BatPost

Installation and primary server setup

Contents

Chapter 1

1 Introduction

This user guide will help you install the BatPost mail server and configure it from scratch. Even if you lack the knowledge on how to configure a mail server the step by step instructions will help you get the mail server up and running in its default configuration.

Sooner or later you will need to make some special settings and by that time you will be able to make brave experiments concentrating on the smallest details instead of trying to grasp everything at once. Making changes to a small number of settings will always allow you to revert to what you initially had, just in case something goes wrong.

If you didn't have a chance to configure a mail server in the past it might turn out to be a tough job to do. In the process a lot of questions arise and sometimes it is very hard to find the answers. This guide provides you with the basic information about e-mail needed to understand the key principles of the way a mail server works.

A mail server can be used in various environments. The internet connection can be permanent, periodic (via the telephone line) or even not available at all. That is why before installing the server you have to make decisions on several important issues so that afterwards you do not have to spend the time on reconfiguring the mail server.

In the FAQ section you will find advices on how to solve the most common issues.

Chapter 2

2 Basic information about e-mail

E-mail was one of the first services available on the internet. At that time messages were text only and it has been an amazing possibility to send them across the globe in just a few seconds. It was much more convenient than telegrams, usual letters are not even worth mentioning. At the same time it became clear that simple text messages were not enough. And so message format has been improved several times allowing the user to decorate the text using different fonts, size, etc, to paste pictures into the message body, attach files, make use of the message encryption and digital signature.

It all became possible thanks to the development of the e-mail clients. Primitive text editors with the mere ability to send and receive messages evolved into powerful programmes with lots of handy features. Now it's possible to create colourful HTML-messages, use templates to speed up the work process, automatically sort incoming messages into different folders and make use of the feature-rich address books. But what happens when we press the "Send" or "Receive" button in the e-mail client? And here the mail server steps into play.

An e-mail client knows nothing of how to deliver a message to the recipient. Its main goal is to correctly compose a message, put in the right sender's and recipient's e-mail addresses, and then hand over that message to the mail server. Incoming messages an e-mail client also receives from a mail server. That's on one hand, on the other one is the mail server, which does not care at all how a message is laid out, but it knows how to deliver it to the respective recipient. In a sense it resembles a usual post-office and people coming in. A post-office stands for the mail server and people represent e-mail clients.

Naturally you would ask why do we need a mail server at all? Can't the e-mail clients send a message directly to the recipient? After all, internet, unlike usual mail, doesn't have any frontiers and transmitting data to another part of the world is as easy as though it were the room next door. So, just imagine you want to send a message and the recipient has already gone to bed or simply shut down the computer. Or someone sent a message, but at that time you were on vacation. It's hard to imagine that the sender will keep on waiting to send you the message until you are back. It's much more convenient to transmit a message to the mail server and it will make sure the message is delivered. Moreover, the message will not be delivered directly to the recipient, but to the respective mail server that will put the e-mail into the user's inbox folder, where it will reside until the user checks his/hers incoming mail.

But how can one know to which server should the message be transmitted to so the recipient finally gets it? The answer to that question lies on the recipient's e-mail address. It consists of two parts split by the "@" sign. What you have to the right of that sign is called domain name and what you have to the left of that sign is the user name within that domain. In order to deliver a message to its recipient it is necessary to determine which server is responsible for that domain's mail. That is why the so called DNS (Domain Name System) is used, it keeps records of each mail domain and the list of servers attached to them.

2.1 E-mail client

Modern e-mail clients offer a lot of useful features, however from the electronic mail's point of view it is required that an e-mail client is able to send and receive mail. Depending on the mail server type and the quality of the internet-connection we can distinguish several ways of how e-mail is handled:

1. Connect to the Internet, download messages, disconnect.
   Read the incoming mail, compose replies to messages, but instead of sending them out right away place them into the Outbox folder.
   Connect to the Internet, send outgoing mail, eventually download new incoming mail, disconnect.
   This is characteristic to circumstances when there is no permanent connection to the Internet, e.g. via a modem. In this case the POP3 protocol is mainly used, though IMAP4 is also no wonder here.
2. Connect to the server, download mail, disconnect.
   Read the incoming mail, compose replies to messages and send them out immediately.
   This is common to permanent connection to the Internet. Both POP3 and IMAP4 protocols can be used to retrieve mail.
3. Permanent connection to the server.
   All the mail is stored on the server and can be downloaded on demand. The server will notify the user when new messages arrive on the server.
   In this case a broadband connection to the Internet (e.g. local network) is required. The use of the IMAP4 protocol is also mandatory.

Each of the methods described above has its advantages and disadvantages. For instance, the first one is the most undemanding to the quality of the Internet connection, but is bound to a PC as all the mail is stored there. The third method allows accessing the same account from different PCs, but it requires a broadband connection to the Internet, it constantly uses server resources and is therefore dependent on the server's capabilities and current load.

2.2 Mail server

A mail server does not have a user interface, it's a set of protocols which allow interacting with the server.

POP3
Its purpose is to check the incoming mail. Though a user might have many folders on the server, they are not checked, the only active folder is the INBOX one. The protocol is very simple and is supported by all e-mail clients. It allows retrieving a list of messages available on the server, you can see their size, either download, delete or leave on the server certain message(s). It is also possible to partially download messages, in other words the message header and a certain number of text lines.

IMAP
It also allows checking incoming mail, but offers a lot more additional features. It appeared mainly as the evolution of the POP3 protocol, almost all of the functions are shifted from the client to the server. This allows using the so called "thin" clients, that are primitive applications and all the work is done by the server.

The user can have quite a rich folder structure where the messages are stored. The server splits e-mails into several parts and you can gain access to any of it. For instance, you can download the attachment only. Unlike POP3, this protocol allows processing more than one message at once requesting only the needed information from the server. For example, you can request senders, recipients and messages' size for e-mails in the list from 5th to 10th inclusively.
It is possible to search for messages with quite narrow conditions, e.g. find all the messages from a certain sender for the last month.
Besides the protocol standards a lot of additional extensions were adopted increasing the protocol's capabilities even more. It all made IMAP4 very complicated and sophisticated. Not all the e-mail clients and servers support it, and what's more none of them really support all the protocol's features.

SMTP
It is used to send messages from users to the server and also for sending messages from one server to another. At the same time the sender's addresses and the addresses of the recipients are indicated separately for each of them. This allows returning an error message at any stage of the process. For instance, the server might refuse to receive a message from a given sender if such a sender address does not exist or if a recipient's e-mail address is not found.
Primarily, the SMTP protocol did not perform authentication, but the ever increasing amount of spam e-mails (unsolicited messages with mainly advertising content) made that feature necessary.
When messages are sent to the external addresses the SMTP-authentication is often required, otherwise there would be an good possibility for the delinquents to misuse the server by sending messages via that mail server (open relay). Instead of authentication the server can impose restrictions based on IP-addresses. For example, sending messages to external addresses is only allowed from within the internal subnet.
The mail exchange between the mail servers is usually carried out without authentication as in most cases different servers know nothing of each other. The only restriction is that the recipient-server has to be the point of destination (at least from the sender-server's point of view). There are exceptions when, for example, a small company makes use of the ISP's server as an intermediate server for sending out mail. In this case it is necessary to use some sort of authentication or impose some restrictions (e.g. base on IP-addresses).

2.3 DNS

An e-mail address provides a server with all the necessary information on the recipient and how the message should be delivered. For instance, a server needs to deliver a message to the following e-mail address: user@example.com. In this case "example.com" is the mail domain and "user" is a username within that domain. First of all, it is necessary to define which server(s) are responsible for mail directed to the "example.com" domain and then initialize a connection to it via the SMTP protocol, after that it's up to the respective server to deliver the message to the user. In order to find the responsible server the DNS (Domain Name System) is used. The DNS-server keeps different sets of records for domain names and returns the respective records for certain domains upon request. In our case (in terms of e-mail) we are interested in records of the MX and A types.

MX records

The MX record contains number which determine the priority of a record followed by host name responsible for the certain domain's mail. Host with lower priority value should be used prior to the others. Hosts with equal priority can be used randomly.
If the server encounters it's own name in the list of the hosts responsible for that domain's mail, it has to establish connections only with hosts possessing lower priority than it's own. This is done on purpose to avoid situations in which an e-mail might stuck within a circle of servers forever.
In our case, for example, the following MX records might be received:

```
example.com  MX  10 mx1.example.com
example.com  MX  20 mx2.example.com
example.com  MX  20 mx3.example.com
```

In this case the server has to try to pass over the message to "mx1.example.com" host first, and in case of a failure randomly to the "mx2.example.com" and "mx3.example.com" hosts.

A records

The MX records provide us with host names only, whereas the A records are used to determine its IP-address. The A record contains the IP-address for a given domain. A single domain can possess several A records, in other words a domain can have several IP-addresses. While sending out mail we have to try to connect to each IP-address (in any order)
In our case we might receive the following A records for the "mx1.example.com" host:

```
mx1.example.com  A  192.5.19.31
mx1.example.com  A  192.5.25.5
```

What we should do next is try to connect to the 192.5.19.31 and 192.5.25.5 IP-addresses in any order and in case it fails move to the next MX record.

Such record structure allows implementing the following:

- Server cluster. Large-scale mail portals might use several servers simultaneously in order to improve its performance and stability, in case a particular server goes down the rest continue to function properly. This procedure also allocates workload between them thus improving the overall server performance.
  Such structure can be realized either with the help of several MX (with equal priorities) or A records. Since the mail servers process records in the order they had been received from the DNS-server it is highly important they be returned in a random order, thus the load will be divided evenly.
- Reserving servers. In such structures there is a main server and backup servers which can receive mail if the main server is not responding. In order to implement that several MX records with different priorities are created. The main server possess the lowest priority value (it will the first one to be connected to) while the backup servers have a higher priority value (they will be connected to if the main server is not responding).
  In our example "mx1.example.com" represents the main server while "mx2.example.com" and "mx3.example.com" are backup servers.

Important. The latter scheme can also be used in case the company decides to maintain the mail domain (formerly maintained by the ISP server) by itself. It is necessary to add the MX record referring to a server within the company. Note that its priority value has to be lower (main server) than the one of the ISP server (backup server).
This will allow setting up the server within a company much easier, as a last resort the

mail will still be received by the ISP's server. When the internal server is finally up and running, you can remove the MX record from the ISP's server (or leave it as it is just to be on the safe side).

Chapter 3

3 Installing the BatPost server

In this section we take a look at the server's installation sequence and its primary set up. The step by step instructions will help you get the mail server up and running in its default configuration. Besides that you will find tips on further server settings.

3.1 Before installing

Since the server is able to operate in several modes it is highly recommended that before its actual installation and primary set up you decide how exactly it will be used. You can obtain some of the technical data (the IP-address of the DNS, provider's mail server settings) by contacting your ISP.

The mail servers widely use the Domain Name System (DNS) for delivering e-mail messages. The DNS allows the server to decide which mail server is responsible for receiving e-mails for a specific domain (using the MX records). Usually there are several servers and each of them has its own priority. Such method makes it possible to considerably increase the reliability of the mail subsystem on the whole, because if a certain server is temporarily unavailable another one can successfully replace it. For even better flexibility in the MX entries instead of the IP-addresses of the mail servers the names of their hosts are used. In order to determine the IP-address by the host names the DNS is used too (the A records).

First of all you have to decide whether the computer on which you plan to install the server will be permanently connected to the internet or will go on-line periodically. This issue is very important since the DNS access often requires permanent connection to the Internet. Sometimes the server has no connection to the Internet, but nevertheless has access to DNS. It is possible if the server is used exclusively for internal needs in the local network, e.g. in order to organize the file exchange within a company. However, in this case there is no need in the DNS access as all operations are carried out within a single server.

Therefore the server can be set up in three different ways:
- permanently on-line with access to DNS
- periodically going on-line using the dial-up connection
- working within the local network without ever connecting to the Internet

The BatPost mail server is flexible enough to be set up in these three different ways and successfully manage the e-mails. All you need is to change several server settings. Our next step is to have a closer look at each of the three ways of setting up the server and providing you with some recommendations.

Permanently on-line
If the connection to Internet is permanent it is preferable that the server distributes the e-mail on its own using the DNS. In this case you have to find out the IP-address of the DNS server. Such data can be provided by your system administrator or your ISP. However, in case of the self-reliant delivery it might be necessary to create your own mail domain, because some servers perform extensive verification and refuse to receive mail from servers not registered in the DNS. If you do not intend to register your own domain then you can distribute mail via the provider's server or via any other relay server.

Periodically going on-line
If the server connects to Internet for short periods of time then the usage of DNS and direct mail distribution becomes ineffective. In this case it is more convenient to use the provider's server both for receiving the incoming mail and sending the outgoing messages.
For retrieving the incoming mail either the Remote POP (RPOP) or ETRN/ATRN technologies can be used. In the first case BatPost will connect to the external server via the POP3 protocol and fetch messages. Then it will analyze the messages' headers to determine their recipients. In the second case the special ETRN/ATRN protocol will be used in order to get messages from the external server's sending queue. The latter one is preferred to the first one, because in that case the message recipients are unambiguously identified, though it is mandatory the external server is specially set up which is not always possible.
For sending the outgoing mail the ISP's server will be used as the relay server. Configuring the server that way is more difficult as there are many settings that have to be taken into account. Anyway, you will have to discuss all the details with your provider and receive all the necessary information on the settings.

No connection to Internet
If there is no connection to Internet at all then the internal DNS can be used. It will allow having several mail servers in the local network. Note that the DNS records are local network specific and are of no practical use outside it. In most cases the DNS is also absent and all the operations with messages are performed within one mail server. In this case, though in reality messages will not be distributed at all, you have to define the relay server as the method of distributing the messages. Relay server settings in this case are inessential.

Server installation
As soon as you decide the way the server will be functioning and as soon as you have all the necessary information on the settings, you can proceed to the installation of the server. Here are the steps you will go through installing BatPost:
1. Installing the BatPost server
2. Primary server configuration
3. Creating global administrators
4. Creating domains
5. Creating users
6. Testing server configuration

In case the server will periodically go on-line check the advanced settings.

3.2 Server installation

The installation file of the server is BatPost_v221xx.exe. During the installation you will be prompted for the directory where all the files will be copied to and the program group where the icons will appear. To make a typical installation all you need is press "Next", thus applying the default programme settings. Alongside the typical installation it is possible to perform a custom installation and choose which components you would like to install. The server components are required by the server itself to be able to operate, while the administrative ones are designed to configure and manage the server. If you

intend to only remotely control a server from a given PC, then you only have to install the administrative components on that computer.

If you do not change the parameters during the installation there will appear the "BatPost" group which will contain the following icons: the Configurator, the Monitor and uninstaller for removing the server from your computer.

3.3 Primary server configuration

After the successful installation the Startup Wizard will be launched. It will help you set the basic parameters required for normal server functioning.

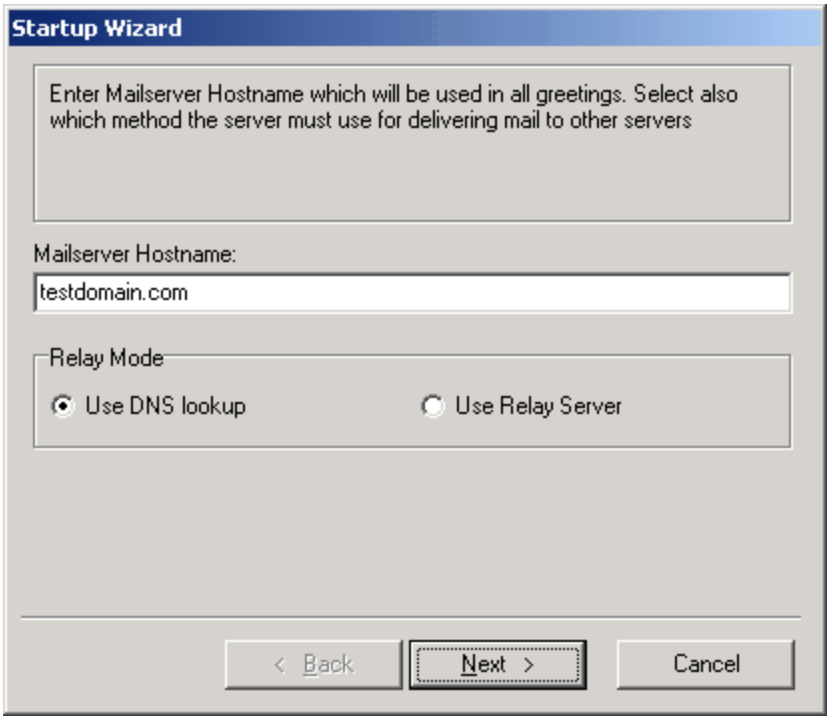

In the first window of the Startup Wizard it is necessary to enter the server name which will appear in the network; this name usually coincides with the name of the main mail server domain. You also have to select the method for delivering the outgoing mail. In the "Before install" section you can find more on selecting which method to use.

In case the DNS for delivering the mail will be used, the following window will appear that will let you define the DNS servers:

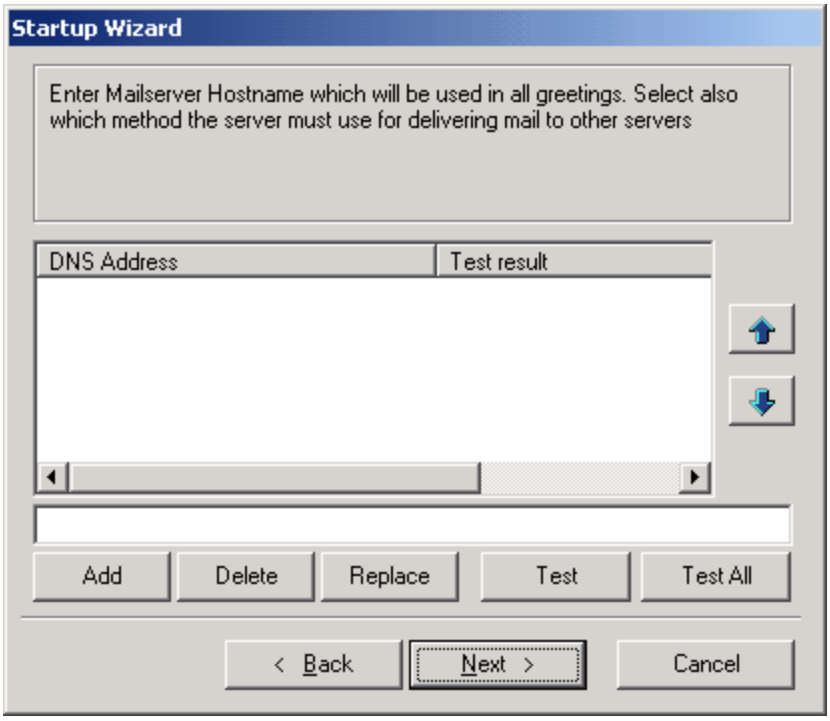

BatPost allows you entering several DNS servers and will try to use the fastest one. Here you can also test the defined DNS servers.

If you choose the relay server for delivering the outgoing mail then the following window will appear, where you have to enter the corresponding data:

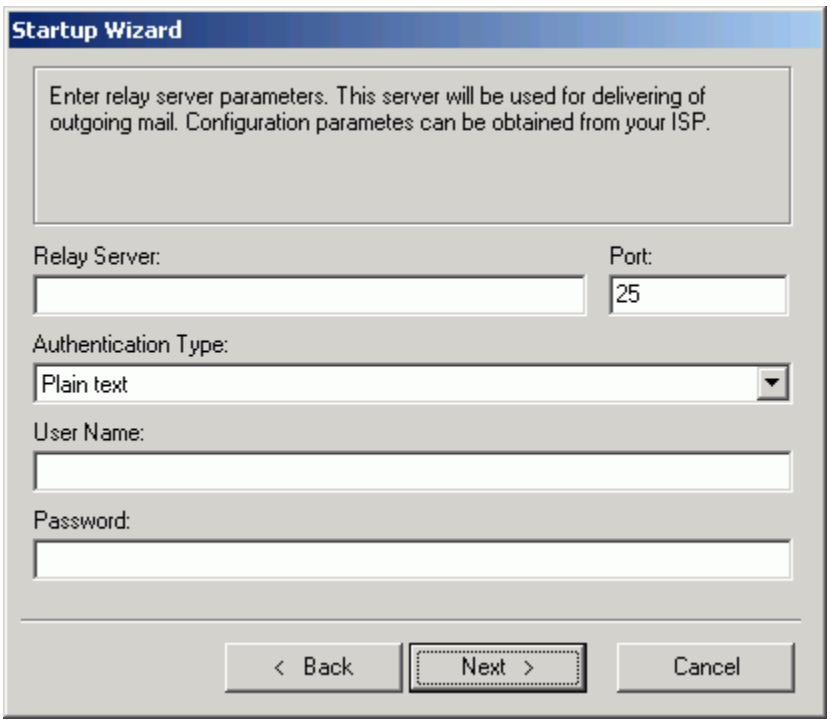

In this window you can enter either the Hostname or the IP-address of the relay server, the Port number for the outgoing SMTP sessions, authentication type and its parameters. If the authentication is not required you can leave the user name and password fields blank.

After pressing the "Next" button the last window of the Startup Wizard will appear:

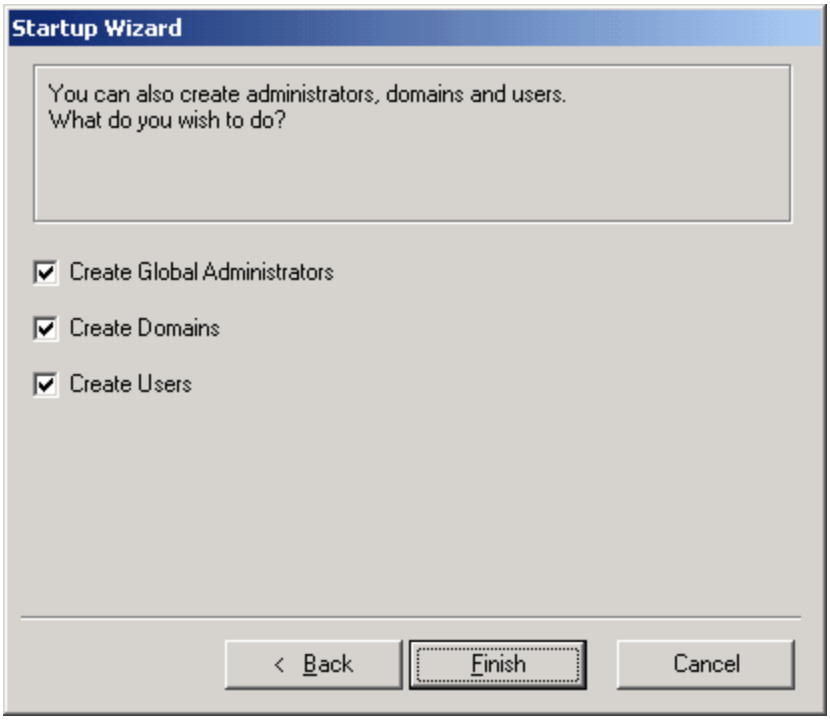

Here you can select whether global administrators, domains or users should be created.

3.4 Creating global administrators

Global administrators are entitled to monitor the server and change any of its settings. These operations can be carried out not only from the local machine, but from the others as well. If you intend to work with the server remotely you will have to create at least one global administrator.

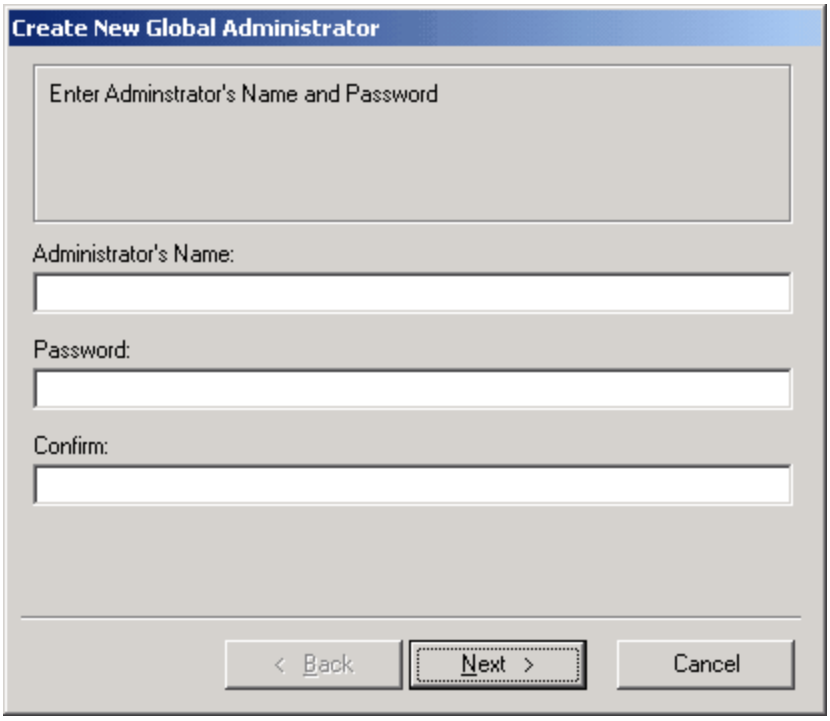

For the global administrator it is necessary to enter only the name and the password.

The next window will either terminate the Wizard or will continue creating global administrators:

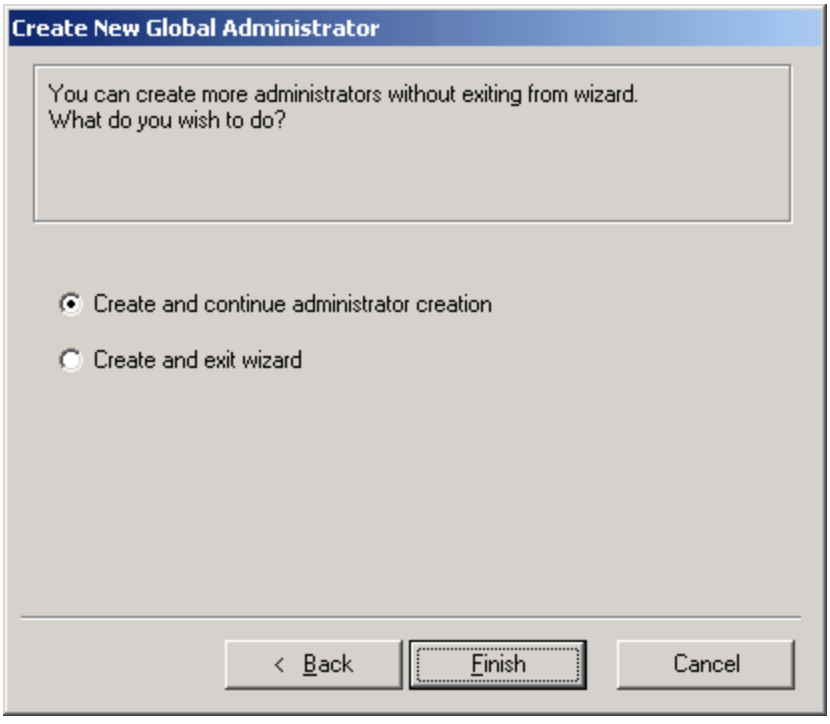

3.5 Creating domains

For a mail server the domain is an essence that joins together many users. From the Internet point of view a domain is a symbolic resources address in the network. In our case this duality means that the mail server has to be correctly configured in the DNS. In other words for this domain name there has to be one or more MX record. For every host present in the MX records there has to be the A record, which will allow finding out its IP-address.

On the server the domain name must coincide to the one entered in DNS.

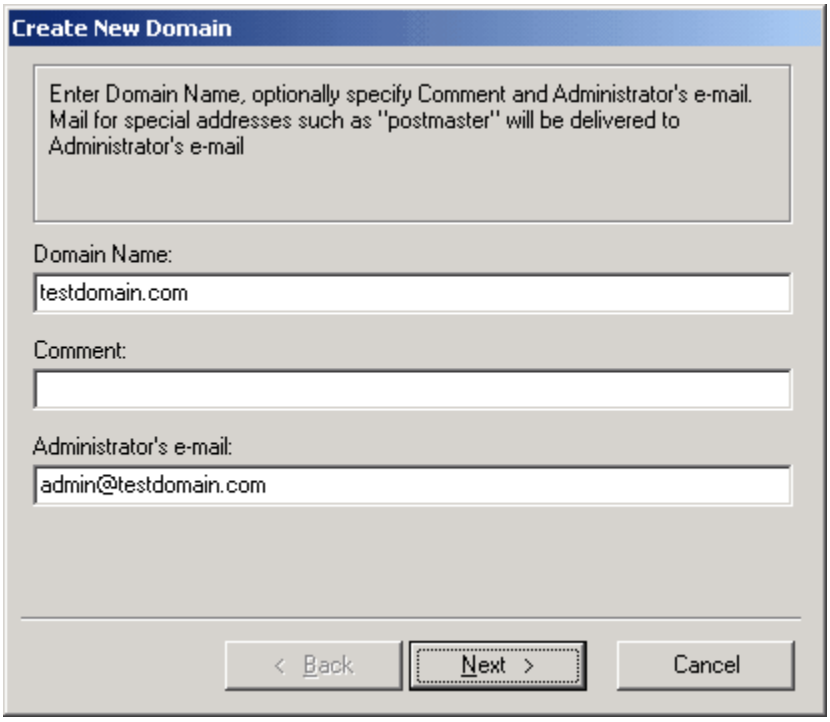

As you create a domain you have to enter its name, optional comment and the administrator's e-mail address. There are common standards for creating mail subsystems that require the presence of some special e-mail addresses such as POSTMASTER. It is possible to create users with the necessary names on the server, but usually it is enough to redirect to the administrator messages that arrive at special e-mail addresses.

The next window will either terminate the Wizard or will continue creating domains:

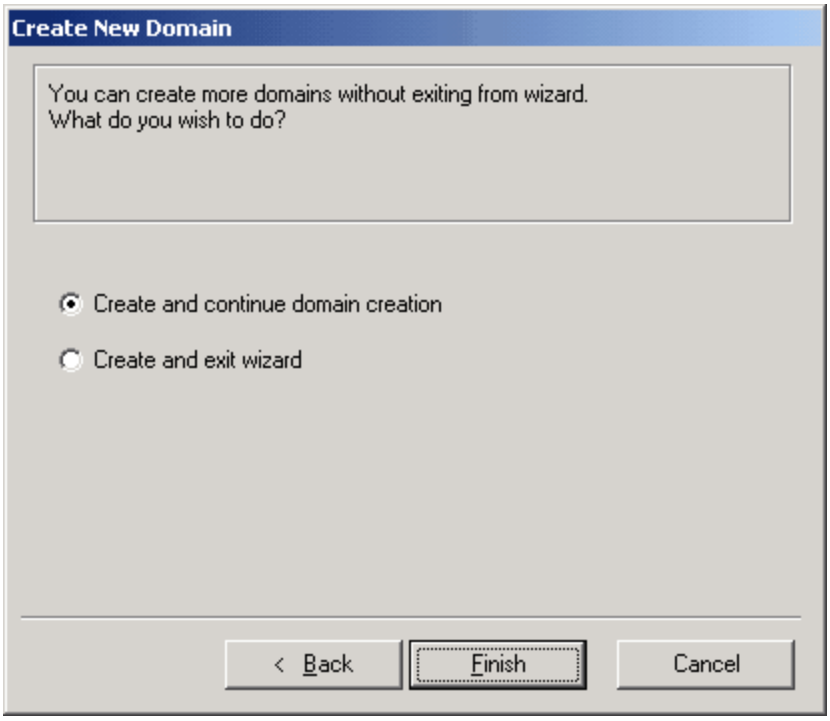

3.6 Creating users

The BatPost mail server supports several types of users or also known as accounts. Messages arriving at different accounts are processed differently. In this section we will take a closer look at "ordinary" users whose messages are stored in their personal mailboxes.

When creating a new user first of all you will be asked to specify the domain the user will belong to:

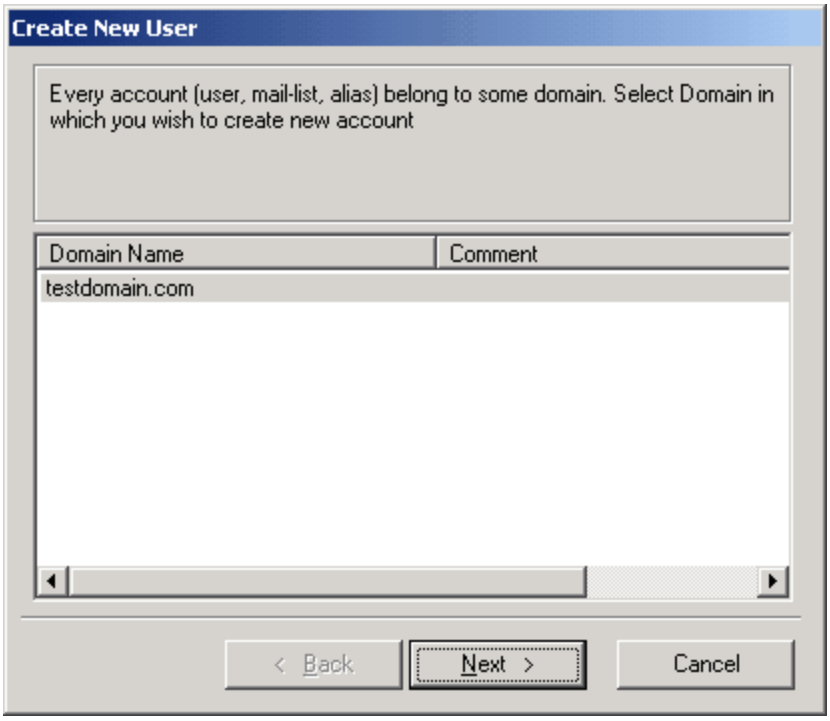

Then a window where all the basic user parameters can be entered will appear:

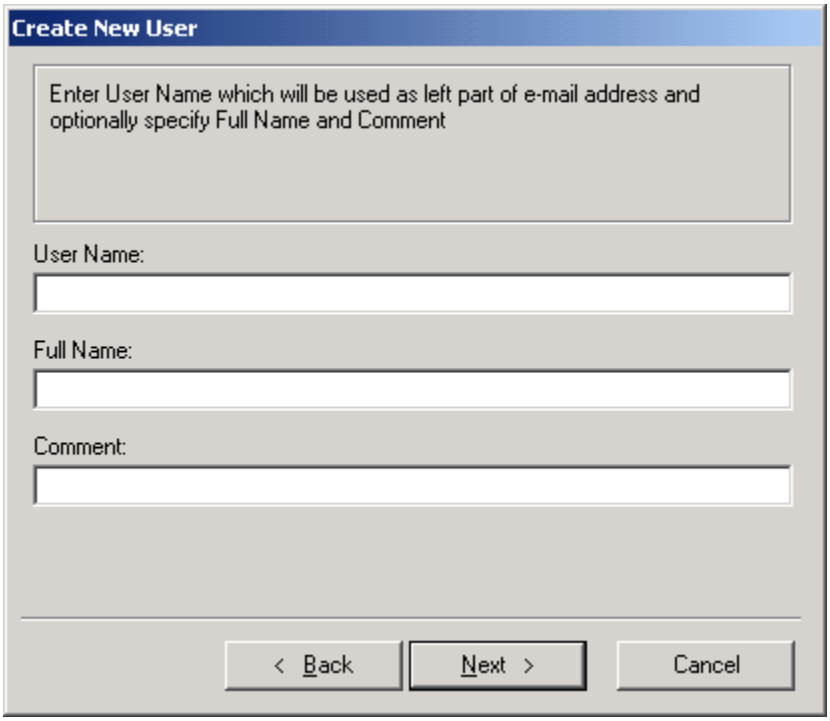

You must enter the user name on the server and optionally his/hers full name and comment.

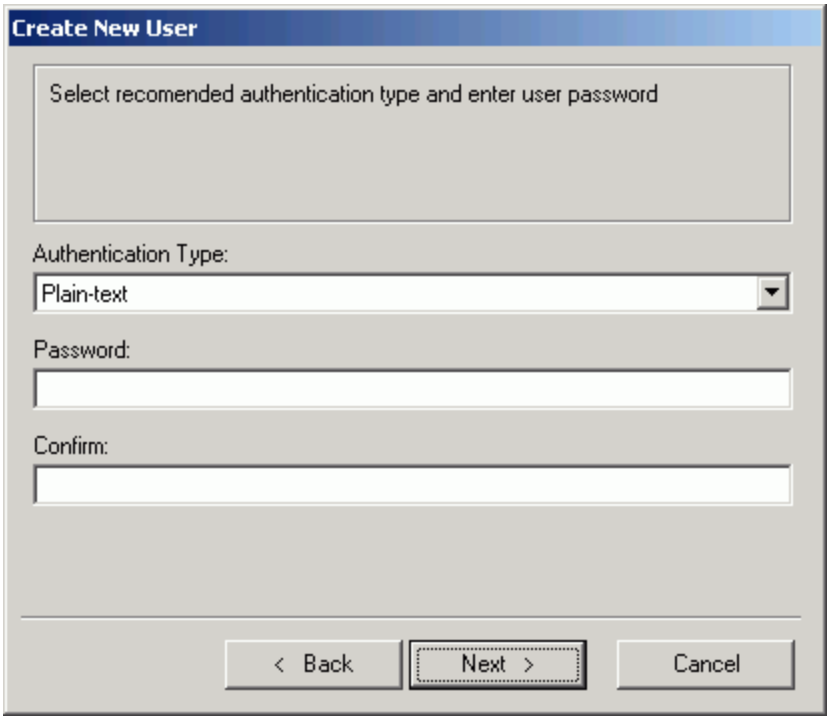

The next window is where the authentication parameters can be set up. If users will connect to the server using e-mail clients supporting strong authentication methods (e.g. The Bat!) then the CRAM-MD5 authentication method is preferable. If Outlook is planned to be used then you should stick to NTLM (MSN) authentication method.

The next window will either terminate the Wizard or will continue creating users:

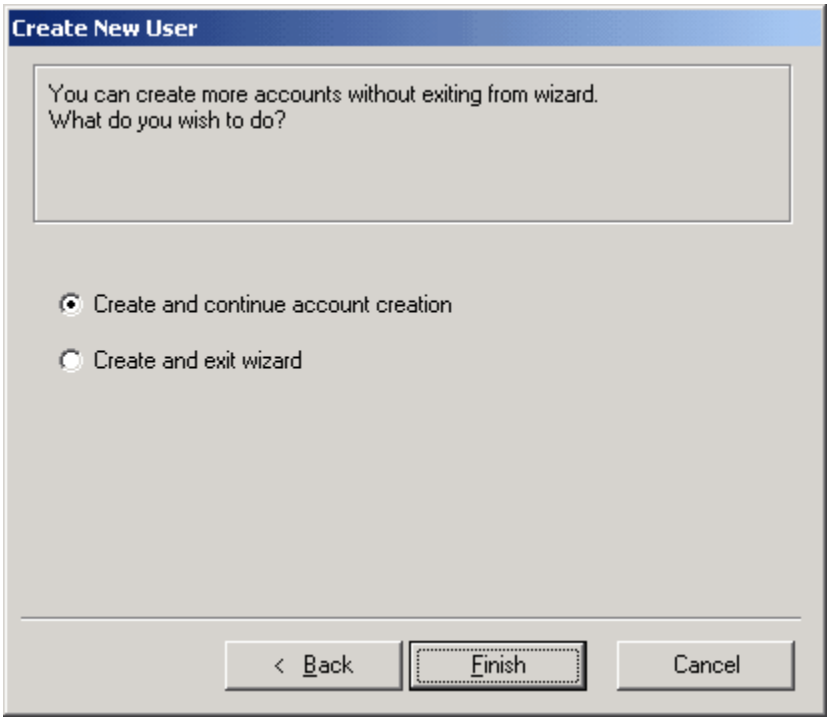

3.7 Testing server configuration

The next step is launching the server in order to test its configuration. After the primary setting wizards are done the mail server will be launched automatically.

The monitor module (BatPostM.exe) is designed to operate and control the server. It allows starting and stopping the server, supervising its activity and CPU usage, as well as checking the Log viewer. The monitor has the ability to supervise several servers at the same time.

The upper part of the monitor window is for controlling the servers and monitoring the level of the CPU and mail protocols' usage:

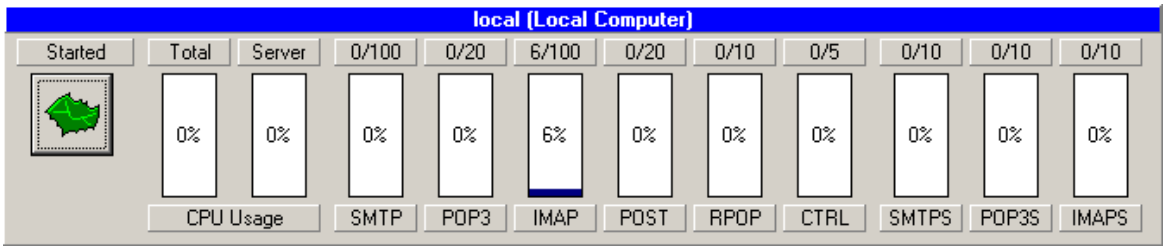

The icon with the bat indicates the current state of the server. The same is displayed in text form above it. Pressing the icon will invoke the following menu for managing the server:

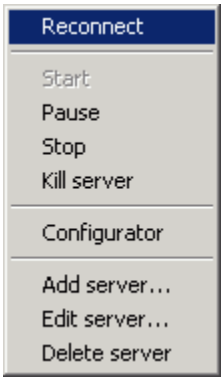

It allows executing all operations concerning the server management and maintaining the list of servers under observation.

In the lower part of the window you will see the Log viewer:

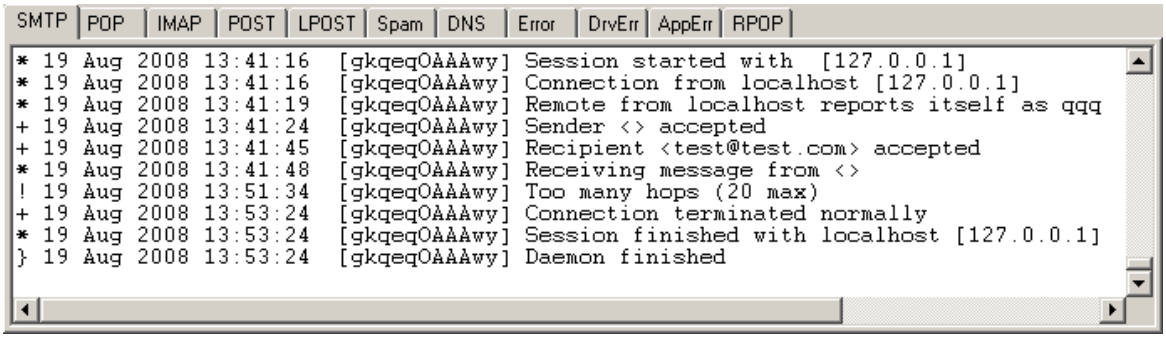

Looking through the server's logs helps detecting problems in the server's operation. Events are sorted according to the categories. The most important is the Error.log which contains general errors in the server's work, and DrvErr.log showing errors in the data accessing.

If the installation and the primary server configuration have been completed successfully, then the server has to be running by now. If not, then the server has not been correctly set up. For more details it is necessary to look into Error.log. E.g. the following error:

```
! 14 Jun 2008 09:59:47  DNS server not specified. Server stopped
```

means that at least one DNS address must be specified. For explanations of the other errors please refer to the "Problems and solutions" section.

If the server is launched we will proceed to checking how it is handling different protocols. For future reference: let's assume that on the server there is the "testdomain.com" domain with the "test" user.

First it is necessary to check the server's working locally. In order to do this run the e-mail client on the same computer where the server is installed. In the e-mail client's settings set "test@testdomain.com" as the e-mail address, the user name will be "test". Make sure you enter the correct password for that user. In the POP3 and SMTP fields enter "127.0.0.1" meaning that the current PC will be used as the server.

Now send a test message from "test@testdomain.com" to the same e-mail address and check for new messages. If everything goes as planned in the Inbox folder you will find that very test message. That is: the server is correctly configured for exchanging messages between local users.

In case error messages appear please refer to the server's logs in order to find out the problem and eventually correct the server settings. In case there are problems sending mail refer to the SMTP.LOG file, whereas the POP.LOG should be considered in case messages cannot be retrieved.

There are several methods of allowing message delivery to external addresses:

- tick the option for checking by mail-from, but be aware this can be easily forged
- define the templates of the addresses or of the host names, which will have the remote delivery allowed
- oblige all the e-mail clients to authorize while sending messages

The last one is preferable and that very option is set as the default one.

In order to check the remote delivery you have to send a test message to some external address. In case there are problems check the SMTP.log in order to make sure that the message has been successfully accepted by the server, and Post.log to make sure that the message has been delivered to the remote server.

3.8 Advanced settings

In case the server is not constantly on-line, it can periodically connect to Internet via dial-up for receiving the incoming mail and sending the outgoing post.

First of all it is necessary to configure the type of the Internet connection in "Options > Networking" section:

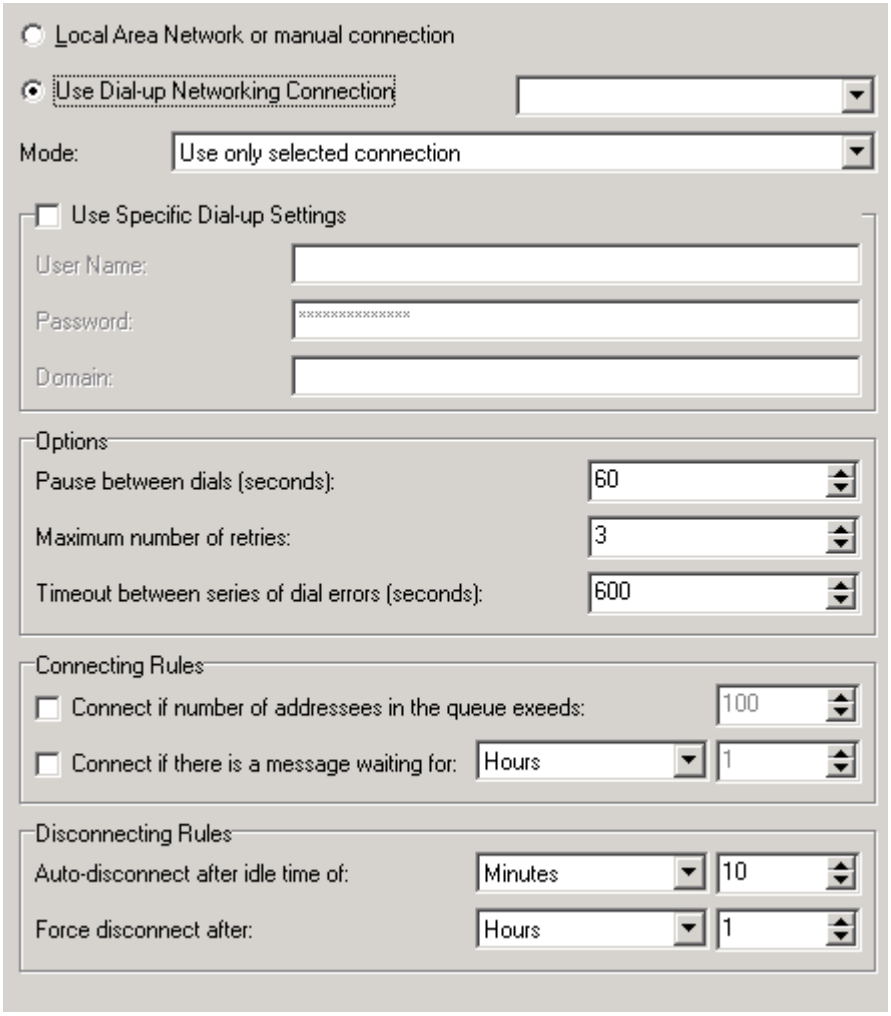

In our case you have to define that the connection to Internet will be via dial-up ("Use Dial-up Networking Connection") and select from the list one of the existing connections. If at the moment the server requires to go on-line the connection to Internet is already established, then there are several options:

- Don't use existing connections
- Use only selected connection

- Use any existing connection

By default the settings entered during the connection configuration will be used, however you can set a specific user name, password and domain.

While establishing the dial-up connection different errors may occur. That is why the server is using the following strategy: it makes a series of dials and in case the connection could not be established, the server makes a break and produces another series of dials. Therefore there are several additional settings: pause between dials within a series, maximum number of retries of a series and pause between series.

As soon as the connection to Internet has been established the server sends the outgoing mail. In order to prevent a mail jam the server can go on-line depending on the amount of messages being queued or on the pending time limit.

Since connecting to Internet does cost a certain amount of money it becomes necessary to decrease that time period. Especially for that there are the following settings available: the automatic disconnect on the preset idle period of time or the forced disconnect upon the timeout.

As the server is not constantly on-line it cannot be responsible for receiving messages to all its domains. In order to keep the mail subsystem working properly it is necessary that someone else takes over the mail retrieval to our domains during that time. In most cases it's the provider's mail server. It receives messages, but does not process them. Instead it keeps them in a separate depository that may be presented as special e-mail account or server queue (spool).

If all messages are collected in a separate e-mail account then it is necessary to use the Remote POP (RPOP) technology in order to retrieve these messages. That means that our server has to connect to the provider's server via the POP3 protocol and receive all messages for that account. At the same time in order to identify the sender and the recipient the message headers are analyzed.

In case the server temporarily keeps the messages in its queue then the ETRN/ATRN technology has to be used. It means that our server connects to the provider's one via the SMTP protocol and sends a special command telling the server it is intending to retrieve mail from the given domain(s). After that, depending on which command was used, the provider's server either establishes a new SMTP connection with our server or the current connection out of the outgoing into an incoming one is transformed. Thus the mail is delivered as though our server had been constantly on-line.

Anyway you have to contact your provider and find out which method they can offer. It is recommended to use the second method, because the message header analysis does not always provide trustworthy information on the sender and on the recipient.

Chapter 4

4 FAQ

In this section you will find answers to the frequently asked questions. Before contacting the support staff please make sure there is no answer to your question here.

4.1 Server settings

In the user settings there is "Authentication type". What should I enter?
In this field the weakest authentication type that is allowed to a user is set. In other words the user will be able to access the server using this method or any other method which is more secure then the initially allowed one.
The following authentication types are supported:

- Plain-text
- NTLM (MSN)
- CRAM-MD5

Authentication types are listed in the order of their security strength (the most secure at the bottom of the list). Using SSL (TLS) connection type increases the security and even though CRAM-MD5 is being used it is acceptable to authenticate in plain text via a secure connection. Authenticating in plain-text transmits the password as open text and therefore the secure SSL (TLS) connection is preferable. Other authentication types make use of the cryptographic methods to confirm the user authenticity without disclosing its password. It is recommended to use CRAM-MD5 as it is a more secure method, though not all e-mail clients support it. If MS Outlook or MS Outlook Express are supposed to be used then it is necessary to use NTLM (MSN) as the only secure methods these e-mail clients can offer.

On the server there are users with identical logins, but belonging to different domains. As they authenticate they get odd results. What is the correct way to authenticate?
If in the user's e-mail client only the name is entered, then the server tries to find the user with the same name across all domains and uses the first one which is found. Therefore a user from a different domain might fall in. In this case authentication will fail if the passwords do not match, but in the event they do, one user will see the messages of a totally different user.
To avoid that problem the whole e-mail address should be used as the user name. In other words, instead of "user" put "user@example.com" This will eliminate any equivocal situations.

What is the right way to configure "Remote POP"?
"Remote POP" is designed to retrieve mail from external mailboxes via the POP3 protocol and to place them into internal server mailboxes or forward to a predefined address. In order to get it function properly it is necessary to enter the external server's settings and to create a schedule for the periodic execution of that task. This schedule should be added to the list of the active ones under the "Schedules" tab.
There are two operating modes:

- "Remote Account" - in that mode messages from external mailboxes are moved into the respective local e-mail accounts or are forwarded to a certain e-mail address.
- Domain POP" - in that mode the message headers are analyzed in order to determine the recipient.

Since any changes in the settings require the server to be restarted in order to take effect, it is quite handy to make use of the "Run Now!" button to execute the tasks with new settings right away. You can check the results in the RPop.LOG file - this is the only way to find out the reasons for the error messages or to make sure the tasks have been carried out as expected. After you have checked the settings you have to make sure the schedule is also set up accordingly and then restart the server.

What is the purpose of the "Forward To" option under the "Remote POP" settings and should it be filled in?
Depending on the mode in use, the "Forward To" option executes different actions. In the "Remote Account" mode into this field you should enter the e-mail address, to which all the messages will be forwarded. In case the recipient of the message could not be identified, it will be forwarded to the e-mail address entered into that field in the "Domain POP" mode.
It is recommended to fill in that field in both modes, otherwise some messages can be left on the server. In future, these e-mails will be downloaded during each session until they are manually deleted. If there are too many such messages on the server it may lead to the slowdown of the "Remote POP" mode and result in the high traffic load.

On our local server we would like to handle only a part of the domain users, the rest are processed by another server. How do we set it up?
To accomplish that it is necessary to create a domain and to register only those users, who will be stored locally. While in the domain settings on the "Advanced" tab one should also enable "Treat unknown users as remote".

How to deprive certain users of the right to send messages to external e-mail addresses?
In the user settings on the "Basic" tab there is the "SMTP mail sending rights" checkbox. If you enable it, then you can determine privileges of every single user regarding the the right to send messages to local and external addresses. Green boxes mean it is enabled, the red ones - disabled. The more boxes are ticked, the stronger the prohibition or permissions are. This might be important since at the same time other rules might be in power (for more details refer to the "Administrator's Manual"). To make sure the highest level of prohibition is reached it is necessary to enable four red boxes.

We have registered a new mail domain name, but still there are messages coming from the old one. What do we have to do to have two different domains on the server, but the users of the domains should be the same?
How do we create several domains with the same content (a main domain and its aliases)?
While creating the domain on the "Advanced" tab in the "Contents Type" field you can enable the "Referral contents" option and in the "Referral" field indicate the domain, which will be used as the source of users for the newly created domain. At the same time the data is not duplicated, because the users of the indicated domain will be used when the users of the new domain are contacted.

In the domain settings on the "Protocols" tab there is an "IP address" option. What do I have to put in there?
In most of the cases this field should be left blank. Only if the server possesses several network interfaces (network cards) located in different subnets and it is required that the

connections only from a certain subnet are allowed, is it necessary to input the IP-address of that very subnet. If that field is left blank, then all the connections from the existing subnets will be accepted.

Which settings' changes require the server to be restarted?
Restarting the server require the changes in the settings of domains, schedules, anti-virus and anti-spam, as well as the changes in the settings of the "Options" section.

4.2 Registration

Why is it necessary to register the programme?
After you register the programme there will be no more limit on the number of users imposed by the trial version. You will also benefit from the free of charge technical support.

Where do I enter the registration key?
In order to enter the registration key you have to invoke the Configurator and go to the **Help > Enter Registration Code** menu. If you go to **Help > About** you will see the list of all registrations. The **Add** and **Delete** buttons allow adding new and removing existing registrations. The number of licences (the maximum number of users) is summed up from all the existing registrations.

Note. In order to register the programme under Windows Vista you have to launch the Configurator with the administrator privileges (Run as Administrator).

What do I receive after the registration?
In the e-mail with registration information you will find:
- the Key
- the Key Checksum
- the Key Password

The password is used to protect the registration key from being misused and is put into that e-mail only if it has been assigned by the dealer. If you entered the password during the registration process then it will not appear in that e-mail.

Who assigns the password to the registration key?
The password is either entered by the user in the registration form or generated by the dealer. In the latter case it is the dealer who has to supply the user with the respective password.

What are the limitations of a unregistered version?
In the trial version it is possible to create no more than 20 accounts (users) across all domains. The trial period is 30 days long. During the trial period all the programme features are accessible.

What happens if the programme runs out of the trial period if it has not been registered yet?
After the trial period expires the programme will not stop working, but there will be some more restrictions imposed: the maximum number of connections for each protocol will be limited to 2 regardless of the current settings.

4.3 Problems and solutions

While trying to send out mail the server replies "We don't relay". What does it mean?
The default settings require the SMTP-authentication is always performed when the mail is sent to external e-mail addresses. This way the server is protected against being used as the means of sending unsolicited e-mail (junk messages).
To eliminate that error message it is necessary to enable SMTP-authentication for the outgoing mail in the e-mail client's settings.

When launching the server the following error messages are written into the Error.log files: "! 15 Oct 2007 15:57:06 POP addresses configuration error" or "! 15 Oct 2007 15:57:06 SMTP addresses configuration error". What's wrong in the settings?
These error messages mean there is a conflict between the domains in the protocol settings. Most likely, for the given protocol, in one domain an IP-address was set for a certain subnet, while in the other domain this field has been left blank. The connections on that port are accepted either from the defined subnet or from all subnets, but not from both simultaneously. That is the reason of the conflict.
The "IP address" field might have been filled in by mistake. In most cases all connections should be accepted and that field has to be blank.

When launching the server the following error messages are written into the Error.log files: "Could not bind to address "0.0.0.0", port 25".
This means that port 25 (SMTP) is already used by another programme, mail server or the system's SMTP service is started.
To find out what occupied the port, run the following command from the command line:

```
telnet localhost 25
```

Usually, in the server's greeting its own name is mentioned.

When launching the server the following error messages are written into the Error.log files: "Could not bind to address "192.168.101.1", port 25".
Probably in the domain protocol settings the "IP address" field had been filled in by mistake or this address does not belong to that computer. In order to find out the IP-addresses of that computer it is necessary to run the following command from the command line:

```
ipconfig
```

To get all the information on the network settings run the following command:

```
ipconfig /all
```

While trying to send out mail the server replies "5.7.1 - unable to relay".
Obviously, this error message does not come from BatPost. While sending the mail out either there was another server in use by mistake or besides BatPost there was another server in play, which did intercept the SMTP port.
To clarify that it is necessary to run the following command on the server:

```
telnet localhost 25
```

and make sure the server's greeting contains this word: "BatPost".

The mail delivery (even local) takes too much time.
High CPU load.
Antivirus monitors may try scanning every incoming message. With a high amount of

messages this might significantly slow down the server. Try disabling the antivirus and, if this solves the problem, it means the antivirus is unable to process such number of e-mails.
In order to find out which application loads the CPU you can make use of the Task Manager. In some cases it indicates that the CPU load is due to the System process. To get a more detailed information use the Process Explorer. You can download it from here: http://technet.microsoft.com/en-us/sysinternals/bb896653.aspx

While trying to receive mail the server replies "Authentication failed. Mailbox locked".
The POP3 protocol requires that the mail account is blocked while it is in use. If you try to access that account while someone has already established a connection with it, you will get the above mentioned error message.
You have to make sure that at a certain point of time a POP3 account is accessed only from one client. For instance, it is of no wonder when a user access his/hers e-mail account from work and from home, but does not do that simultaneously. On the other side there is IMAP, which allows multiple concurrent connections.

The server is unable to send mail to the external servers. What might be the problem?
First of all it is necessary to look through the Post.LOG file, it may contain error messages which could clarify the situation. If the DNS is used to deliver mail to external addresses it is necessary to make sure that its address is set accordingly and that it is functioning properly. In case you use a relay server you have to make sure the message has been successfully passed over to it.
Quite often while using DNS you may successfully connect to an external server, but it refuses to accept mail. The reason for that might go down to the extended check-ups of e-mail performed by the external server, which tries to combat spam this way. And so, the following might be checked:

- Whether our server is registered in the DNS (this implies we use a static IP-address).
- Whether the PTR record is registered in the DNS reverse zone. This is characteristic to the mail.ru server, in case there is no PTR record, it returns the "550 Unroutable address" error message.

Sometimes, the only solution to the problem is to use a relay server to send messages to external addresses.

We do not receive mail from outer world. What might be the problem?
There are two ways for the mail from external servers to arrive at our server:

- If we have our own mail domain and in the DNS it is prescribed that our server is responsible for mail in that domain, then other servers will contact us and hand over the mail. In this case the A record describing our IP-address and the MX record referring to the A record, have to be set for our server.
- We can connect to the other servers and collect the corresponding mail ourselves. In this case we can even lack our own domain and fetch mail from certain addresses via "Remote POP".
  In case we do have a mail domain, it can point to a server that temporarily collects mail for our domain (e.g. ISP's server). This intermediate server can store e-mails in its message queue and give them away via ETRN/ATRN or pile up messages into a special account so that we can fetch them via "Remote POP". The first instance is preferable as in this case the information on the initial message recipient is not lost, but

not all servers support that. Anyway, the decision is up to the intermediate server owner.

Authentication fails. In the logs I see "Incorrect password", though I am sure the password is correct. What is the problem?
Most likely, in the e-mail client settings you entered only the user name without indicating the domain, while on the server such user exists in several domains. In this case the server will look for the user in all domains and will check the first one it finds possessing that name. To avoid such ambiguous situations you have to enter the whole username (coinciding with the e-mail address). In other words, instead of mere "user" put "user@example.com".

Authentication fails. In the logs I see"User not found", though such user exists. What is the problem?
There are two main reasons:

- A new domain with users had been created, but the server has not been restarted. Creating, deleting and changing the domain settings require the server to be restarted.
- In the domain, to which the user belongs, the needed protocol is disabled. Or this protocol has been incorrectly set up. You have to make sure that in the domain settings on the "Protocols" tab the needed protocol is enabled and that in the Error.log file there are no messages regarding errors with that protocol. After altering the settings it is mandatory to restart the server.

Changing settings does not take effect. In the Configurator the changes are visible, but the server functions as before. What do I have to do?
Changing some settings requires the server to be restarted in order for the changes to take effect. This refers to the domain settings, schedules, antivirus and anti-spam as well as to the settings under the "Options" section.

4.4 Logs

Which logs are maintained on the server and what are their purpose?
The server itself does not have a graphical interface and begins operating before a user logins under his/hers account. Therefore, the only way to supervise server functioning and diagnose related problems are the log files. To make it more convenient different events are logged into different files. By default, the logs are stored in plain text, but it is also possible to save logs as a database. However, you have to keep in mind that processing database takes longer compared to plain text files and can slow down the server.

The server stores the following logs:

- SMTP.LOG - stores the SMTP-session events. SMTP is used by e-mail clients and by other servers to send mail.
- POP.LOG - stores the POP-session events. POP is used to receive mail.
- IMAP.LOG - stores the IMAP-session events. IMAP is used to receive mail.
- Post.LOG - stores the events related to delivering mail to other servers.
- LPost.LOG - stores the events related to the local mail delivery.
- Session.LOG - stores sessions which have the "Log Sessions" option enabled on the "Parameters" tab in the "Options > Logs" section.
- RPOP.LOG - stores the events of the "Remote POP" sessions.

- Control.LOG - stores the events of the control protocol sessions. The control protocol is used by the Configurator to set up the server.
- PCP.LOG - stores the events of the PCP-sessions. PCP (Password Change Protocol) is used to change user passwords. This protocol is supported by The Bat! e-mail client.
- Spam.LOG - stores the events related to the functioning of the anti-spam mechanisms.
- DNS.LOG - stores the events related to the requests at the DNS.
- Error.LOG - stores the events related to the errors in the server operation. You should pay more attention to these errors as they mean that either there are problems with the configuration or with the server itself. In the latter case you should notify the developers.
- DrvErr.LOG - stores the events related to the errors in the data access driver operation. They should also be handled with care, because there might be errors in the data access driver or the databases are damaged.
- AppErr.LOG - currently not used.
- Virus.LOG - stores the events related to the functioning of the anti-virus mechanisms.
- Dialup.LOG - stores the events related to the functioning of the dial-up. If the telephone line is used to connect to the internet then here you will see the records regarding the established and ended connections.

What are the exception logs for?
An exceptional situation in most cases notifies of a problem in one of the server components. In a case like this the following log file is created:

```
<component_name>_Exceptions.log
```

For instance, if an exception occurred in the BatPost.exe module, then the following log file is created:

```
BatPost_Exceptions.log
```

For most of the components the log files are saved into the following folder:

```
%allusersprofile%\Application Data\BatPost\
```

The data access driver log file can be found under:

```
%allusersprofile%\Application Data\BatPost\Drivers\
```

The value of the `%allusersprofile%` variable depends on the system. E.g. under Windows XP it is usually:

```
C:\Documents and Settings\All Users\
```

In order to find out the value in your system you have to run the following command from the command line:

```
echo %allusersprofile%
```

These logs allow developers better understand the cause and the location of the error. Therefore, when reporting the problem to the developers do attach the respective log file.

The size of the log files increases dramatically, what can I do about that?
Since the log files store a lot of events regarding the server operation they quickly grow in size. How fast this happens depends on the server load, but sooner or later you might run across a situation when the disk has run out of free space and the server is not responding. Even if there is enough free space it might turn out to be quite inconvenient to work with such large files.
In order to avoid such situations the server supports the automatic splitting and purging of the logs. You can set it up on the "Splitting/Purging" tab in the "Options > Logs" section. When the log is split into a separate file it gets the corresponding date in the file name. E. g. when splitting POP.LOG the following file might be created:

```
POP20080602.LOG
```

which relates to the year 2008, month 06 (June) and day 02. This way of setting the date allows sorting the files by name and have them arranged according to the time of their

creation.
Logs can be split into parts: every day, every week or when they reach the maximum allowed size. If the log files does not increase in size too fast, then splitting it once a week should do; on the other hand for the fast growing logs (e.g. SMTP.LOG) the daily splitting is preferable.
Splitting the log files keeps their size at the desired level, but does not care about the risks of the hard-drive being overfilled. However, the automatic purging will solve this problem. Logs that are older than the defined number of days can be automatically deleted from the disk or moved into a certain folder. In case you delete logs it is advisable to keep them for a month at least, because sometimes one has to deal with problems that appeared way too earlier then you first noticed them. If the log files are moved into a special folder, then we recommend you back up the old ones onto a CD/DVD, so that the new location does not get overfilled either.

Which is the better application to view logs?
You can use the Monitor to view the last records in the main log files. This would be enough to supervise the current server operation, but for the advanced tasks it is necessary to process the log files themselves. Any application capable of viewing large files and possessing a search facility will be enough. Many file managers (Far, Total Commander) and text editors will do. A file manager in this sense is preferable as it allows displaying all log files, search within several files and execute commands from the command line.

How do I find in the logs what I need?
Since the log files can be very large, often it is quite hard to find the necessary fragment of the log. However, if you stick to a certain strategy then finding the needed info will become quite easy.
First of all you have to define the key phrase which will help us find what we are looking for. For instance, if we are looking for something related to the user of "user@example.com" then this is our key phrase.
Then you have to define to which date the event in case belongs. If the logs were split, then you can concentrate on the files corresponding to that date. Sometimes, the precise date is unknown and therefore we can search through several files. For example, we can search in the "Logs" folder using the "*.*" mask to search in all files, or make use of the "SMTP*.*" mask to search only in files related to the SMTP protocol.
When the key phrase is found we are usually interested in the events related to the same session. Here is how the POP session looks like:

```
{ 03 Jun 2008 17:38:01  [gURXxDAAAgA] Daemon started
* 03 Jun 2008 17:38:01  [gURXxDAAAgA] Session started with  [127.0.0.1]
* 03 Jun 2008 17:38:01  [gURXxDAAAgA] Connection from localhost [127.0.0.1]
+ 03 Jun 2008 17:38:01  [gURXxDAAAgA] User test <test@test.com> logged on
(Plain-text)
+ 03 Jun 2008 17:38:01  [gURXxDAAAgA] User test <test@test.com> logged off
+ 03 Jun 2008 17:38:01  [gURXxDAAAgA] Connection terminated normally
* 03 Jun 2008 17:38:01  [gURXxDAAAgA] Session finished with localhost
[127.0.0.1]
} 03 Jun 2008 17:38:01  [gURXxDAAAgA] Daemon finished
```

You have probably noticed that all these records possess one and the same unique session ID. In this case it's "gURXxDAAAgA". This ID can be used as an identifier in searching records of that session. In simple cases you can do that by sight, but in case the records of different sessions are mixed up you can execute the following command from the command line:

```
grep gURXxDAAAgA POP.LOG
```

in order to output the "gURXxDAAAgA" session's entries directly onto the screen or this command:

```
grep gURXxDAAAgA POP.LOG >Ses.txt
```

to put them into the "Ses.txt" file.

Note. By default your system might lack the grep utility. Initially it had been developed for UNIX and under Windows it is absent. However, you can download it as part of the UnxUtils package from the following web-site:http://unxutils.sourceforge.net
This package is distributed as a zip-archive which you need to unpack into a folder on your hard-drive (e.g. C:\UnxUtils). After that you should add that folder to the environment variable called PATH, so that the utility can be accessed from everywhere.

How do I trace the message going through the server?

Once a message arrives on the server it gets a unique ID. Make sure you do not confuse it with the session ID, they do look differently. You will notice that ID in the logs where the message is being referred to. Here is an example of how it looks like in the SMTP.LOG:

```
> 02 Jun 2008 17:05:45  [g0Q9fIAAAAA] Message 4843FD8700000001 (4438 bytes)
received
```

in this case 4843FD8700000001 is the unique ID assigned to the received message. Here is an example of how it looks like in the POP.LOG when the message is read and deleted:

```
+ 02 Jun 2008 17:05:52  [g0Q9fIAAAAB] Message 4843FD8700000001 (4438 bytes)
was read
+ 02 Jun 2008 17:05:52  [g0Q9fIAAAAB] Message 4843FD8700000001 (4438 bytes)
marked as deleted
+ 02 Jun 2008 17:05:52  [g0Q9fIAAAAB] Message 4843FD8700000001 (4438 bytes)
permanently deleted
```

This ID is also put into the message headers and namely into the Received: field:

```
Received: from userdomain.com (userdomain.com [192.168.101.1])
     by testdomain.com with BatPost v2.21r4 ESMTP Daemon
     id 4843FD8700000001
     for <test@testdomain.com>; Mon, 2 Jun 2008 17:05:45 +0300
```

The Received field is usually written by every server which the message has gone through, that is why you have to find the one put by the BatPost server.
Since the unique ID is written into the logs upon every action on the message it is possible to trace the life-cycle of every message on the server. For instance, in the SMTP.LOG file you can find entries on how the message had been accepted by the server, whereas in the POP.LOG file (or IMAP.LOG) there are infos on how the message had been read by the user and, maybe, deleted.
If the message ID from the Received field in the message headers is known, then it is possible to search through all the log files to find all records related to a certain message and have an idea what had been happening to the message on the server all way long.

4.5 Maintaining the server

Where does the server store its files?

Prior to version 2.21 the server stored all its files in the %programfiles%\BatPost folder by default. Since version 2.21 the programme files are stored in the %programfiles%\BatPost folder, whereas the data files are saved in the %allusersprofile%\Application Data\BatPost folder. Such division has been done in order to make the server compatible with Window Vista.
For compatibility with the previous versions of the server it checks the location of the

Server.ini file, if it is in the %programfiles%\BatPost folder, then it is presumed that the data files are also there. In case the file is not there, then it is presumed that the data files are stored under %allusersprofile%\Application Data\BatPost.
In the root of the server data structure the following files are available:

- Server.ini - containing the general server settings.
- GlobalAdmins - containing the list of global administrators. They have the right to remotely monitor and configure the server.

The following subfolders are stored there as well:

- Archive - the archive of messages can be stored here.
- DB - here is the server configuration stored, it comprises groups, domains, users, folders, messages, as well as the "black" and "white" lists.
- Drivers - the data access driver settings are stored here.
- Logs - here are the server operation logs stored.
- PEMs - server certificates can be stored here.
- Spool - this is the place for storing the server sending queue. It temporarily keeps messages (MSG-files) until they reach the needed base or are sent out to an external server. Also here the status of their sending off is stored (QUE-files).
- Stream - temporary files are stored here, they are too large to be kept in RAM.

The place for storing the message archive can be altered in the Configurator under "Options > Archive&Audit". The location of the logs, sending queue and streams can be changed in the Configurator under the "Options > Common" section.

Note. In order to change the location where the configuration settings are stored it is not enough to only alter the settings under the "Options > Common" section. It is also necessary to manually change the RootDir parameter in the \Drivers\DefDrv.ini file in the Common section.

To change the location of the certificate database go to "Options > Security".

How do I move the server to another computer?
Owing to the fact that the server stores all its settings in form of files on the disk and does not make use of the Windows registry, moving the server from one PC to another is reduced to mere copying of files with settings.
First of all you have to install the server on the new PC. At the end of the installation process you can skip the server set up, because the server configuration will be replaced anyway. Then you have to transfer to the new place the Server.ini and GlobalAdmins files, as well as the contents of the DB, Drivers, Spool and PEMs (if you have any certificates) folders. If needed, the contents of the rest of the folders can also be transferred.
After that you have to manually launch the server to make sure the transfer of the server has been completed successfully.

What is the correct procedure of updating the server to a newer version?
To accomplish that it is necessary to close the Configurator and Monitor (if launched). After that install the new version. At the end of the installation process you will be notified that the server is already configured and prompted to make changes, the latter can be canceled.

4.6 Other questions

How can I set up the web-interface access to the server?
The server itself does not have a web-interface, but there are certain kits by third-party developers which allow adding it. Such packages usually refer to the server via the IMAP protocol. The greater part of such solutions, in order to function properly, require the Apache web-server (http://www.apache.org/) with PHP (http://www.php.net/) be installed.
Examples of such solutions are the following kits:

- SquirrelMail (http://www.squirrelmail.org/)
- Horde IMP (http://www.horde.org/imp/)
- RoundCube (http://roundcube.net/)

Which anti-virus programmes are supported by the server?
The server can make use of the anti-virus plug-ins suitable for The Bat! E-mail Client (http://www.ritlabs.com/en/products/thebat/plugin.php). Some plug-ins are already included in the anti-virus programmes. Not all the plug-ins function properly with the server. This may be related to the fact that the server checks for viruses many messages at once and not all the anti-virus plug-ins are capable of doing that. Among the tested and trusted plug-ins we can point out the NOD32 plug-in.

Which anti-spam solutions are supported by the server?
The server has got a built-in support for SpamAssassin (http://spamassassin.apache.org/) - one of the most powerful anti-spam solutions. SpamAssassin can be installed on the same PC where the server is running or on a standalone computer (e.g., under Unix). The server communicates with SpamAssassin via the spamd module, so it also has to be installed. Installing and, mainly, configuring it under Windows can become a difficult task to complete. You can find its detailed description at: http://wiki.apache.org/spamassassin/UsingOnWindows.
The server supports the anti-spam plug-ins suitable for The Bat (http://www.ritlabs.com/en/products/thebat/plugin.php). Some of them are bound to The Bat! and do not work with the server. Among the tested plug-ins there is the Bayes Filter Plugin.

Bayes Filter Plugin v2.0.4 for The Bat!
http://www.ritlabs.com/download/files3/the_bat/plugins/antispam/bayesfilter2.0.4.exe

The "No Report Dialog Below" parameter has to be set to 0 in order to disable the on-screen messages. We recommend to set full paths to files to be sure of their location. The statistics on the "About" tab is not refreshed at once, the differences take effect after the server restart. Probably, this had been done to increase the processing speed and therefore the information is not written into the database at once, but first is cached in memory. Since in the Configurator, unlike the server, another copy of the plug-in is loaded, the changes in statistics are not seen at once.

In order to function properly many of the anti-spam solutions require prior training. In other words, it is necessary to teach the plug-ins which messages, according to the user, are spam and which ones are not. This is needed to work out the characteristic features of such messages, so that in future it will be possible to distinguish junk and not-junk e-mails. Training should be carried out regularly as new sorts of spam-messages permanently appear. For training there are two special addresses: $SPAM and $NONSPAM to which it is necessary to send spam and not-spam messages respectively.

It is necessary to authenticate in order to be able to send messages to these addresses. Another condition is that the user has to be a member of the group with the right to train the anti-spam modules.

How does one create a certificate for the server?
In order to be able to use the secure SSL/TLS connection it is necessary that the certificate is installed on the server. The sub-folder called PEMs is suitable for storing the certificate and the respective private key.
The server supports SSL/TLS with the help of the OpenSSL library (http://www.openssl.org). That is why the pem-format is used for storing certificates and private keys. This format allows keeping the certificates in plain text which is quite handy for sending them via e-mail.
It is possible to use a self-signed certificate as well as the one issued by the certifying authority. In case of the self-signed certificate it is will be necessary to add it to the list of the trusted certificates in the e-mail client; on the other hand if the certificate had been received from the certifying authority, then, most likely, it will be already trusted, because the root certificates of the certifying authorities are already in the list of the trusted certificates.
The openssl utility, as part of the OpenSSL library, will also be needed. It is possible to build the library from source codes or download a precompiled version. A Windows version can be found here: http://www.openssl.org/related/binaries.html.

First of all it is necessary to generate a private RSA key and a request for certification. First we create a .rnd file and enter random data there, this will be the basis for generating the key later. You can insert into that file any random text, the size of the file is of no importance. After that you have to execute the following from the command line:

```
openssl req -newkey rsa:1024 -keyout key.pem -out req.pem -config openssl.cnf
```

This command will create a 1024-bit private RSA key and will save it into the key.pem file. The key has to look like this:

```
-----BEGIN RSA PRIVATE KEY-----
MIIBOwIBAAJBALtv55QyzG6i2PlwZ1pah7++Gv8L5j6Hnyr/uTZE1NLG0ABDDexm
q/R4KedLjFEIYjocDui+IXs62NNtXrT8odkCAwEAAQJAbwXq0vJ/+uyEvsNgxLko
nWmM1KvqnAo5uQIhALqEADu5U1Wvt8UN8UDGBRPQulHWNycuNV45d3nnskWPAiAw
ueTyr6WsZ5+SD8g/Hy3xuvF3nPmJRH+rwvVihlcFOg==
-----END RSA PRIVATE KEY-----
```

While creating the key you will be prompted for a password to it. Later you will have to enter that password in the server settings.

The certification request will also be created. That request contains the description of the future certificate: Country, State/Region, City, Organisation, Department, the server's domain name, etc. While the command is executed you will be asked for the attributes of the future certificate. Do not forget to enter the server's domain name as the "Common Name" attribute. On basis of the private key from the key.pem file and data entered by the user, this command will create a request for certification and will save it into the req.pem file. The request has to look like this:

```
-----BEGIN CERTIFICATE REQUEST-----
MIIBGzCBxgIBADBjMQswCQYDVQQGEwJBVTETMBEGA1UECBMKUXVlZW5zbGFuZDEa
MBgGA1UEChMRQ3J5cHRTb2Z0IFB0eSBMdGQxIzAhBgNVBAMTGkNsaWVudCB0ZXN0
2NNtXrT8odkCAwEAATANBgkqhkiG9w0BAQQFAANBAC5JBTeji7RosqMaUIDzIW13
oO6+kPhx9fXSpMFHIsY3aH92Milkov/2A4SuZTcnv/P6+8klmS0EaiUKcRzak4E=
-----END CERTIFICATE REQUEST-----
```

Now you have a private key in the key.pem file and a request for certification in the req.pem file. Make sure you save the private key in a safe place. It will be used for proper

server operation over secure connections. Without the private key the certificate will become almost useless. For security reasons the certification authorities will not be able to issue you a new certificate (e.g. if the attributes have been changed) if you do not have the access to the private key.

Now it's time to send your request for certification to the certifying authority. It can be VeriSign (http://www.verisign.com) or Thawte (http://www.thawte.com).After your data has been verified you will either receive the certificate or the reason of the rejection. Then you have to put the certificate into the cert.pem file. It has to look like this:

```
-----BEGIN CERTIFICATE-----
MIICLjCCAZcCAQEwDQYJKoZIhvcNAQEEBQAwWzELMAkGA1UEBhMCQVUxEzARBgNV
BAgTClF1ZWVuc2xhbmQxGjAYBgNVBAoTEUNyeXB0U29mdCBQdHkgTHRkMRswGQYD
dp7jnmWZwKZ9cXsNUS2o4OL07qOk2HOywC0YsNZQsOBu1CBTYYkIefDiKFL1zQHh
8lwwNd4NP+OE3NzUNkCfh4DnFfg9WHkXUlD5UpxNRJ4gJA==
-----END CERTIFICATE-----
```

If you need to create a self-signed certificate, then you have to execute the following from the command line:

```
openssl req -new -key key.pem -out cert.pem -x509 -config openssl.cnf
```

Such certificate is mostly used in testing purposes. It is also possible to receive test certificates at VeriSign: http://digitalid.verisign.com/test_server_ids.html (select "C2Net (Apache-SSL-US)") or at Thawte: https://www.thawte.com/cgi/server/test.exe (select "Generate an X.509v3 certificate & Use the most basic format").

By default the server looks for the private key and the certificate in the ".\PEMs\server.pem" file. It is possible to copy them into that file using a text editor or via executing this command:

```
copy key.pem+cert.pem server.pem
```

Or it is also possible to point out the necessary files in Configurator in the "Options > Security" section. It is also mandatory to enter the password for decrypting the private key.