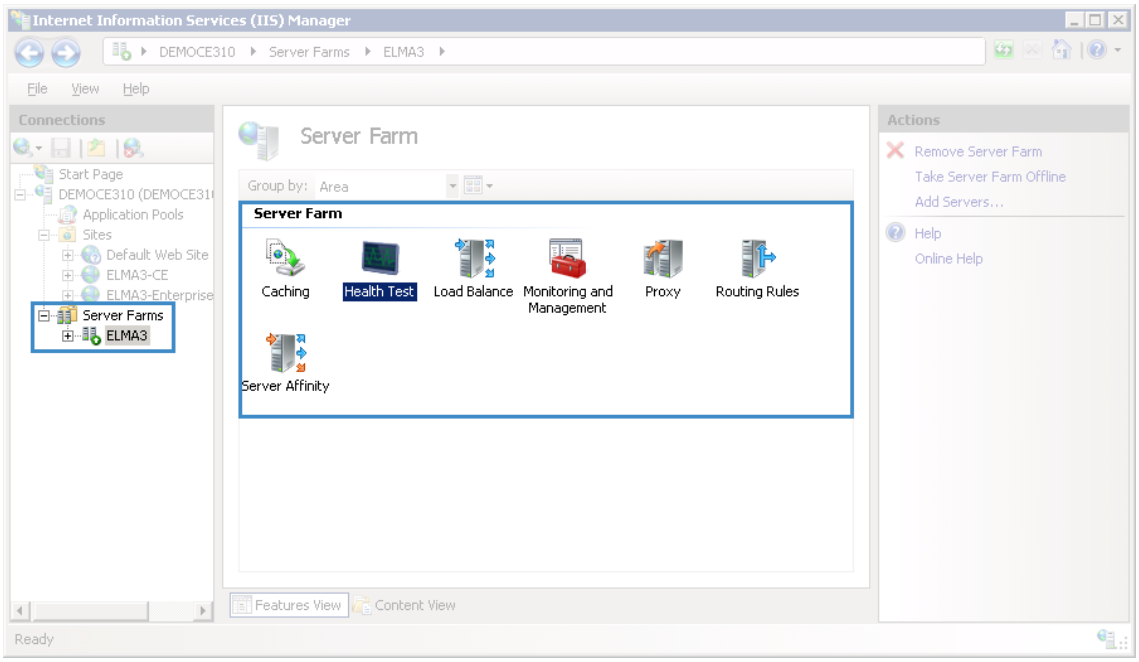

Fig. 93 IIS Manager. Server farm. Health Test icon

In the opened section (Fig. 94), fill in the following fields:

- **URL** – string: http://192.168.19.41/StartInfoHandler.ashx?type=Availability where 192.168.19.41 – the controller server address;
- **Interval (seconds)** – set to 5;
- **Time-out (seconds)** – set to 30;
- **Acceptable status codes** – set to 200.

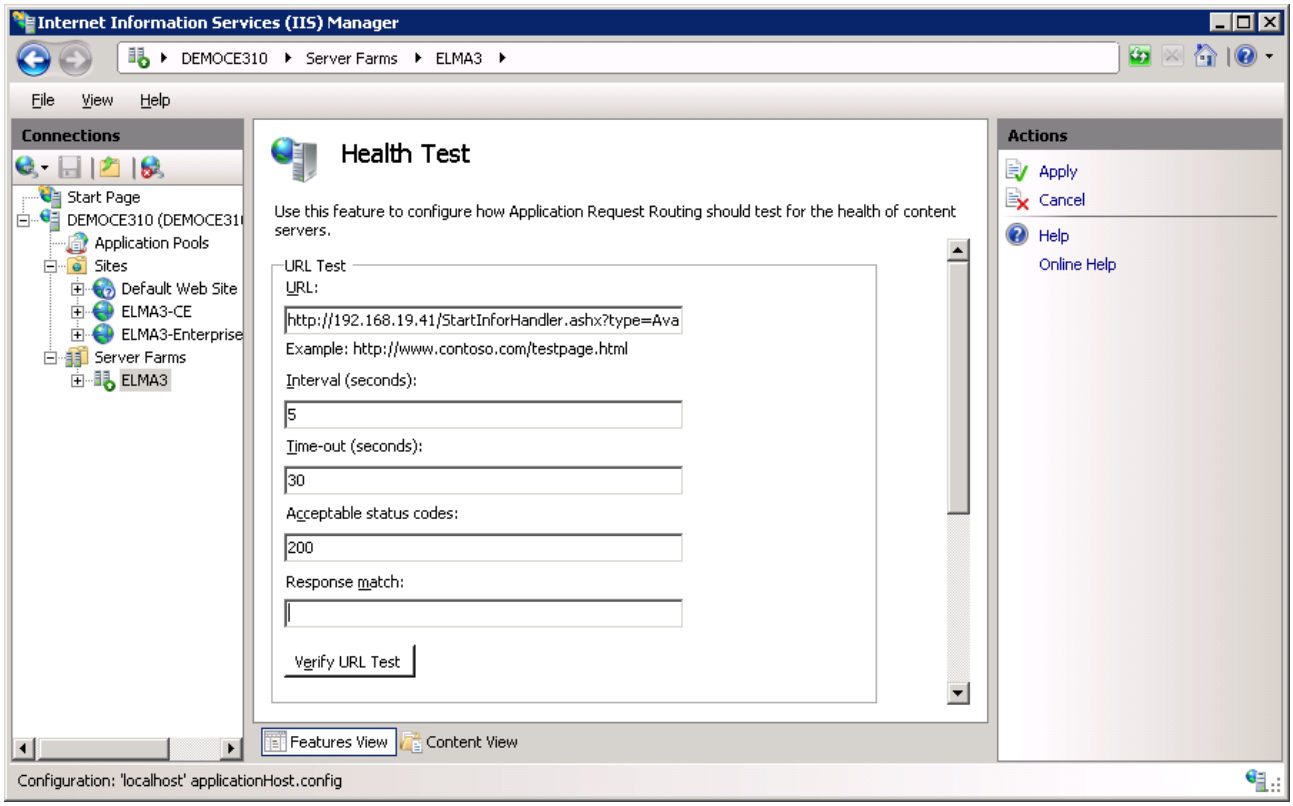

Fig. 94 IIS Manager. Health Test section

To save the changes, click **Apply**.

Open the ELMA3 farm control panel and go to the **Server Affinity** section on the **Server Farm** panel (Fig. 95).

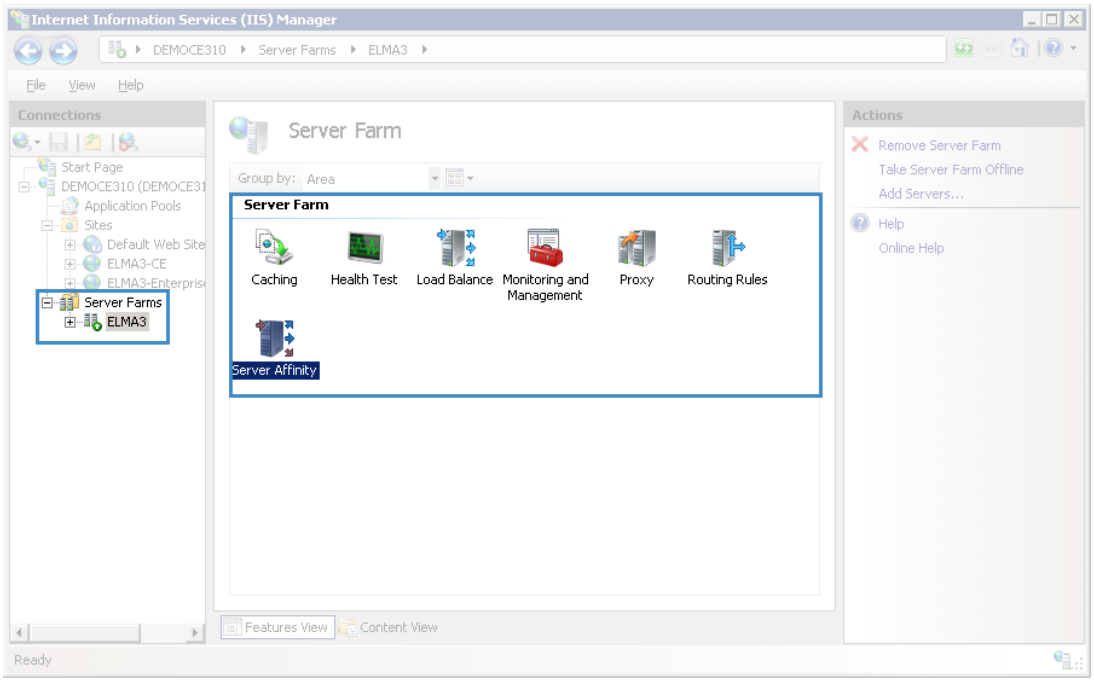

Fig. 95 IIS Manager. Server farm. Server Affinity icon

In the opened section (Fig. 96), uncheck the **Client Affinity** box.

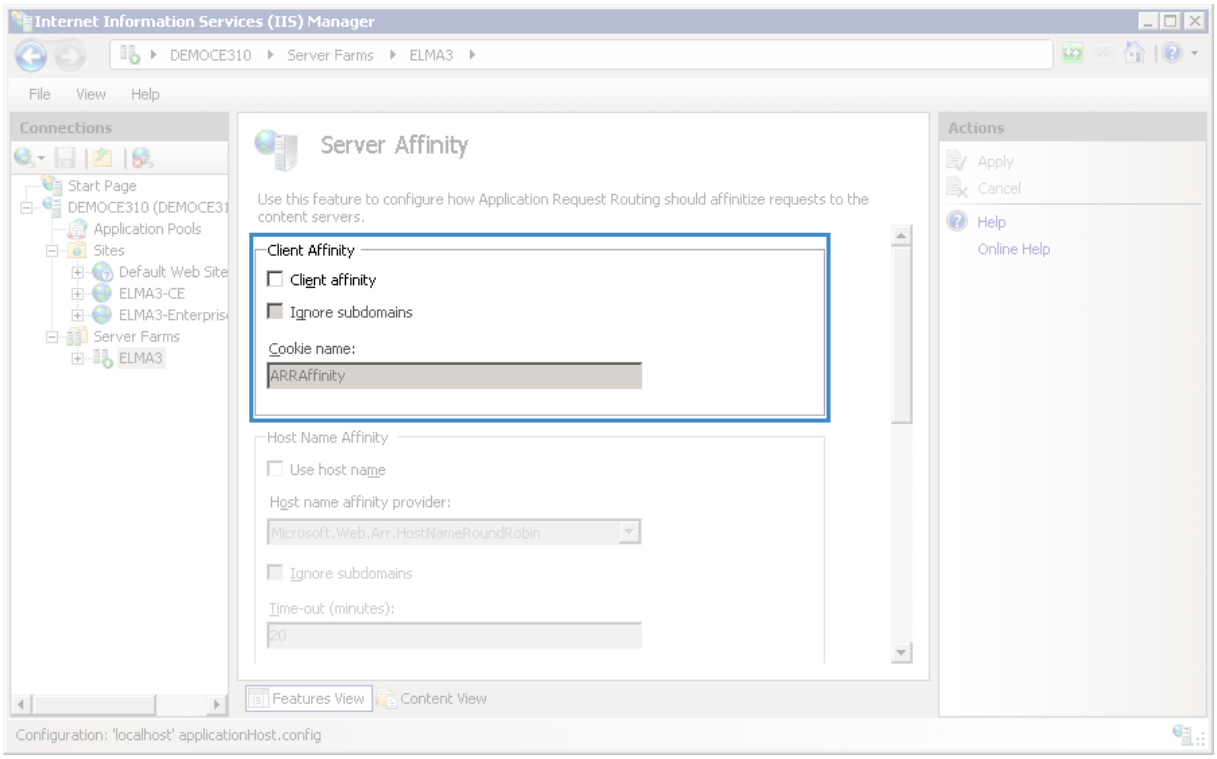

Fig. 96 IIS Manager. Server Affinity section

To save the changes, click **Apply**.

Attention! This parameter is not required, however, we recommend unchecking this box to ensure even load distribution.

If this box is checked, all the requests will be sent to the server, to which the first request from this client was sent; otherwise – to the first unoccupied server.

If the box is checked, the load is distributed less evenly, but the server may respond faster in some cases. Since the increase in the speed is insignificant, it is recommended that you uncheck this box, because there is a risk that most users will be tied to the same server, while all the other servers of the farm will remain idle.

Next, open the ELMA3 farm management panel and go to the **Proxy** section in the **Server Farm** unit (Fig. 97).

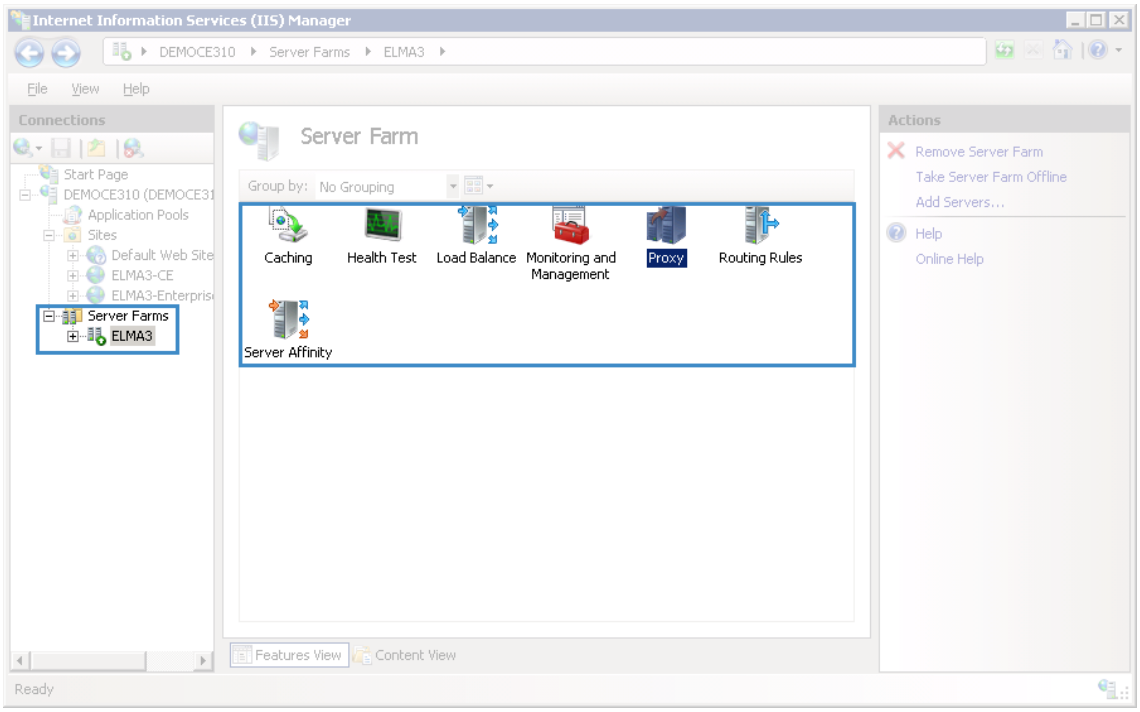

Fig. 97 IIS Manager. Server farm. Proxy icon

In the opened section (Fig. 98), increase the value in the **Time-out (seconds)** field to 600 seconds.

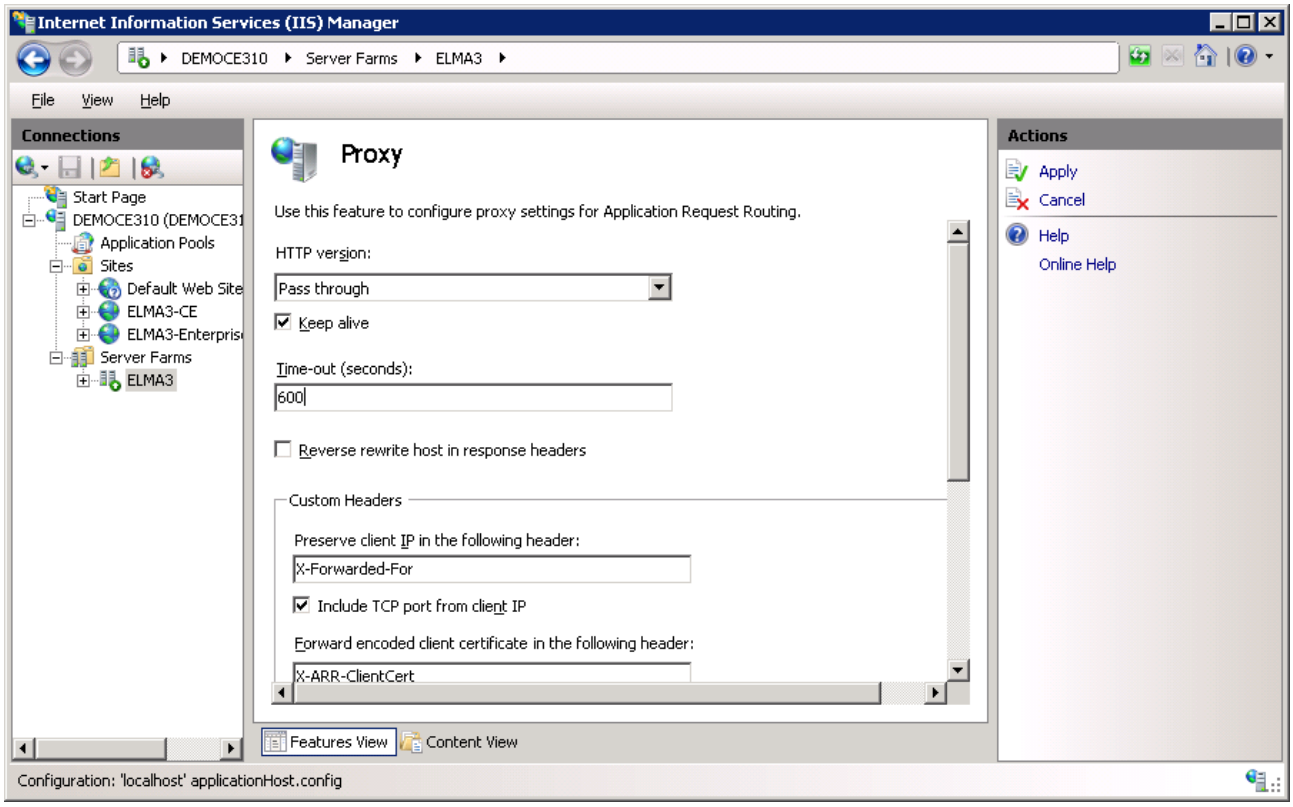

Fig. 98 IIS Manager. Proxy section

To save the changes, click **Apply**.

If this value is exceeded during the page generation, the error 502 will be returned.

Step 7. It is recommended that you disable IIS logging.

If you intend to use the IIS log, you can skip this step; otherwise, it is recommended that you follow the described procedures.

If there are many users, the IIS log may take up to several gigabytes of the controller's space by creating a log file every day. If you do not disable the IIS log, remember to check the controller's free space from time to time.

To disable the log, open the main page of the IIS Manager and select **Logging** in the **IIS** unit (Fig. 99).

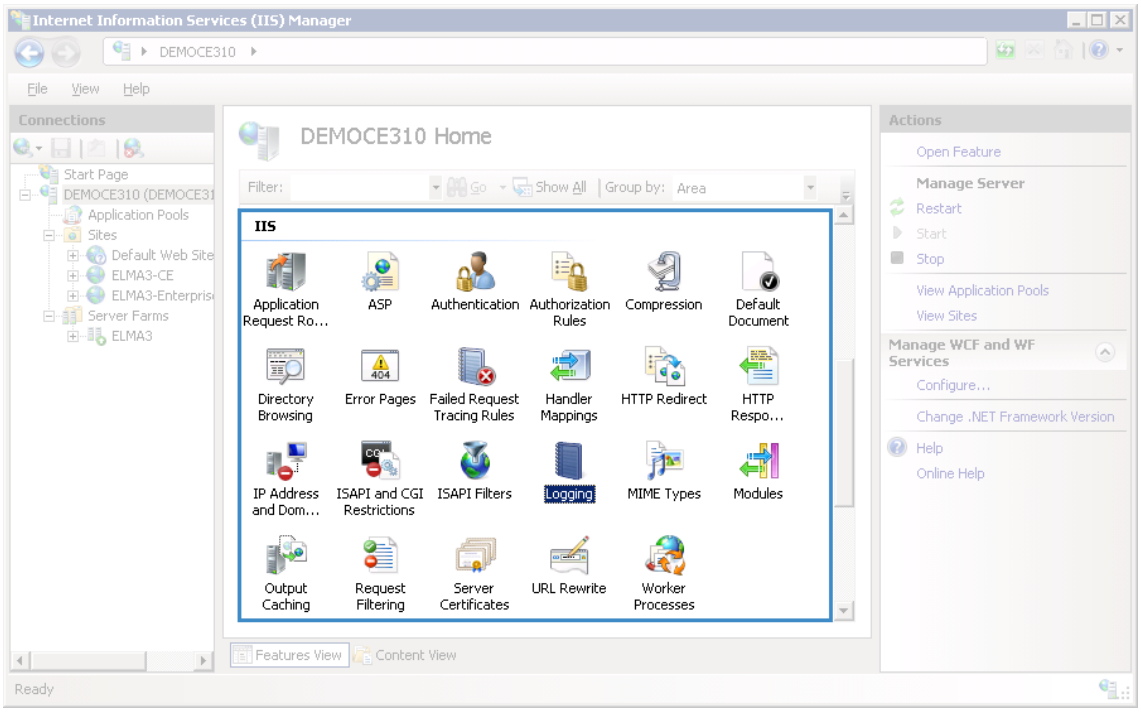

Fig. 99 IIS Manager. Server farm. Logging icon

In the opened section (Fig. 100), select **Disable**.

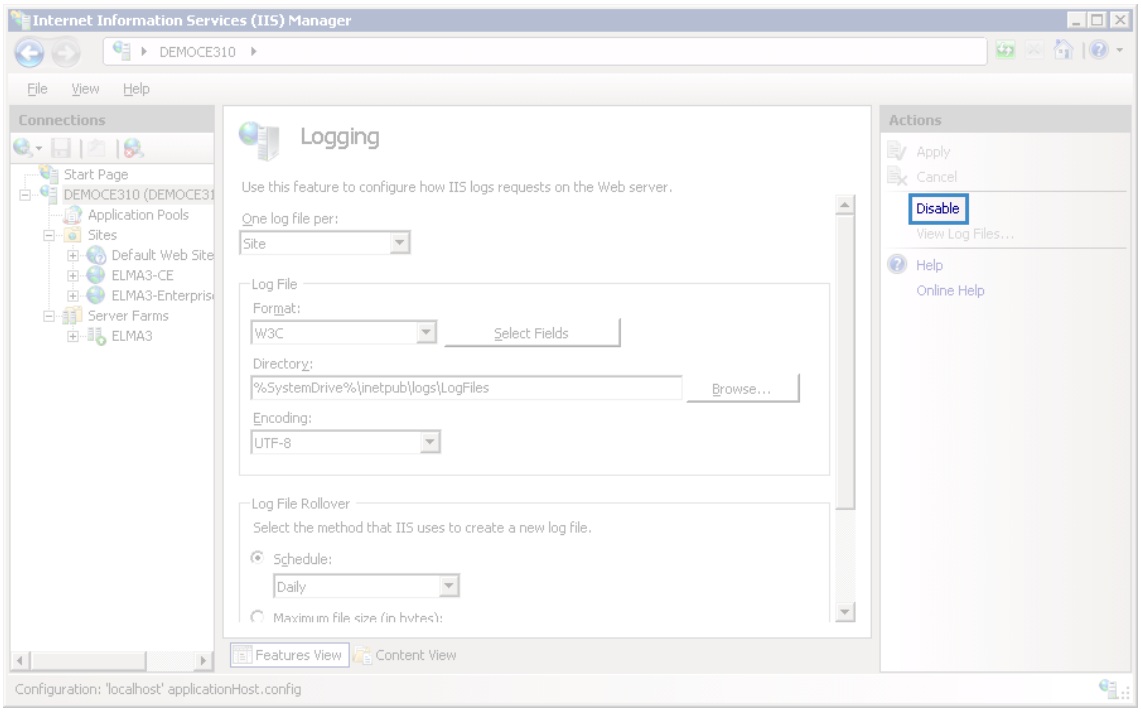

Fig. 100 IIS Manager. Logging section

Chapter 5. Maintenance

This section describes the main procedures for maintaining the system components and the automation of these procedures.

5.1. Database maintenance

Database maintenance usually includes creating backup copies and automating this process to avoid uncontrolled growth of the number of backup copies.

5.1.1. Creating a database backup copy manually

Before making any changes to ELMA, it is very advisable that you manually create a backup copy even if you are confident, that a backup copy has been created automatically.

You should store manually created backup copies for a week, and delete them only if there are newer backup copies.

To create a database backup copy, follow this procedure:

1. Connect to the main database server – 192.168.18.230

If you are using the failover cluster and the MS SQL AlwaysOn availability group, creating backup copies may be allowed only on a particular server, all the others will display a respective warning.

In this case, it is specified in the cluster that backup copies will be created from the main replica – 192.168.18.230.

2. Make sure that there is enough disk space for storing database backup copies.

   If there is not enough disk space:

- delete unnecessary files;
- delete old backup copies if there are newer copies;
- archive old backup copies. In case of archiving, take into account that the backup copies, created with compression, will not be compressed after archiving.

3. Start MS SQL Management Studio and connect to the instance of MS SQL Server (Fig. 101).

Fig. 101 MS SQL Server Management Studio. Authentication dialog box

In the context menu of the database, select **Tasks – Back Up...** (Fig. 102).

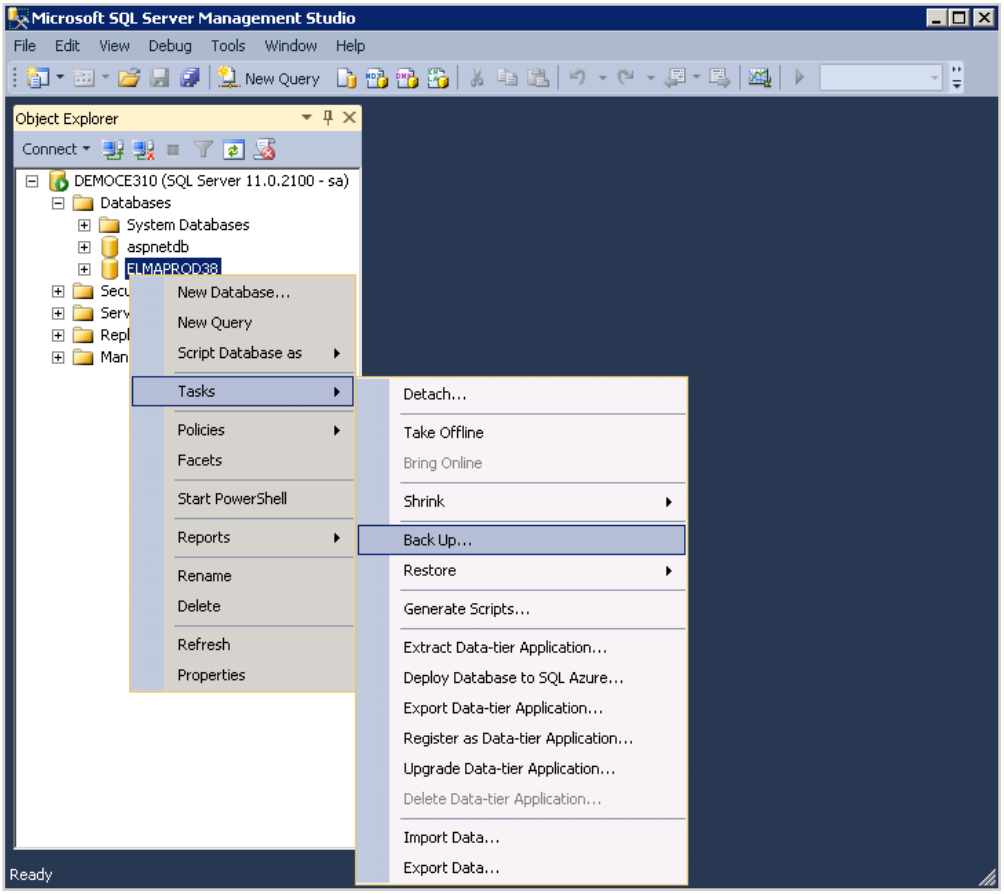

Fig. 102 MS SQL Server Management Studio. Database context menu

In the opened dialog box (Fig. 103), on the **General** page, make sure that you have selected **Full** in the **Backup type** field and specified a backup location

(the shared folder for backup copies, created earlier) in the **Destination** unit. To specify a location, click **Add...**

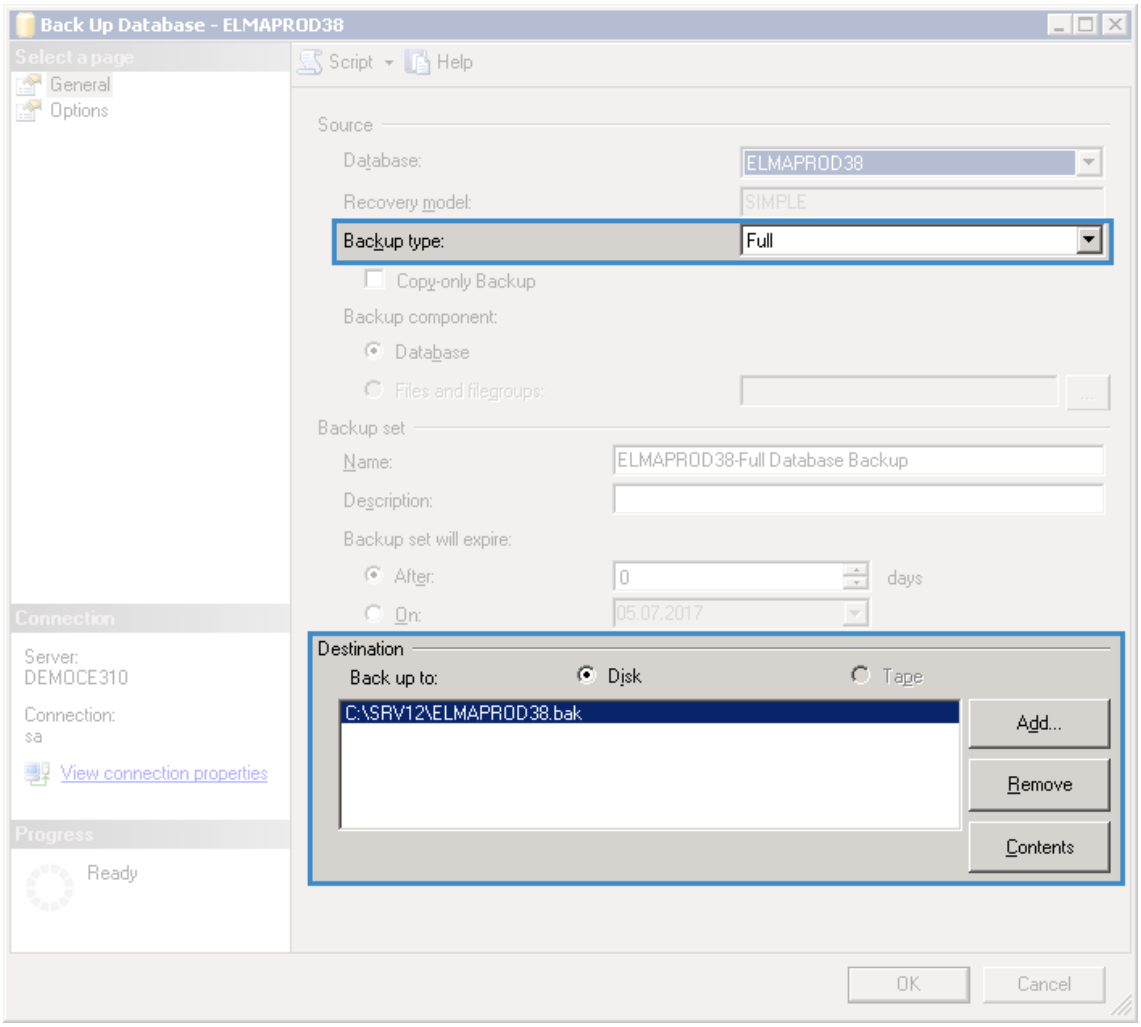

Fig. 103 Back Up Database window. General page

Next, on the **Options** page (Fig. 104) in the **Compression** unit select **Compress backup** from the drop down list.

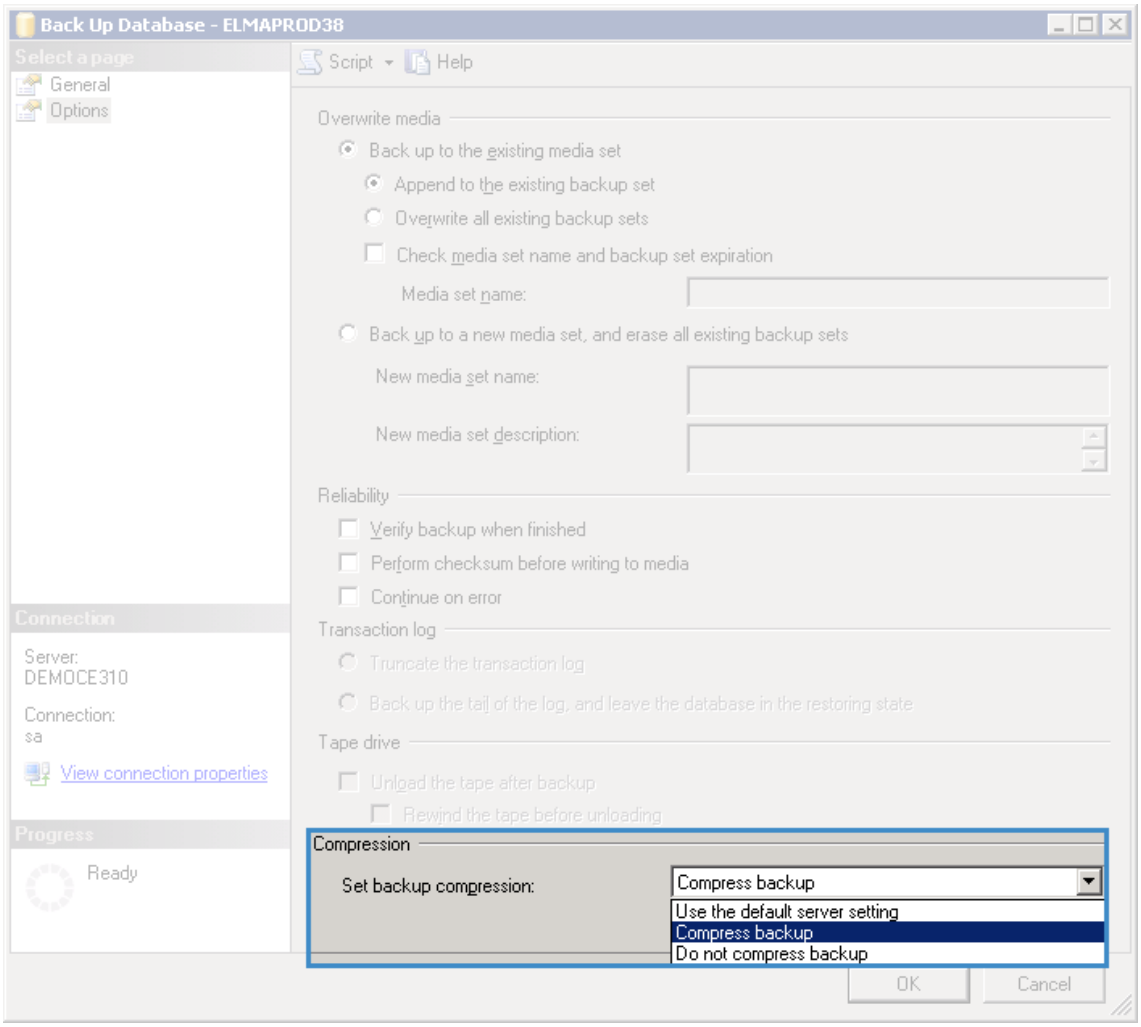

Fig. 104 Back Up Database window. Options page

Experience has shown that using compression reduces the backup file size more than ten-fold.

To save the changes, click **OK**. After that, the progress indicator will be displayed.

If any errors occur at this stage, you should eliminate them according to general recommendations.

Once the backup copy has been created, you will see a respective dialog box.

It is recommended that you automate the backup process. To learn more, see section 5.1.3.

5.1.2. Compressing transaction log manually

Since the Full recovery model is selected in the database settings, the transaction log size will be gradually increasing in operation. Limiting its

maximum size is highly discouraged, since the database will stop executing requests upon reaching the maximum size.

The nature of MS SQL files implies, that the transaction log file may take a significant amount of disk space, while on the inside being 97% reusable. Therefore, backing up the transaction log will free the file space for rewriting, but it will take the same disc space as before these actions.

The correct solution is automatic backup of the transaction log (see section 5.1.3). In this case, the transaction log file will always be around the same size, depending on ELMA load.

You can compress the transaction log file on the hard drive at any moment by creating two backup copies of the transaction file (to learn more, see section 5.1.1). It is important that both backup copies were created with the **Transaction log** type (Fig. 105) and were available simultaneously (you cannot create the first one, delete it, and then create the second).

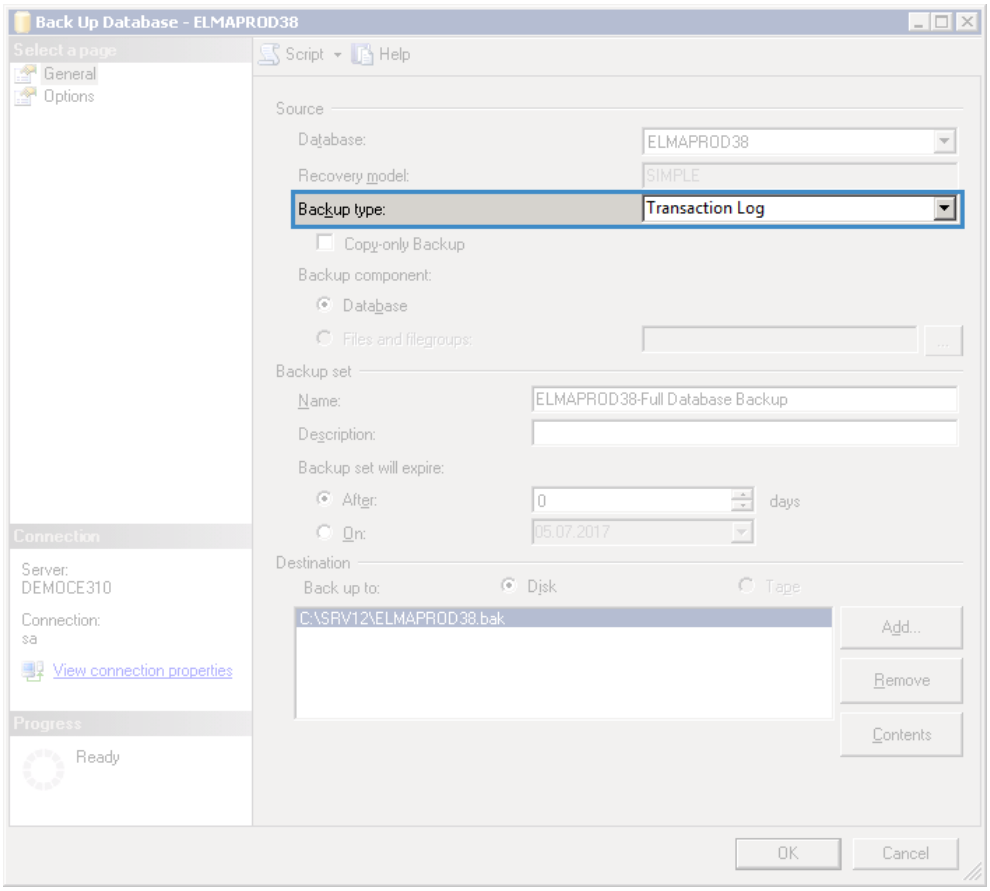

Fig. 105 Back Up Database window. General page

After that, in the database context menu select **Tasks – Srink – Files** (Fig. 106).

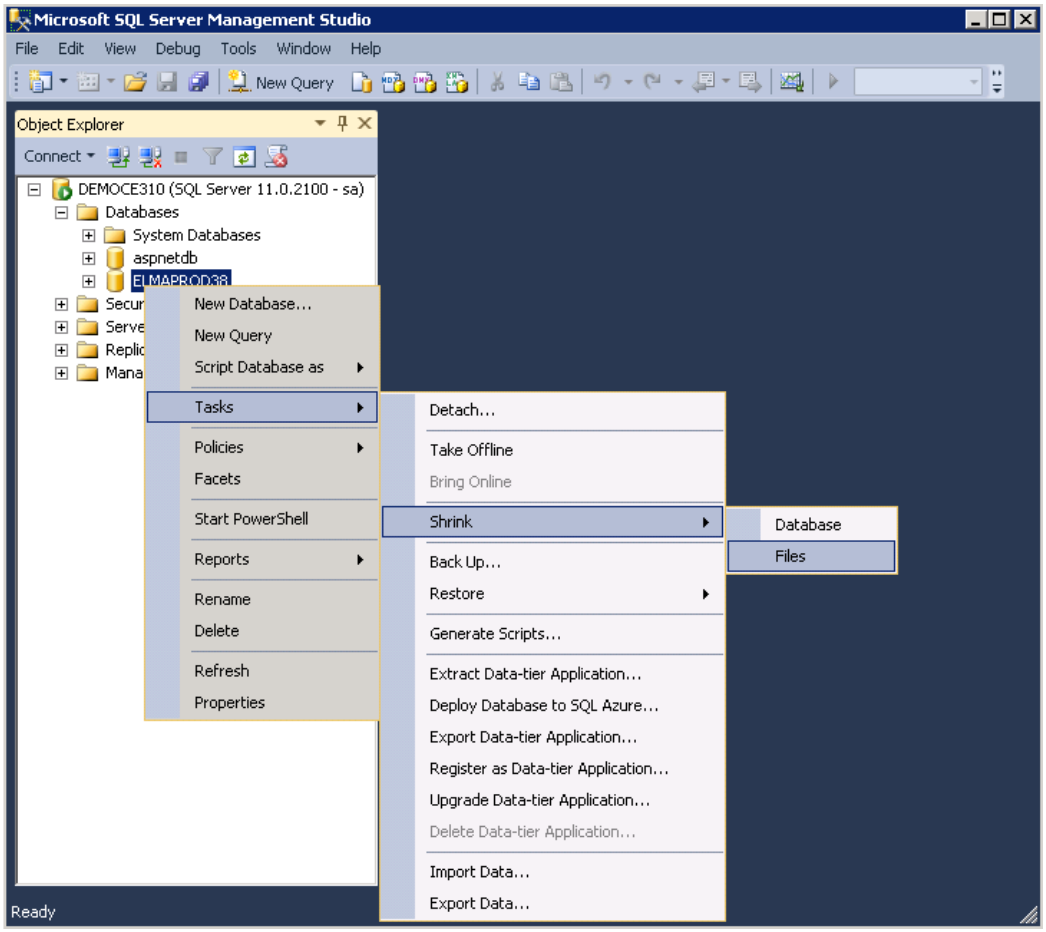

Fig. 106 MS SQL Server Management Studio. Database context menu

In the **File type** field, select **Log** and click **OK** (Fig. 107).

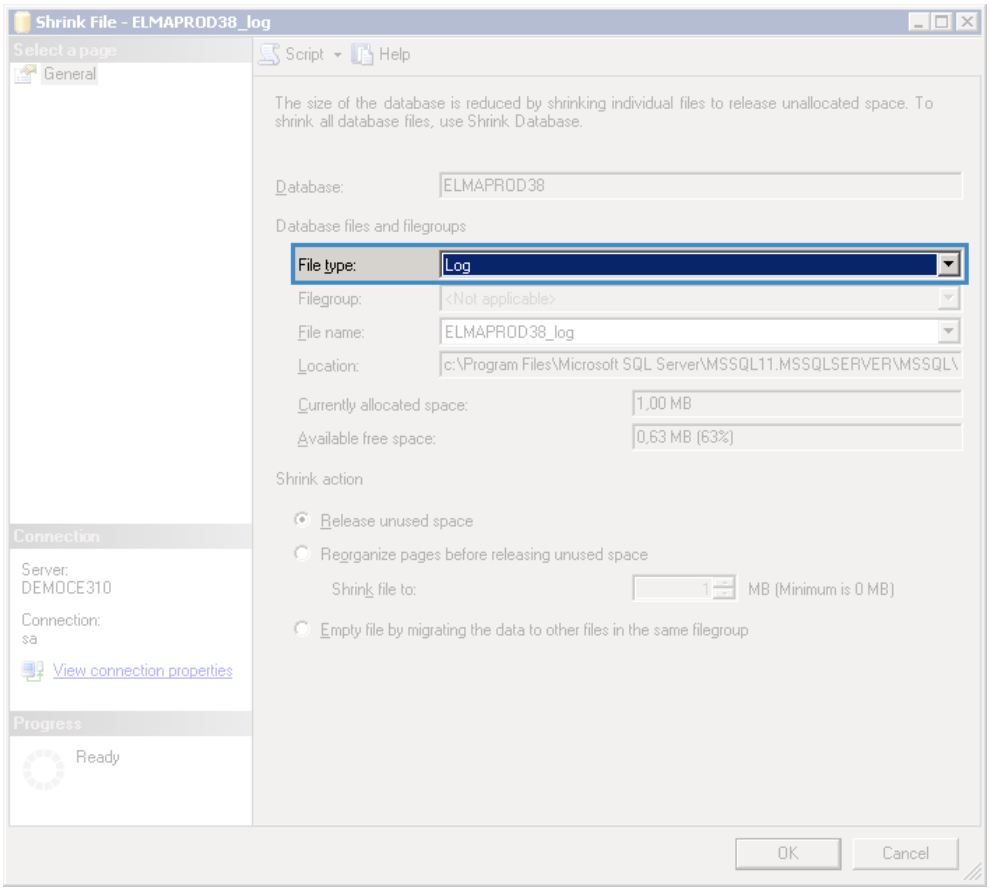

Fig. 107 Shrink File window

Attention! Compressing the data themselves (File type – Data) is highly discouraged.

5.1.3. Creating automatic maintenance plan for ELMA databases

Automatic maintenance plan for databases is required in order to automate certain obligatory operations, such as:

- backing up ELMA database to provide the possibility to recover in case of critical errors while updating the configuration or unexpected server shutdowns and data corruption;
- backing up the transaction log of ELMA database and ASPState database to prevent uncontrolled growth of log-files. This operation ensures, that the log-file is 90-95% reusable and the system automatically uses this space. To learn more read the article "Manage the Size of the Transaction Log File".

We recommend that you simultaneously store from three to seven backup copies for the last days, i.e. you will need to create maintenance plans for each day (Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, Sunday) or at

least three copies for several weekdays (Monday – Thursday, Tuesday – Friday – Sunday, Wednesday – Saturday). Backing up is configured to rewrite the existing copies. i.e. the backup copy of the last Monday will be automatically deleted and a new copy will be created on the next Monday.

Attention! Make sure that the hard drive for backup copies has enough space to spare. If the hard drive runs out of free space, a new backup copy will not be created. It may lead to the growth of transaction log files in the database and to failure to recover after a malfunction. You may also need to manually create backup copies, which will also take up the disk space.

To create a maintenance plan for one weekday, follow the procedure described below.

Step 1. Open MS SQL Management Studio on the primary replica. In the tree, go to **Management – Maintenance Plans** and select **Maintenance Plan Wizard** in the context menu (Fig. 108). The Maintenance Plan Wizard will open (Fig. 109).

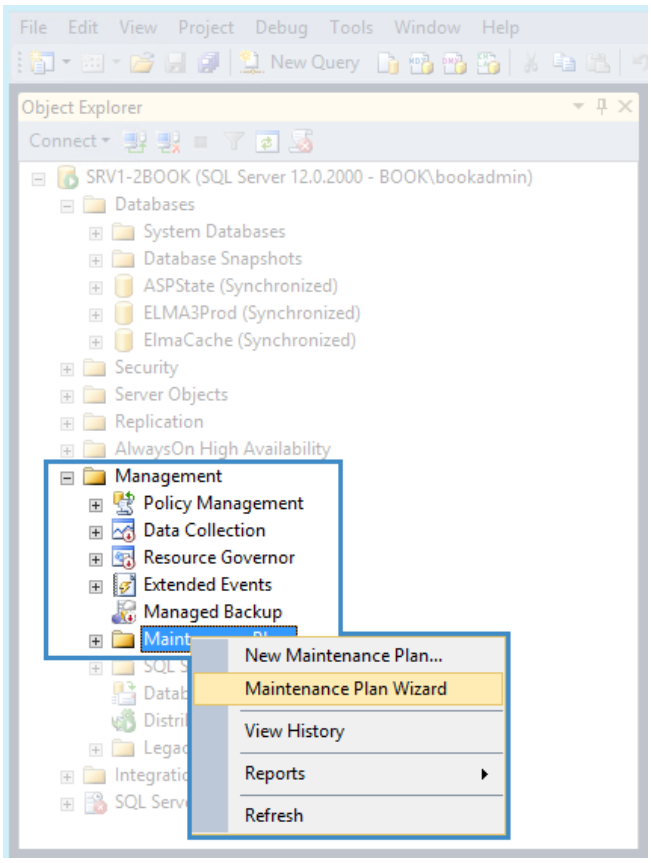

Fig. 108 MS SQL Server Management Studio. Management – Maintenance Plans

Step 2. Create a maintenance plan using the Wizard.

1. Skip the first step of the wizard by clicking **Next**.
2. At the next step (Fig. 109), fill in the required fields and click **Next**.

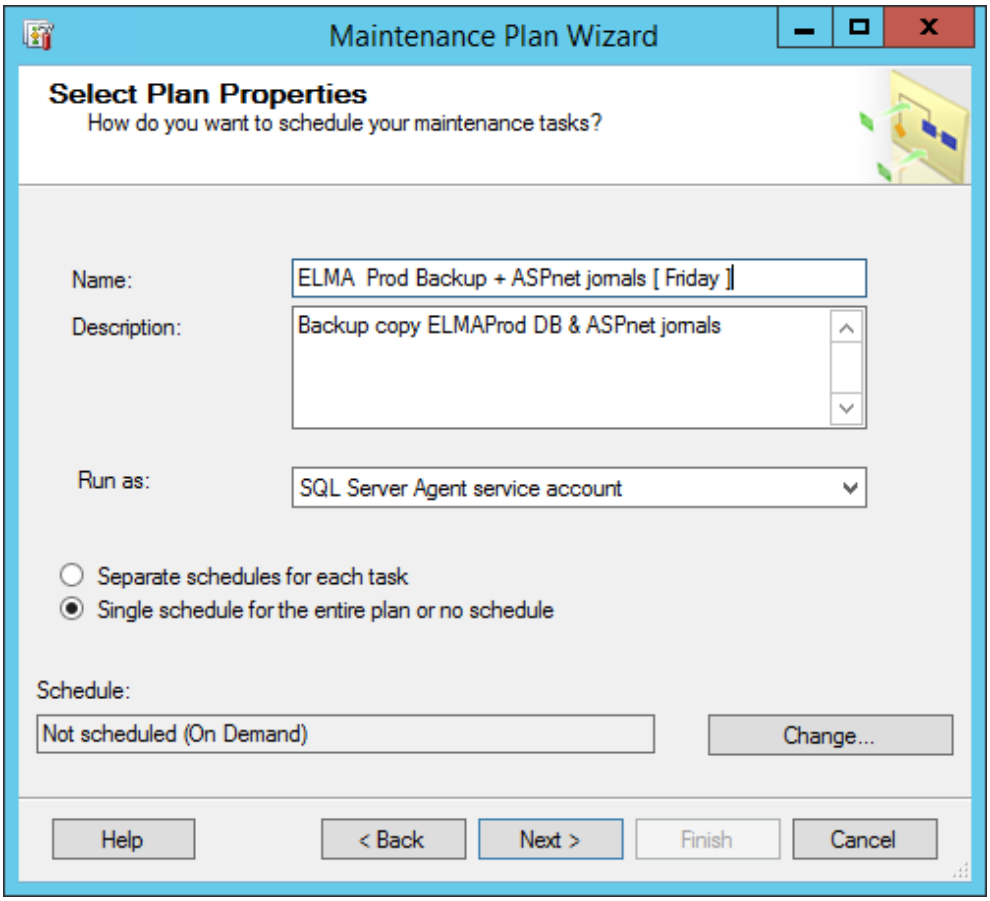

Fig. 109 Maintenance Plans Wizard. Step 2

In the **Name** field, enter the maintenance plan name. It is recommended that you give an understandable name that specifies the weekdays, on which the maintenance will be performed.

In the **Run as** field, select **SQL Server Agent service account**. Below this field, select **Single schedule for the entire plan or no schedule**. This way, when creating backup copies, the inbuilt account will be used. Permissions of this account has been configured earlier, when installing a DBMS. You can specify a different account.

Click **Change...** next to the **Schedule** field. In the opened dialog box (Fig. 110), configure the schedule for this maintenance plan and click **OK** to apply the changes.

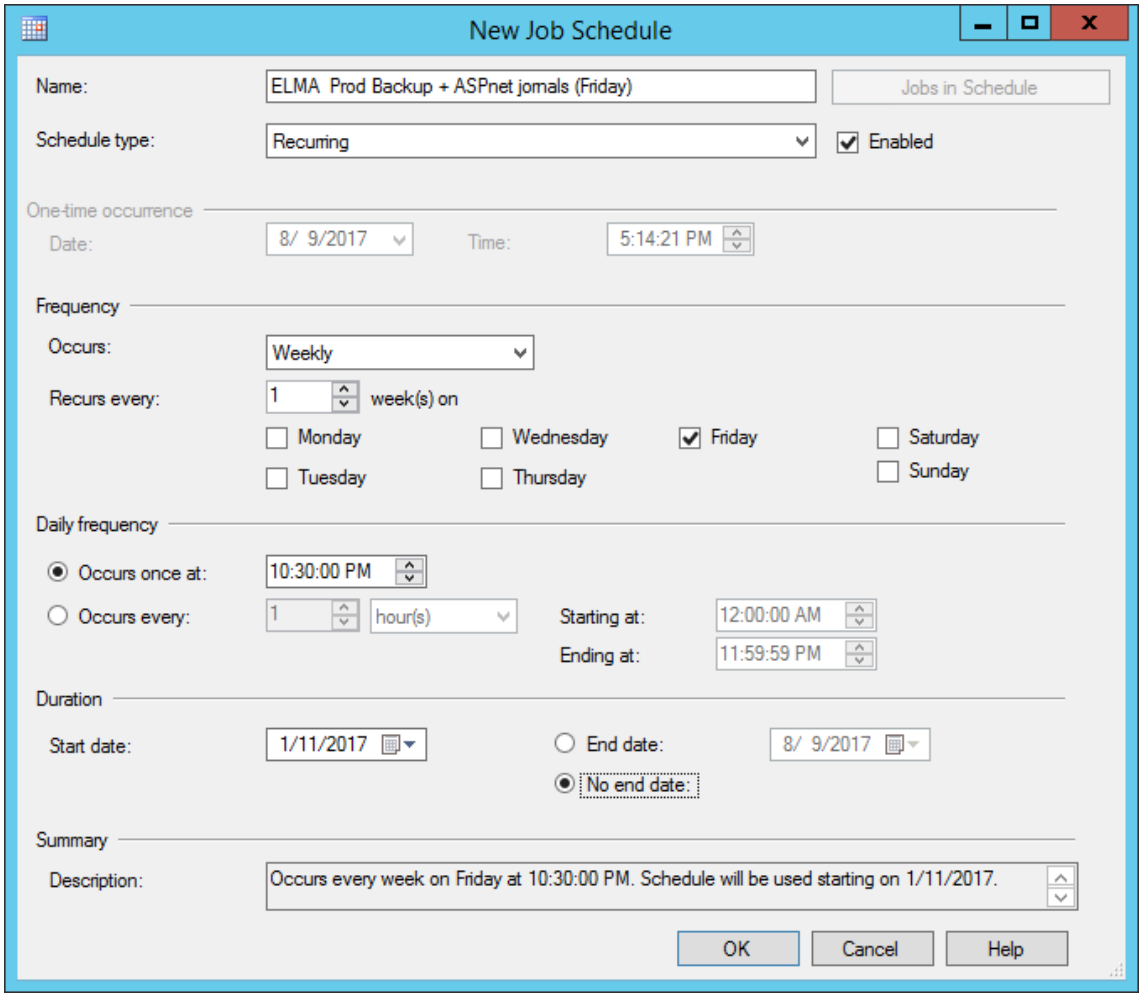

Fig. 110 Creating a schedule

First, specify the frequency and weekdays in the **Frequency** unit. In this example, we intend to create seven similar maintenance plans for each weekday:

- **Occurs** – Weekly;
- **Recurs every** – 1 week(s) on Friday.

Next, define the time when a backup copy will be created. It is recommended that you select time after the working hours with 2-4 spare hours (for the system maintenance). In the **Daily frequency** unit, switch to **Occurs once at** and specify a time. In this case, it is 22:30.

Attention! Do not select a very late time unless it is necessary, since in case of an emergency an employee will have to wait until backing up is complete.

3. At the next step (Fig. 111), select the actions that should be performed as part of maintenance. Check the boxes next to **Back Up Database (Full)** and **Back Up Database (Transaction Log)**.

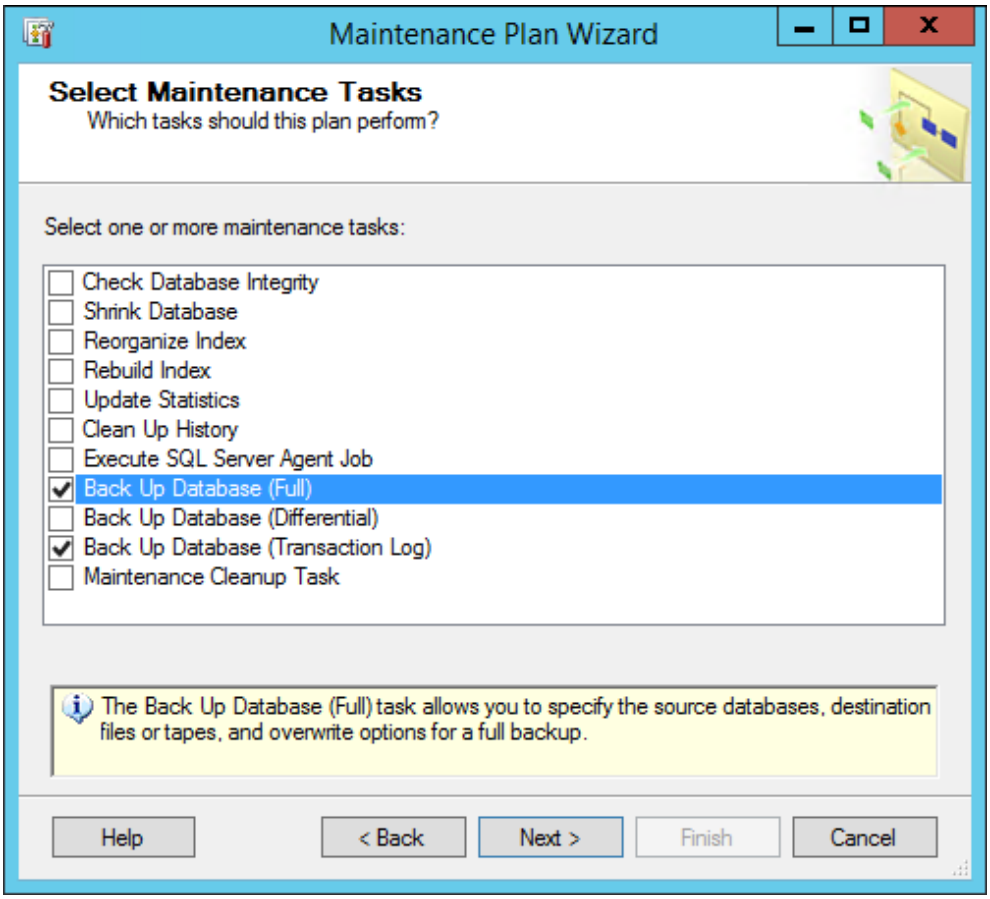

Fig. 111 Maintenance Plan Wizard. Step 3

To go to the next step of the wizard, click **Next**.

4. At the next step (Fig. 112), select the order of the tasks, selected at the previous step (Fig. 111). It is recommended to perform full database backup first and then the transaction log backup. Such order requires more disk space, but it is more reliable.

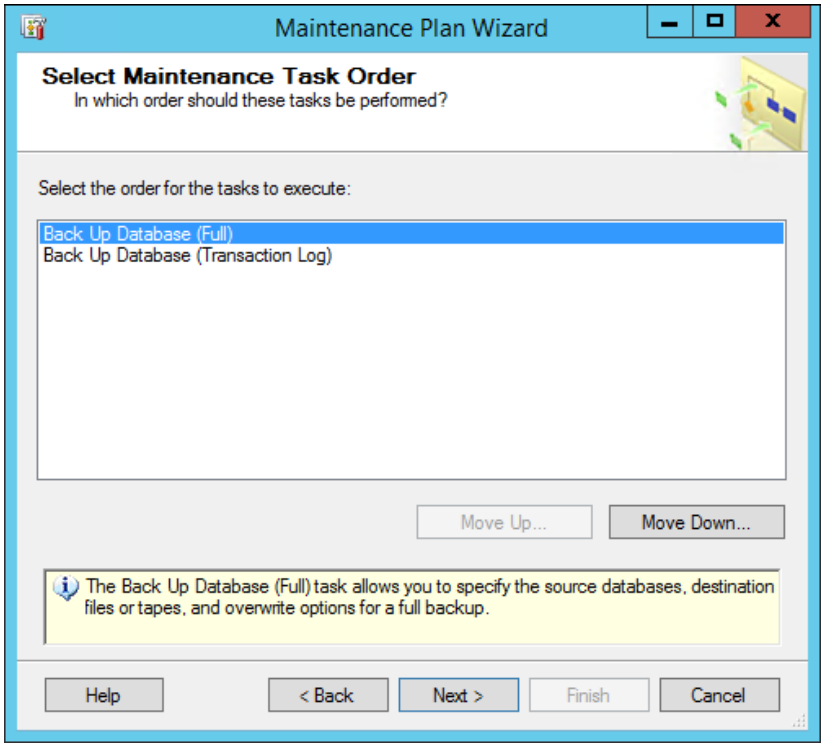

Fig. 112 Maintenance Plan Wizard. Step 4

If the disk space is insufficient, change the maintenance task order: first, back up the transaction log and second, back up the database. With this order, the size of the full back up copy would be smaller.

To go to the next step of the wizard, click **Next**.

5. At this step (Fig. 114), configure full backup of the database.

In the **Database(s)** field (Fig. 113), open the drop down list, select **Specific databases** and select the ELMA database by checking the box next to the required database. In this case, it is **ELMAPROD38**. After that, click **OK**.

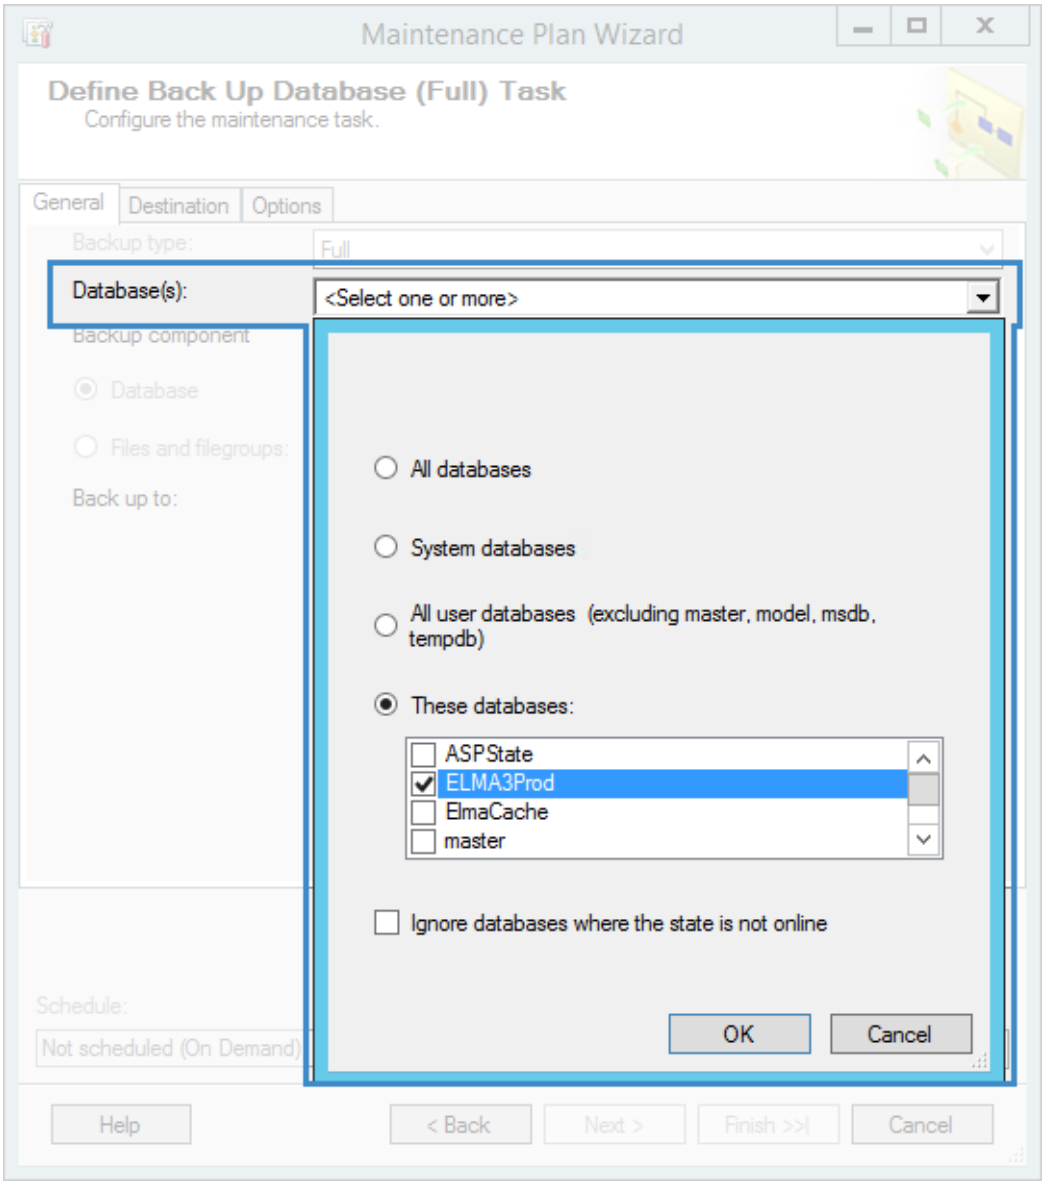

Fig. 113 Maintenance Plan Wizard. Step 5

Next, select **Back up databases across one or more files** (Fig. 114) and specify a location for creating backup copies.

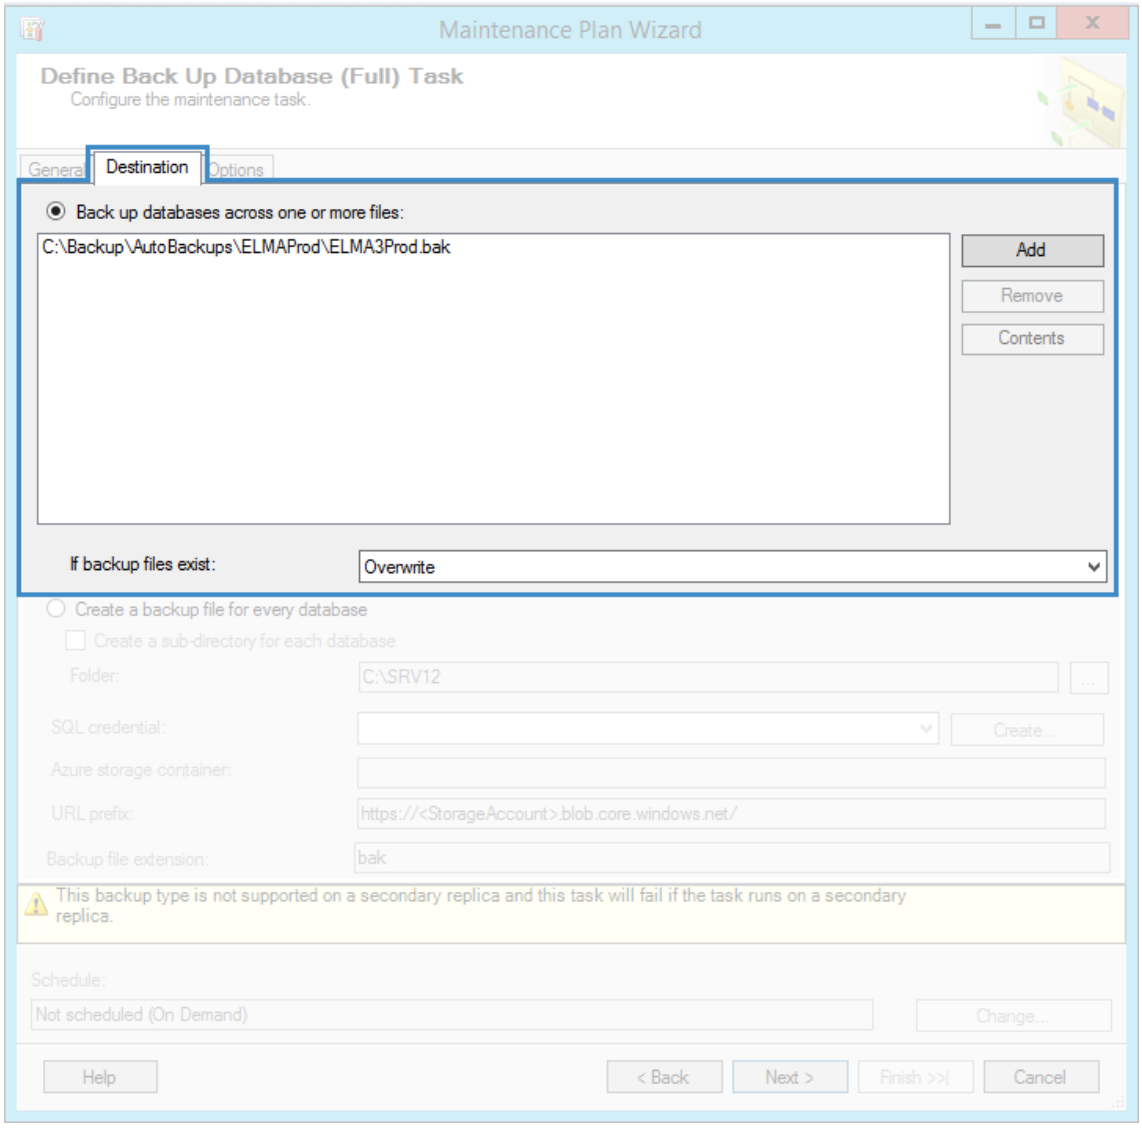

Fig. 114 Maintenance Plan Wizard. Step 5

To specify a location, click **Add...** In the opened dialog box (Fig. 115) select **File name** and specify a database files location. To save the changes, click **OK**.

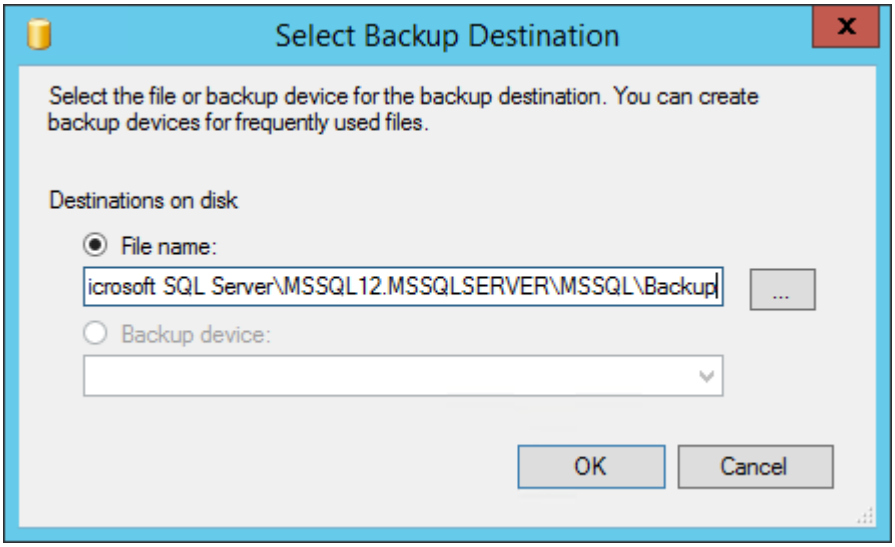

Fig. 115 Selecting a backup location

It is important that you specify a location and a file name without overlapping with other maintenance plans, so that a backup copy for another day was not deleted accidently.

After that, in the **If backup files exist** field, select **Overwrite** (Fig. 114). Thus, this maintenance plan will always back up to the same file.

If you are using a failover cluster, the bottom of the page may display a warning (Fig. 116), that this operation can be performed only on the main replica. If server roles are frequently swapped, you can duplicate all the maintenance plans for the second replica to increase reliability. The maintenance plan has an inbuilt check, so if it is used on a secondary replica, the action is not performed.

Usually, you don’t need to do that, since in case of often failures of replicas it is better to find the reason behind these issues, instead of adapting to them.

Next, in the **Set backup compression** field, select **Compress backup** (Fig. 116), so that database backup copies were compressed on the hard drive. Usually, compression allows reducing the database backup file size up to ten-fold.

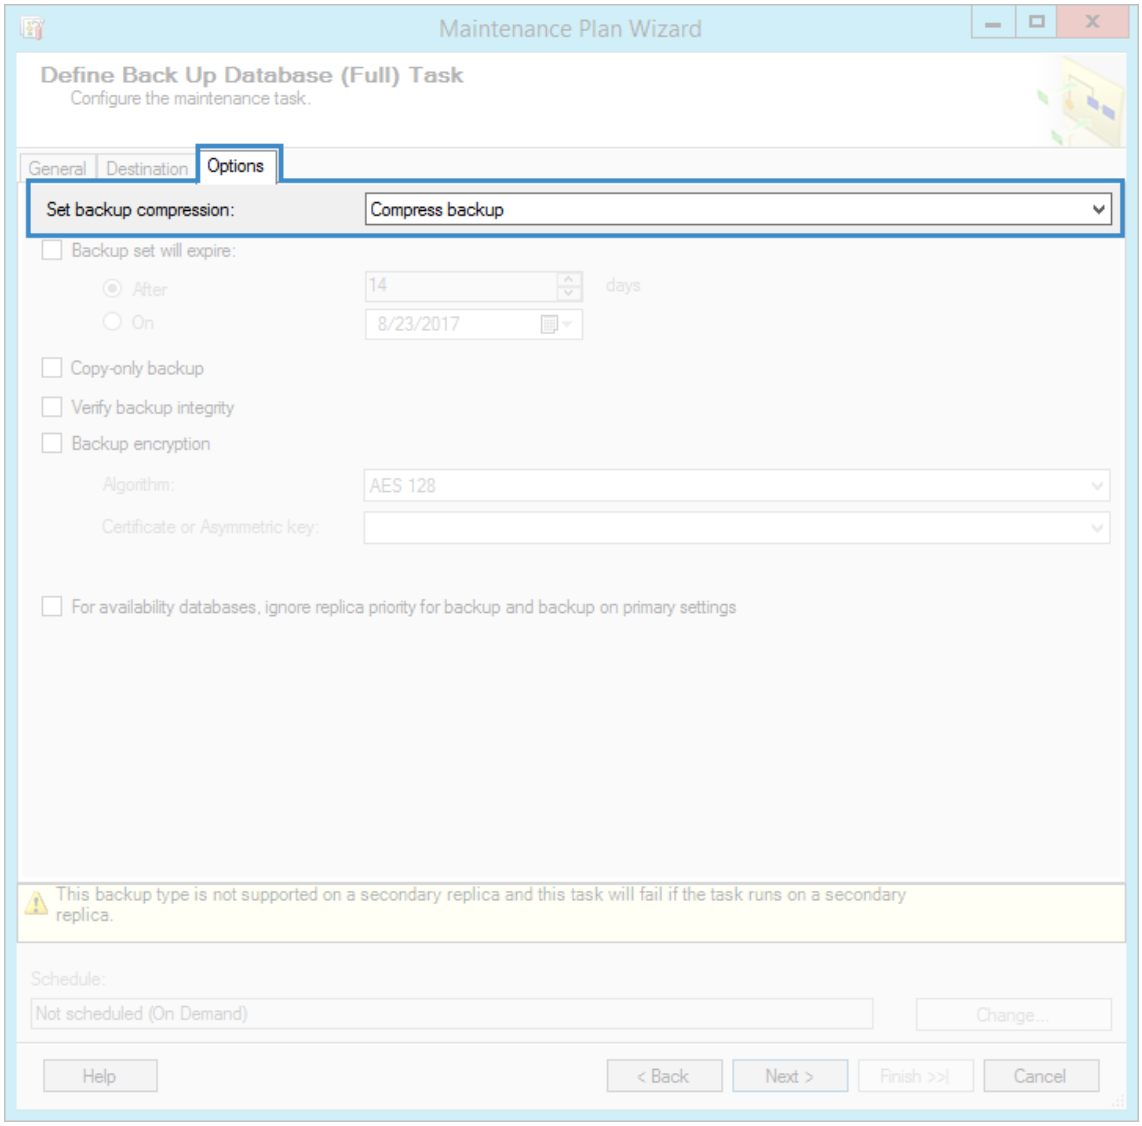

Fig. 116 Maintenance Plan Wizard. Step 5

To go to the next step of the wizard, click **Next**.

6. At this step (Fig. 117), configure transaction log backup. This step is similar to the full database backup (see step 5 above).

However, in the **Database(s)** field, you need to select **Specific databases** and select two databases: **ASPState** and **ELMAPROD38** (Fig. 117).

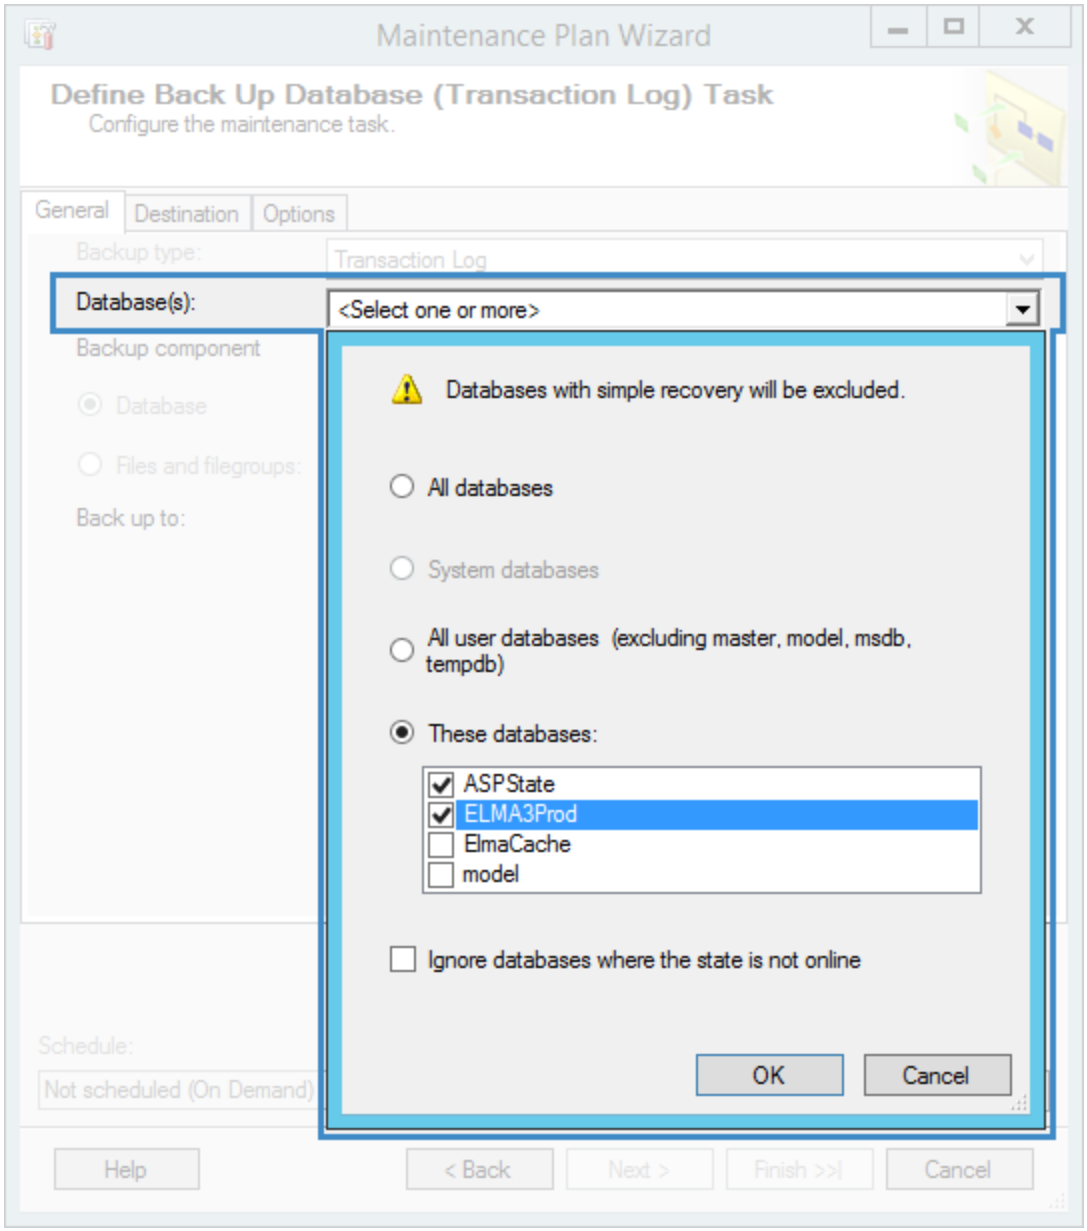

Fig. 117 Maintenance Plan Wizard. Step 6

The ASPState database's size may be around 1 MB, but since it is included in the availability group, it requires the full recovery mode; therefore, its transaction log will grow. To avoid it, back up the transaction log of this database as well as the ELMA database.

Next, select **Back up databases across one or more files** (Fig. 118) and specify a location for saving backup copies. You must specify a file, different from the one above, so that the transaction log backup did not overwrite the data backup.

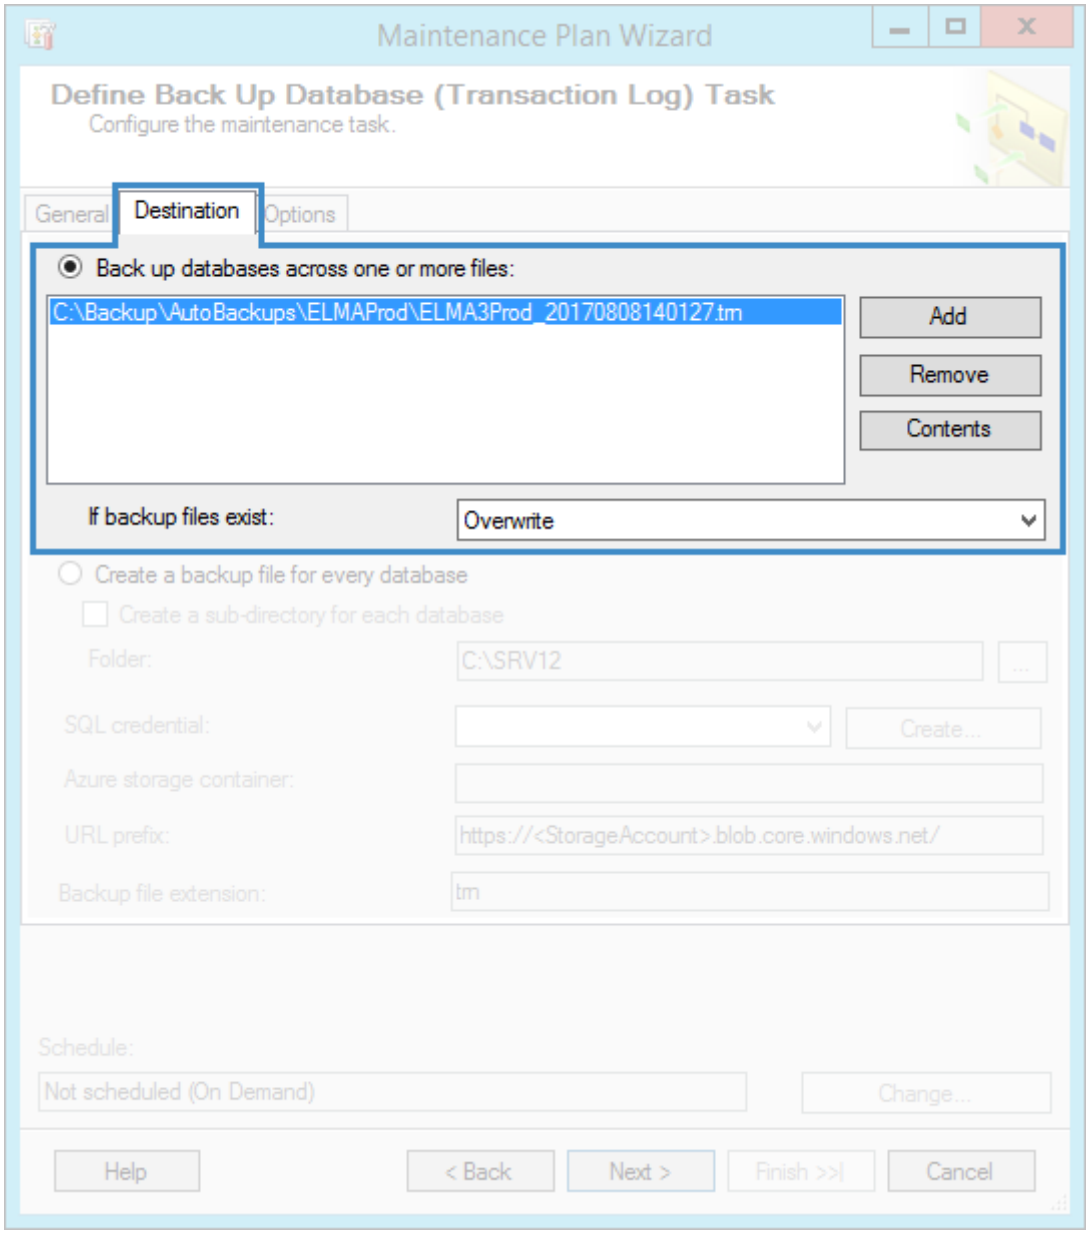

Fig. 118 Maintenance Plan Wizard. Step 6

In the **Set backup compression** field (Fig. 116) select **Compress backups**.

Attention! If a backup file has already been created without compression (e.g. for the last Monday), it cannot be overwritten with a compressed backup of the transaction log.

When you change the **Set backup compression** parameter, you also must move or rename the previous backup copy version (e.g. rename "ProdBackup_Monday.bck" into "ProdBackup_Monday_old.bck").

Unfortunately, it is not always possible to compress the transaction log backup. If possible, it will be compressed after the data backup, and if impossible, an error may be returned.

To go to the next step of the wizard, click **Next**.

7. At this step (Fig. 119), configure the report options. It is recommended that you save a text log of the maintenance plan execution, so that you could see possible errors. To do so, check the box **Write a report to a text file**.

   In the **Folder location** field, you can specify a folder for saving the file.

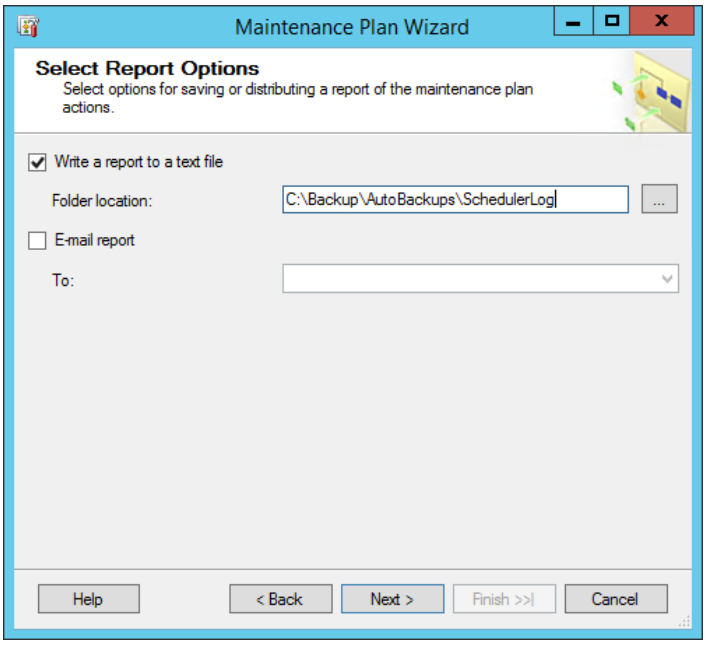

Fig. 119 Maintenance Plan Wizard. Step 7

To go to the next step of the wizard, click **Next**.

8. At the last step (Fig. 120), confirm all the selected options and click **Finish**.

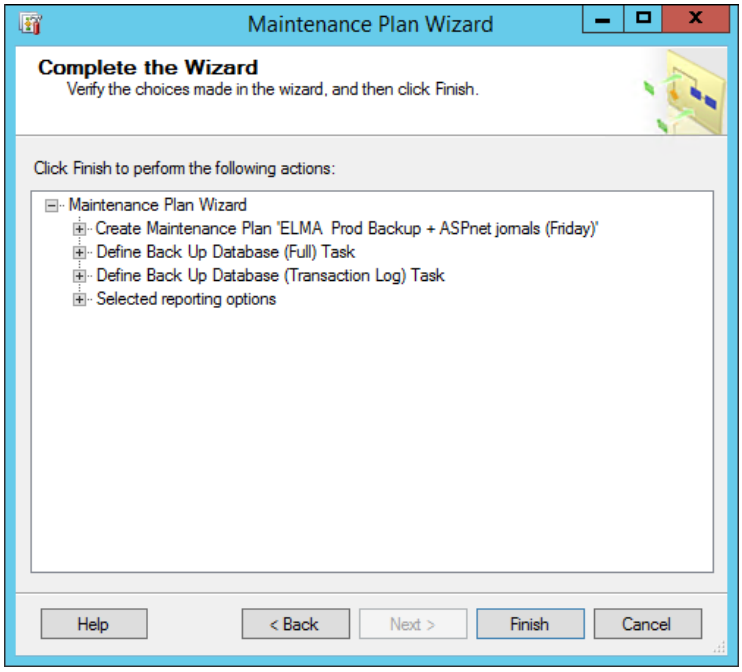

Fig. 120 Maintenance Plan Wizard. Step 8

Wait until the maintenance plan is created, it may take a while.

Once this procedure is finished, the wizard will display the respective notification (Fig. 121), and you will be able to view a report on the results.

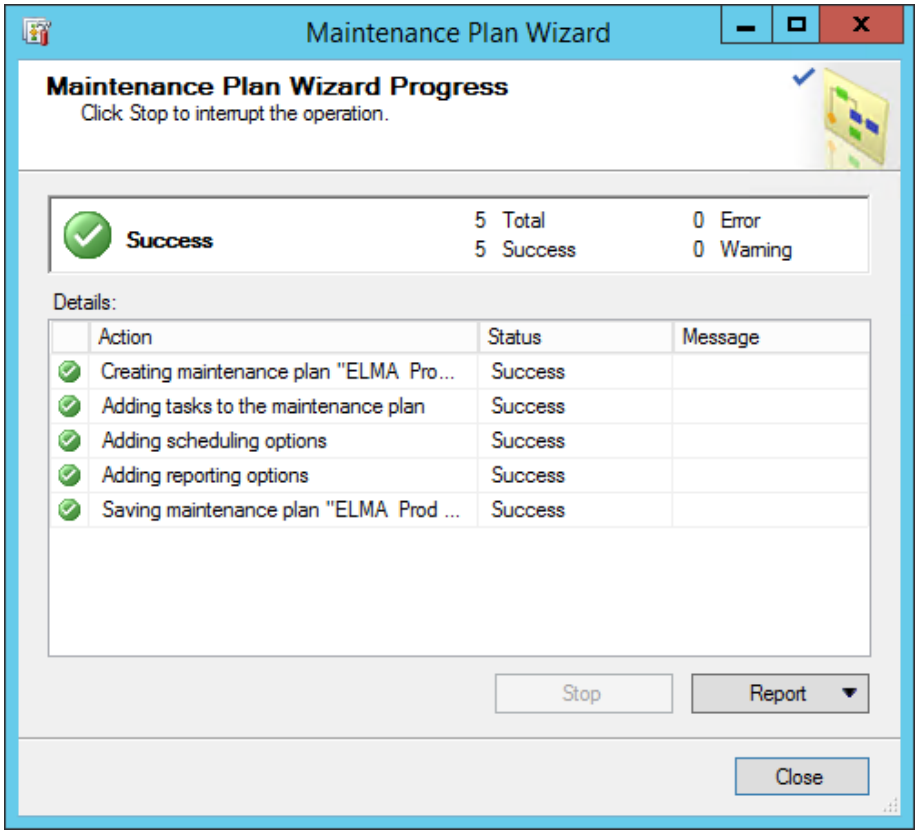

Fig. 121 Maintenance Plan Wizard. Last step

If necessary, create maintenance plans for the other weekdays in a similar way.

5.1.4. Disabling automatic backup upon updating

Every time the structure of the ELMA database is updated, the database backup copy is created. Usually, disabling this procedure is inadvisable; however, when you use a database cluster, it is recommended that you back up according to a schedule and manually when making changes to the configuration.

It is connected, first, to certain peculiarities of creating backups in a cluster, and second, to managing empty disk space.

Since full recovery model is obligatory for database clusters, the size of backup copies may be significant by default. Check if there is enough disk space.

Attention! This step is not obligatory. When disabling automatic back up upon restarting, you must set up a respective maintenance plan and/or create backup copies manually.

To disable database back up upon starting the ELMA server, edit the configuration file **configuration.config** – in **connectionStrings**, in the **main** line add backupEnabled=”false” after EleWise.ELMA.Extensions.MSSQL (after the closing quotation mark).

5.1.5. Recovering a database from a backup copy / moving a database

To recover a database in an availability cluster from a backup copy, you need to temporary disable databases in the availability group. In addition, we recommend that you delete the database from the secondary replica.

Attention! When planning to recover a database from a backup copy, take into account that during this procedure the database will be unavailable as well as all the applications using it. The recovery time depends on the database size. When you estimate the time the procedure takes, consider the time for secondary data synchronization between replicas.

Recovering a database of about 800 GB on high capacity servers with a good connection channel may take around five hours.

It is recommended that you schedule database recovery for Friday evening (after users completed their work in the system). This way you will have maximum amount of time for all the necessary procedures.

1. To start recovering a database from a backup copy, connect to all the DBMS replicas using MS SQL Management Studio.
2. Make sure that the backup copy is available on the main replica (Fig. 122). If you use network drives, having high data transfer speed between servers is advisable, since it directly affects the duration of the restoration procedure.

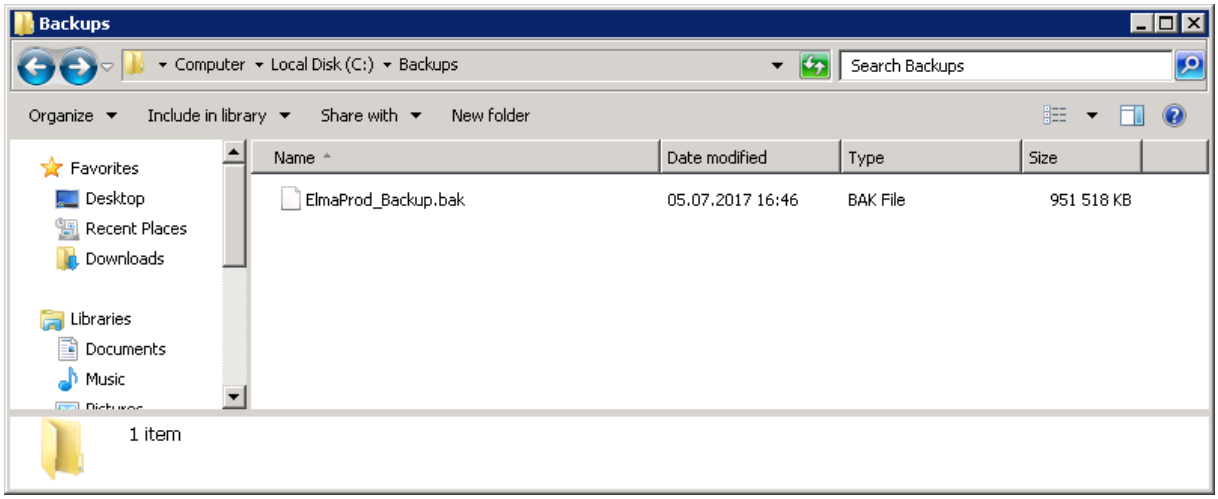

Fig. 122 Main replica. Backup copy

3. Delete the database you intend to recover from the availability group. To do so, in MS SQL Management Studio open the **AlwaysOn High Availability**, **Availability Groups – Availability Databases**, right click on the required database and select **Remove Database from Availability Group**... (Fig. 123).

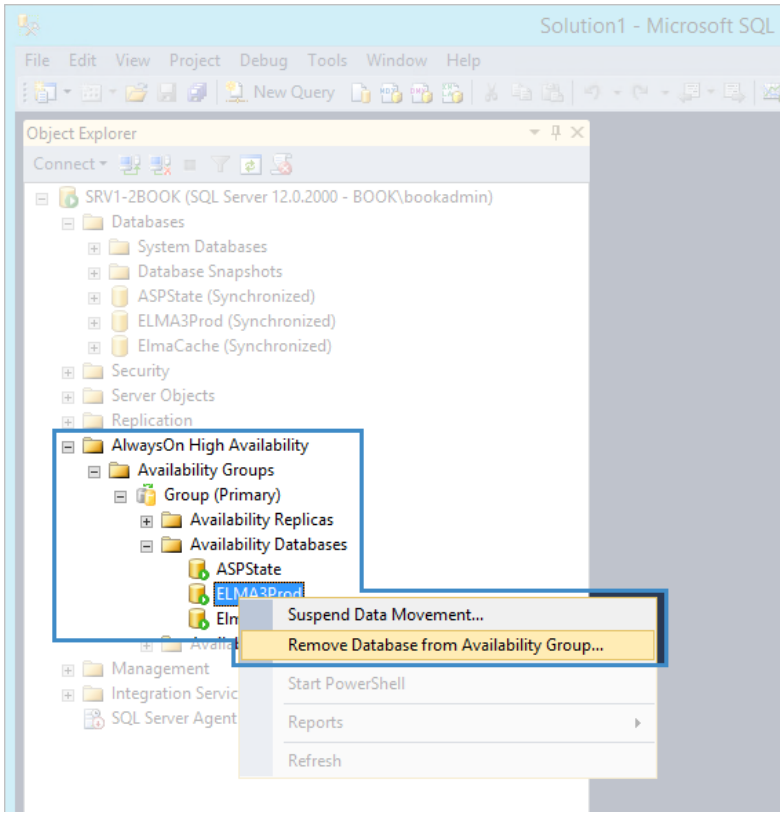

Fig. 123 MS SQL Management Studio. Database context menu

In the opened dialog box (Fig. 124), click **OK**.

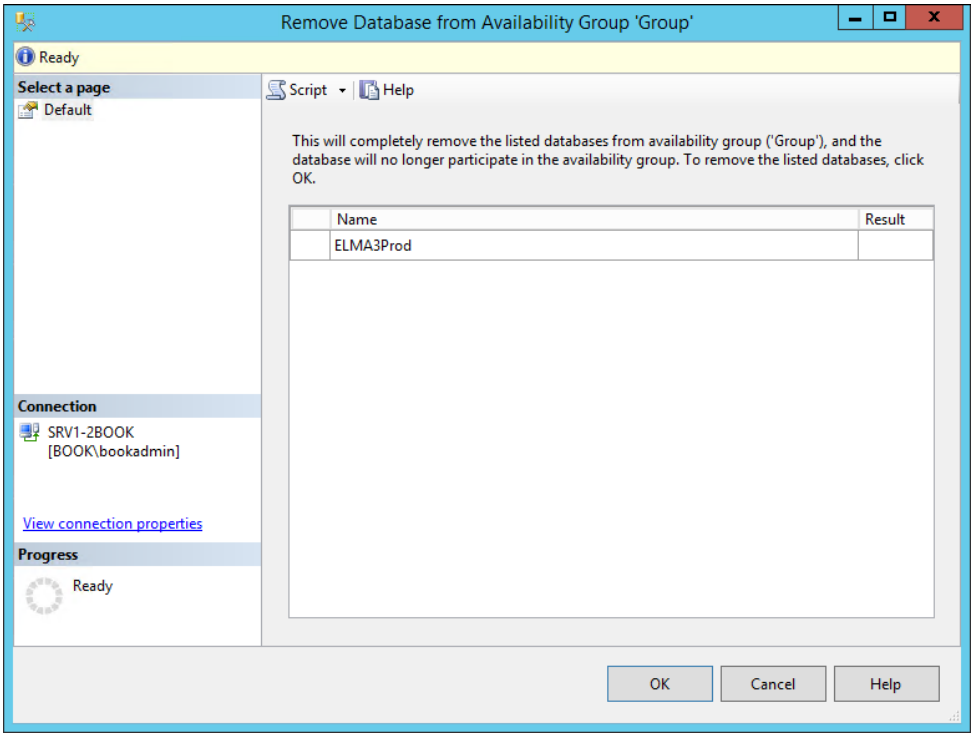

Fig. 124 Removing a database from an availability group

4. Make sure, that this database has the “Not Synchronizing” status in the list of the secondary replica’s databases (Fig. 125).

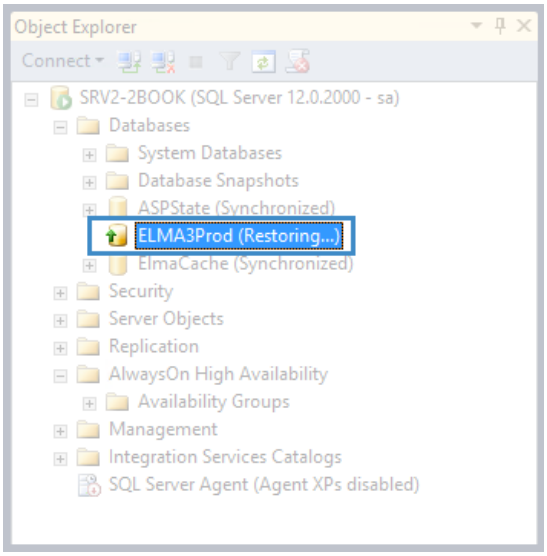

Fig. 125 MS SQL Management Studio. Database with the "Not Synchronizing" status

Remove this database from the secondary replica. In the context menu of this database (Fig. 126) select **Remove**.

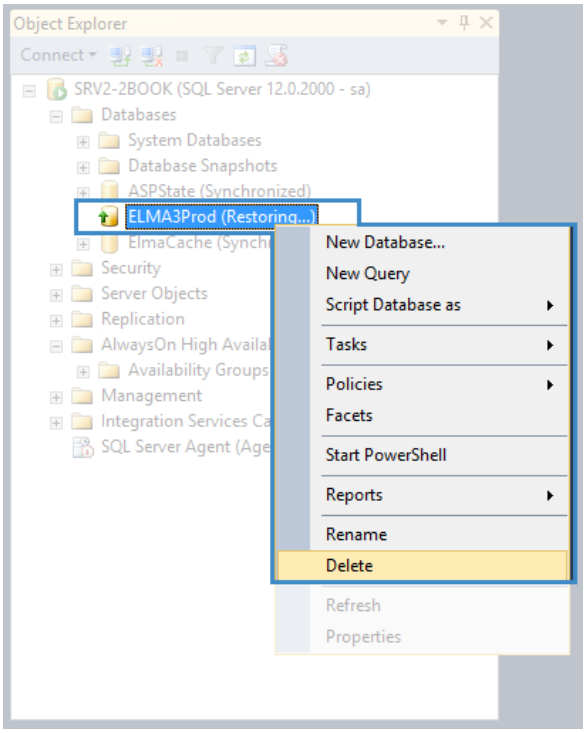

Fig. 126 MS SQL Management Studio. Database context menu

In the opened dialog box (Fig. 127) check the box **Close existing connections** and click **OK**.

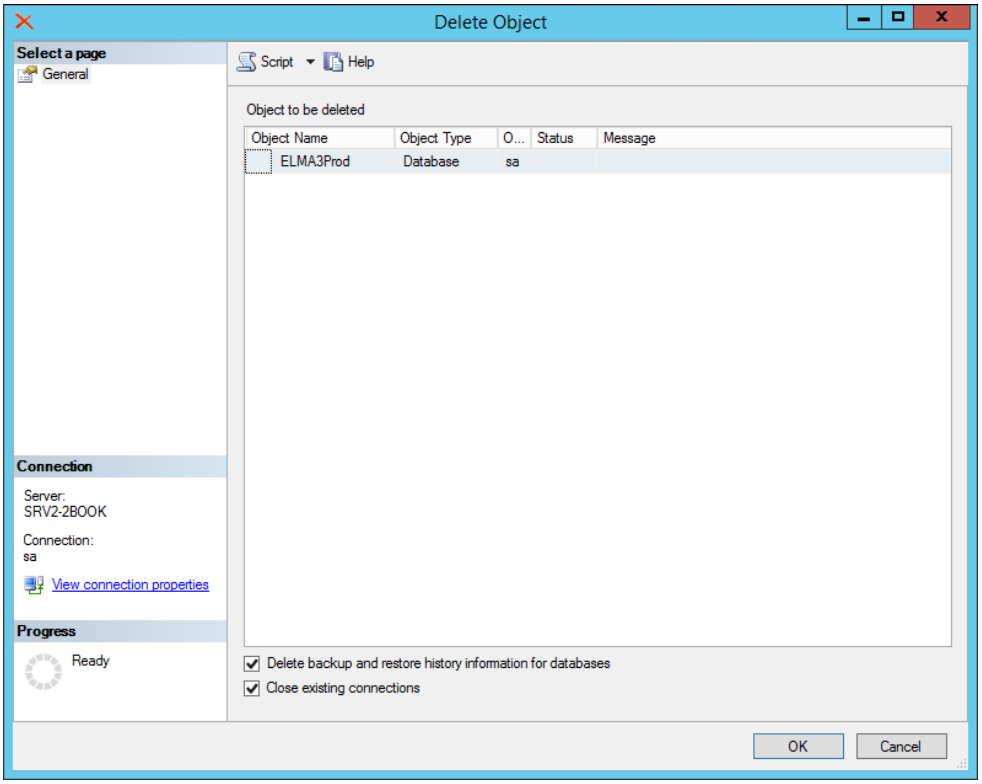

Fig. 127 Delete Object dialog box

5. Next, you can recover the database on the main replica as usual – the database does not participate in synchronization.
6. In the main DBMS replica, open the database context menu and select **Tasks – Restore – Database**... (Fig. 128).

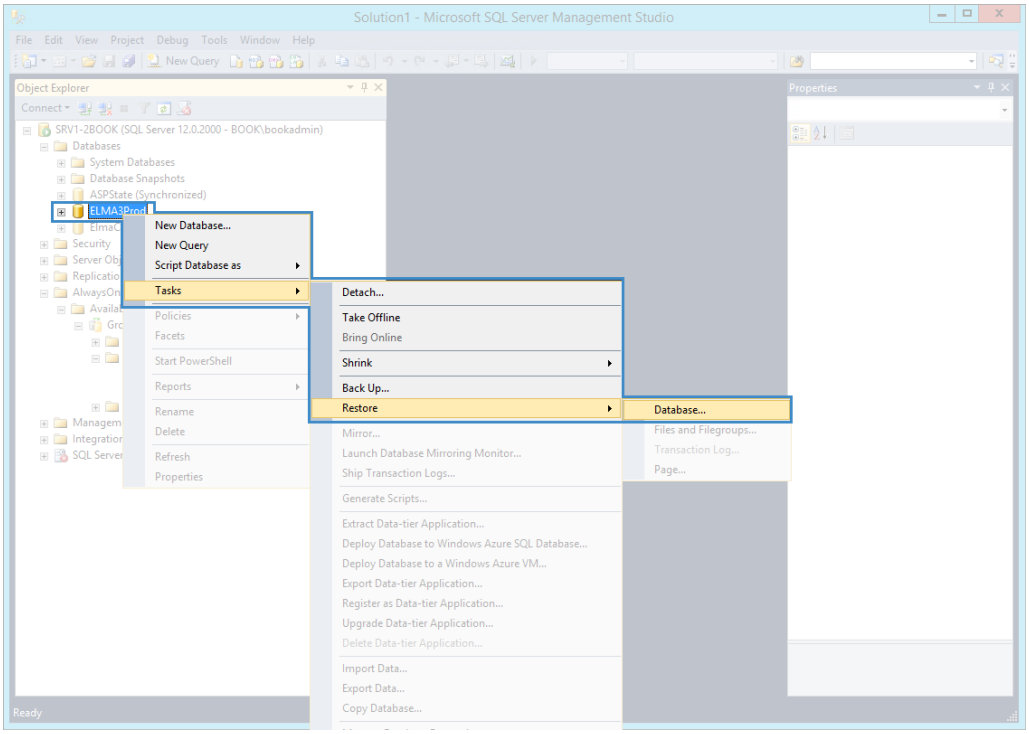

Fig. 128 MS SQL Management Studio. Database context menu

In the opened dialog box (Fig. 129), on the **General** page set the switch to **From device** in the **Source for restore** unit. Click on the "..." button, to select a backup file for restoring.

Attention! If you are restoring a database not from a file, choose the respective options.

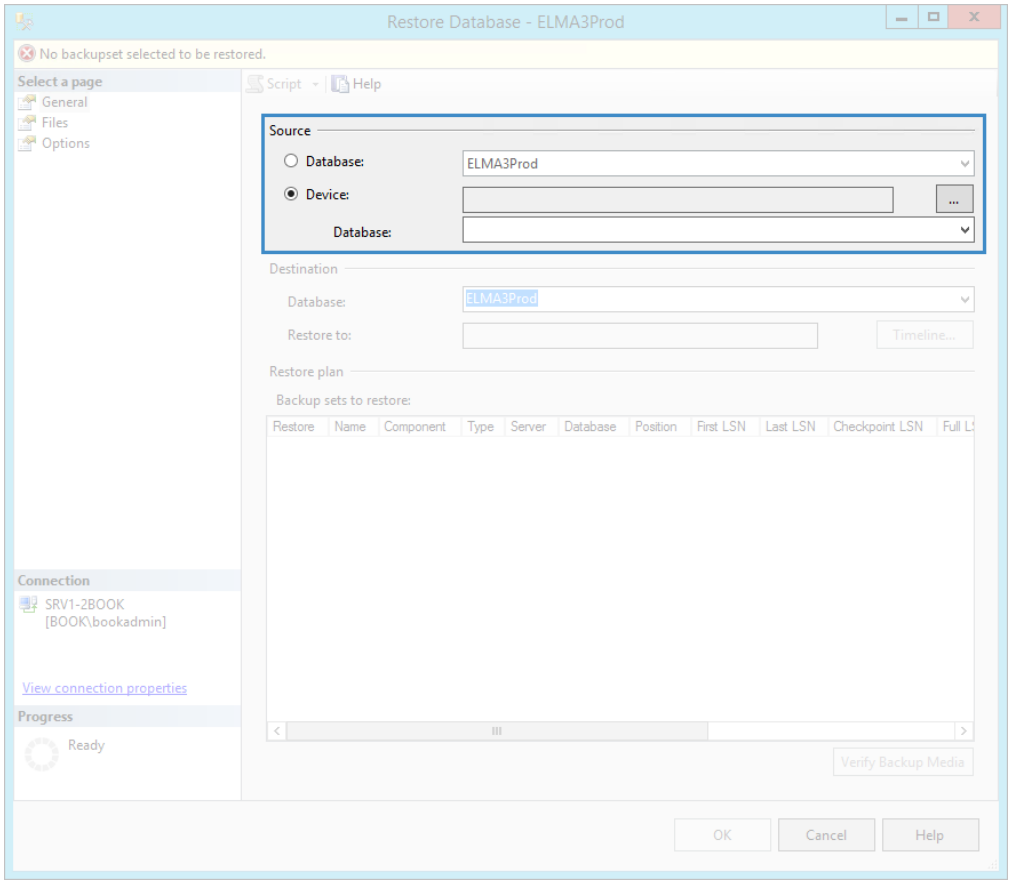

Fig. 129 Restore Database dialog box

When you click on the "..." button, a dialog box opens (Fig. 130), where you need to select **File** in the **Backup media type** field and click **Add.**

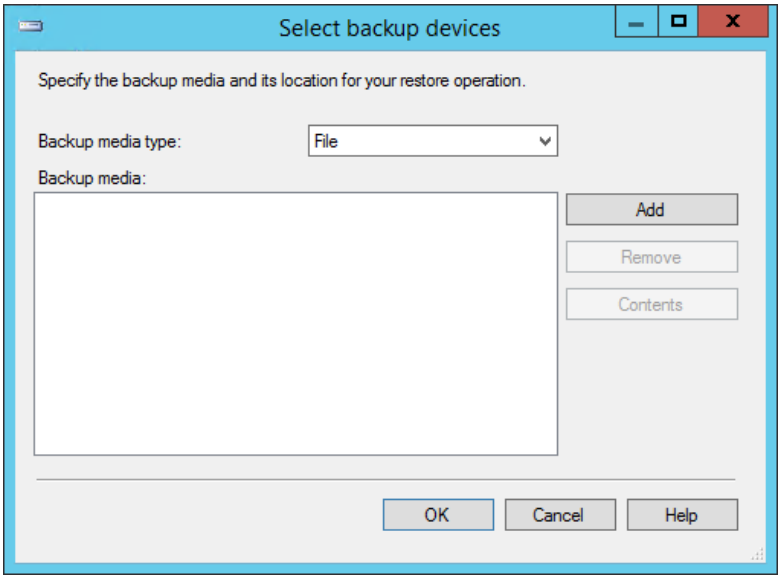

Fig. 130 Selecting a backup device

In the opened dialog box (Fig. 131), select a backup file.

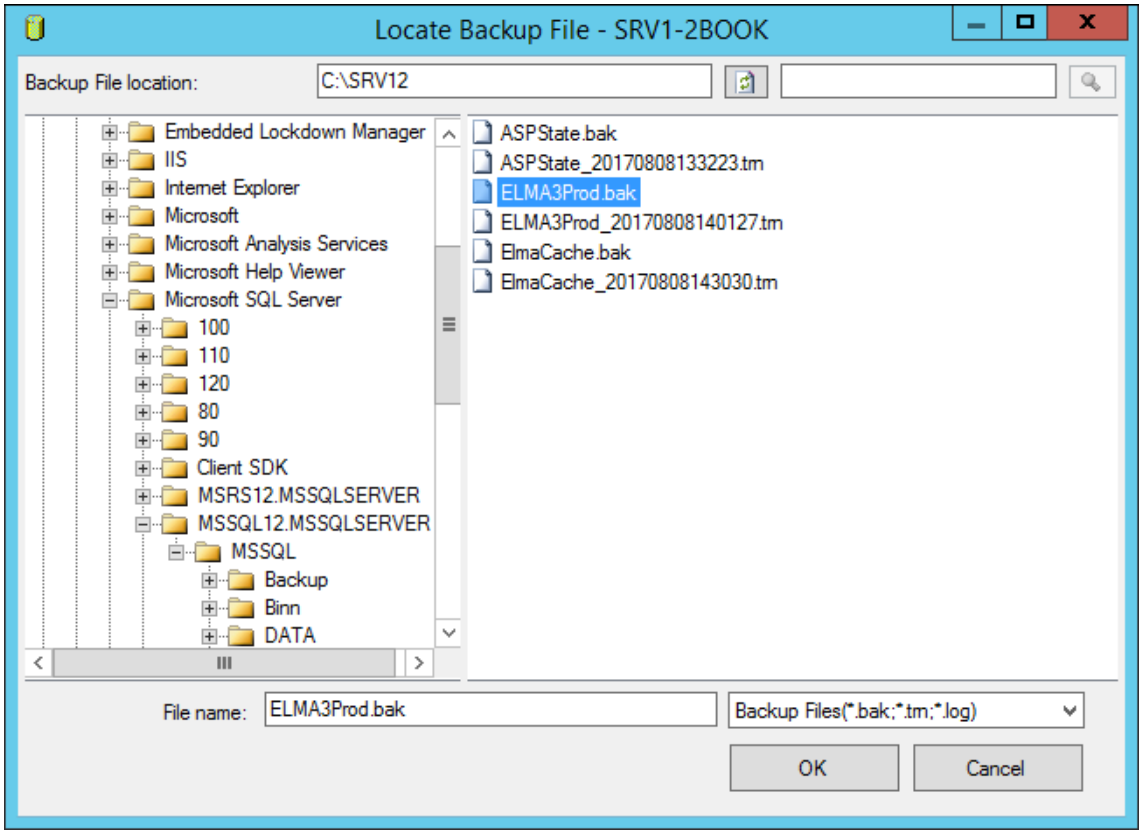

Fig. 131 Selecting a backup file

If you cannot find the required backup copy in the folder, set the file extension filter above the **OK** and **Cancel** buttons to **All Files**.

Confirm the changes by clicking **OK** (Fig. 130 and Fig. 131).

The database recovery wizard will run several checks and automatically suggest the appropriate database.

Attention! Make sure, that the suggested database is the one you intend to recover.

This is very important, since by default, the wizard suggests the database with the same name, as in the backup copy (Fig. 132), and that is not always what is required.

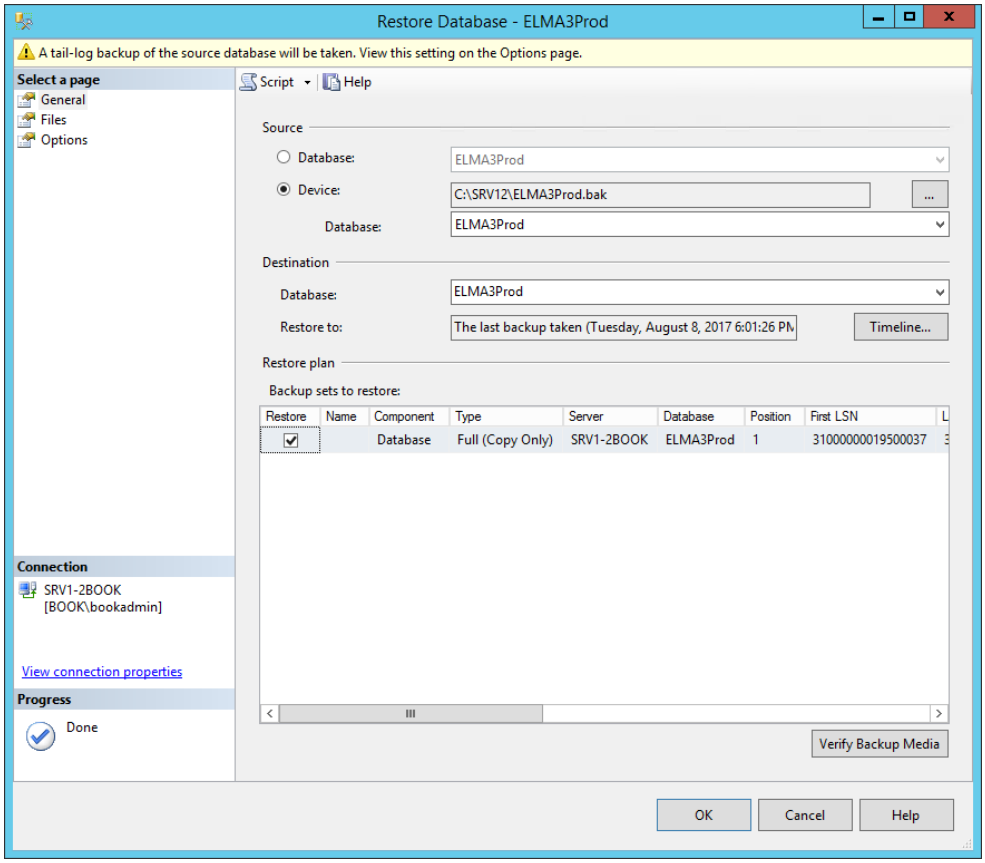

Fig. 132 Restore Database dialog box. General page

After that open the **Files** page (Fig. 133) and make sure, that the database files are located where they are supposed to be. If necessary, check the box **Relocate all files to folder** and configure the respective settings.

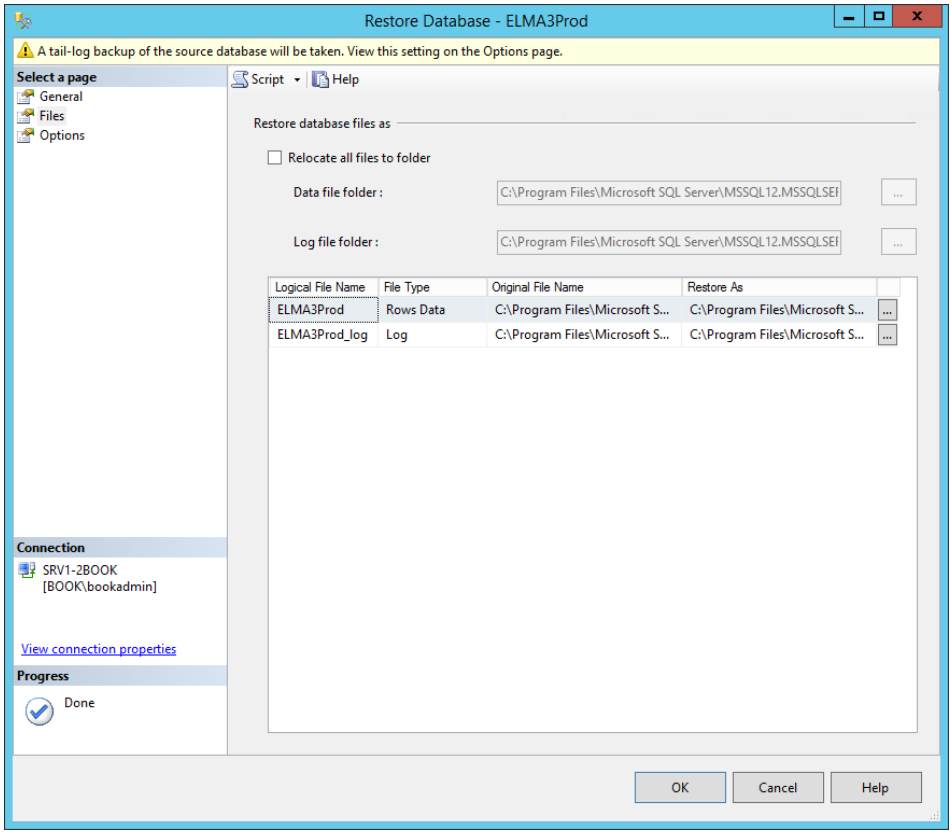

Fig. 133 Restore Database dialog box. Files page

On the **Options** page (Fig. 134), check the boxes **Overwrite the existing database (WITH REPLACE)** and **Close existing connections to destination database**.

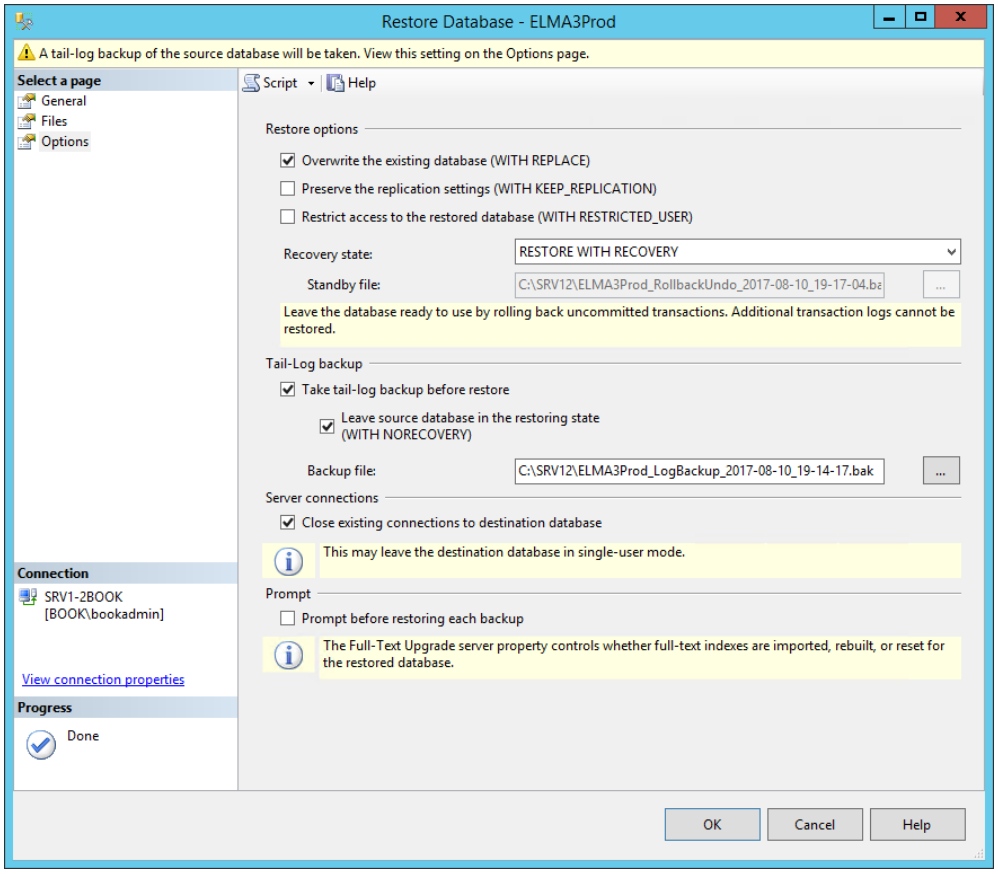

Fig. 134 Restore Database dialog box. Options page

Click **OK** and wait until recovery is complete, it may take a while. During the recovery, the dialog box will be displaying the recovery progress indicator.

Sometimes database servers may have limited disk space, while the database size is significant. For example, the hard drive size is 1 TB, and the database size is 600 GB. Recovering this database may take 600 GB more, i.e. 1.2 TB total.

It seems not logical, that the system takes into account the size of the current database, since it will be overwritten during recovery. In this case, the solution is to delete the database and recover it from grounds up or recover an empty database into it, which will significantly reduce its size.

After the recovery process is completed, you will see a respective dialog box (Fig. 135).

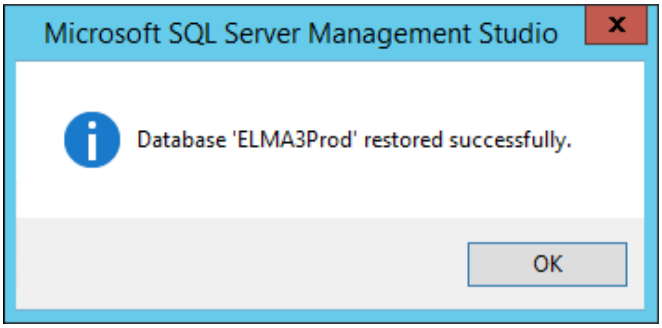

Fig. 135 Completing database recovery

7. Now you need to add the database back to the availability group. To do so, in the main replica open **AlwaysOn High Availability – Availability Groups – Availability Databases**, open the context menu and select **Add Database**... (Fig. 136).

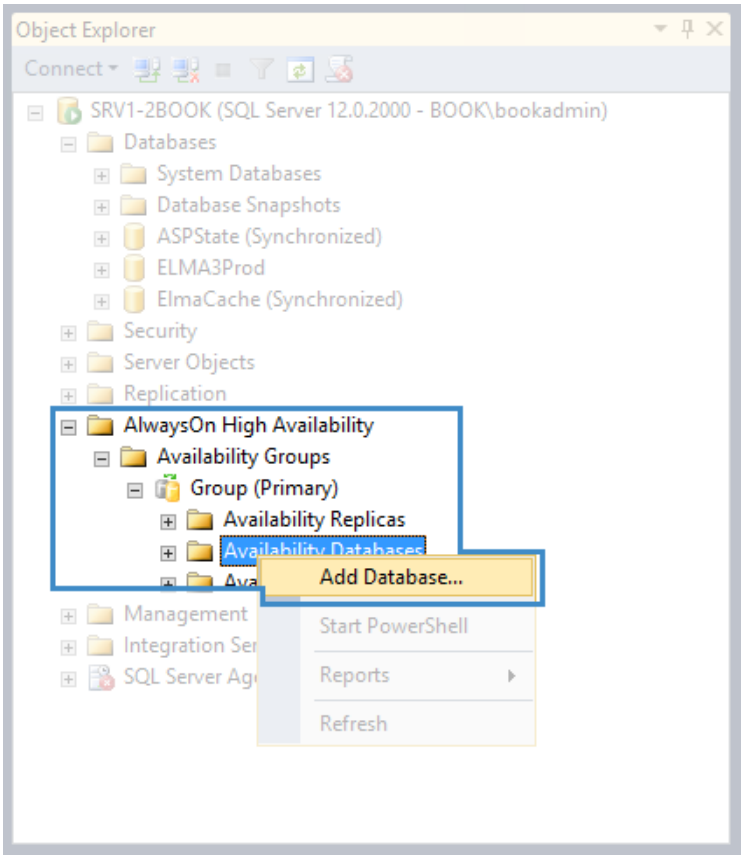

Fig. 136 MS SQL Management Studio. Availability Databases context menu

The wizard for adding a database to the availability group will open. Skip the first step of the wizard by clicking **Next**.

At the **Select Databases** step (Fig. 137), check the box next to the recovered database. It must meet all the requirements (see the **Status** column).

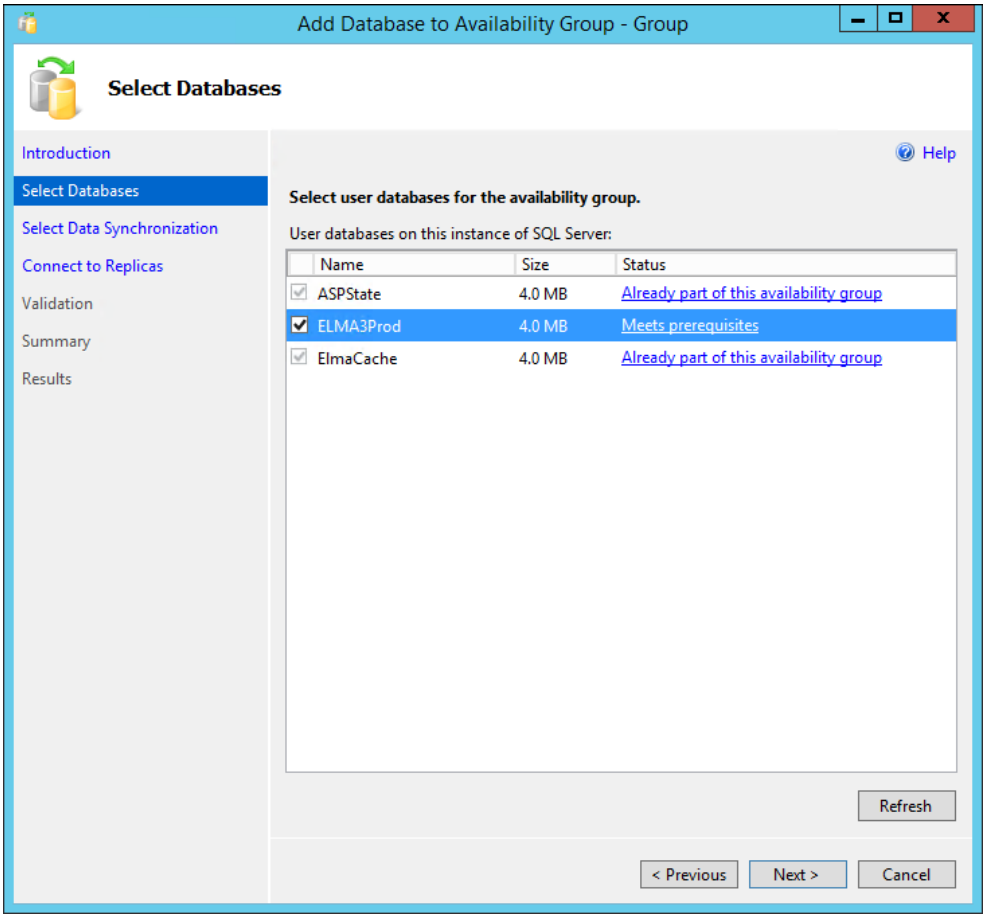

Fig. 137 Adding a database to the availability group. Step 2

To go to the next step of the wizard, click **Next**.

At the **Select Data Synchronization** step (Fig. 138) select **Full** and specify the network folder, available for all replicas. Search the required folder by clicking **Browse**...

After selecting the required folder, click **Next** to go to the next step of the wizard.

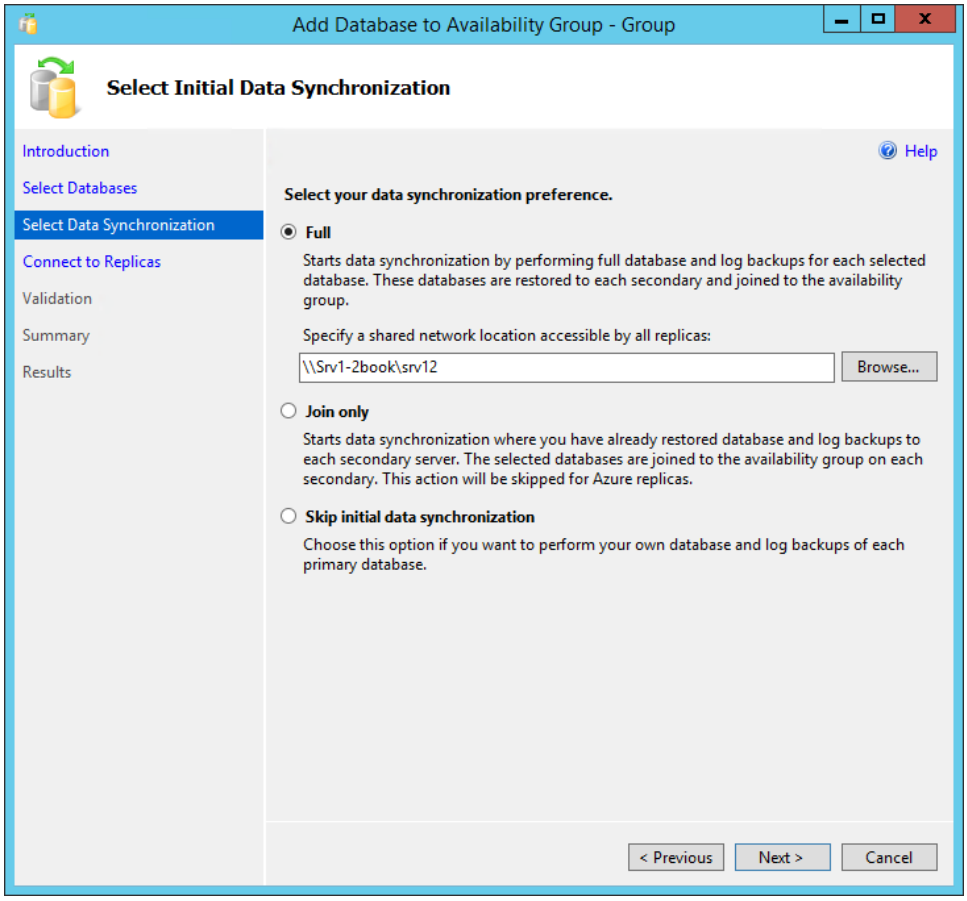

Fig. 138 Adding a database to the availability group. Step 3

At the next step (Fig. 139), connect to all the secondary replicas. To do so, click **Connect all**...

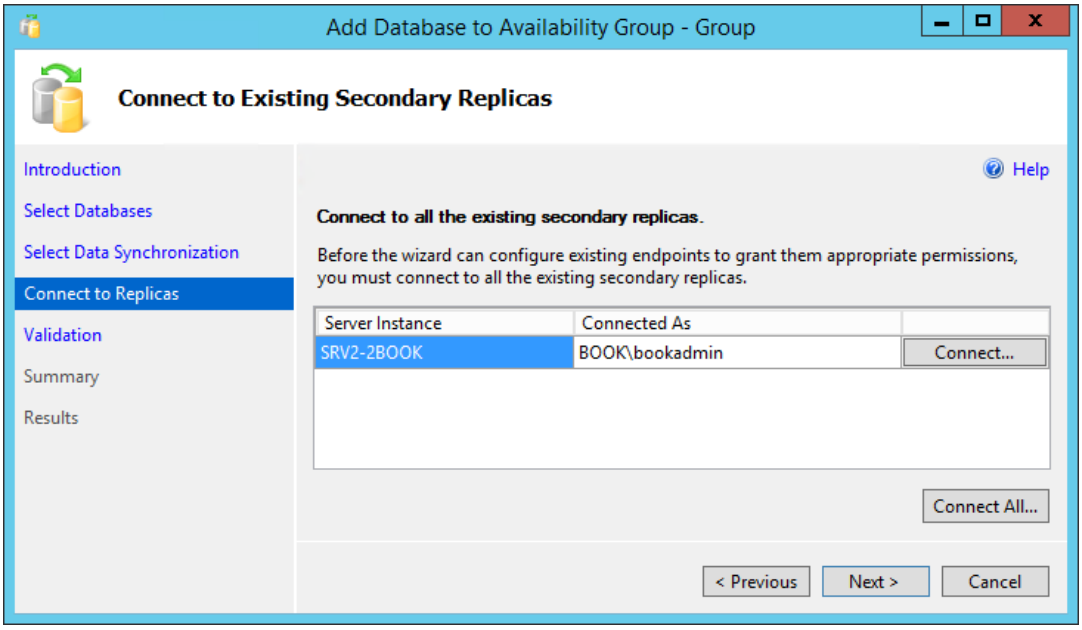

Fig. 139 Adding a database to the availability group. Step 4

To go to the next step of the wizard, click **Next**.

Complete the remaining steps of the wizard and click **Finish** (Fig. 140).

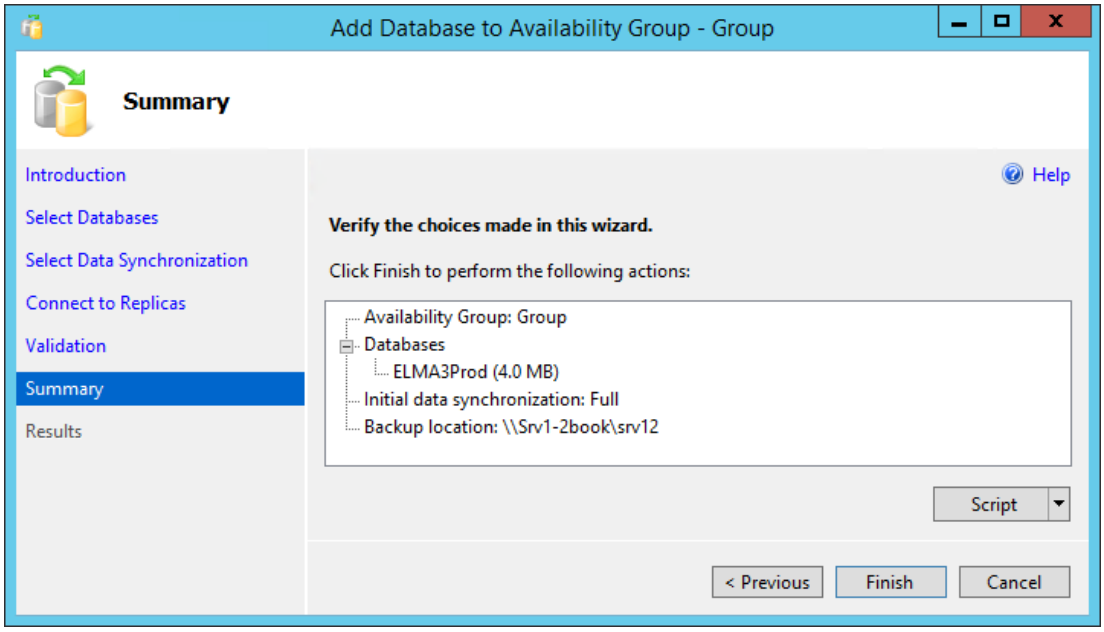

Fig. 140 Adding a database to the availability group. Step 6

Once the wizard completed its work, a respective notification will be displayed (Fig. 141), close the wizard.

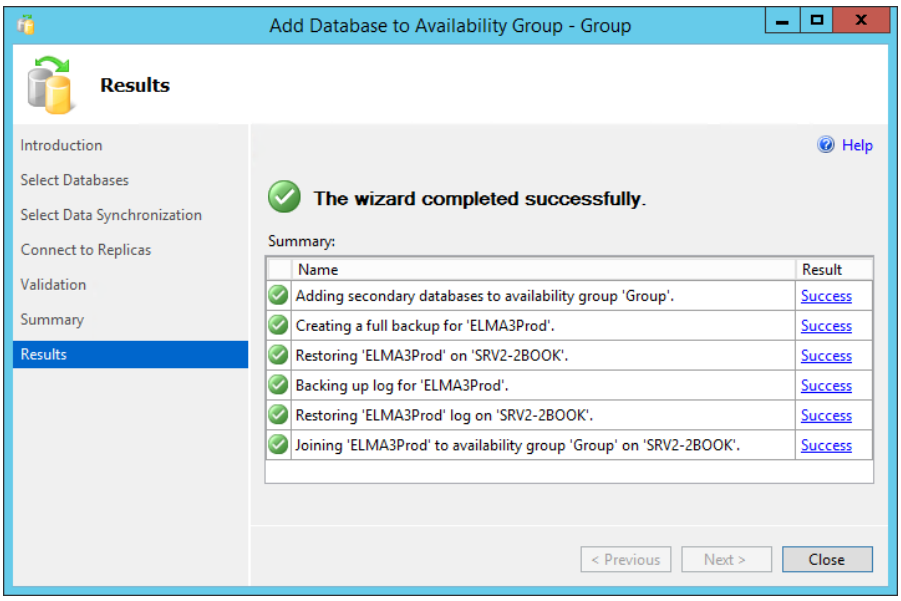

Fig. 141 Adding a database to the availability group. Step 7

8. Open the list of the secondary replica databases. For a while, the recovered database will be synchronizing and will not be available.

This process may take up to several hours. After that, the database status will automatically change to **Synchronized** (Fig. 142) and the database will become accessible.

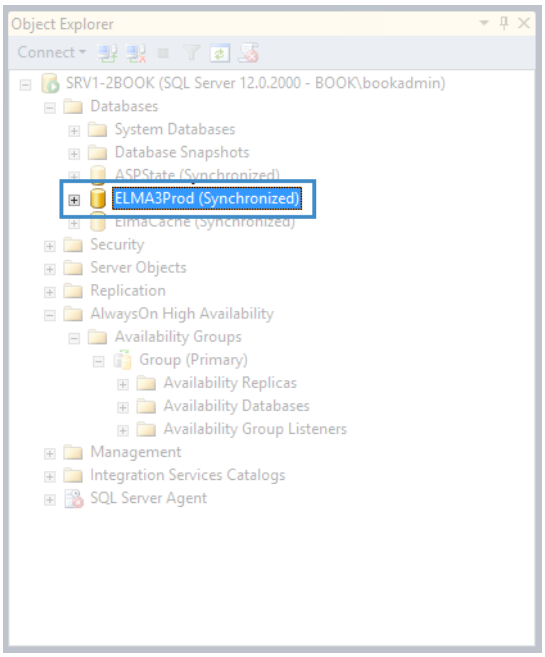

Fig. 142 MS SQL Management Studio. Synchronized database

5.1.6. Restoring a database from the transaction log / rolling back the state to specific time

It may be necessary to return the database state to specific time, e.g. before publishing a certain object, or executing a request to delete certain important data. It is highly recommended that you back up the database manually as described in section 5.1.5 before performing potentially harmful operations and making significant changes.

Since the database works with the full recovery model, it is possible to restore the state of any selected time, if there are backup files. In addition, it is possible to restore using TailLog – the part of the transaction log, saved at the current moment and not used for creating backup copies.

Restoring a database implies removing it from the AlwaysOn availability group and closing all the connections. You must stop ELMA while carrying out these operations. Depending on the database size and connection speed, this process may take up to 15 minutes.

1. Make sure, that ELMA is stopped. In case of a farm, you must stop the ELMA application pool in the IIS Manager on each server.

2. Back up the database manually. This may come in handy, if the database is not restored to the required data in the first attempt.

When restoring with TailLog, MS SQL automatically backs up the log, which may be lost (it is recommended that you create the backup copy manually).

3. First, stop data movement in the AlwaysOn High Availability.

To do so, in MS SQL Management Studio go to **AlwaysOn High Availability – Availability Groups**, select the required availability group; open **Availability Databases**, and select the database you need to restore and remove it from the availability group (Fig. 143).

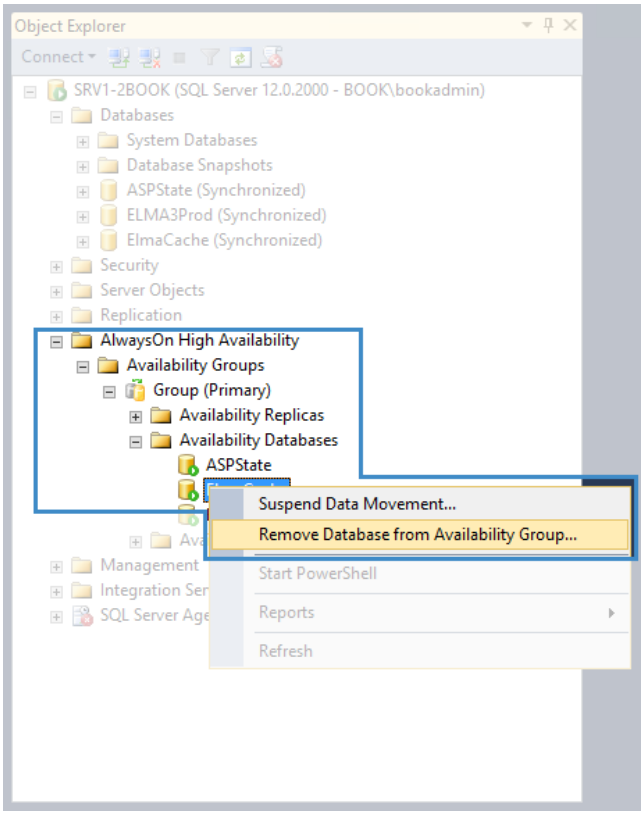

Fig. 143 MS SQL Management Studio. Availability group context menu

In the appeared dialog box, click **OK**.

4. In MS SQL Management Studio, open the context menu of the required database and select **Tasks – Restore – Database...** (Fig. 144).

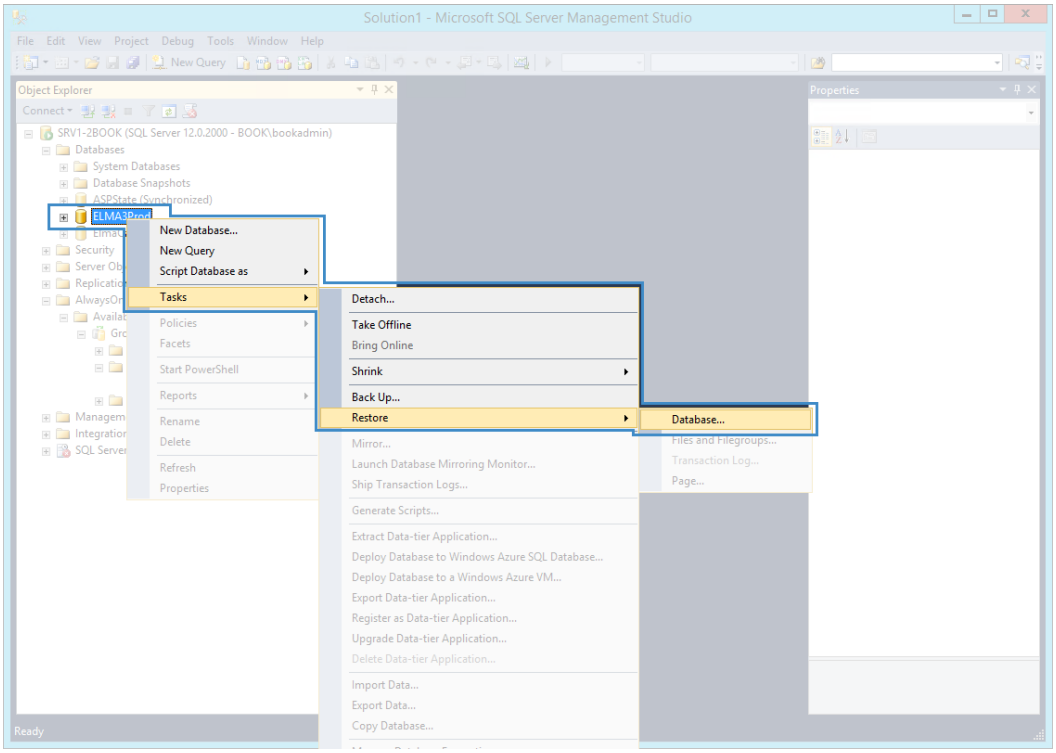

Fig. 144 MS SQL Management Studio. Database context menu

In the opened dialog box, on the **General** page (Fig. 145), select **Database** in the **Source** unit and select the required database.

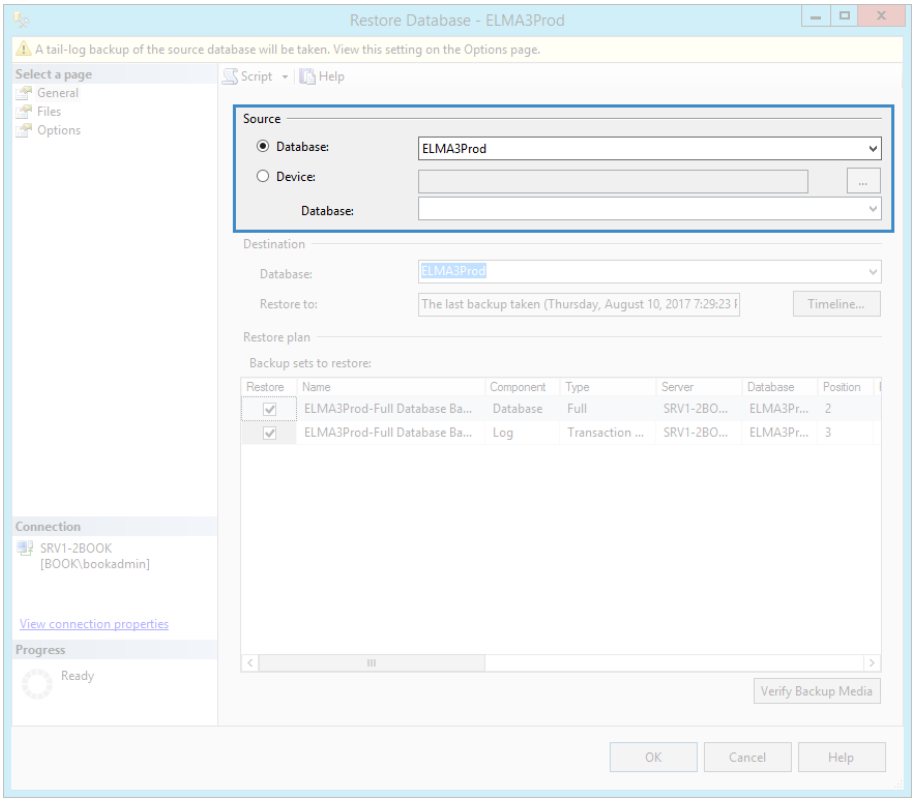

Fig. 145 Restore Database dialog box. General page

In the **Destination** unit, make sure that the required database is selected in the **Database** field and click **Timeline**...

In the appeared dialog box (Fig. 146), select **Specific date and time**, and specify the required date and time. You can select time on the timeline. Here you can also see all the available backup copies.

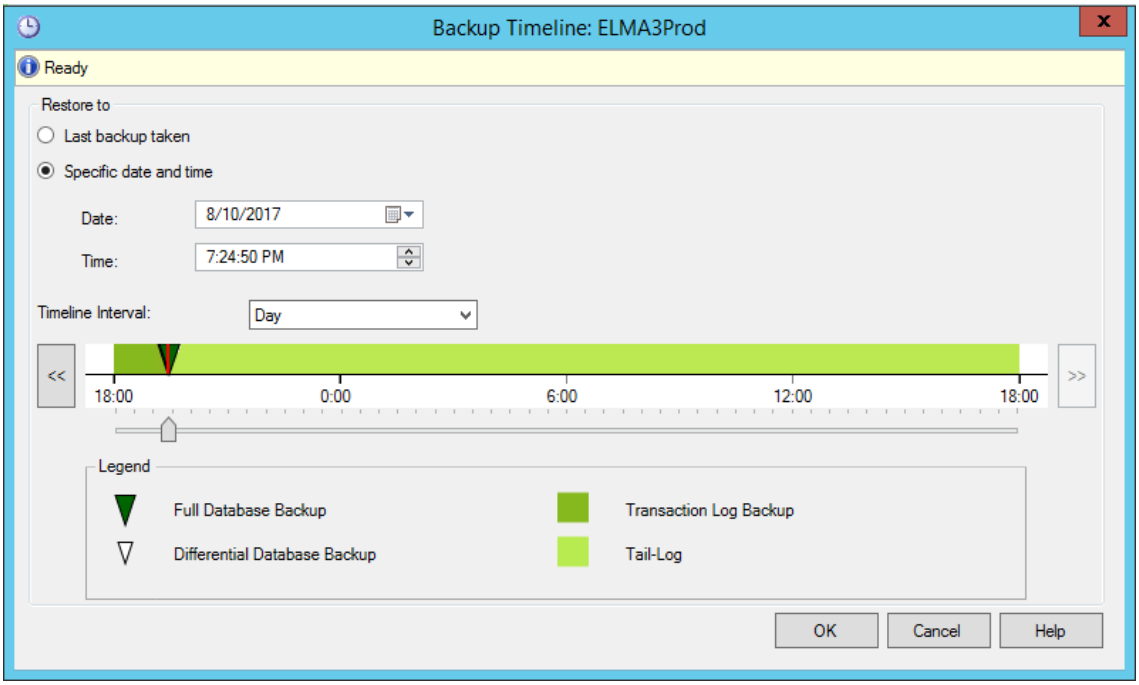

Fig. 146 Backup Timeline dialog box

After selecting the required date and time for restoring the database state, click **OK**.

In the **Restore Database** dialog box, open the **Options** page (Fig. 147) and check the box **Overwrite the existing database (WITH REPLACE)** in the **Restore options** unit.

In the **Server connections** unit, check the **Close existing connections to destination database**. It is necessary if the ELMA application pool, copying in the availability group or another user are connected to the database.

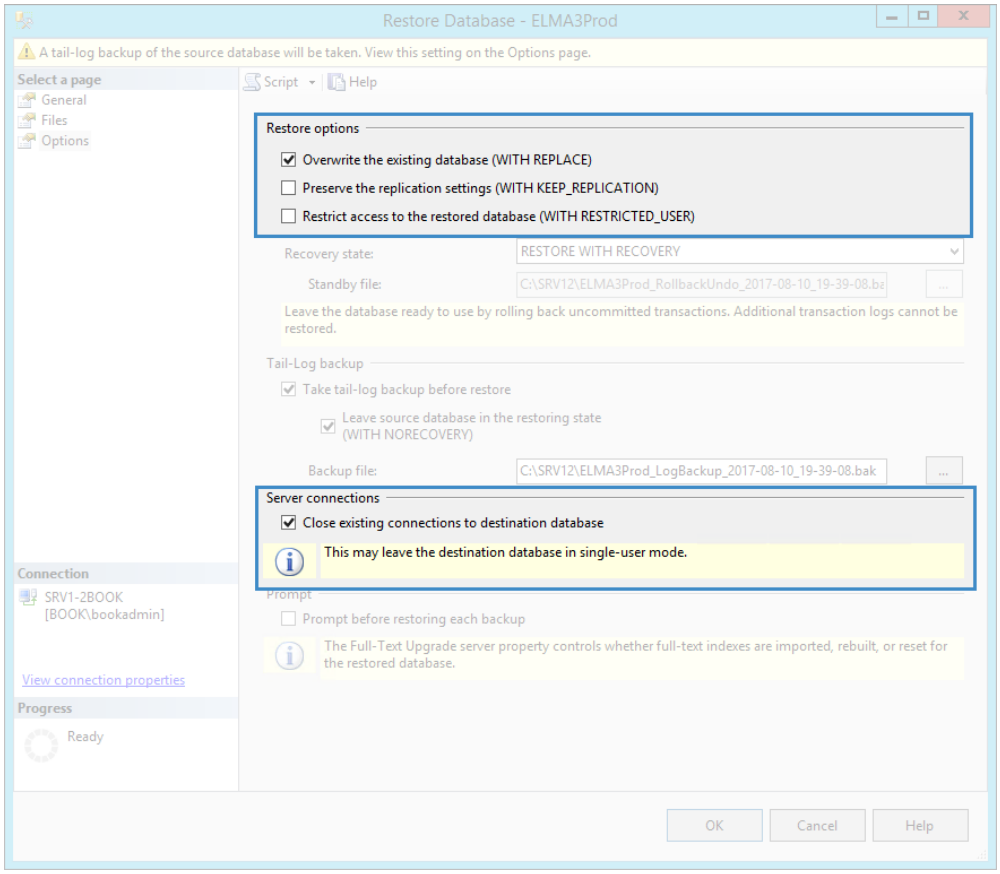

Fig. 147 Restore Database dialog box. Options page

5. Once the database is restored, add it back to the availability group (Fig. 136). You will also have to delete the database from the secondary replica.

5.2. Maintaining ELMA

5.2.1. Configuring automatic system diagnostics

This step is not obligatory; however, it is recommended that you follow through it for further operation and system maintenance.

You can configure system diagnostics parameters in ELMA Web Application, in **Administration – System – System Diagnostics** (Fig. 148).

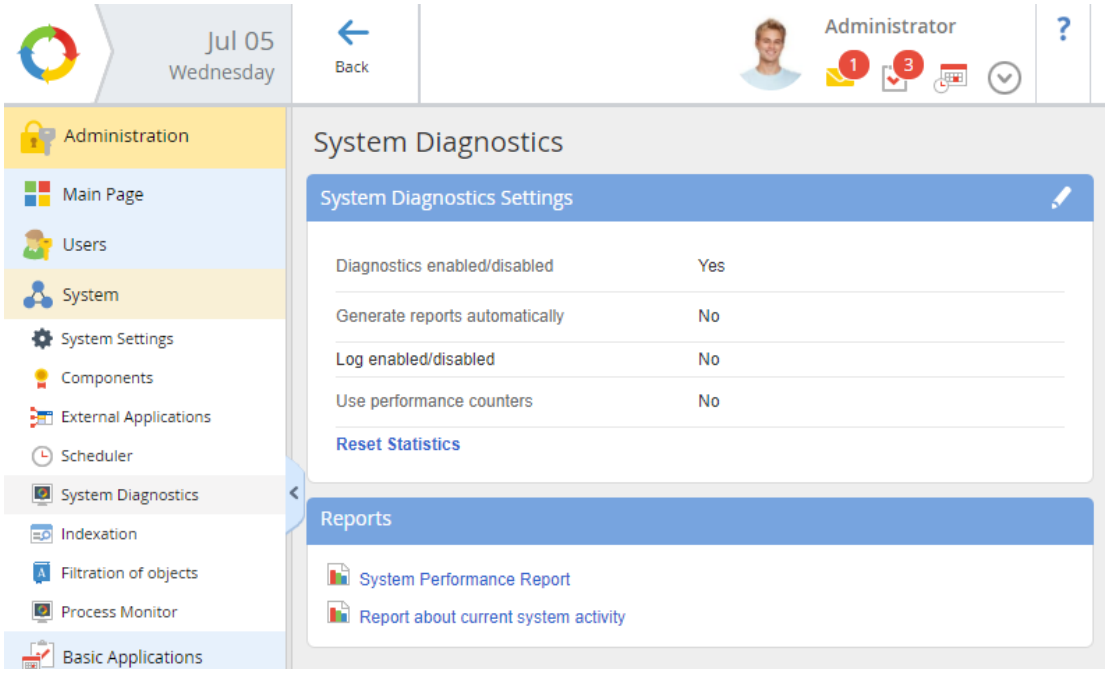

Fig. 148 Administration – System – System Diagnostics

To edit the system diagnostics settings (Fig. 149) click on the icon, next to the unit name.

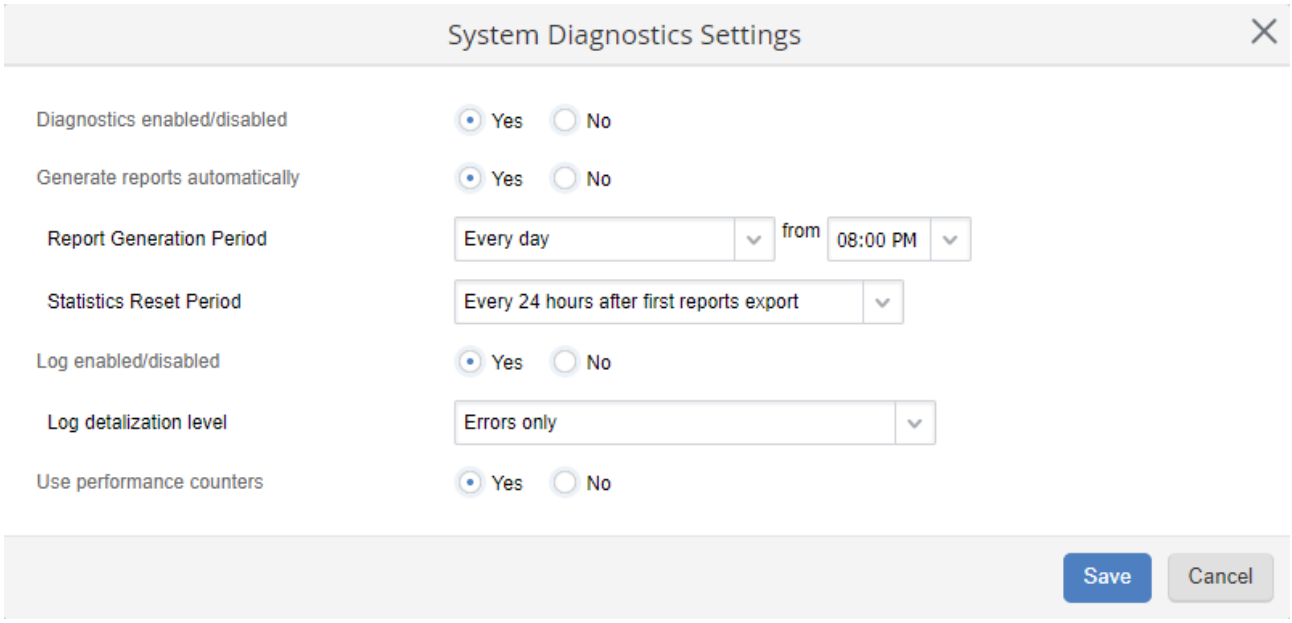

Fig. 149 System diagnostics settings

The **Generate reports automatically** parameter activates automatic generation or reports on ELMA performance with the defined periodicity. Specify the period for generating reports, first report generation time, which will be used as the start point for the period, and the statistics reset period.

The following parameters are recommended:

- **Report Generation Period** – Every day, starting at 20:00. Specify the time two hours after the end of the business hours.
- **Statistics Reset Period** – Every 24 hours after first reports export. This way, statistics will be gathered every day.

Configure other settings as necessary. To save the changes, click **Save**.

5.2.2. Gathering ELMA error-logs from several servers

Despite the possibility to configure logs from all the servers to be saved in a shared folder while specifying the server name, we discourage doing so for systems that work under high load, since it may increase the load on the network, which is one of the essential parts of the system.

It is recommended that you provide the Administrator account with shared access to the Web folder of ELMA (C:\\ELMA3-Enterprise\Web) on each server (access to read will be required to gather logs).

Attention! Any changes you make in the Web folder may become the reason to restart the system on a particular server. For this reason, it is

inadvisable to grant permissions to the employees, who do not need them.

This way, the administrator can quickly gather log files from all the servers by accessing the following folders:

- \\SRV12-1\Web\Logs\Error
- \\SRV12-2\Web\Logs\Error

5.2.3. Restarting ELMA farm

Restarting ELMA farm may be required in different situations. For example, for database and server maintenance.

It is important to understand that a farm involves several ELMA servers and restarting it means stopping each particular server and then starting it.

There may be three situations:

- Restarting servers with full stop of the farm (section 5.2.3.1);
- Restarting servers during operation (section 5.2.3.2);
- Restarting servers to import objects or update (sections 5.2.4 – 5.2.6).

5.2.3.1 Restarting servers with full stop of the farm

This restart type means that the farm will be unavailable to the users during the restart.

1. Open IIS Manager and stop ELMA pool on each server (Fig. 150).

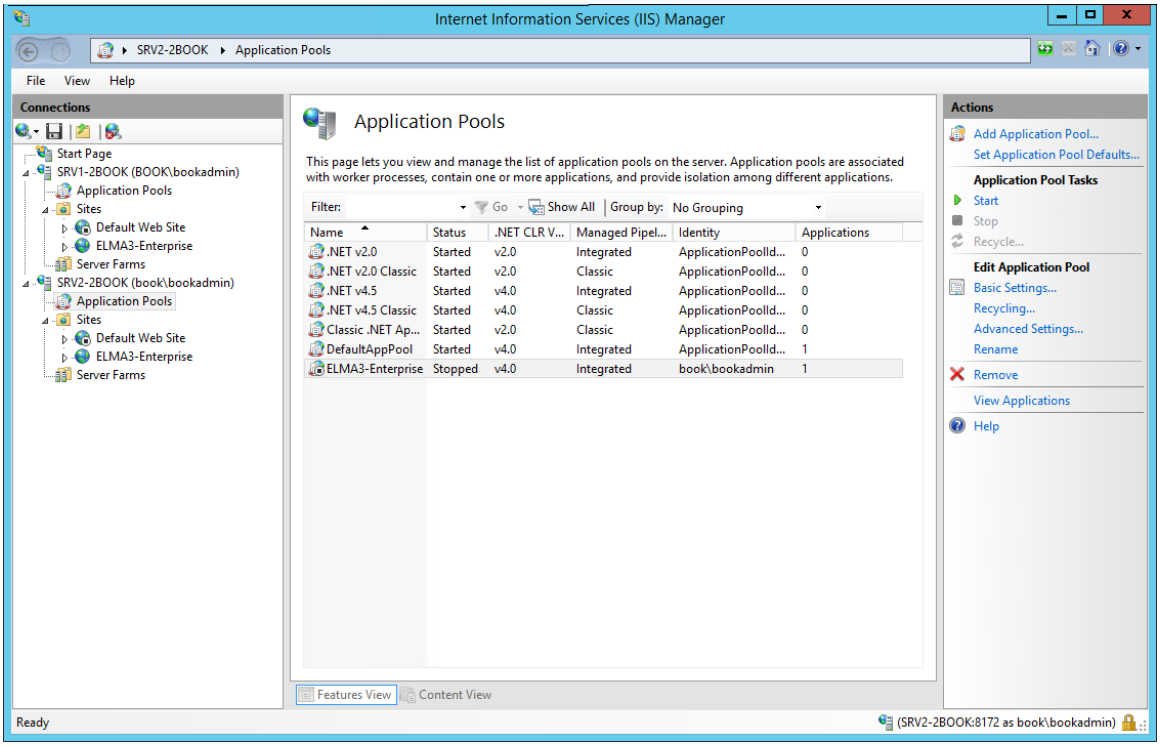

Fig. 150 IIS Manager

2. Next, you can restart the cache cluster. It is recommended that you do that periodically. You can do that during the system maintenance.

To stop the cache cluster on any application server (e.g. **SRV12-1**) start the **PowerShell** console as the administrator. Go to **Start -> All Programs -> AppFabric for Windows Server -> Caching Administration Windows PowerShell** (Fig. 76).

In the command prompt, execute the "stop-cachecluster" command. You will see a progress bar and the caching cluster stop result (Fig. 151).

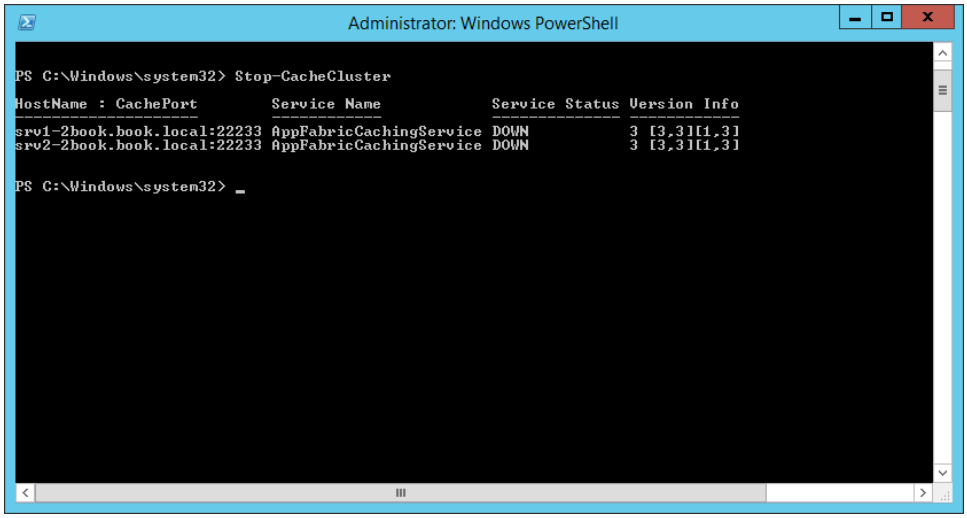

Fig. 151 Windows PowerShell. Stopping the caching cluster

After stopping the caching cluster, you can immediately start it with the "start-cachecluster" command (Fig. 152).

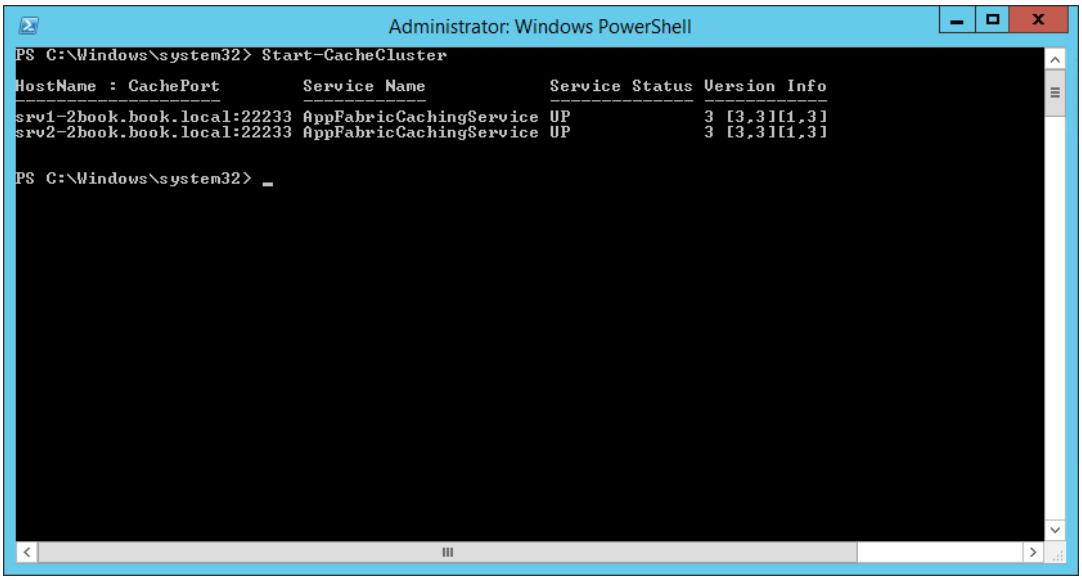

Fig. 152 Windows PowerShell. Starting the caching cluster

3. Start ELMA pool on each server one by one.

Start the pool on one of the servers. Open ELMA Web Application on this server and wait until it has been started.

Attention! During restart, do not try to open ELMA Web Application using the controller address, it will return an error until at least one server is started.

4. Repeat step 3 for the next server.

Although there is a mechanism that automatically starts all servers one by one ("Waiting for another ELMA server to start"), we recommend that you control this process yourself, so that you could timely react to possible issues.

5. As the result, all servers must be started.

5.2.3.2 Restarting a server without stopping the farm

In case of an emergency, you may need to partially restart the farm – specific servers or all servers but without stopping the farm.

Attention! If the situation allows waiting till the end of the business day or restarting with full stop (see выше), restarting without stopping the farm is not recommended.

The idea of this approach is in "disconnecting" a server from the farm and restarting it, while the other servers continue working in the farm.

1. Open IIS Manager on the farm controller (93.158.134.30 – WFCONTROLLER).
2. Go to **Server Farms**; find the name of your farm (in this case, **ELMA3**) and open **Servers** (Fig. 153). Here you will see the list of all application servers, included in the farm, and their statuses.

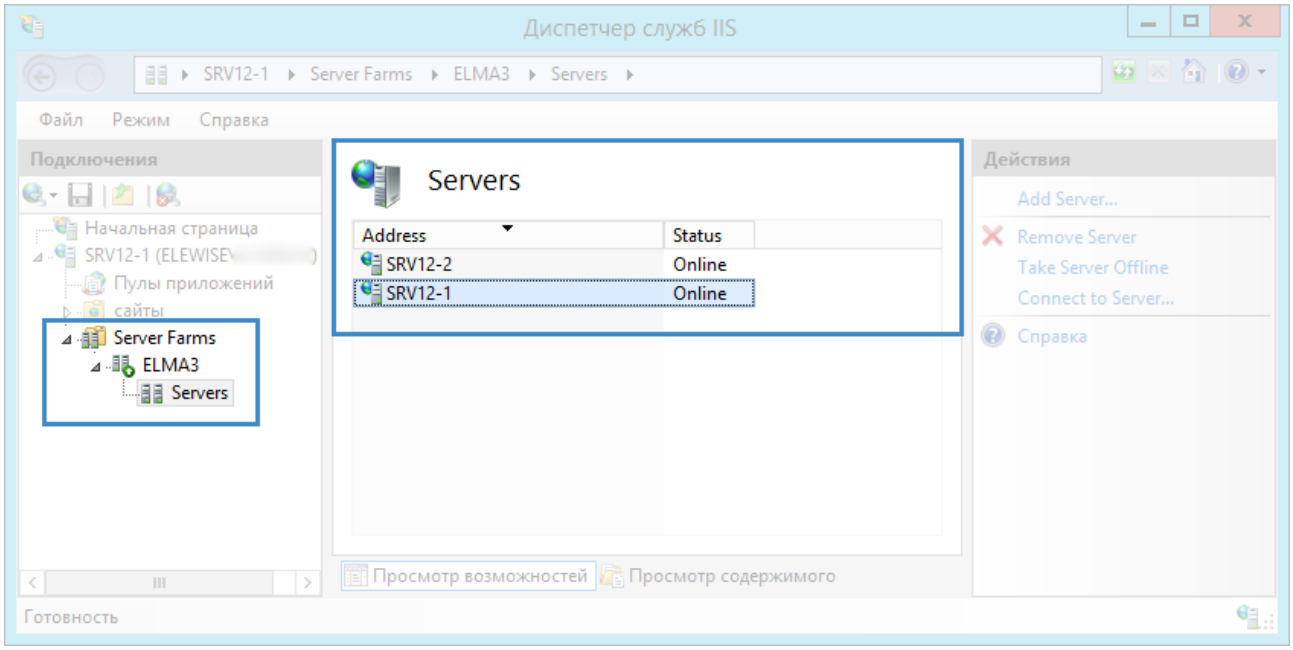

Fig. 153 IIS Manager

3. Select the server you need to restart (e.g. SRV12-1), and click **Take Server Offline** in the right menu. In the opened dialog box (Fig. 154), click **Yes**.

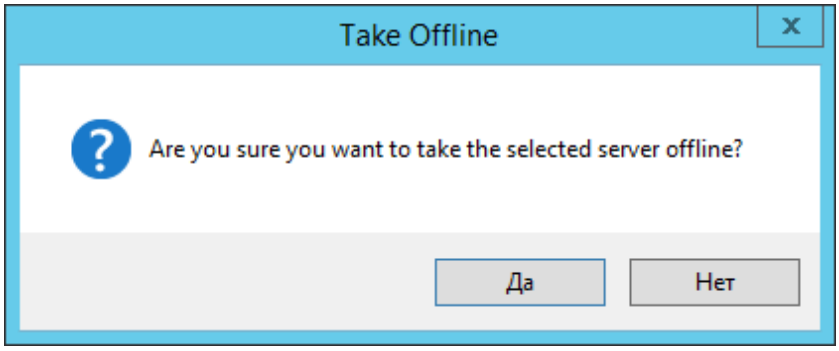

Fig. 154 "Take Server Offline" dialog box

This way you will "disconnect" the server from the farm, and it will no longer receive requests from users, while the server will remain enabled.

Attention! When disconnecting a server from a farm, make sure, that there are still active servers in the farm. If you disable the only online or available (on the "Monitoring and Management" panel) server, users will receive an error when accessing the Web Application.

4. After that, you can perform any actions with the SRV12-1 server, including restarting ELMA pool or even restarting the operating system.

Attention! Do not stop the caching cluster under any circumstances – it is required for all the servers at the same time and stopping it will restart all the other servers.

5. Once the server is started, go to **Server Farms**, find your farm (in this case, **ELMA3**) and open **Servers** (Fig. 153).
6. Select the SRV12-1 server and click **Bring Server Online** in the right menu. After that, the server will continue receiving requests.
7. If necessary, repeat all the steps for other servers.

5.2.4. Importing configuration via ELMA Designer

Importing a configuration to a farm has peculiarities, related to object publishing.

1. First, stop all the ELMA application servers except for the one you will import to.
2. Connect ELMA Designer to this server and perform import (as usual).

If the imported configuration contains metadata that requires restart (objects, documents, global modules, etc.), then confirm restart at the import step that will suggest restarting. Other servers were stopped beforehand.

3. Once the import is complete, start all the ELMA servers.

5.2.5. Installing and updating packages/components

Installing and updating modules via the package manager has peculiarities. When installing via web interface, the system checks the package compatibility, installs the module and makes an entry in the database. In case of the farm, the procedure is following:

1. Copy the package to install/update to the **Packages** folder on each server, located in the ELMA folder, e.g. **C:\\ELMA3-Enterprise\Packages**.
2. Stop all the ELMA servers except for one (e.g. SRV12-1).
3. Open ELMA Web Application on the SRV12-1 server and go to **Administration – System – Components** (Fig. 155).

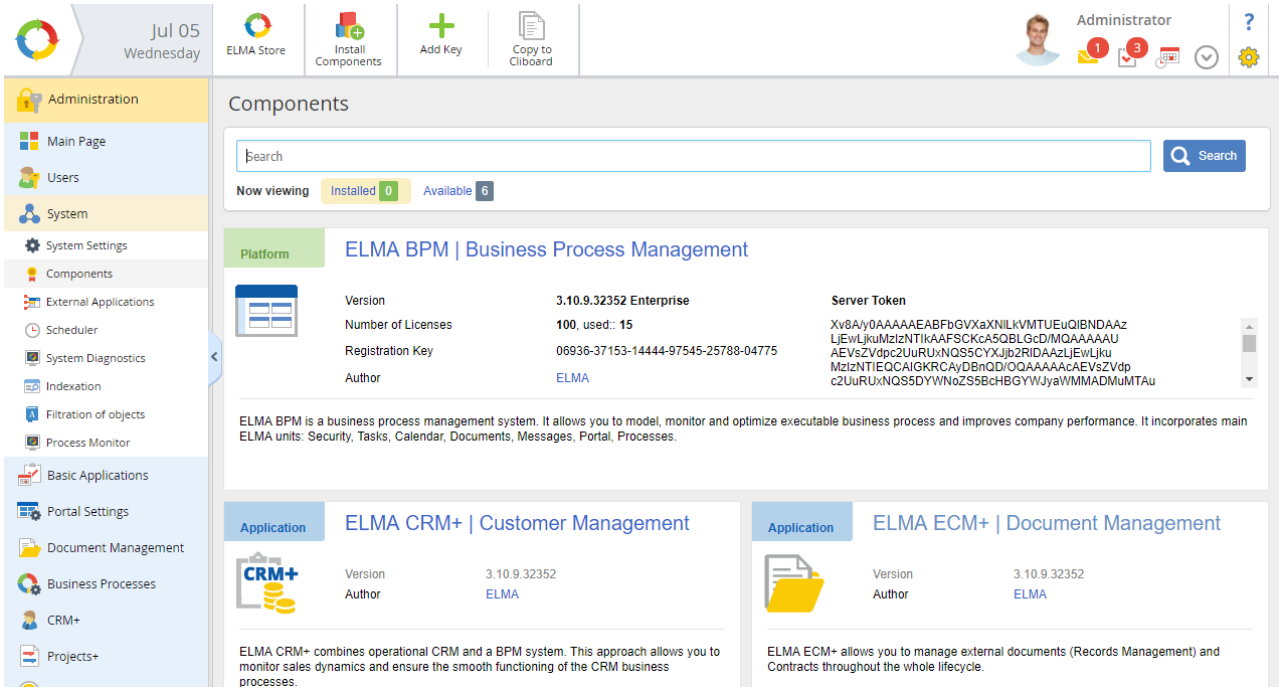

Fig. 155 ELMA Web Application. Administration – System – Components

Click **Install Components** or **Update Components** in the top menu. In the opened dialog box, select the required components, click **Install** and wait until the installation is complete. It may take a while. During this process, the ELMA server will be restarted.

You can monitor the components installation/update progress using the address of ELMA Web Application that will be temporarily replaced with the web part of the package manager.

If you open IIS Manager, you can see that the website path has changed. After the package manager completed its work, path will change back to the ELMA website path.

4. Once the SRV12-1 server is started after package installation/update, proceed with installing/updating packages on the other servers in the console mode.

On each of the remaining servers, open the command prompt as the administrator and run **\PackageManager\EleWise.ELMA.Packaging.Console.exe** with a parameter – name of your package and installation key in Web (otherwise it will also be installed to the Designer).

The installation command will look as follows:

EleWise.ELMA.Packaging.Console.exe install EleWise.ELMA.Project -c Web

Here is an example of the full installation command:

C:\ELMA3-Enterprise\PackageManager\EleWise.ELMA.Packaging.Console.exe install EleWise.ELMA.Project -c Web

Note, that the package name is specified without **.Web** and version (Fig. 156).

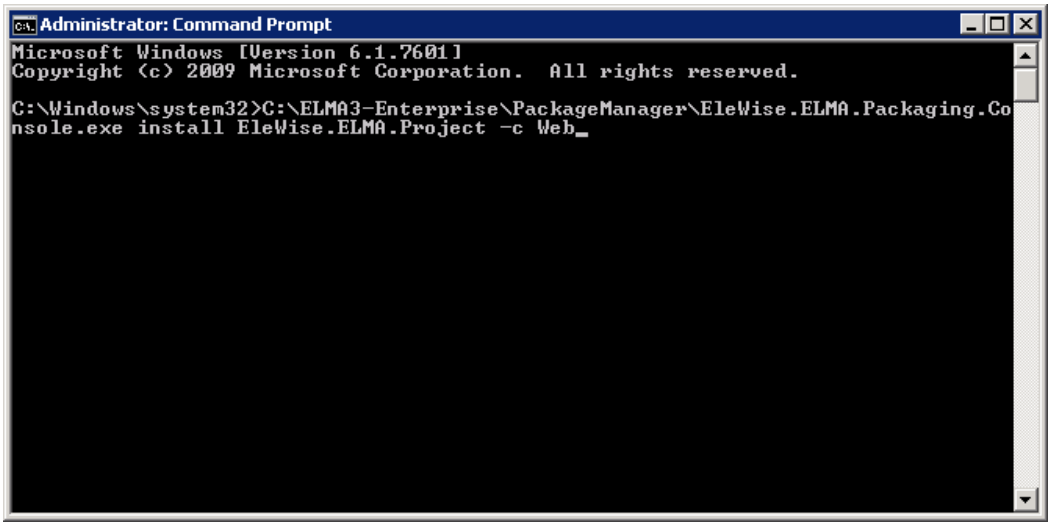

Fig. 156 Command prompt

If you receive an error, make sure you have started the console as the Administrator.

5. Start the ELMA server in the IIS Manager.
6. Open Web Application of the started server, go to **Administration – System – Components** and make sure, that the component is displayed in the list and its version corresponds to the installed one.

Installation to this server is complete. Apply steps 4-6 to the next stopped server. As the result, all the ELMA servers must be started.

5.2.6. Updating ELMA farm

Updating ELMA in a farm is not much different from regular update. Take into account, that all the servers must be stopped before updating, and started only after updating ELMA on each server.

To update ELMA farm, follow this procedure.

1. Stop ELMA pools on each server.
2. Install system updates on each server as usual.
3. Restart the caching cluster.

To stop the caching cluster on any application server (e.g. 93.158.134.3 – SRV12-1) start the **PowerShell** console as the Administrator. Go to **Start -> All Programs -> AppFabric for Windows Server -> Caching Administration Windows PowerShell** (Fig. 76).

In the command prompt run the "stop-cachecluster" command. You will see the stopping progress bar and the result (Fig. 151).

After stopping the caching cluster, you can start it with the "start-cachecluster" command (Fig. 152).

4. Next, start ELMA pool on each server one by one.

Start the pool on one of the servers. Open ELMA Web Application on this server and wait until it is started.

5. Repeat step 4 for the next server. As the result, all the servers must be started and work with the new ELMA version.

5.3. Maintaining Application Servers

Application servers rarely require any maintenance; however, there are several rules.

1. Monitor CPU, memory and network usage on the servers, especially in case of peak load.
2. If CPU or memory usage exceeds 80% during peaks, you should start increasing capacity. Usually it is topical after the system load is balanced and users' work is even.
3. Monitor empty disk space, and if necessary clear it. Don't forget to disable IIS logging, if it is not required for trouble shooting.
4. If the **Windows\Temp** folder grows, you can clear it when ELMA is offline.

5.4. Maintaining File Storage

These are the main recommendations for maintaining the file storage:

1. Allocate a hard drive with plenty of disc space.
2. Make sure that there is always free space – absence of free space may cause the system or some functions to slow down or stop.
3. Do not enable antimalware scanners. It will significantly slow down work with files. If necessary, schedule daily scanning for nighttime (after 2:00 AM).
4. Do not try to edit or change files in the **Files** folder – ELMA checks hash-sums of files, uploaded via the system.

5.5. Maintaining Controller Server

The controller server is least demanding in terms of resources and maintenance, however there are several recommendations:

1. If CPU or memory usage exceeds 80% during peaks, you should start increasing capacity. Usually it is topical after the system load is balanced and users' work is even.
2. Monitor empty disk space, and if necessary clear it.

The controller uses very little disk space (only for some logs), but if the disk runs out of free space, the response time will increase by several times.

3. Disable IIS logging, if it is not required for trouble shooting. Practice shows, that it is the only issue with free space on the controller.

Chapter 6. Possible Issues and Trouble Shooting

This chapter describes possible issues you may face when installing, configuring, operating and maintaining ELMA web farm with an MS SQL AlwaysOn cluster. If you have faced an issue, which is not described in this chapter, try searching on the Internet, since most issues are related not to ELMA, but to other components (such as .NET Framework, VMWare, MS SQL Server, ASP.Net).

6.1. Recommendations on VM Resource Reservation

If the servers for deploying the system (controller, application servers, DBMS server, file server) are virtual machines, pay attention to resource allocation.

The nature of virtual machines is so, that you can allocate 128 GB of RAM and 28 CPU cores that a user will see in the operating system, while placing all that on an overloaded or "outdated" host. In this case, the 28 CPU cores and 128 GB RAM in fact will be very slow. If such a situation occurs, look for the reasons of low performance on the host or VM controller.

1. First, make sure that the VM host server has enough resources.
2. Reserve CPU and RAM resources for virtual machines. Reserve maximum of resources (within the recommended system requirements), especially before putting the system under the maximum load, when there are no exact data.
3. Make sure that the memory allocation on the VM controller is not in the **ballooned** or **swapped** mode.
4. Check total host server load. Situations may occur when the host load is high enough, and resources simply will not be allocated with sufficient speed and in sufficient amount.

6.2. Solving Issues with MS SQL Server Installation

Default instance name already exists

Most likely, the previous installation completed with an error and you need to delete it if you already have an operational MS SQL Server.

1. If an MS SQL Server already exists, you don’t need to install it again – it is ready for creating databases.
2. If the previous installation completed with an error, it would be better to delete the installed components and reinstall it from scratch.

6.3. Solving Issues with Logging in to MS SQL Server

No MS SQL Management Studio

You can download this component from the Microsoft website. Later on, we recommend that you download and install **MS SQL Server with advanced tools** distribution packages, which already include this component.

Impossible to specify an LDAP account

Most likely, the server is not included in the domain. This is an obligatory requirement for the farm and the cluster. Contact the system administrator to include the server to the domain.

You can do it yourself, if you have an account with the required permissions:

1. Go to **Start –> Computer** and select **Manage** in the context menu.
2. In the opened window, select the **Local Server** tab and click **Domain**.
3. In the appeared dialog box click **Change...** on the **Computer Name** tab (Fig. 157) and log in to the domain.

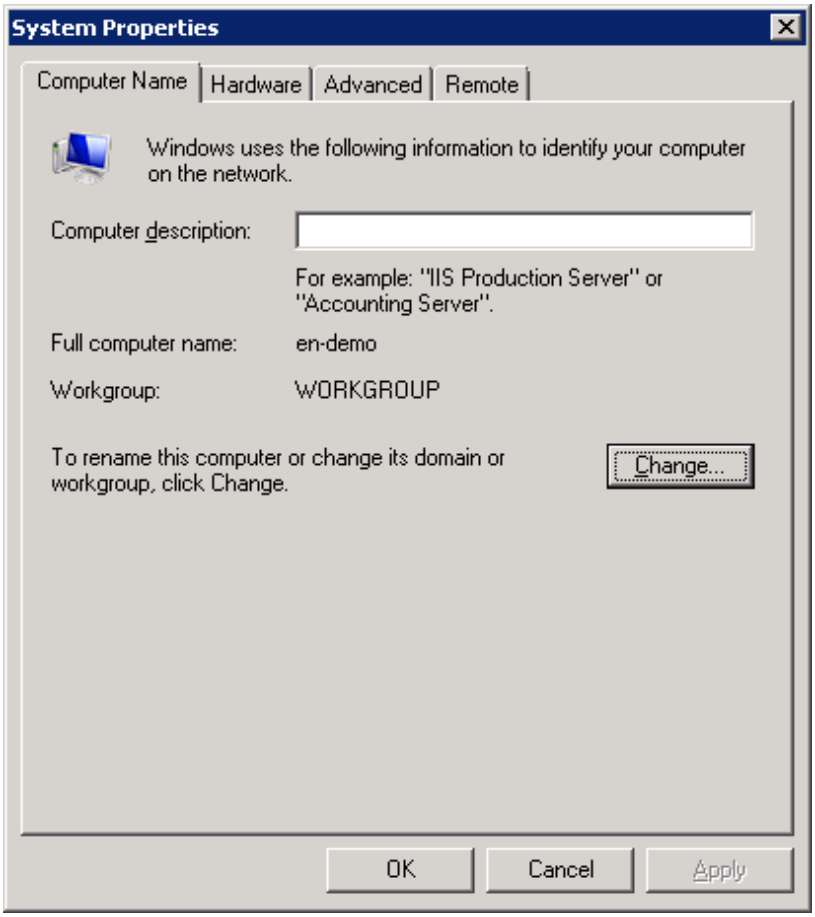

Fig. 157 System properties dialog box

6.4. Solving Issues with Creating ASPState databases

Database with this name already exists

Stop all the ELMA servers, delete the ASPState database and try creating again.

After restarting DBMS, ELMA returns an error that objects are missing in ASPState

Most likely, the ASPState database was created with incorrect parameters, specifically, with saving the structure to RAM, and not as tables.

Delete the ASPState database and create it following the instructions in 3.1.5.

6.5. Solving Issues with AppFabric Installation

Cannot start AppFabric installation (error)

Enable automatic update in the system settings. Without it, installation will not start.

It does not matter if you have Internet connection to receive updates. During the installation, the parameters in the Control Panels themselves are checked.

Port is occupied

If you install AppFabric normally, for the first time, the port should not be occupied. Make sure, that you have not installed AppFabric earlier. If you have an installed AppFabric, uninstall it.

6.6. Solving Issues with Configuring Controller Server

During the installation, the URL Rewrite component has not been installed or when creating the farm, automatic redirect settings were rejected

First, install URL Rewrite; without it, it is impossible to distribute requests among application servers.

Second, open the IIS Manager and select **URL Rewrite** on the server main page in the **IIS** unit (Fig. 158).

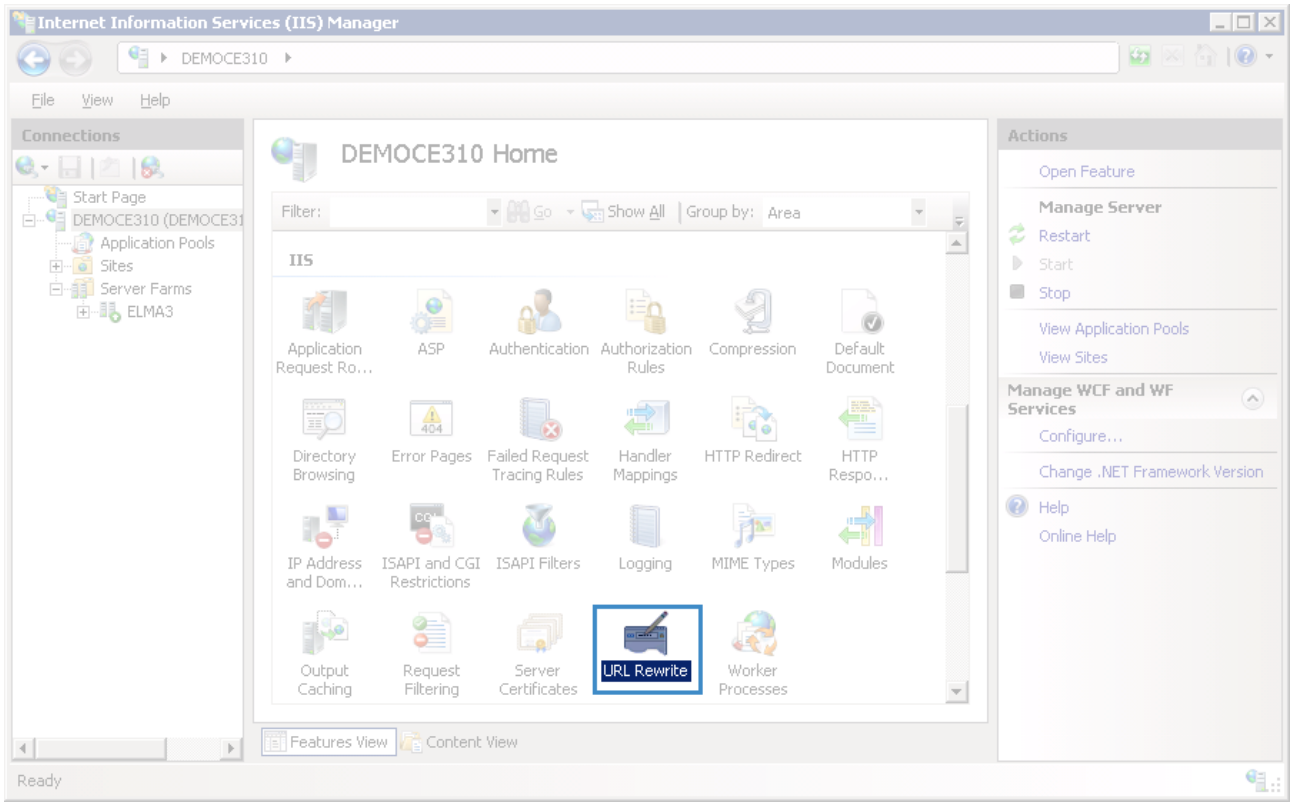

Fig. 158 IIS Manager. URL Rewrite

Create the following rule with a request filtering condition (Fig. 159):

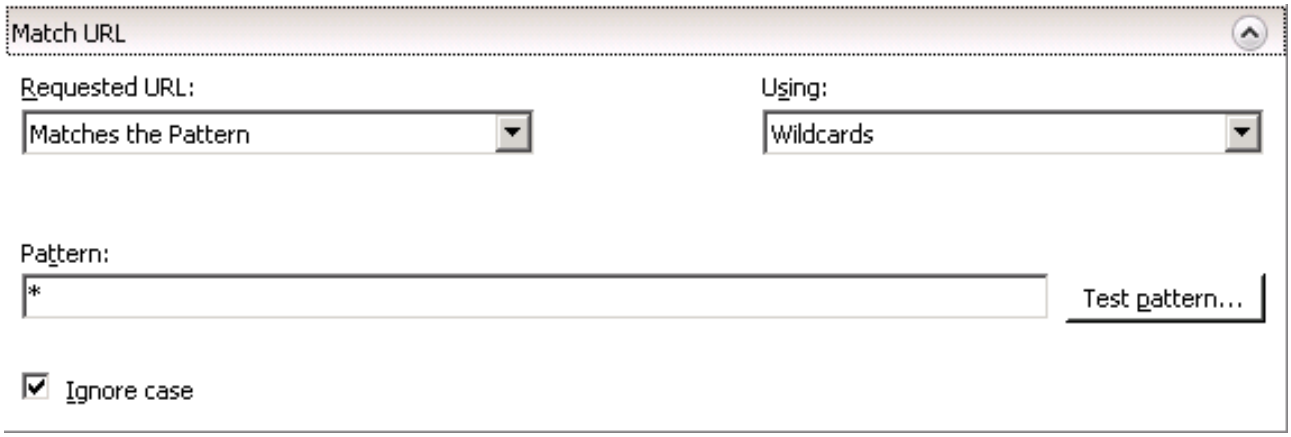

Fig. 159 Configuring URL rewrite

And a rule with an action (Fig. 160):

Fig. 160 Configuring URL rewrite

These settings are configured as part of one rule.

Web farm cannot get access to the Internet during installation/configuration via proxy.

When installing or configuring a farm/farm controller, IIS service or one of dependent components may require temporary access to the Internet.

In this case, IIS does not inherit custom proxy server settings in the web browser or Control Panel.

To configure Internet access via a proxy server for IIS, run the following command in the command prompt as the Administrator:

```
C:\Windows\System32\inetsrv\appcmd.exe set config -section:webFarms
/[name='Farm_name'].defaultProxy.enabled:"True"
```

```
/[name='Farm_name'].defaultProxy.proxyaddress:"Proxy_Address"
/[name='Farm_name'].defaultProxy.userName:"User"
/[name='Farm_name'].defaultProxy.password:"Password" /commit:apphost
```

In this example, the command looks like this:

```
C:\Windows\System32\inetsrv\appcmd.exe set config -section:webFarms
/[name='ELMA3'].defaultProxy.enabled:"True"
/[name='ELMA3'].defaultProxy.proxyaddress:"93.158.134.321:8888"
/[name='ELMA3'].defaultProxy.userName:"ELMAAdmin"
/[name='ELMA3'].defaultProxy.password:"12345678" /commit:apphost
```

6.7. Solving Issues with Failover Cluster Installation

Check, if the current machine is a failover cluster node

1. Open SQL Server Configuration Manager (Fig. 15), go to **SQL Server Services**, select the required instance of **SQL Server (MSSQLSERVER)** and open its **Properties** (Fig. 161).

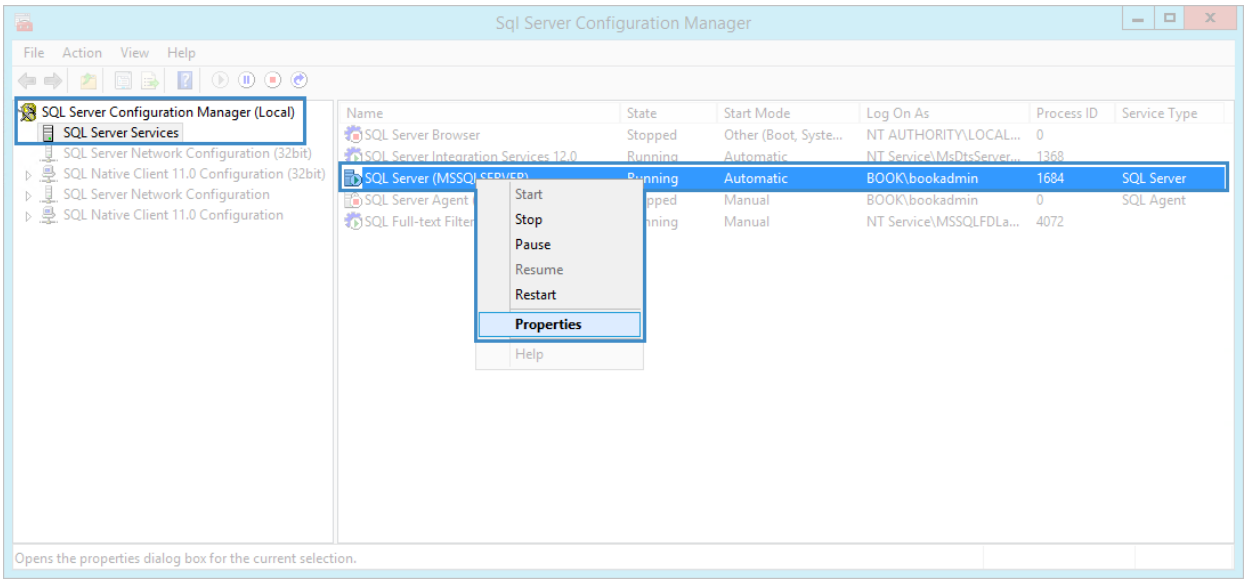

Fig. 161 SQL Server Configuration Manager. SQL Server instance context menu

2. In the opened dialog box, go to the **AlwaysOn High Availability** tab. The **Windows failover cluster name** field (Fig. 162) will display the status of the current machine.

If it is impossible to select the **Enable AlwaysOn Availability Groups** box, close this dialog box and follow the steps described in 3.2.5.

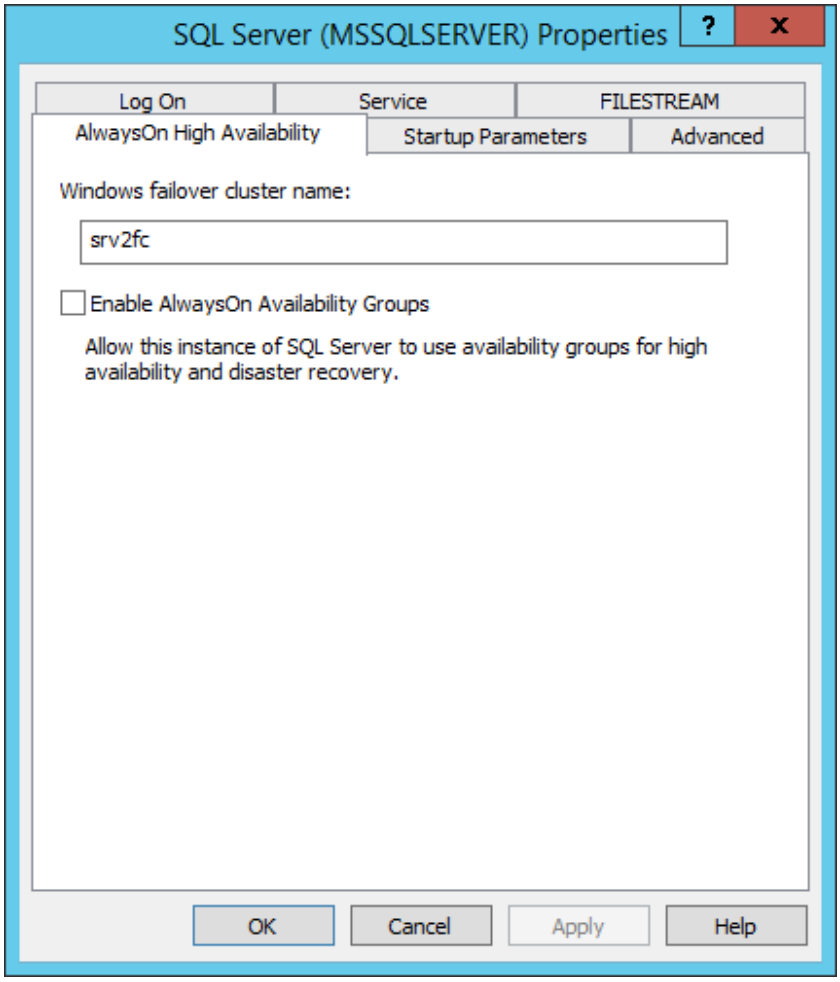

Fig. 162 SQL Server (MSSQLSERVER) Properties

If there is no such user, you need to create it. To do so, open the Server Manager and go to **Configuration – Local Users – Groups**. Open the context menu of the Administrators group and click **Add group...** In the opened dialog box, click **Add...** on the **Member of** tab (Fig. 163).

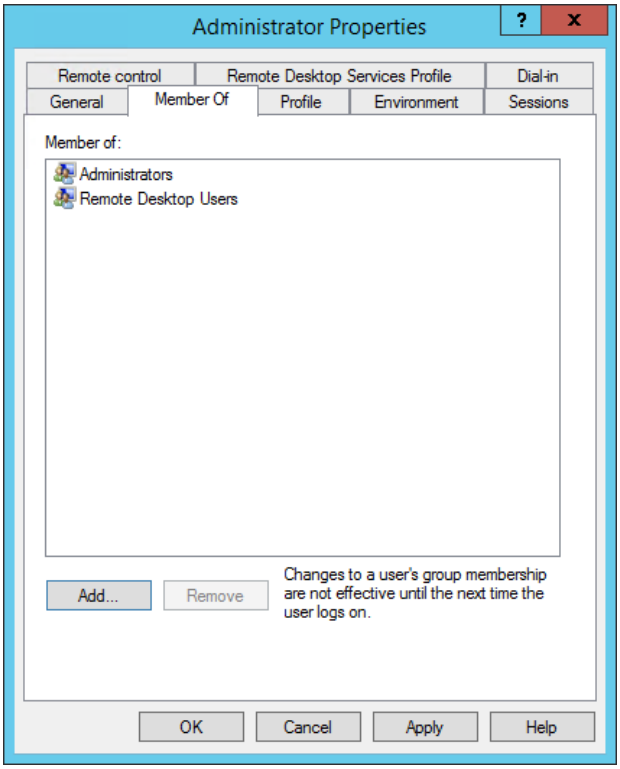

Fig. 163 Administrator Properties

In the opened dialog box (Fig. 164), enter the name of the required user and click **OK**.

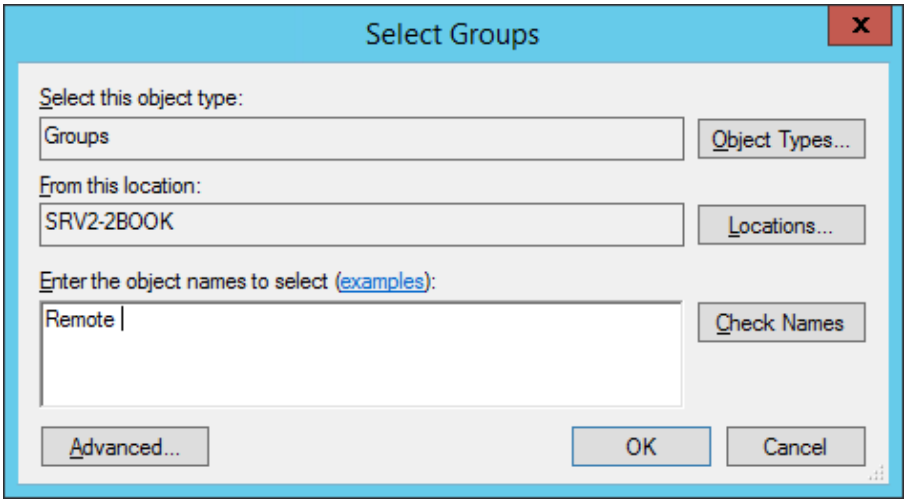

Fig. 164 Selecting a group

6.8. Solving Issues with Creating MS SQL AlwaysOn Availability Group

Issues with creating an availability group are mostly related to incorrect order of actions and not installed components.

6.8.1. The computer is not a failover cluster node

Check, if the current machine is a failover cluster node. To do so:

1. Open the SQL Server configuration manager (Fig. 15), go to **SQL Server Services**, select the required **SQL Server (MSSQLSERVER)** instance and open its **Properties** (Fig. 161).
2. In the opened dialog box (Fig. 162), go to the **AlwaysOn High Availability** tab. The **Windows failover cluster name** field will display the status of the current machine.

If it is impossible to check the box **Enable AlwaysOn availability groups**, and follow the steps described in **Ошибка! Источник ссылки не найден.**.

6.8.2. Error opening the failover cluster manager

Open the Server Manager, go to **Dashboard**, click tools and select **Failover Cluster Manager** (Fig. 23). If an error occurred (Fig. 24), log in to the system as a domain user. Note, that the server must be included in the domain.

If there is no such user, you need to create it or have the system administrator create it.

If you have sufficient permissions in the domain, you can add an account yourself. To learn more, see section 6.7.

6.8.3. Cannot log in to MS SQL Server with a domain account or via single sign on

Create a login in the MS SQL Server Management Studio. To do so:

1. Log in to MS SQL Management Studio under the current account (e.g. **sa**) or the one specified during installation (by default, single sign on under the same user).
2. In **Security – Logins** open the context menu and click **New Login...** (Fig. 12).

3. In the opened dialog box, go to the **General** tab and specify the domain and the name of the required users (Fig. 12).
4. On the **Server Roles** tab (Fig. 14), check all the available boxes and click **OK.**

Not all the permissions may be required when working with the system; however, they may be required during installation, deployment and maintenance configuration.

After that, you can log in to MS SQL Management Studio under the created account.

6.9. Solving Issues with Starting ELMA

Attention! Note, that if the system works properly and during a planned restart (update, configuration import, maintenance) you see an error, try restarting IIS pool.

The causes of most start errors are temporary. For example:

- Designer is running and being used – the configuration file was busy during the start;
- Antivirus or another monitor temporary blocked access to one of the configuration files;
- During the start, a request with blocking was executed in the database.

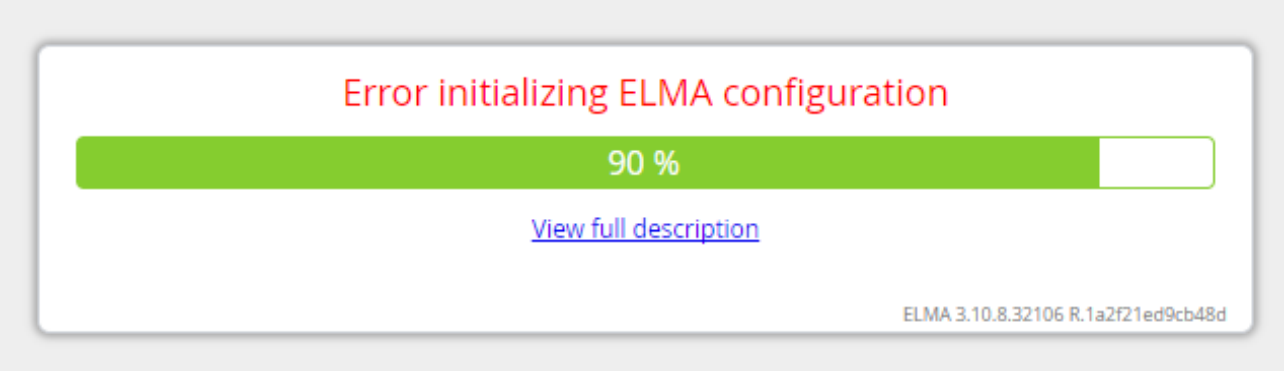

Fig. 165 ELMA Web Application. Start error

In any case, first, open the full description of the error and try restarting ELMA pool in IIS Manager.

6.9.1. No access to the ConfigurationModel.dll file

This error may have different descriptions, stating that there is no access to EleWise.ELMA.DynamicModel.dll, ConfigurationModel.dll libraries.

First, close all the Designers, connected to this server. It is very likely, that something is opened in the Designer, and, therefore, the file is busy.

The second common cause of this error is that files are being scanned by anti-virus software. If possible, stop scanning and try again. Note that scanning may get in the way of starting, if it is active on both the ELMA application servers and the file storage.

Usually, if the error description specifies the file, which could not be accessed, follow these steps:

1. Open Windows Task Manager, go to the **Performance** tab and click **Resource Monitor...** (Fig. 166).

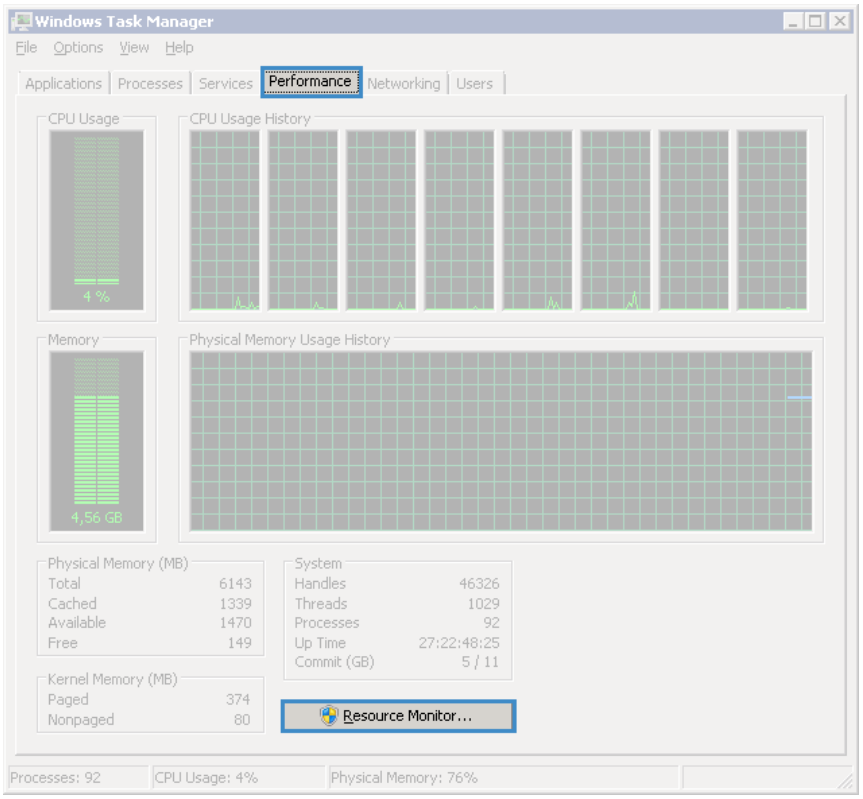

Fig. 166 Windows Task Manager. Performance tab

2. In the opened dialog box, go to the **Disk** tab and expand the **Disk Activity** panel (Fig. 167).

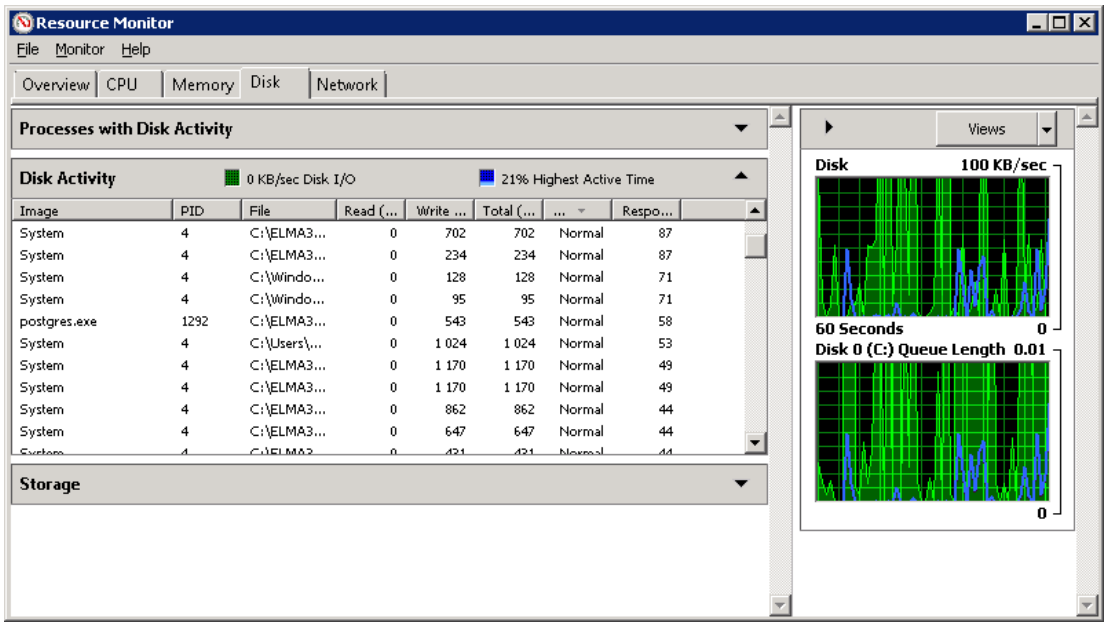

Fig. 167 Resource monitor. Disk tab

3. In the **File** column, find the blocked file and identify, which process is using it.
4. Stop the process – close the applications, stop scanning.
5. Start the ELMA pool in IIS.

6.9.2. Connection with an MS SQL Server mirror instance with the MultiSubnetFailover parameter is not supported

Ошибка при запуске системы

```
*EleWise.ELMA.Runtime.Exceptions.ConfigurationInitializeException: Ошибка инициализации конфигурации ELMA ---> System.InvalidOperationException: Не удалось подключиться к базе данных ---
> System.Data.SqlClient.SqlException: Соединение с зеркальным экземпляром SQL Server с параметром MultiSubnetFailover не поддерживается.
в System.Data.ProviderBase.DbConnectionPool.TryGetConnection(DbConnection owningObject, UInt32 waitForMultipleObjectsTimeout, Boolean allowCreate, Boolean onlyOneCheckConnection,
DbConnectionOptions userOptions, DbConnectionInternal& connection)
в System.Data.ProviderBase.DbConnectionPool.TryGetConnection(DbConnection owningObject, TaskCompletionSource`1 retry, DbConnectionOptions userOptions, DbConnectionInternal& connection)
в System.Data.ProviderBase.DbConnectionFactory.TryGetConnection(DbConnection owningConnection, TaskCompletionSource`1 retry, DbConnectionOptions userOptions, DbConnectionInternal
oldConnection, DbConnectionInternal& connection)
в System.Data.ProviderBase.DbConnectionInternal.TryOpenConnectionInternal(DbConnection outerConnection, DbConnectionFactory connectionFactory, TaskCompletionSource`1 retry,
DbConnectionOptions userOptions)
в System.Data.SqlClient.SqlConnection.TryOpenInner(TaskCompletionSource`1 retry)
в System.Data.SqlClient.SqlConnection.TryOpen(TaskCompletionSource`1 retry)
в System.Data.SqlClient.SqlConnection.Open()
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.qYaTnWRqS3JPCYueuQ(Object )
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.InitInternal()
--- Конец трассировки внутреннего стека исключений ---
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.InitInternal()
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.G2B8hDbHfMk1Zes5rCW(Object )
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.CreateTransformationProvider()
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath, Dictionary`2 extenderParams)
--- Конец трассировки внутреннего стека исключений ---
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath, Dictionary`2 extenderParams)
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath)
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.RegisterComponents(Object componentManager)
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.wxSM8dxaDgDxkIKFaE(Object )
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.StartApplicationInThread()
Ошибка инициализации конфигурации ELMA
EleWise.ELMA.SDK
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath, Dictionary`2 extenderParams)
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath)
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.RegisterComponents(Object componentManager)
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.wxSM8dxaDgDxkIKFaE(Object )
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.StartApplicationInThread()*
```

Fig. 168 ELMA Web Application. Start error

This error (Fig. 168) is caused by incorrect database connection settings; specifically a path to a replica was specified, instead of a failover cluster.

To eliminate this error, follow these steps:

1. Stop ELMA pool in IIS Manager.
2. Open the configuration file (in this case, \\ELMAConfig\ELMAShared\Config\configuration.config).
3. Find the **<connectionStrings>** line and check the database connection string. In this case, connection must be established to the failover cluster, to the availability group listener – SRV12DBLst, not to its replicas: DBCLUSTERTEST01, DBCLUSTERTEST02.
4. Correct the connection string and start the ELMA pool in IIS Manager.

6.9.3. Target database "ELMA3" is included in the availability group and currently is not available for requests

Ошибка при запуске системы

```
*EleWise.ELMA.Runtime.Exceptions.ConfigurationInitializeException: Ошибка инициализации конфигурации ELMA ---> System.InvalidOperationException: Не удалось подключиться к базе данных ---
> System.Data.SqlClient.SqlException: Целевая база данных "ELMA3" участвует в группе доступности и в настоящее время недоступна для запросов. Перемещение данных приостановлено, либо в
реплике доступности не включен доступ для чтения. Чтобы включить доступ только для чтения к этой и другим базам данных в данной группе доступности, включите доступ для чтения к одной или
нескольким вторичным репликам доступности в группе. Дополнительные сведения см. в описании инструкции ALTER AVAILABILITY GROUP в электронной документации по SQL Server.
в System.Data.ProviderBase.DbConnectionPool.TryGetConnection(DbConnection owningObject, UInt32 waitForMultipleObjectsTimeout, Boolean allowCreate, Boolean onlyOneCheckConnection,
DbConnectionOptions userOptions, DbConnectionInternal& connection)
в System.Data.ProviderBase.DbConnectionPool.TryGetConnection(DbConnection owningObject, TaskCompletionSource`1 retry, DbConnectionOptions userOptions, DbConnectionInternal& connection)
в System.Data.ProviderBase.DbConnectionFactory.TryGetConnection(DbConnection owningConnection, TaskCompletionSource`1 retry, DbConnectionOptions userOptions, DbConnectionInternal
oldConnection, DbConnectionInternal& connection)
в System.Data.ProviderBase.DbConnectionInternal.TryOpenConnectionInternal(DbConnection outerConnection, DbConnectionFactory connectionFactory, TaskCompletionSource`1 retry,
DbConnectionOptions userOptions)
в System.Data.SqlClient.SqlConnection.TryOpenInner(TaskCompletionSource`1 retry)
в System.Data.SqlClient.SqlConnection.TryOpen(TaskCompletionSource`1 retry)
в System.Data.SqlClient.SqlConnection.Open()
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.qYaTnWRqS3JPCYueuQ(Object )
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.InitInternal()
--- Конец трассировки внутреннего стека исключений ---
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.InitInternal()
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.G2B8hDbHfMk1Zes5rCW(Object )
в EleWise.ELMA.Extensions.MSSQL.MSSQLProvider.CreateTransformationProvider()
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath, Dictionary`2 extenderParams)
--- Конец трассировки внутреннего стека исключений ---
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath, Dictionary`2 extenderParams)
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath)
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.RegisterComponents(Object componentManager)
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.wxSM8dxaDgDxkIKFaE(Object )
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.StartApplicationInThread()
Ошибка инициализации конфигурации ELMA
EleWise.ELMA.SDK
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath, Dictionary`2 extenderParams)
в EleWise.ELMA.Runtime.RuntimeApplication..ctor(String configurationFileName, Type sessionProviderType, String[] assembiesPath)
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.RegisterComponents(Object componentManager)
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.wxSM8dxaDgDxkIKFaE(Object )
в EleWise.ELMA.BPM.Mvc.Application.Orchard.ELMAEnviromentStarter.StartApplicationInThread()*
```

Fig. 169 ELMA Web Application. Start Error

There may be two causes of this error (Fig. 169).

In the first case, this error is caused by incorrect database connection settings; specifically, the path to the secondary replica is specified, instead of the failover cluster path. To eliminate this error, follow these steps:

1. Stop the ELMA pool in IIS Manager.
2. Open the configuration file (in this case \\ ELMAConfig\ELMAShared\Config\configuration.config).
3. Find the **<connectionStrings>** line and check the database connection string. In this case, the connection must be established to the failover cluster SRV12DBLst, not to its secondary replica DBCLUSTERTEST02.
4. Correct the connection string and start the ELMA pool in IIS Manager.

In the second case, the cause is a malfunction in the failover cluster – at least the primary replica failed. You need to start the cluster recovery mechanism in the automatic mode. To eliminate this error, follow these steps:

1. Stop ELMA pool in IIS Manager.
2. Go to the secondary replica server – DBCLUSTERTEST02.
3. Restart the operating system on the secondary replica DBCLUSTERTEST02.

This solution may seem strange, but it is the simplest and the most effective way to start the automatic recovery of connection between replicas.

There are two more ways to do this:

1. Run **Manual Failover** in MS SQL Management Studio. Unfortunately, if the problem is in the network or the servers, it will not work.
2. Restart the failover cluster replica in **Failover Cluster Manager**. Unfortunately, if the problem is in the MS SQL, it will not work.

Once DBCLUSTERTEST02 is started, you can start the ELMA pool in IIS.

6.9.4. Invalid object name 'ASPState.dbo.ASPStateTempApplications'

This error is caused by the fact that ASPState database was created with incorrect parameters, specifically, with saving the structure in RAM, not as tables.

To eliminate this error, delete the ASPState database and create it again following the instructions in section 3.1.5.

6.9.5. No styles and images in ELMA Web Application

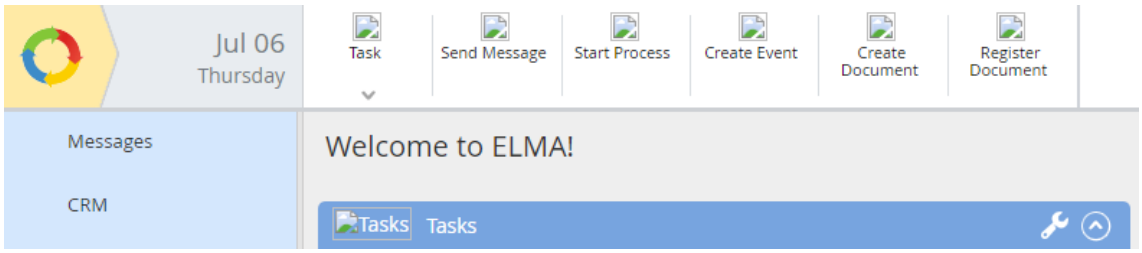

Fig. 170 ELMA Web Application

The underlying reason is usually connected with access to ELMA files. Make sure, that in the settings of the **config** and **Web** folders the Full Access box is checked (Fig. 171) for the account, used for starting the ELMA system.

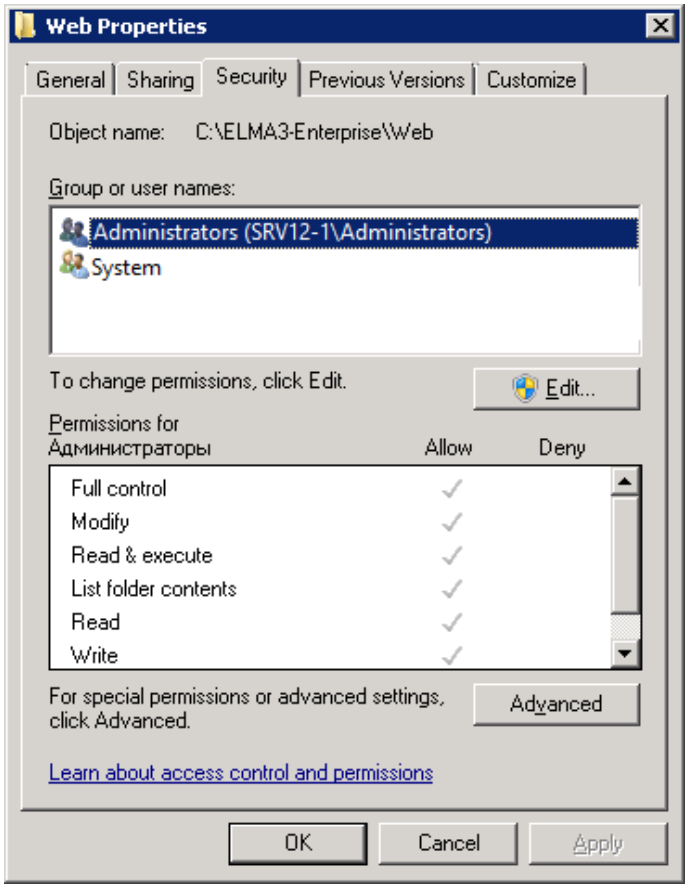

Fig. 171 Folder properties. Security tab

If permissions are configured correctly, configure IIS as described in steps 1 and 2 of section 4.2.2.

6.9.6. Error HTTP 500.21 – Internal Server Error

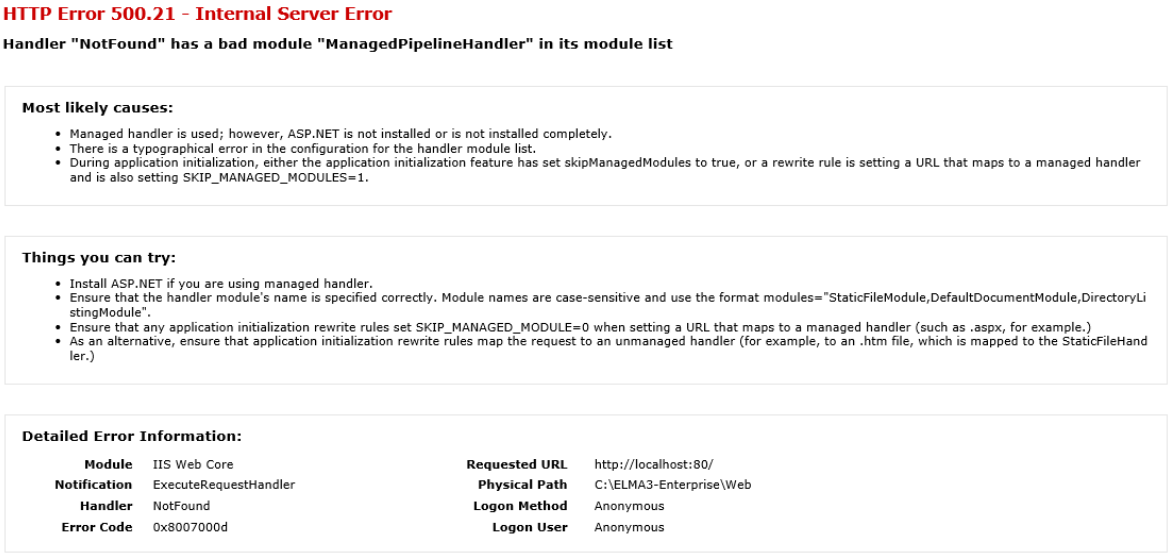

Fig. 172 Error HTTP 500.21

This error (Fig. 172) indicates incorrect configuration of the manageable code in the IIS pool.

1. Open IIS Manager, **Application Pools**.
2. Select the ELMA pool and click **Basic settings**... in the right menu
3. In the opened dialog box (Fig. 173), specify the correct settings.

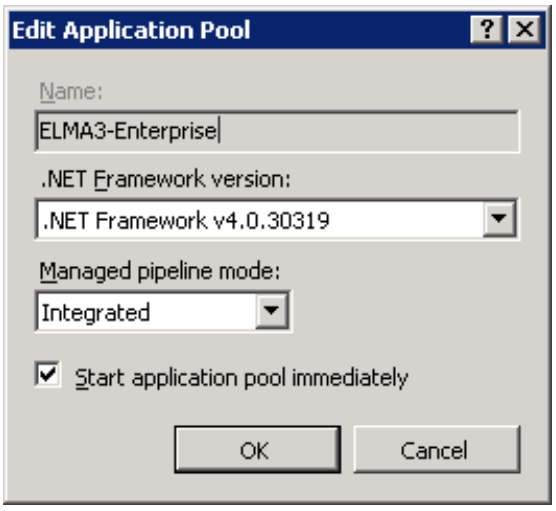

Fig. 173 Edit Application Pool dialog box

The .Net environment version must be 4.0 or higher, Managed pipeline mode – **Integrated**.

6.10. Solving Issues with Working in ELMA

6.10.1. Warm up

There are several widely spread reasons of ELMA working slowly.

First, you should keep in mind, that as soon as the system is started/restarted, pages are opened after compilation, which takes more time. When you open a task page right after restarting, the page is first compiled on the server. When you open the task for the second time or refresh the page, it will be passed instantly and it will take less time.

In case of a farm, pages are opened "for the first time" on each server independently.

It means that when you open a task page on a farm, the request is sent, for example, to the first server. The task page is then compiled on the first server and you will see the page in about 5-45 seconds.

When you try to refresh the page, the request may be sent to the second server, where the operation will be executed in the same way for the first time.

To avoid this process taking time of the users, you can manually warm the system up after the start. On each server, open the most popular pages, such as:

- Main page;
- A task page;
- Task list;
- Process instance;
- Messages;
- Process monitor.

6.10.2. The system or its components work slowly after warm up

There may be many reasons for that. First, you need to configure automatic diagnostics, collect results for a period of user work (1-2 hours of active usage) and ask a question on the technical support website.

In most cases, the cause is not in the software components (MS SQL Server, .NET, AppFabric), but in servers and their configurations.

6.10.3. The system works slowly on virtual machines without apparent issues

It is important to check the real load on the host server and the resource reservation. In the simplest case, it is enough to follow recommendations on reserving resources for virtual machines. Contact the system administrator for performance diagnostics.

Unfortunately, the specificity of virtual machines is so, that you can create a "fast" machine on a problematic disk space and end up with unexpected restarts and freezes, you can reserve resources and at the peak the performance will drop, you can create a VM on a slow host-server and no settings will make it work well.

6.10.4. MS SQL Server response time is 500-501 ms.

This issue occurs on virtual machines under VMWare with certain virtual network interfaces. You can be sure that the reason is this, if the MS SQL Server diagnostics results (Fig. 174) show, that request execution waits for the network (ASYNC_NETWORK_IO – item 3 on Fig. 174).

The MS SQL profiler analysis results show that the time to execute SQL queries is insignificant in comparison to the statistics that include the passing of data to the client. Database processes most queries in 0-1 ms.

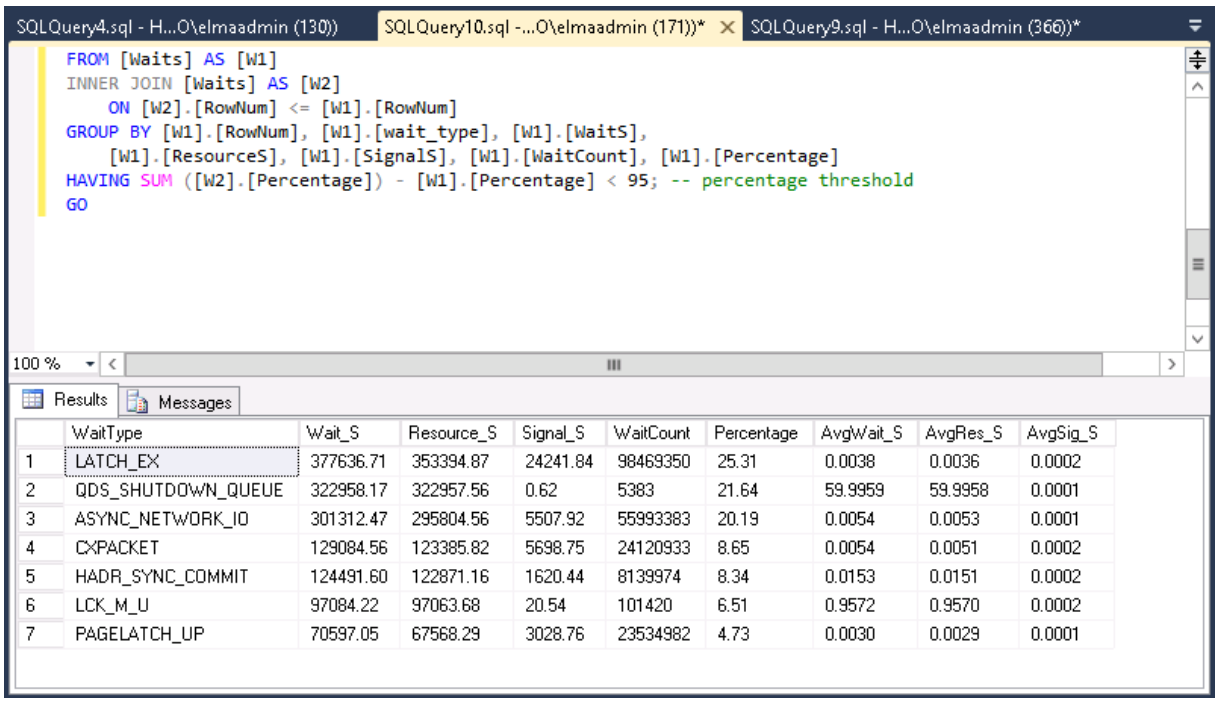

	WaitType	Wait_S	Resource_S	Signal_S	WaitCount	Percentage	AvgWait_S	AvgRes_S	AvgSig_S
1	LATCH_EX	377636.71	353394.87	24241.84	98469350	25.31	0.0038	0.0036	0.0002
2	QDS_SHUTDOWN_QUEUE	322958.17	322957.56	0.62	5383	21.64	59.9959	59.9958	0.0001
3	ASYNC_NETWORK_IO	301312.47	295804.56	5507.92	55993383	20.19	0.0054	0.0053	0.0001
4	CXPACKET	129084.56	123385.82	5698.75	24120933	8.65	0.0054	0.0051	0.0002
5	HADR_SYNC_COMMIT	124491.60	122871.16	1620.44	8139974	8.34	0.0153	0.0151	0.0002
6	LCK_M_U	97084.22	97063.68	20.54	101420	6.51	0.9572	0.9570	0.0002
7	PAGELATCH_UP	70597.05	67568.29	3028.76	23534982	4.73	0.0030	0.0029	0.0001

Fig. 174 MS SQL Server diagnostics results

The issue occurs under the following conditions:

- VMWare 6.0 or higher;
- vmxnet3 Ethernet adapter network interfaces;
- slowdown is in MS SQL and it does not affect large file transfer speed;
- slowdown of 500 ms for each request.

This issue is described on the vmware website. There is also a discussion and solution for it.

To bypass this issue, follow the recommendations on all the virtual machines that interact with the MS SQL database.

Application servers:

- 93.158.134.3 – SRV12-1
- 213.180.204.3 – SRV12-2

Database servers:

- 192.168.18.230 – main replica
- 192.168.18.23 – secondary replica

Here are some specific recommendations:

1. In the network settings on the servers, open the **Network** tab (Fig. 175), click **Configure**... and go to the **vmxnet3 Ethernet Adapter** parameters (Fig. 176).

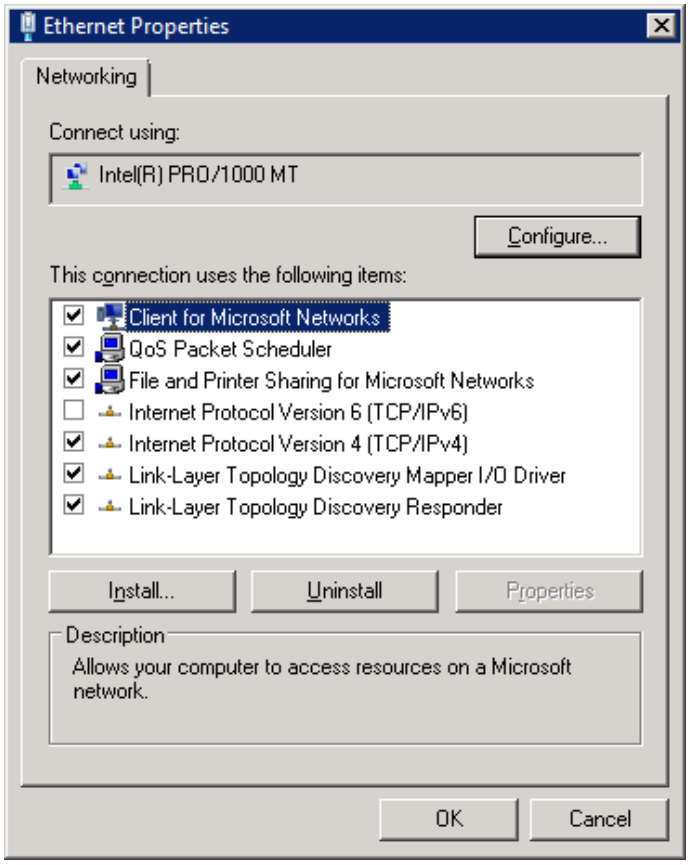

Fig. 175 Network connection settings

In the opened dialog box, on the **Advanced** tab (Fig. 176), set the Disabled value for the following properties:

- IPv4: Checksum Offload
- IPv4: TSO offload
- Large Send Offload
- Offload TCP Options
- Offload tagged traffic
- TCP Checksum offload

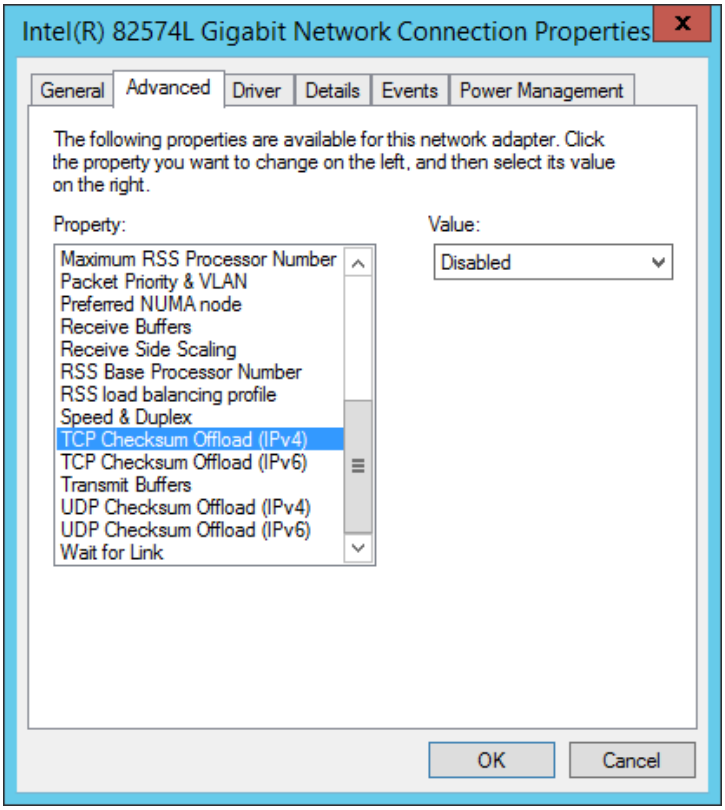

Fig. 176 Network connection settings

It is also recommended that you disable the IPv6 protocol for these servers.

6.10.5. FastReport reports do not work

If the FastReport reports do not work, read this knowledge base article and follow the instructions. Do not forget to stop the ELMA farm and make changes to the configuration file in the shared folder (see section 4.5).

6.10.6. Everything is OK, but requests are executed slowly

Make sure that there is free disk space on the controller server.

In this case, you will not see any obvious reasons for performance drop, since ELMA servers will work at normal speed, but the users will be receiving responses with a long delay.

You may have run out of disk space because of the enabled IIS logging. Clear the disk of unnecessary files and disable IIS logging as described in section 4.6.

Chapter 7. Additional information

7.1. Migrating user session storage to Redis

If there is a significant number of users working simultaneously and intensively in the system (more than 1000) and if the database report shows that there is an excessive load on the ActiveUserSessions table, it becomes necessary to use the network data storage with user sessions in order to make ELMA work properly.

This Chapter deals with the configuration of the Redis network data storage.

Step 1. It is necessary to start three more new Redis and Sentinel servers that are not connected to the Redis cache. Installing Redis is described above.

Attention! If Redis has already been installed, start the installation with the sudo ./install_server.sh command. Please note that during installation you need to change the default values to new ones (e.g. port 6381, configuration file 6381.config, service redis_6381).

Step 2. Copy the Redis libraries to the **C:\ELMA3-Enterprise\Web\bin\folder**.

Step 3. In the configuration file **C:\ELMA3-Enterprise\Web\Web.config**, the following actions should be performed:

1. Change the line `<sessionState mode="InProc" timeout="30"/>` for the lines:

```
<sessionState mode="Custom" customProvider="MySessionStateStore">
  <providers>
    <add type = "Microsoft.Web.Redis.RedisSessionStateProvider"
      name = "MySessionStateStore"
 connectionString = "machine1:6379, password=
MyPassword,abortConnect=False"/>
  </providers>
</sessionState>
```

   machine1:6379 is the server address and port, **MyPassword** is the server password.
2. Comment out the blocks of the lines, as follows:

```
<!--<remove name="Session" />-->
```

```
<!--<add name="Session"
type="EleWise.ELMA.Web.Mvc.Modules.ELMASessionStateModule,
EleWise.ELMA.SDK.Web" />-->
```

Note that there are several such blocks in the configuration file.

To learn more about migrating user session storage to Redis, see the respective Knowledge Base article. Redis lock settings

You can use Redis networked data storage to store locks.

Step 1. It is necessary to start three more non-synchronized Redis servers without Sentinel that are not connected with those that were previously configured. Redis installation is described above.

Step 2. Make the following changes to the configuration file **C:\ELMA3-Enterprise\UserConfig\configuration.config**:

1. Add the lines to the **<configSections>** section:

```
<section name="lock"
type="EleWise.ELMA.Configuration.GenericProviderFeatureSection`1[[EleWis
e.ELMA.Locking.LockServiceManager, EleWise.ELMA.SDK]],
EleWise.ELMA.SDK"/>
<section name="RedisDistributedLockServiceSettings"
type="EleWise.ELMA.DistributedLock.Redis.RedisDistributedLockServiceSett
ings, EleWise.ELMA.DistributedLock.Redis"/>
```

2. Add the lines at the end of the configuration file:

```
<lock defaultProvider="redis">
   <providers>
      <clear/>
      <add name="redis"
type="EleWise.ELMA.DistributedLock.Redis.RedisDistributedLockService,
EleWise.ELMA.DistributedLock.Redis"/>
   </providers>
</lock>
<RedisDistributedLockServiceSettings>
   <connections>
      <add
connection="machine1:6379,password=your_redis_password"/>
      <add
connection="machine2:6379,password=your_redis_password"/>
      <add
connection="machine3:6379,password=your_redis_password"/>
```

```
    </connections>
</RedisDistributedLockServiceSettings>
```

machine1:6379, machine2:6379 and **machine3:6379** are the addresses and ports of the servers, **your_redis_password** are the passwords of the respective servers.

To learn more about migrating user session storage to Redis, see the respective Knowledge Base article.

Chapter 8. Useful References

Along with this user manual, the following sources describe the functions of **ELMA** applications:

- ELMA BPM Platform user manual
- ELMA Web Portal user manual
- ELMA ECM+ user manual
- ELMA Projects+ user manual
- ELMA KPI user manual
- ELMA Administrator user manual

These manuals walk you through the key features of the system. You can find a more detailed description of ELMA functions in the system Help, available via the following link: https://kb.elma-bpm.com/help.

General description and purchase conditions of the applications are available on ELMA website: http://www.elma-bpm.com. You can also **Ask a question** on this website, using a respective link.

An **Online Demo** http://demo.elma-bpm.com demonstrates the main functions and utilization of the applications.

We continuously develop **ELMA** and Platform-based components for coping with more specific tasks. You can find the list of these components and their purchase conditions at **ELMA Store**: http://store.elma-bpm.com/.

If you are experiencing technical difficulties, please visit ELMA technical support website: http://support.elma-bpm.com.

If you need assistance with the system or have questions about partnership with **ELMA** Company, contact us:

- Luxemburg: + (352) 20-30-11-40
- http://www.elma-bpm.com/contacts/