

Exchange 2003 Mailflow – Part I

Written by Marc Grote - <mailto:grotem@it-training-grote.de>

Abstract

In this article I will show you how e-mail messages will flow through Exchange Server 2003 for Intrasite- and Intersite Messaging. This is part one of a two part article. Part one deals with the basics of message delivery and transmitting from Outlook to Exchange and between Exchange Servers, part two of the article will try to help you troubleshooting e-mail delivery and message flow.

Let's begin

There are several components that are involved in the Mail delivery process.

Information Store (Store.exe)

The Microsoft Exchange Server Information Store (Store.exe) is the end point for e-mails send to users on this server and the start point for e-mails which are send by MAPI clients like Microsoft Outlook 2003 which directly connects to the MSExchangeIS.

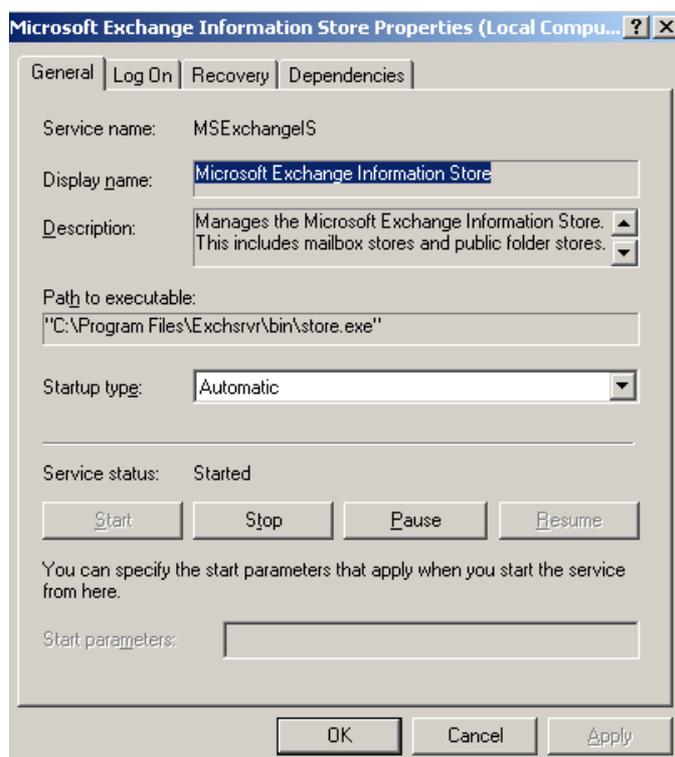


Figure 1: MSExchangeIS

Exchange InterProcess Communication (EXIPC)

EXIPC is responsible for Data Transfer between Internet Information Server 6.0 (IIS) and the Microsoft Exchange Server Information Store (MSExchangeIS). EXIPC provides a layered service between both components to achieve the best possible Performance between IIS dependant components and the Exchange Databases. As you might know, all Internet Client Access Protocols like HTTP/S, SMTP, POP3 and IMAP4 are configured and managed by IIS with some exceptions.

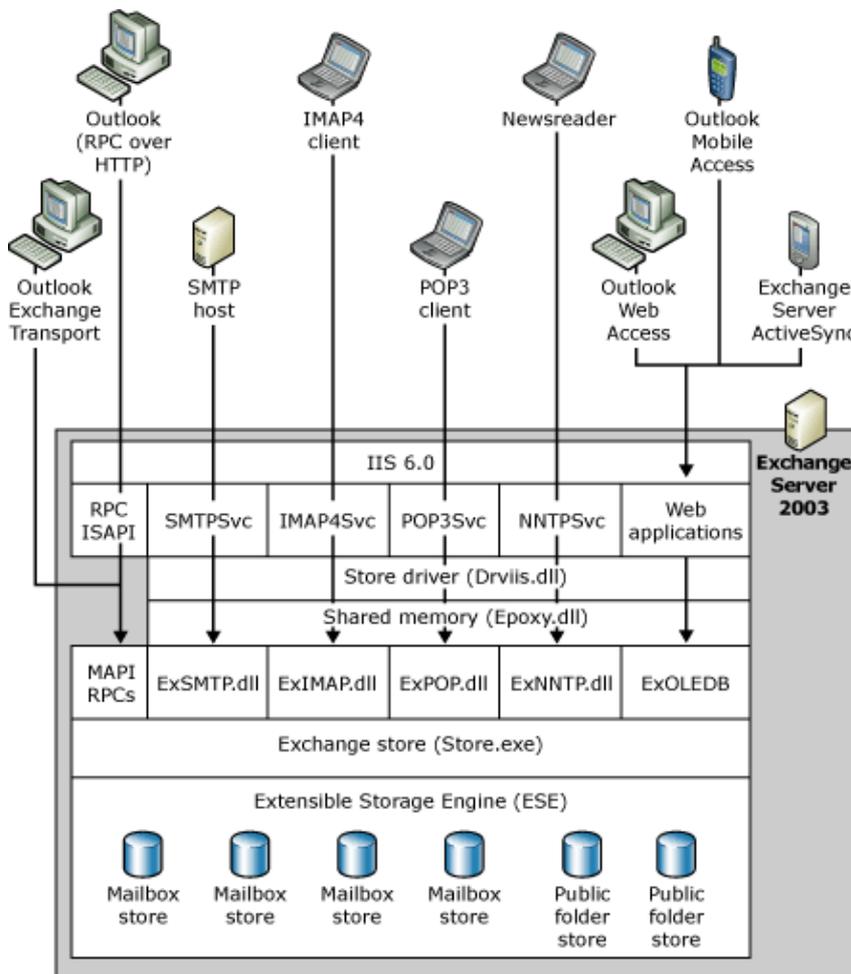


Figure 2: EXIPC Layer

This interaction allows Exchange to be in a FrontEnd- and BackEnd- Server scenario.

Through Virtual Servers, multiple configurations of the same protocol can exist on a single Exchange Server.

Advanced Queuing Engine (AQE)

The Advanced Queuing Engine (AQE) is responsible for creating and managing Message queues for E-Mail delivery. When AQE receives a Simple Mail

Transfer Protocol (SMTP) Mailmsg object, this object will be forwarded to the Message Categorizer. The Advanced Queuing Engine queues then the Mailmsg object for message delivery based on the Routing information provided by the Routing Engine process of Exchange Server 2003.

The Message Categorizer is part of the Advanced Queuing Engine and is responsible for address resolution on every Mailmsg object that flows through the AQE. The Message Categorizer is implemented as a Event Sink. The Message Categorizer is also responsible for splitting messages into RTF or MAPI.

Routing Engine

The Exchange Routing Engine uses Link State information for E-Mail routing. The Routing Engine will forward these information to the Advanced Queuing Engine.

Please note:

The SMTP Stack from Windows Server 2003 will be extended through the Exchange Server installation process with several enhancements. One of these enhancements is the implementation of the XLINKSTATE protocol.

The Routing Engine creates and maintains the Link State information for every Exchange Server and is also responsible for routing the messages to inbound or outbound destinations.

SMTP Service

The SMTP Service processes incoming traffic from any SMTP hosts. SMTP is also used in most communication between Exchange Servers (except Exchange 5.x Servers which uses RPC for Message Transferring). SMTP is also responsible for some advanced Exchange Server functions like Message Journaling. During the Exchange installation, the built in SMTP Service from Windows Server 2003 will be extended with several new functions. Some of the Enhancements are:

- Moving the Message Queue Directories to the Exchange installation Directory
- Providing support for the LSA (Link State Algorithm) in SMTP
- Moving SMTP Messaging from IIS to the Exchange System Manager

Message Flow

Because the understanding of E-Mail message flow is important I will list some High Level steps in Message flow:

- MAPI Client sends a message to a remote recipient
- Information Store (Store.exe) receives the Message
- The created MailMsg object is forwarded to the Advanced Queue Engine (AQE)
- The Message Categorizer from the AQE processes the MailMsg object and splits it into MIME or RTF as necessary.

- The Message Categorizer expands groups, checks defined Message limits on Exchange
- The MailMsg object is then transferred to the Remote Destination Domain within the AQE
- The AQE passes the Destination address to the Exchange Routing Engine
- SMTP initiates an SMTP session with the remote SMTP host
- After the SMTP session with the remote host has been established, the Information Store retrieves the body of the message and converts the Message as necessary.
- SMTP sends the Message from the Queue to the Remote Host.

The following Exchange Features require the use of SMTP

- Intra Server Message Delivery
- Inter Server Message Delivery
- Message Delivery to the Internet
- Exchange of Routing Information

Intra Server Message Delivery

SMTP will be used for Intra Server Message Delivery for several components like Message Journaling and Message categorization. Exchange Server in the same Routing Group uses SMTP to communicate with each other.

Message delivery to the Internet

SMTP is often used to delivery E-Mail to other Exchange Organizations or other Messaging systems. Exchange Server 2003 can use the Virtual SMTP Server to deliver Messages or one or more Exchange SMTP Connectors or Routing Group Connectors.

Exchange of Routing Information

SMTP is also used to exchange Link State information between Routing Groups.

MX Record

A Mail Exchanger Record (MX Record) is a special DNS record that the e-mail server names for the mail domain. When a Message should be send to that domain a DNS lookup into the Destination DNS domain occurs and is looking for an MX record and a responding A Record. The E-Mail will then be send to the specified Exchange Frontend or Backend Server for Message Delivery.

```
C:\Windows\system32\cmd.exe - nslookup
Microsoft Windows [Version 6.0.5384]
(C) Copyright 1985-2005 Microsoft Corp.

C:\Users\groten>nslookup
Default Server: UnKnown
Address: 192.9.200.240

> set q=mx
> microsoft.com
Server: UnKnown
Address: 192.9.200.240

Non-authoritative answer:
microsoft.com MX preference = 10, mail exchanger = maila.microsoft.com
microsoft.com MX preference = 10, mail exchanger = mailb.microsoft.com
microsoft.com MX preference = 10, mail exchanger = mailc.microsoft.com
>
```

Figure 3: MX Record in NSLOOKUP

Relaying

SMTP Relaying occurs when one SMTP host forwards e-mail to another SMTP host. Open SMTP relaying occurs when the SMTP host accepts messages from recipients outside the Organization and forwards the Messages to other Recipients that are also outside the Organization.

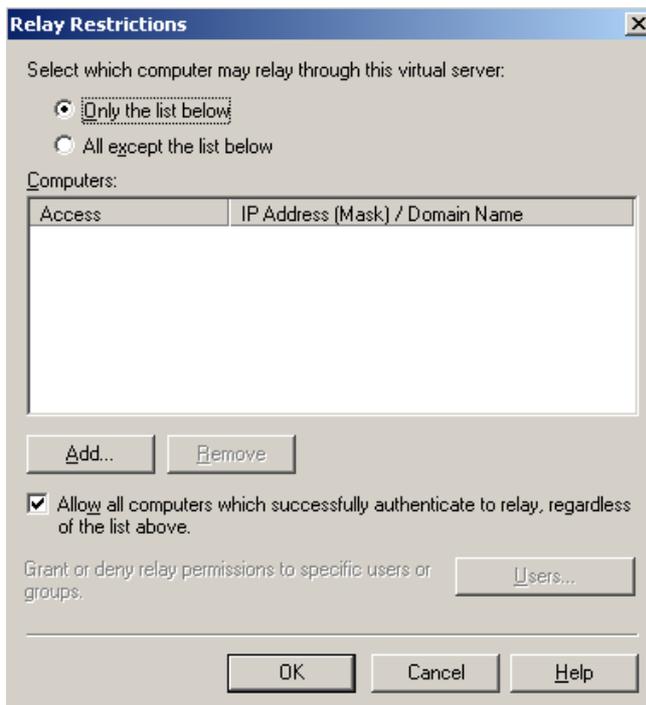


Figure 4: Relaying

If the Exchange Servers allows everyone without Authentication to deliver messages the Server is called a Open Relay. Open Relays can be used to send UCE (Unsolicited Commercial E-Mail). By default Exchange Server 200x is not an Open Relay.

The following steps describe the process:

The unauthorized User sends an e-mail message to the SMTP Server and addresses multiple recipients in the Message and the recipients in the E-Mail are in domains external to the Exchange Servers Messaging Organization.

The Exchange Server accepts the Message

After Exchange has accepted the Message, Exchange delivers this Message to outside because there is no match in the Recipient Policies in the Exchange Organization.

Routing Groups

Exchange Server 2003 support the concept of Routing Groups to control the Message Flow between Exchange Servers. Routing groups are Groups of Servers running Exchange Server 2003 that are connected over permanent Highspeed network links. Within Routing Groups, Exchange Server always transfer Messages over SMTP.

There is one special Server called the Routing Group Master which is responsible for Tracking and Maintaining the Routing Information which are necessary for determining the best path for Message Delivery. The Default Routing Group Master is the first Server in the Routing Group. If you wish to transfer the Routing Group Master Role you must do that manually in the Exchange System Manager.

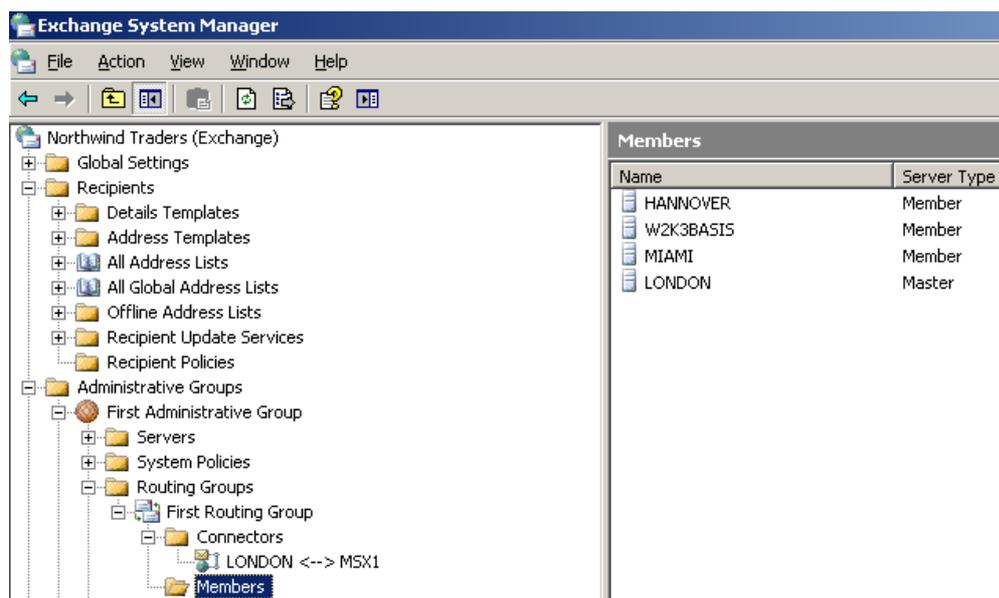


Figure 5: Routing Groups

If your organization has more than one Routing Group you must install a connector between the two or more Routing Groups. The preferred Connector is the Routing Group Connector but you can also use a SMTP- or X.400-Connector.

By default, all Exchange Server Organizations include only a single Routing Group called First Routing Group. All Servers in the Organization are members of the First Routing Group, unless routing membership is modified by you as an Exchange Server Administrator.

You should plan to implement multiple Routing Groups when one ore more of the following conditions occur:

- Network connections are slow or not permanent.
- The network is unreliable or unstable.
- Message transmission is complex and indirect, requiring multiple physical network hops.
- Message transmission must be scheduled between different locations.
- The Routing Group structure is created to prevent users from accessing public folder replicas.

Link State Algorithm (LSA)

Exchange Server 2003 determines the route that an e-mail must take based on the Status and availability of connectors between different Routing Groups and to external Messaging systems through an SMTP connector or other connectors.

Every Exchange Server stores its status information in a Link State Table (LST). The Link State Table is small table which requires about 32 Byte per entry which is hold in the RAM of the Exchange Servers.

All information's will be collected by the Routing Group Master (RGM) of the Routing Group. The Routing Group Master talks with TCP Port 691 with other Exchange Servers in the Routing Group and is responsible for generating / updating the LST and for the distribution of the LST to each Exchange Server in the Routing Group. The updated LST is propagated to other Routing Groups through Bridgehead Servers. The Routing Group Master (RGM) then sends the updated information to the Bridgehead Server, and then the Bridgehead Server sends the information to Bridgehead Servers in other Routing Groups over TCP Port 25.

Address Space	Address Type	Cost	Non-default Restrictions	Connector	Routing Group	Administrative Group
domaene.internal	SMTP	1	No	LONDON <-> MSX1	First Routing Group	First Administrative Group

```

NODEINFO (5) {3B}Version: 6.5.7638.138.1 Microsoft Routing Server ready

MULTI (5) ((0000040C) ORGINFO_RESPONSE a2dffab5a688834f9bcd18dcf07436b4 4f3ac3b9940e30c99ccf134c8691c557 ( 9285b53c9bed854ea08d06a287d1eaf8
db9c4e9e1e01224ab1f593eb6fa3c417 1 0 5 357078cb2a2f6741b27f18ff60232a72 {26} * 9285b53c-9BED-854E-A08D-06A287D1EAF8 {50};c=US;a= ;p=Northwind
Trader;o=Exchange;cn=9285b53c-9BED-854E-A08D-06A287D1EAF8; * {59};o=Northwind Traders;ou=First Administrative Group/* /9285b53c-9BED-854E-A08D-06A287D1EAF8
(20bb6172ffe6364f96c9c242a28a9672 NO 0 1c3a {0} 63eb2da48028254daeca460b26c4c779 NO 0 1dc7 {0} b283d05795841642ba5982fe9f52d9a6 NO 0 1b20 {0}
db9c4e9e1e01224ab1f593eb6fa3c417 YES 1 1dd6 {10}070100000000101 ) (04834f140f84c5409df099d77a31c9ac ( CONFIG {4}SMTP {23}
04834f140f84c5409df099d77a31c9ac_S {f}[192.9.200.101]{66};o=Northwind Traders;ou=First Administrative Group;cn=Configuration;cn=Connections;cn=LONDON <->
MSX1 0 0 0 0 ffffffff 0 1 0 0 0 0 0 0 ) ARROWS ( {4}SMTP {10}domaene.internal 1 ) BH ( 6e3ddce01c724e419a524a0536467788 CONN_AVAL {14}
ToPSecret.domain.tid ) TARGBH ( STATE UP))) )

```

Figure 6: Link State Table

The Link State Table lists all connectors, and their status, in an Exchange Server 2003 organization.

The following information is included in the LST:

Link status

There are only two states for any given link: up or down. For this reason, connection information, such as whether a link is active or in a retry state, is not propagated between servers running Exchange Server 2003, and it is only available on the server involved in the message transfer. Exchange Server 2003 only considers routing messages by using Connectors with a link status of up.

Link cost

The Link State Table stores costs for each Connector. Exchange Server 2003 uses the cost values stored in the link state table to select the least cost route for a message. Costs are configured on each Connector, and Exchange Server 2003 records them in the Link State Table.

Conclusion

In this article I tried to show you how Exchange Server 2003 Message flow works. In the second part of this article I will show you how to use Message Tracking, Message Queues and some other tools such as Winroute to troubleshoot Exchange Message flow.