



MonitorWare Agent 10.0

© 2015 Adiscon GmbH

Table of Contents

Part I Introduction	1
1 About MonitorWare Agent	1
2 Features	2
3 Components	6
Core Components	6
Add-On Components	7
4 System Requirements	9
Part II Getting Started	10
1 Setup	10
2 Creating an Initial Configuration	10
3 Installing LogAnalyzer	11
4 Obtaining a Printable Manual	11
5 Export Settings	12
6 MonitorWare Agent Tutorial	13
Filter Conditions	13
Ignoring Events	14
Logging Events	21
Time-Based Filters	25
Email Notifications	28
Alarming via Net Send	29
Starting Scripts and Applications in Response to an Event	31
Monitoring Hard Disk Space	33
Monitoring External Devices via PING	36
Monitoring FTP Server via a FTP Probe	38
Monitoring SMTP Server via a SMTP Probe	40
Monitoring IMAP Server via a IMAP Probe	42
Monitoring NNTP Server via a NNTP Probe	44
Monitoring External Devices via a Port Probe	46
Part III Common Uses	48
Part IV Step-by-Step Guides	48
Part V Using InterActive SyslogViewer	51
1 About InterActive SyslogViewer	51
Features	51
Requirements	51
2 Options & Configuration	52
Launching InterActive SyslogViewer	52
Using InterActive SyslogViewer	52
Options & Menus	52

File Menu.....	53
Options	53
General Options.....	54
Notifications & Questions.....	56
License	57
Edit Menu.....	58
View Menu.....	60
Help Menu.....	60
Live Syslog View	61
Database View	63
Part VI Configuring MonitorWare Agent	66
1 Client Options	68
2 General Options	71
License Options	71
General	72
Debug	73
Engine	74
QueueManager	77
3 Services	78
Understanding Services	78
Syslog Server	79
SETP Server	85
Event Log Monitor	88
Event Log Monitor V2 (for Vista)	98
SMTP Listener	102
SNMP Monitor	104
RELP Listener	107
Database Monitor	108
Serial Port Monitor	112
File Monitor	116
Heartbeat	123
Ping Probe	124
Port Probe	127
SMTP Probe	130
POP3 Probe	132
FTP Probe	134
NT Services Monitor	136
HTTP Probe	138
IMAP Probe	141
NNTP Probe	143
Disk Space Monitor	144
SNMP Trap Receiver Service	146
CPU/Memory Monitor	148
MonitorWare Echo Reply	153
MonitorWare Echo Request	153
4 Filter Conditions	155
Filter Conditions	155
Global Conditions	158
Operators	159
Filters	160
General	161

Date/Time	164
InformationUnit Type	165
Syslog	167
FTP	169
IMAP	170
HTTP	171
NNTP	172
POP3	173
SMTP	174
Event Log Monitor	175
Event Log Monitor V2	178
NT Service Monitor	180
DiskSpace Monitor	182
File Monitor	183
CPU / Memory Monitor	184
SNMP Traps	185
SerialPort Monitor	188
Custom Property	189
File Exists	191
Extended IP Property	192
Store Filter Results	194
5 Actions	195
Understanding Actions	195
Resolve Hostname Action	195
File Options	196
Database Options	202
OLEDB Database Action	207
Event Log options	210
Mail Options	212
Forward Syslog Options	218
Forward SETP Options	225
Send RELP	226
Send MSQueue	228
Start Program	229
Play Sound	230
Send to Communications Port	232
Post-Process Event	234
Net Send	243
Set Status	244
Set Property	245
Call RuleSet	246
Discard	246
Control NT Service	247
HTTP Request	249
Send SNMP Trap	250
Compute Status Variable	253
Normalize Event	254
Part VII Getting Help	255
Part VIII MonitorWare Concepts	257
Part IX Purchasing MonitorWare Agent	258

Part X Reference	258
1 Comparison of properties Available in MonitorWare Agent, EventReporter and WinSyslog	
2 Event Properties	259
Accessing Properties	259
Property.....	260
FromPos.....	260
ToPos	261
Options.....	262
Examples.....	264
System Properties	265
Custom Properties	266
Event-Specific Properties	266
Standard Properties.....	267
Windows Event Log Properties.....	268
Windows Event Log V2 Properties.....	268
Syslog Message Properties.....	269
Disk Space Monitor.....	270
CPU/Memory Monitor.....	270
File Monitor.....	271
Windows Service Monitor.....	271
Ping Probe.....	271
Port Probe.....	272
Database Monitor.....	272
Serial Monitor.....	273
MonitorWare Echo Request.....	273
FTP Probe.....	273
IMAP Probe.....	273
NNTP Probe.....	273
SMTP Probe.....	273
POP3 Probe.....	273
HTTP Probe.....	273
3 Complex Filter Conditions	273
4 MonitorWare Agent Shortcut Keys	277
5 Command Line Switches	278
6 Version Comparison	279
7 Connect to Computer	279
Part XI Copyrights	280
Part XII Glossary of Terms	280
1 EventReporter	280
2 Millisecond	281
3 Monitor Ware Line of Products	281
4 Resource ID	281
5 RELP	282
6 SETP	282
7 SMTP	283

8	Syslog Facility	283
9	TCP	283
10	UDP	284
11	Upgrade Insurance	284
12	UTC	284
13	NNTP	284
14	SNMP	285
15	FTP	285
16	HTTP	285
17	POP3	285
18	IPv6	285
19	IMAP	286
	Index	287

1 Introduction

1.1 About MonitorWare Agent

MonitorWare is a line of products for network monitoring, management and analysis by Adiscon. MonitorWare has been designed with a philosophy that small to mid sized enterprises have same critical network security needs and limited resources to purchase security solutions.

MonitorWare Agent is a solution for those who are interested in centralizing their monitoring needs. It provides all capabilities of WinSyslog and EventReporter plus many classic features of its own. It has been aided with a reliable event delivery SETP protocol and remotely device monitoring features. It has features to help network administrators/planners, desktop support professionals, system analysts, system administrators, internet managers and technical support personals to make their lives easy.

MonitorWare Agent runs on the systems to be monitored and provides the core functionality. It can gather the data from various sources, like the Windows event log, routers, switches, firewalls and many more. It supports very flexible and powerful local filtering and processing of these gathered events based on a powerful rule processor, events can be forwarded, acted on or discarded - all at the discretion of the system administrator. Even a stand-alone MonitorWare Agent can play a vital role in network management by performing a role like generating alert emails at the occurrences of specific events.

Larger environments consolidate MonitorWare Agent data in a central repository like the MonitorWare event database or combined log files. Database is the source of information for all reporting and analysis modules of the MonitorWare system. By default, database can be created with MySQL, Microsoft Access or Microsoft SQL Server (also available as cost-free MSDE). As standard SQL and ODBC are being used, it is easily adaptable to other database systems. For example, we know that many customers use it successfully with Oracle databases.

A number of different modules work on this consolidated database or the log files to carry out various activities. These modules include scheduled reporting facilities like [MonitorWare Console](#) for analysis, a web interface or [MoniLog](#) reporting.

Currently under development is an enterprise configuration manager, which facilitates configuration of MonitorWare system on an enterprise scope. With the MonitorWare Enterprise Manager, groups of configurations can be created (e.g. for Syslog servers, NT event log monitors, consolidation servers and the like). These function-focused groups can then be automatically applied to machine groups. By doing this, you can monitor a very large MonitorWare system with a single MonitorWare Enterprise Manager. If you are interested in the enterprise configuration system, please mail us at support@adiscon.com to become enrolled in our beta program.

MonitorWare Agent can also integrate with other network monitoring and management related Adiscon products like [EventReporter](#) and [WinSyslog](#). In fact, it uses common terms and methods wherever possible, so upgrading from these solutions to the full MonitorWare system is easy.

For a complete overview over the [MonitorWare line of products](#), please visit

www.monitorware.com.

1.2 Features

Complete Windows Event Monitoring

MonitorWare Agent automatically monitors Windows Event Logs. All Event Logs including the Windows 2000 specific extensions are fully processed. Application log file monitoring provides support for virtually any application that logs to a text file like Web server log files, Oracle error logs files or Windows application log files (like the DHCP log files).

Active Network Probes

Ping and Port Probe services allow monitoring of both local and remote systems and services. These services are not restricted to Windows machines only – virtually any existing service can be used with these probes. Good examples are LINUX based web and mail servers or firewalls. But our probes don't restrict you to an OS – even if you have a server running on a mainframe, MonitorWare can check its operational state.

Failing systems and services are detected and alert be generated.

Monitor Windows' Services and Disk Space Monitor

The Windows service monitor and disk space monitor services check the local machine. Failing services and low disk space are quickly detected and can be used to trigger notifications or even corrective actions before problems arise.

External Events

Events are accepted via a standard Syslog server and hence all of the Syslog-enabled devices can be included in the MonitorWare system. This includes popular devices like routers and switches as well as printers and a large number of UNIX / Linux based systems and applications. Virtually all currently existing network devices support Syslog – so MonitorWare Agent can monitor all of them.

To reach an even broader device range, MonitorWare Agent not only supports standard compatible Syslog but also it supports popular extensions like Syslog over.

Scalability

The MonitorWare system is modular and highly scalable. If a single server is to be monitored, MonitorWare Agent can provide all monitoring and alerting needs. However, multiple MonitorWare Agents in a complex, hierarchical network can talk to each other and provide both local and central alerting and event archiving.

Event Archiving

All incoming events – no matter what source they came from – can be stored persistently. Options include archiving in databases as well as log files.

Alerting

All incoming events – no matter what source they came from – can be stored persistently. Options include archiving in databases as well as log files.

Powerful Event Processing

MonitorWare Agent is powerful and flexible rule engine processes all events based on a configured set of actions. An unlimited number of rules and actions allows tailoring to the specific needs.

Zero-Impact Monitoring

MonitorWare Agent has no noticeable impact on system resources. It is specifically written with minimal resource usage in mind. In typical scenarios, its footprint is barely traceable. This ensures it can also be installed on heavily loaded servers.

Robustness

MonitorWare Agent is written to perform robust even under unusual circumstances. The reliability of the product is proven since 1996.

Ease of Use

MonitorWare Agent is easy to install and configure. Comprehensive step-by-step guides and wizards help administrators with setting up even complex systems.

Firewall Support

Does your security policy enforce you to use non-standard ports? MonitorWare Agent can be configured to listen on any TCP/IP port for Syslog messages.

Syslog Support

NT Event Messages can be forwarded using standard Syslog protocol. NT severity classes are mapped to the corresponding Syslog classes. Codes are fully supported.

Send Syslog Test Message

The MonitorWare Agent client comes with "Send Syslog Test Message". This option

enables you to check if Syslog Messages being sent properly to the destination or not.

SETP Support

NT Event Messages can be forwarded using Adiscon proprietary SETP protocol. Windows Event Logs are monitored successfully as well.

SNMP Trap Receiver

A new feature added in Monitor Ware Agent. It allows receiving SNMP messages.

FTP Probe

It connects to the FTP server and on receiving the response sends the QUIT command to terminate the connection. It saves the connection status and response replies.

HTTP Probe

It connects to HTTP server, receives response and sends the QUIT command to terminate the connection. The connection status and response are saved. It also keeps some additional properties to configure like URL and QueryString, Request Type, Use Secure HTTPs Protocol, Referer and User Agent.

IMAP Probe

It connects to an IMAP server, receives response and sends the QUIT command to terminate the connection. The connection status and response are saved.

NNTP Probe

It connects to NNTP (Usenet) server, receives the response and sends the QUIT command to terminate the connection. The connection status and response are saved.

POP3 Probe

It connects to POP3 (Usenet) server, receives and sends the QUIT command to terminate the connection. The connection status and response are saved.

SMTP Probe

It connects to SMTP (Usenet) server and sends the HELLO command that is automatically constructed by MonitorWare Agent on startup using the full qualified DNS name. The SMTP probe then receives the response and sends the QUIT command to terminate the connection. The connection status and response are saved.

IPv6

Support for IPv6 is available in all network related facilities of the engine. All network related actions will automatically detect IPv6 and IPv4 target addresses if configured. You can also use DNS resolution to resolve valid IPv6 addresses. Network related Services can either use IPv4 or IPv6 as internet protocol. In order to support both protocols, you will need to create two services. The only exception is the RELP Listener, which uses IPv4 and IPv6 automatically if available.

Runs on large Variety of Windows Systems

Windows 4.0, 2000, XP, 2003, Vista, 2008, 7, 8 and Windows 2012; Workstation or Server – MonitorWare Agent runs on all of them.

Multi-Language Client

The MonitorWare Agent client comes with multiple languages ready to go. Out of the box, English, French and German are supported. Other languages will be added shortly. Languages can be switched instantly. Language settings are user-specific; so multiple users on the same machine can use different languages.

Friendly and Customizable User Interface

New Skinning feature has been added to MonitorWare Agent Client. New Cloning feature added to MonitorWare Agent Client helps to clone a Ruleset, a Rule, an Action or a Service with one mouse click. Move up and Move down function has been added for Actions in the MonitorWare Agent Client. Wizards have been enhanced for creating Actions, Services and RuleSets. And other minute changes!

Handling for low-memory cases

MWAgent allocates some emergency memory on startup. If the system memory limit is reached, it releases the emergency memory and locks the queue. That means not more items can be queued, this prevents a crash of the Agent and the queue is still being processed. Many other positions in the code have been hardened against out of memory sceneries.

Multithreaded Queue Engine

The Action processing engine is multithread enabled, which means that the overall processing performance will increase in larger environments and MWAgent will benefit from smp machines.

1.3 Components

1.3.1 Core Components

MonitorWare Agent Configuration Client

The MonitorWare Agent Configuration Client – called "the client" - is used to configure all components and features of the MonitorWare Agent. The client can also be used to create a configuration profile on a base system. That profile can later be distributed to a large number of target systems.

MonitorWare Agent Service

The MonitorWare Agent Service – called "[the service](#)" - runs as a Windows service and carries out the actual work.

The service is the only component that needs to be installed on a monitored system. The MonitorWare Agent service is called the product "engine". As such, we call systems with only the service installed "[engine-only](#)" installations.

The service runs in the background without any user intervention. It can be controlled via the control panel "services" applet or the "Computer Management" MMC under Windows 2000. The client can also be used to control service instances.

x64 Build

The installer inherits the 32bit as well as the 64bit edition. It determines directly, which version is suitable for your operating system and therefore installs the appropriate version. Major compatibility changes for the x64 platform have been made in the Service core. For details see the changes listed below:

- ODBC Database Action fully runs on x64 now. Please note that there are currently very few ODBC drivers for x64 available!
- Configuration Registry Access, a DWORD Value will now be saved as QWORD into the registry. However the Configuration Client and Win32 Service Build can handle these data type and convert these values automatically into DWORD if needed. The Configuration Client will remain a win32 application. Only the Service has been ported to the x64 platform.

A note on cross updates from Win32 to x64 Edition of MonitorWare Agent!

It is not possible to update directly from Win32 to x64 Edition using setup upgrade method. The problem is that a minor upgrade will NOT install all the needed x64 components. Only a full install will be able to do this. Therefore, in order to perform a cross update, follow these instructions:

1. Create a backup of your configuration, save it as registry or xml file (See the Configuration Client Computer Menu)
2. Uninstall MonitorWare Agent.
3. Install MonitorWare Agent by using the x64 Edition of the setup.
4. Import your old settings from the registry or xml file.

1.3.2 Add-On Components

There are a number of optional components available as free downloads.

All optional components work with the MonitorWare common database format.

InterActive SyslogViewer

The InterActive SyslogViewer is a Windows GUI application receiving and displaying Syslog events. It is a Syslog server in its own right. Typically, it is used in conjunction with the WinSyslog service, but it can also be used as a stand-alone Syslog server.

The InterActive SyslogViewer replaces the Interactive display from the pre 4.0 release WinSyslog Client. It was brought into a separate program because there was some confusion about the interactive display in the past.

Though it is not a core component, it is included in the MonitorWare Agent install set.

Adiscon LogAnalyzer

Adiscon LogAnalyzer is a convenient facility to access MonitorWare gathered events over the web. All major browsers are supported.

Adiscon LogAnalyzer is an easy to use solution for browsing Syslog messages, Windows event log data and other network events over the web. Adiscon LogAnalyzer enables the system administrator to quickly and easily review his central log repository. It provides views typically used on log data. It integrates with web resources for easy analysis of data found in the logs.

Mainly it helps to have quick overview over current system activity and accessing the log data while not being able to access the administrator workstation (e.g. being on the road or roaming through the enterprise). While originally initiated to work in conjunction with Adiscon's MonitorWare product line, it can easily be modified to work with other solutions as well.

Adiscon LogAnalyzer is included in the MonitorWare Agent install set. It gets copied onto machine but not installed. For installation of Adiscon LogAnalyzer, refer to the installation instructions in the doc folder of Adiscon LogAnalyzer or see the online manual at <http://loganalyzer.adiscon.com>

MonitorWare Console

MonitorWare Console facilitates the network administrators to gather valuable information about their networks and offers them strong analytical abilities with which they can examine their network proficiently against countless problems including security breaches. Using the Views and Reporting Modules of the MonitorWare Console, you can find the problematic areas in your network very efficiently and promptly. As a network administrator, you would not only like to find the problems but also their solutions. MonitorWare Console's Knowledge Base Module is exactly

meant for this purpose. In short, MonitorWare Console is a very powerful tool that facilitates the Network Administrators to scrutinize their networks from tip to toe and gives an in-depth perspective about what's going on in their system.

For further details, please visit the MonitorWare Console website at www.mwconsole.com

These tools will be available from the tools folder

Logger

Adiscon logger is an UNIX-like logger command line tool for Windows. It is a re-write of the UNIX logger tool with enhanced functionality. All the popular UNIX options are supported. Also, it supports reliable syslog transport via RFC 3195 and plain tcp as found in other Adiscon's products as well as some others like syslog-ng. In addition to that, it includes options specifically needed for the Windows Environment.

For further details, please visit the MonitorWare Console website at www.monitorware.com/logger

LogZip

Adiscon LogZip is a command line tool for Windows for zipping logfiles. Its sole purpose is to pick up log files and put them into a zip file in a location that is specified. By running on the command line, it can be easily used with the Windows Task Scheduler. That way you can automatically archive and store unneeded logfiles in a different location, only keeping the recent logfiles available for review.

LogZip is written in .NET, so please ensure, that you have the .NET framework from Microsoft installed. Else, the tool will not work.

For further details, please visit the MonitorWare Console website at www.monitorware.com/logzip

LogViewer

The Log Viewer tool is able to process very large log files. Other tools often have problems with files large than 100 MB. With Adiscon's Log Viewer there isn't even an issue with files larger than 1 GB. In fact it works with 5 GB almost as fast as with 5 MB files. A special help for reviewing the logs is the highlighting option. Define rules for highlighting any keyword or phrase which should be especially marked while reviewing. These terms can be associated with a color, bold text or similar. E.g. "error" appear in bold text and red color.

For further details, please visit the MonitorWare Console website at www.monitorware.com/log-viewer

1.4 System Requirements

The MonitorWare Agent has minimal system requirements. The actual minimum requirements depend on the type of installation. If the client is installed, they are higher. The service has minimal requirements, enabling it to run on a large variety of machines – even highly utilized ones.

Client

- The client can be installed on Windows 2000 SP3 and above. This includes Windows XP, Windows 2003/2008/2012 servers, Windows Vista and Windows 7/8. The operating system variant (Workstation, Server ...) is irrelevant.
- The client is suited for 32bit and 64bit operating systems. The installer determines the correct version for the operating system by itself.
- The client uses XML technology. Unfortunately, operating system XML support is only available if at least Internet Explorer 4.01 SP1 is installed.
- The client requires roughly 8 MB RAM in addition to the operating system minimum requirements. It also needs around 10 MB of disk space.
- The client is available for Intel based systems, only.

Service

- The service has fewer requirements. Most importantly, it does not need Internet Explorer to be installed on the system.
- It works under the same operating system versions.
- At runtime, the base service requires 5 MB of main memory and less than 2 MB of disk space. However, the actual resources used by the agent largely depend on the services configured.
- If the MonitorWare Agent shall just monitor the local systems event log, impact on the monitored system is barely noticeable, if at all visible.
- If the MonitorWare Agent acts as a central Syslog server receiving hundreds of messages per second, it needs many more resources. Even then, the actual load is depending on the actions carried out. Storing the messages to text files is much less performance intense than writing them to a database table – especially if the database engine is located on the same machine. As such, there is no single guideline for hardware sizing. It needs to be adapted to the expected workload. We have created an article on [performance optimization for syslog server operations](#), which you may want to read.
- Please note however, that the service is specifically optimized to handle high throughput including message bursts (for example received via syslog).
- If you expect high volume burst and carry out time consuming actions (for example database writes), we highly recommend adding additional memory to the machine. Even 64 MB additional memory works nicely. A typical Syslog message (including overhead) takes roughly 1.5 KB. With 64 MB, you can buffer up to 50,000 messages in 64 MB.
- Please note MonitorWare Agent is capable of storing such bursts temporarily in memory even if the machine would otherwise be too slow to process the messages.

Adiscon LogAnalyzer

- Adiscon LogAnalyzer requires Microsoft Internet Information Server (IIS) version 4 or higher to be present on the machine where Adiscon LogAnalyzer is to be installed but it is not mandatory, we recommend to go on with Apache in conjunction with PHP5 such as included in the package [WAMP](#).

2 Getting Started

MonitorWare Agent can be used for simple as well as complex scenarios.

This chapter provides a quick overview of the MonitorWare Agent and what can be done with it.

Most importantly, it contains a tutorial touching many of the basic tasks that can be done with MonitorWare Agent as well as pointer on how to setup and configure.

Be sure to at least briefly read this section and then decide where to go from here - it is definitely a worth time spent.

2.1 Setup

[Installing MonitorWare Agent](#) is simple and easy. A standard setup program installs the application.

A number of different [Download Versions](#) of the product is available. The main difference is whether or not a current version of the Microsoft Windows Installer program is included. If you use recent software (e.g. Windows XP or Windows 2003 Server), you can typically use the small install set. Install sets have different names. Those ending in "max" are typically the version for older operating systems without a current installer. If in doubt, use an install set whose name ends in "max". All files are direct install sets, so there is no need to unzip them or to find a setup.exe or such.

Depending on the download directory, the setup program may also be supplied in a ZIP file.

Furthermore the installer adds a Windows Firewall exception for the service process automatically during the installation routine.

2.2 Creating an Inital Configuration

MonitorWare Agent actually contains the features of five products in it. In order to get MonitorWare Agent running be sure to go through the [Tutorial](#) section, once you have read this section. MonitorWare Agent can work as:

Data gatherer

Here, it gathers event data from important sources like Windows event logs, text files, ping and port probes and the like.

Real Time Alerter

Alert conditions can be detected in real time and alerts be issued. Alerts can be sent

via email and various other means. Alerts based on data gathered by the data gatherers can be configured with respect to different parameters.

Automatic Admin Actions

Depending on certain events, administrative actions can be automatically initiated, for example the deletion of temporary files in a low-disk space condition.

Relay Server

MonitorWare Agent can be used to build, highly scalable, complex systems with relay servers between locations or networks. As a relay server, it forwards incoming events to another instance of MonitorWare Agent or a Syslog daemon.

Event Repository

All gathered event data can be stored in a repository. The repository is a database providing the base for all other MonitorWare products. Events can also be stored in text files. With a specific configuration, a secure log repository can be created for auditing purposes.

MonitorWare Agent can perform any mix of the five functions on a given machine. There are no limits inside the product. Right after installation, however, it is not configured for any of the above functions. So in order to have it do some useful work, it needs to be configured.

2.3 Installing LogAnalyzer

[Adiscon LogAnalyzer](#) is a convenient facility to access MonitorWare gathered events over the web. All major browsers are supported. Adiscon LogAnalyzer is included in the MonitorWare Agent install set. It gets copied onto machine but not installed.

For installation of Adiscon LogAnalyzer, refer to the installation instructions in the doc folder of Adiscon LogAnalyzer or see the online manual at <http://loganalyzer.adiscon.com/doc/install.html>. Please email support@adiscon.com, if you want some more help in this regard.

2.4 Obtaining a Printable Manual

A printable version of the manual can be obtained at <http://www.mwagent.com/help/manual/>

The manuals offered on this web page are in printable (in PDF format) or HTML Versions for easy browsing and printing. The manual is also included as a standard Windows help file with all installations. So if you have the product already installed, there is no need to download these documents.

The version on the web might also include some new additions, as we post manual changes frequently – including new samples and as soon as they become available.

Past manual versions are also available for those customers in need of it.

2.5 Export Settings

When working on a support incident, it is often extremely helpful to re-create a customer environment in the Adiscon lab. To aid in this process, we have added functionality to export an exact snapshot of a configuration. This is done via standard Windows registry files. Please note that when we have received your file, we are also able to make adjustments (if needed) and provide those back to you. This is a very helpful support tool.

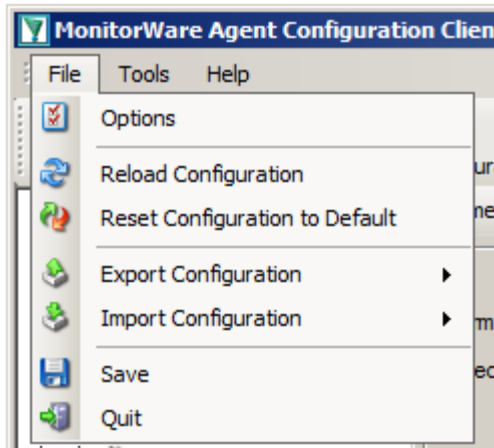


Figure1: Export Settings to a file

To use it, please do the following:

1. Go to "Computer Menu"
2. Choose "Export Settings to Registry-File" (be sure **NOT** to select a binary format - they are only for special purposes. You can also **NOT** review binary files for security-relevant data.) Please also note that you can export in Win32 or x64 format so please choose the right one for your system.
3. Save this registry file.

You may be reluctant to send the registry file because of security reasons. We recommend you to review the contents of the registry file for security purposes with a notepad or any other text editor.

Please Note: We have a 1 MB limit on our mail account. Please zip the registry file and then send it to us. If the file size doesn't reduce after compressing it you should contact Adiscon Support for further instructions.

Fully XML Import & Export of Settings

It is now possible to save the whole configuration as XML. You can edit this XML, duplicate Services, Rules or Actions and reimport the Settings. This is very useful to sort and order large configurations.

2.6 MonitorWare Agent Tutorial

This tutorial is to provide an overview of MonitorWare Agent and some of its typical uses.

It is not a complete product documentary but helps enough to understand and target the application according to your needs.

In the tutorial, we start by describing and focusing on the filter conditions, as these are often needed to understand the usage scenarios that follow below.

MonitorWare Agent gathers network events – or "information units" as we call them – with its services.

Each of the events is then forwarded to a rule base, where the event is serially checked against the different rule's filter conditions.

If such a condition evaluates to true ("matches"), actions associated with this rule are carried out (for example, storing the information unit to disk or emailing an administrative alert).

Debug Logging

When having Debug Logging enabled, the Username used for the Service will be printed below the Version Build number now. This is useful for debugging purposes.

2.6.1 Filter Conditions

For every rule, filter conditions can be defined in order to guarantee that corresponding actions are executed only at certain events.

These filter conditions are defined via logical operators. Boolean operators like "AND" or "OR" can be used to create [complex filter conditions](#).

If you are not so sure about the Boolean operators, you might find the following brush-up helpful:

AND – All operands must be true for the result to be true. Example: AND (A, B): Only if both A and B are true, the result of the AND operation is true. In all other cases, it is false.

OR – If at least one of the operands is true, the end result is also true. Example: OR (A, B): The end result is only false if A and B are false. Otherwise, it is true.

XOR – It yields true if exactly one (but not both) of two operands is true. Example: XOR (A, B): The end result is false if A and B both are True or False. Otherwise, it is true.

NOT – Negates a value. Example: NOT A: If A is true, the outcome is false and vice versa. There can only be a single operand for a NOT operation.

TRUE – Returns true.


FALSE – Returns false.

2.6.2 Ignoring Events

There are some events which occur often and you do not want them to be stored in your log files or either take any action on those.

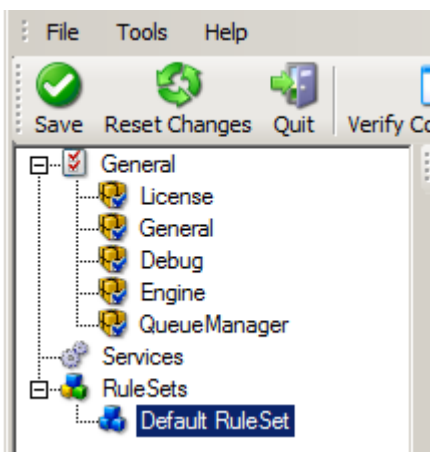
We handle these events on top of our rule set. This ensures that only minimal processing time is needed and they are discarded as soon as possible.

In this tutorial, we define a filter that discards such events. In our example, we assume that Events with the ID 105, 108 and 118 are not required. Please note that for simplicity reasons we only apply filter based on the event ID. In a production environment, you might want to add some additional properties to the filter set. In this sample, no service or rule set is yet defined. It is just a "plain" system right after install, as can be seen in the following screen shot:




Ignoring Events - Figure 1

We begin by defining a rule set. Right-click on "RuleSets" and choose "Add RuleSet" from the context menu. Type in a name of your choice. In this tutorial, we use the name "Defaults". Click on "Next". Leave all as it is in the next dialog. Click "Next", then "Finish". As can be seen in following screen shot, the rule set "Defaults" has been created but is still empty.



Ignoring Events - Figure 2

Of course we can only use a rule if we configure a corresponding service. To do so, right-click on "Running Services" and then select "Add Services". Choose the desired "Service" from the context menu i.e. "Event Log Monitor" in this sample. Provide a name of your choice. In our sample, we call the service "Event Log Monitor". Leave all defaults and click "Next", then "Finish". Now click on "Event Log Monitor" under "Running Services". Your screen should look as follows:

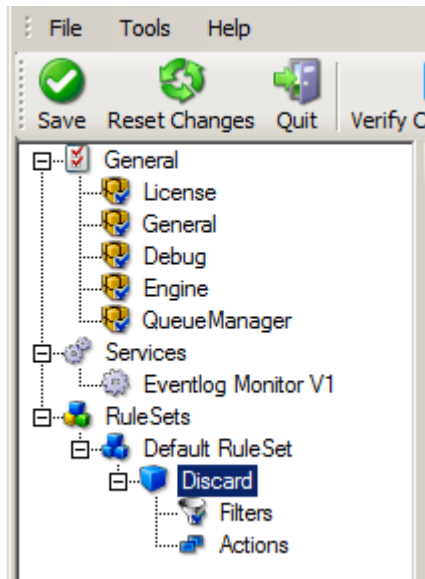


Ignoring Events - Figure 3

As we had created the "Defaults" rule set initially, it is shown as the rule set to use for this service. For our purposes, that is correct. To learn more on the power of rule set assignments, see other sections of this manual.

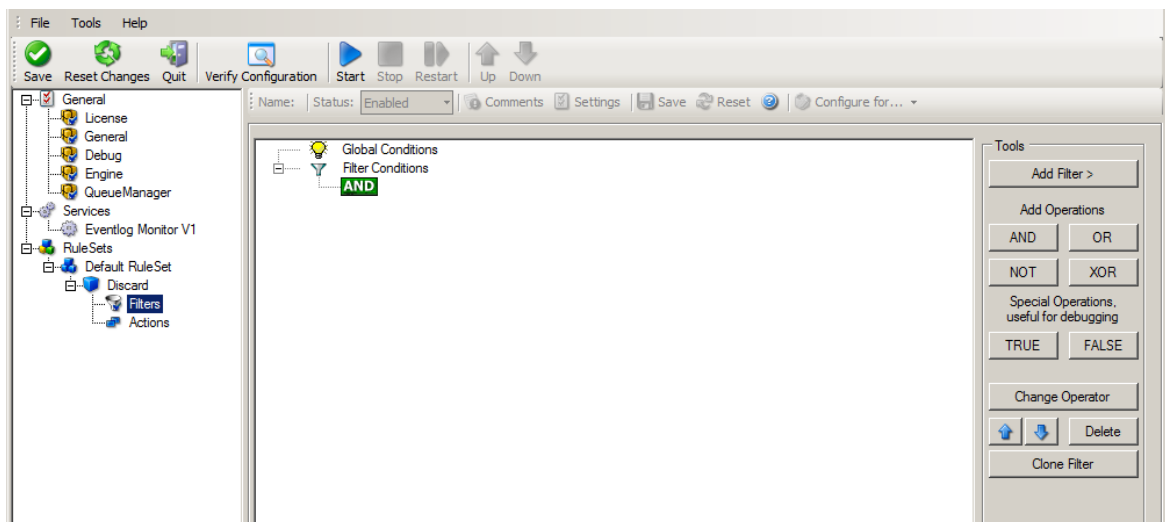
Now we do something with the data that is generated by the event log monitor. To do so, we must define rules inside the rule set.

In the tree view, right-click "Defaults" below "RuleSets". Then, click "Rules", select "Add Rule". Choose any name you like. In our example, we call this rule "Discard". Then, expand the tree view until it looks like the following screen shot:



Ignoring Events - Figure 4

Click on "Filter Conditions" to see this dialog:

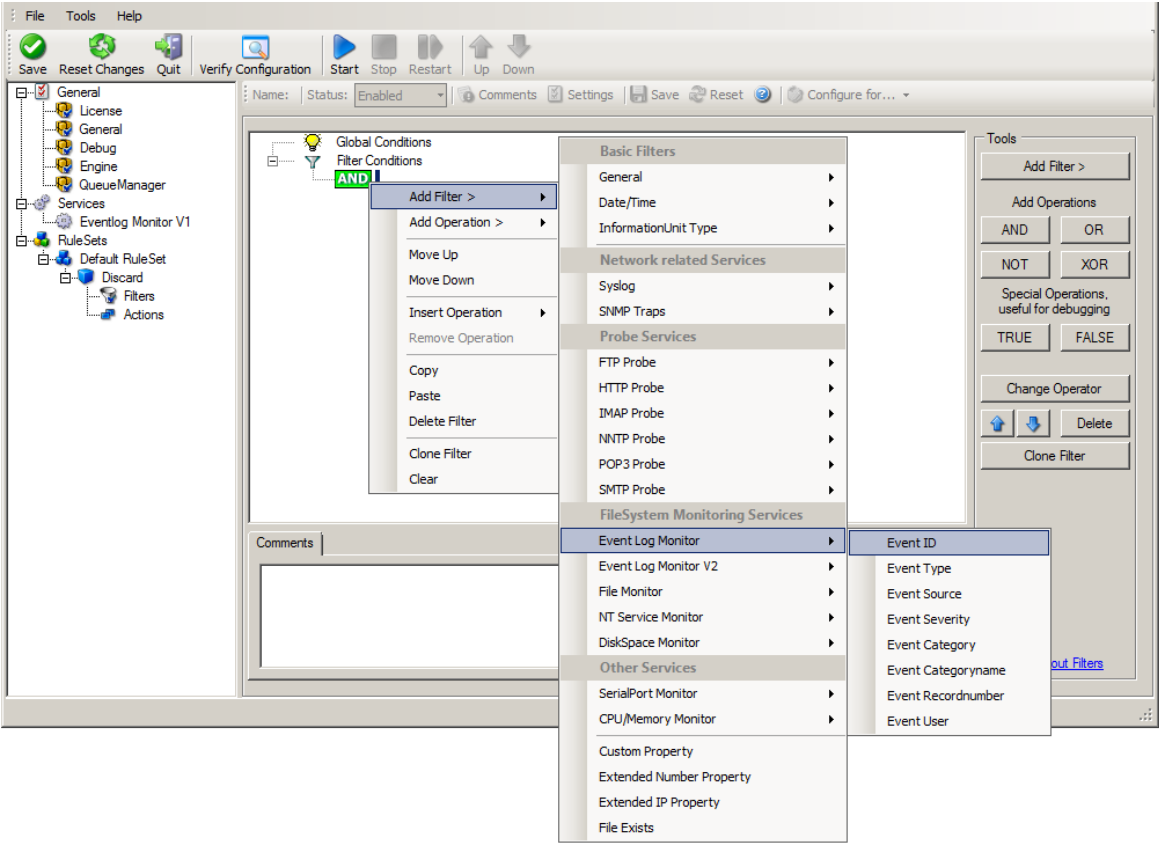


Ignoring Events - Figure 5

In that dialog, we define our filter. Remember: we are about to filter those events, which we are **not** interested in. As we would like to discard multiple events, we need the Boolean "OR" operator in the top-level node, not the default "AND". Thus, we need to change the Boolean operator.


There are different ways to do this. Either double-click the "AND" to cycle through the supported operations or select it and click "Change Operator". In any way, the Boolean operation should be changed to "OR".

We filter out "uninteresting" events via their event id. Again, there are different ways to do this. In the sample, we do it via right clicking the "OR" node and selecting "Add Filter" from the pop up menu. Or you can use the Add Filter Button. This can be seen in the screen shot below:




Ignoring Events - Figure 6

I prefer to add all three Event ID property filters first and later on change the Event ID to the actual value I am looking for. When you have added them, it should look as follows:



Ignoring Events - Figure 7


In order to enter the actual values, select each of the three filters. A small dialog opens at the bottom of the screen. There you enter the values you are interested in. In our sample, these are IDs 105, 108 and 118. As we are only interested in exactly these values, we do a comparison for equality, not one of the other supported comparison modes. When you have made the updates, your screen should look as follows:



Ignoring Events - Figure 8


Save the settings by clicking the (diskette-like) "Save" button. We have now selected all events that we would like to be discarded. In reality, these are often far more or a more complicated filter is needed. We have kept it simple so that the basic concept is easy to understand – but it can be as complex as your needs are.

Now let us go ahead and actually discard these events. This is done via an action. To do so, right-click on "Actions" and select "Discard."



Ignoring Events - Figure 9

Again, name the action as you like in the following dialog. We use "Discard" as this is quite descriptive. Select "Next" and then "Finish" on the next page. Your screen should look like follows:



Ignoring Events - Figure 10


This concludes the definition of our first rule.

If we would start MonitorWare Agent service now, all events with IDs 105, 108 and 118 would be handled by this rule and thus be discarded. All other events do not cause the filter condition to evaluate to true and thus those would be left untouched. Consequently, only these other events flow down to rules defined behind the "Discard" rule. Obviously, our configuration effort is not yet completed. We just finished a first step, excluding those events that we are not interested in. And of course, in reality you need to decide which ones to discard in a real rule set.

2.6.3 Logging Events

Often, a broad range of events (or information units as we call them) needs to be stored persistently so that you can review and analyze them if the need arises. As such, we are in need of a rule that persists the events. In our sample, we choose to work with a text log file (not a database, which we also could use). We now create a rule to store all those events not discarded by the previous rule into a log file.

To do so, right click the "Defaults" rule set as shown below. Then, select "Rules" and "Add Rule":




Logging Events - Figure 1

Use a name of your choosing. In our sample, we call this rule "Write To file". This rule should process **all** events that remained after the initial discard rule. As such, we do not need to provide any filter condition (by default, the filter condition matches always).


Since we want to store all still open Events with help of this rule, we do not require any filter rules here. However, a corresponding action must be defined. Therefore, we just need to define the action:

To do so, expand "Write To file" and right-click "Actions". Select "Add Action", then "Write to File" as can be seen below:



Logging Events - Figure 2

Again, choose a name. Do not modify the defaults. In our sample, we call this action "Records". Click "Next", then "Finish." Now the tree view contains a node "Records", which we select:



Logging Events - Figure 3

Important


If the configured directories are missing, they are automatically created by MonitorWare Agent i.e. the folder specified in "File Path Name".

In our sample, we also change the file base name to "logdata". This was just done out of personal preference. There is no need to do so, but it may be convenient for a number of reasons.

Summary

What did we do so far? All events from the Windows event log are passed through our rule engine and rule filters. Certain events are discarded and the remaining events are stored to a text file on the local disk (for later review or post-processing).

We can now do a quick test: Start MonitorWare Agent by hitting the start button seen below:



Logging Events - Figure 4

The log file should be created in the path you have specified. Open it with notepad. You should see many events originating from the event log. When you re-open the log file, new events should appear (if there were any new events in the Windows event

log). The file is not easily readable. Most probably, you have created it for archiving purposes or to run some external scripts against it. For review, we recommend using either the web interface or the [MonitorWare Console](#).

Please note that the current date is appended to the log file. This facilitates file management in archiving. The format is "logdata-YYYY-MM-DD.log".

You have now learned to define rules and actions. The following chapters thus does not cover all details of this process. If in doubt, refer back to these chapters here.

2.6.4 Time-Based Filters

Time based filters are especially useful for notifications. For example, a user login is typically a normal operation during daytime, but if there are no night shifts, it might be worth generating an alert if a user logs in during night time. Another example would be a backup run that routinely finishes during the night. If we see backup events during the day, something might be wrong.


Similarly, there are a number of other good reasons why specific actions should only be applied during specific time frames. Fortunately, MonitorWare Agent allows defining complex time frames. In this tutorial, though, we focus on the simple ones.

Let us first define a sample time-based filter that applies a nightly time frame. In fact, there are many ways to do this. We have used the method below, because it is straightforward and requires the least configuration work.

To make matters easy, we use this filter condition just to write nightly event log data to a different log file. In reality, time based filters are often combined with other conditions to trigger time based alerts. However, this would complicate things too much to understand the basics.

In the sample below, an additional rule called "Timing Control" has been added. It includes a time-based filter condition. Only if that condition evaluates to "true", the corresponding action is executed. This action can be "Write to Database" or "Write to File". Here we had selected "Write to File" action and renamed it as "Write at Night".

Please note: we use the 12-hour clock system.



Time-Based Filters - Figure 1

All events generated by services binding to our rule set "Defaults" are now also be passed along the "Timing Control" rule set. If these events come in night times between 07:00:01 PM and 5:59:50 AM, the action "Write at Night" is executed.

Please note that the use of the "OR" operator is important because either one of the time frames specified does apply. This is due to the midnight break.


If an event comes in at 08:00:00 AM in the morning, the action is not called – it is outside of the specified time frame:

08:00:00 AM > 07:00:00 PM = *false*
 08:00:00 AM < 06:00:00 AM = *false*

If the very same event comes in at 08:00:00 PM in the filter condition, It evaluates to true and the action is executed.

08:00:00 PM > 07:00:00 PM = *true*
 08:00:00 PM < 06:00:00 AM = *false*

As stated earlier, time frames are most often used in combination with other filters. Here is a more complex example:



Time-Based Filters - Figure 2

In this example, we call the configured actions if events with ID 592 occur between 01:00:01 PM and 08:59:59 (roughly 9 PM). We also execute the configured actions if event ID 593 occurs. Please note that in the case of 593 events, the time filter does not apply due to the used Boolean operations.

In this sample, you also notice that we use an "AND" condition to build the time frame. The reason is that there is no implicit midnight boundary for our time frame as was in the first sample. As such, we need to employ "AND" to make sure the events are WITHIN the specified range.

Now let us look at some sample data:

We receive a 592 event at 07:00:00 AM sharp:

Event ID = 592	= <i>true</i>
07:00:00 AM > 01:00:00 PM	= <i>false</i>
07:00:00 AM < 09:00:00 PM	= <i>false</i>
"AND" Branch	= <i>false</i>
Event ID = 593	= <i>false</i>

In all, the filter condition is false.

Now, the same event comes in at 02:00:00 PM:

Program start ID = 592	= <i>true</i>
Event ID = 592	= <i>true</i>
02:00:00 PM > 01:00:00 PM	= <i>true</i>
02:00:00 PM < 09:00:00 PM	= <i>true</i>
"AND" Branch	= <i>true</i>
Event ID = 593	= <i>false</i>

This time, the time frame is correct, yielding to an overall true condition from the "AND" branch. That in turn yields to the filter condition as whole to evaluate to true.

In this example still is another Event ID. All events with event ID 593 is grasped. This happens independently from the timing control when grasping the Events 592.

One last sample. At this time, event 593 comes in at 07:00:00 AM:

Program start ID = 593	= true
Event ID = 592	= false
07:00:00 AM > 01:00:00 PM	= false
07:00:00 AM < 09:00:00 PM	= false
"AND" Branch	= false
Event ID = 593	= true

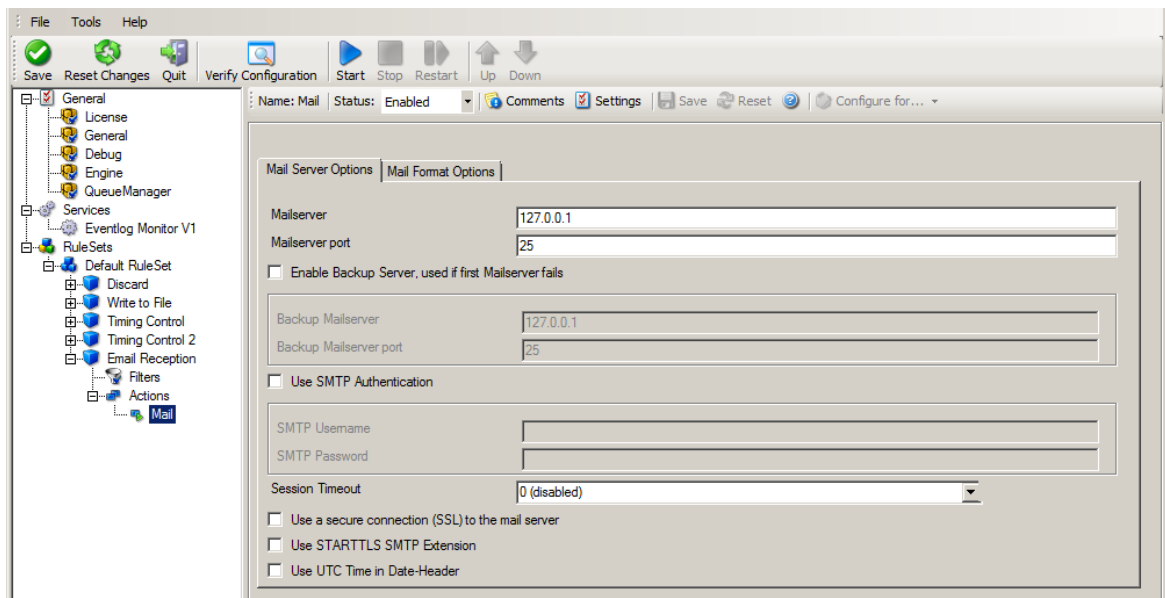
This time the filter condition evaluates to true, too. The reason is that the (not matched) time frame is irrelevant as the other condition of the top-level "OR" branch evaluates to true (Event ID = 593).

2.6.5 Email Notifications

In this example, we would like to receive email notifications when certain events happen.

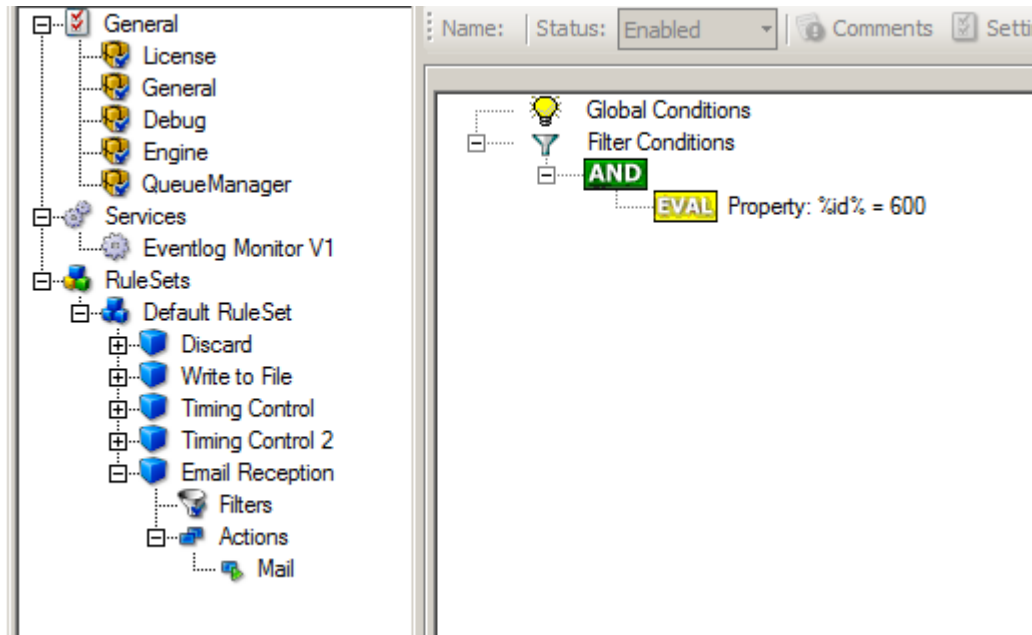
So let us create an additional rule for that purpose: Right-click the "Defaults" rule set and select "Rule Sets", "Add Rule" from the pop up menu. Provide a name. We call it "Email Reception" in this example. Then, add a "Forward via Email" action. In the action details, be sure to configure at least the mail server, recipient and subject properties.

Please note that many mail servers also need a valid sender mail address or otherwise deny delivery of the message.



Email Notifications - Figure 1

Then, select the filter conditions. Let us assume we are just interested in events of ID 600. Then the filter conditions should look as can be seen below:



Email Notifications - Figure 2

When you have finished these steps, be sure to save the configuration and re-start the MonitorWare Agent service. After the restart, the newly extended rule set is executed. In addition, the rules defined so far, the new one is carried out, emailing all events with ID 600 to the specified recipient.

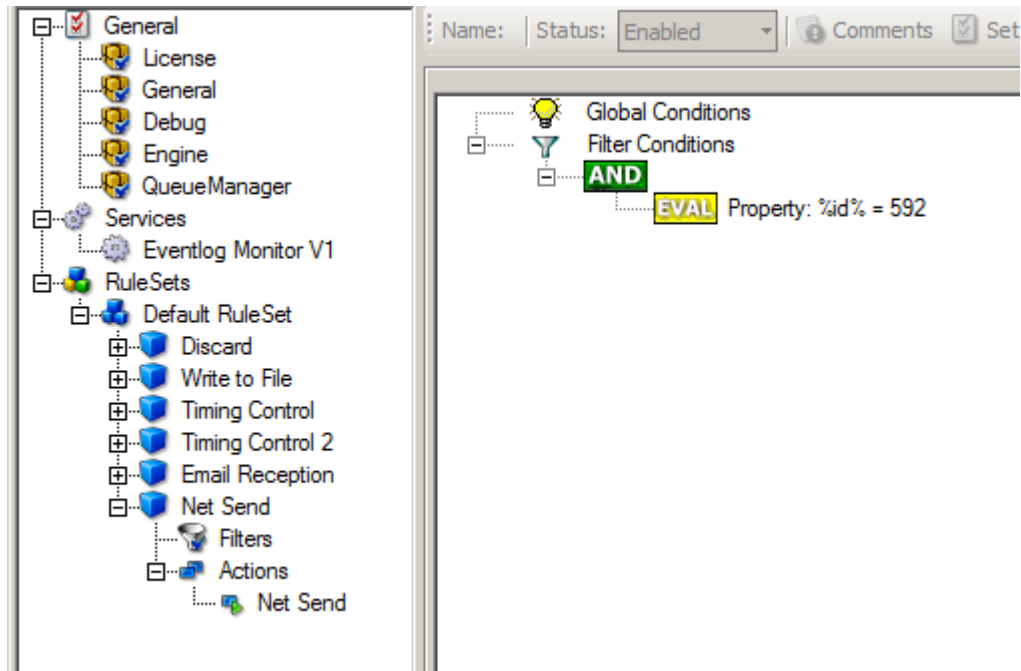
2.6.6 Alarming via Net Send

Again, we add another rule to our rule set.

This time, we would like to receive notification via the Windows messenger service (aka "net send").


Please bear in mind that the Windows messenger service is not the instant messaging service that many people nowadays associate with it. The messenger service is meant for administrator notifications. If a windows workstation (or server) receives a message via that service, a message box pops up on that workstation and the user needs to press an "OK" button to continue. No interaction is possible.

We create a new rule in our rule set "Defaults". In this case, we assume that we receive messenger notifications for all events with Event ID 592. In a real use case, you make sure that this is a real important event, or chances are good you become overwhelmed with messaging windows. A better example could be a filter that checks for a server running low on disk space (using the disk space monitor).



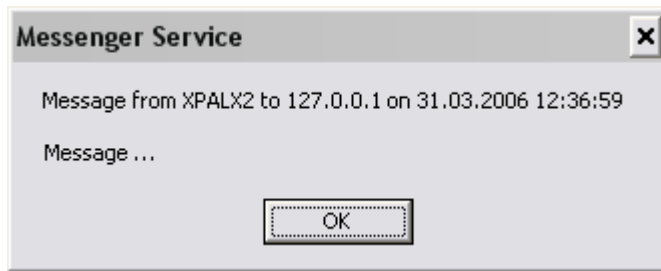
Alarming via Net Send - Figure 1

This time, we use the "Net Send" action as can be seen below. The target field holds either the name or IP-Address of the workstation this message should be sending to. The message text itself goes into "Message to send".



Alarming via Net Send - Figure 2

After saving the configuration and restarting the MonitorWare Agent, we receive notifications if the filter condition evaluates to true. A sample message might look like this (slightly obscured in this sample):

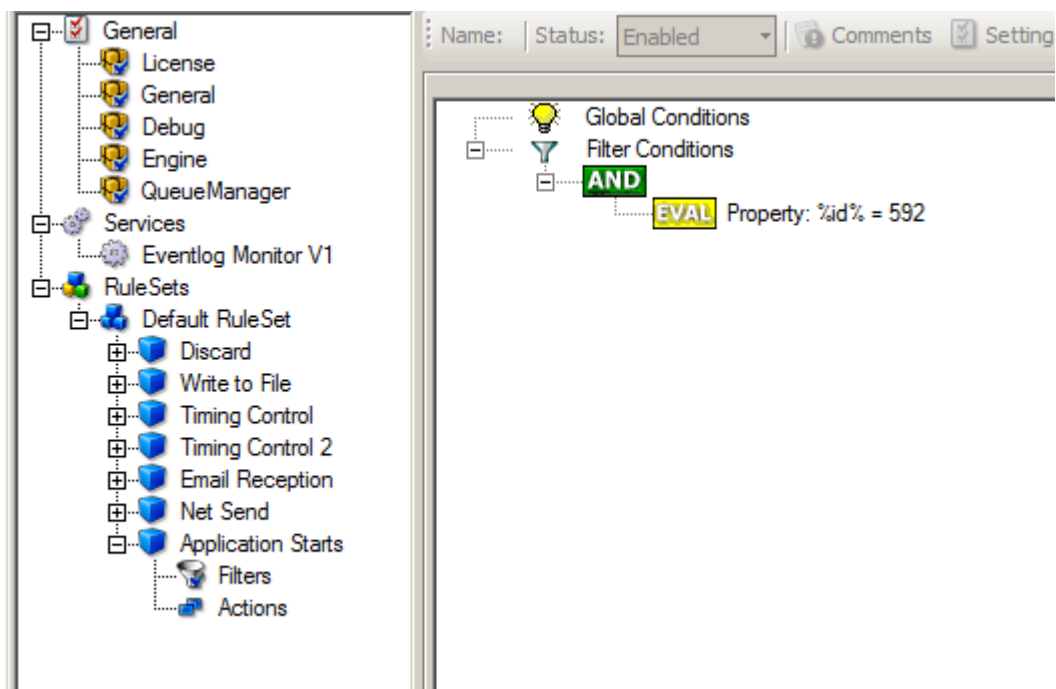


Alarming via Net Send - Figure 3

2.6.7 Starting Scripts and Applications in Response to an Event


We now want to start an application or a script when certain events occur. Typically, this is done to start administrative scripts or corrective action. For example, if a disk runs low on space, you could start a script that deletes temporary files, or if a service fails, a script could restart it.

Our sample, on the other hand, is kept quite simple again. We just show how to generically start an exe file. To do so, we define a new rule, name "Application starts" below. Again, we use the imaginary event 592 as a filter condition. Therefore, the application starts whenever event 592 comes in.




Starting Scripts and Applications in Response to an Event - Figure 1

The start program action is just a "normal" action:



Starting Scripts and Applications in Response to an Event - Figure 2

In the "Start Program" action's parameters, select the file to run as well as all parameters to be supplied to it (if any):



Starting Scripts and Applications in Response to an Event - Figure 3

Once this configuration is done, the program is executed as soon as an event matching the filter condition comes in.


2.6.8 Monitoring Hard Disk Space

Monitoring hard disk space solves at least two purposes: it can be used to generate alerts or trigger corrective actions if a system runs out of free space. It can also be used as a statistical tool to monitor disk space utilization over time.

In our tutorial, we configure a simple disk space monitor and define a rule that stores the results into a text file that can be analysed later. Of course, we could have added trigger conditions for alerts and such. We have not done this, to keep things simple.

As always, we create the needed rule set first. In our sample, we call it "DiskSpace". **Please note that this time we actually create a "Rule Set", not just an additional rule in the "Defaults" rule set.** The reason is that for our purpose, it is much easier to define a specialised rule set and then bind this specialised rule set to the disk space monitor. If we would use the generic "Defaults" rule set, we had to make sure that our filter conditions would only match when an event of type disk space monitor would come in. In addition, it would require more processing time, as all rules and condition filters would be processed – a process that is not needed as we deal with a specific case. As such, it is more appropriate to define a specific rule set, which is then only used for the disk space monitor. What is appropriate in your environment depends on your needs. There is no general rule.

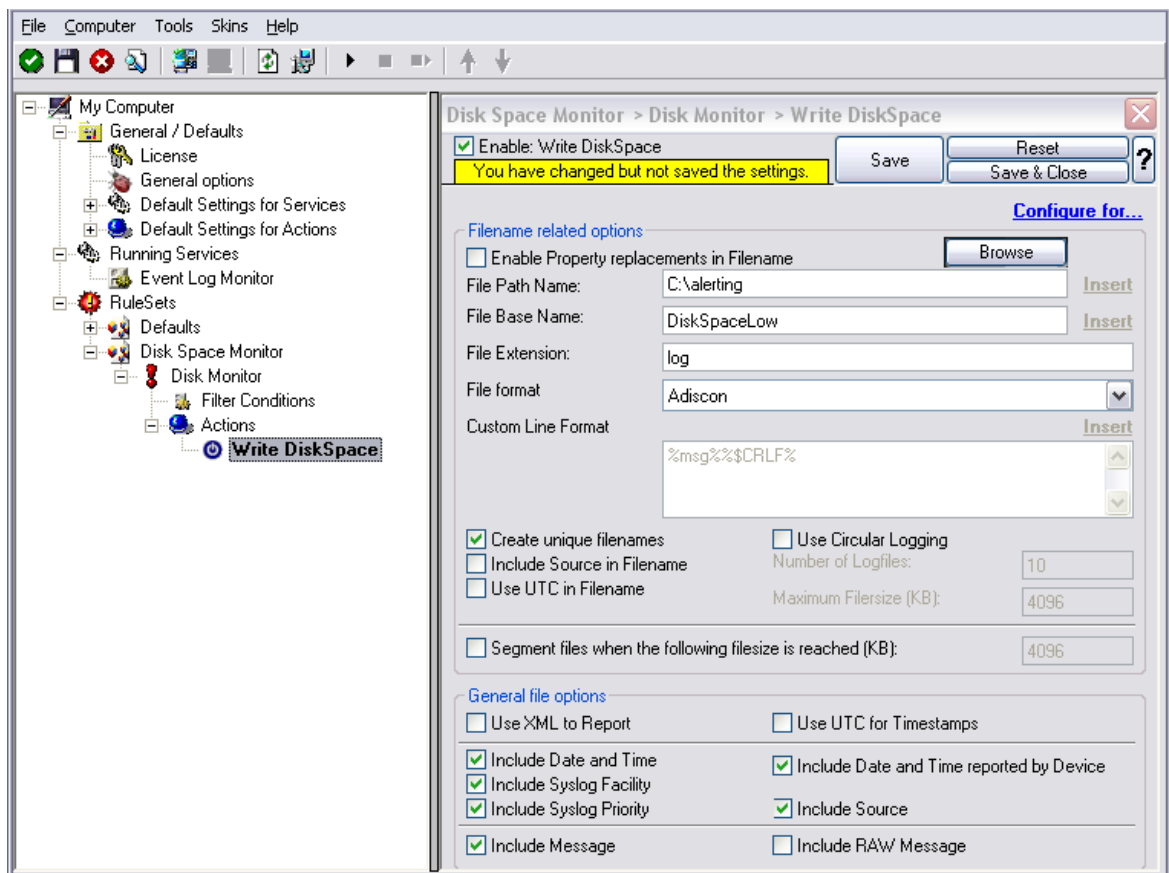
Inside the new rule, we create a filter condition that evaluates to true only if the reported disk space has less than 20 gigabytes of free space. So we log date only when we potentially have constrained disk space. The filter looks as follows:



Monitoring Hard Disk Space - Figure 1

To create this filter, select "Disk Space Monitor", then "Disk Space Left" when pressing the "Add Filter" button.

As I said initially, we use the "Write to File" action in this sample. The action is called "Write DiskSpace" as can be seen below. We could also have used other actions, including emailing, to alert an administrator or start a script to delete temporary files.



Monitoring Hard Disk Space - Figure 2

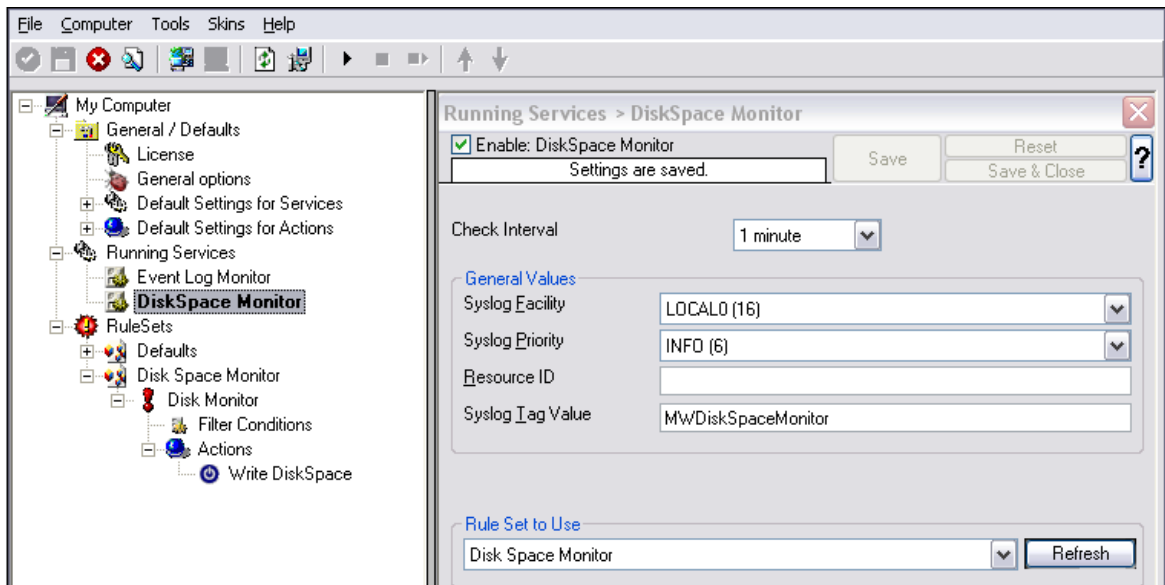
Please note: you should make sure that the base name is different from other "write to file" actions. Otherwise data might get mixed up in the files.

Having created the new rule set, we now need to create the disk space monitor

service itself. It is the part of the software that actively goes out and monitors the disk space. To create it, right-click "Running Services" and select "Add Service", then "DiskSpace Monitor" as seen below:

Monitoring Hard Disk Space - Figure 3


When the wizard starts, you need to name the new service. We use "DiskSpace Monitor" in our sample. Leave the default settings and click "Next" and "Finish".



Monitoring Hard Disk Space - Figure 4

Please note that when you select the new service, it is typically bound to the "Defaults" rule set (as seen above). We need to change this, as we have created the specific "Disk Space Monitor" rule set. Change the "Rule Set to Use" to update it to the new binding.

Save the configuration and restart the service. After a few moments, the disk space log file should fill up (**if there is less than 20 GB of free space on the monitored system**). In notepad, it looks like follows:




Monitoring Hard Disk Space - Figure 5

2.6.9 Monitoring External Devices via PING

In this sample, we use the ping probe to monitor the availability of external devices. The ping probe issues a standard IP "PING". Each system that is "pinged" provides a reply to the system initiating the ping. When the reply comes back, the initiator knows that the pinged system is up and running. **Please note that this does not imply that all services on that machine are running.** To check this, a port probe must be used. At least the ping probe can detect failing systems. It can also be used in any case, whereas the port probe can only be used with TCP based services.


As first step, we create a new rule set. Please see the previous example for the reasoning of doing so. We call the new rule set "Ping Probe". We would like to receive email notifications if the ping probe fails. So we add a "Send Email" action. After doing so, the screen looks as follows:



Monitoring External Devices via PING - Figure 1

Please note we do not customize the Send Email action properties in this sample. In your environment, you need to use some meaningful settings.

Now that we have defined the rule set, we need to create the corresponding service. To do so, right-click "Running Services" and follow the screen shot below:




Monitoring External Devices via PING - Figure 2

Use a name of your choosing, leave the defaults as is and click "Next" and then "Finish". We have used the name "Ping Probe" in our sample.

Click the newly created service. We need to uncheck the "Generate an event if Ping Probe was successful" check box. If it is checked, an event is generated every time. If unchecked, it is generated only when the ping fails. As we are just interested in failed systems, we uncheck it. Therefore, we do not need to apply any other filters. If you forget to uncheck this option, you receive multiple emails – one each time the Ping Probe runs and probes the configured system.

Your screen should now look as follows:



Monitoring External Devices via PING - Figure 3

Now save the settings and restart the service. Whenever the Ping Probe fails, you receive mail. This mail looks as follows:

Event message:

Facility: 16

Priority: 6

Source: 172.19.0.1

Message:

PingProbe Status="error" remoteip="172.19.0.1" PingStatus="11003"
ErrorMessage="Destination Host Unreachable"

A ping probe service can monitor a single device in this version of MonitorWare Agent. Therefore, if you would like to monitor multiple devices, you need to create multiple ping probe services.

2.6.10 Monitoring FTP Server via a FTP Probe

This sample is very similar to the ping probe and port probe samples directly above. Thus, we describe it briefly, only.


FTP probe is used to make a connection to the FTP server and then it receives the

response from FTP server and sends the QUIT command to terminate the connection. The connection status is saved in the property **ftpstatus** and the response in the property **ftprespmsg**.

In our sample, we probe a FTP server, which typically listens to port 21 (the default port for FTP). We send an email alert if the FTP probe cannot connect successfully to the FTP server.

Because this sample is so close to the previous ones, we do not create a new rule set specifically for email alerting. Please view "Ping Probe" for it. This is a good sample of rule set re-use. It might be clever to rename the rule set in such a case. For simplicity reasons, we have not done this here.


Therefore, we begin by creating the new service, done by right-clicking "Running Services":



Monitoring FTP server via a FTP Probe - Figure 1

Use a name of your choosing, leave the defaults as is and click "Next" and then "finish". We have used the name "FTP Probe" in our sample.

After doing so, select the newly created service in the tree view to look at its properties. Be sure to bind it to the "FTP Probe" rule set as seen below:



Monitoring FTP server via a FTP Probe - Figure 2

Save the configuration and restart the service. From now on, the following mail alert is generated when the port cannot be connected to:

Event message:

Facility: 16

Priority: 6

Source: 192.168.1.1

Message:

FTPProbe status="fail" target="192.168.1.1" port="21" netstate="10061"
message="Couldn't connect to host"

2.6.11 Monitoring SMTP Server via a SMTP Probe

This sample is also very similar to the ones directly above. Thus, we describe it briefly, only.


SMTP probe is used to make a connection to the SMTP server and then it receives the response from SMTP server and sends the QUIT command to terminate the connection. The connection status is saved in the property **smtpstatus** and the response in the property **smtprespmsg**.

In our sample, we probe a SMTP server, which typically listens to port 25 (the default port for SMTP). We send an email alert if the SMTP probe cannot connect successfully to the SMTP server.

Because this sample is so close to the previous ones, we do not create a new rule set specifically for email alerting. Please view "Ping Probe" for it. It might be clever to rename the rule set in such a case. For simplicity reasons, we have not done this

here.

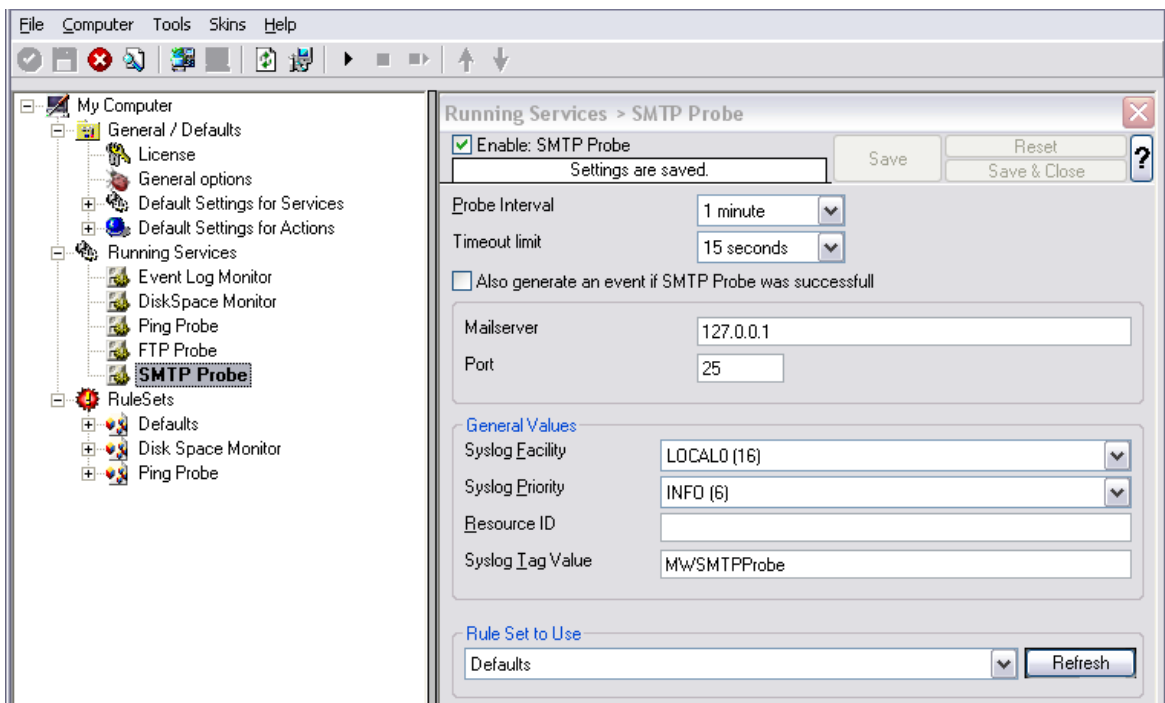
Therefore, we begin by creating the new service, done by right-clicking "Running Services":



Monitoring SMTP server via a SMTP Probe - Figure 1

Use a name of your choosing, leave the defaults as is and click "Next" and then "finish". We have used the name "SMTP Probe" in our sample.

After doing so, select the newly created service in the tree view to look at its properties. Be sure to bind it to the "SMTP Probe" rule set as seen below:



Monitoring SMTP server via a SMTP Probe - Figure 2

Save the configuration and restart the service. From now on, the following mail alert is generated when the port cannot be connected to:

Event message:

Facility: 16

Priority: 6

Source: 192.168.1.1

Message:

SMTPProbe status="fail" target="192.168.1.1" port="25" netstate="10054"

message="Receive call failed"

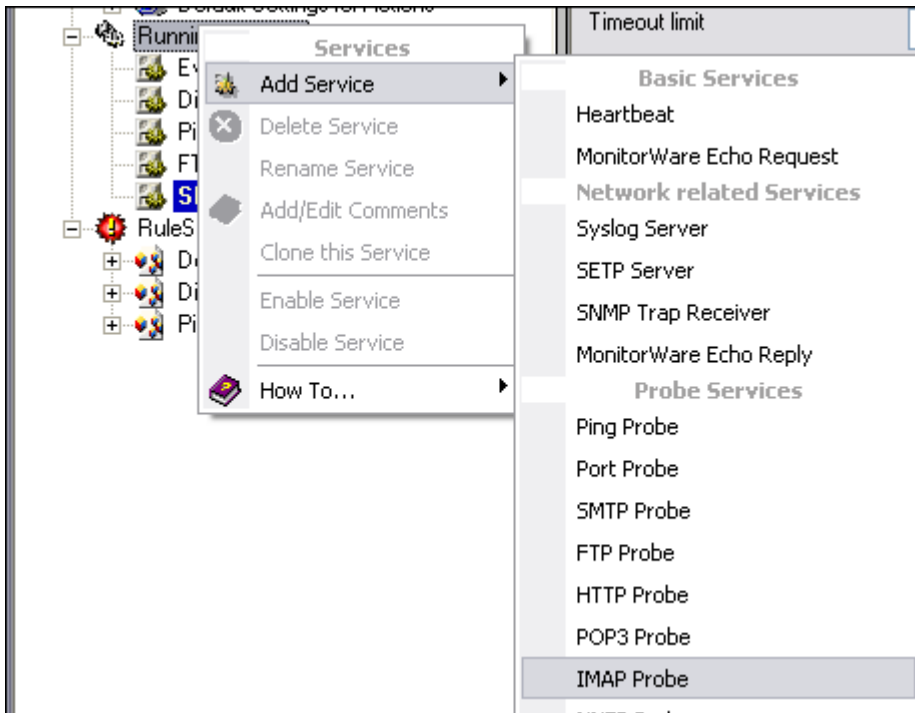
2.6.12 Monitoring IMAP Server via a IMAP Probe

IMAP probe is used to make a connection to the IMAP server and then it receives the response from IMAP server and sends the QUIT command to terminate the connection. The connection status is saved in the property **imapstatus** and the response in the property **imaprespmsg**.

In our sample, we probe a IMAP server, which typically listens to port 143 (the default port for IMAP). We send an email alert if the IMAP probe cannot connect successfully to the IMAP server.

Because this sample is so close to the previous ones, we do not create a new rule set specifically for email alerting. Please view "Ping Probe" for it. It might be clever to rename the rule set in such a case. For simplicity reasons, we have not done this here.

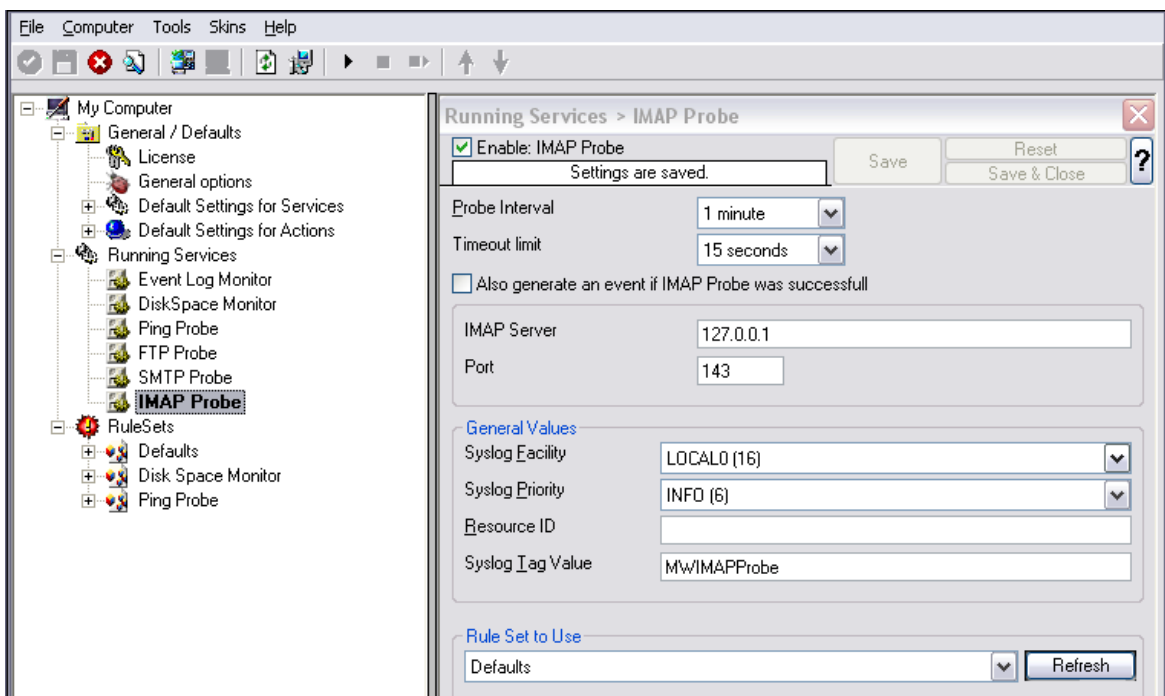
Therefore, we begin by creating the new service, done by right-clicking "Running Services":



Monitoring IMAP server via a IMAP Probe - Figure 1

Use a name of your choosing, leave the defaults as is and click "Next" and then "finish". We have used the name "IMAP Probe" in our sample.

After doing so, select the newly created service in the tree view to look at its properties. Be sure to bind it to the "IMAP Probe" rule set as seen below:



Monitoring IMAP server via a IMAP Probe - Figure 2

Save the configuration and restart the service. From now on, the following mail alert is generated when the port cannot be connected to:

Event message:

Facility: 16

Priority: 6

Source: 192.168.1.1

Message:

IMAPProbe status="fail" target="192.168.1.1" port="143" netstate="10061"
message="Couldn't connect to host"

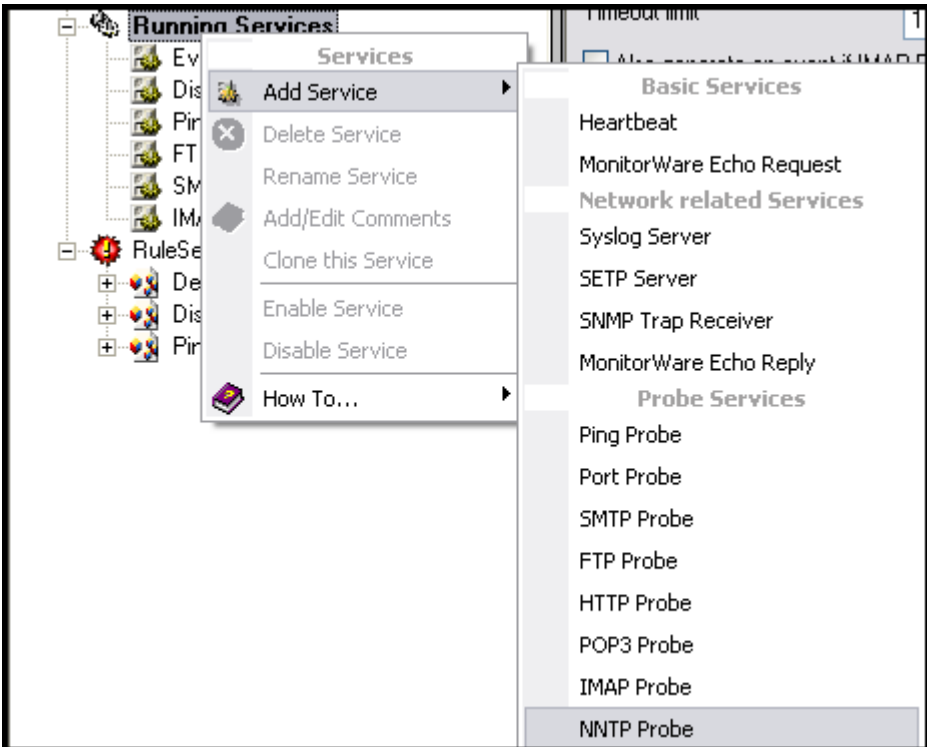
2.6.13 Monitoring NNTP Server via a NNTP Probe

NNTP probe is used to make a connection to the NNTP server and then it receives the response from NNTP server and sends the QUIT command to terminate the connection. The connection status is saved in the property **nntpstatus** and the response in the property **nntprespmsg**.

In our sample, we probe a NNTP server, which typically listens to port 119 (the default port for NNTP). We send an email alert if the NNTP probe cannot connect successfully to the NNTP server.

Because this sample is so close to the previous ones, we do not create a new rule set specifically for email alerting. Please view "Ping Probe" for it. It might be clever to rename the rule set in such a case. For simplicity reasons, we have not done this here.

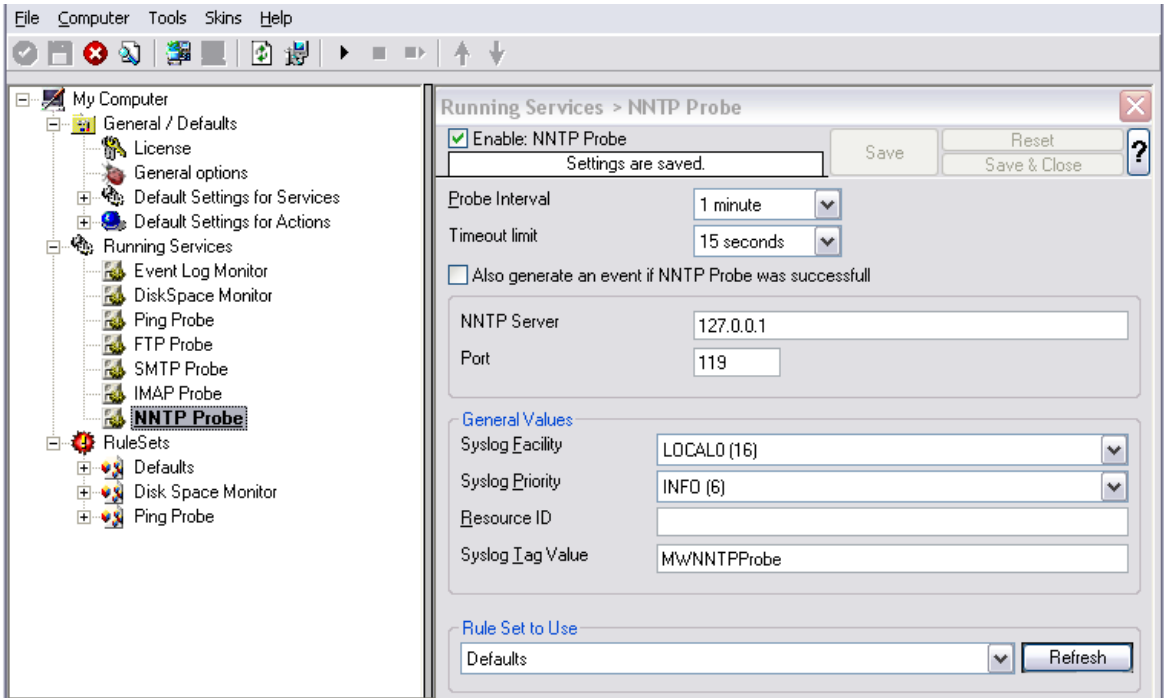
Therefore, we begin by creating the new service, done by right-clicking "Running Services":



Monitoring NNTP server via a NNTP Probe - Figure 1

Use a name of your choosing, leave the defaults as is and click "Next" and then "finish". We have used the name "NNTP Probe" in our sample.

After doing so, select the newly created service in the tree view to look at its properties. Be sure to bind it to the "NNTP Probe" rule set as seen below:



Monitoring NNTP server via a NNTP Probe - Figure 2

Save the configuration and restart the service. From now on, the following mail alert is generated when the port cannot be connected to:

Event message:

Facility: 16

Priority: 6

Source: 192.168.1.1

Message:

NNTPProbe status="fail" target="192.168.1.1" port="143" netstate="10061"

message="Couldn't connect to host"

2.6.14 Monitoring External Devices via a Port Probe

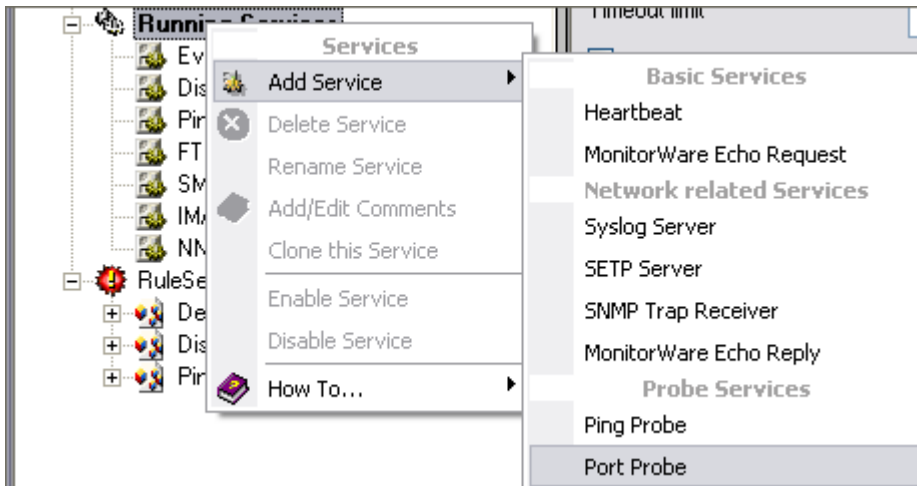
This sample is very similar to the ping probe sample directly above. Thus, we describe it briefly, only.

The main difference between the ping probe and the port probe is that the port probe tries to connect to a specific TCP port. As such, it can only be used with TCP based services like mail server, web servers or ftp servers. For the very same reason, the port probe does not only check the status of the machine it is connecting to but rather if a specific, **service** is available. Let us assume you are interested in monitoring a mail server. If you do a ping probe, the mail server itself might have died while the machine is still running. The ping probe cannot detect this. The port probe, on the other hand, directly connects to the mail server, e.g. on port 25 (the default SMTP port). If the mail server has died, it will probably not answer this connection request and thus the port probe is able to detect the failing state of the service.

In our sample, we probe a web server, which typically listens to port 80 (the default port for http). We will send an email alert if the port probe cannot connect successfully to the web server.

Because this sample is so close to the previous one, we do not create a new rule set specifically for email alerting. It is already covered in the "Ping Probe". This is a good sample of rule set re-use. It might be clever to rename the rule set in such a case. For simplicity reasons, we have not done this here.

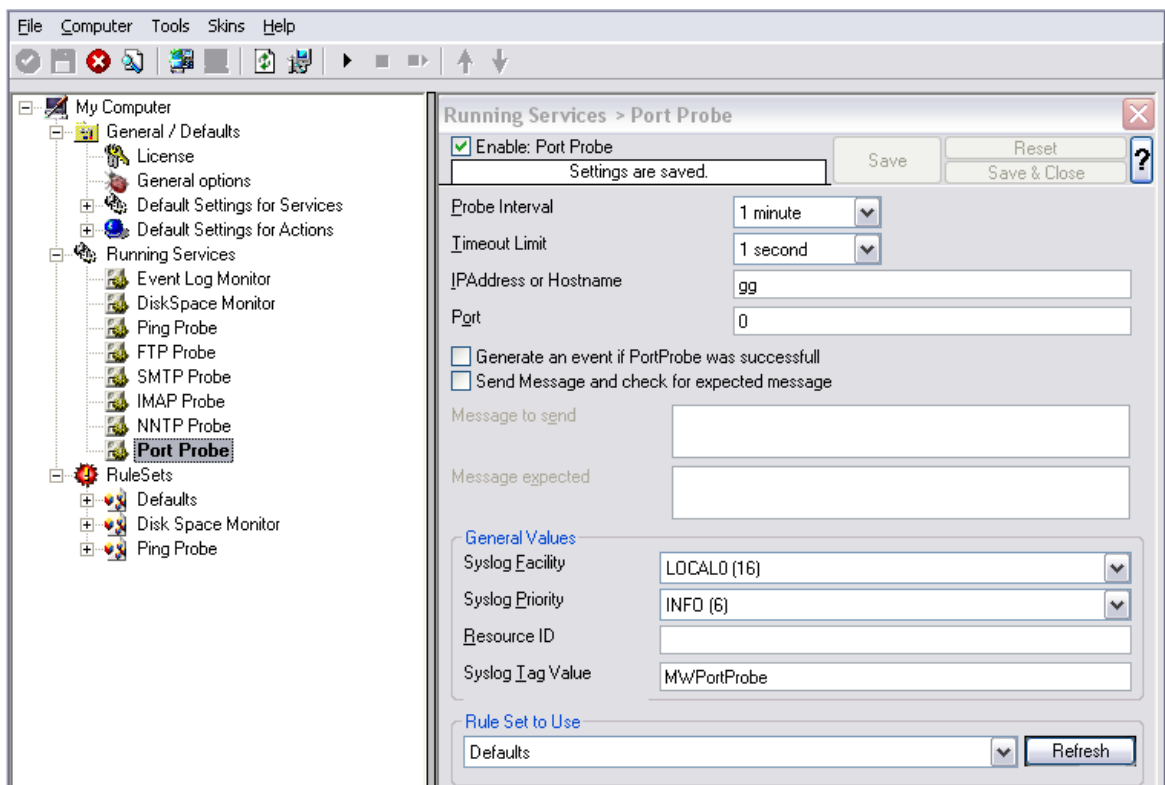
Therefore, we begin by creating the new service, done by right-clicking "Running Services":



Monitoring External Devices via a PortProbe - Figure 1

Use a name of your choosing, leave the defaults as is and click "Next" and then "finish". We have used the name "Port Probe" in our sample.

After doing so, select the newly created service in the tree view to look at its properties. Be sure to bind it to the "Ping Probe" rule set as seen below:



Monitoring External Devices via a PortProbe - Figure 2

Save the configuration and restart the service. From now on, the following mail alert will be generated when the port cannot be connected to:

Event message:
Facility: 16
Priority: 6
Source: 192.168.1.1

Message:
PortProbe status="fail" target="192.168.1.1" port="80" netstate="10065"
message="Couldn't connect to host"

3 Common Uses

MonitorWare Agent can be used in a multitude of ways to perform well in many different environments serving many different needs. We have set up some web pages to address these questions. This allows us to add timely comments, should need arise. Please follow the links below to access the web pages with detailed descriptions.

In general, there are four main use cases for MonitorWare Agent:

[Analysis](#)
[Event Archival](#)
[Alerting](#)
[Solving Problems](#)

Besides these main cases, there are also some other scenarios, like [relaying event data](#).

While reading the scenarios, please keep in mind that MonitorWare Agent is extremely flexible. A single instance on a single machine can be configured to perform all actions and functions concurrently. They are grouped here for easier lookup, but this in no way implies that the Agent can do only one thing or the other.

4 Step-by-Step Guides

The step-by-step guides are meant to get you started quickly. They provide information on how to configure the product in common scenarios. Each section includes the information necessary to complete a specific task.

The information is presented in an easy to follow "step by step" way (hence the name). Each section begins with the intended result and then explains the steps to achieve it in the correct order. They are documented together with hardcopies, so they should be easy to follow. For best results, please be sure to follow the exact order of the steps.

The step-by-step guides do eventually not include all information that might be relevant to the situation. Please use your own judgment if the scenario described sufficiently matches your need.

In the step-by-step guides, we assume the product is already successfully installed

but no configuration has been done. If it is not installed, please do so first.

All step-by-step guides assume that the client is running. This is kind of a step 0 for all the guides.

To keep download times reasonable, the step-by-step guides are not included in this manual. They are kept as separate web pages. This also allows us to modify and add step-by-step guides. Additions are made all the time, so it is probably a good idea to check the [Step-by-step guide page](#) for new guides.

As of this writing, the following step-by-step guides were available:

MonitorWare Agent Specific

Installations and Configurations

- [How to perform a mass update rollout?](#)
- [How to perform a mass rollout?](#)
- [Store IIS Logfiles into a Database](#)
- [How do I export the configuration and create a debug file?](#)
- [How do I monitor the Windows 2003 DHCP Logfiles?](#)
- [How do I monitor the Windows Update log?](#)
- [How do I enter the license information from the product delivery email?](#)
- [How to create a failover syslog server](#)
- [Forwarding filtered IIS Logfiles](#)
- [How do I can do Database Logging with MSSQL in MonitorWare Agent](#)
- [How do I apply filters in MonitorWare Agent?](#)
- [How to apply filters to only get interactive logon/logoff events?](#)
- [How To setup php-syslog-ng with MonitorWare Products?](#)
- [How to monitor ISA logfiles?](#)
- [Forwarding NT event logs to a Syslog server](#)
- [Forwarding NT event logs to an SETP server](#)

Services

- [How do I create a simple Syslog server](#)
- [How To Configure a Syslog Server](#)
- [How To Configure a SETP Server](#)
- [How To setup the EventLogMonitor Service](#)
- [How To setup the EventLogMonitor V2 Service](#)
- [How To setup the File Monitor Service](#)
- [How To setup the NT Service Monitor Service](#)

Actions

- [How To setup the Forward via Syslog Action](#)
- [How To setup an SETP Action](#)
- [How To setup a Write to File Action](#)
- [How To setup the Forward via EMail Action](#)
- [How To setup the Set Property Action](#)
- [How To setup the Set Status Action](#)

- [How To setup the Start Program Action](#)
- [How To setup the Control NT Services Action](#)
- [How To Create a Rule Set for Database Logging](#)
- [How to store custom properties of a log message in a database](#)

Centralized Monitoring / Reporting

- [How To setup PIX centralized Monitoring \(MonitorWare Agent 5.x & MonitorWare Console 3.x\)](#)
- [How To setup Windows centralized Monitoring \(MonitorWare Agent 5.x & MonitorWare Console 3.x\)](#)
- [How To setup a central log server for Windows machines and syslog sending devices \(MonitorWare Agent 4.x & EventReporter 8.x\)](#)
- [Centralized Event Reports with MonitorWare Console](#)

MonitorWare Agent All Versions

Installations and Configurations

- [How to create a failover syslog server](#)
- [MonitorWare Agent Database Formats](#)
- [Database Logging with MSSQL](#)
- [How do I apply filters in MonitorWare Agent, WinSyslog and EventReporter?](#)
- [How To Setup MonitorWare Agent/ WinSyslog/ EventReporter](#)
- [Configuring Windows for the Event Log Monitor](#)
- [Sample Syslog device configurations](#)
- [Intrusion detection via the Windows event log](#)
- [Firewall setup for MonitorWare Agent](#)
- [Creating a hardened log host](#)

Services

- [How To Configure a Syslog Server](#)
- [How To setup SETP Server Service](#)
- [How To setup EventLogMonitor Service](#)
- [Creating a simple Syslog server](#)
- [Forwarding NT event logs to a Syslog server](#)
- [Forwarding NT event logs to an SETP server](#)

Actions

- [How To setup an SETP Action](#)
- [Creating a rule set for database logging](#)

Centralized Monitoring / Reporting

- [How To setup Windows centralized Monitoring \(MonitorWare Agent 4.x & MonitorWare Console 2.x\)](#)
- [How To setup PIX centralized Monitoring \(WinSyslog 7.x, MonitorWare Agent 4.x & MonitorWare Console 2.x\)](#)
- [How to setup PIX centralized Monitoring](#)

- [How to setup Windows centralized Monitoring \(Common\)](#)
- [How to setup PIX centralized Monitoring](#)
- [Centralized Event Reports with MonitorWare Console](#)
- [Centralized event reports with Monilog](#)
- [How To Report Log Truncation](#)

You may also want to visit our syslog device configuration pages at <http://www.monitorware.com/en/syslog-enabled-products/>. They contain instructions on setting up several devices for syslog.

5 Using InterActive SyslogViewer

5.1 About InterActive SyslogViewer

InterActive SyslogViewer is a tool that let's you review your syslog data very easy. It is a separate syslog server, that simply displays all incoming data. By this you can see directly what is happening.

5.1.1 Features

Fast and Easy syslog Viewing

The SyslogViewer allows you to directly view and review syslog messages. Therefore you can react much better on occuring problems or check if everything is ok.

Review stored logs from a database

You can as well directly review log entries in a database. Simply enter the login details and thats it. You can then review your logs and even filter the view. That helps you to find the important data in an easy way.

Export selected data

You can export selected data for further manual processing, like sending an email to your colleague for informing them about what is happening.

5.1.2 Requirements

Any Windows-NT based operating system like Windows 2000, XP or Vista.

You need .NET 2.0 framework installed in order to run Adiscon's Syslog Viewer.

Hardware requirements:

- 32MB RAM

5.2 Options & Configuration

InterActive SyslogViewer is an add-on to the MonitorWare Agent and WinSyslog. **Please note that it is a utility program, with a primary focus on real-time troubleshooting.**

InterActive SyslogViewer is **not** meant to continuously monitor a system. This is what the service is designed for. While Interactive SyslogViewer allows to view current syslog traffic, the service should be used for all other purposes, like creating log files.

5.2.1 Launching InterActive SyslogViewer

To run the InterActive SyslogViewer, click the "SyslogViewer" icon present in the Programs Folder -> MonitorWare Agent/WinSyslog located in the Start menu.

It can also be launched from the command prompt:

- Open a Command Prompt window
- Change to the drive and directory where the MonitorWare Agent is installed.
- Type "InteractiveSyslogViewer.exe" and hit enter.

Available Command Line parameters are:

`/?` = Show Options
`/autolisten` = Start Syslog Server automatically
`/port=10514` = Overwrites the configured port
`/windowpos 0,0,512,800` = Sets default window positions

5.2.2 Using InterActive SyslogViewer


InterActive SyslogViewer is an add-on to the MonitorWare Agent and WinSyslog. **Please note that it is a utility program with a primary focus on real-time troubleshooting.**

Interactive Syslog Server is **not** meant to continuously monitor a system. This is what the service is designed for. While Interactive Server allows to view current Syslog traffic, the service should be used for all other purposes, like creating log files.

5.2.3 Options & Menus

Please find more information about the different menus and options in the respective sub-category.

5.2.3.1 File Menu



File Menu

Options

This will open the Options dialog. Please see the sub-chapters for more details on this.

Run Configuration Client

This option will open the configuration client of MonitorWare Agent/WinSyslog. Here you can do detail configuration of the service.


Minimize to tray

This will minimize the InterActive SyslogViewer window and remove it from the taskbar. You can open it again by double-clicking on the icon in the system tray.

Quit

By clicking here, InterActive SyslogViewer stops receiving data and it will close the application.

5.2.3.1.1 Options



General Options Tab

Enable Resolving Syslog RealSource (Adiscon specific)

With this option enabled, you can see the real source in multiply forwarded messages. That means, you can see the system that forwarded the message and the system where the message originates from.

Allow multiple SyslogViewer instances

You can have multiple instances of the InterActive SyslogViewer by activating this option. This allows you to have multiple forwarding servers sending on different ports and receive their messages separately.

Enable Trayicon (Hide Main Windows on Close)

Enable this to have a tray icon. This enables a soft-close. InterActive SyslogViewer will stay active, but the window will be completely hidden except the tray icon. By double-clicking on the icon, the window will show again.

Autostart the SyslogServer on Startup

Enable this to start the syslog server directly when starting InterActive SyslogViewer.

Switch between two colors for each row in the Datagrid

To have a better overview over the syslog data, activate this option.

Display Syslog message counter in the Datagrid


You can enable a counter by checking the box here. It will count further, even if the maximum of messages is already exceeded.

Max Messages in the Datagrid

Here you can adjust the maximum messages that will be available in the datagrid. By increasing this value, you can store more messages for direct review. **Please note, that increasing the maximum number of messages will have a severe impact on your memory.**

Select UI Language

Here you can choose your favorite language for the InterActive SyslogViewer. By default it is english. You can choose german or japanese as well.



Notifications & Questions Tab

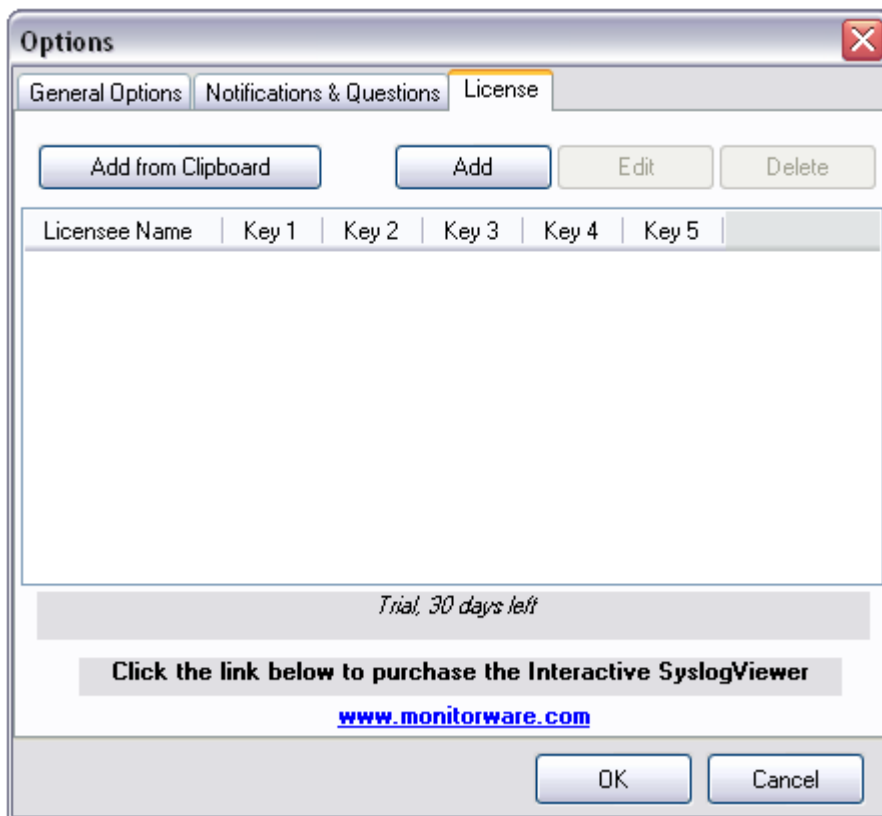
Always ask to reload application on language change

While the box is checked, InterActive SyslogViewer will ask to reload the application on a language change. This is, because the language file can only be loaded while starting the application and not while it is running.

Show warning if more than 10000 records are taken from the database

By activating this option, you will be warned, if the records in the database are just too much. This is to prevent the machine from receiving too much load. Polling lots of messages from a database can have a severe impact on the performance of the machine.

5.2.3.1.1.3 License



License Tab

Here you can insert the license. You have several options:

Add from Clipboard

This will insert the license you have currently on your clipboard.

Add

This button is to manually add a license manually. A new window will open, which shows you the form for entering the license information. This consists of a license name and five blocks of numbers.

Edit

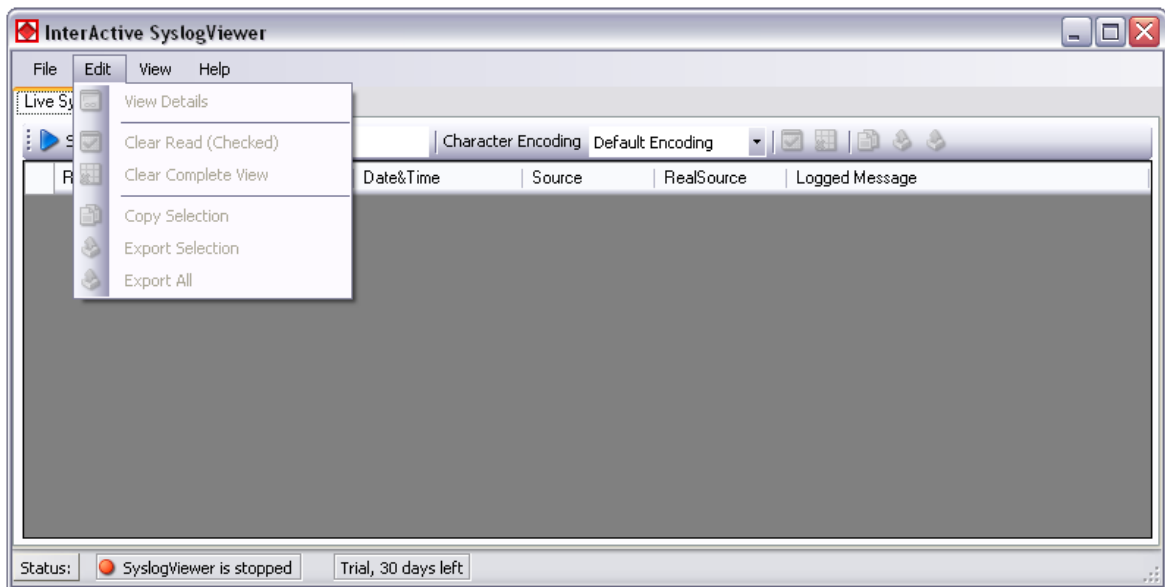
Once a license is entered, it can be changed afterwards. This is done with this button. Mark the license you want to edit and click the button. A window will open which looks just like when adding a license, but the marked license details are inserted already. You can edit every field separately.

Delete

If a license is not needed anymore, you can delete it from the license screen. Mark the license and hit the button. The license will be deleted directly.

Please note, that the screen will give you additional information. You have an overview of the licenses used and if not entered correctly it will show how long your trial period still is.

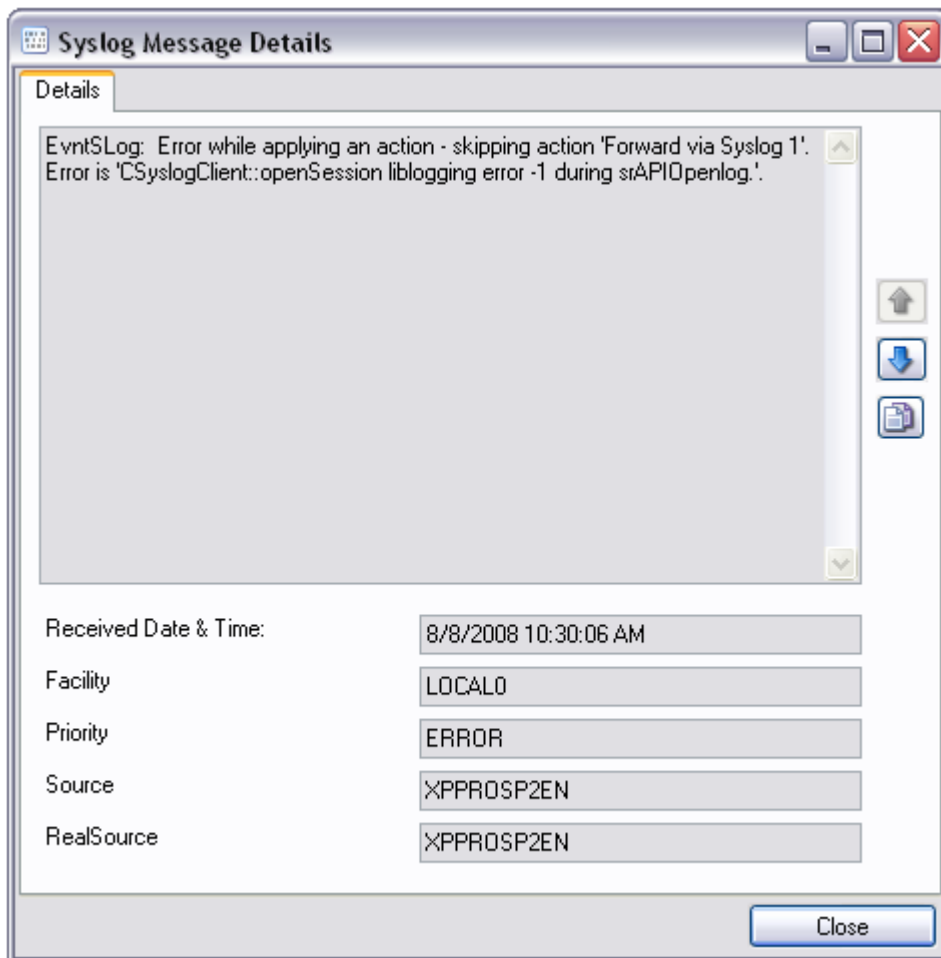
5.2.3.2 Edit Menu



Edit Menu

View Details

When using this option, another window will open up, which shows the details of this event in a more readable view. This could look like this:



Syslog Message Details

Clear Read (Checked)

By activating this, you can clear the checkboxes of the items your marked as read.

Clear Complete View

This option will clear the screen and remove all received data from the view.

Copy Selection

Having selected one or mutiple entries, you can copy them using this function.


Export Selection

Instead of copying you can extract the selected data into a text file.

Export All

Or you directly export all the data that is currently in the list.

5.2.3.3 View Menu

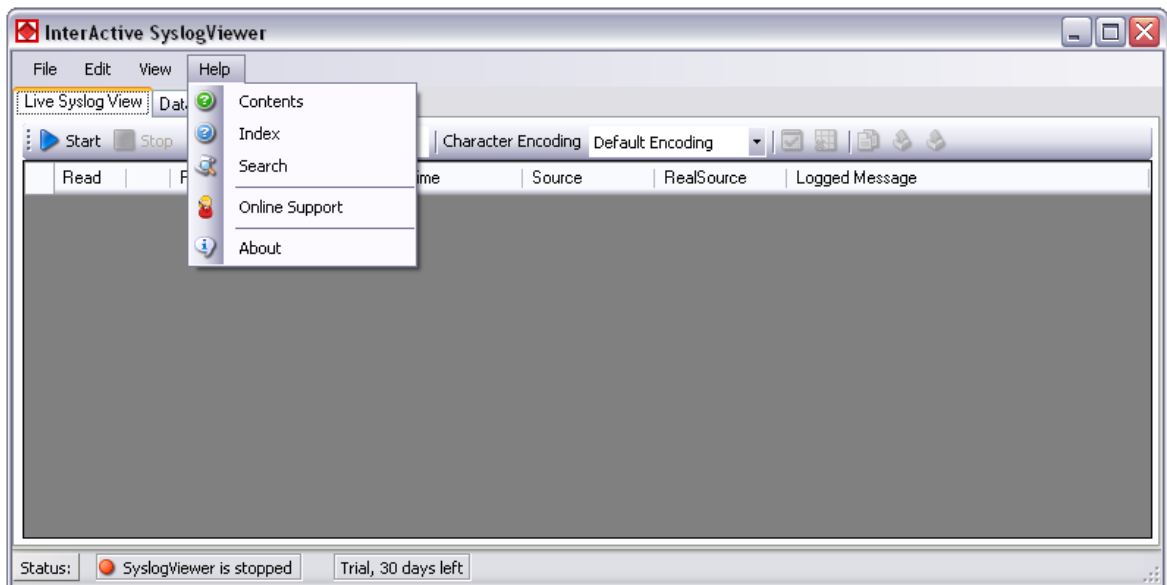


View Menu

Always on Top

This option is very self-explanatory. While activated, the InterActive SyslogViewer window will stay on top of all other applications, so you will have all incoming log data directly in your point of view.

5.2.3.4 Help Menu



Help Menu

Contents

Show the manual.

Index

Show the manual index.

Search

Search the manual.

Online Support

By clicking here, a browser window will open and you will be directed to our support website.


About

The About-window will give you additional information to the tool, like the program version.

5.2.4 Live Syslog View

Primarily, the InterActive SyslogViewer is used for viewing current syslog traffic. All messages are shown in a list with the most important information. These are the Priority, Facility, Date&Time, Source, RealSource and the Message. At the beginning of each line you can see the number of the logged event and a checkbox, for you to track if a message has been read.

The status bar at the bottom of the screen shows you, if the SyslogViewer is running or stopped and how much time you have left for the trial or your licensing status.



Live Syslog View

The toolbar provides you with direct access to the most important functions. These are described here:

Start

With the start button, you start the receiving service. Now the InterActive SyslogViewer will receive and display all incoming messages. If messages were sent before starting the service, they will be dropped.

Stop

Here you can stop the receiving server.

Syslog Port

Here you can define the syslog port where the Viewer should be receiving the syskigness

Character Encoding

Here you can define how characters will be decoded. You can choose from Default Encoding (depending on OS), Ascii, Unicode, UTF8 or UTF32.

Clear checked

With this button, you can clear all the checkboxes in front of the messages.

Clear View

By clicking on this button, all data will be deleted from you datagrid.

Copy Selection

This helps you copying the selected messages.

Export Selection

You can export the selected data directly by using this button.


Export All

Export the complete data that is in the data grid.

5.2.5 Database View

Another feature is the possibility to review log messages which are stored in a database.

The status bar at the bottom of the screen shows you, if the SyslogViewer is running or stopped, how many records are currently shown and how much time you have left for the trial or your licensing status.



Database View

The toolbar in this case is for entering the login information for the database.

System DSN

Specify the System DSN of your database here.

Username

The username for the database.

Password

The appropriate password for the database.

Store Username and Password

With the checkbox you can tell the InterActive SyslogViewer to keep the username and password or not. This is to make usage easier for you.

Limit Records

This limits the maximum of the shown records. The default value is 10000. If changed, this can have a enormous impact on your machine.

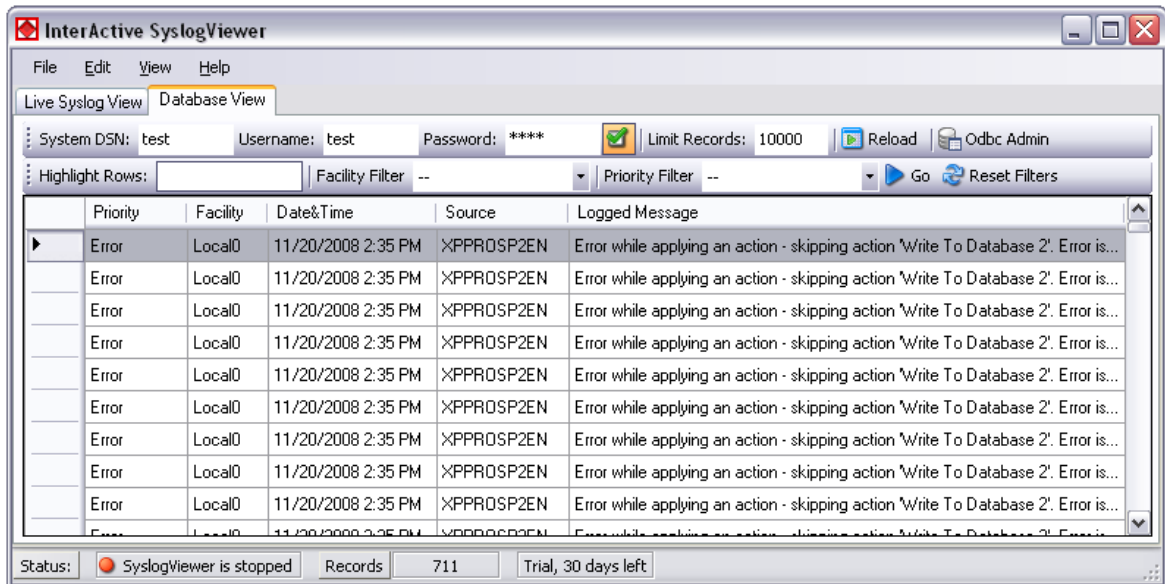
Reload

This button is to reload the database. This is needed to view if there are new log messages in the database.

Odbc Admin

This button opens the Administration Panel for ODBC Data Source connections

Once a database connection is successfully established, you can see another toolbar with the filter options:



Active Database View

Highlight Rows

You can enter a keyword into the field, the rows containing this keyword will be highlighted. You can then find the messages much easier,

Facility Filter

Allows you to only show messages with a certain facility. You can use the dropdown menu to specify the facility.

Priority Filter

Allows you to only show messages with a certain priority. You can use the dropdown menu to specify the priority.

Go

With this button, you apply the filter settings to the current view. Depending on the filter settings you chose you will see either colored lines and/or only the lines from the category you wish to see.

Reset Filters

Resets the filter settings and returns you to the default view of your database.

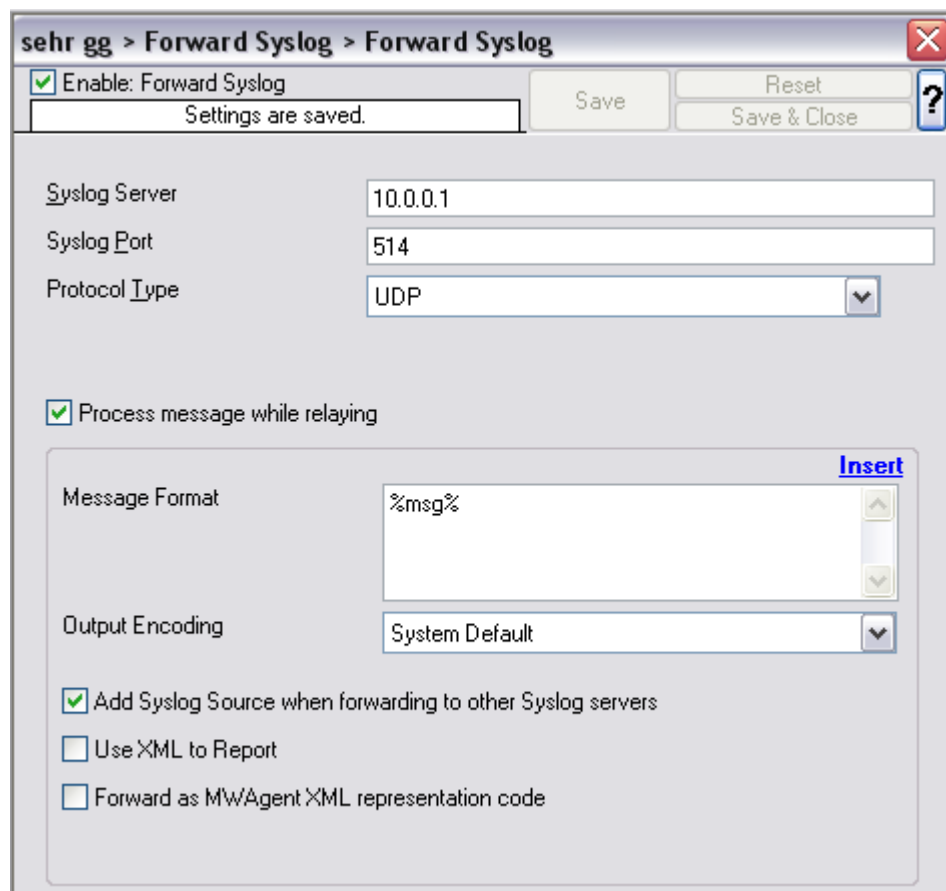
6 Configuring MonitorWare Agent

MonitorWare Agent is easy to use and is powerful.

In this chapter, you see how to configure the MonitorWare Agent Service.

The MonitorWare Agent service runs in the background once it is configured. There is no manual intervention needed to operate it. As such, this chapter focuses on the MonitorWare Agent configuration client application. It is used to configure the service settings.

To run the MonitorWare Agent Configuration client, simply click its icon present in the MonitorWare Agent program folder located in the Start menu. Once started, a Window similar to the following one appears:



MonitorWare Agent Configuration Client

The configuration client ("the client") has two elements. On the left hand side is a tree view that allows you to select the various elements of the MonitorWare Agent system. On the right hand side, there are the parameters specific to the element selected in the tree view. In the sample above, the right hand side displays the specific parameters for a rule action.


The tree view has three top-level elements: **General / Defaults**, **Running Services** and **RuleSets**.

Under **General / Defaults**, basic operational parameters as well as defaults for actions and services are defined. The default themselves do not activate anything. However, the parameters in here are used each time an actual service or action needs a configuration parameter and none is defined in that specific instance. We highly recommend putting the most common parameters into the defaults. That reduces the amount of data entry in the specific elements dramatically. Please note that each default can be overwritten in a specific service or action.

The tree view's **Running Services** area lists all configured services as well as their parameters. There is exactly one service entry for each service created. **Please note that there can be as many instances of a specific service type as your application requires.** Typically, there can be multiple instances of the same service running, as long as their configuration parameters do not conflict. For example the Syslog service: there can be multiple Syslog servers on a given system as long as they listen to different ports. Consequently, there can be multiple instances of the Syslog service be created. For example, there could be three of them: two listen to the default port of 514, but one with TCP and one with UDP and a third one listens to UDP, port 10514. All three coexist and run at the same time. If these three services are listening to the same port then an error message is logged into Windows Event log that more than one instance of Syslog Server is running. After which MonitorWare Agent wouldn't be able to perform the desired action.

Theoretically, you can run a few hundred services in a single service instance. However, both from a usage scenario point of view as well as concerning operating system resources, we recommend limiting the services to a maximum of 20 to 30. Of course, there are some applications where more than this limit is useful. MonitorWare Agent does not restrict this number. If there is a need for a large number of services and the hardware is capable of managing all these tasks, there is nothing in the MonitorWare Agent that limits from doing so.

The service definition looks like this:



MonitorWare Agent Configuration Client - Service Definition View

The actual parameters depend on the service type. Common to all services is the

capability to enable or disable a service. A service is started only if it is enabled. Otherwise, it does not run, but the configuration data can still be present. That way, it is easy to temporarily disable a service without deleting it.

Also common to all service types is the association to a rule set seen at the bottom of the right hand configuration dialog. This specifies which of the rule sets have to apply to information units generated by this service.

To create a new service, right click on "**Running Services**". Then select "**Add Service**" and the respective service type from the pop up menu. Then follow the wizard. To delete an existing service, right click it and select "**Delete Service**". This removes the service and its configuration is irrecoverable. To temporarily "**Remove a service**", simply disable it in the property sheet.

The tree view's last main element is **RuleSets**. Here, all rule sets are configured. Directly beneath "Rules" are the individual rule sets. Each set is completely independent from each other. They are just centrally stored so they can be associated with services (see above for an explanation).


Beneath each rule set are the individual rules. As described in [Rules](#), a rule's position in the list is vitally important. Rules at the top of the rule set are executed before those further down. To move a rule up or down, simply right click it and select "move up" or "move down" from the pop up menu.

In the tree view, filter conditions and actions are beneath the rule they are associated with. Finally, beneath actions are all actions to carry out.

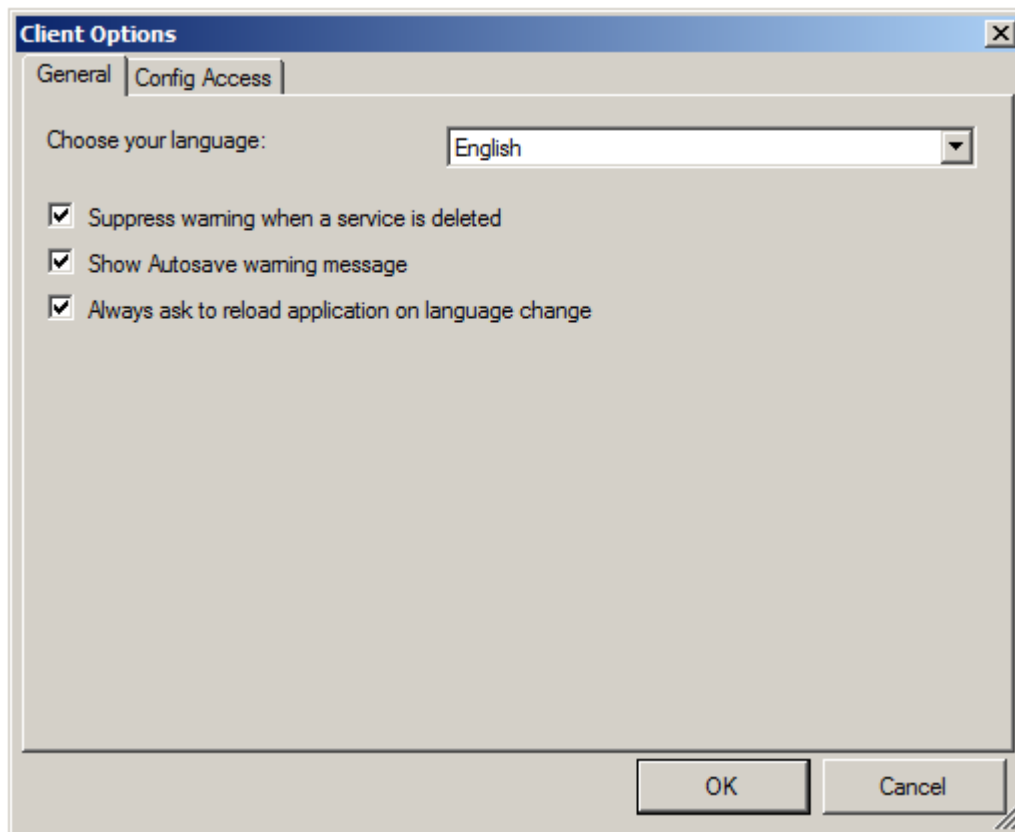
The following sections describe each element's properties.

6.1 Client Options

There are several options, that refer to the configuration client and not the service. These can be found under File -> Options



General Options



Choose your language

You can choose from various language packs, delivered with the client. Please note, that some languages are not fully supported and "English" is the default and suggested language.

Suppress warning when a service is deleted

If this option is checked, warnings when deleting a service will be suppressed. Such a warning can occur when you try deleting a service and there is no other service using the connected ruleset.


Show autosave warning message

If you make changes in the configuration and switch to another component, a warning will occur if you haven't saved the changes. This warning will also allow you to directly enable auto-saving the configuration.

Always ask to reload application after language change

When you change the language, a popup will ask you to reload the configuration client to properly apply the changes and load with the set language.

Config Access



The screenshot shows the 'Client Options' dialog box with the 'Config Access' tab selected. The dialog has two radio buttons: 'Load Configuration from Registry' (selected) and 'Load Configuration from File'. Under the 'Load Configuration from Registry' option, there are three text boxes: 'Root Registry Key' with the value 'SOFTWARE\Adiscon\MonitorWare', 'Base Registry Key' with 'SOFTWARE\Adiscon\MonitorWare\Agent', and 'Data Registry Key' with 'SOFTWARE\Adiscon\MonitorWare\AgentData'. Under the 'Load Configuration from File' option, there are two text boxes: 'Configuration Filename' with 'C:\Program Files (x86)\MonitorWare\Agent\mwa' and 'Data Directory' with 'C:\Program Files (x86)\MonitorWare\Agent\'. Each text box has a 'Browse' button to its right. At the bottom of the dialog are 'OK' and 'Cancel' buttons.

Load Configuration Registry Path

The Configuration Client can be switched to a different registry path for configuration. The registry path change can be made permanent here. The changed registry path is the saved within the Parameters key of the Service.


Load Configuration from File

Alternatively, you can configure the service to load the configuration from a file. You can set the paths with the two fields below.

6.2 General Options

6.2.1 License Options

This tab can be used to enter the MonitorWare Agent license after purchase.



The screenshot shows the 'License Options' configuration window. At the top, there is a toolbar with buttons for 'Comments', 'Settings', 'Save', 'Reset', and 'Configure for...'. Below the toolbar, the 'License Informations' section is expanded. It contains a 'Registration Name' text box, a 'Registration Number' section with five 'Key' text boxes (Key1 to Key5), and two buttons: 'Import from Clipboard' and 'Verify License'. At the bottom, there is a help message: 'More information can be found at our Homepage' with a link to 'http://www.mwaagent.com/'.

License Option Parameters

Registration Name

The user chooses the registration name. It should correspond to your organization name, e.g. a company called "AA Carpenters, Inc." should not choose "AA" as registration name. This can easily be mistaken and most probably be rejected by Adiscon for that reason. With the above scenario, we recommend using the full company name "AA Carpenters, Inc.".

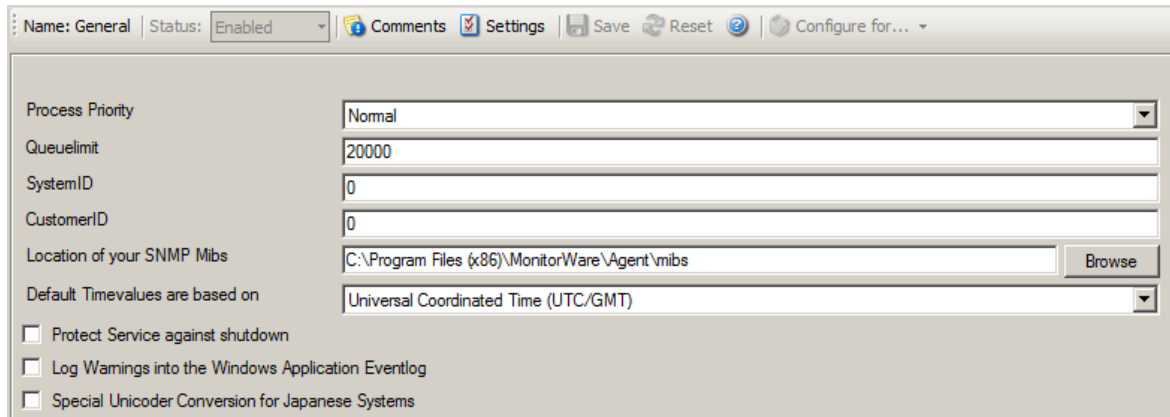
Please note: The registration name is case sensitive. It must be entered exactly as given. Leading and trailing spaces are also part of the registration name, so be sure to enter none.

Registration Number

Adiscon provides this number. It is valid for a specific registration name. Be sure to enter the correct registration number. The client detects invalid registration numbers and report the corresponding error.

6.2.2 General

The General Options available on this form are explained below:



The screenshot shows the 'General' configuration window for the MonitorWare Agent. The window title is 'Name: General | Status: Enabled'. The interface includes several fields and options:

- Process Priority:** A dropdown menu set to 'Normal'.
- QueueLimit:** A text input field containing '20000'.
- SystemID:** A text input field containing '0'.
- CustomerID:** A text input field containing '0'.
- Location of your SNMP Mibs:** A text input field containing 'C:\Program Files (x86)\MonitorWare\Agent\mibs' with a 'Browse' button to its right.
- Default Timevalues are based on:** A dropdown menu set to 'Universal Coordinated Time (UTC/GMT)'.
- Checkboxes:** Three unchecked checkboxes are listed below: 'Protect Service against shutdown', 'Log Warnings into the Windows Application Eventlog', and 'Special Unicoder Conversion for Japanese Systems'.

Figure1: General Options

Process Priority

Configurable Process Priority to fine-tune application behavior.

Queue Limit

The application keeps an in-memory buffer where events received but not yet processed are stored. This allows the product to handle large message bursts. During such burst, the event is received and placed in the in-memory queue. The processing of the queue (via rule sets) itself is de-coupled from the process of receiving. During traffic bursts, the queue size increases, causing additional memory to be allocated. At the end of the burst, the queue size decreases and the memory is freed again.

Using the queue limit, you can limit that maximum number of events that can be in the queue at any given time. Once the limit is reached, no further enqueueing is possible. In this case, an old event must first be processed. In such situations, incoming events might be lost (depending on the rate they come in at). A high value for the queue size limit (e.g. 200,000) is recommended, because of the risk of message loss. It is also possible to place no limit on the queue. Use the value zero (0) for this case. In this case, the queue size is only limited by virtual memory available. However, we do not recommend this configuration as it might cause the product to use up all available system memory, which in turn could lead to a system failure.

SystemID

SystemID is of type integer to be used by our customer. In addition, it is user configurable.

CustomerID

CustomerID is of type integer provided for customer ease. For example if someone monitors his customer's server, he can put in different CustomerIDs into each of the clients. Let us say someone monitors servers A and B. A has 5 servers all of them with CustomerID = 1 and B has 2 servers all of them with CustomerID = 2. Both A and B happen to have a server named "SERVER". Together with the customerID, these machines are now uniquely identifiable. This is user configurable.

Location of your MIBS

Click the Browse button to search for your MIBS location or enter the path manually.

6.2.3 Debug

Debug Options Tab

This tab can be used to debug rule bases. Especially with complex bases, it might be necessary to learn what application is internally doing while it is processing them. With the debug log, the service tells you some of these internal workings.

Other than rule basis testing, the debug log is also helpful when contacting Adiscon support. An Adiscon support engineer might ask you to set the debug log to a specific level while doing troubleshooting.

Important: Debug logging requires considerable system resources. The higher the log level, the more resources are needed. However, even the lowest level considerable slows down the service. As such, **we highly recommend turning debug logging off for normal operations.**

Name: Debug | Status: Enabled | Comments | Settings | Save | Reset | Configure for...

Enable Debug output into file

File and path name: C:\Program Files (x86)\MonitorWare\Agent\MonitorWare Agent.txt | Browse

Errors Warnings
 Minimal Output
 Information Output
 Ultra Verbose Output
 Rule Filter Engine Output

Use circular Logging

Number of logfiles: 10
Maximum filesize (KB): 51200

Figure3: Debug Options

Enable Debug output into file

If checked, the debug log is enabled and written as the service operates. If unchecked, no debug log is written. For performance reasons, it is highly recommended that this box is unchecked during normal operations.

File and path name

The full name of the log files to be written. Please be sure to specify a full path name **including** the driver letter.

If just the file and/or path name is specified, that information is local to the service default directory. As this depends on a number of parameters, it might be hard to find the actual log file. So for consistency purposes, be sure to specify a fully qualified file name including the drive.

Note: If the configured directories are missing, they are automatically created by application i.e. the folder specified in "File and Path Name".

Debug Levels

These checkboxes control the amount of debug information being written. We highly recommend only selecting "Errors & Warnings" as well as "Minimum Debug Output" unless otherwise instructed by Adiscon support.

Circular Debug Logging

Support for circular Debuglogging has been added as the debuglog can increase and increase over time. This will avoid an accidental overload of the harddisk. Of course you can also customize the amount of files used and their size or disable this feature.

6.2.4 Engine

Engine specific Options Tab

The Engine specific Options are explained below:

The screenshot shows the configuration window for the MonitorWare Agent engine. The window title is "Name: Engine" and the status is "Enabled". The interface includes several sections for configuration:

- Action specific:**
 - Enable retry of Actions on failure
 - Retry Count: 1
 - Retry period (ms): 100
- Rule Engine specific:**
 - Abort Rule Execution when one Rule fails?
 - Enable internal DNS Cache
 - How long should dns names be cached?: 1 hour
 - How many DNS records can be cached?: 1024
 - Internet Protocoltype: IPv4
- Resource Library Cache Options:**
 - How long should libraries be cached?: 30 minutes

Figure2: Engine specific Options

Action specific

Enable retry of Actions on failure

If enabled, the Agent retries Actions on failure (until the retry counter is reached). Note that the Event error 114 will only be written if the last retry failed, previous error's will only be logged in the debug log (With the error facility). Note that you can customize the Retry Count and the Retry Period in *ms* as well.

Rule Engine specific

Abort Rule Execution when one Rule fails?

If checked, and an action fails, the execution will be aborted.
If unchecked, and an action fails, simply the next action in this rule will be executed.

DNS Cache Options

Enable internal DNS Cache

The DNS cache is used for reverse DNS lookups. A reverse lookup is used to translate an IP address into a computer name. This can be done via the [resolve hostname action](#). For each lookup, DNS needs to be queried. This operation is somewhat costly (in terms of performance). Thus, lookup results are cached. Whenever a lookup needs to be performed, the system first checks if the result is already in the local cache. Only if not, the actual DNS query is performed and the result then stored to the cache. This greatly speeds up reverse host name lookups.

However, computer names and IP addresses can change. If they do, the owner updates DNS to reflect the change. If we would cache entries forever, the new name would never be known (because the entry would be in the cache and thus no DNS lookup would be done). To reduce this problem, cache records expire. Once expired, the record is considered to be non-existing in the cache and thus a new lookup is done.

Also, cache records take up system memory. If you have a very large number of senders who you need to resolve, more memory than you would like could be allocated to the cache. To solve this issue, a limit on the maximum number of cache records can be set. If that limit is hit, no new cache record is allocated. Instead, the least recently used record is overwritten with the newly requested one.

How long should DNS names be cached?

This specifies the expiration time for cache records. Do not set it too high, as that could cause problems with changing names. A too low-limit results in more frequent DNS lookups. As a rule of thumb, the more static your IP-to-hostname configuration is, the higher the expiration timeout can be. We suggest, though, not to use a timeout of more than 24 to 48 hours.

How many DNS records can be cached?

This is the maximum number of DNS records that can be cached. The system allocates only as many memory, as there are records required. So if you have a high limit but only few sending host names to resolve, the cache will remain small. However, if you have a very large number of host names to resolve, it might be useful to place an upper limit on the cache size. But this comes at the cost of more frequent DNS queries. You can calculate about 1 to 2 KBytes per cache record.

Preferred protocol for name resolution

Select if you wish to prefer IPv4 or IPv6 addresses for name resolution. Note that this only has an effect on names which return both, IPv4 and IPv6 addresses.

Ressource Library Cache Options

How long should libraries be cached?

This feature will be mainly useful for EventLog Monitor. For events with the same reoccurring event sources, this will be a great performance enhancement. The cache will also work for remote system libraries (requires administrative default shares). All libraries will be cached for 30 minutes by default.

6.2.5 QueueManager

Queue Manager Tab

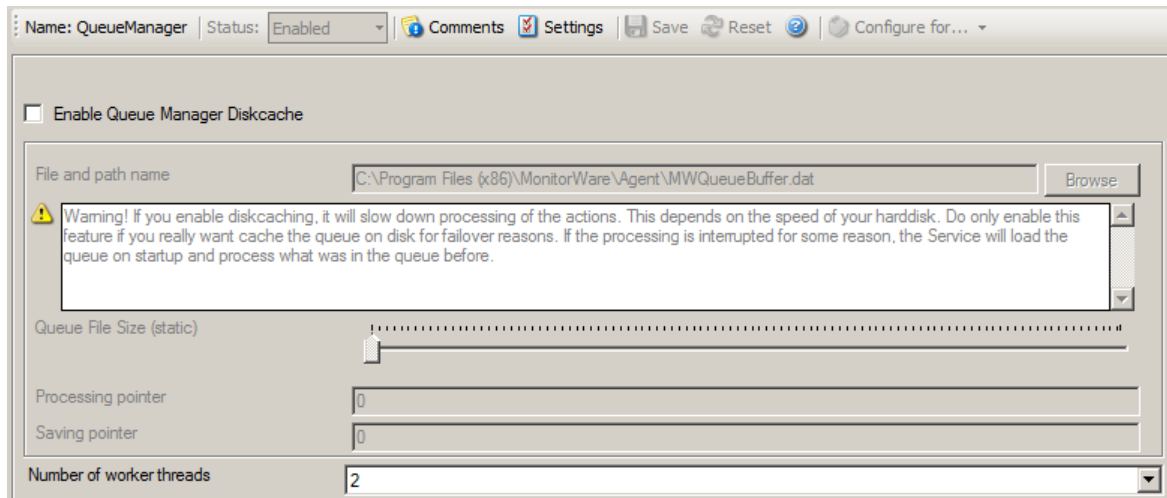


Figure 4: Queue Manager Options

Queue Manager DiskCache

This feature enables the Agent to cache items in its internal queue on disk using a fixed data file. **First of all a Warning. Only use this feature if you really need to!** Depending on the speed of your hard disks, it will slow down processing of the actions, in worst case if the machine can't handle the IO load, the Queue will become full sooner or later. The DiskCache is an additional feature for customers, who for example want to secure received Syslog messages which have not been processed yet.

The diskcache will not cache infounits from services like EventLog Monitor, as this kind of Service only continues if the actions were successfully. All other information sources like the Syslog Server will cache it's messages in this file. If the Service or Server crashes for some reason, the queue will be loaded automatically during next startup of the Agent. So messages which were in the queue will not be lost. Only the messages which was currently processed during the crash will be lost.

File and Pathname

As everywhere else, you can define here, where the queue file should be stored.

Queue Manager specific

Number of worker threads

Defines the number of worker background threads that MWAgent uses to process it's queue.

6.3 Services

6.3.1 Understanding Services

Services gather events data. For example, the Syslog server service accepts incoming Syslog messages and the Event Log Monitor extracts Windows event log data. There can be unlimited multiple services. Depending on the service type, there can also be multiple instances running, each one with different settings.

You must define at least one service, otherwise the product does not gather event data and hence does not perform any useful work at all. Sometimes, services are mistaken with service defaults those are pre-existing in the tree view. Service defaults are just the templates that carry the default properties assigned to a service, when one of the respective type is to be created. Service defaults are NOT executed and thus can not gather any data.

6.3.2 Syslog Server

Configures a Syslog Server service. It can be set to listen to any valid port. UDP and TCP communication is supported.

The screenshot shows the 'Syslog Server Properties' configuration window. At the top, there is a status dropdown set to 'Enabled', and buttons for 'Comments', 'Save', 'Reset', and a help icon. Below this, there are four main configuration fields: 'Internet Protocoltype' (set to 'IPv4'), 'Protocol Type' (set to 'UDP'), 'IP Address' (set to '0.0.0.0'), and 'Listener Port' (set to '514'). A tabbed interface below these fields includes 'General', 'Encoding', 'UDP Options', 'TCPOptions', and 'Syslog TLS'. The 'General' tab is active and contains several checkboxes: 'Use original message timestamp (RFC 3164)', 'Take source system from Syslog message', 'Resolve Hostnames', 'Escape control characters', 'Enable RFC3164 Parsing' (checked), 'Enable RFC5424 Parsing' (checked), and 'Append ProcessID to Syslogtag if available' (unchecked). At the bottom, there is a 'RuleSet to use' dropdown menu and a 'Refresh' button.

Syslog Server Properties

Internet Protocol Type

Select the desired protocol type. IPv4 and IPv6 are available. The IPv6 protocol needs

to be properly installed in order to be used. Note that one Service can only handle IPv4 or IPv6, so if you want to use both protocols, you will need to create two separate services.

Protocol Type

Syslog messages can be received via [UDP](#), [TCP](#) or [RFC 3195](#) RAW. One listener can only listen to one of the protocols. Typically, Syslog messages are received via UDP protocol, which is the default. The syslog server also can receive Syslog messages via TCP and reliable Syslog messages via TCP using the new [RFC 3195](#) RAW standard.

IP Address

The Syslog Server can now be bound to a specific IP Address. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. This feature is useful for multihome environments where you want to run different Syslog Servers on different IP Addresses. Please note that the default IP Address 0.0.0.0 means ANY IP Address.

Listener Port

The port the Syslog server listens on. The typical (standard) value is 514. This should be changed only if there is a definite need for it. Such a need typically arises from security concerns. If the port is changed, all reporting devices (routers, printers ...) must also be configured to use the non-standard port.

General Options

Use Original Message Timestamp

If this box is checked, the timestamp is retrieved from the Syslog message itself (according to [RFC 3164](#)). If left unchecked, the timestamp is generated based on the local system time. The Syslog message timestamp does not contain time zone information. Thus, we strongly recommend unchecking this box if messages from devices in multiple time zones are to be received.

Take source system from Syslog message

If this box is checked, the name or IP address of the source system is retrieved from the Syslog message itself (according to [RFC 3164](#)). If left unchecked, it is generated based on the address, the message was received from.

Please note that there are many devices, which do NOT generate RFC 3164 compliant messages. If you check this option here, you might see a very strange value as the event source!

Resolve Hostnames

If this box is checked, the name of the source system is retrieved via DNS reverse name resolution. If unchecked, the IP address itself is used as the name.

Please note that this setting does have any effect if the "Take source system from Syslog message" setting is checked. In this case, the message is always taken from the Syslog message itself.

Escape Control Characters

Control characters are special characters. They are used e.g. for tabulation, generating beeps and other non-printable uses. Typically, syslog messages should not contain control characters. If they do, control characters could eventually affect your logging. However, it might also be that control characters are needed.

With this setting, you can specify how control characters received should be handled. When checked, control characters are replaced by a 5-byte sequence with the ASCII character ID. For example, a beep is the ASCII BEL character. BEL is assigned the numerical code 7. So if a BEL is received, it would be converted to "<007>" inside your syslog message. When the box is left unchecked, no conversion takes place.

In any case, ASCII NULs are converted to "<000>" to prevent security issues in the log files.

Please note: if you used double-byte character sets, control character escaping can cause your message to become clobbered. So be sure to leave it unchecked in that case.

Enable RFC 3164 Parsing

If this box is checked, [RFC 3164](#) compliant message parsing is enabled. If unchecked, "traditional" Adiscon message parsing is selected. If you experience trouble with the sender host name or the timestamp, we suggest that you turn off RFC 3164 compliant message parsing. Many existing devices do not fully comply with RFC 3164 and this can cause those issues.

Enable RFC 5424 Parsing

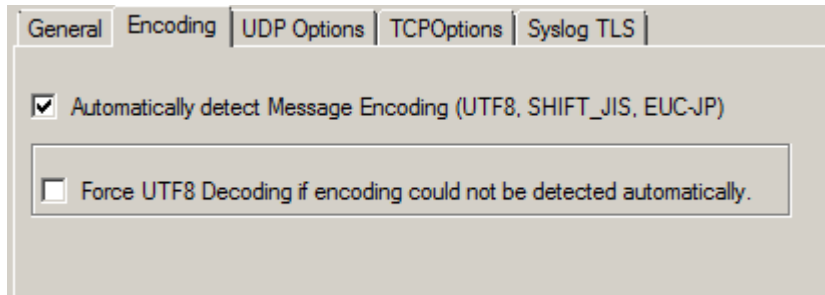
If this box is checked, RFC 5424 compliant message parsing is enabled for Syslog RFC5424 Header detection and decoding. This also involves new useable Syslog properties.

If unchecked, "traditional" Adiscon message parsing is selected. If you experience trouble with the sender host name or the timestamp, we suggest that you turn off RFC 5424 compliant message parsing. Many existing devices do not fully comply with RFC 5424 and this can cause those issues.

Append ProcessID to SyslogTag if available

This option is related to RFC5424 header parsing and was default in previous versions. However the default now is off in order to separate the Syslogtag from the ProcessID.

Encoding options



Encoding Options

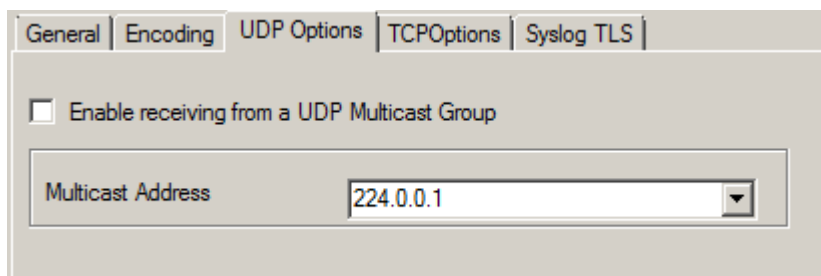
Automatically detect Message Encoding (UTF8, SHIFT_JIS, EUCJP)

If enabled, the message will be checked for different encodings. This is important if you have syslog messages with multibyte characters. Once an encoding is detected, it will automatically be converted into UTF16 internally.

Force UTF8 Decoding

This option forces UTF8 Decoding of all incoming messages. This is also useful for syslog messages encoded in UTF8 but missing the BOM withing the Syslog message.

UDP Options



UDP Options

UDP Options - Enable receiving from a UDP Multicast Group

This option supports receiving Syslog messages via multicast IP Addresses like 224.0.0.1 for example.

TCP specific options

The screenshot shows the 'TCP Options' configuration window. It features a tabbed interface with 'TCPOptions' selected. The configuration includes a 'Session Timeout' dropdown menu set to '15 Minutes', a checked checkbox labeled 'Messages are separated by the following sequence', a 'Message separation' text input field containing the text '\n', and a 'Message Completion Timeout' dropdown menu set to '15 seconds'.

TCP Options

TCP Options - Session Timeout

One of the TCP-specific options is the session timeout. This value declares, how long a TCP session may be kept open, after the last package of data has been sent. You can by default set values between 1 second and 1 day. Or you can use a custom value with a maximum of 2147483646 milliseconds. If you wish to disable the session timeout, you can use a custom value of 0 milliseconds to disable it.

TCP Options - Messages are separated by the following sequence

If this option is checked, you can use multiple messages in the same transmission and the following options are enabled:

Message separation sequence - determines, how you want to separate the messages. By default "\r\n" is the value for this, as most times a message ends with a carriage return and/or a line feed. But, you can choose your own separation sequence here as well.

Message Completion Timeout - here you can set the time that is allowed to complete a message. If the time is exceeded, but the message not yet completed, the rest will be treated as a new message. The counter is resetted each time, a new message begins. You can choose from multiple values between 1 second and 1 day, or choose a custom value in milliseconds (0 = disable, maximum = 2147483646)

Syslog TLS

SSL/TLS Options

Enable SSL / TLS Encryption

This option enables SSL / TLS encryption for your syslog server. Please note, that with this option enabled, the server only accepts SSL / TLS enabled senders.

TLS Mode

The TLS mode can be set to the following:

Anonymous authentication

Default option, which means any client certificate will be accepted, or even none.

x509/name (certificate validation and name authentication)

When this mode is selected, the subject within the client certificate will be checked against the permitted peers list. This means the Syslog Server will only accept the secured connection if it finds the permitted peer in the subject.

509/fingerprint (certificate fingerprint authentication)

This mode creates a SHA1 Fingerprint from the client certificate it receives, and compares it to fingerprints from the permitted peers list. You can use the debuglog to see fingerprints of client certificates which were not permitted.

x509/certvalid (certificate validation only)

A Syslog Sender is accepted when the client certificate is valid. No further checks are done.

Select common CA PEM

Select the certificate from the common Certificate Authority (CA), the syslog receiver should use the same CA.

Select Certificate PEM

Select the client certificate (PEM Format).

Select Key PEM

Select the keyfile for the client certificate (PEM Format).

Permitted Peers

This list contains all permitted peers. If *x509/name* is used, this can contain parts of the client certificate subject. For example if you have CN = secure.syslog.msg in the certificate subject, you can add "secure.syslog.msg" as permitted peer. When using *x509/fingerprint*, this list holds a list of permitted SHA1 fingerprints. The fingerprints can either be generated with OpenSSL Tools, or grabbed from the debug logfile. The format is like described in RFC 5425, for example:
"SHA1:2C:CA:F9:19:B8:F5:6C:37:BF:30:59:64:D5:9A:8A:B2:79:9D:77:A0".

Default Ruleset Name


Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

Please Note

Updated the OpenSSL components and libraries with the latest Version openssl-1.0.1j.

6.3.3 SETP Server

Configures a [SETP](#) server service. A SETP server is used inside the [MonitorWare line of products](#) to ensure reliable receiving of events from other systems. There are only few configuration options, as SETP takes the original message from the sender and uses the exact settings that the sender was configured for. No alteration occurs at the SETP server side; as such, no values need to be configured for the message format.



SETP Server Properties

Internet Protocoltype

Select the desired protocol type. IPv4 and IPv6 are available. The IPv6 protocol needs to be properly installed in order to be used. Note that one Service can only handle IPv4 or IPv6, so if you want to use both protocols, you will need to create two separate services.

Listener Port

The port the [SETP](#) server listens on. The default value is 5432. This should be changed only if there is a definite need for it. Such a need typically arises from security concerns. If the port is changed, all reporting devices must also be configured to use the non-standard port. SETP operates over [TCP](#).

Listener IP Address

The SETP server service can be bound to a specific IP Address. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. This feature is useful for multihome environments where you want to run different Syslog Servers on different IP Addresses. Please note that the default IP Address 0.0.0.0 means ANY IP Address.

Session Timeout

This controls how long a session is to be opened from the server side.

Enable SSL/TLS

If this option is enabled then this action connects to SSL / TLS [SETP](#) servers. Please make sure that you want this option to be enabled.

Please note: If this option is enabled, this action will not be able to connect to NON-SSL SETP Servers.

Options

Under this group box, you can see three more options as discussed below:

Use zLib Compression to compress the data

When enabled, MonitorWare Agent decompresses the zLib compressed data sent by the SETP senders. It is still be able to receive normal data. zLib compression is useful to reduce traffic in WAN environments.

Session Timeout

It controls how long a session is to be opened from the server side.

Notify Sender about Rule Action Errors?

Enable this option to communicate the outcome of an action back to the the sender of the SETP message.

This communicates back the status of actions carried out on the receiver to the sender of the event. In essence, the sender system will know if the action failed or succeeded on the remote machine. It can then act exactly like the action was carried out on the local machine. The exact handling of failure states is depending on the event source.

An example: you have a machine running an EventLog Monitor and sending these events via SETP, and on the other side have all incoming events written into a database. If the database would be offline and the events not being written into it, the SETP server would return as the last message that the action failed (as long as this option is enabled) and generate a error event with ID 1005 (and generate a Success Event with ID 1012 if successful again). The sender would then halt and retry sending the event. This is because SETP is built somehow like TCP which ensures data transfer, but additionally can return a status to the sender if the following action was successful.

This happens because the event log monitor (as well as the file monitor and others) is a restartable event source. It uses the outcome of actions to decide if the action is to be retried in another run of the same source. Other event sources have different behavior. The syslog server, for example, does not retry failed actions. This is due to the lossy nature of syslog, in which loosing syslog messages is explicitly permitted (and favourable over taking up too many system resources by trying to buffer them).

Please Note: If you enable this feature, older MonitorWare Agent Versions (4.2.x and below, as well as WinSyslog 7.2.x and EventReporter 8.2.x and below) may have trouble sending data over SETP once a Rule Exception occurs! If you intend to use this feature, make sure all MonitorWare Agent Installations are at least Version 4.3.x (This applies for WinSyslog 7.3.x and EventReporter 8.3.x as well).

Default Ruleset Name

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

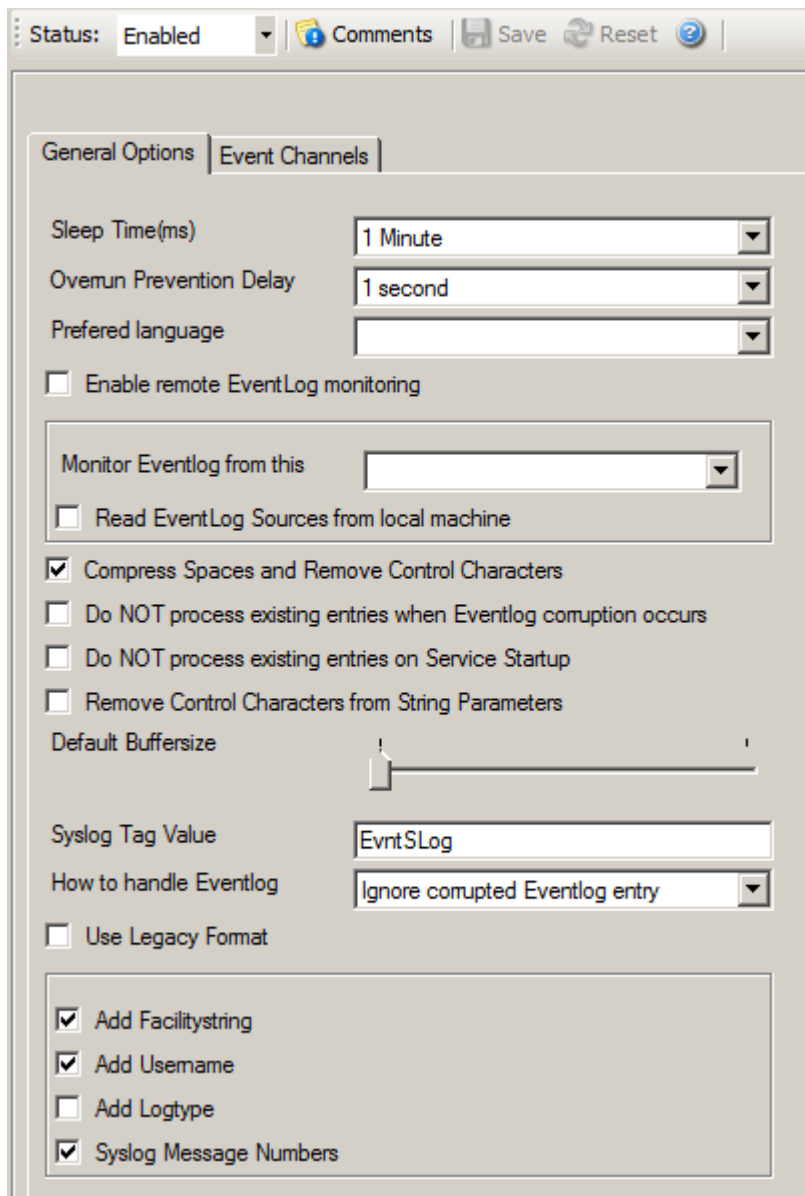
6.3.4 Event Log Monitor

This dialog configures the Windows Event Log Monitor service.

This service was initially introduced by Adiscon's [EventReporter](#) product. To allow previous EventReporter customers seamless upgrades, there are a number of compatibility settings to support older message formats.

Windows Vista comes with a considerably changed event logging system. In theory, the Event Log Monitor works with Vista. However, we know of some incompatibilities. For best results, **we recommend using the [Event Log Monitor V2](#) service, which was specifically written for Windows Vista.** The Event Log Monitor described here is applicable for Windows 2000, 2003 and XP (where the new Vista-like event logging system is not available). The Client will automatically detect and load available EventLog types during the first startup of the Event Log Monitor.

Event Log Monitor	Event Log Monitor V2
Windows 2000	Windows Vista
Windows XP	Windows 2008
Windows 2003	Windows 7
	Windows 8
	Windows 2012



Event Log Monitor Properties

The most important part of this dialog is the table in the middle. It specifies which event logs are to be monitored. In the screenshot above, the monitor is set to all Windows-native event log types that can possibly occur. However, there might also be custom event logs. Such custom logs can be created by any application. For example, an application "MySuperApp" might create an event log "MySuperAppLog". Then, it might log its messages into this log and not the Windows application event log.

In order to support such custom event logs, the log monitor allows an unlimited number of additional logs to be added to it. In order to do so, press the "Insert" button. A new log is added to the bottom of the list. Then, you can edit its name in the "Eventlogtype Name" combobox (yes, you can overwrite the provided values!).

Logs checked in the table are actually processed. Those unchecked are kept in the configuration, but are not processed.

General Options

The General Options available on this form are explained below:

Sleep Time

The event log monitor periodically checks for new event log entries. The "Sleep Time" parameter specifies how often this happens. This value is in [milliseconds](#).

We suggest a value of 60,000 milliseconds for the "Sleep Time". With that setting, the event log monitor checks for new events every 60 seconds. Larger periods can be specified for occasionally connected systems or if email delivery with few emails per day is intended.

Very security-aware environments might use a shorter interval. The event log monitor service is specifically designed to limit the burden on the monitored system. As such, resource usage is typically low, even with frequently run event log checks. However, we recommend not running the event log monitor more often than once a second.

Overrun Prevention Delay

This property allows configuring a delay after generating an event. The time is the delay in [milliseconds](#).

If run at a value of zero, the event log monitor service generates events as fast as the machine permits. We have seen scenarios where routers and receivers are not able to keep up with this rate, resulting in packet loss. In addition, the CPU of the reporting machine is run at 100% - which is not a problem because the service runs at a low priority. However, with even a 1-millisecond delay, there is no noticeable CPU activity even when large bursts of events are forwarded. At one millisecond, the service can still generate 1000 events per second.

The default setting is an overrun protection of five millisecond, which allows roughly 200 events per second. This should be sufficient for even very busy servers.

Preferred Language

You can select a preferred language, and the Eventlog Monitor will send the message in this language. It will only work if these languages are installed and message libs are available with the preferred language. If this fails, it will automatically fall back to the system default language.

Enable Remote EventLog Monitoring

If enabled, the EventLog Monitor will read and process the EventLog from a remote

machine. Use the verify button to make sure that the network connection is working correctly

Please make sure that the machine, which you are going to monitor, does have File and Print Services enabled and is accessible by this machine.

This is important as the EventLog Service will read the message libraries on the remote machine by using the default administrative shares. For this reason, the Service must be configured to run with a user who has administrative privileges/permissions on the local **and** remote machine. If File and Print services remain disabled, the local message libraries will be used automatically instead. Note that you may experience a lot of missing message libraries in this case.

Additionally you have an option to **read the EventLog Sources** from the local machine. If enabled, the local message libraries will be used instead of the remote machines ones. Sometimes local Event Sources are more reliable or are required for thirdparty EventLog implementations.

Compress Control Characters

This option allows you to control the control character removal and space compression. If checked, control characters (e.g. CR, LF, TAB - non printable characters in general) are removed. Also multiple spaces are compressed to a single one. By default this is checked. We recommend keeping it checked for most cases as it provides better display.

Please note that it should be unchecked if events should be forwarded via email. And it MUST be turned off if double-byte character sets are being processed (e.g. Japanese).

Do NOT process existing entries when Event log corruption occurs

When this option is checked, it prevents from reprocessing of the whole Windows event log when it is [corrupted or truncated](#). So EventReporter / MonitorWare Agent do not process all entries again.

Do NOT process existing entries on Service Startup

When this option is checked, it prevents from reprocessing of the whole Windows event log when the EventReporter / MonitorWare Agent service is restarted.

Remove Control Characters from String Parameters

Enable this option to remove control characters like carriage return or line feeds from parameter strings and category names in Windows Events.

Default Buffersize

The default Buffersize is 10k. This value will be increased or decreased dynamically if

necessary. If you want to use thirdparty applications like Netapp you must increase the Buffersize manually (minimum 65k), because dynamic adjusting is not possible with them.

SyslogTag Value

The SyslogTag Value determines the SyslogTag that is used when forwarding Events via syslog. This is useful, if you want to determine later, what kind of syslog message this is, perhaps because you log EventLogs and syslog into the same database.

How to handle Eventlog Corruption

Sometimes it can occur that Eventlog messages are corrupted and cannot be read correctly. This usually happens if someone tampered with the Eventlog or if you are processing the Eventlog for the first time. In cases like this, you can automatically handle the situation with this option. You have the following options:

- **Retry processing Eventlog from the beginning** - in this case the complete Eventlog will be processed again.
- **Ignore corrupted Eventlog entry (default)** - the affected Eventlog entry will be ignored and processing will continue.
- **Clear all Events from the Eventlog** - the Eventlog will be cleared completely and new Events hopefully don't get corrupted before they are processed.

Use Legacy Format

This option enhances compatibility to scripts and products working with previous versions of EventReporter. The legacy format contains all Windows event log properties within the message itself.

The new format includes the plain text message only. The additional information fields (like event ID or event source) are part of the XML formatted event data. If the new format is used, we highly recommend sending or storing information in XML format. This is an option in each of the action properties (of those actions that support it – the write to database option for example always stores the fields separated, so there is no specific option to do so).

Legacy format is meant to support existing parser scripts. We encourage you to use the new, XML-bound format for new implementations. Legacy format will be maintained in future releases to support backward compatibility, but it is no longer actively being developed. There are some shortcomings in legacy format, which can lead to issues when building or operating a log parser. These shortcomings are by design. We will not change this in legacy format - the solution is to use the new format. After all, the new format was created in order to address the issues with legacy format.

Add Facility String

If checked, facility identification is prepended to the message text generated. This parameter enhances compatibility with existing Syslog programs and greatly

facilitates parsing the generated entries on the Syslog server. We strongly encourage users to use this enhancement.

This setting only applies if the "Use Legacy Format " option is checked.

Otherwise, it does not have any meaning and consequently cannot be configured in that case.

Add Username

If checked, the NT user that generated the event log entry is transmitted. If unchecked, this information is not forwarded. This is a configurable option for customers who have written parsing scripts for a previous format which did not contain Usernames. This option must also be unchecked if MoniLog is being used.

This setting only applies if the "Use Legacy Format " option is checked.

Otherwise, it does not have any meaning and consequently cannot be configured in that case.

Add Log Type

If checked, then log types e.g. system, security etc. etc. is appended to the generated message.

This setting only applies if the "Use Legacy Format " option is checked.

Otherwise, it does not have any meaning and consequently cannot be configured in that case.

Syslog Message Numbers

If checked, a continuously advancing message number is appended to the generated message. This is useful for Syslog delivery to make sure that all messages have been received at the remote server.

This setting only applies if the "Use Legacy Format " option is checked.

Otherwise, it does not have any meaning and consequently cannot be configured in that case.

Event Log Channels Tab

The "Event Log Channels" are basically a list of the different log types. The corresponding log is only be processed if the respective "Enable" checkbox is checked. The parameters are common to all logs and are explained only once.

	Enable	Eventlog Channel
▶	<input checked="" type="checkbox"/>	HardwareEvents
	<input checked="" type="checkbox"/>	Windows PowerShell
	<input checked="" type="checkbox"/>	Application
	<input checked="" type="checkbox"/>	Key Management Service
	<input checked="" type="checkbox"/>	System
	<input checked="" type="checkbox"/>	Internet Explorer
	<input checked="" type="checkbox"/>	Security
	<input checked="" type="checkbox"/>	Active Directory Web Services
	<input checked="" type="checkbox"/>	DFS Replication
	<input checked="" type="checkbox"/>	Directory Service
*	<input type="checkbox"/>	

Eventlog Channels

Report Truncated Log

Do NOT process existing entries

Try to convert Security IDs (SID) to Object Names

Try to convert Active Directory Object Classes

Use Checksum to verify the last processed event

Always search for the last processed Event using

Syslog

Last

Read Eventlog from File

File Path

Type of

Enable date replacement characters (See manual)

Offset in

szEvent TypesToLog

Success

Information

Warning

Error

Audit

Audit

Event Log Monitor - EventLog Types General Options

Report Truncated Log

Windows event logs can be truncated programmatically or via the Windows Event Viewer program. When a log is truncated, all information is erased from it. Any entries not already processed by the service are lost.

The service detects event log truncation. If "Report Truncated Log" is checked, it generates a separate message stating the truncation. This option is most useful in environments where truncation is not expected and as such might be an indication of system compromise.

If you regularly truncate the NT event logs as part of your day-to-day operation, we suggest you turn this option off. In this case, we also recommend using a short sleep

period (for example 10,000 which is 10 seconds) to avoid losing log entries.

Do not Process Existing Entries

If you don't want to get a dump of an existing specific Windows event log then use this option. When MonitorWare Agent / EventReport are restarted it will start processing after that last entry and do not look for the previous entries.

Try to convert Security IDs (SID) to Object Names

With this option you can convert Security ID's (SIDs) to object names. You can enable this feature in the advanced configuration of each event log type in the Event Log Monitor service. Simply check the "Try to convert Security IDs (SID) to Object Names" option.

Note that only the Security event log has this feature enabled by default. For all other event log types this feature is disabled by default.

Try to convert Active Directory Object Classes

With this option you can convert ActiveDirectory Schema GUID's from Security Events on Domain Controllers to object names. For Example Event 565, which usually has a lot of these Schema GUID's! The GUID's are internally cached to speed up EventLog processing operations.

Note that only the Security event log has this feature enabled by default. For all other event log types this feature is disabled by default.

Use Checksum to verify the last processed Event

A checksum of the last processed Event will be stored along with the LastRecord of an eventlog. This checksum is checked during each iteration. If the checksum does not match, we consider the EventLog has been altered, cleared or something else happened. In this case the EventLog is being reprocessed from the beginning.

Please note: This option will prevent you from modifying the LastRecord value. If you do, the whole EventLog will be reprocessed! Please note that this behavior is by design and cannot be avoided. So we recommend to use this feature only if you intend to double check if the LastRecord value is valid.

Always search for the last processed Event using this Checksum

Usually, the last processed Event is determined by the LastRecord value. If the Checksum to verify the last processed Event is activated, then this option to always search for the last processed Event using the Checksum is available. When activated, the last processed Event will also be always determined by the Checksum, not the LastRecord value.

Syslog Facility

The [Syslog facility](#) to map information units stemming from this log to. Most useful if the message is to forward to a Syslog daemon.

Last Record

Windows event log records are numbered serially, starting at one. The service records the last record processed. This textbox allows you to override this value. **Use it with caution!**

If you would like a complete dump of a specific Windows event log, reset the "Last Record" to zero. If you missed some events, simply reset it to some lower value than currently set. It is possible to set "Last Record" to a higher value. This suspends event reporting until that record has been created. We strongly discourage to use this feature unless definitely needed.

Read Eventlog From File

When enabled, the Eventlog is read from a file instead from the system.

File&Path Name

It defines that which file to be read, only available when "Read Eventlog From File" is enabled.

Type of Event Log

It defines as which type of Event log from file is handled. This is important to read the correct message libs from the system.

Enable date replacement characters

Allow the use of dynamic files/paths when using evt files. The same replacement characters as in the FileMonitor apply to this feature. A configured filename may look like this: `C:\temp\evt_%Y%m%d.evt` and would be replaced with `C:\temp\evt_20130101.evt`.

To support changing log file names, there are replacement characters available within the file name. These are:

Character	Meaning
%y	Year with two digits (e.g. 2002 becomes "02")
%Y	Year with 4 digits
%m	Month with two digits (e.g. March becomes "03")
%M	Minute with two digits

%d	Day of month with two digits (e.g. March, 1 st becomes "01")
%h	Hour as two digits
%S	Seconds as two digits. It is hardly believed that this ever be used in reality.
%w	Weekday as one digit. 0 means Sunday, 1 Monday and so on.
%W	Weekday as three-character string. Possible values are "Sun", "Mon", "Tue", "Wed", "Thu", "Fri", "Sat". This replacement character is most useful for DHCP log files.
%generatedfilename%	It contains the fully generated filename (Can be useful for filtering).
%msgsep%	Only available if enable in the advanced option of the File Monitor. This value contains the current used messageseparator. This is usefull if you want to reconstruct messages where the seperator is part of the message.
%msgseplast%	Only available if enable in the advanced option of the File Monitor. This value contains the last used messageseparator. This is usefull if you want to reconstruct messages where the seperator is part of the message.

Character Replacement Table

Offset in Seconds

When "Enable date replacement characters" is enabled, the current date will be used to generate the filenames. However in certain cases, there is a need to generate filenames with past or future dates. Negative values will generate past filenames, while positive values will generate filenames in the future. For example if you want to generate filenames from yesterday (24 hours back), use -84600 as offset.

Event Types to Log

These checkboxes allow local filtering of the event log. Filtering is based on the Windows event type. There is a checkbox corresponding to each Windows event type. Only checked event types will be processed. Unchecked ones will be ignored.

Filtering out unnecessary log types at this level enhances system performance because no information units will be generated and passed to the rule engine. Thus, Adiscon strongly recommends dropping unnecessary log types.

Default Ruleset Name

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

Please Note if you intend to make the Event ID part of the actual Syslog message while forwarding to a Syslog Server then you have to make some changes in the Event Log Monitor Settings. [Click here](#) to know the settings.

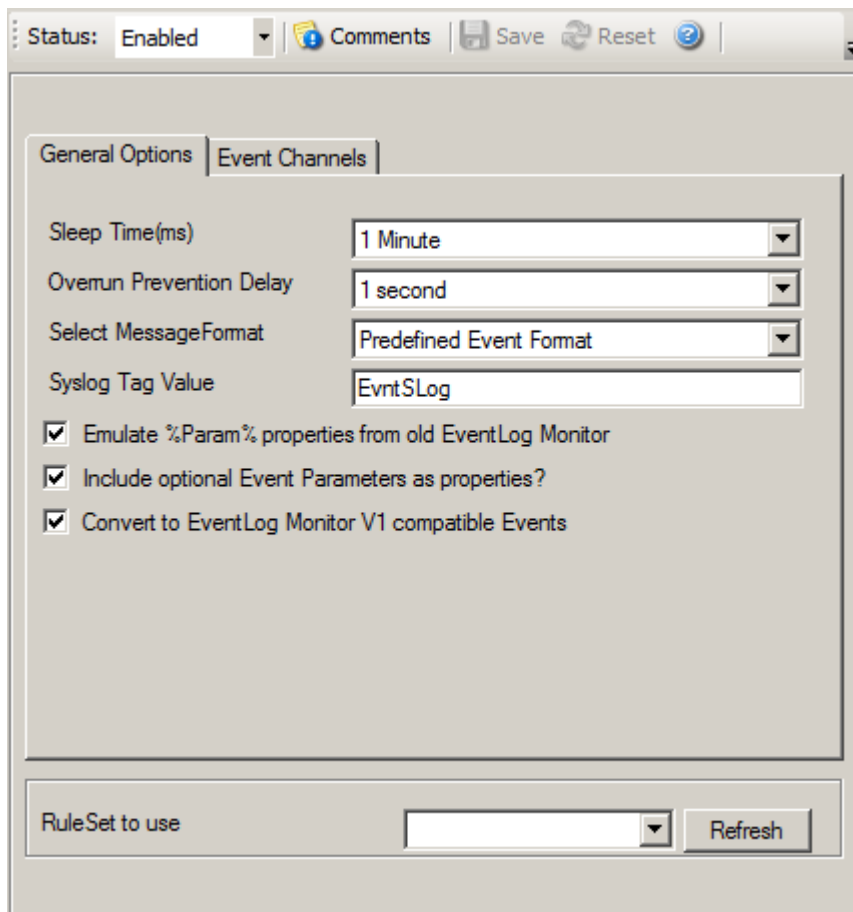
The event log monitor caches messages libraries. This greatly speeds up processing, but causes memory consumption for the cached libraries. By default, libraries are cached for 30 minutes. If memory consumption is too high, you may consider to lower the cache period. The cache is global to all event log monitors. As such, its size must be changed in the general settings.

6.3.5 Event Log Monitor V2 (for Vista)

This dialog configures the Windows Event Log Monitor V2 service for Windows Vista, Windows 2008, Windows 7, Windows 8 and Windows 2012. For Windows 2000, 2003 and XP use the [classical event log monitor](#).

Event Log Monitor	Event Log Monitor V2
Windows 2000	Windows Vista
Windows XP	Windows 2008
Windows 2003	Windows 7
	Windows 8
	Windows 2012

Due to the vast changes to the Windows EventLog in Windows Vista, it was necessary to create a new edition of the EventLog Monitor. This one is specifically designed to process the Windows Vista event logs. The log entries have been split up and are now shown in so-called Channels. These Channels can be considered as categories. First we have the classic EventLog Channels. These consist of the Application-, Security- and System-EventLog etc., which were already known in Windows XP. Then there are the serviced and the direct Channels. The serviced Channels are processed by the EventLog framework for a reliable delivery of the messages, while direct channels are meant for debugging purposes. ConsLogging them may cause a high performance impact. As direct channels are typically not used in practical logging scenarios, they are not yet implemented in the event log monitor. If you have a need to process them, please let us know at support@adiscon.com.



Event Log Monitor Properties

Sleep Time

The event log monitor periodically checks for new event log entries. The "Sleep Time" parameter specifies how often this happens. This value is in [milliseconds](#).

We suggest a value of 60,000 milliseconds for the "Sleep Time". With that setting, the event log monitor checks for new events every 60 seconds. Larger periods can be specified for occasionally connected systems or if email delivery with few emails per day is intended.

Very security-aware environments might use a shorter interval. The event log monitor service is specifically designed to limit the burden on the monitored system. As such, resource usage is typically low, even with frequently run event log checks. However, we recommend not running the event log monitor more often than once a second.

Overrun Prevention Delay

This property allows configuring a delay after generating an event. The time is the delay in [milliseconds](#).

If run at a value of zero, the event log monitor service generates events as fast as the machine permits. We have seen scenarios where routers and receivers are not able to keep up with this rate, resulting in packet loss. In addition, the CPU of the reporting machine is run at 100% - which is not a problem because the service runs at a low priority. However, with even a 1-millisecond delay, there is no noticeable CPU activity even when large bursts of events are forwarded. At one millisecond, the service can still generate 1000 events per second.

The default setting is an overrun protection of five millisecond, which allows roughly 200 events per second. This should be sufficient for even very busy servers.

Select Message Format

With this option you can choose whether the Events will be extracted in "Raw XML Format" or in the "Predefined Event Format".

The XML format is the exact representation of the XML Stream returned by the EventLog System.

Please note that it only contains EventLog data and not a formatted message.

The "Predefined Event Format" is what the Event in the event viewer looks like.

SyslogTag Value

The SyslogTag Value determines the SyslogTag that is used when forwarding Events via syslog. This is useful, if you want to determine later, what kind of syslog message this is, perhaps because you log EventLogs and syslog into the same database.

Emulate %Param% properties from old EventLog Monitor

This option emulates the %Param% properties, which were often used in the old EventLog Monitor. The new EventLog implementation (e.g Windows 7, Windows Server 2008 Windows 8, Windows Server 2012) does not support them in the same way anymore. The Event Log Monitor V2 is still able to provide parameters in the "old style" format, what means that log analysis scripts can receive a consistent stream of data for both new style and old style Windows events.

Include optional Event Parameters as properties

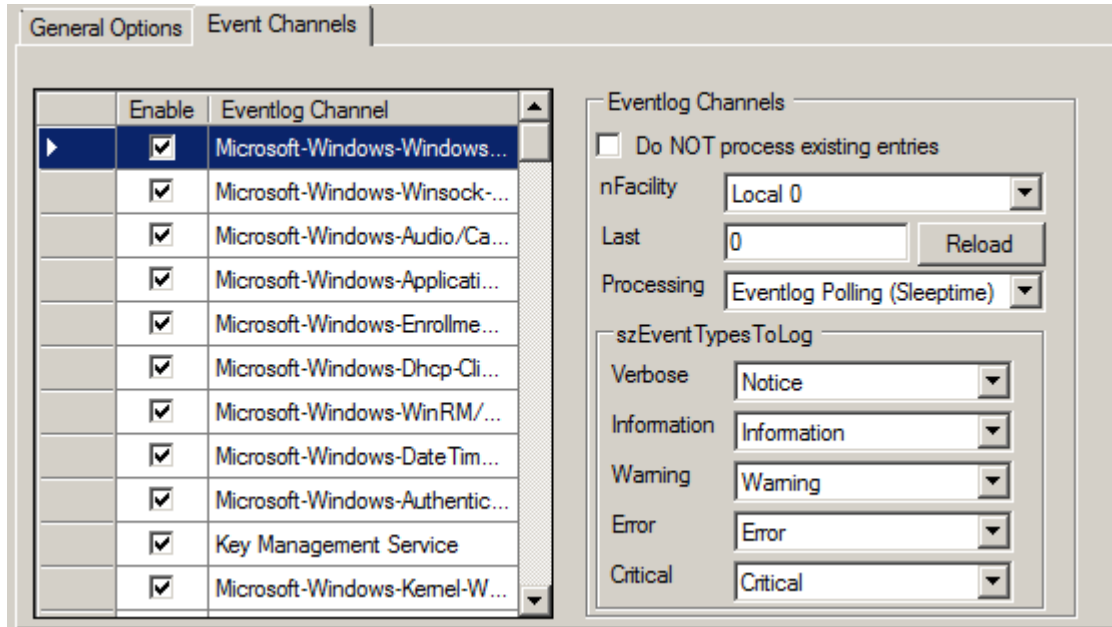
If enabled, the <EventData> node from the raw XML stream (Eventlog entry) will be searched for variables. If variables with names are found, they will be set as Properties with their variable name automatically. If the variable does not have a name, it will be set to a common name like "Param1, Param2 ParamX".

Convert to EventLog Monitor V1

This option maps the EventID's from the Security EventLog back to V1 (Windows 2000/2003). The internal InforUnitID is also changed to V1. This option helps

postprocessing EventLog V1 and V2 events equally.

Event Channels Tab



The most important part of this dialog is the treeview of available Channels. It specifies which event logs are to be monitored. In the screenshot above, the monitor is set to all Channels that are currently available. There happen to be custom Channels, too, due to Applications creating them on their own. They will be added to the treeview automatically every time you re-open this configuration window.

Channels checked in the table are actually processed. Those unchecked are kept in the configuration, but are not processed.

Here you can adjust the syslog facility and the event log types. You are also able to overwrite all existing custom advanced channel configurations with your new ones.

Facility

The [Syslog facility](#) to map information units stemming from this log to. Most useful if the message is to forward to a Syslog daemon.

Last Record

Windows event log records are numbered serially, starting at one. The service records the last record processed. This textbox allows you to override this value. **Use it with caution!**

If you would like a complete dump of a specific Windows event log, reset the "Last Record" to zero. If you missed some events, simply reset it to some lower value than

currently set. It is possible to set "Last Record" to a higher value. This suspends event reporting until that record has been created. We strongly discourage to use this feature unless definitely needed.

Processing Mode

There are two processing modes available, first the default processing mode is "EventLog Subscription" which processes Events in realtime. This means events are send to MonitorWare Agent by the OS as they happen, there is no delay at all. The other processing mode called "Eventlog Polling" and is similar to the method used in the old EventLog Monitor. The EventLog is checked and processed periodically controlled by the sleeptime. However using the polling method, you enable the "Read EventLog From File" option.

Event Types to Log

These checkboxes allow local filtering of the event log. Filtering is based on the Windows event type. There is a checkbox corresponding to each Windows event type. Only checked event types will be processed. Unchecked ones will be ignored.

Filtering out unnecessary log types at this level enhances system performance because no information units will be generated and passed to the rule engine. Thus, Adiscon strongly recommends dropping unnecessary log types.

6.3.6 SMTP Listener

The SMTP Listener is a service, which allows you to convert emails to syslog messages.

Status: Enabled | Comments | Save | Reset

Internet Protocoltype: IPv4

Listener Port: 25

Listener IPAddress: 0.0.0.0

Connection Timeout Limit: 15 seconds

General Values

Syslog Facility: Local 0

Syslog Priority: Notice

Syslog Tag Value: MWSMTPProbe

Resource ID:

RuleSet to use: | Refresh

SMTP Listener Services

Internet Protocoltype

Select the desired protocol type. IPv4 and IPv6 are available. The IPv6 protocol needs to be properly installed in order to be used. Note that one Service can only handle IPv4 or IPv6, so if you want to use both protocols, you will need to create two separate services.

Listener Port

This port is to be probed. Please see your server's reference for the actual value to use. For example, mail servers typically listen to port 25.

Listener IP Address

Either the IP address or resolvable host name of the SMTP server, the SMTP probe is to be run against. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. This system has been called "remote host" in the description above. Please note that specifying a host name can cause the SMTP probe to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address.

Connection Timeout limit

The Timeout limit specifies the time the listener waits for the sender.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

6.3.7 SNMP Monitor

SNMP Monitor will help monitoring SNMP capable devices. There are many devices that support SNMP and can be queried for information by SNMP GET. This can be printers, router, managed switches, linux / windows servers and so on.

The SNMP Monitor Service runs continuously based on the configuration mentioned below:

The screenshot shows the 'SNMP Trap Receiver Properties' dialog box. At the top, there is a status bar with 'Status: Enabled', a 'Comments' icon, and 'Save' and 'Reset' buttons. The main area contains the following fields:

- Probe Interval: 1 Minute
- Timeout limit: 5 seconds
- Internet Protocoltype: IPv4
- Protocol Type: UDP
- Remote host: (empty)
- SNMP Port: 161

Below these is a section titled 'SNMP Query related' with the following fields:

- Community: public
- SNMP Version: All supported Versions
- Query OID (Object ID): .1.3.6.1.2.1.1.5 (with a 'Browse' button)

At the bottom, there is a 'RuleSet to use' dropdown menu and a 'Refresh' button.

SNMP Trap Receiver Properties

Probe Interval

This is the interval of the queries. After each probe, the MonitorWare Agent SNMP Monitor process goes "to sleep". This period is specified in milliseconds.

Timeout limit

The amount of time (in [milliseconds](#)) the remote system is expected to answer in. If no response is received within this period, the ping fails and an event is generated. The default value of 1000 milliseconds is a proper value for most well connected networks. If the ping probe runs against a heavily loaded system and/or slow network link, the amount must be adjusted accordingly.

Internet Protocoltype

Select the desired protocol type. IPv4 and IPv6 are available. The IPv6 protocol needs to be properly installed in order to be used. Note that one Service can only handle IPv4 or IPv6, so if you want to use both protocols, you will need to create two separate services.

Protocol Type

Syslog messages can be received via [UDP](#), [TCP](#) or [RFC 3195](#) RAW. One listener can

only listen to one of the protocols. Typically, Syslog messages are received via UDP protocol, which is the default. The syslog server also can receive Syslog messages via TCP and reliable Syslog messages via TCP using the new [RFC 3195](#) RAW standard.

Remote host

Either the IP address or resolvable host name of the system you want to monitor. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. Please note that specifying a host name can cause the SNMP monitor to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address.

SNMP Port

The port the SNMP listener is listening to. If in doubt, leave it at the default of 161.

Community

Specify the SNMP community to which the messages belong too.


SNMP Version

Can be used to restrict the SNMP versions. The available values are:

1. SNMP Version 1 only
2. SNMP Version 2c only

Query OID (Object ID)

This is the Object ID you will query the device for. You can use Browse option to select your OID. If you click the Browse link, the screen similar to shown below is appeared:



MIB Browser

You can select your OID here.

Instance Subidentifier

The Instance Subidentifier defines the message you query, if the selected OID has multiple data.

Rule Set to Use

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.8 RELP Listener

The [RELP](#) listener support the new reliable event logging protocol (RELP), which enables a more reliable transmission of messages than plain tcp syslog protocol. The service permits to accept messages from senders who themselves support RELP.

Other than that it is using a different communications protocol, the RELP listener is functionally equivalent to the syslog listener. The RELP Listener will automatically listen on all available IP Addresses which includes IPv4 and IPv6. This is due the librelp implementation method.




Figure 1: RELP Listener Properties

Listener Port

The port the RELP Listener listens on. The typical (standard) value is 20514. This should be changed only if there is a definite need for it. Such a need typically arises from security concerns. If the port is changed, all reporting devices must also be configured to use the non-standard port.

Session Timeout

It controls how long a session is to be opened from the server side.

Rule Set to Use

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.9 Database Monitor

The database monitor is used to monitor database tables. It periodically checks a database table for new records and if it finds them, generates an event from each record. A table that is to be monitored by the database monitor **must** have an integer ID field that auto-increments.

Please note that the database monitor transmits all of the data obtained within its [event properties](#). This means that you **must** use an output format suitable to show event properties if you intend to process the record with a third party application. We strongly recommend using XML based formatting for this. Alternatively, you can also select a format for the msg property itself (which is the default message). To do so, you need to configure the database monitor's advanced option's msg field settings.

Name: DatabaseMonitor | Status: Enabled | Comments | Settings | Save | Reset | Configure for...

Check Interval (ms)

Overrun Prevention Delay (ms)

Database Properties | MessageField (%msg%)

DSN

User-ID

Password

ID Field name

szSelectStatement

Select Statement

FROM

Database Table

Maximum char length (Bytes)

Maximum text length (Bytes)

General Values

Syslog Facility

Syslog Priority

Syslog Tag Value

Resource ID

RuleSet to use Refresh

Database Monitor Properties

Check Interval

The database monitor runs periodically. This specifies, how often it should run. Please note that the Database Monitor waits for the configured amount of time after the current run is finished. The time is the delay in [milliseconds](#).

Overrun Prevention Delay

This property allows configuring a delay after generating an event. The time is the delay in [milliseconds](#).

If run at a value of zero, the service generates events as fast as the machine permits. We have seen scenarios where routers and receivers are not able to keep up with this rate, resulting in packet loss. In addition, the CPU of the reporting machine is run at 100% - which is not a problem because the service runs at a low priority. However, with even a 1-millisecond delay, there is no noticeable CPU activity even when large bursts of events are forwarded. At one millisecond, the service can still generate 1000 events per second.

The default setting is an overrun protection of five millisecond, which allows roughly 200 events per second. This should be sufficient for even very busy servers.

Database Properties Tab

Here you can configure the access data for the database.

DSN

The data source name of the database to access. All databases that support standard SQL syntax and have an ODBC driver support.

User-ID

The User ID to log on to the database system. Please note that the database system may not require this setting or may ignore it (e.g. Microsoft SQL Server in "integrated" security mode) - please check with your database vendor or your administrator if in doubt.

Password

The password to access the database. If the "Enable Encryption" check box is checked, a **weak** symmetrical encryption is applied on the password. Thus, we highly recommend to create a specific account with very limited permissions if you store a password. This account does only need to have "select" permissions.

Select Statement

You can configure a Select Statement to be issued to the database. This way, you have full control over what is fetched from the database.

Please note that if you specify specific fields, the ID field must be present in the select clause - otherwise the service can not process the records.

Also be sure to enter the name of the table from which the records are to be taken.

Message Field Tab

If you click on the "Message Field (%msg%)" tab, you can configure the following additional properties:

The screenshot shows the 'Database Properties' window with the 'MessageField (%msg%)' tab selected. It features three radio buttons: 'None', 'CSV (Comma-separated values)', and 'Custom'. The 'None' option is selected. Below these are two checkboxes: 'Use custom separator' (unchecked) and 'Custom' (checked). A text field labeled 'Custom separator' contains the value '%\$TAB%'. At the bottom, there is a 'Custom Message Content' section with a text area containing the text: 'Actual values are in XML stream - see http://www.monitorware.com/Common/en/FAQ/dbmon-values.php for more information.' and an 'Insert' button.

Database Monitor Message Field Tab

Message Field

This field specifies the content of the "[msg](#)" property. By default, msg does NOT contain any useful information. This is because all data is provided via the event properties. If you actually need this as part of the msg, you can either select a custom format or CSV format.

If CSV is selected, msg contains all field values (not names) in comma-separated format. Instead of a comma you can also use any other custom separator. The field order is as it was in the select statement. We recommend **not** to use "select *" in this case (but specify the fields inside the select statement, so that they have a definite order).

We recommend **not** to rely on msg with the database monitor. Access via properties is much better. If you do not need msg, we recommend to set the msg content to "none" as this removes the unnecessary default message.

General Values

Syslog Facility

The [Syslog facility](#) to be assigned to events created by the service. Most useful if the message shall be forwarded to a syslog daemon.

Syslog Priority

The Syslog priority to be assigned to events created by the service. Most useful if the message shall be forwarded to a syslog daemon.

Resource ID

The [Resource ID](#) to be assigned to events created by the service. Most useful if the message shall be forwarded to a Syslog daemon.

Syslog Tag Value

The Syslog tag value to be assigned to events created by the service. Most useful if the message shall be forwarded to a Syslog daemon.

Default Ruleset Name

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.10 Serial Port Monitor

The serial port monitor allows you to monitor devices attached to local communication ports. Actually, this is not limited to serial (RS232) devices - devices connected via e.g. LPT ports can also be monitored as long as the device provides a proper interface to the port device.

For Example - uses for the serial port monitor may be interfacing to data loggers, "strange" log sources (e.g. PBX call logs) or out-of-band log retrieval (e.g. setting a router to log to the serial port instead to the network and then picking the data from that serial line). Out-of-band log retrieval can also be used to hide the fact that logging is actually taking place.

The serial port monitor works as follows: it listens to the configured port. With each received character, it checks if a configured "message end sequence" is received. If it isn't it continues listening until either another character is received or a timeout occurs.

If either the "message end sequence" is received or the timeout occurred, the message is considered to be complete. In this case, an event is generated and that event is scheduled for processing.

Status: Enabled | Comments | Save | Reset | ?

Timeout Limit: 1 Minute

How the message is ended: \r\n

Send this on startup:

Which Port do you want to monitor: COM1:

Port Settings

Bits per second: 57600

Data bits: 8

Parity: No Parity

Stop bits: 1 Stop bit

DTR Control Flow: DTR Control Disable

RTS Control Flow: RTS Control Disable

General Values

Syslog Facility: Local 0

Syslog Priority: Notice

Syslog Tag Value: MWSerialPortMonitor

Ressource ID:

RuleSet to use: Refresh

Filter Conditions - Serial Port Monitor

Timeout Limit

This is the maximum amount of time the service waits to receive the "message end characters" from the attached device.

How the message is ended

This is the sequence that, when received, denotes the end of the message. Most often, this is either "\r\n" or "\n" ("\r" represents a CR characters, "\n" represents a LF character).

Which Port do you want to monitor

This is the port that the device is attached to. Most often, this is either COM1: or COM2:. All locally-existing ports can be used. When working locally, the configuration program enumerates the locally available ports. This can be one of the following

values:

1. MSFAX
2. COM1
3. COM2
4. COM3
5. COM4
6. FILE
7. LPT1
8. LPT2
9. LPT3
10. AVMISDN1
11. AVMISDN2
12. AVMISDN3
13. AVMISDN4
14. AVMISDN5
15. AVMISDN6
16. AVMISDN7
17. AVMISDN8
18. AVMISDN9

Port Settings

These settings must be set as expected by your device. If in doubt, consult your device manual.

Bits per Seconds

Bits per second can be 110 and go up to 256000, by default 57600 is selected.

Databits

Databits define the number of bits in the bytes transmitted and received.

Parity

With Parity you can configure the Parity scheme to be used. This can be one of the following values:

1. Even
2. Mark
3. No parity
4. Odd
5. Space

Stop bits

You can configure the Number of stop bits to be used. This member can be one of the following values:

1. 1 stop bit
2. 1.5 stop bits
3. 2 stop bits

DTR Control Flow

DTR (data-terminal-ready) flow control. This member can be one of the following values:

1. DTR_CONTROL_DISABLE - Disables the DTR line when the device is opened and leaves it disabled.
2. DTR_CONTROL_ENABLE - Enables the DTR line when the device is opened and leaves it on.
3. DTR_CONTROL_HANDSHAKE - Enables DTR handshaking.

RTS Control Flow

RTS (request-to-send) flow control. This member can be one of the following values:

1. RTS_CONTROL_DISABLE - Disables the RTS line when the device is opened and leaves it disabled.
2. RTS_CONTROL_ENABLE - Enables the RTS line when the device is opened and leaves it on.
3. RTS_CONTROL_HANDSHAKE - Enables RTS handshaking. The driver raises the RTS line when the "type-ahead" (input) buffer is less than one-half full and lowers the RTS line when the buffer is more than three-quarters full.
4. RTS_CONTROL_TOGGLE - Specifies that the RTS line will be high if bytes are available for transmission. After all buffered bytes have been sent, the RTS line will be low.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by the service. Most useful if the message shall be forwarded to a syslog daemon.

Syslog Priority

The Syslog priority to be assigned to events created by the service. Most useful if the message shall be forwarded to a syslog daemon.

Resource ID

The [Resource ID](#) to be assigned to events created by the service. Most useful if the message shall be forwarded to a Syslog daemon.

Syslog Tag Value

The syslog tag value to be assigned to events created by the service. Most useful if the message shall be forwarded to a syslog daemon.

Default Ruleset Name

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.11 File Monitor

The file monitor monitors the content of a text file just as the event monitor monitors the NT event log. Its purpose is to gather vital information that is stored in system text files. Many applications do not write events to the event log but to a text file. This is also the case with many Microsoft applications (for example the WINS log).

The file monitor can also gather Internet Information Server (Windows' web server) log files. This is very useful for monitoring web activity and detecting attacks.

General Tab

Figure 1: General Options

File and path name

Here, you must type the name of the file to be monitored. To select a file from any specified location, press the browse button. Once a complete file name is specified, exactly that file is monitored.

The file name is never changed automatically. However, many systems generate changing log files. For example, Internet Information Server (IIS) can be configured to change the log file every day. Therefore, each day's log file has a different name.

To support changing log file names, there are replacement characters available within the file name. These are:

Character	Meaning
%y	Year with two digits (e.g. 2002 becomes "02")

%Y	Year with 4 digits
%m	Month with two digits (e.g. March becomes "03")
%M	Minute with two digits
%d	Day of month with two digits (e.g. March, 1 st becomes "01")
%h	Hour as two digits
%S	Seconds as two digits. It is hardly believed that this ever be used in reality.
%w	Weekday as one digit. 0 means Sunday, 1 Monday and so on.
%W	Weekday as three-character string. Possible values are "Sun", "Mon", "Tue", "Wed", "Thu", "Fri", "Sat". This replacement character is most useful for DHCP log files.
%generatedfilename%	It contains the fully generated filename (Can be useful for filtering).
%msgsep%	Only available if enable in the advanced option of the File Monitor. This value contains the current used messageseparator. This is usefull if you want to reconstruct messages where the seperator is part of the message.
%msgseplast%	Only available if enable in the advanced option of the File Monitor. This value contains the last used messageseparator. This is usefull if you want to reconstruct messages where the seperator is part of the message.

Character Replacement Table

Please note: the replacement characters are case sensitive!

For example, daily Internet Information Server log files are named "exyymmdd.log", with yy being the 2 digit year, mm the month and dd the day of month. To generate the same name with file monitor, use the following name "ex%y%m%d.log".

Please note that there is no replacement character for the monthly week number (1st week, 2nd week). As such, the weekly log file setting of IIS is not supported.

TimeMode Used for Filename

Select the time mode used for the log file to be monitored with this drop-down list. Available options are:

1. Local Time: log file is monitored based on local time.
2. [UTC](#): log file is monitored based on universal coordinated time. UTC was formerly referred to as "GMT" (Greenwich Mean Time) and is the basis of the international time zone system.

Allow Directories or read multiple files

This is the new Multiple Files feature which means you can now read an array of files. This will require a wildcard in the filename. If using directories, the amount of subdirectories is unlimited.

Use Wildcard in Filename

This option allows you use * as random character sequence in the filename.

Please note: this character can only be used in Filename and not in the filepath.

Keep reading the current opened file until a new is created

This has been added to define if the Service shall continuously read an open logfile until a new file (depending on the configured filename) is available. This Options is helpful for such cases where you don't know when a new logfile is generated and the old one is closed.

Report an Error if the File was not found

As the name says, if this setting is enabled, an error is reported in the Windows Eventlog if the file was not found.

Skip all lines on Startup

If this option is enabled, the File Monitor will skip all new lines of a logfile during startup. This will work in singlefile mode as well as multifile mode.

Check Interval

This interval is in [milliseconds](#). After the specified interval the file monitor checks the file for new records.

We recommend a value of 60000 milliseconds for the "Check Interval". With that setting, the file monitor checks for new records every 60 seconds. Larger periods can be specified for occasionally connected systems or if email delivery with few emails per day is intended.

Very security-aware environments might use a shorter interval. The MonitorWare Agent 3.0 is specifically designed to limit the burden on the monitored system. As such, resource usage is typically low, even with frequently run file monitor checks. However, we recommend not running the file monitor more often than once a second.

Overrun Prevention Delay

This property allows configuring a delay after generating an event. The time for the delay is in [milliseconds](#).

If run at a value of zero, the MonitorWare Agent 3.0 generates events as fast as the machine permits. We have seen scenarios where routers and receivers are not able to keep up with this rate, resulting in packet loss. In addition, the CPU of the reporting machine is run at 100% - which is not a problem because MonitorWare Agent 3.0 runs at a low priority. However, with even a 1-millisecond delay, there is no noticeable CPU activity even when large bursts of events are forwarded. At one millisecond,

MonitorWare Agent 3.0 can still generate 1000 events per second.

The default setting is an overrun protection of five millisecond, which allows roughly 200 events per second. This should be sufficient for even very busy servers.

Logfile Type

Select the type of the log file to be monitored with this drop-down list. Available options are:

1. Standard - a standard text log file
2. W3C Web Server logfile - log files in the W3C web server compliant format.

Advanced Options Tab

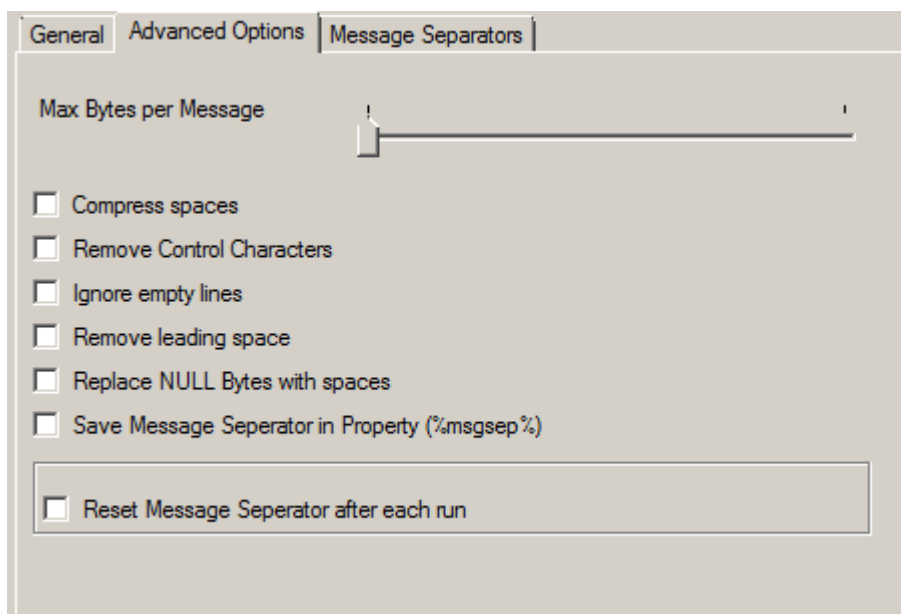


Figure 2: Advanced Options

Max Bytes per Message

Maximum value of bytes that the file monitor reads per line. If a message is larger than this value, the message splits into multiple parts.

Compress spaces

This option compresses sequences of spaces found inside the message to a single one.

Remove Control Characters

Removes control characters like CR and LF(carriage return and line feed).

Ignore empty lines

As the name already says, this option discards empty lines within the logfile.

Remove leading space

If there are any leading spaces in the file, this option removes them.

Replace NULL Bytes with spaces

If this option is enabled, the FileMonitor will replace NULL Bytes within files with spaces. These spaces again can be compressed with the "Compress spaces" options.

Save Message Separator in Property

If this option is enabled, the current and last used message separator will be saved into the properties %msgsep% and %msgseplast%.

Reset Message Separator after each run

If enabled, the message separator values will be resetted after the File Monitor has finished a run (reached the end of a file).

Message Separators Tab

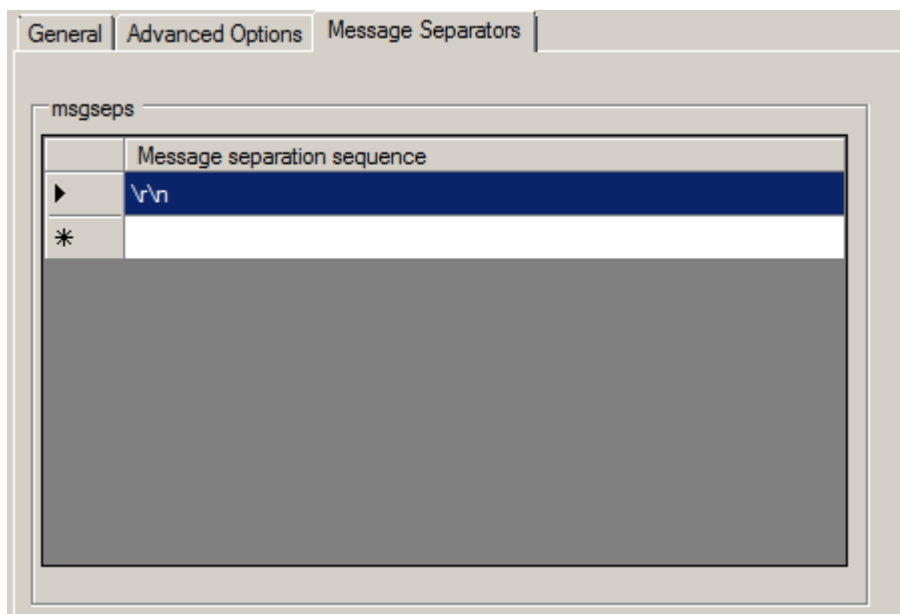


Figure 2: Message Separators

Message separation sequence

The customizable separation sequence when this option is enabled. The file monitor splits messages by this value. If it is disabled crlf (carriage return line feed) is used. If using multiple separation sequences, the comparison operation will be held as an OR operation. That means, that either this value or another value has to be true, so a message can be split. This is especially important for logfiles with different log formats.

To date, the following characters are available:

Character	Meaning
\r	carriage return
\n	line feed

General Values

Last Line Value

This value contains the last read line of the FileMonitor Service. The file Monitor Service reads a configured file continuously line by line and everytime there is a new line, this value is incremented.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by the service. Most useful if the message shall be forwarded to a Syslog daemon.

Syslog Priority

The Syslog priority to be assigned to events created by the service. Most useful if the message shall be forwarded to a Syslog daemon.

Resource ID

The [Resource ID](#) to be assigned to events created by the service. Most useful if the message shall be forwarded to a Syslog daemon.

Syslog Tag Value

The Syslog tag value to be assigned to events created by the service. Most useful if the message shall be forwarded to a Syslog daemon.

Default Ruleset Name

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

Further Reading

Please visit our white paper on [monitoring IIS logs](#).

6.3.12 Heartbeat

The heartbeat process can be used to continuously check if everything is running well. It generates an information unit every specified time interval. That information unit can be forward to a different system. If it does not receive additional packets within the configured interval, it can doubt that the sender is either in trouble or already stopped running.

The screenshot shows the 'Heartbeat Properties' configuration window. At the top, the status is 'Enabled'. There are buttons for 'Comments', 'Save', 'Reset', and a help icon. The main configuration area includes a text box for the message to send, currently containing 'I am still running'. Below this is a dropdown menu for the heartbeat clock (sleeptime), currently set to '1 Minute'. A section titled 'General Values' contains several fields: 'Syslog Facility' (dropdown, 'Local 0'), 'Syslog Priority' (dropdown, 'Notice'), 'Syslog Tag Value' (text box, 'MWHeartbeat'), and 'Ressource ID' (empty text box). At the bottom, there is a dropdown for 'RuleSet to use' and a 'Refresh' button.

Heartbeat Properties

Message to Send

This is the message that is used as text inside the information unit. Use whatever value is appropriate. The message text does not have any special meaning, so use whatever value you seem fit.

Sleep Time

This is the interval, in [milliseconds](#), that the heartbeat service generates information units in. **Please note that the receiving side should be tolerant.** The interval specified here is the minimum time between packets. Under heavy load, the interval might be slightly longer. It is good practice to allow twice this interval before the service is considered suspect by the system monitoring the services health.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by the heartbeat service. Most useful if the message shall be forwarded to a syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by the heartbeat process. Most useful if the message shall be forwarded to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by the heartbeat process. Most useful if the message shall be forwarded to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by the heartbeat process. Most useful if the message shall be forwarded to a Syslog server.

RuleSet to Use

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.13 Ping Probe

The ping probe can be used to check the health of a remote system. The ping probe process sends ping messages (more precisely: ICMP Echo Requests) to a configured system. If configured properly, the remote system sends a response. If this response is received, the machine and its IP stack are operating. This does not indicate, however, that all services on this machine are alive.

If no response is received, the remote system or its IP stack is most probably not operating properly. However, the ping message might have been lost in transit or the round-trip time might have been too long so that a timeout occurred. Therefore, a single failing ping makes a system suspect, but it alone cannot be used to confirm problems at the remote system. If multiple successive pings fail, it is relatively safe to assume that the remote system has failed

Please note that most firewall setups do not allow ping messages. As such, a system behind a firewall typically cannot be pinged and the ping probe cannot be used in this configuration. If in doubt, please check with your firewall administrator.

The ping probe is typically used to check the availability of a remote system. The ping probe periodically sends the ping messages. As long as responses are received, nothing happens. If no response is received, it generates an event and passes it to the rule engine. As ping messages can get lost, the ping probe retries failed probes before it reports an error. Both the number of retries and the retry interval can be specified.

Status: Enabled | Comments | Save | Reset | ?

Probe Interval: 1 Minute

Timeout Limit: 5 seconds

IP Address or Hostname:

Number Of Retries: 3

Retry Interval (ms): 5 seconds

Generate an event if PingProbe was successful

General Values

Syslog Facility: Local 0

Syslog Priority: Notice

Syslog Tag Value: MWPingProbe

Resource ID:

RuleSet to use: | Refresh

Ping Probe Properties

Probe Interval

This is the interval of the ping probes. After each probe, the MonitorWare Agent 3.0 ping probe process goes "to sleep". This period is specified in milliseconds.

Timeout Limit

The amount of time (in [milliseconds](#)) the remote system is expected to answer in. If no response is received within this period, the ping fails and an event is generated. The default value of 1000 milliseconds is a proper value for most well connected networks. If the ping probe runs against a heavily loaded system and/or slow network link, the amount must be adjusted accordingly.

IP Address / Hostname

Either the IP address or resolvable host name of the system, the ping probe is to be run against. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. This system has been called "remote host" in the description above. Please note that specifying a host name can cause the ping probe to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address.

Please provide the IP address or the hostname according to your environment. We have left it empty by intention.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message is forwarded to a Syslog server.

Number of Retries

If a ping fails, it is first retried to see if it is a persistent problem. The "Number of Retries" controls how many retries to be made. If this is set to zero, no retries are made and a ping probe fail event is immediately generated. For typical systems, we recommend a setting of three retries. This is also the default value.

Retry Interval

If there is a temporary network issue like network congestion, it most probably takes some seconds to resolve it. As such, an immediate retry might not be appropriate. To delay it, configure a retry interval. This value is in [milliseconds](#). If a ping fails, the next retry is after a pause specified in this property. The default and recommended value is 5 seconds (5000 milliseconds).

Generate an event if Ping Probe was successful

When checked, an event is generated every time. If unchecked, it is generated only when the ping fails. The most common option is to leave it unchecked to catch events upon a failed ping.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message is forwarded to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message is forwarded to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message is forwarded to a Syslog server.

Default Ruleset Name

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.14 Port Probe

The port probe is very similar to the ping probe described above. The main difference is that it does not check the IP stack availability but rather a specific [TCP](#) port.

The difference here is that using this method a specific service on the remote machine is monitored, for example a mail ([SMTP](#)) server. The port probe tries to connect to the service port (25 in our example). If that fails, the service is definitely not running. In this case, an event is generated. A single event is a definite indication of problems, as such there is no need for repetitive failures before initiating action on this (although this can be configured in the rule engine).

Being able to connect to the remote machine and service, [TCP](#) port most probably means that the remote service is running. However, more certainty can be gained by actually initiating some communication with the service. The exact application protocol needs to be known to try this test. Thus, this step is optional. If turned on, a single command can be send to the remote service and a single response is expected back and can be compared to a pre-defined response. This does not take care of all possible application protocols, but provides an additional layer of confidence for important services like SMTP. It is up to the user to know the command sequences that a given service can understand and reply with.

As a rule of thumb, the port probe provides superior protection against service failure even without checking the message exchange. So if in doubt, use it without this advanced feature.

Please note that the port probe can probe TCP based services only. Most application services are TCP based, but there are some – mostly system – services out there, that are not. One of the most notable exceptions is DNS, which is operated primarily over [UDP](#). In UDP, there is no notion of a session and as such, it is not possible to probe the session setup, which essentially is what the port probe does. As such, a port probe can unfortunately not be used to check the status of those services. However, the majority of services like application server, databases, mail, web and a large number of others can be used with the port probe.

Status: Enabled | Comments | Save | Reset | ?

Probe Interval: 1 Minute

Timeout Limit: 1 second

IP Address or Hostname:

Port: 0

Generate an event if PingProbe was successful

Send Message and check for expected message

Message to send:

Message to expect:

General Values

Syslog Facility: Local 0

Syslog Priority: Notice

Syslog Tag Value: MWPportProbe

Resource ID:

RuleSet to use: | Refresh

Port Probe Properties

Probe Interval

This is the interval of the port probes. After each probe, the MonitorWare Agent 3.0 port probe process goes "to sleep". This period is specified in [milliseconds](#).

Timeout Limit

The amount of time (in [milliseconds](#) the remote system is expected to answer in. If no response is received within this period, the ping fails and an event is generated. The default value of 1000 milliseconds is a proper value for most well connected networks. If the ping probe runs against a heavily loaded system and/or slow network link, the amount must be adjusted accordingly.

IP Address / Hostname

Either the IP address or resolvable host name of the system the ping probe is to be run against. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. This system has been called "remote host" in the

description above. Please note that specifying a host name can cause the ping probe to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address. **Please note that you typically can use 127.0.0.1 (the so-called loop back address) to check a service that is running on a local machine. This ability might be limited by service configuration, because the service must listen to that IP address.**

Please provide the IP address or the hostname according to your environment. We have left it empty by intention.

Port

This port is to be probed. Please see your server's reference for the actual value to use. For example, mail servers typically listen to port 25 and web servers to 80.

Generate an event if Port Probe was successful

When checked, an event is generated every time. If unchecked, it is generated only when the port probe fails. The most common option is to leave it unchecked to catch events upon a failed port probe.

Send Message and wait for expected Message

If left unchecked, the port probe checks the [TCP](#) session setup to the remote service only. As stated above, a successfully completed session setup most probably means the service is healthy. As an extra measure, some actual message exchange can be enabled. This is done by checking this box.

Message to Send

This message text is sent to the service after the [TCP](#) session has been established.

Message Expected

This is the message expected to be received from the service. Reception starts after sending the "Message to Send". Please note that the "Message Expected" is compared against the **first** message sent from the service on the [TCP](#) session. With some protocols, this means the message compared is an initial greeting message and **not** a response to the "Message to Send".

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Default Ruleset Name

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.15 SMTP Probe

SMTP probe does a connection to SMTP server and sends the HELLO command. The HELLO command is automatically constructed by MonitorWare Agent on startup by using the fully qualified DNS (Domain name server) name. SMTP server sends response in reply to SMTP probe. On getting the response from SMTP server, SMTP probe sends the QUIT command to terminate the connection.

The connection status is saved in the property **smtpstatus** and the response in the property **smtprespmsg**.

The screenshot shows the configuration window for SMTP Probe Services. At the top, the status is set to 'Enabled'. There are buttons for 'Comments', 'Save', 'Reset', and a help icon. The main configuration area includes:

- Probe Interval:** A dropdown menu set to '1 Minute'.
- Timeout Limit:** A dropdown menu set to '5 seconds'.
- Generate an event if SMTP Probe was successful**
- SMTP Server:** A text input field containing '127.0.0.1'.
- SMTP Port:** A text input field containing '25'.
- General Values:** A section containing:
 - Syslog Facility:** A dropdown menu set to 'Local 0'.
 - Syslog Priority:** A dropdown menu set to 'Notice'.
 - Syslog Tag Value:** A text input field containing 'MWSMTPProbe'.
 - Ressource ID:** An empty text input field.
- RuleSet to use:** A dropdown menu with a 'Refresh' button next to it.

SMTP Probe Services

Probe Interval

This is the interval of the SMTP probes. After each probe, the MonitorWare Agent 3.0 SMTP probe process goes "to sleep". This period is specified in milliseconds.

Timeout Interval

The amount of time (in [milliseconds](#)) the remote system is expected to answer in. If no response is received within this period, the SMTP probe fails and an event is generated. The default value of 1000 milliseconds is a proper value for most well connected networks. If the SMTP probe runs against a heavily loaded system and/or slow network link, the amount must be adjusted accordingly.

Also generate an event if SMTP Probe was successful

When checked, an event is generated every time. If unchecked, it is generated only when the SMTP probe fails. The most common option is to leave it unchecked to catch events upon a failed SMTP probe.

SMTP Server

Either the IP address or resolvable host name of the SMTP server, the SMTP probe is to be run against. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. This system has been called "remote host" in the

description above. Please note that specifying a host name can cause the SMTP probe to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address.

SMTP Port

This port is to be probed. Please see your server's reference for the actual value to use. For example, mail servers typically listen to port 25.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

6.3.16 POP3 Probe

POP3 probe does a connection to POP3 server. It receives the response from POP3 server and sends the QUIT command to terminate the connection.

The connection status is saved in the property **pop3status** and the response in the property **pop3respmsg**.

The screenshot shows the 'POP3 Probe Properties' configuration window. At the top, there is a status bar with 'Status: Enabled', a 'Comments' icon, and buttons for 'Save', 'Reset', and a help icon. The main configuration area includes:

- Probe Interval:** A dropdown menu set to '1 Minute'.
- Timeout Limit:** A dropdown menu set to '5 seconds'.
- Generate an event if POP3 Probe was successful**
- POP3 Server:** A text input field containing '127.0.0.1'.
- POP3 Port:** A text input field containing '110'.
- General Values:** A section containing:
 - Syslog Facility:** A dropdown menu set to 'Local 0'.
 - Syslog Priority:** A dropdown menu set to 'Notice'.
 - Syslog Tag Value:** A text input field containing 'MWPOP3Probe'.
 - Resource ID:** An empty text input field.
- RuleSet to use:** A dropdown menu and a 'Refresh' button.

POP3 Probe Properties

Probe Interval

This is the interval of the POP3 probes. After each probe, the MonitorWare Agent 3.0 POP3 probe process goes "to sleep". This period is specified in milliseconds.

Timeout Interval

The amount of time (in [milliseconds](#)) the remote system is expected to answer in. If no response is received within this period, the POP3 probe fails and an event is generated. The default value of 1000 milliseconds is a proper value for most well connected networks. If the POP3 probe runs against a heavily loaded system and/or slow network link, the amount must be adjusted accordingly.

Also generate an event if POP3 Probe was successful

When checked, an event is generated every time. If unchecked, it is generated only when the POP3 probe fails. The most common option is to leave it unchecked to catch events upon a failed POP3 probe.

Pop3 server

Either the IP address or resolvable host name of the POP3 server, the POP3 probe is to be run against. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. This system has been called "remote host" in the description above. Please note that specifying a host name can cause the POP3 probe

to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address.

Port

This port is to be probed. Please see your server's reference for the actual value to use. For example, mail servers typically listen to port 25.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

6.3.17 FTP Probe

FTP probe does a connection to FTP server. It receives the response from FTP server and sends the QUIT command to terminate the connection.

The connection status is saved in the property **ftpstatus** and the response in the property **ftprespmsg**.

The screenshot shows the 'FTP Probe Properties' configuration window. At the top, there is a status dropdown set to 'Enabled', and buttons for 'Comments', 'Save', 'Reset', and a help icon. The main configuration area includes:

- Probe Interval:** A dropdown menu set to '1 Minute'.
- Timeout Limit:** A dropdown menu set to '1 second'.
- Generate an event if FTP Probe was successful**
- FTP Server:** A text input field containing '127.0.0.1'.
- FTP Port:** A text input field containing '21'.
- General Values:** A section containing:
 - Syslog Facility:** A dropdown menu set to 'Local 0'.
 - Syslog Priority:** A dropdown menu set to 'Notice'.
 - Syslog Tag Value:** A text input field containing 'MWFTPProbe'.
 - Ressource ID:** An empty text input field.
- RuleSet to use:** A dropdown menu with a 'Refresh' button next to it.

FTP Probe Properties

Probe Interval

This is the interval of the FTP probes. After each probe, the MonitorWare Agent 3.0 FTP probe process goes "to sleep". This period is specified in milliseconds.

Timeout Interval

The amount of time (in [milliseconds](#)) the remote system is expected to answer in. If no response is received within this period, the FTP probe fails and an event is generated. The default value of 1000 milliseconds is a proper value for most well connected networks. If the FTP probe runs against a heavily loaded system and/or slow network link, the amount must be adjusted accordingly.

FTP server

Either the IP address or resolvable host name of the FTP server, the FTP probe is to be run against. This system has been called "remote host" in the description above. Please note that specifying a host name can cause the FTP probe to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address.

Port

This port is to be probed. Please see your server's reference for the actual value to

use. For example, mail servers typically listen to port 25.

Also generate an event if FTP Probe was successful

When checked, an event is generated every time. If unchecked, it is generated only when the FTP probe fails. The most common option is to leave it unchecked to catch events upon a failed FTP probe.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

RuleSet to Use

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.18 NT Services Monitor

The NT Services Monitor is used to monitor if vital operating services are running. The monitor continuously checks all services set to "automatic" startup. If such a service does not run, an event is generated and passed to the rule engine.

Name: NTServiceMonitor | Status: Enabled | Comments | Settings | Save | Reset | Configure for...

Probe Interval: 0 (disabled) milliseconds

Delay on Startup in Milliseconds (1000 ms): 0 (disabled) milliseconds

Generate an event if a Service is in the running state

General Values

Syslog Facility: Local 0

Syslog Priority: Notice

Syslog Tag Value: MWNTServiceMonitor

Ressource ID:

RuleSet to use: Default RuleSet Refresh

NT Service Monitor Properties

Check Interval

This is the interval in which the service status is checked. This period is specified in [milliseconds](#). The default is 60,000 ms, which is one minute. We recommend to lower this interval only if the server is performing very critical operations and service stops need to be detected in close real-time.

For performance reasons, we do not recommend using an interval of less than 2000 ms.

Delay on Startup

During system boot, the monitoring service eventually starts before all other services have been started. As such, the service monitor probably finds some services not running – simply because they are to be started very soon. Nevertheless, the service monitor still generates a "service not running" event.

To avoid this situation, use the startup delay setting. It specifies an amount of time (in [milliseconds](#)) that the service monitor is to hold right after startup. So during system boot, the operating system has a chance to start all other services before the service monitor comes into action.

The actual delay is very much depending on the number of services and hardware sizing of a particular server. Typically, a value 60,000 ms (one minute) should be a good value. But a busy server with many services might require a much higher value.

Also generate an event if a Service is in the running state

When checked, an event is generated every time. If unchecked, it is generated only when the Service probe fails. The most common option is to leave it unchecked to catch events upon a failed Service startup.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.

Rule Set to Use

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.19 HTTP Probe

The HTTP Probe connects to a HTTP Server, and sends a valid HTTP request as configured. It then either receives the header or header and content of a website, depending on how the service is configured (See Request Type).

Name: HTTPProbe | Status: Enabled | Comments | Settings | Save | Reset | Configure for...

Probe Interval: 1 Minute
Timeout Limit: 5 seconds

Generate an event if HTTP Probe was successful

General Values

HTTP Server: 127.0.0.1
URL Querystring: /index.html
HTTP Port: 21
Request Type: HEAD

Use secure https protocol

Referer:
UserAgent (Browser): Mozilla/4.0

URL Preview
<http://127.0.0.1:21/index.html>

General Values

Syslog Facility: Local 0
Syslog Priority: Notice
Syslog Tag Value: MWHTTPProbe
Resource ID:

RuleSet to use: Default RuleSet | Refresh

HTTP Probe Properties

Probe Interval

This is the interval of the HTTP probes. After each probe, the MonitorWare Agent 3.0 HTTP probe process goes "to sleep". This period is specified in milliseconds.

Timeout Limit

The amount of time (in [milliseconds](#)) the remote system is expected to answer in. If no response is received within this period, the HTTP probe fails and an event is generated. The default value of 1000 milliseconds is a proper value for most well connected networks. If the HTTP probe runs against a heavily loaded system and/or slow network link, the amount must be adjusted accordingly.

Generate an event if HTTP Probe was successful

When checked, an event is generated every time. If unchecked, it is generated only when the HTTP probe fails. The most common option is to leave it unchecked to catch events upon a failed HTTP probe.

HTTP Server

Either the IP address or resolvable host name of the HTTP server, the HTTP probe is to

be run against. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. Please note that specifying a host name can cause the HTTP probe to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address.

Url & QueryString

By default this is /index.html. This value is used to construct an URL which is previewed in a rectangular field under Use secure https Protocol option.

Port

This port is to be probed. Please see your server's reference for the actual value to use. For example, mail servers typically listen to port 25 and web servers to 80.

Request Type

The Request Type can be HEAD or GET. HEAD just receives the header of a website where GET receives the whole website content. When probing a web server, you should use HEAD in order to reduce network and processing overhead.

Use secure https Protocol

You can enable this option, if you want to query a web server using SSL (Secure Socket Layer). Note that the default port is changed from 80 to 443 here.

Referrer

An optional configuration option where you can specify a Referrer that is send in the HTTP header.

UserAgent (Browser)

It is also an optional value which can be used to specify an UserAgent that is send in the HTTP header.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

6.3.20 IMAP Probe

IMAP probe does connection to IMAP server. After receiving the response from IMAP server it sends the QUIT command to terminate the connection.

The connection status is saved in the property **imapstatus** and the response in the property **imaprespmsg**.

Name: NNTPProbe | Status: Enabled | Comments | Settings | Save | Reset | Configure for...

Probe Interval: 1 Minute

Timeout Limit: 5 seconds

Generate an event if IMAP Probe was successful

NNTP Server: 127.0.0.1

NNTP Port: 119

General Values

Syslog Facility: Local 0

Syslog Priority: Notice

Syslog Tag Value: MWNNTTPProbe

Resource ID:

RuleSet to use: Default RuleSet | Refresh

IMAP Probe Properties

Probe Interval

This is the interval of the IMAP probes. After each probe, the MonitorWare Agent 3.0 IMAP probe process goes "to sleep". This period is specified in milliseconds.

Timeout Interval

The amount of time (in [milliseconds](#)) the remote system is expected to answer in. If no response is received within this period, the IMAP probe fails and an event is

generated. The default value of 1000 milliseconds is a proper value for most well connected networks. If the IMAP probe runs against a heavily loaded system and/or slow network link, the amount must be adjusted accordingly.

Also generate an event if IMAP Probe was successful

When checked, an event is generated every time. If unchecked, it is generated only when the IMAP probe fails. The most common option is to leave it unchecked to catch events upon a failed IMAP probe.

IMAP Server

Either the IP address or resolvable host name of the IMAP server, the IMAP probe is to be run against. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. Please note that specifying a host name can cause the IMAP probe to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address.

Port

This port is to be probed. Please see your server's reference for the actual value to use. For example, mail servers typically listen to port 25 and web servers to 80.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

6.3.21 NNTP Probe

NNTP probe does a connection to NNTP server. After receiving the response from NNTP server it sends the QUIT command to terminate the connection.

The connection status is saved in the property **nntpstatus** and the response in the property **nntprespsmsg**.

The screenshot shows the configuration interface for the NNTP Probe. At the top, the status is 'Enabled'. Below this are several configuration options: 'Probe Interval' is set to '1 Minute', 'Timeout Limit' is set to '5 seconds', and there is an unchecked checkbox for 'Generate an event if IMAP Probe was successful'. The 'NNTP Server' is set to '127.0.0.1' and the 'NNTP Port' is set to '119'. A 'General Values' section contains 'Syslog Facility' set to 'Local 0', 'Syslog Priority' set to 'Notice', 'Syslog Tag Value' set to 'MWNNTTPProbe', and an empty 'Resource ID' field. At the bottom, there is a 'RuleSet to use' dropdown menu and a 'Refresh' button.

NNTP Probe Properties

Probe Interval

This is the interval of the NNTP probes. After each probe, the MonitorWare Agent 3.0 NNTP probe process goes "to sleep". This period is specified in milliseconds.

Timeout Interval

The amount of time (in [milliseconds](#)) the remote system is expected to answer in. If no response is received within this period, the NNTP probe fails and an event is generated. The default value of 1000 milliseconds is a proper value for most well connected networks. If the NNTP probe runs against a heavily loaded system and/or slow network link, the amount must be adjusted accordingly.

Also generate an event if NNTP Probe was successful

When checked, an event is generated every time. If unchecked, it is generated only when the NNTP probe fails. The most common option is to leave it unchecked to catch events upon a failed NNTP probe.

NNTP Server

Either the IP address or resolvable host name of the NNTP server, the NNTP probe is to be run against. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. This system has been called "remote host" in the description above. Please note that specifying a host name can cause the NNTP probe to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address.

Port

This port is to be probed. Please see your server's reference for the actual value to use. For example, mail servers typically listen to port 25.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message is to forward to a Syslog server.

6.3.22 Disk Space Monitor

This monitor checks the available and used space on all hard disks in the system. All hard disks present in the system are automatically checked. New disks are automatically detected. One event specifying the maximum size and the used size is generated per disk. The Disk Space Monitor runs continuously based on an interval set in the configuration.

The screenshot shows a configuration window for the DiskSpace Monitor. The status is 'Enabled'. The check interval is set to 1 minute. Under the 'General Values' section, the syslog facility is 'Local 0', the priority is 'Notice', and the tag value is 'MWDisk.SpaceMonitor'. The resource ID field is empty. There is also a 'RuleSet to use' dropdown and a 'Refresh' button.

DiskSpace Monitor Properties

Check Interval

This is the interval in which the service status is checked. This period is specified in [milliseconds](#). The default is 60,000 ms, which is one minute. This should be sufficient for a typical server. If you would like to have the disk space check run less often, you might for example use the value of 3,600,000 for one hour (or a multiple for multiple hours).

For performance reasons, we do not recommend using an interval of less than 30,000 ms.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.


Rule Set to Use

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.23 SNMP Trap Receiver Service

SNMP Trap Receiver allows you to receive SNMP messages. A rough description of a Trap is that it is somewhat like a Syslog message, just over another protocol (SNMP). A trap is generated when the device feels it should do so and it contains the information that the device feels should be transmitted. It also contains some (few) standard items, as the version, community etc. [Click here](#) to know more about the SNMP Trap Receiver Service.

The SNMP Trap Receiver Service runs continuously based on the configuration mentioned below:



The screenshot shows the configuration window for the 'SNMP Trap Receiver' service. The status is 'Enabled'. The configuration fields are: Internet Protocoltype (IPv4), Protocol Type (UDP), Listener Port (162), and SNMP Version (All supported Versions). There is an unchecked checkbox for 'Fully resolve Mibnames (Long Format)'. At the bottom, the 'RuleSet to use' is set to 'Default RuleSet' with a 'Refresh' button.

SNMP Trap Receiver Properties

Internet Protocoltype

Select the desired protocol type. IPv4 and IPv6 are available. The IPv6 protocol needs to be properly installed in order to be used. Note that one Service can only handle IPv4 or IPv6, so if you want to use both protocols, you will need to create two separate services.

Protocol Type

You can select to listen on UDP or TCP protocol for SNMP Traps.

Listener Port

The port the SNMP listener is listening to. If in doubt, leave it at the default of 162, which is the standard port for this.

SNMP Version

Can be used to restrict the SNMP versions. The available values are:

1. All Supported Versions (i.e. SNMP Version 1 and SNMP Version 2c only)
2. SNMP Version 1 only
3. SNMP Version 2c only

Fully Resolve Mibnames (Long Format)

This Option fully resolves the Mibnames like in the Client Mibbrowser Application.


Rule Set to Use

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

Please Note:

Managing incoming Traps works the same way as with a Syslog server for example. Incoming Traps will be forwarded to the corresponding Ruleset and pass by rule after rule. There it can be filtered for general information like the "Community", the "Version" or "Value" for example. Finally it will be processed by an action, which you can select to your needs. The SNMP Agent service will co-exist peacefully next to the Windows SNMP Agent and will not hinder it in its functionality. The Windows SNMP Agent listens to port 161, while MonitorWare Agent and WinSyslog listen to port 162.

For internal processing, the variables of incoming SNMP messages will be added to a new property. Those properties will be named %snmp_var_x% with the x being a number starting with 1. You can use these custom properties for filtering and everywhere you can use or print properties. For example, you can create a "send mail"-action. Here you can specify complete freely how the message will look like. You can use an introductory text and then let it show the error message in some context. This could look like this:



The result will be, that the 5th property of the snmp trap will be inserted into

the message text.

6.3.24 CPU/Memory Monitor

The CPU/Memory monitor has two parts of monitoring, a CPU and a Memory part. Both parts are checked frequently in a specified check interval. By default this interval is configured to 60 seconds so a check is done after every 1 minute. Both CPU and memory part can be disabled and enabled. The service becomes useless if you disable both parts.

Note for Windows NT4 users: This service uses the WMI (Windows Management Instrumentation) to query CPU and memory utilization. By default WMI is not installed on NT4 but it can be obtained and installed [from here](#).

Status: Enabled | Comments | Save | Reset | ?

Check Interval (ms) | 1 Minute

CPU | Total Memory | Physical Memory | Virtual Memory

Enable CPU Check

CPU usage alarm level | [Slider]

Occurrences until CPU alarm | 10

Also report if CPULevel is below the alarm level

For Multiprocessor Systems only: Report for each CPU

General Values

Syslog Facility | Local 0

Syslog Priority | Notice

Syslog Tag Value | MWCPUMonitor

Resource ID |

RuleSet to use | [Dropdown] | Refresh

CPU/Memory Monitor Properties

You can see a series of tabs at the top of screen shot above, namely CPU, Total Memory, Physical Memory and Virtual Memory.

CPU

When you click CPU tab, you are shown the options as shown in the screen shot above and explained below:

Enable CPU Check

If this option is checked then it allows you to monitor the CPU.

CPU usage Alarm level

If the CPU usage reaches this level, an event is generated internally. Depending on how the occurrences value is configured, an Information Unit is generated.

Note: this value is in terms of percentage.

Occurrences until alarm is raised

Defines how often the CPU usage level has to occur in a row until an event is raised.

Also report if CPU level is below the Alarm level


If this option is checked it generates an Information Unit during each run whether it is below the alarm level or above. This option is useful for statistics and debugging.

For Multiprocessor Systems only: Report for each CPU

This option can only be used if you have a Multiprocessor System. When this option is checked the CPU usage is analyzed for each CPU and for the whole amount of CPU usage.

Total Memory

When you click Total Memory tab, you are shown the options as shown in the screen shot below:



The screenshot shows a configuration window with four tabs: CPU, Total Memory, Physical Memory, and Virtual Memory. The 'Total Memory' tab is selected. The window contains the following options:

- Enable Memory Check
- Memory usage alarm level: A slider control with a vertical bar in the middle.
- Occurrences until Memory: A text input field containing the value '10'.
- Also report if Memory Level is below the alarm level

Total Memory Properties

Enable Memory Check

If this option is checked then it allows you to monitor the memory.

Memory usage Alarm level

Defines the alarm level for the memory usage. Note that virtual and physical memory are calculated together.

Note that this value is in terms of percentage.

Occurrences until alarm is raised

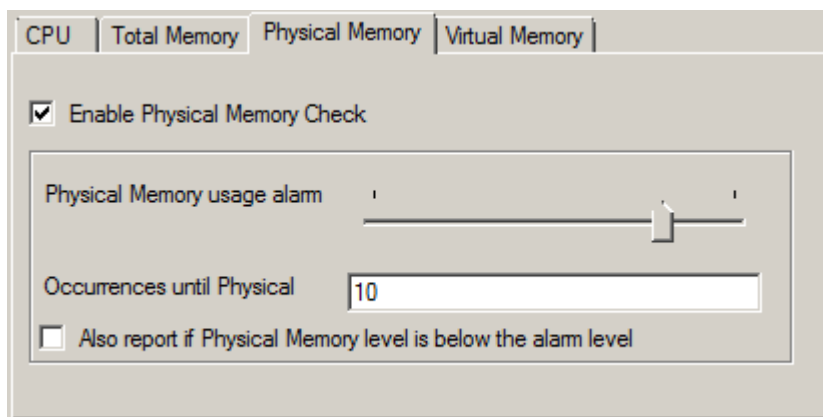
Defines how often the memory usage has to be over the memory usage alarm level in a row.

Also report if Memory level is below the Alarm level

This also generates an event if the memory usage is below the alarm level. A useful option for testing and debugging.

Physical Memory

When you click Physical Memory tab, you are shown the options as shown in the screen shot below:



Physical Memory Properties

Enable Physical Memory Check

If this option is checked then it allows you to monitor the physical memory.

Physical Memory usage Alarm level

Defines the alarm level for the physical memory usage.

Note that this value is in terms of percentage.

Occurrences until alarm is raised

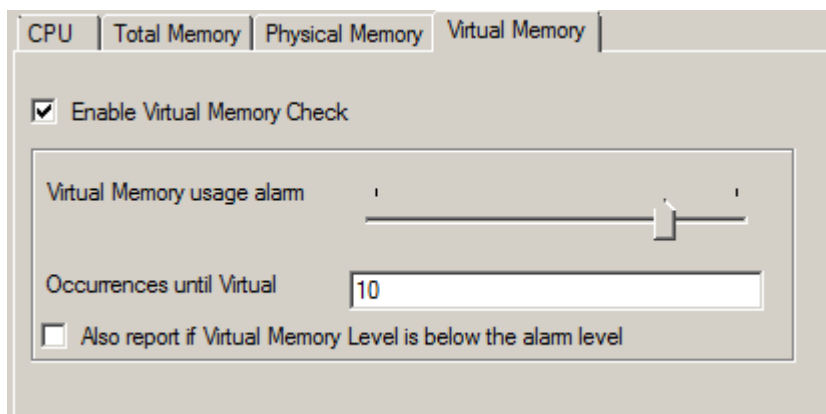
Defines how often the physical memory usage has to be over the physical memory usage alarm level in a row.

Also report if Physical Memory level is below the Alarm level

This also generates an event if the physical memory usage is below the alarm level. A useful option for testing and debugging.

Virtual Memory

When you click Virtual Memory tab, you are shown the options as shown in the screen shot below:



Virtual Memory Properties

Enable Virtual Memory Check

If this option is checked then it allows you to monitor the virtual memory.

Virtual Memory usage Alarm level

Defines the alarm level for the virtual memory usage.

Note that this value is in terms of percentage.

Occurrences until alarm is raised

Defines how often the virtual memory usage has to be over the virtual memory usage alarm level in a row.

Also report if Virtual Memory level is below the Alarm level

This also generates an event if the virtual memory usage is below the alarm level. A useful option for testing and debugging.

Check Interval

The CPU/Memory Monitor runs periodically. This specifies, how often it should run. Please note that the CPU / Memory Monitor waits the configured amount of time after the current run is finished. The time is the delay in [milliseconds](#).

General Values

In General Values group box, you can see different fields discussed below:

Syslog Facility

The [Syslog facility](#) to be assigned to events created by this service. Most useful, if the message shall be forwarded to a Syslog server.

Syslog Priority

The Syslog priority to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.

Resource ID

The [Resource ID](#) to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.

Syslog Tag Value

The Syslog tag value to be assigned to events created by this service. Most useful if the message shall be forwarded to a Syslog server.


Default Ruleset Name

Name of the rule set to be used for this service. The Rule Set name must be a valid Rule Set.

6.3.25 MonitorWare Echo Reply

The MonitorWare Echo Reply service is a somewhat unusual service. It by itself does **not** generate any events.

It is the passive counterpart of the [MonitorWare Echo Request](#) service. These together are used to detect failing agents. In this services configuration, only the listening port can be specified. This port must be the same to which the Echo Request service tries to connect to.



MWEchoReplyService

Listener Port

Specify the listener port here.

IP Address

The MonitorWare Echo Reply service can be bound to a specific IP Address. This feature is useful for multihome environments where you want to run different Syslog Servers on different IP Addresses. Please note that the default IP Address 0.0.0.0 means ANY IP Address. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address

6.3.26 MonitorWare Echo Request

The MonitorWare Echo Request service is used to check the availability of MonitorWare agents (either MonitorWare Agent, WinSyslog or EventReporter). It works in conjunction with the MonitorWare Echo Reply service, which needs to be running on the agents which are to be monitored.

Echo Request service tries to connect to the reply service on remote agents. If it can connect, it checks if the remote agent is alive. If either the connect fails or the remote response is not satisfactory, an event is generated (which could be used for alerting as well as corrective action). Optionally, an event can also be generated when the remote agent can be contacted successfully.

A single MonitorWare Echo Request service can check an unlimited number of remote agents. Please note, however, that all checks are done in sequence. So with a large number of systems to be checked, there may be a longer delay between the checks than you expect. This is especially the case over slow network links (like found in wide

area networks). If this is not acceptable, multiple Echo Request services can be configured. They then run independent of each other.

Enable: MonitorWare Echo Request

Settings are saved.

Save Reset

Save & Close ?

Check Interval (ms) 1 minute

Also generate an event if echo reply was successful

Insert Delete ↑ ↓ IP / Hostname

Port

IP	Port

*General Values

Syslog Facility LOCAL0 (16)

Syslog Priority INFO (6)

Resource ID

Syslog Tag Value MWEchoRequest

Rule Set to Use Defaults Refresh

Check Interval

The Echo Request service runs periodically. This specifies how often it should run. Please note that the Echo Request service waits for the configured amount of time after the current run is finished. The time is the delay in [milliseconds](#).

Also generate an event if echo reply was successful

If checked, an event is to be created each time the probe runs, even when it is successful. If unchecked, events are only created when the remote system fails.

IP / Port Pairs

This table contains the systems that are to be checked in each interval. Once in each run, each remote system is checked. The checks are carried out in the exact same order that the systems appear in the table - from top to bottom. Use "insert" to create a new entry (you can edit it after it has been inserted), "delete" to delete the current entry and the arrow buttons to change the order of the entries.

Syslog Facility

The [Syslog facility](#) to be assigned to events created by the service. Most useful if the message is to be forwarded to a Syslog daemon.

Syslog Priority

The Syslog priority to be assigned to events created by the service. Most useful if the message is to be forwarded to a Syslog daemon.

Resource ID

The [Resource ID](#) to be assigned to events created by the service. Most useful if the message is to be forwarded to a Syslog daemon.

Syslog Tag Value

The Syslog tag value to be assigned to events created by the service. Most useful if the message is to be forwarded to a Syslog daemon.

Rule Set to Use

Name of the rule set to be used for this service. The rule set name must be a valid rule set.

6.4 Filter Conditions

6.4.1 Filter Conditions

Filter conditions specify **when** to apply a rule. If the filter condition evaluates to true, the rule containing those conditions is treated as matching and the actions specified in that rule are carried out.

Filter conditions can be as complex as needed. Full support for Boolean operations and nesting of conditions is supported.

By default, the filter condition is empty, respective tree contains only a single "AND" at the top level. This is to facilitate adding filters (the top level-node is typically "AND" and thus provided by default). A filter condition containing only the "AND" always evaluates as true. A sample screenshot can be found below:

Status: Enabled | Comments | Save | Reset | Name:

Global Conditions

Filter Conditions

AND

Tools

Add Filter >

Add Operations

AND OR

NOT XOR

Special Operations, useful for debugging

TRUE FALSE

Change Operator

↑ ↓ Delete

Clone Filter

Global Conditions

Threat not found filters as TRUE

Fire only if Event occurs times within

Minimum Wait Time seconds.

Global Conditions relative to this property

[Learn about Filters](#)

Filter Conditions - Display form

The default filter condition means that the actions associated with the rule are to be carried out for every information unit received. It is often used for actions that should be broadly taken, for example to write all incoming information units to a database or text file.

On the other hand, there are actions that should only be executed under very special conditions. They may even require a complex filter condition including multiple levels of Boolean operations. Below is a sample of such a condition:

Filter Conditions - Complex Filter

This filter condition is part of an intrusion detection rule set. Here, Windows file system auditing is used to detect a potentially successful intrusion via Internet Information Server (IIS). This is done by enabling auditing on all executable files. Internet Information Server accesses them under the IUSR_<machinename> account, which in our sample is "P15111116\IUSR_ROOTSERVER". If that user runs any unexpected executables, chances are good that someone was able to intrude the machine via IIS. Please note that Perl and PHP scripts need to run the Perl and PHP engine. This is reflected by specifically checking, if perl.exe and php.exe is executed – and if so, no alarm is triggered.

Here is how the above sample works: first, the message contents are checked if it contains either the full path name to perl.exe or php.exe. This is done in the "OR" branch at the bottom. We now need to keep in mind that when a filter condition evaluates to "true", the actions are executed. In case of perl.exe and php.exe, this is just the opposite of what we want. We need it to be executed, when other files are executed. Consequently, we negate (Boolean "NOT") the result of the OR. The end result of the "NOT" operation is then combined via a "AND" with some other properties describing the event we need.

First, we check if the specific event really occurred. For this, we need to make sure we deal with an Event Log Monitor information unit. Then, these information units are identified by the event source as well as the Event ID. We also check for the Event User to identify only IIS generated requests. Lastly, we check if the message contains the string ".exe".

In order to avoid too frequent alerts, we also have specified a minimum wait time of

60 seconds. Therefore, the filter condition evaluates as "true" at most every 60 seconds, even if all other conditions are true.


Note: If you want to know more about [complex filter conditions](#) you can click on the "Learn about Filters" link.

String comparison in Filter Conditions are "Case Sensitive"! For example, if the Source System name is "ws01" and you had written "WS01" while applying the filter, then this filter condition would **"NEVER"** evaluate to True! Please double check before proceeding further!

If you are not still sure about what to do, you can drop a word about your requirements to support@adiscon.com, and we look into it!

6.4.2 Global Conditions

Global Conditions apply to the rule as whole. They are automatically combined with a logical "AND" with the conditions in the filter tree.



Filter Form - Global Conditions

Treat not found Filters as TRUE

If a property queried in a filter condition is not present in the event, the respective condition normally returns "FALSE". However, there might be situations where you would prefer if the rule engine would evaluate this to "TRUE" instead. With this option, you can select the intended behaviour. If you check it, conditions with properties not found in the event evaluates to "TRUE".

Fire only if Event occurs

This is kind of the opposite of the "Minimum WaitTime". Here, multiple events must come in before a rule fires. For example, this time we use a ping probe. Ping is not a very reliable protocol, so a single ping might be lost. Thus, it may not be the best idea to restart some processes just because a single ping failed. It would be much better to wait for repetitive pings to fail before doing so.

Exactly this is why the "Fire only if Event Occurs" filter condition is made for. It waits until a configured amount of the same events occurs within a period. Only if the count is reached, the filter condition matches and the rule can fire.

Note: If you used previous versions of the product, you might remember a filter called "Occurrences". This has just been renamed.

Minimum Wait Time

This filter condition can be used to prevent rules from firing too often. For example, a rule might be created to check the status of a port probe event. The port probe probes an [SMTP](#) server. If the event is fired and the rule detects it, it spawns a process that tries to restart the service. This process takes some time. Maybe the SMTP gateway need some more time to fully start up so that the port probe might fail again while the problem is already taken care of. The port probe as such generates an additional event.

Setting a minimum wait time prevents this second port probe event to fire again if it is – let's say – within 5 minutes from the original one. In this case, the minimum wait time is not yet reached and as such, the rule is not match. If, however, the same event is generated 5 hours later (with the mail gateway failing again), the rule once again fired and corrective action taken.

Global Conditions relative to this property

This feature enables you to control the Global Conditions based on a property. For example take the source of a message as property. In this case, the Minimum WaitTime for example would be applied individual on each message source.

6.4.3 Operators

In general, operators describes how filter conditions are linked together. The following operators can be used.

AND

All filters placed below must be true. Only then AND returns true.

OR

Even if one of the filter placed below OR is true, OR returns true.

NOT

Only one Filter can be placed below NOT operator, and if the filter evaluation is true, NOT returns false.

XOR

Only one of the two filters are possible in the XOR Operator.

TRUE

Useful for debugging, just returns TRUE.

FALSE

Useful for debugging as well, returns FALSE.

6.4.4 Filters

Filters can be added under each Operation node. There are a few common filters which can be used for all services, and there are special filters which only apply if a special kind of Information Unit is evaluated.

What happens with Filters that are not available in an "Information Unit"?

Every filter that is not found in an Information Unit is ignored in the filtering process. If you want to create filters specialized for types of Information Units, always make sure to add an "Information Unit Type" filter.

An example, you have one ruleset, rule and action. In the filters you have one EventID filter. Then you have two services, one Eventlog Monitor and the other is Heartbeat monitor both pointing to this ruleset. The Information Units from the Eventlog Monitor would be filtered correctly, but those from the Heartbeat monitor would not be filtered as they don't have an EventID property. The EventID filter would be ignored and the actions would be executed every time.

Note, if a filter is used that does not apply to the evaluated Info Unit, it will be just ignored. This gives you the possibility to build one filter set for several types of Information Units.

There are different types of filters, and so there are different ways in which you can compare them to a value. The following Types exist:

String

Can be compared to another String with "=", "Not =" and "Range Match".

Number

Can be compared with another number with "=", "Not =", "<" and ">"

Boolean

Can be compared to either TRUE or FALSE with "=" and "Not ="

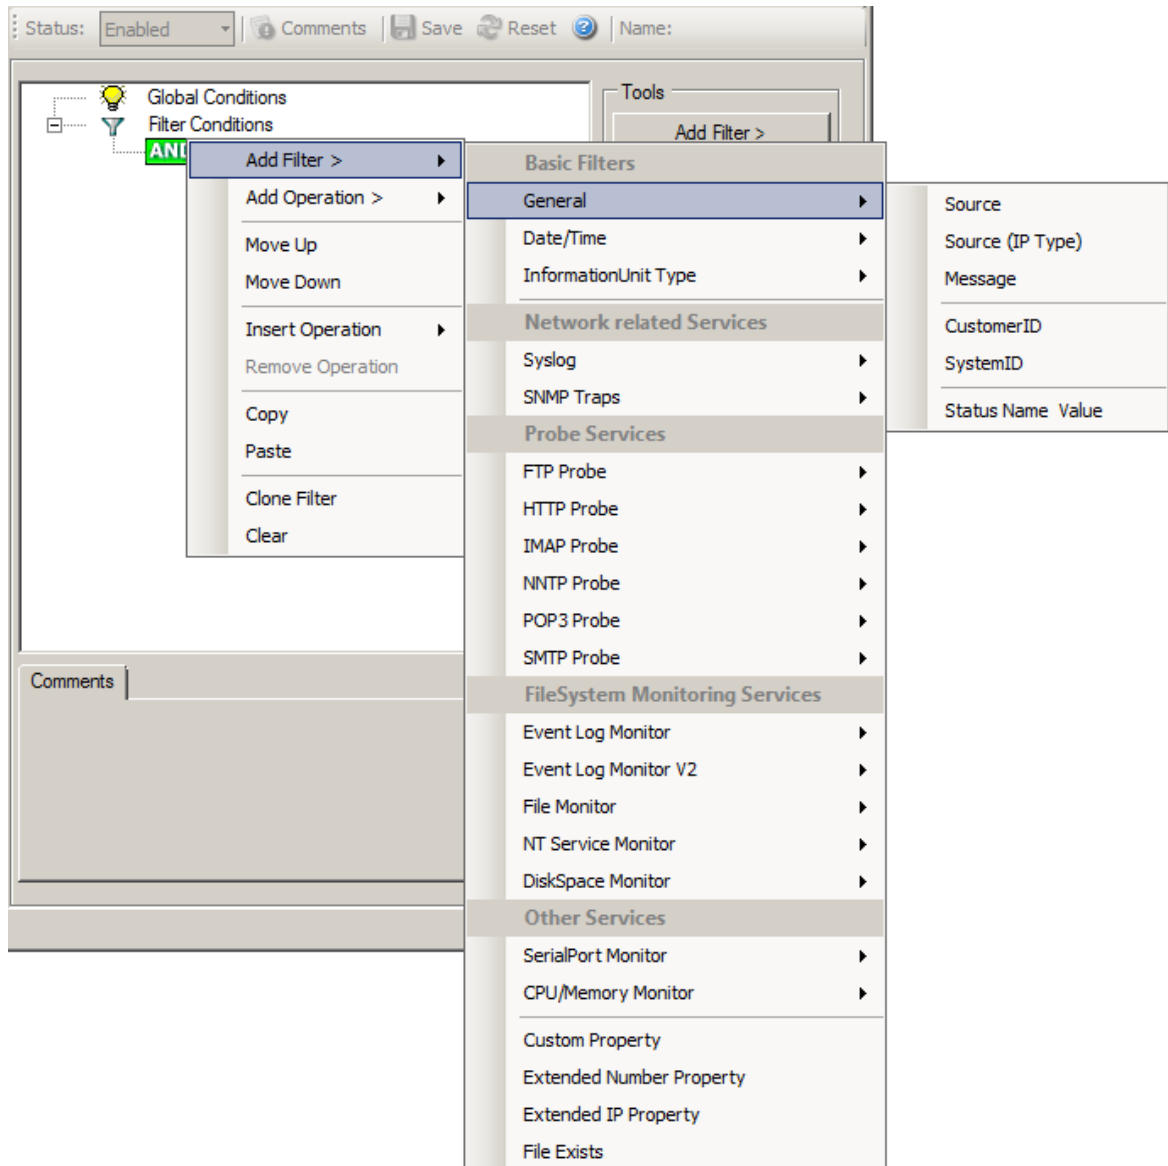
Time

Can be compared with another time but only with "="

The list of possible filters, which can be evaluated is described in the upcoming sections.

6.4.5 General

These are non-event log specific settings.



Filter Conditions - General

Source System

This filter condition checks the system that generated the information unit. For

example, in case of the Syslog server, this is the Syslog device sending a Syslog message.

This filter is of type string and should contain the source system name or IP address.

Source System (IP)

The IP Filter can basically work on any property, but we recommend to only use it on the %source% property, as we usually can be sure that this contains a valid IP Address or hostname. The IP Filter can filter against hostnames and IP Addresses, hostnames are automatically resolved using the internal DNSCache (for obvious performance reasons).

This filter is of type string and should contain the source system name or IP address.

Please see the description for "Extended IP Property" for more information on how to use this property.

Message Content

The message content filter condition is very powerful. It evaluates to true if the specified content is found anywhere within the message. As there is implicit wildcarding, there is no need for extra wildcards to be specified.

The content search can be limited to a region within the message. To do so, select a starting and ending position within the string by choosing the "**contains within range**" compare operation. This can be done by specifying the start range and end range into the respective boxes.

Please note that you can enter the character position you desire in these fields. The default "Start Range" and "End Range" are set to 0.

If you would like to search for a string just between positions 10 and 50, specify these values as start and end values, respectively. Similarly if you want to receive all logs from 192.168.0.1 then set this as:

```
Property value = 192.168.0.0  
Range Start = 0  
Range End = 10
```

Which means 10 characters starting at zero ("192.168.0."). Please note that the final DOT must be included. If you just used range "9", then 192.168.010 would also match.

This filter is of type string.

CustomerID

CustomerID is of type integer provided for customer ease. For example if someone monitors his customer's server, he can put in different CustomerIDs into each of the

agents. Let us say someone monitors servers A and B. A has 5 servers all of them with CustomerID = 1 and B has 2 servers all of them with CustomerID = 2. Both A and B happen to have a server named "SERVER". Together with the customerID, these machines are now uniquely identifiable. This is user configurable.

CustomerID (Type=Number).

SystemID

SystemID is of type integer to be used by our customer. In addition, it is user configurable.


SystemID (Type=Number).

Status Name and Value

These filter type corresponds to "[Set Status](#)" Action. Status Name and Value (Type=String)

6.4.6 Date/Time

This filter condition is used to check the time frame and / or day of week in which an event occurred.



Filter Conditions - Date / Time

Time

This filter condition is used to check the period in which an event occurred. For example, a Syslog message from a Cisco router saying that it dialed up is normal if it occurs during office hours. If it occurs at night, so, it is an alerting signal and an administrator might receive notification of this event (while he might otherwise decide to discard it). This can be done with the time setting.

You can also set the timezone setting (DefaultTimemode, UTC or Localtime) for the TimeMode's (DeviceReportedTime/ReceivedTime).


Weekdays

This is closely equivalent to the time filter condition, except that it is applied on a per-day basis. So it can be used to detect for example events occurring on weekends and act differently on them. The following filters are available:

1. Run on Monday (Type=Boolean)
2. Run on Tuesday (Type=Boolean)
3. Run on Wednesday (Type=Boolean)
4. Run on Thursday (Type=Boolean)
5. Run on Friday (Type=Boolean)
6. Run on Saturday (Type=Boolean)
7. Run on Sunday (Type=Boolean)

6.4.7 InformationUnit Type

Select the specific information if a rule should just be processed for some information unit types. This is especially useful if a specific type needs non-standard processing. There is one pre-defined filter for each possible InformationUnit Type available (shown below).



Filter Conditions - InformationUnit Type

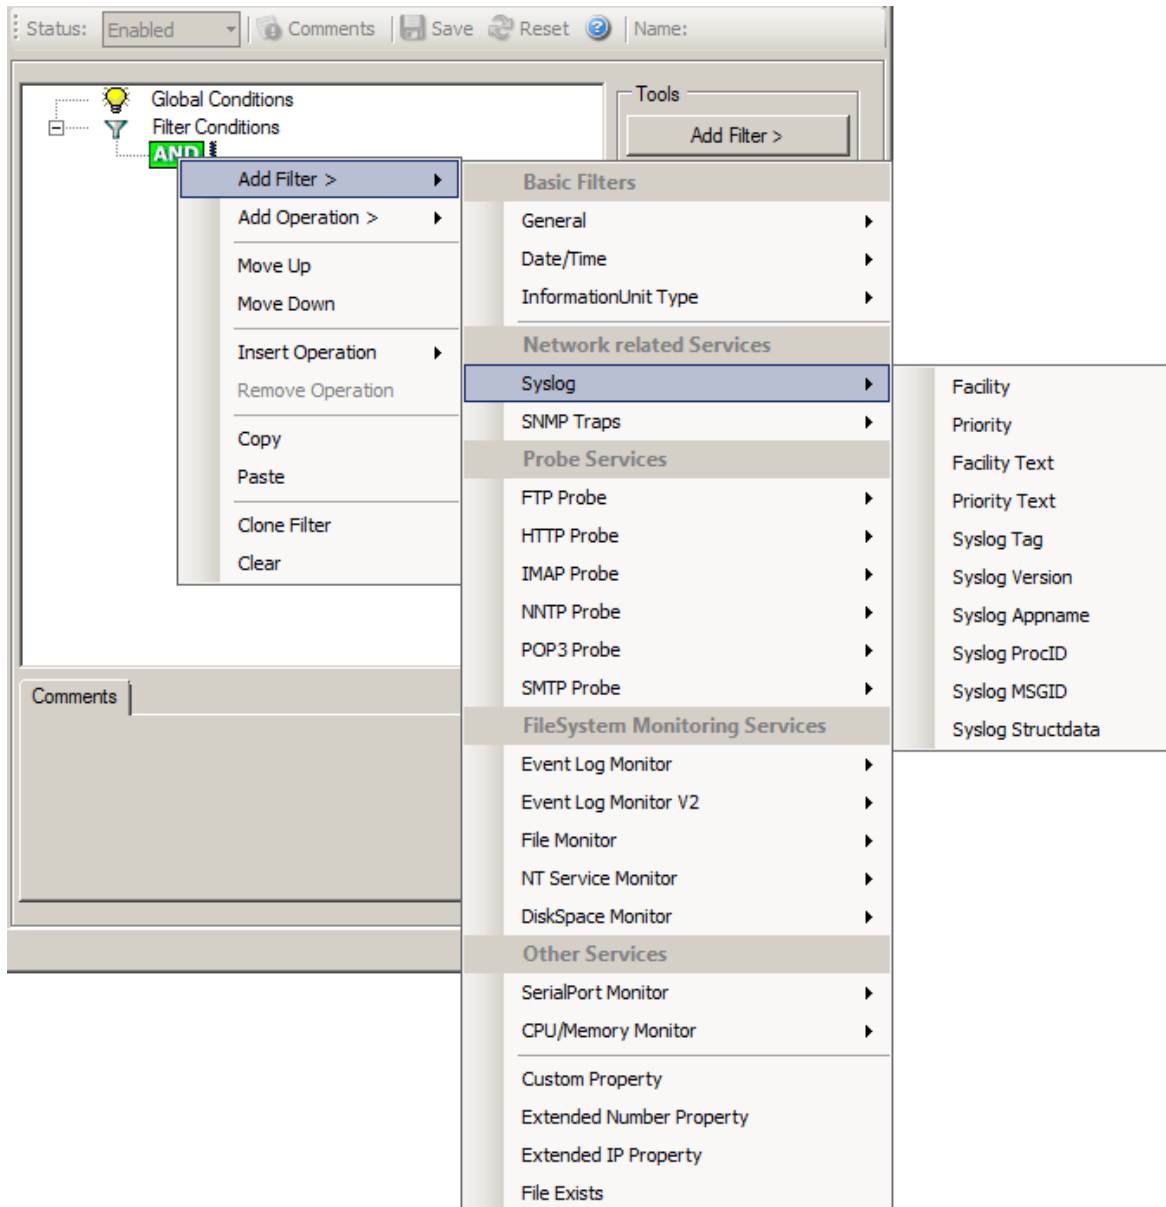
The following filters are available:

1. Syslog (Type=Boolean)
2. Heartbeat (Type=Boolean)
3. SNMP Traps (Type=Boolean)
4. Event Log Monitor (Type=Boolean)
5. File Monitor (Type=Boolean)
6. Ping Probe (Type=Boolean)
7. Port Probe (Type=Boolean)
8. NT Services Monitor (Type=Boolean)
9. Disk Space Monitor (Type=Boolean)
10. Database Monitor (Type=Boolean)
11. Serial Port Monitor (Type=Boolean)
12. CPU/Memory Monitor (Type=Boolean)

13. FTP Probe (Type=Boolean)
14. HTTP Probe (Type=Boolean)
15. IMAP Probe (Type=Boolean)
16. NNTP Probe (Type=Boolean)
17. POP3 Probe (Type=Boolean)
18. SMTP Probe (Type=Boolean)

6.4.8 Syslog

Syslog related filters are grouped here. Please keep in mind that every Information Unit has assigned a Syslog priority and facility and thus these filters can be used with all Information Units.



Filter Conditions - Syslog

Syslog Facility

The information unit must have the specified [Syslog facility](#) value. For Syslog type information units, it is the actual Syslog priority code, for all others it is a value mapped on a best effort basis.

This filter is of type number.

Syslog Priority

The information unit must have the specified Syslog priority value. For Syslog type information units, it is the actual Syslog priority code, for all others it is a value mapped on a best effort basis.

The first list box allows to set a matching mode. The operations "less than" (<), "greater than" (>) and "equal" (=) can be selected. The match is made depending on these operations, so a "less than" operation means that all priorities below the specified priority match. Please note that the specified priority is **not** a match. If you would like to include it, be sure to specify the next higher one.

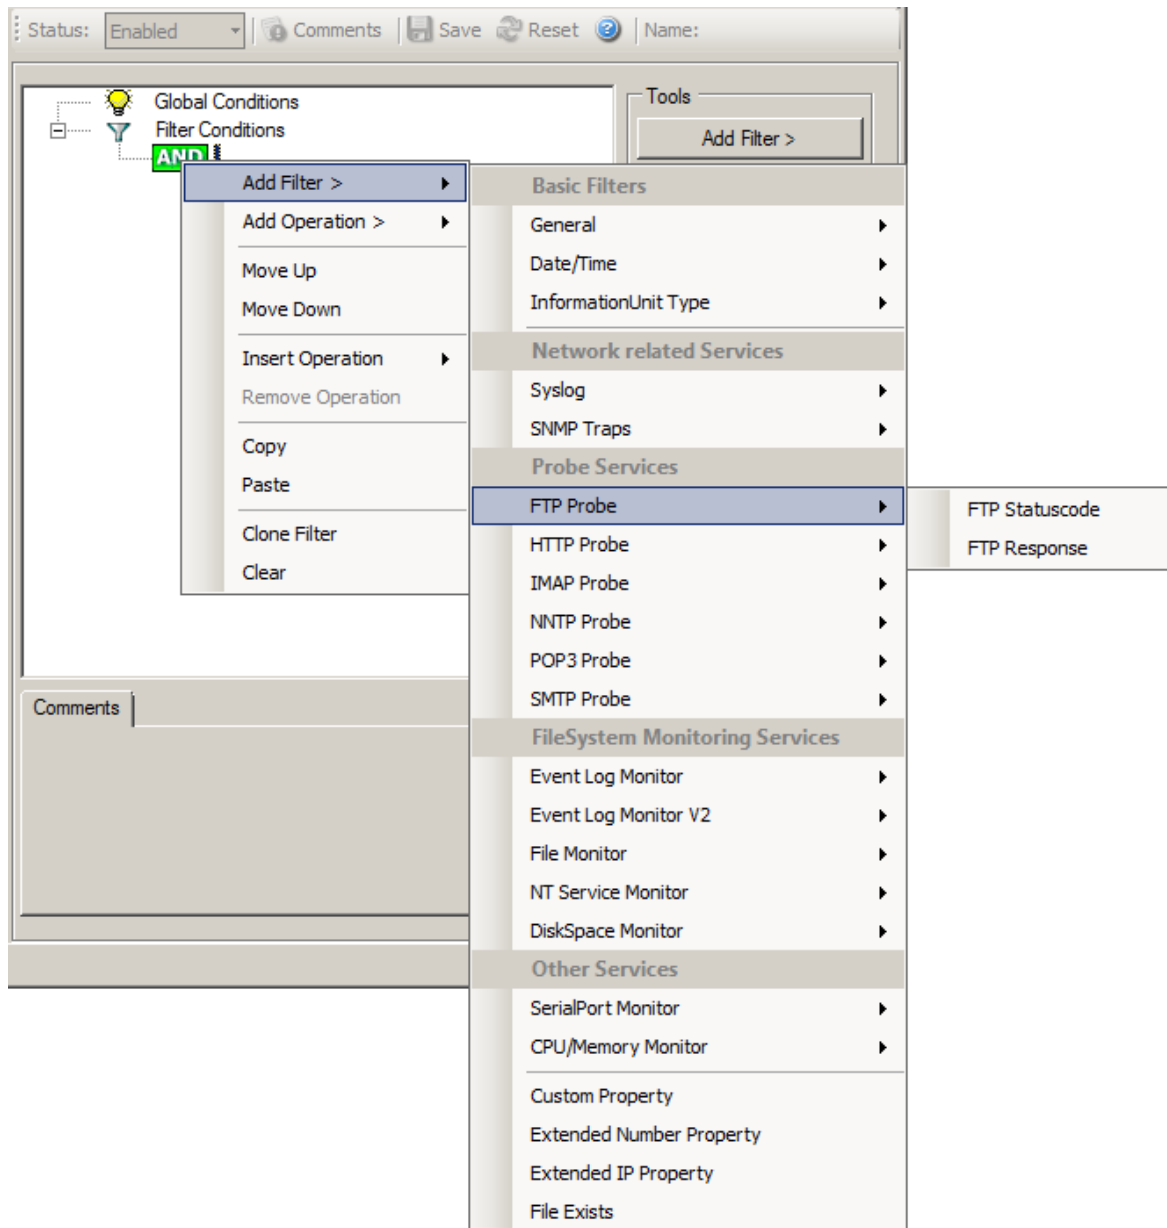
This filter is of type number.

Syslog Tag

This filter is of type string.

6.4.9 FTP

FTP related filters are grouped here.



Filter Conditions - FTP

FTP Statuscode

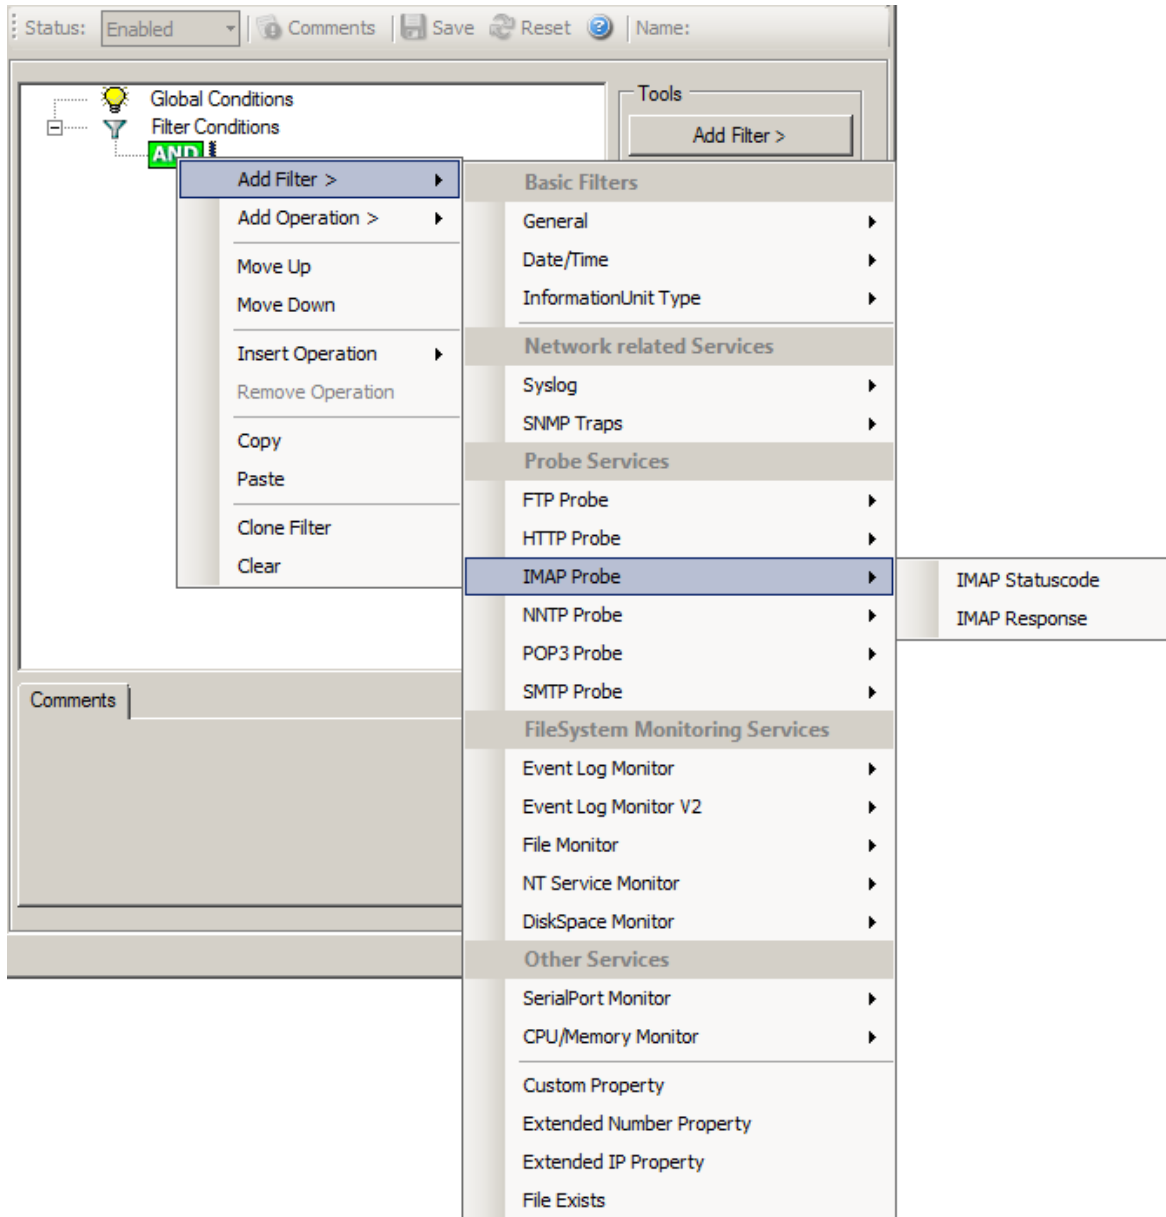
It contains the FTP success or error code. This filter is of type number.

FTP Response

It contains the FTP response.

6.4.10 IMAP

IMAP related filters are grouped here.



Filter Conditions - IMAP

IMAP Statuscode

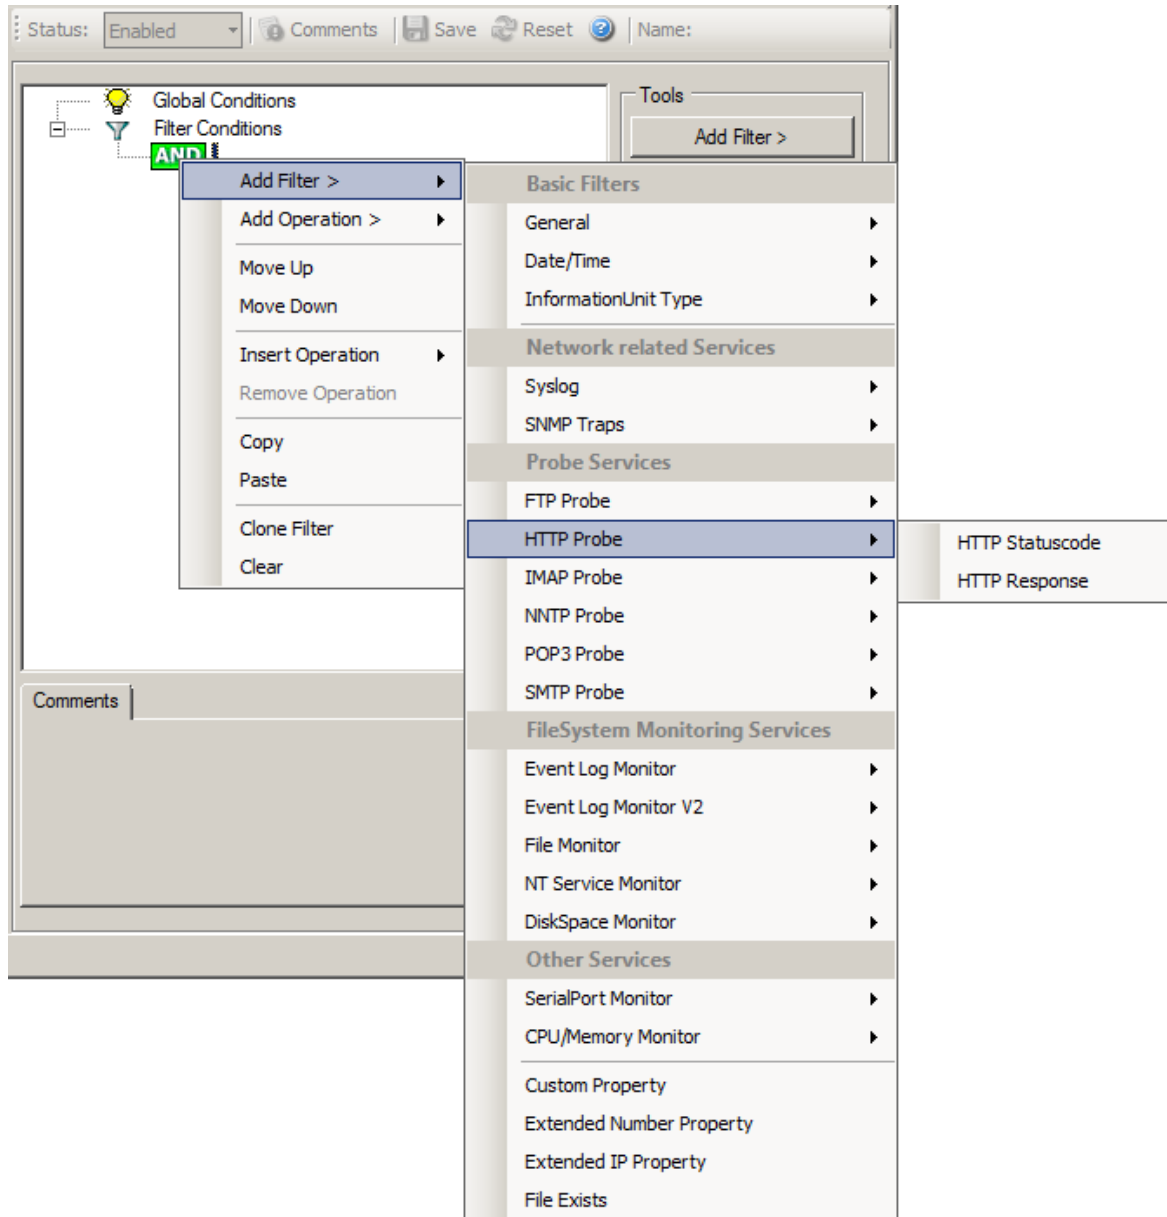
It contains the IMAP success or error code. This filter is of type number.

IMAP Response

It contains the FTP response.

6.4.11 HTTP

HTTP related filters are grouped here.



Filter Conditions - HTTP

HTTP Statuscode

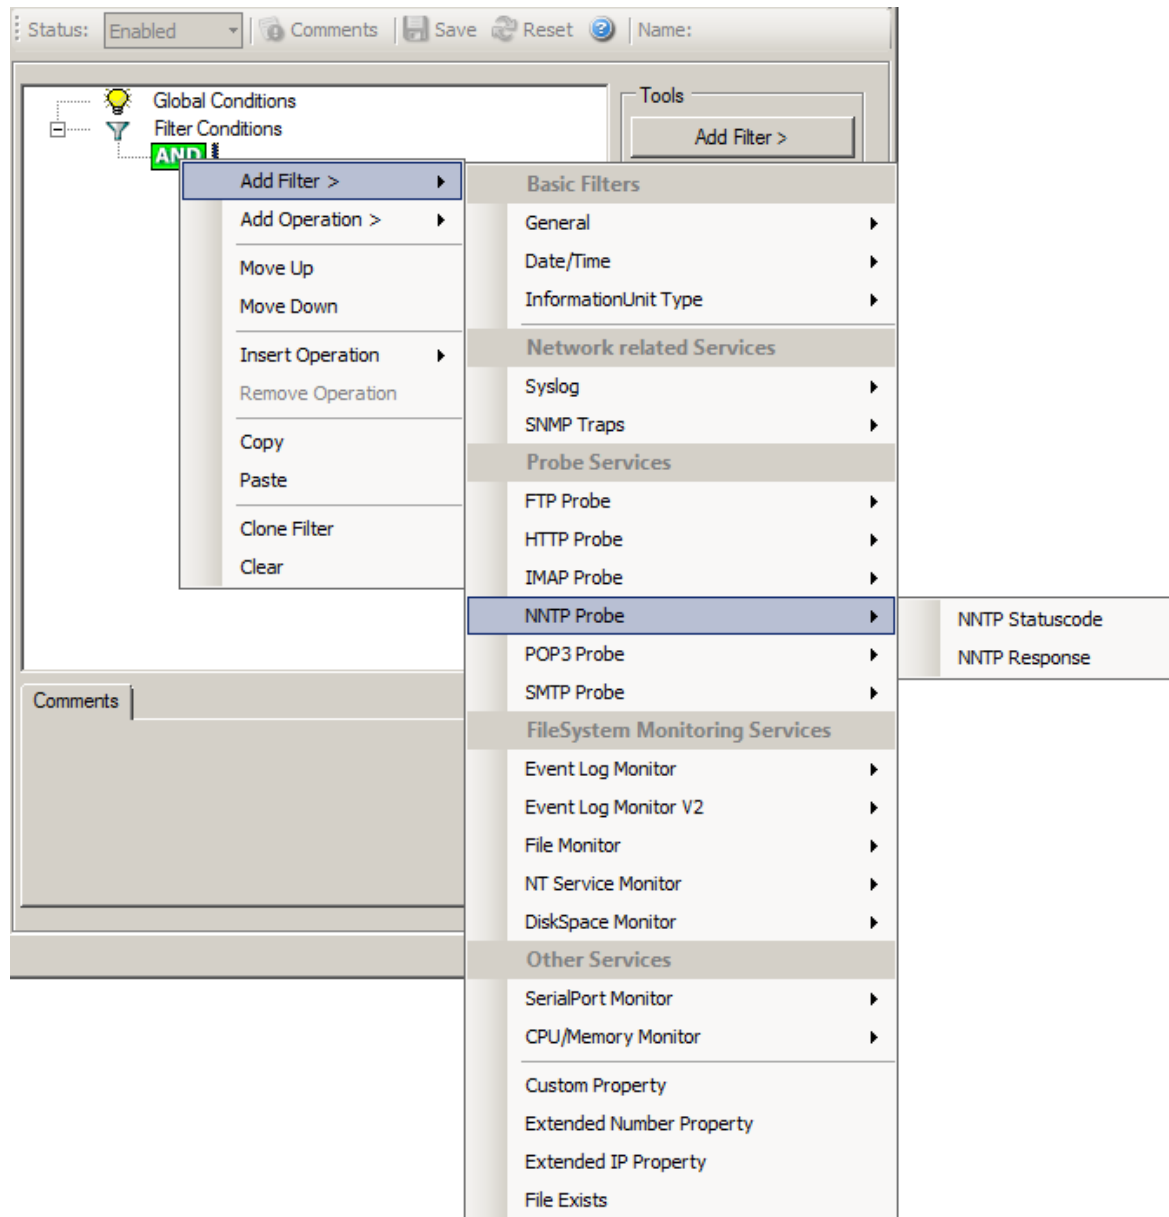
It contains the HTTP success or error code. This filter is of type number.

HTTP Response

It contains the FTP response.

6.4.12 NNTP

NNTP related filters are grouped here.



Filter Conditions - NNTP

NNTP Statuscode


It contains the NNTP success or error code. This filter is of type integer.

NNTP Response

It contains the NNTP response.

6.4.13 POP3

POP3 related filters are grouped here.



Filter Conditions - POP3

POP3 Statuscode

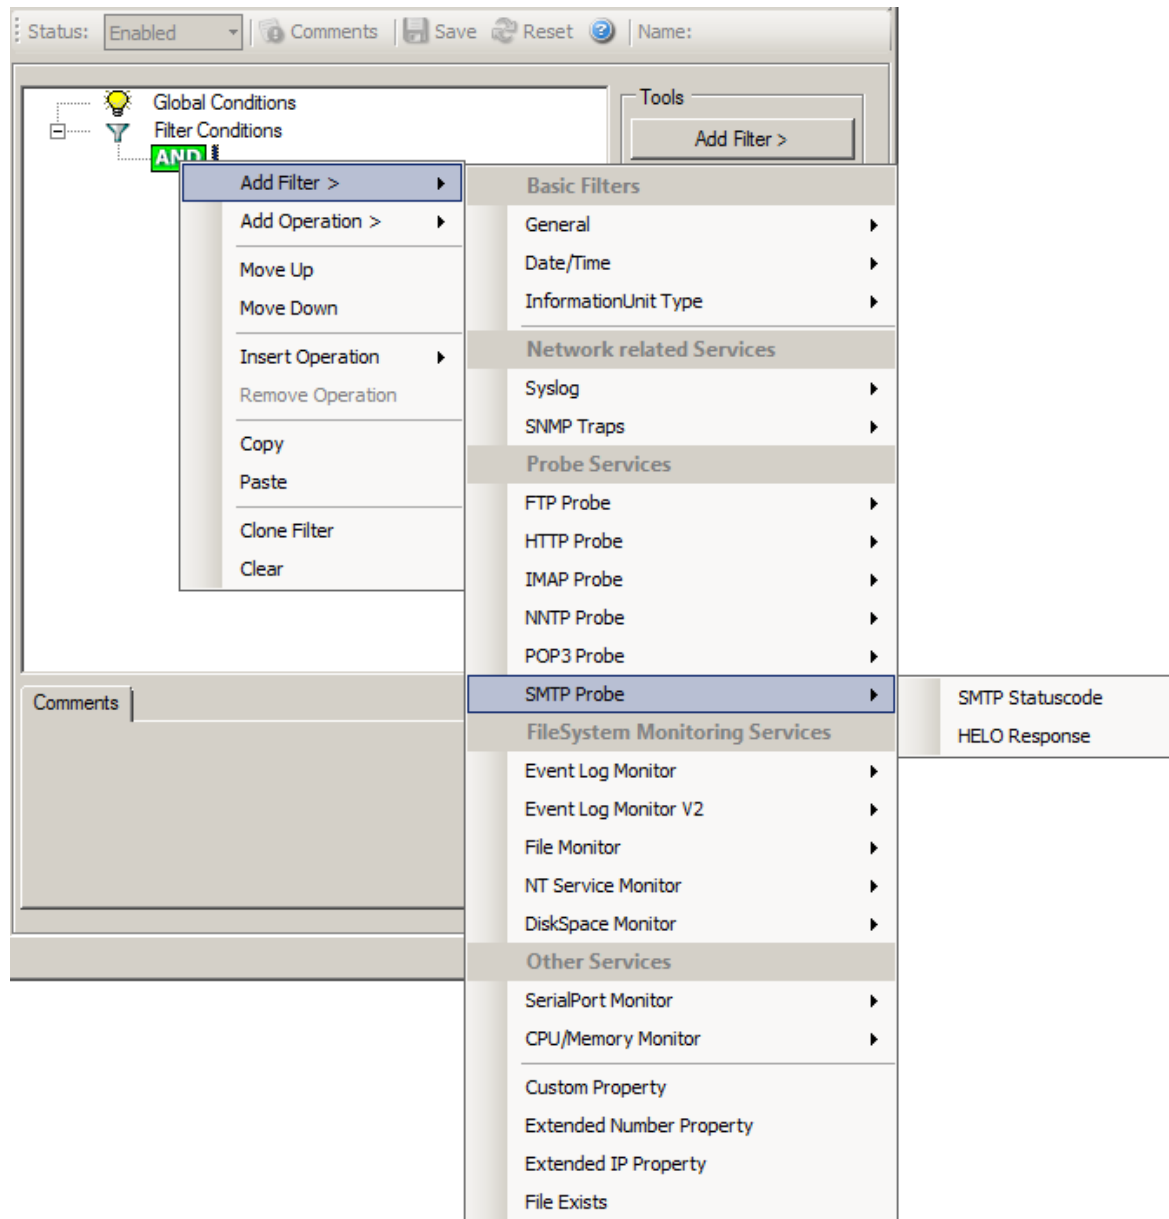
It contains the POP3 success or error code. This filter is of type number.

POP3 Response

It contains the POP3 response.

6.4.14 SMTP

SMTP related filters are grouped here.



Filter Conditions - SMTP

SMTP Statuscode

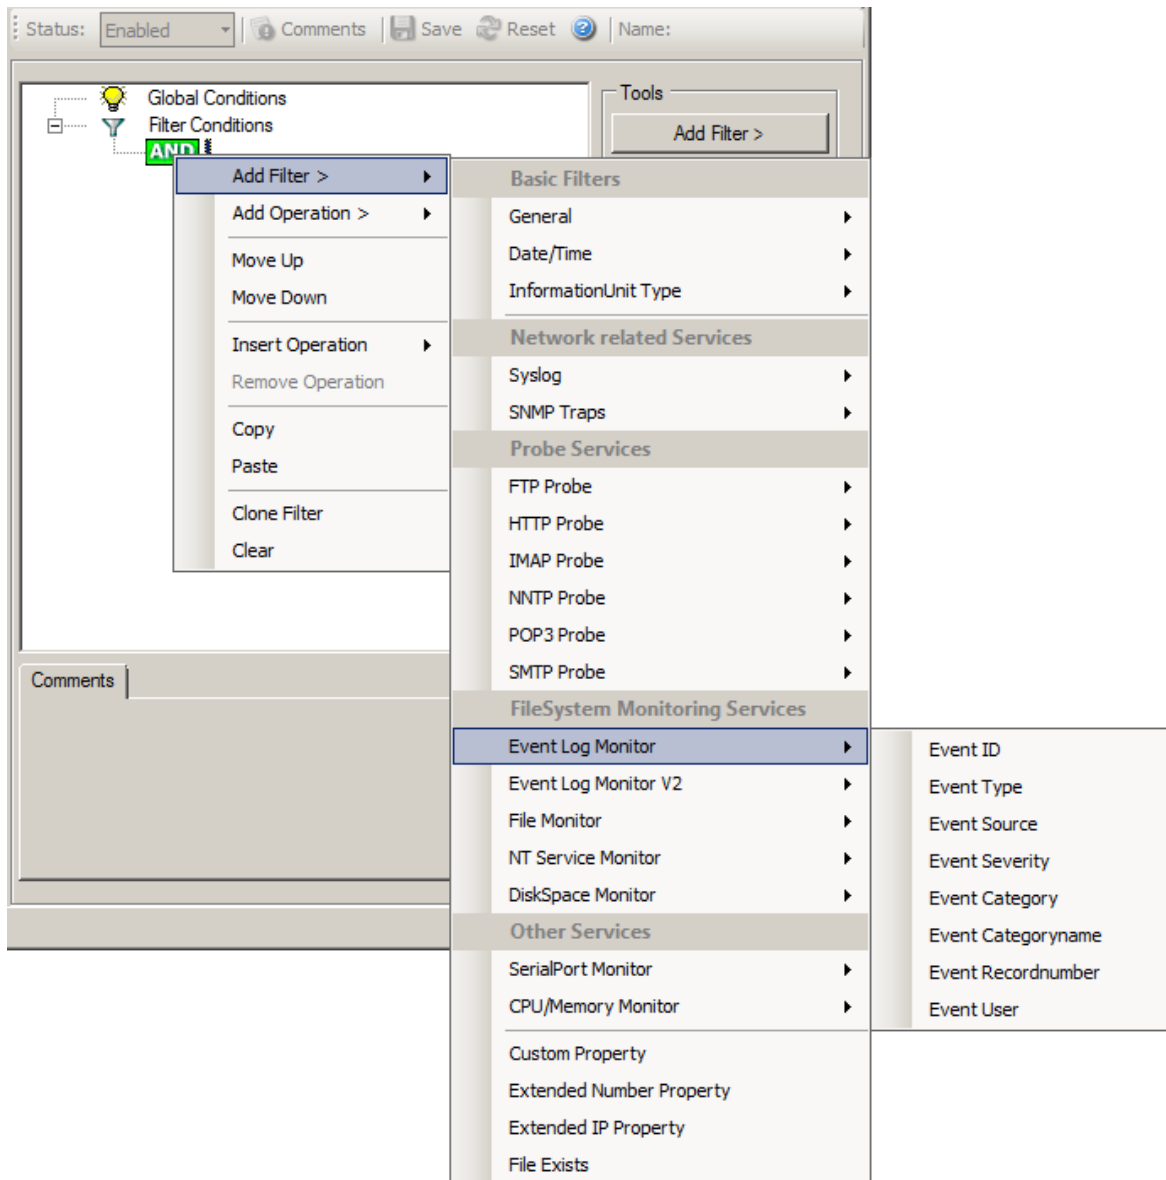
It contains the SMTP success or error code. This filter is of type number.

HELO Response

It contains the SMTP response.

6.4.15 Event Log Monitor

Event log monitor specific filters are grouped here.



Filter Conditions - Event Log Monitor

Event ID

This is the event log ID as specified in the NT event log. If enabled, the event must have the configured event ID or the rule will not match. This is an integer value.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type number.

Event Type

This is the event log type as specified in the NT event log. If enabled, the event must have the configured event type or the rule will not match. The supported values can be selected from the list box.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type string.

Event Source

This is the event log source as specified in the NT event log. If enabled, the event must have the configured event source or the rule will not match. This is a string value. There must be an exact match. Please note that this value is case-sensitive.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type string.

Event Severity

This is the event log severity as specified in the NT event log. If enabled, the event must have the configured severity or the rule will not match. The supported values can be selected from the list box.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type number.

Event Category

This is the event log category as specified in the NT event log. If enabled, the event must have the configured event category or the rule will not match.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual

value.

This filter is of type number.

Event Categoryname

This value contains the Category value as string if it can be resolved. Otherwise it contains the category number.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type string.

Event Recordnumber

This value contains the internal event record number. Please note that if the event log has been truncated before, it may not start with 0 or 1 but a higher number.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type number.

Event User


This is the event log user as specified in the NT event log. If enabled, the event must have the configured event user or the rule will not match. Since it's a string value there must be an exact match. Please note that this value is case-sensitive.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type string.

6.4.16 Event Log Monitor V2

Event log monitor V2 specific filters are grouped here.



Filter Conditions - Event Log Monitor V2

Event Channel

The channel property for event log entries, for classic Event logs they match the % nteventlogtype% property, for new event logs, they match the "Event Channel". If enabled, the event must have the configured event type or the rule will not match. There must be an exact match. Please note that this value is case-sensitive.

This filter condition should only be used with event log information units. If used with

others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type string.

Event Raw Source

This contains the full internal name of the event source for new event logs, for classic event logs it contains the same value as in %sourceproc%. If enabled, the event must have the configured event source or the rule will not match. This is a string value. There must be an exact match. Please note that this value is case-sensitive.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type string.

Event SeverityID

This is the internal ID of the event log level as number. This is a integer value. There must be an exact match. Please note that this value is case-sensitive.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type number.

Event Level

This is a textual representation of the eventlog level (which is stored as number in %severityid%). This property is automatically localized by the system. If enabled, the event must have the configured level or the rule will not match. There must be an exact match. Please note that this value is case-sensitive.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type string.

Event Keyword

This is a textual representation of the event keyword. This property is automatically localized by the system. If enabled, the event must have the configured event keyword or the rule will not match. There must be an exact match. Please note that this value is case-sensitive.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type string.

Event KeywordID

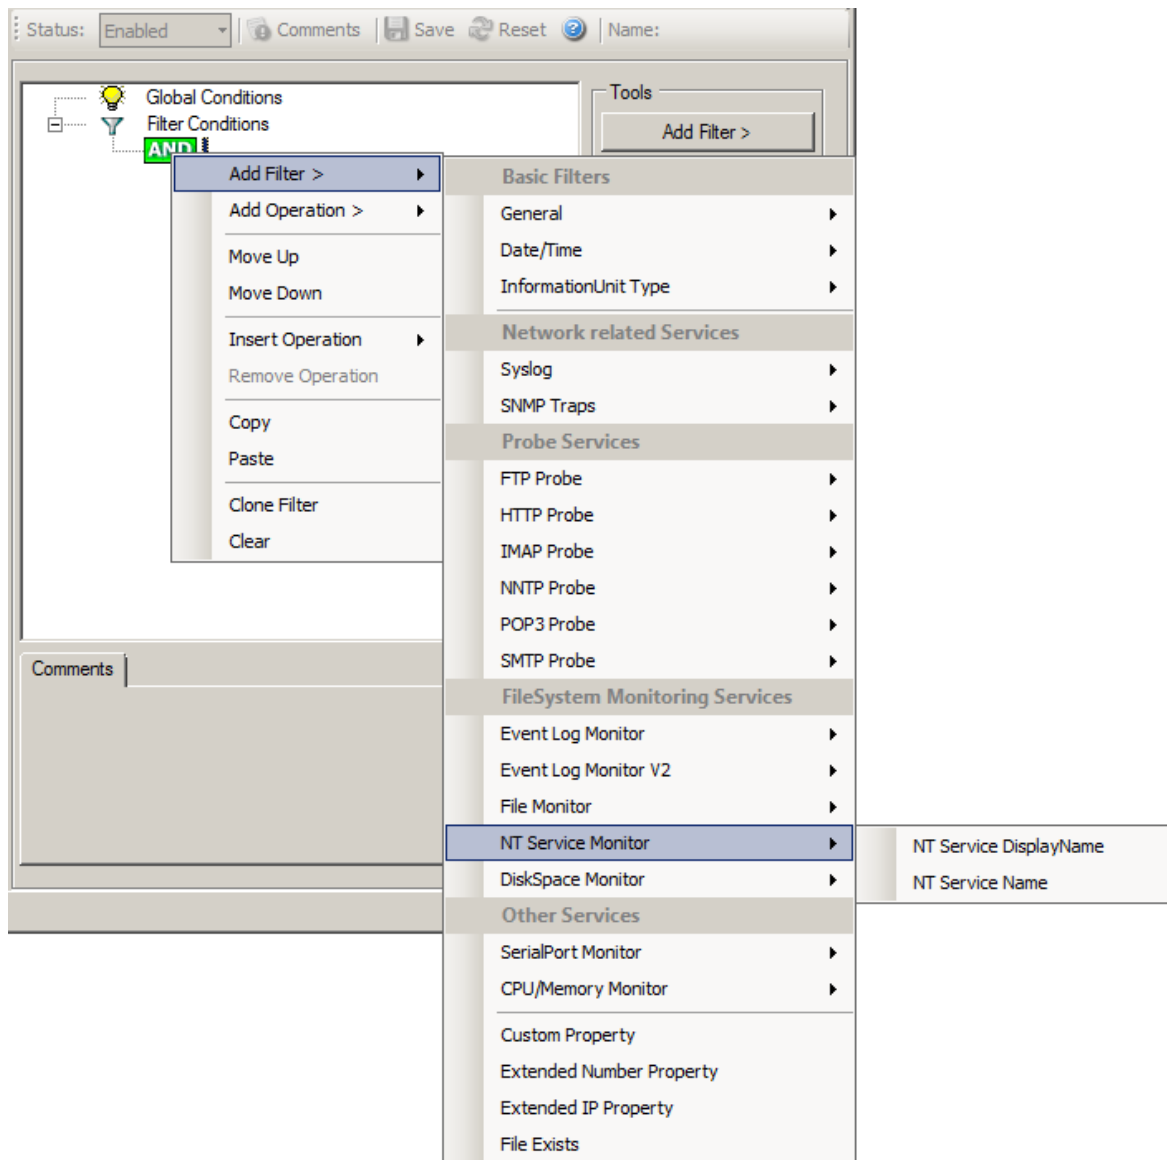
This is the internal keyword ID as string. If enabled, the event must have the configured event keyword ID or the rule will not match. There must be an exact match. Please note that this value is case-sensitive.

This filter condition should only be used with event log information units. If used with others, a mapped value is to be used which might not properly reflect the actual value.

This filter is of type string.

6.4.17 NT Service Monitor

The NT Service Name is used to check if vital operating services are running continuously. By default these services set to "automatic" startup. If the value returned isn't true then corrective measures can be taken e.g. alerts can be generated. See [NT Services Monitor](#) for more details.




Filter Conditions - NT Service Monitor

NT Service DisplayName
NT Service Name (Type=String).

6.4.18 DiskSpace Monitor

This filter works with the disk space report, only. It can be used to trigger actions when disk space is running low and / or becoming free again.



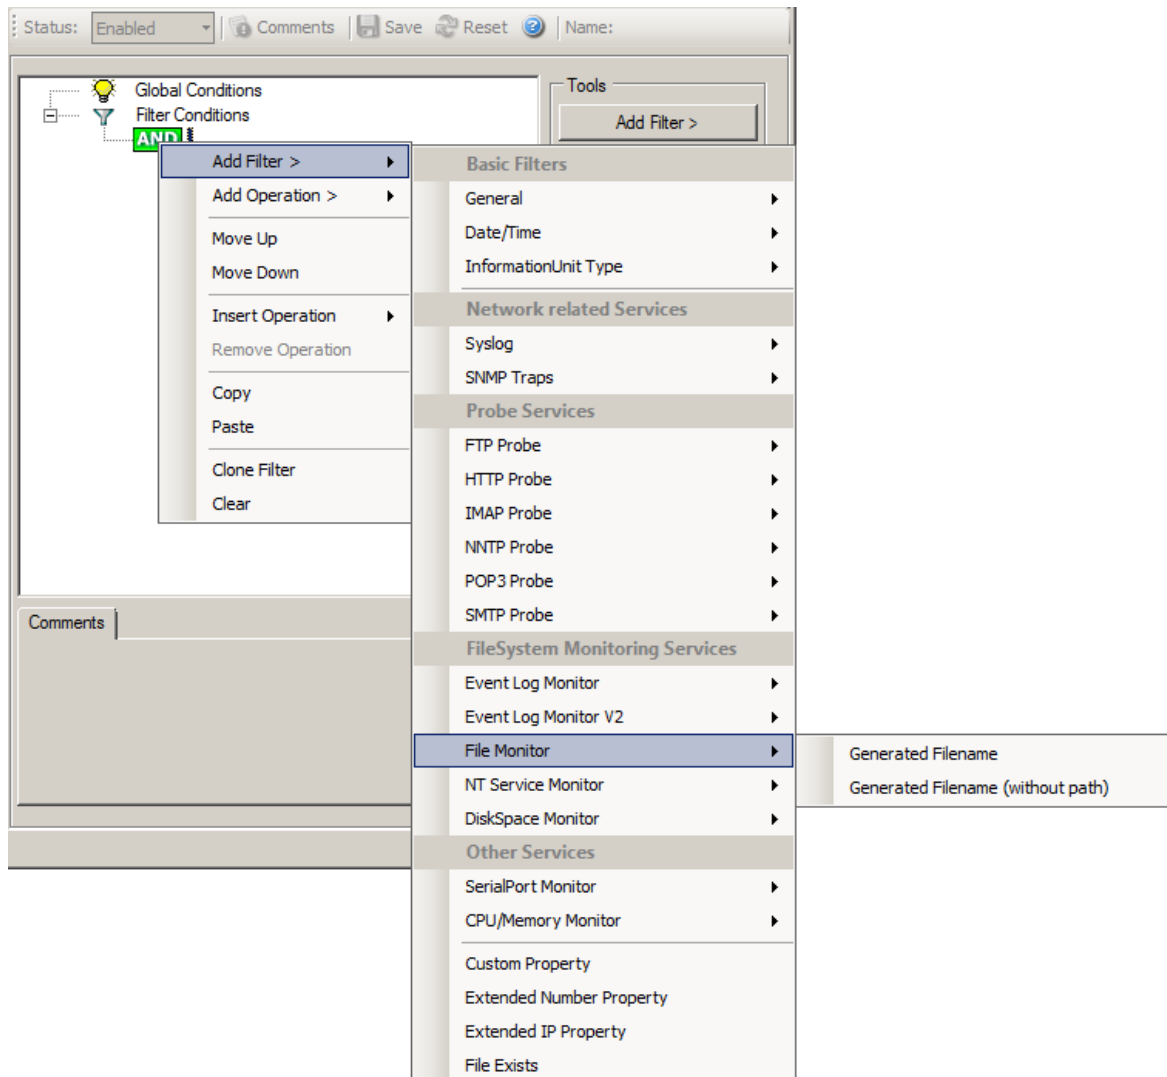
Filter Conditions - DiskSpace Monitor

The following filters are available:

1. Disk Space left (MB) (Type=Number)
2. Disk Space left (GB) (Type=Number)
3. Disk Space left (%) (Type=Number)

6.4.19 File Monitor

File Monitor specific filter is described here.




Filter Conditions - File Monitor

Generated Filename

The configured generic name of the file being reported. Filter has to match exactly to work.

6.4.20 CPU / Memory Monitor

CPU and Memory Monitor specific filter is described here.



Filter Conditions - CPU/Memory Monitor

CPU number

This lets you filter for the number of the monitored CPU.

CPU load

The workload of the CPU as number, can be 0 to 100

Virtual memory load

How much virtual memory is used (MB)

Max virtual memory

How much virtual memory is max available (MB)

Free virtual memory

How much virtual memory is free (MB)

Physical memory load

How much physical memory is used (MB)

Max physical memory

How much physical memory is max available (MB)

Free physical memory

How much physical memory is free (MB)

Total memory load

How much total(Virtual+Physical) memory is used (MB)

Max total memory

How much total(Virtual+Physical) memory is max available (MB)

Free total memory


How much total(Virtual+Physical) memory is free (MB)

6.4.21 SNMP Traps

Using SNMP Traps, since MonitorWare Agent 3.0 now can be used to manage and monitor all sorts of equipment including computers, routers, wiring hubs, toasters and jukeboxes.

A trap is generated when the device feels it should do so and it contains the

information that the device feels should be transmitted.



Filter Conditions - SNMP Traps

Community

It corresponds to the respective SNMP entity.

This filter is of type string.

Enterprise

It corresponds to the respective SNMP entity.

This filter is of type string.

Generic name

It corresponds to the respective SNMP entity.

This filter is of type string.

Version

It corresponds to the respective SNMP entity.

This filter is of type number.


Uptime

It corresponds to the respective SNMP entity.

This filter is of type string.

6.4.22 SerialPort Monitor

SerialPort Monitor specific filter is described here.



Filter Conditions - SerialPort Monitor

Portname


There can be unlimited number of ports in a system - there is no limitation i.e. port names are not specified. But you can use the default list from the "[SerialPort Monitor](#)" configuration window as sample of values, some examples are COM1, LPT1, FILE etc. The MonitorWare Agent Client dynamically reads the properties from the local

machine.

This filter is of type string.

6.4.23 Custom Property

Custom Property specific filter is described here.



Filter Conditions - Custom Property


Custom Property

As the name suggests it is a "Custom Property". Internally in MonitorWare Agent all values are stored in properties. For example the main message is stored in a property called "msg". By using this dialog you can access properties which are dynamic (Like those from SNMP Trap Monitor when using V2 protocol).

This filter is of type string.

6.4.24 File Exists

Filter setting by string



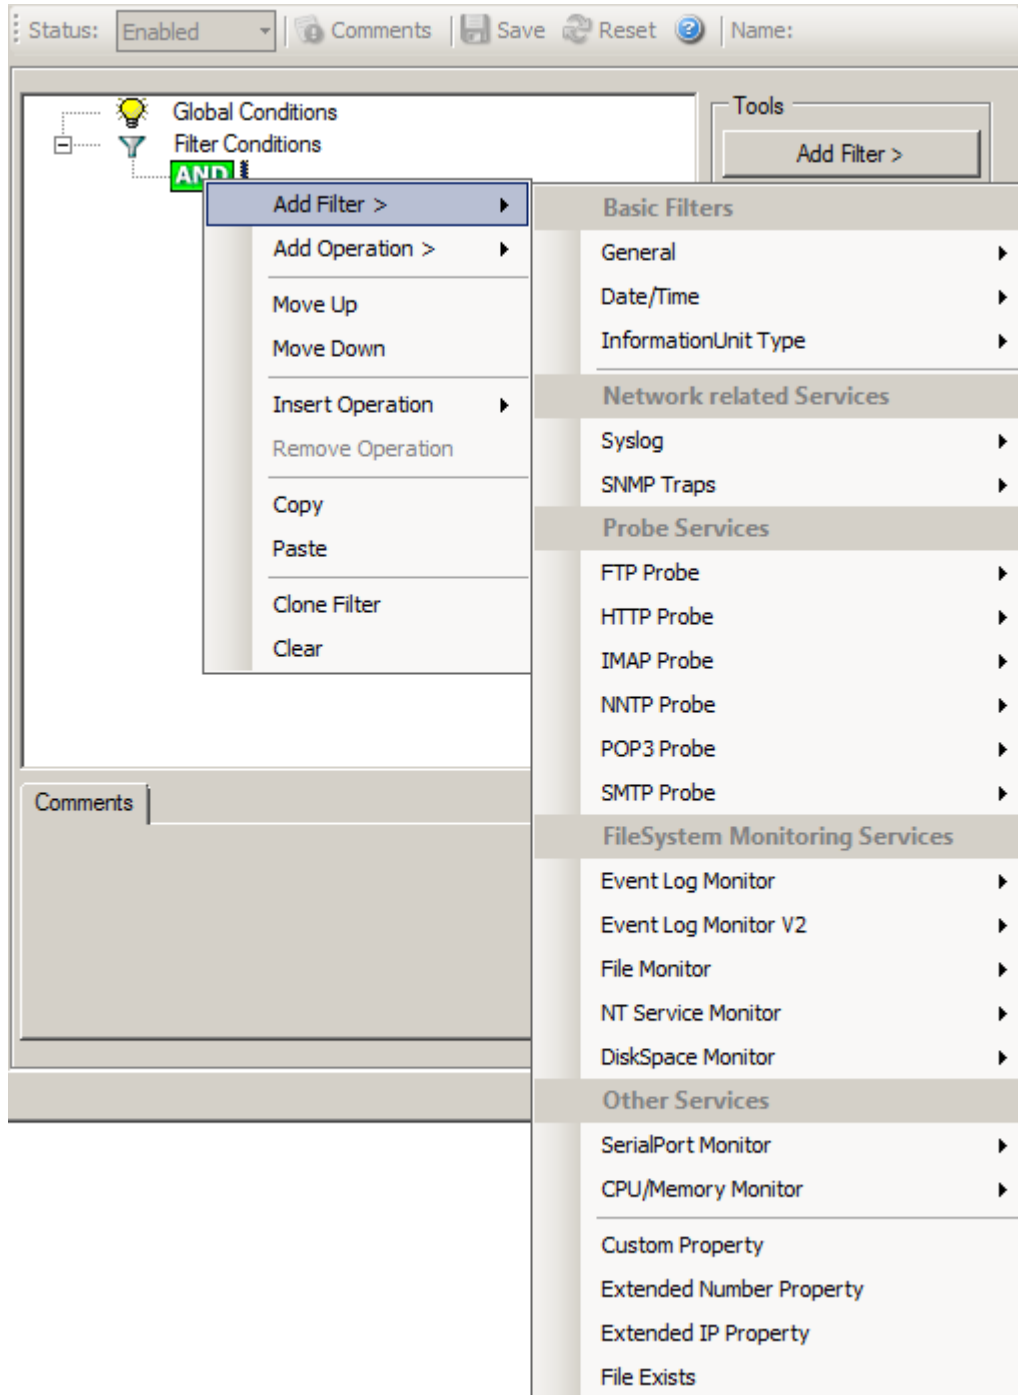
Filter Conditions - File Exists

File Exists

With this Filter you can simply check if a file exists or not. You can directly enter the file and its location or you can use the browse-button to find it.

6.4.25 Extended IP Property

Extended IP Property filter settings



Filter Condition - Extended IP Property

The IP Filter can basically work on any property, but we recommend to only use it on the %source% property, as we usually can be sure that this contains a valid IP Address or hostname. The IP Filter can filter against hostnames and IP Addresses, hostnames are automatically resolved using the internal DNSCache (for obvious performance reasons). If you are going to use a different or custom property, please make sure, that the data in the property is a valid IP Address.

Available compare operations for the IP Filter Type are:

Equal (=): The IP Address must match the one you configured in the Property Value field.

Not Equal (!=): The IP Address must not match the one you configured in the Property Value field.

Higher (>): The IP Address must be higher than the one you configured in the Property Value field. You can use IP Address Formats like 192.168.0.10, 192.168.0, 192.168 or even 192. It depends on what IP Ranges you are going to filter for.

Lower (<): The IP Address must be lower than the one you configured in the Property Value field. You can use IP Address Formats like 192.168.0.10, 192.168.0, 192.168 or even 192. It depends on what IP Ranges you are going to filter for.

If you want to filter for IP Ranges, I recommend to use two filters to define the range, one filter with the "Higher (>)" compare operation, and one with the "Lower (<)" compare operation. This could look like the following:

Filter Condition - Filtering for an IP Range

The filter you can see here will accept all IPs which lie between 172.16.0.110 AND 172.16.0.130. That means, that for every IP that matches these two conditions, the whole filter will evaluate to true and therefore the message will be processed. If the filter does not evaluate to true, the rule will be aborted and the message is sent to the next rule.

6.4.26 Store Filter Results

How to store Filter Results is described here.

Filter Conditions - Store Filter Results Property

Store Filter Results

If a filter matches, you can now store the result of the match into a custom property. This custom property can be used in Actions later.

6.5 Actions

6.5.1 Understanding Actions

Actions tell the application that what to do with a given event. With actions, you can forward events to a mail recipient or Syslog server, store it in a file or database or do many other things with it.

There can be multiple actions for each rule. Actions are processed in the order they are configured. **However you can change the order of the actions by moving them Up or Down.**

6.5.2 Resolve Hostname Action

Many Customers asked for resolve hostname options in different services. This feature has now been implemented as an action. An action can be used with every service, and it doesn't delay the work of a service. See the Screenshot and Descriptions below on how to configure it correctly:




Figure1: Resolve Hostname Action with opened up "Insert" Menu.

Select Source Property from which the name will be resolved:

Click on the Insert menu link on the right side of the textfield to customize the source property from which the name will be resolved.

Destination Property in which the resolved name will be saved to:

Same as above, please click on the Insert menu link on the right side of the textfield to customize the destination property in which the resolved name will be saved to.

Also resolve name if the source property is already a name.

Activates the feature that the name will also be resolved if there is already a source property with that name.

Cache resolved host entry

If activated this will, as it says, cache the resolved host entry.

6.5.3 File Options

This configuration dialog is available both in the defaults section as well as with file logging actions.

File logging is used to write text files of received messages. One file per day is written. New entries are appended to the end of the file.

File locks are released when currently no data is written. Therefore, other applications can access the files while the service is running. However, please be sure that the other applications do not place a file-lock onto it. Popular WordPad does so. In this case, the service will not be able to log any further messages (an error event is written to the NT Event Log in this case). We recommend copying the file when accessing it at runtime - or use notepad.exe, which does not place file-locks on the files it opens.

The filename is build as follows:

<FilePathName><FileName>-year-month-day.<FileExtension>

Parameters in the brackets can be configured via dialog shown below:

The screenshot shows the 'File Logging Options' section of the MonitorWare Agent configuration interface. The window title is 'Name: File | Status: Enabled'. The 'Filename related options' section includes:

- Output Encoding:** System Default (dropdown menu)
- Enable Property replacements in Filename**
- File Path Name:** C:\Program Files (x86)\MonitorWare\Agent (text field with 'Browse' button)
- File Base Name:** MonitorWare Agent (text field with 'Insert' button)
- File Extension:** log (text field)
- Continuous Logging**

Below these options are several checkboxes:

- Create unique filenames**
- Include Source in Filename**
- Use UTC in Filename**
- Segment files when the following filesize is reached (KB)**

At the bottom, there is a slider for **Segment Filesize (KB)** with a value of 1.

File Logging Options

Enable Property replacements in Filename

By activating this option, you can use properties within the file or pathname like % Source% and all the others. For example:

File Path Name can be **F:\syslogs\%source%**

File Base Name can be **IIS-%source%**

If your source is 10.0.0.1, that writes the following file:

F:\syslogs\10.0.0.1\IIS-10.0.0.1.log

Please note that the path f:\syslogs\10.0.0.1 was generated because the source property was used inside the path.

Note: You can use ANY property inside the path and base name. [Event properties](#) are described in the [property replacer section](#).

File Path Name

The base path (directory) of the file. Please see above for exact placement. Default is "c:\temp". The Insert Menu entry allows you to create "**Dynamic Directories**". For example:

File Path Name can be **F:\syslogs\%source%**

[Event properties](#) are described in the [property replacer section](#).

File Base Name

The base name of the file. Please see above for exact placement. Default is "MonitorWare". The Insert Menu entry allows you to recreate "Dynamic Base Filenames". For example:

File Base Name can be **IIS-%source%**

File Extension

The extension to be used when writing the file. Please see above for exact placement. Default is ".log".

Create unique Filenames

If checked, MonitorWare Agent 3.0 creates a unique file name for each day. This is done by adding the current date to the base name (as can be seen above).

If left unchecked, the date is not added and as such, there is a single file with consistent file name. Some customers that have custom scripts to look at the file name use this.

Include Source in Filename

If checked, the file name generation explained above is modified. The source of the Syslog message is automatically added to the file name.

This feature has been introduced because many customers would like to have separate log files for each device. While this can be achieved with multiple rules, it is much more straight forward with this single checkbox. If it is checked, the messages are automatically written to separate files and the file name includes the originating device information.

Use UTC in Filename

This works together with the "Create unique Filenames" setting. If unique names are to be created then select the "Use [UTC](#) in Filename" option, in this case the file name is generated on the basis of universal co-ordinated time (UTC) or on local time. UTC was formerly referred to as "GMT" and is the basis of the time zone system. For example, New York, USA is 5 hours behind UTC. Therefore, if it is 12 noon in New York, the UTC time is 5pm.

When it comes to log file creation, it means that the date is computed on UTC. Taking the same example, if the "Use UTC in Filename" is checked, the log file name would roll over to the next date at 7 pm New York time. If it were unchecked, the rollover would occur exactly at midnight New York time (5 am UTC).

Using UTC for file name creation can be helpful if log files are written among different time zones and later consolidated. Using UTC ensures a consistent time notation across all log files.

Please note that this setting does affect the file name creation only. A different setting controls the dates recorded inside the file.

Segment files when the following file size is reached (KB)

Files are segmented when the defined file size is reached. The file name will have a sequence number appended (_1 to _n).

[Event properties](#) are described in the [property replacer section](#).

Circular Logging

Number of Logfiles: 10

Maximum Filesize (KB): 1

Clear logfile instead of deleting (File will be reused)

File format

Adiscon

Use XML to Report

Include Date and Time

Include Syslog Facility

Include Syslog Priority

Include Date and Time reported by Device

Use UTC for Timestamps

Include Source

Include Message

Include RAW Message

Raw Syslog message

Webtrends syslog compatible

Custom format

Custom Line Format: %msg%%\$CRLF%

Insert

File Logging Options #2

Use Circular Logging

When enabled log files are created and over written in a cycle.

Number of Log files

Once the last logfile is reached, circular logging begins and over write the first log file

again.

Maximum File size

Max filesize of a log file, once this size is reached a new logfile is created.

Clear logfile instead of deleting (File will be reused)

This option causes the File Action to truncate the logfile instead of deleting and recreating it.

File Format

This controls the format that the log file is written in. The default is "Adiscon", which offers most options. Other formats are available to increase log file compatibility to third party applications.

The "Raw Syslog message" format writes raw Syslog format to the log file. That is, each line contains the Syslog message as of RFC 3164. No specific field processing or information adding is done. Some third party applications require that format.

The "WebTrends Syslog compatible" mimics the format that WebTrends applications expect. Please note that we only mimic the log file format. It is still the job of the reporting device (most notable firewall) to generate the correct WebTrends WELF format. The "WebTrends" format is supported because many customers would like to use MonitorWare Agent 3.0 enhanced features while still having the ability to work with WebTrends.

The "Custom" format allows you to customize formats to increase log file compatibility for third party applications. When you choose this option then Custom line format is enabled.

Please note that any other format besides "Adiscon Default" is a fixed format. As such, if it is selected, all other formatting options do not apply and consequently are turned off.

General file options

Under this group box, you can see two options discussed as under:

Use XML to Report

If checked, the message part includes a complete XML-formatted information record. It includes additional information like timestamps, [Syslog facility](#) and priority and others in an easy to parse format. If XML output format is selected, you might consider turning all other information fields off, as they are already included in the XML stream. However, this is not a requirement.

Use UTC for Timestamps

Please see the definition of [UTC](#) above at "Use UTC in Filename". This setting is very similar. If checked, all time stamps are written in UTC. If unchecked, local time is used instead. Again, UTC is useful if logs written in multiple time zones are to be consolidated.

Include <Fieldname>

The various "include" settings controls at the bottom are used to specify the fields which are to be written to the log file. All fields except the message part itself are optional. If a field is checked, it is written to the log file. If unchecked, it will not be written. All fields are comma-delimited.

Please note the difference between the "Date and Time" and "Date and Time reported by Device". Both are timestamps. Either both are written in local time or [UTC](#) based on the "Use UTC for Timestamps" check box. However, "Date and Time" is the time when MonitorWare Agent 3.0 received the message. Therefore, it is always a consistent value.

In contrast, the "Date and Time Reported by Device" is a timestamp taken from the actual message. As such, it is dependent on the reporting device clock, which might be off. In addition, in the case of Syslog messages, there is no time zone information within the device reported timestamp. As such, if devices from multiple time zones are reporting, the timestamp information is not consistent. This is due to Syslog design as of [RFC 3164](#). The Syslog server can be configured to ignore the RFC in this case and provide a consistent time stamp. However, from the view of the log file writer, the "Date and Time Reported by Device" might not be as trustworthy as the "Date and Time" field. Nevertheless, it might also be more useful than the former one. This is the reason both timestamps are present and can individually be selected.

The "Include Message" and "Include RAW Message" fields allow customizing the message part that is being written. The raw message is the message as MonitorWare Agent 3.0 – totally unmodified, received it. This might be useful if a third party application is expecting raw Syslog entries. The message itself is just that part of the Syslog message that is being parsed as message that is without e.g. host information or a tag value. Please note that we recommend selecting only one of these options, as otherwise two message fields are written. Similarly, if none is selected no message is written at all. Please note that we support these configurations, too – there might be a legitimate need for them.

Custom Line Format

Custom Line Format enables you to fully customize the output for the log file. The Insert Menu entry provides further options and they only work in custom line format. Default value is "%msg%%\$CRLF%".

Configure For ...

If you want to generate the reports on log files using [Monilog](#) or [MonitorWare Console](#), then its absolutely necessary that the log files are in a specific format. This option allows you to configure the file logging format for Monilog and MonitorWare Console.

If the log file entries are not in the correct format for MonitorWare Console (for PIX or Windows Reports), then it writes error messages for first 50 lines in Windows event log and ignores them for the generation of report, resulting in a generation of empty report.

And, if the log file entries are not in the correct format for Monilog, then an empty report would be generated.

Following three options are available:

1. Configure for MonitorWare Console PIX Reports
2. Configure for MonitorWare Console Windows Reports
3. Configure for Monilog

Configure for MonitorWare Console PIX Reports

This option changes the file logging format of MonitorWare Agent to the correct format expected by MonitorWare Console for PIX report generation.

Configure for MonitorWare Console Windows Reports

This option changes the file logging format of MonitorWare Agent to the correct format expected by MonitorWare Console for Windows report generation.

Configure for Monilog

This option changes the File Logging format of MonitorWare Agent (i.e. custom line format) to the correct format that is expected by Monilog for report generation.

6.5.4 Database Options

Use database logging to store messages into a database.

Database logging allows writing incoming events directly to any ODBC - compliant database (virtually any database system currently available for the Windows operating system supports ODBC). Adiscon directly supports Microsoft JET databases (as used by Microsoft Access), Microsoft SQL Server and MySQL. We also know of many customers who run it successfully with Oracle and Sybase as well as a variety of other systems.

Once stored inside the database, different message viewers as well as custom applications can easily browse them. The defaults for the write database action are suitable for Adiscon [MonitorWare Console](#) product as well as the web interface.

The database format can be fine-tuned. This is most useful if you intend to run some additional analysis on the database. Also, in high volume environments, tuning the

database action to exactly those fields need helps getting best performance out of the database.

Connection Options

DSN | Verify Database

User-ID: test

Password: **** Enable Password

SQL Connection Timeout: 0 (disabled)

szSQLOptions

Table Name: SystemEvents

Statement Type: CALL (MSSQLStored Procedure)

Output Encoding: System Default

Insert NULLValue if string is empty

Enable Detail Property Logging

Detaildata Tablename: SystemEventsProperties

Maximum value length (Bytes): 512

Datafields

Fieldname	Fieldtype	Fieldcontent
CurUsage	int	curusage
CustomerID	int	CustomerID
DeviceReportedTime	Date Time UTC	timereported
EventBinaryData	text	%bdata%
EventCategory	int	category
EventID	int	id
EventSource	varchar	sourceproc
EventUser	varchar	user
Facility	int	syslogfacility

Database Logging Options

The main feature of the "Write To Database" property sheet is the field list. The default reflects the typical assignment of event properties to database columns. However, you can modify this assignment in any way you like. You only need to keep in mind that Adiscon analysis products (like MonitorWare Console) need the database contents as specified. As such, malfunctions may occur if you modify the database assignments and then use these tools.

The "**fieldname**" is the database column name. It can be any field inside the table. The provided names are those that Adiscon's schema uses - you can add your own if you have a need for this. "**Fieldtype**" is the data type of the database column. It must reflect the column type selected in the database. It must also be consistent in type with the actual property that must be stored. For example, an integer type property like the syslogpriority can be stored in a varchar column. A string data type like the syslogtag can - for obvious reasons - not be stored in an integer column. Finally, the "**Fieldcontent**" is the event property. For a complete list of supported

properties, see [Event properties](#).

You can edit the field list by selecting a row and then modifying the text fields above the table. You can insert and delete rows by selecting the respective button. If you press delete, the currently selected row is deleted. You can move rows up and down by using the arrow keys. Moving them up and down is cosmetic - it will not affect the write to database action.

For string data types, you can use the property replacer. This can be helpful if you would like to store a substring. For example, if you intend to store only the first 200 characters of each message, you can use "%msg:1:200%".

The rest of this section describes the labelled paramters.


Data Sources (ODBC)

If you click on this button, it starts the ODBC administrator of the operating system where you can add, edit or remove a data source(s).

Please Note: The DSN must be a System DSN.

Create Database

If you click on this button, it opens a form as shown below:



Create Database Form

In this form, you have to provide your DSN, User-ID, Password and select your underlying database. After this you have to click Create button to create the table in your database. You can also click Show SQL button to see the SQL query that is to be executed. Close button is to close the form.

DSN

This is the name of the system data source (DSN - data source name) to be used when connecting to the database. Create this in ODBC manager (can be found in control panel under Windows NT). Press the "Data Sources (ODBC)" button to start the operating system ODBC administrator where data sources can be added, edited and removed.

Important:The DSN must be a system DSN, not a user or file DSN. The DSN must be configured to have the correct connection parameters (for example database type and name, server name, authentication mode etc.).

User-ID

The User-ID used to connect to the database. It is dependant on the database system used if it is to be specified (e.g. Microsoft Access does not need one, while Microsoft SQL Server can force you to use one). If in doubt, please see your database administrator.

Password

The password used to connect to the database. It must match the "User-ID". Like the User ID, it is dependent on the database system if a password is needed. Passwords can be stored either encrypted or unencrypted. We highly recommend storing them encrypted.

Enable Encryption

Check to store the ODBC password encrypted. If left unchecked, the password is stored unencrypted. We strongly recommend checking this box.

If you store the password unencrypted for some reason, please be aware of the security implications. In this case, we recommend using an account with limited access privileges. Even when stored encrypted, we recommend using limited privileges accounts. We are not applying strong cryptography here.

Table Name

The name of the table to log to. This name is used to create the SQL insert statement and must match the database definition. Default is "SystemEvents".

Please note that the default table name must be used when other members of the MonitorWare family (like the web interface or the MonitorWare Console) should work with the database. This customization option is meant for those customers that use third-party or custom software.

SQL Statement Type

You can select between a INSERT and Call Statement, which is Microsoft specific for Stored Procedures. This means also this type of SQL Statement will only work if MSSQL is used as database. If you select MSSQL Call Statement, the tablename field will automatically be used as stored procedure name.

Output Encoding

This setting is most important for Asian languages. A good rule is to leave it at "System Default" unless you definitely know you need a separate encoding. "System Default" works perfect in the far majority of cases, even on Asian (e.g. Japanese) Windows versions.

Connection Timeout

Defines the Timeout for the connection

Enable Detail Property Logging

This option logs event properties other than the standard properties to the SystemEventProperties table. A single event can potentially have multiple properties, so selecting this option can result in multiple writes. With Syslog data, however, there are seldom any additional properties. They most often occur when you use the "Post Process" action to define your own properties. Additional properties are typically found in SETP received data originating from an event log monitor, file monitor or database monitor (plus other monitors, but these are the most prominent ones).

For example, with Event Log data received via SETP, these properties contain the actually Windows event properties and the event data. Please note that this does not apply to event log messages received via Syslog, because they are no native events but rather Syslog data.

Please make sure you actually need this before activating it. As a side note, some of the MonitorWare Console reports may need detail logging.

Connection Retry

If a connection is broken, MWAgent gracefully shutdowns the DB Connection and tries to reopen the Connection with the next Actioncall.

Insert NULL Value if string is empty

This option inserts a NULL value, if a property is empty.

6.5.5 OLEDB Database Action

Due the changes to x64, it became more important to also support the newer database layer from Microsoft called OLEDB. The OLEDB Action works similar to the ODBC Action from configuration point of view. The MS SQL OLEDB Provider and JET4.0 OLEDB Provider have been successfully tested in the Win32 environment. Unfortunately, the JET4.0 Provider has not been ported to the x64 platform yet. In our internal performance tests, there was an enhancement of up to 30% compared to ODBC. So this action may also be interesting for people with a huge amount of incoming data.

This Action allows writing incoming events directly to any OLEDB - compliant database.

Once stored inside the database, different message viewers as well as custom applications can easily browse them. The defaults for the write database action are suitable for Adiscon [MonitorWare Console](#) product as well as the web interface.

The database format can be fine-tuned. This is most useful if you intend to run some additional analysis on the database. Also, in high volume environments, tuning the database action to exactly those fields need helps getting best performance out of the database.

RuleSet 1 > Write To OLEDB Database > Write To OLEDB Database 1

Enable: Write To OLEDB Database 1

Settings are saved. Save Reset Save & Close ?

Configure Data Source Verify Database Access

Main Table Name: SystemEvents

SQL Statement Type: INSERT Statement

Output Encoding: System Default

Connection Timeout: 60 seconds

Detail data logging

Enable Detail Property Logging

Detail data TableName: SystemEventsProperties

Maximum value length (Bytes): 512

Insert Delete

Fieldname: Facility

Fieldtype: int

Fieldcontent: syslogfacility [Insert](#)

Fieldname	Fieldtype	Fieldcontent
Facility	int	syslogfacility
Priority	int	syslogpriority
FromHost	varchar	source
Message	text	%msg%
ReceivedAt	DateTime UTC	timegenerated
DeviceReportedTime	DateTime UTC	timereported
CustomerID	int	CustomerID
SystemID	int	SystemID

OLEDB Database Action Options

The main feature of the "OLEDB Database Action" property sheet is the field list. The default reflects the typical assignment of event properties to database columns. However, you can modify this assignment in any way you like. You only need to keep in mind that Adiscon analysis products (like MonitorWare Console) need the database contents as specified. As such, malfunctions may occur if you modify the database assignments and then use these tools.

The "**fieldname**" is the database column name. It can be any field inside the table. The provided names are those that Adiscon's schema uses - you can add your own if you have a need for this. "**Fieldtype**" is the data type of the database column. It must reflect the column type selected in the database. It must also be consistent in type with the actual property that must be stored. For example, an integer type property like the syslogpriority can be stored in a varchar column. A string data type like the syslogtag can - for obvious reasons - not be stored in an integer column. Finally, the "**Fieldcontent**" is the event property. For a complete list of supported properties, see [Event properties](#).

You can edit the field list by selecting a row and then modifying the text fields above the table. You can insert and delete rows by selecting the respective button. If you

press delete, the currently selected row is deleted. You can move rows up and down by using the arrow keys. Moving them up and down is cosmetic - it will not affect the write to database action.

For string data types, you can use the property replacer. This can be helpful if you would like to store a substring. For example, if you intend to store only the first 200 characters of each message, you can use "%msg:1:200%".

The rest of this section describes the labelled paramters.

Configure Data Source

If you click on this button, it starts the OLEDB administrator of the operating system where you can add, edit or remove a data source(s).

Verify Database Access

This button verifies if your indicated data source works fine.

Main Table Name

The name of the table to log to. This name is used to create the SQL insert statement and must match the database definition. Default is "SystemEvents".

Please note that the default table name must be used when other members of the MonitorWare family (like the web interface or the MonitorWare Console) should work with the database. This customization option is meant for those customers that use third-party or custom software.

SQL Statement Type

You can select between a INSERT and Call Statement, which is Microsoft specific for Stored Procedures. This means also this type of SQL Statement will only work if MSSQL is used as database. If you select MSSQL Call Statement, the tablename field will automatically be used as stored procedure name.

Output Encoding

This setting is most important for Asian languages. A good rule is to leave it at "System Default" unless you definitely know you need a separate encoding. "System Default" works perfect in the far majority of cases, even on Asian (e.g. Japanese) Windows versions.

Connection Timeout

Defines the Timeout for the connection

Enable Detail Property Logging

This option logs event properties other than the standard properties to the SystemEventProperties table. A single event can potentially have multiple properties, so selecting this option can result in multiple writes. With Syslog data, however, there are seldom any additional properties. They most often occur when you use the "Post Process" action to define your own properties. Additional properties are typically found in SETP received data originating from an event log monitor, file monitor or database monitor (plus other monitors, but these are the most prominent ones).

For example, with Event Log data received via SETP, these properties contain the actually Windows event properties and the event data. Please note that this does not apply to event log messages received via Syslog, because they are no native events but rather Syslog data.

Please make sure you actually need this before activating it. As a side note, some of the MonitorWare Console reports may need detail logging.

Connection Retry

If a connection is broken, MWAgent gracefully shutdowns the DB Connection and tries to reopen the Connection with the next Actioncall.

6.5.6 Event Log options

This tab is used to configure the logging to the Windows NT / 2000 or XP event log. It is primarily included for legacy purposes.

The screenshot shows the configuration window for the EventLog service. The status is 'Enabled'. The window title is 'Name: EventLog'. The configuration options are as follows:

- Use logsource from service
- Replace Event Log Source
- Custom Eventlog Source:
- Use logsource from service
- Custom Eventlog Type:
- Use Custom Eventlog Type:
- Event ID:
- Message to log:

Event Logging Options

Use logsource from service

Takes the service name as logsource for the log entry. This option is enabled by default.

Replace Event Log Source

If checked, a special mapping mechanism is activated. In this mode, the Windows event source is set to the IP address of the system sending the Syslog message. In addition, the ID is set to [Syslog facility](#). This mode helps to quickly gather information about the system state in Windows event viewer.

However, this mode has its drawbacks. Effectively, we are writing invalid event source information to the event log. This does not harm any application, but Windows event viewer will try to locate the matching message libraries. Of course, this is impossible. As such, event viewer will warn the user that the message library could not be found. Nevertheless, it will display the complete logged message. This happens only in detail view.

Users should fully understand the implications of this mapping mechanism for their environment before turning this option on.

Custom Event Source

EventSource is now fully configurable with all possibilities the property engine gives you. **Please note that content of this field can be configured. [Event properties are described in the property replacer section](#).**

Use custom Eventlog Type

EventType

The type – or severity – this log entry is written with. Select from the available Windows system values.

EventID

The ID to be used when writing to the event log. Different IDs can be used to provide other processes with a consistent interface to specific messages. WinSyslog does not restrict the IDs that can be used. However, if an ID is written that is not registered with the operating system, Windows event viewer places an error message pointing this out before the actual message text. To avoid this text, event IDs 10,000 to 10,100 have been registered with the OS. We highly recommend that these IDs

should be used for all custom messages. IDs below 10,000 should not be used as they might potentially interfere with events generated by MonitorWare Agent 3.0 itself.

Message to Log

It is the message which will be logged into the Windows event log. It is fully configurable what is logged into the Eventlog.

Please note that Insert Menu entry allows you to add replacement characters e.g. % msg% - you can write the actual message of an event into the Windows event log.

Please note that The message content of the message field can be configured. [Event properties](#) are described in the [property replacer section](#).

6.5.7 Mail Options

This tab is used to configure mail (SMTP) parameters. These are the basic parameters for email forwarding. They need to be configured correctly, if mail message should be sent by the service.

Status: Enabled | Comments | Save | Reset | Help | Name: Mail

Mail Server Options | Mail Format Options

Mailserv: 127.0.0.1
Mailserv port: 25

Enable Backup Server, used if first Mailserv fails

Backup Mailserv: 127.0.0.1
Backup Mailserv port: 25

Use SMTP Authentication

SMTP Username:

Session Timeout: 1 second

Use a secure connection (SSL) to the mail server
 Use STARTTLS SMTP Extension
 Use UTC Time in Date-Header

Forward Email Properties - Server Options

Mailserv

This is the Name or IP address of the mail server to be used for forwarding messages. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address. Please note that this server must be able to relay messages if the recipient is not hosted at this server. Be sure to contact your mail server's administrator if in doubt on this issue.

The service expects to talk to a standard SMTP mail server. Message relaying to the final destination must be permitted.

Mailserver Port

Port the mail server is to be contacted at. Usually, this is 25. It might, however, be changed in your system. Then, specify the port your mail server uses. If in doubt, try the default of 25 - or contact your mail server administrator.

Enable Backup Server, used if first Mailserver fails

When enabled, you can configure a second Mailserver that will be used if the regular Mailserver is not available/accessible.

Backup Mailserver

In case that the connection to the main configured mail server can not be established, the backup mail server is tried. Note that an error is only generated, if the connection to the backup server fails as well.

Backup Mailserver Port

Port the mail server is to be contacted at. Usually, this is 25. It might, however, be changed in your system. Then, specify the port your mail server uses. If in doubt, try the default of 25 - or contact your mail server administrator.

Use SMTP Authentication

Check this box if your server requires SMTP authentication. To fight SPAM, more and more server operators allow relaying only for authenticated users. It might even happen that an existing account does no longer work because the server has been reconfigured to disallow anonymous posting.

If your server requires (or supports) SMTP authentication, check this box and enter your User ID and password in the boxes below. The exact values will be provided by your server operator – if in doubt, please ask the mail server support.

If the mail server does not support authentication, leave this box unchecked.

We recommend using authentication if it is available. Even when the current server configuration allows unauthenticated relay, this potentially will change in the future

(as the SPAM problem grows). If you already use authentication, such a server configuration change will not affect you. Otherwise, it will disrupt mail service.

Session Timeout

This option controls if multiple rapidly incoming messages should be combined to a single email message. The SMTP session with the server is held open for the specified timeout period. Please note that the period is specified in milliseconds, not seconds.

If a new event arrives within the specified timeout period, that event will be included in the same email message as the previous one. Then, the timeout is re-started. As such, any events coming in within successive timeout periods will be combined in a single mail.

This is most appropriate when large burst of messages are expected and these should be combined in few mail messages. Otherwise, multiple mail messages can easily overflow the administrator's mailbox.

The session timeout is user configurable between 0 and 4000 milliseconds. Larger values are not supported as they probably affect the SMTP server performance and can lead to unpredictable results.

The session timeout of zero milliseconds has a special meaning: if it is selected, every event will be sent in a separate message, no matter how fast two messages occur after each other.

Use a secure connection (SSL) to the mail server

This option enables SSL-secured traffic to the mail server. Please note, that this only works, if the receiving mail server supports SSL-secured transmission of emails.

Use STARTTLS SMTP Extension

This extension is required for SMTP Servers which can optionally enable encryption during communication.

Use UTC time in Date-Header

Some Email Readers do not support UTC time in date-headers. Therefore here is a switch to turn the UTC time on or off.

Mail Server Options | Mail Format Options

Sender Emailaddress: sender@example.com

Recipient Emailaddress: receiver@example.com

Use legacy subject line processing

Subject: Email for you [Insert]

Mail Priority: Normal Priority

Mail Message Format: Event message:
Facility: %syslogfacility%
Priority: %syslogpriority%
Source: %source% [Insert]

Output Encoding: System Default

Include message / event in email body

Use XML to Report

Forward Email Properties - Format Options

Sender

Email address used as the sender address for outgoing messages. In order for your SMTP server to accept it, it probably must be a valid address.

Recipient

The recipient emails are addressed to. To send a message to multiple recipients, enter all recipient's email addresses in this field. Separate addresses by spaces, semicolons or commas (e.g. "receiver1@example.com, receiver2@example.com"). Alternatively, you can use a single email address and define a distribution list in your mail software. The distribution list approach is best if the recipients frequently change or there is a large number of them. Multiple recipients are also supported. They can be delimited by space, comma or semicolon.

Use legacy subject line processing

This checkbox specifies which type of subject line processing will be done. If it is checked, the old-style processing using single character replacement sequences is applied. If it is left unchecked, the far more powerful event property based method is used.

In legacy mode, the following replacement characters are recognized inside the subject line:

%s

IP address or name (depending on the "resolve hostnames" setting) of the source system that sent the message.

%f

Numeric facility code of the received message

%p

Numeric priority code of the received message

%m

the message itself. Please note: this is the complete message text and can be rather lengthy. As such, it is most probably subject to truncation. If that occurs, all other information after the %m replacement character is also truncated. As such, we strongly recommend using the %m replacement at the end of the subject line only.

%%

It represents a single % sign.

As an example, you may have the subject line set to "Event from %s: "m" and enabled legacy processing. If a message "This is a test" were received from "172.16.0.1", the resulting email subject would read: "Event from 172.16.0.1: This is a test"

In non-legacy mode, the Property Replacer can be used. With it, you can include any property from the event message and also modify it. Please visit the Property Replacer documentation for details.

As an example, in non-legacy mode, you can set the subject line to "Msg: '%msg:1:15%' From: %fromhost%". If the message "This is a lengthy test message" were received from "172.16.0.1", the resulting email subject would read: "Msg: 'This is a lengt' From: 172.16.0.1". Please note that the message is truncated because you only extracted the first 15 characters from the message text (position 1 to 15).

Subject

Subject line to be used for outgoing emails and it is used for each message sent. It can contain replacement characters or "Event Properties" to customize it with event details. This is especially useful when sending email to cellular phones or pagers, which often display only the subject line and not the actual message body. The subject line – after expansion of the any replacement sequences – can hold a maximum of 255 characters. Characters beyond this will be truncated. Please note that many email systems impose a more strict limit and truncation may occur before the 255-character limit. It is advisable to limit the subject line length to 80 characters or less.

The mail body will also include full event information, including the source system, facility, priority and actual message text as well as any other information that came with this event. As there is no size limitation for message bodies, the body always contains the full message received (except otherwise configured – see below).

Please note that Insert Menu entry allows you to add replacement characters e.g. %msg% - you can send out the actual message of an event in the subject line.

There will be one email for each received message. Email delivery is meant for urgent

notifications and actions (e. g. calling pagers and such). It is not meant to provide an email report.

Please note that The message content of the Message field can be configured. Event properties are described in the property replacer section.

Mail Priority

Here you can adjust the priority with which the mail will be sent. You can choose between "low", "normal" and "high" priority. With this you can give your setup some complexity, being able to send some events as "important" and others with less importance.

Mail Message Format

This is the format of the message body. Properties from the event can be included by using the Property Replacer. Please note that the message body is only sent if "Include Message/Event in Email Body" is checked.

Output Encoding

Determines the character encoding mode.

Include message / event in email body

This checkbox controls whether the Syslog message will be included in the message body or not. If left unchecked, it will not be included in the body. If checked, it will be sent.

This option is useful for pagers and mobile phones, especially those with WML support. These devices are often capable of displaying only limited amounts of data. Some do not display the message body at all. As such, it makes limited sense to send a message body. As such, it can be turned off with this option. With these devices, use a subject line with the proper replacement characters.

Even if your WML enabled phone supports receiving message bodies, it might be a good idea to turn them off. WML and WAP are relatively expensive. Generated messages can become lengthy (depending on the message source). As such, it might be appropriate to disable the message body in such a scenario.

This option is must useful together with a well-formatted subject line in non-legacy mode.

Use XML to Report

If checked, the received event will be included in XML format in the mail. If so, the event will include all information, like the original timestamp, the facility, priority etc. XML format is especially useful if the mail is sent to an automated system, which will

then parse the message.

If unchecked, just the plain text message will be included in the mail. This format is more readable for a human reader.

6.5.8 Forward Syslog Options

This dialog controls Syslog forwarding options.

Forward Syslog Properties

Protocol Type

There are various ways to transmit syslog messages. In general, they can be sent via [UDP](#), [TCP](#) or [RFC 3195](#) RAW. Typically, syslog messages are received via UDP protocol, which is the default. UDP is understood by almost all servers, but doesn't guarantee transport. In plain words, this means that syslog messages sent via UDP can get lost if there is a network error, the network is congested or a device (like a router or switch) is out of buffer space. Typically, UDP works quite well. However, it should not be used if the loss of a limited number of messages is not acceptable.

TCP and RFC 3195 based syslog messages offer much greater reliability. RFC 3195 is a special standardized transfer mode. However, it has not received any importance in practice. Servers are hard to find. As one of the very few, Adiscon products support RFC 3195 also in the server implementations. Due to limited deployment, however, RFC 3195 is very little proven in practice. Thus we advise against using RFC 3195 mode if not strictly necessary (e.g. part of your requirement sheet).

TCP mode comes in three flavours. This stems back to the fact that transmission of syslog messages via plain TCP is not yet officially standardized (and it is doubtful if it ever will be). However, it is the most relevant and most widely implemented reliable transmission mode for syslog. It is a kind of unwritten industry standard. We support three different transmission modes offering the greatest compatibility with all existing implementations. The mode "TCP (one message per connection)" is a compatibility mode for Adiscon servers that are older than roughly June 2006. It may also be required for some other vendors. We recommend not to use this setting, except when

needed. "TCP (persistent connection)" sends multiple messages over a single connection, which is held open for an extended period of time. This mode is compatible with almost all implementations and offers good performance. Some issues may occur if control characters are present in the syslog message, which typically should not happen. The mode "TCP (octet-count based framing)" implements algorithms of an upcoming (but not yet finalized) IETF standard. It also uses a persistent connection. This mode is reliable and also deals with embedded control characters very well. However, there is only a limited set of receivers known to support it. As of this writing (January 2007), there were no non-Adiscon receivers supporting that mode. We expect progress once the IETF standard is officially out.

As a rule of thumb, we recommend to use "TCP (octet-count based framing)" if you are dealing only with (newer) Adiscon products. Otherwise, "TCP (persistent connection)" is probably the best choice. If you select one of these options, you can also select a timeout. The connection is torn down if that timeout expires without a message being sent. We recommend to use the default of 30 minutes, which should be more than efficient. If an installation only occasionally sends messages, it could be useful to use a lower timeout value. This will free up connection slots on the server machine.

Syslog Server

This is the name or IP address of the system to which Syslog messages should be sent to. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address.

Syslog Port

The remote port on the Syslog server to report to. If in doubt, please leave it at the default value of 514, which is typically the Syslog port. Different values are only required for special setups, for example in security sensitive areas. Set the port to 0 to use the system-supplied default value (which defaults to 514 on almost all systems).

Instead of the port number, a service name can be used. If so, that name is looked up via the socket service database functions.

Use this backup syslog server if first one fails

The backup server is automatically used if the connection to the primary server fails. The primary server is automatically retried when the next Syslog session is opened. This option is only available when using TCP syslog.

Session Timeout

Timeout value for TCP persistent and octet-count based framing connections.

Syslog Message Options

Syslog Message Options

Syslog processing

With this settings you can assign how your syslog messages will be processed. For processing syslog you can choose out of four different options. You can use [RFC3164](#) or RFC5424 (recommended) which is the current syslog standard, you are able to customize the syslog header or you do not process your syslog and forwards it as it is.

Custom Header Format

In this field you can specify the contents of your syslog header. This option is only available when you choose "Use Custom Syslog Header" in the Syslog Processing menu. The contents can be either a fixed message part which you can write into the field yourself or you use properties as dynamic content. By default the Header field is filled with the content of the RFC 5424 header.

Please note that the header content of the Header field can be configured. [Event properties](#) are described in the [property replacer section](#).

Output Encoding

This setting is most important for Asian languages. A good rule is to leave it at "System Default" unless you definitely know you need a separate encoding. "System Default" works perfect in the far majority of cases, even on Asian (e.g. Japanese)

Windows versions.

Used Message Format

You can use several different message formats for forwarding messages via syslog.

Use Custom Format

The custom format lets you decide how the content of a syslog message looks like. You can use properties to insert content dynamically or have fixed messages that appear in every message. [Event properties are described in the property replacer section.](#)

Use XML to Report

If this option is checked, the forwarded Syslog message is a complete XML-formatted information record. It includes additional information like timestamps or originating system in an easy to parse format.

The XML formatted message is especially useful if the receiving system is capable of parsing XML data. However, it might also be useful to a human reader as it includes additional information that cannot be transferred otherwise.

Forward as MW Agent XML Representation Code

MonitorWare supports a specific XML-Representation of the event. If it is checked, that XML representation is used. It provides additional information (like informationunit type, original source system, reception time & many more) but is harder to read by a human. At the same time, it is obviously easier to parse. **Please note that this option is only "experimental" and is not an official standard.**

Use CEE enhanced Syslog Format

If enabled, the new CEE enhanced Syslog format will be used (work in progress). All useful properties will be included in a JSON Stream. The message itself can be included as well, see the "Include message property in CEE Format" option. Here is a sample how the format looks like for a security Eventlog message:

```
@cee: {"source": "machine.local", "nteventlogtype": "Security",
"sourceproc": "Microsoft-Windows-Security-Auditing", "id": "4648",
"categoryid": "12544", "category": "12544", "keywordid":
"0x8020000000000000", "user": "N\\A", "SubjectUserSid": "S-1-5-11-
22222222-33333333-4444444444-5555", "SubjectUserName":
"User", "SubjectDomainName": "DOMAIN", "SubjectLogonId":
"0x5efdd", "LogonGuid": "{00000000-0000-0000-0000-
000000000000}", "TargetUserName": "Administrator",
"TargetDomainName": " DOMAIN ", "TargetLogonGuid": "{00000000-
0000-0000-0000-000000000000}", "TargetServerName":
"servername", "TargetInfo": " servername ", "ProcessId": "0x76c",
"ProcessName": "C:\\Windows\\System32\\spoolsv.exe", "IpAddress":
"-", "IpPort": "-", "catname": "Logon", "keyword": "Audit Success",
"level": "Information", }
```

Additionally to this format you can set *Include message property in CEE Format*

If enabled, the message itself will be included in the JSON Stream as property. Disable this option if you do not want the message itself in the CEE Format.

Please note you can also make Event ID part of the actual Syslog message while forwarding to a Syslog Server then you have to make some changes in the Forward Syslog Action. [Click here](#) to know the settings.

Message Format

You can change the message format. By default the original message is forwarded.

Please note that the message content of the Message field can be configured. [Event properties](#) are described in the [property replacer section](#).

Add Syslog Source

If this box is checked, information on the original originating system is prepended to the actual message text. This allows the recipient to track where the message originally came from.

Please note: This option is not compatible with [RFC 3164](#). We recommend selecting it primarily when message forwarding to a WinSyslog Interactive Server is intended.

Use zLib Compression to compress the data

With this option you can set the grade of compression for your syslog messages. For more information please read the note at the bottom of this page.

SSL/TLS related Options

The screenshot shows a configuration window with four tabs: 'Syslog Message Options', 'SSL/TLS related Options' (selected), 'TCP related Options', and 'UDP related Options'. Under the 'SSL/TLS related Options' tab, there is a checkbox for 'Enable SSL / TLS Encryption. Note if this Option is enabled, this action will not be able to connect to NON-SSL Syslog Servers.' Below this, there is a 'TLS Mode' dropdown menu set to 'Anonymous authentication'. There are three rows for selecting PEM files: 'Select common CA PEM', 'Select Certificate PEM', and 'Select Key PEM', each with a text input field and a 'Browse' button.

SSL/TLS related Options

Enable SSL / TLS Encryption

If this option is enabled, the action will not be able to talk to a NON-SSL secured server. The method used for encryption is compatible to RFC5425 (Transport Layer Security (TLS) Transport Mapping for Syslog).

TLS Mode

Anonymous Authentication

Default option. This means that a default certificate will be used.

Use Certificate

If this option is enable, you can specify your own certificate. For further authentication solutions, you will need to create your own certificates using OpenSSL Tools for example.

Select common CA PEM

Select the certificate from the common Certificate Authority (CA). The syslog receiver should use the same CA.


Select Certificate PEM

Select the client certificate (PEM Format).

Select Key PEM

Select the keyfile for the client certificate (PEM Format).

TCP related Options



The screenshot shows the 'TCP related Options' tab in a configuration window. It includes a checkbox for 'Use Diskqueue if connection to Syslog Server fails', a slider for 'Split files if this size is reached', a text input for 'Diskqueue Directory' with a 'Browse' button, and a dropdown menu for 'Waittime between connection tries' set to '15 seconds'.


TCP related Options

When using TCP-based syslog forwarding, you have the additional option to use the diskqueue. Whenever a connection to a remote syslog server fails, the action starts caching the syslog messages into temporary files. The folder for these files can be configured. The filenames are generated using a unique GUID which is automatically generated for each Action, thus enabling you to use this feature in multiple Actions. Once the syslog server becomes available again, the cached messages are being sent automatically. If you restart the Service while the Syslog Cache was active, it cannot be checked during service startup if the syslog server is available now. Once the action is called again, the check is done and if the syslog server is available, the messages are being sent. The size of this cache is only limited by the disk size. Files are splitted by 10MB by default, but this can also be configured. The maximum

supported file size is 2GB.

Please Note: This option is not available for UDP or RFC3195.

UDP related Options



UDP related Options

Enable IP Spoofing for the UDP Protocol

This option enables you to spoof the IP Address when sending Syslog messages over UDP. Some notes regarding the support of IP Spoofing. It is only supported the UDP Protocol and IPv4. IPv6 is not possible yet. Due system limitations introduced by Microsoft, **IP Spoofing is only possible on Windows Server 2003, 2008 or higher**. It is NOT possible in Windows XP, VISTA, 7 or higher. For more information see the Microsoft explanation. Also please note that most routers and gateways may drop network packages with spoofed IP Addresses, so it may only work in local networks.

Fixed IP or single property

You can either use a static IP Address or a property. When using a property, the IP Address is tried to be resolved from the content of the property. For example by default the %source% property is used. If the name in this property cannot be resolved to an IP Address, the default local IP Address will be used.

Note on Using Syslog Compression

Compressing syslog messages is an **experimental** feature. There is only a very limited set of receivers who is able to understand that format. Turning on compression can save valuable bandwidth in low-bandwidth environments. Depending on the message, the saving can be anything from no saving at all to about a reduction in half. The best savings ratios have been seen with Windows event log records in XML format. In this case, 50% or even a bit more can be saved. Very small messages do not compress at all. Typical syslog traffic in non-xml format is expected to compress around 10 to 25%.


Please note that compression over TCP connections requires a special transfer mode. This mode bases on an upcoming IETF standard (syslog-transport-tls) that is not yet finalized. That transfer mode is highly experimental in itself. As a result, future releases of our product might not be able to work with the current implementation. So there is a chance that you need to exchange all parts of the syslog/TCP system in future releases. Backwards compatibility can not be guaranteed.

Besides the fact that the mechanisms behind compression are experimental, the

feature itself is solid.

6.5.9 Forward SETP Options

This dialog controls the Send options. With the "Send SETP" action, messages can be sent to a SETP server.



Send SETP Dialog

Servername

The MonitorWare Agent sends [SETP](#) to the server / listener under this name. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address.

Default SETP Port

The Send [SETP](#) sends outgoing requests on this port. The default value is 5432. Set the port to 0 to use the system-supplied default value (which defaults to 5432 if not modified by a system administrator).

Instead of the port number, a service name can be used. If so, that name is looked up via the socket service database functions. The lookup is for protocol TCP.

Please note: The SETP port configured here must match the port configured at the listener side (i.e. MonitorWare Agent 3.0 or WinSyslog Enterprise edition). If they do not match, a Send SETP session cannot be initiated. The rule engine will log this to the NT Event Log.

Options

Under this group box, you can see different options as discussed below:

Enable SSL/TLS

If this option is enabled then this action will be able to connect to SSL/TLS [SETP](#) servers. Please make sure that you want this option to be enabled.

Use zLib Compression to compress the data

It enables zLib compression support. Note that the SETP receiver must have zLib Compression support and enabled, otherwise it does not work.

Compression level

Higher level results in better compression but slower performance.

Session Timeout

The maximum time a session to a SETP server is to be kept open.

Advanced Connection Options

In this group box, you can find the options discussed below:

Connection Timeout

Maximum time a connection can take to connect or disconnect.

Send / Receive Timeout

When sending or receiving data, this timeout applies.

Please note: If this option is enabled, this action is not be able to connect to NON-SSL SETP servers.

6.5.10 Send RELP

This action is roughly equivalent to the "send syslog" action, except that it utilizes the new reliable event logging protocol ([RELP](#)) for message transmission. It can only be used together with a [RELP](#)-enabled receiver but then provides enhance reliability in the communications process.

Note that for reliable operation where messages should be preserved over a service shutdown, queue cache mode must be activated. This is because RELP guards only the transmission channel, but not local processing.

The screenshot shows a configuration window titled 'Send RELP Properties'. At the top, there is a status bar with 'Status: Enabled', a 'Comments' icon, 'Save', 'Reset', and 'Help' icons, and the text 'Name: SendRelp'. Below this are five configuration fields: 'RELP Servername' (empty text box), 'RELP Port' (text box containing '20514'), 'Session Timeout' (dropdown menu showing '30 seconds'), 'Send /Receive Timeout' (dropdown menu showing '1 Minute'), and 'Message Format' (text box containing '%msg%' with an 'Insert' button to its right).

Figure 1: Send RELP Properties

RELP Servername

This is the name or IP address of the system to which RELP messages should be sent to. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address.

RELP Port

The remote port on the RELP server to report to. If in doubt, please leave it at the default value of 20514, which is typically the RELP port. Different values are only required for special setups, for example in security sensitive areas.

Instead of the port number, a service name can be used. If so, that name is looked up via the socket service database functions.

Session Timeout

The maximum time a session to a SETP server is to be kept open.

Send / Receive Timeout

The maximum time a server waits for a response of a remote server. When the timeout expires without receiving a response, the connection is broken and (based on rule settings) being reestablished. This can be a useful option if the remote system drops connections for whatever reason AND the sender system is not notified about this (which, for example, can happen due to some firewall configurations).


Messageformat

You can change the message format. By default the original message is forwarded.

Please note that the message content of the Message field can be configured. [Event properties](#) are described in the [property replacer section](#).

6.5.11 Send MSQueue

In order to use this Action, the "Microsoft Message Queue (MSMQ) Server" needs to be installed. This Action can be used to send a message into the Microsoft Message Queue.



Send MSQueue Properties

Server Computername/IP

Sets the computername or IP which contains the MSQueue you want to query. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address.

Queue name

Specify the Queue name into which you want to write.

Queue Priority

Configure or set the priority property here.

Queue Message Label

Sets the Label text of a queue item.

Queue Message Body


The text here will be set to the body of a queue item.

6.5.12 Start Program

This dialog controls the start process options.

With the "Start Program" action, an external program can be run. Any valid Windows executable can be run. This includes actual programs (EXE files) as well as scripts like batch files (.BAT) or VB scripts (.vbs).

Start process can, for example, be combined with the service monitor to restart failed services. Another example application is a script that deletes temporary files if the disk space monitor detects a low space condition.



Start Process Dialog

Command to execute

This is the path of actual program file to be executed. This can be the path of any valid executable file. A relative file name can be specified if it can be found via the operating system default search path.

Use legacy parameter processing

When enabled, old style parameter processing is used. Otherwise all properties can be used.

Parameters

These parameters are passed to the program executed. They are passed as command line parameters. There is no specific format – it is up to the script to interpret them.

Parameters can contain replacement characters to customize it with event details. This allows passing event data to the script. The following replacement characters can be used:

%d	Date and time in local time
-----------	-----------------------------

%s	IP address or name (depending on the "resolve hostnames" setting) of the source system that sent the message.
%f	Numeric facility code of the received message
%p	Numeric priority code of the received message
%m	The message itself
%%	Represents a single % sign.

In the example above, replacement characters are being used. If a message "This is a test" were received from "172.16.0.1", the script would be started with 3 parameters:

Parameter 1 would be the string "e1" – it is assumed that this has some meaning to the script. Parameter 2 would be the IP address, 172.16.0.1. Parameter 3 would be "This is a test". Please note that due to the two quotes ("), the message is interpreted as a single parameters. If they were missing, it would typically be split into several ones, with parameter 3 being "This", 4 being "is" and so on. So these quotes are very important!

Sync Timeout

Time Out option is under Sync. Processing. When a program is executed, the service waits for it to finish before it carries on further actions. This is needed in order to ensure that all actions are carried out in the correct sequence.

The external program should only run for a limited amount of time. If it would block for some reason, the agent would be prevented from carrying out any further processing. As such, a timeout value must be specified. If the program still runs after the configured timeout, the rule engine cancels it, flags the action as unsuccessful and then carries on with processing.


Important: Even though the timeout value can be as high as 30 seconds, we strongly recommend limiting the run time of external program to below 5 seconds. Otherwise, they could affect the overall performance too much. If the average run time is 5 seconds, the default timeout of 10 seconds ensures that the program can finish even when there is high system activity.

For performance reasons, we also strongly recommend to use the "Start Program" action only for rules that apply relatively seldom.

6.5.13 Play Sound

This action allows you to play a sound file. Since Windows VISTA/2008/7, Microsoft has disabled any interaction between a system service and the user desktop. This includes playing sounds as well. So if you want to use the Play Sound Action on any of this Windows Version, you will need to run the service in console mode (From

command prompt with the -r option).



Play Sound Dialog

Please note: if your machine has multiple sound cards installed, the "Play Sound" action will always use the card, that was installed first into the system.

However there is a work around if you want to use [Play Sound Action](#) for a second sound card!

Filename of the Soundfile

Please enter the name of the sound file to play. **This must be a .WAV file**, other formats (like MP3) are **not** supported. While in theory it is possible that the sound file resides on a different machine, we highly recommend using files on the local machine only. Using remote files is officially not supported (but currently doable if you are prepared for some extra effort in getting this going).

If the file can either not be found or is not in a valid format, a system beep is emitted instead (this should - by API definition - be possible on any system).

Playcount

This specifies how many times the file is played. It can be re-played up to a hundreded times.

Please note: Playing sounds is performance intense and MonitorWare Agent will block all other actions while sounds are being played. As such, we recommend to limit the duration and repeat count of sounds played.

Delay between Plays

If multiple repeats are specified, this is the amount of time that is to be waited for between each individual play.

6.5.14 Send to Communications Port

This action allows you to send a string to an attached communications device, that is it sends a message through a Serial Port.

Status: Enabled | Comments | Save | Reset | Name: SendComPort

Timeout Limit: 1 Minute

Send message to this: COM1:

Port Settings

Bits per second: 57600

Data bits: 8

Parity: No Parity

Stop bits: 1 Stop bit

DTR Control Flow: DTR Control Disable

RTS Control Flow: RTS Control Disable

Message to send: %msg% | Insert

Send to Communications Port Options

Timeout Limit

The maximum time allowed for the device to accept the message. If the message could not be send within that period, the action is aborted. Depending on the device, it may be left in an unstable state.

Port to Send To

Specify the port to which your device is being attached. Typically, this should be one of the COMx: ports. The listbox shows all ports that can be found on your local machine. You may need to adjust this to a different value, if you are configuring a remote machine.

1. MSFAX
2. COM1
3. COM2
4. COM3
5. COM4
6. FILE
7. LPT1
8. LPT2

9. LPT3
10. AVMISDN1
11. AVMISDN2
12. AVMISDN3
13. AVMISDN4
14. AVMISDN5
15. AVMISDN6
16. AVMISDN7
17. AVMISDN8
18. AVMISDN9

Port Settings

Use those settings that your device expects. Please consult your device manual if in doubt.

Bits per Seconds

Bits per second can be 110 and go up to 256000, by default 57600 is selected.

Databits

Databits defines that how many bits you want to send and receive to the communication port.

Parity

With Parity you can configure the Parity scheme to be used. This can be one of the following values:

1. Even
2. Mark
3. No parity
4. Odd
5. Space

Stop bits

You can configure the number of stop bits to be used. This can be one of the following values:

1. 1 stop bit
2. 1.5 stop bits
3. 2 stop bits

DTR Control Flow

DTR (data-terminal-ready) flow control. This member can be one of the following values:

1. DTR_CONTROL_DISABLE - Disables the DTR line when the device is opened and leaves it disabled.
2. DTR_CONTROL_ENABLE - Enables the DTR line when the device is opened and leaves it on.
3. DTR_CONTROL_HANDSHAKE - Enables DTR handshaking.

RTS Control Flow

RTS (request-to-send) flow control. This member can be one of the following values:

1. RTS_CONTROL_DISABLE - Disables the RTS line when the device is opened and leaves it disabled.
2. RTS_CONTROL_ENABLE - Enables the RTS line when the device is opened and leaves it on.
3. RTS_CONTROL_HANDSHAKE - Enables RTS handshaking. The driver raises the RTS line when the "type-ahead" (input) buffer is less than one-half full and lowers the RTS line when the buffer is more than three-quarters full.
4. RTS_CONTROL_TOGGLE - Specifies that the RTS line will be high if bytes are available for transmission. After all buffered bytes have been sent, the RTS line will be low.

Message to Send

This is the message that is to be send to the device. You can enter text plainly and you can also include all properties from the current event. For example, if you have a serial audit printer and you would just plainly like to log arrived messages to that printer, you could use the string "%msg%%\$CRLF%" to write the actual message arrived plus a CRLF (line feed) sequence to the printer.


Please note that the message content of the Message field can now be configured. [Event properties](#) are described in the [property replacer section](#).

6.5.15 Post-Process Event

The post process action allows you to re-parse a message after it has been processed e.g. **Tab Delimited** format.

Such re-parsing is useful if you either have a non-standard Syslog format or if you would like to extract specific properties from the message.

The post process action takes the received message and parses it according to a parse map. The parse map specifies which properties of which type are present at which position in the message. If the message actually matches the parse map, all properties are extracted and are set as part of the event. If the parse map does not match the message, parsing stops at the first-non matching entry.



Post Process Dialog

Templates

Parse maps can be quite complex. In order to facilitate exchange for parse maps, they can be persisted to XML files. Adiscon also plans to provide parse maps for some common devices.

We know that creating a parse map is often not a trivial task. If you are in doubt how to proceed, please contact support@adiscon.com - we will happily assist you with your needs. In this case, you will probably receive a parse map file that you can import here.

The Parse Map Editor

In this dialog, you can edit only in the text boxes above the data grid. When you select an entry in the grid, its values are updated in the textboxes. Any edits made there will automatically be reflected to the grid. Pressing Insert or Delete will create a new entry or delete the currently selected one.

Property

The property name that is to be parsed. The list box is pre-populated with standard and event properties. However, you can add any property name you like. If you create your own properties, we highly recommend prefixing their name with "u-" so that there will be no duplicates with standard properties. Adiscon will never prefix any properties with "u-". For example, if you would like to create a custom property "MyProperty", we highly suggest that you use the property name "u-MyProperty" instead.

The property name "Filler" is reserved. Any values assigned to the Filler-property will be discarded. This is the way to get rid of fill-characters that you do not really need.

Type

This is the format that will be parsed from the message. For example, an integer type will parse one integer from the message while a word type will parse the next word.

Value

Some types need an additional value. If that is needed, you can provide it here.

Message Preview

This is a read-only box. It shows a hypothetical message that would match the configured parsing rules.

Parsing log messages

This article describes how to parse log message via "Post-Process". It illustrates the logic behind Post-Process action.

Get relevant information from logs

Log files contain a lot of information. In most cases only a small part of the log message is of actual interest. Extracting relevant information is often difficult. Due to a variety of different log formats a generic parser covering all formats is not available.

Good examples are firewalls. Cisco PIX and Fortigate firewalls both use syslog for logging.. But the content of their respective log messages are very different. Therefore a method is needed to parse the logs in a generic way. Here Post-Process action of Adiscon's MonitorWare comes into play.

Tool kit for parsing

Post-Process action provides an editor for creating a log format template. A template consists of as many rules as necessary to parse out the relevant information.

Determine necessary information

In order to parse out information it is vital to know the exact structure of the message. Identifying the position of each relevant item is essential. Assuming for auditing purposes the following items are needed:

Timestamp | Source IP-Address | SyslogTag | MessageID | Username | Status | Additional Information

A sample message looks like:

Mar 29 08:30:00 172.16.0.1 %Access-User: 12345: rule=monitor-user-login user=Bob status=denied msg=User does not exist

In order to extract the information let us examine each item within the message. Splitting the message makes it easier to explain. So here we go.

Pos = Position of the character.

*p = Points to the position the parser stands after parsing the rule.

Log = Message subdivided into its characters.

Pro = Property. In the term of Adiscon a property is the name of the item which is parsed out.

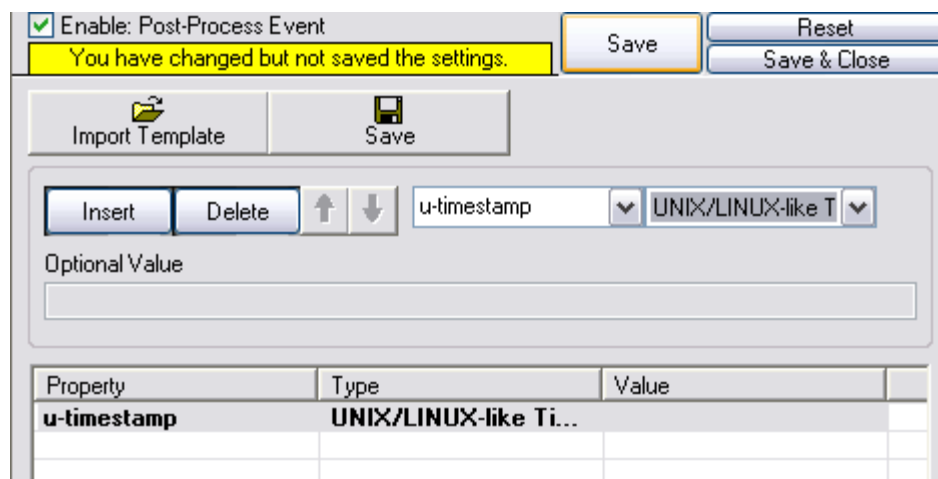
Pos	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Log	M	a	r		2	9		0	8	:	3	0	:	0	0		1	7	2	.
*p	*																			
Pro																				

Note that at beginning of the parse process the parser's pointer points to the first character. Each parse type starts parsing at the current position of the pointer.

Parsing out a Timestamp

The first identified item is a so called Unix/Timestamp. It has always a length of 15 characters. 'UNIX/LINUX-like Timestamp' parse type exactly covers the requirement to parse this item. Therefore insert a rule and select 'UNIX/LINUX-like Timestamp' type. This rule parses out the timestamp and moves the pointer to the next character after the timestamp. Name the property 'u-timestamp' [1].

Note: There is a second timestamp-type, the **ISO-like-timestamp**. It has the format **2006-07-24 13:37:00**.



Post-Process Editor: Inserted a 'UNIX/LINUX like timestamp' rule

Pos	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	
Log	M	a	r		2	9		0	8	:	3	0	:	0	0		1	7	2	.	
*p																*					
Pro	u-timestamp																				

Get the IP-Address

Next item is the IP address. Note that after the timestamp follows a space and then the IP address. Therefore insert a 'Character Match' rule with a space as value. Select the 'Filler' [2] property for this rule. 'Character Match' requires a user defined value. This parse type compares the given value with the character at the current position of the message. The character has to be identical with the given value otherwise the parse process will fail. After applying this parse type the parse pointer is moved to the position immediately after the given value. In our sample this is the start position of the IP Address (Pos 17).

After that the address can be obtained. Place in a 'IP V4 Address' type. This type parses out a valid IP regardless of its length. No need to take care about the characters. Select 'Source' property or name it to whatever you prefer. The parser will automatically move the pointer to the position next to the address.

Property	Type	Value
u-timestamp	UNIX/LINUX-like Timest...	
Filler	Character Match	← space
Source	IP V4 Address	

Message preview of your rules

Jul 24 11:39:36 192.168.0.1

Note the value of 'Character Match' rule is a space.

Pos	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	
Log	0		1	7	2	.	1	6	.	0	.	1		%	A	c	c	e	e	s	
*p													*								
Pro	Filler	Source																			

Obtain the syslogtag

Behind the IP it is a blank followed by a percent sign. The percent indicates that the syslogtag is following. To move the pointer to the syslogtag position once again a 'Character Match' rule is necessary. It has to match the space (actual position of the pointer) and the percent sign. This content is not needed therefore assign it to the 'Filler' property.

A colon is immediately behind the syslogtag. So all characters between the percent sign and the colon are needed. The 'UpTo' type can do this job. Insert an 'UpTo' rule. As value enter ':' (without the quotes) and select the syslogtag property. Note that after parsing the pointer stands on the first character of the 'UpTo' value.

Property	Type	Value
u-timestamp	UNIX/LINUX-like Timest...	
Filler	Character Match	
Source	IP V4 Address	
Filler	Character Match	%
syslogtag	UpTo	:

Message preview of your rules

```
Jul 24 11:45:13 192.168.0.1 %:
```

Pos	26	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	
Log	l		&	A	c	c	e	s	s	-	U	s	e	r	:		l	2	3	4	
*p															*						
Pro		Filler		syslogtag																	

Important: It points to the colon not to the blank.

Take the MessageID

The next interesting item is the MessageID. Move the pointer to start position of the MessageID part. Again, do this by using a 'Character Match' rule. Keep in mind that the pointer points to the colon. Behind the colon is a space and then the MessageID starts. Thus, the value of the rule has to be ':'.

MessageID consist of numbers only. For numeric parsing the 'Integer' parse type exist. This type captures all characters until a non-numeric character appears. The pointer is moved behind the number. Note that numeric values with decimal dots can not be parsed with this type (because they are not integers). This means trying to parse 1.1 results in 1, because the dot is a non-numeric value.

Property	Type	Value
u-timestamp	UNIX/LINUX-like Timest...	
Filler	Character Match	
Source	IP V4 Address	
Filler	Character Match	%
syslogtag	UpTo	:
Source	Character Match	: ← space
u-messageid	Integer	

Message preview of your rules

```
Jul 24 12:19:39 192.168.0.1 %: 12345
```

Pos	39	40	41	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58
Log	r	:		1	2	3	4	5	:		r	u	l	=	m	o	n	i	t	
*p									*											
Pro				u-messageid																

Find the username and status

Looking at the remainder of the message indicates that the username is not immediately after syslogtag. Thankfully though, the username always starts with 'user='. Consequently the 'UpTo' type can be used to identify the username. To get the start position of the username we have to use 'UpTo' together with 'Character Match'. Remember that 'UpTo' points to the first character of the given value. For this reason the 'Character Match' rule is necessary.

After locating the start position of the username 'Word' parse type can be used. 'Word' parses as long as a space sign is found. Enter 'u-username' as property.

Property	Type	Value
u-timestamp	UNIX/LINUX-like Timest...	
Filler	Character Match	
Source	IP V4 Address	
Filler	Character Match	%
syslogtag	UpTo	:
Source	Character Match	:
u-messageid	Integer	
Filler	UpTo	user=
Filler	Character Match	user=
u-username	Word	

Message preview of your rules

```
Jul 24 12:23:53 192.168.0.1 %: 12345user=user=aWord
```

Pos	70	71	72	73	74	75	76	77	78	79	80	81	82	83	84	85	86	87	88	89
Log	i	n		u	s	e	r	=	B	o	b		s	t	a	t	u	s	=	d
*p	Filler		Filler				u-username		*											
Pro																				

Notice: After parsing a word the pointer stands on the space behind the parsed word.

The steps to get the status are very similar to the previous one.

Property	Type	Value
u-timestamp	UNIX/LINUX-like Timest...	
Filler	Character Match	
Source	IP V4 Address	
Filler	Character Match	%
syslogtag	UpTo	:
Source	Character Match	:
u-messageid	Integer	
Filler	UpTo	user=
Filler	Character Match	user=
u-username	Word	
Filler	UpTo	status=
Filler	Character Match	status=
u-status	Word	

Message preview of your rules

```
Jul 24 12:27:38 192.168.0.1 %: 12345user=user=a\Wordstatus=status=a\Word
```

The last rule - Additional Information

One item of interest is left. The last part of the message contains additional information. It starts after 'msg='. So the combination of 'UpTo' and 'Character Match' is used to go to the right position. All characters after 'msg=' until the end of the message are interesting. For this purpose the 'Rest of Message' parse type is available. It stores all characters from the current position until the end of the message. This also means that this rule can only be used once in a template and is always the last rule.

Enable: Post-Process Event

Settings are saved. Save Reset Save & Close

Import Template Save

Insert Delete ↑ ↓ Rest of Message

Optional Value

Property	Type	Value
u-timestamp	UNIX/LINUX-like Timest...	
Filler	Character Match	
Source	IP V4 Address	
Filler	Character Match	%
syslogtag	UpTo	:
Source	Character Match	:
u-messageid	Integer	
Filler	UpTo	user=
Filler	Character Match	user=
u-username	Word	
Filler	UpTo	status=
Filler	Character Match	status=
u-status	Word	
Filler	UpTo	msg=
Filler	Character Match	msg=
msg	Rest of Message	

Message preview of your rules

```
Jul 24 12:29:03 192.168.0.1 %: 12345user=user=aWordstatus=status=aWordmsg=msg=$R$E
$M$A$I$N$D$E$R$$$M$S$G$$$$$$$$$$$$$$$$
```

Complete parse template.

What happens if the parser fails?

If a rule does not match processing stops at this point. This means all properties of rules which were processed successfully until the non-matching rule occurs are available.

Let's assume the fourth rule of the following sample does not match.

Property	Type	Value
u-timestamp	UNIX/LINUX-like Timest...	
Filler	Character Match	
Source	IP V4 Address	
Filler	Character Match	%
syslogtag	UpTo	:
Source	Character Match	:
u-messageid	Integer	

rule does not match

these rules are never reached

The first three rules were processed successfully. Therefore u-timestamp and Source are available. But syslogtag and u-messageid are always empty due to the parser never process this rules.

The Post-Process template which was created in this article is available for [download](#). If you have further question on Post-Process, please contact our [support](#).

[1] Using the "u-" prefix is recommended to differentiate between MonitorWare-defined properties and user defined one. It is not required, but often of great aid. A common trap is that future versions of MonitorWare may use property names that a user has also used. MonitorWare will never use any name starting with "u-", so the prefix also guards against such a scenario.


[2] Filler is a predefined property which acts as a bin for unwanted characters. Essentially, the data is simply discarded.

Please Note: There's also a StepByStep Guide available which describes how the PostProcessAction works, you can find it [here](#).

6.5.16 Net Send

This dialog controls the net send options.

With the "Net Send" action, short alert messages can be sent via the Windows "net send" facility. These messages are delivered on a best-effort basis. If the recipient can be reached, they will pop up in a message box on the recipient's machine. If the recipient cannot be reached, they will simply be discarded. No buffering takes place. Consequently, the rule engine does not check if the message can be delivered. It will never flag an action to be in error due to a reported delivery problem with "net send".



Net Send Dialog

Target

This is the Windows user name of the intended recipient, a NETBIOS machine name or even an IP address (in the form of 10.1.1.1). You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address.

Message to Send

This is the message that is sent to the intended target.

Please note that the message content of the Message to send field can now be configured. [Event properties](#) are described in the [property replacer section](#).


6.5.17 Set Status

This dialog controls the set status options.

Each information unit have specific properties e.g. EventID, Priority, Facility etc. These properties have some values. Lets suppose that EventID has property value 01. Now, If you want to add "**a new property of your own choice**" in the existing set of properties then Set Status action allows you to accomplish this!

You can create a new property and assign any valid desired value to it e.g. we had created a new property as CustomerID and set its value to 01 in the screen-shot below. After you have created the property through this action, then you can define filters for them. There is an internal status list within the product which you can use for more complex filtering.

Please note: when you change a property, the value will be changed as soon as the set status action is carried out. It will not change before that happens and the old value is no longer available thereafter. That means all actions and filter conditions will use the new value after it is set. So if you would like e.g. rename a system, make sure the set status actions are at the top of the rule base!



The screenshot shows a dialog box titled "Set Status". At the top, there is a status indicator "Status: Enabled" with a dropdown arrow, followed by icons for "Comments", "Save", "Reset", and a help icon. The name of the action is "Name: SetStatus". Below this, there are two input fields. The first is labeled "Status variable name" and the second is labeled "Status variable value". Each input field has an "Insert" button to its right.

Set Status Dialog

Status Variable Name

Enter the Property name. That name will from now on be used inside the rule base. More precisely, it will be used in the filter conditions and actions.

Status Variable Value

The value to be assigned to the property. Any valid property type value can be entered.

Insert


Click here to get a list of predefined variables/values to insert.

6.5.18 Set Property

You can set every property and custom properties using this action.

This dialog controls the set property options. With the "Set Property" action, some properties of the incoming message can be modified. This is especially useful if an administrator would like to e.g. rename two equally named devices.

Please note: when you change or create a property, the value will be changed as soon as the set property action is carried out. It will not change before that happens and the old value is no longer available thereafter. That means all actions and filter conditions will use the new value after it is set. So, if you would like e.g. rename a system, make sure the set property actions are at the top of the rule base!



Set Property Dialog

Select Property Type

Select the property type to be changed. The list box contains all properties that can be changed. By default it is set to nothing.

Set Property Value

The new value to be assigned to the property. Any valid property value can be entered. Please use the "Insert Button".

In the example above, the SourceSystem is overridden with the value "newname". That name will from now on be used inside the rule base. More precisely, it will be used in the filter conditions and actions.

Insert

Click here to get a list of predefined variables/values to insert.


6.5.19 Call RuleSet

The dialog shown below controls the Call RuleSet options.

A Call RuleSet action simply calls another rule set in some existing rule set. When this action is encountered, the rule engine leaves the normal flow and go to the called rule set (which may contain many rules as well). It executes all the rules that have been defined in the called Rule Set. After the execution of all of them, it will return to its point from where it left the original flow. Let's take an example to clarify it a little further.

Let's say that Rule 1 has two actions - Action 1 and Action 2. The Action 1 of Rule 1 is an include (Call Ruleset) action. If the filter condition result of Rule 1 evaluates to true, it will execute the Action 1. Since Action 1 is the include action in this example, it will go to the included rule set and will execute its filter condition. If that filter condition evaluates to true, it will execute all of its actions and will return to Action 2 of Rule 1 (of normal flow) and if on the other hand, the filter condition of the included rule set evaluates to false, it will skip all of its actions and will come back to the Action 2 of Rule 1 (of normal flow).

Note that there is no limit on including the rules which means that a rule that has been included in another rule may contain another rule in it which might contain another rule in it and so on.



Call Ruleset Dialog

Ruleset to Call

Select the Ruleset to be called.

Note: Call RuleSet stays disabled until you have more then "One" RuleSet!

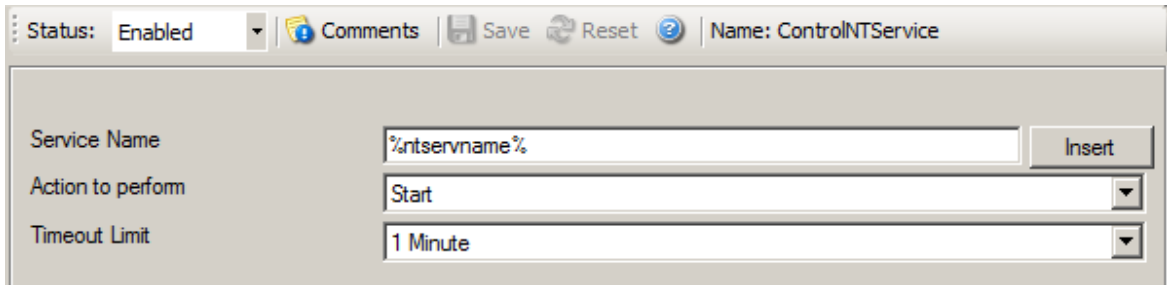
6.5.20 Discard

A Discard Action immediately destroys the current Information Unit and any action of any rule that has been defined after the Discard action execution. When this action is been selected then no dialog appears as nothing needs to be configured for this.

[Click here](#) to see an example about the Discard action.

6.5.21 Control NT Service

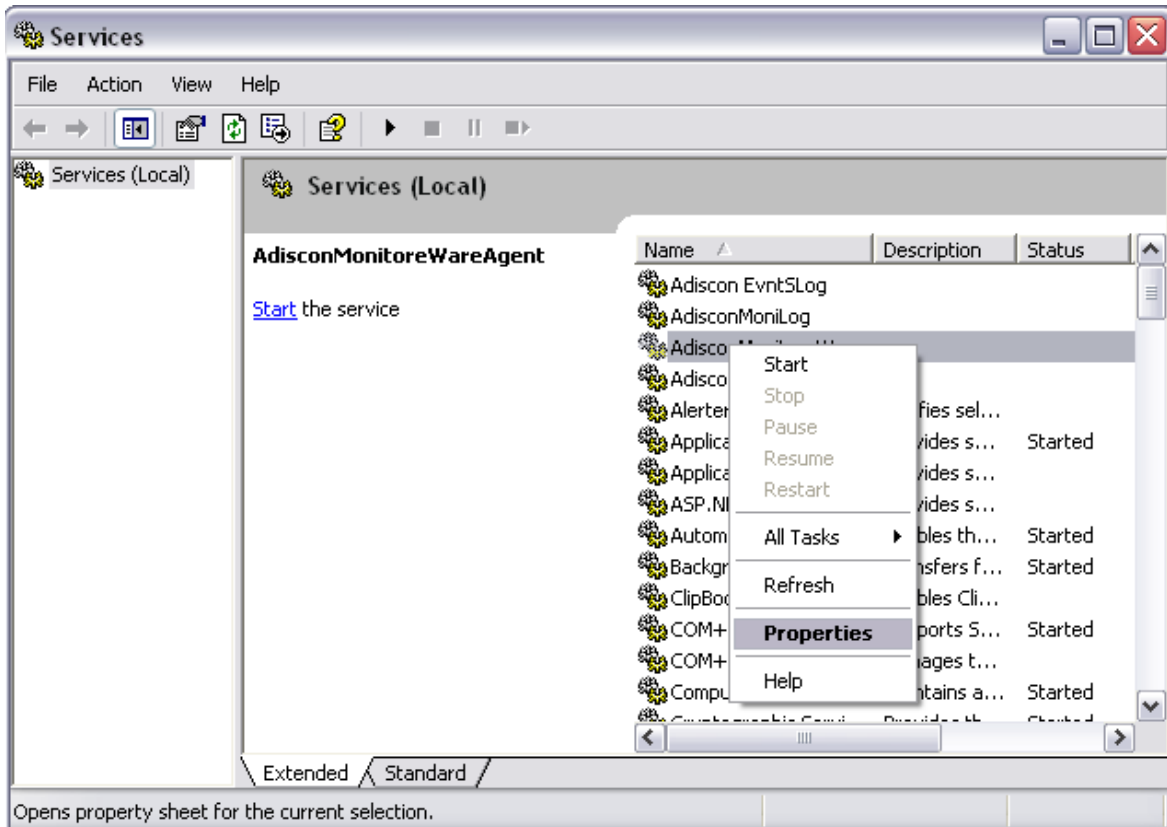
This dialog controls the Control NT Service options.



Control NTService Dialog


Service Name

Specify the service name which you want to control. Please note that it is the internal name, not the display name. You can see the service internal name when you view the properties in the services control panel, as shown below:



Services Control Panel

If you click on properties, you get a properties form as shown below:



Service Properties Form

Here, you can see the Service name as well as the service Display name.

Action to perform


Few Actions are available in the drop down namely Start, Stop, Restart, Pause and Unpause. The selected action is performed on the configured service.

Timeout limit

The amount of time (in [milliseconds](#)) the service is expected to wait for the configured action to take place.

6.5.22 HTTP Request

This dialog controls the HTTP request options.



The screenshot shows a dialog box titled "Name: HTTPRequest" with a status of "Enabled". It features a toolbar with icons for "Comments", "Settings", "Save", "Reset", and "Configure for...". The main configuration area includes the following fields and buttons:

- Host:** Input field containing "www." with an "Insert" button.
- Uri Querystring:** Input field containing "/index.html" with an "Insert" button.
- HTTP Port:** Input field containing "80".
- URL Preview:** A rectangular preview area showing "http://".
- Referer:** Empty input field with an "Insert" button.
- User Agent (Browser):** Input field containing "Mozilla/4.0" with an "Insert" button.

HTTP Request Dialog

Host

Specify the targeted host here.

URL & Querystring

By default this is /index.html. This value is used to construct an URL which is previewed in a rectangular field under Use secure https Protocol option.

Port

This port is to be probed. Please see your server's reference for the actual value to use. For example, mail servers typically listen to port 25 and web servers to 80.

Referrer

An optional configuration option where you can specify a Referrer that is send in the HTTP header.

UserAgent (Browser)

It is also an optional value which can be used to specify a UserAgent that is send in the HTTP header.

6.5.23 Send SNMP Trap

This dialog controls the send SNMP trap options.

Status: Enabled | Comments | Save | Reset | Name: SendSnmp

General SNMP Options

Internet Protocoltype: IPv4
 Protocol Type: UDP
 Agent Receiver: 127.0.0.1
 SNMP Port: 162
 Community: public

SNMP Version 1 Only

Enterprise OID: .1.3.6.1.4.1.3.1.1 Browse
Generic I

Specific Type: 0

SNMP Version 2c Only

Trap OID: .1.3.6.1.4.1.19406.1.2.2 Browse

SNMP Variables

	Variable OID	Variable Type	Variable Value
▶	.1.3.6.1.4.1.19406.1.1.1.7	Octet String	%msg%
*			

Send SNMP Trap Dialog

SNMP Version

Specify the SNMP version here.

Internet Protocoltype

Select the desired protocol type. IPv4 and IPv6 are available. The IPv6 protocol needs to be properly installed in order to be used. Note that one Service can only handle IPv4 or IPv6, so if you want to use both protocols, you will need to create two separate services.

Protocol Type

You can select to listen on UDP or TCP protocol for SNMP Traps.

Agent receiver

Specify the agent that has to receive the SNMP trap. Please note that specifying a host name can cause the SMTP probe to fail if DNS name resolution fails (for example due to a failing DNS server). To avoid this, specify an IP address. You can either use an IPv4, an IPv6 Address or a Hostname that resolves to an IPv4 or IPv6 Address.

SNMP Port

This port is to be probed. Please see your server's reference for the actual value to use. For example, mail servers typically listen to port 25 and web servers to 80.

Community

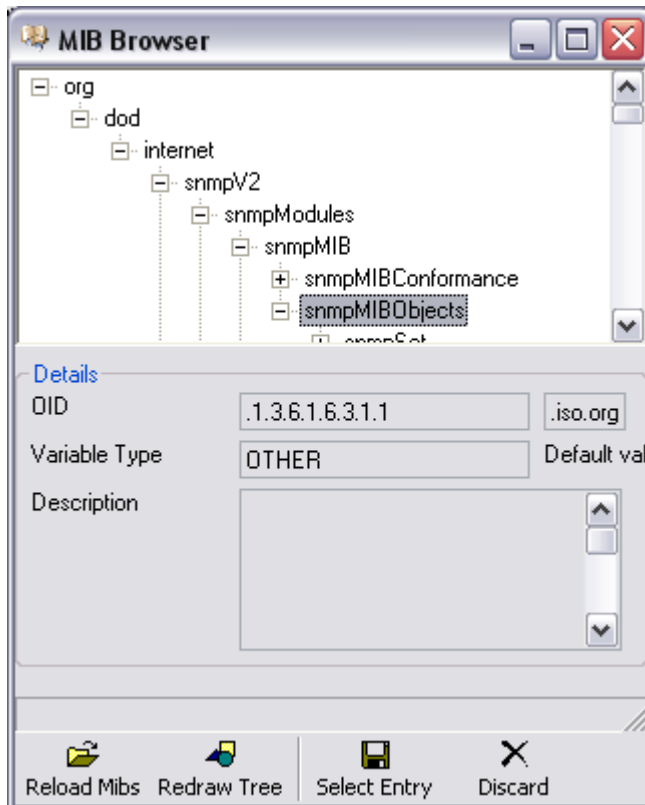
Specify the SNMP community to which the messages belong too.

SNMP V1 Specific Parameters

Under this group box you can see the parameters related to SNMP version 1.

Enterprise OID

Specify the enterprise object ID here. You can use Browse option to select your OID. If you click the Browse link, the screen similar to shown below is appeared:



MIB Browser

You can select your MIB here.

Generic Name

You can specify the generic name of the trap which can be one of the coldStart(0), warmStart(1), linkDown(2), linkUp(3), authenticationFailure(4), egpNeighborLoss(5) and enterpriseSpecific(6).

Specific Type

You can define an additional code for the trap. It is also an Integer value.

SNMP Variables

These are the variables to send in the SNMP Trap. If you know the trap codes, you can enter them manually, otherwise use the inbuild SNMP MIB Browser. Under this group box, you see the following fields:

Variable OID

OID of the SNMP Trap. Use the inbuild SNMP Mib Browser for a list of known and available OIDs.

Variable Type

The variable type of the variable, usually OCTETSTRING or INTEGER. Depending on

this type, the Variable value needs to be formatted correctly (Like for the type IPADDR).

Variable Value


The value of the Variable. It needs to be formatted depending on the variable type.

Please Note:

The "Send SNMP Trap"-Action is capable of sending all kinds of Traps. You can choose the whole variety of the MonitorWare Products' Properties as a value for the messages. With that, you can send SNMP Traps to the Windows internal SNMP Agent or any other device that is able to receive SNMP Traps. Of course you have full enterprise support, too. This gives you the possibility to involve every machine on your network into your security plan or whatever purpose it should serve.

6.5.24 Compute Status Variable

An internal action used to compute a status variable. This is needed for RuleSets which operate on a counter basis. This dialog controls the compute status variable options.



Compute Status Variable Dialog

Status variable

Name of the status variable. You can use property replacement variables here.

Operation Type

In this group box, you can see two options:

Increment Value

It increments the value by the operation value.

Decrement Value

It decrements the value by the operation value.

Operation value

The operation value that is to be used.

6.5.25 Normalize Event

Parameters can be normalized and converted into XML, CSV and JSON formats. The normalization result is stored into a internal properties which can be used for filtering decisions as well as or for output actions.

The action uses liblognorm (<http://www.liblognorm.com/files/manual/index.html>) which is also used by rsyslog. Rulebases created for liblognorm can easily be used and adapted.

Normalize Event Dialog

Normalize Parameter:

Specifies the property that you want to normalize, by default this is the %msg% property.

Select Rulebase File:

The textfile that contains the rulebase definitions (See liblognorm documentation for more).

Lognorm Output Format:

- * Disabled: No additional output format.
- * JSON Format: Creates a string formatted in JSON which is stored in the output property.
- * XML Format: Creates a XML formatted string which is stored in the output property.

* CSV Format: Creates a CSV (Comma separated values) string which is stored in the output property.

Output Property:

The property where the normalized format is saved to.

7 Getting Help

MonitorWare Agent is very reliable. In the event you experience problems, find here how to solve them.

Please note that all options (except priority support) are also open to evaluating customers. So do not hesitate to try them. Help is available in English and German language. Our local resellers may provide local language support. Please check with them.

Frequently Asked Questions

For a current list of Frequently Asked Questions (FAQ), please visit <http://www.mwagent.com/en/FAQ/>. The FAQ area is continuously being updated.

Support Forum

Share questions and answers with your peers! The forum is also monitored by Adiscon support staff. To access the forum, point your browser at <http://forum.adiscon.com/forum,1.html>.

Customer Support System

Our customers service and support system is available at <http://custservice.adiscon.com>. With it, you can quickly open a support ticket via a web-based interface. This system can be used to place both technical support calls as well as general and sales questions. We would appreciate if you select the appropriate category when opening your ticket.

Please note: the customer service system asks you for a User ID and Password when you open it. If you do not have a User ID yet, you can simply follow the "register" link (in the text part) to create one. You can also open a ticket without registering first, in which case the system creates one for you. You receive the generated User ID as part of the email notifications the system generates.

Why using the customer support system? As you see further below, we also offer support by email. In fact, email is just another way to create a ticket in the customer support system. Whenever we reply to your ticket, the system automatically generates an email notification, which includes a link to your ticket as well as the answer we have provided. So for the most cases, you can use email, only. However, there are some situations where the support system should be used:

- **Email notifications do NOT include attachments!** If we provide an attachment,

you must login into the ticket in order to obtain this. For your convenience, each email notification contains an active link that allows you to login immediately.

- **If you seem to not receive responses from us, it is a very good idea to check the web interface.** Unfortunately, anti-SPAM measures are being setup more and more aggressive. We are noticing an increasing number of replies that simply do not make it to your mailbox, because some SPAM filter considered it to be SPAM and removed it. Also, it may happen that your support question actually did not get past our own SPAM filter. We try very hard to avoid this. If we discard mail, we send a notification of this, so you should at least have an indication that your mail did not reach us. Using the customer support system via its own web interface removes all SPAM troubles. So we highly recommend doing this if communication otherwise seems to be disturbed. In this case, please remember that notification emails may also get lost, so it is a good idea to check your ticket for status updates from time to time.

MonitorWare Agent Web Site

Visit the support area at www.mwagent.com/en/support/ for further information. If for any reason that URL ever becomes invalid, please visit www.adiscon.com for general information.

Email

Please address all support requests to support@adiscon.com. An appropriate subject line is highly appreciated.

Please note: we have increasingly seen problems with too-aggressive SPAM filtering, resulting in loss of our replies. If you do not receive a response from us within two working days, we highly recommend re-submitting your support call via the [customer support system](#).

Online Seminars

Adiscon offers a selection of online seminars. This selection is continuously being expanded. All available seminars can be found at <http://www.adiscon.com/Common/en/SeminarsOnline/>

Please note: Windows Media Player is required to view the seminars.

Phone

Phone support is limited to those who purchased support incidents. If you are interested in doing so, please email info@adiscon.com for further details.

Fax

Please direct your faxes to

+49-9349-928820

Toll free in the US: 1-888-900-3772

with "+" being the international dialing prefix, e.g. 011 in the US and 00 in most other countries.

Software Maintenance

Adiscon's software maintenance plan is called [UpgradeInsurance](#). It offers unlimited free upgrades and priority support during its duration. It can be purchased for a period between 1 and 5 years.

[Click here](#) to learn more about UpgradeInsurance.

Non-Technical Questions

Please address all non-technical questions to info@adiscon.com. This email alias answers all non-technical questions like pricing, licensing or volume orders.

Please note: we have increasingly often problems with too-aggressive SPAM filtering, resulting in loss of our replies. If you do not receive a response from us within two working days latest, we highly recommend re-submitting your question via the [customer support system](#).

Product Updates

The [MonitorWare line of products](#) is being developed since 1996. New versions and enhancements are made available continuously.

Please visit www.mwagent.com for information about new and updated products.

8 MonitorWare Concepts

MonitorWare Agent offers advanced monitoring capabilities. It can not only monitor the system it is installed on; it can also include information received from Syslog-enabled devices. To fully unleash MonitorWare's power, you need to learn a bit about its concepts. These web resources (provided links) describe each element in detail.

MonitorWare Agent operates on a set of elements. These are

- [Services](#)
- [Information Units](#)
- [Filter Conditions](#)
- [Actions](#)
- [Rules](#)
- [Rule Engine](#)
- [The SETP Protocol](#)

It is vital to understand each element and the way they interact. MonitorWare Agent has multiple and very powerful capabilities. This enables very quick configuration of highly efficient and comprehensive systems. On the other hand, the concepts must be fully understood to make such complex systems really work.

9 Purchasing MonitorWare Agent

All MonitorWare Agent features can be used for 30 days after installation without a license. However, after this period a valid license must be purchased. The process is easy and straightforward.

The License

The end user license agreement is displayed during setup. If you obtained a ZIP file with the product, there is also a file license.txt inside that ZIP file. If you need to receive a copy of the license agreement, please email info@adiscon.com.

Pricing & Ordering

Please visit <http://www.mwagent.com/en/intermediate-order.php> to obtain pricing information. This form can also be used for placing an order online. If you would like to place a purchase order, please visit <http://www.adiscon.com/Common/en/OrderByPO.asp> to obtain details.

If you would like to receive assistance with your order or need a quote, please contact info@adiscon.com.

10 Reference

The following references provide in-depth information to some very specific things. You may want to review them if you are looking for one of these. Some references are placed on the web and some other are directly contained in this manual. We decided to provide web-links wherever we considered them useful.

- [The MonitorWare Agent Service](#)
- [Support for Mass Rollouts](#)
- [Formats \(XML and Database\)](#)
- [Version History](#)
- [ICMP Codes](#)
- [Property Replacer](#)
- [Comparsion Sheet of Property Replacer](#)

10.1 Comparison of properties Available in MonitorWare Agent, EventReporter and WinSyslog

The property replacer is a reference - the actual properties are very depending on the edition purchased. We have just included information on what is available in which products for your ease and convenience.

Properties Available	MonitorWare Agent WinSyslog	EventReporter
Standard Property	Yes	Yes
Windows Event Log Properties	Yes	Yes
Syslog Message Properties	Yes	Yes
Disk Space Monitor	Yes	
File Monitor	Yes	
Windows Service Monitor	Yes	Yes
Ping Probe	Yes	
Port Probe	Yes	
Database Monitor	Yes	
Serial Port Monitor	Yes	
MonitorWare Echo Request	Yes	
System Properties	Yes	Yes
Custom Properties	Yes	Yes
NNTP Probe	Yes	
HTTP Probe	Yes	
FTP Probe	Yes	
SMTP Probe	Yes	
POP3 Probe	Yes	

10.2 Event Properties

Events have certain properties, for example the message associated with the event or the time it was generated. Each of these properties has an assigned name. The actual properties available depend on the type of event. The following sections describe both how to access properties as well as properties available.

Knowing about event properties is important for building complex filter conditions, customized actions as well as for integrating into a third-party system. Event properties provide a generic way to look at and process the events generated. Thus we highly recommend that you at least briefly read this reference section.

10.2.1 Accessing Properties

Properties are accessed by their name. The component used for this is called the "property replacer". It is a generic component that allows you to merge properties from the event processed to e.g. the email subject line or a log file line. It is a central component that is used as often in the product as possible. The idea behind the property replacer is that there is often a need to specify a value from the event processed.

The property replacer provides very powerful ways to access the properties: they can not only be accessed as one full property. They can also be accessed as substrings and even be reformatted. As such, the property replacer provides a specific syntax to access properties:

`%property:fromPos:toPos:options%`

The percent-signs ("%") indicates the start of a special sequence. The other parameters have the following meanings

FromPos and ToPos can be used to copy a substring from a lengthy property. The options allow to specify some additional formatting.

Within the properties, all time is based on UTC regardless if your preferred time is UTC or localtime. So if you want to display localtime instead of UTC, you have to use the following syntax: `%variable:::localtime%`

10.2.1.1 Property

This is the name of the property to be replaced. It can be any property that a given event posses. If a property is selected that is empty for the event processed, an empty string is returned.

A property is either an [event property](#), a [custom property](#), a dynamic property or a [system property](#).

If a property is selected that is **not** present, the result will always be an empty string, no matter which other options have been selected.

10.2.1.2 FromPos

If you do not want to use the full string from the property, you can specify a start position here. There are two ways to specify the start location:

Fixed Character position

If you know exactly on which position the string of interest begins, you can use a fixed location. In this case, simply specify the character position containing the first character of interest. Character positions are counted at 1.

Search Pattern

A search pattern is specified as follows:

`/<search-pattern>/<options>`

If a search pattern is specified, the property value is examined and the first occurrence of `<search-pattern>` is detected. If it is not found, nothing is returned. If it is found, the position where the pattern is found is the start position or, if the option "\$" is specified, the position immediately after the pattern.

The search pattern may contain the "?" wildcard character, which represents any character. Other wildcards are not supported with the property replacer.

Please note that a slash inside the search pattern will terminate the search field. So

pure slashes can not be used. However, they can be escaped by prefixing them with a backslash (\). The same applies to the '?' character. For example, if you intend to search for "http://" inside a search pattern, you must use the following search string: "/http:\\\\\\".

Default Value

If the FromPos is not specified, the property string is copied starting at position 1.

10.2.1.3 ToPos

If you do not want to use the full string from the property, you can specify the highest character position to be copied here.

Absolute Position

Specify a simple integer if you would like to specify an absolute ending position.

Relative Position

This is most useful together with the search capabilities of **FromPos**. A relative position allows you to specify how many characters before or after the FromPos you would like to have copied. Relative positions are specified by putting a plus or minus ("+"/"-") in front of the integer.

Please note: if you specify a negative position (e.g. -20), FromPos and ToPos will internally be swapped. That is the property value will not be (somehow) reversely copied but they will be in right order. For example, if you specify %msg:30:-20% actually character positions 10 to 30 will be copied.

Search Pattern

Search pattern support is similar to search pattern support in **FromPos**.

A search pattern is specified as follows:

```
/<search-pattern>/<options>
```

If a search pattern is specified, the property value is examined and the first occurrence of <search-pattern> is detected. The search is only carried out in the string that follows FromPos. If the string is not found, nothing is returned. If it is found, the position where the pattern is found is the ending position or, if the option "\$" is specified, the position immediately after the pattern.

The search pattern may contain the "?" wildcard character, which represents any character. Other wildcards are not supported with the property replacer.

Please note that a slash inside the search pattern will terminate the search field. So

pure slashes can not be used. However, they can be escaped by prefixing them with a backslash (\). The same applies to the '?' character. For example, if you intend to search for "http://" inside a search pattern, you must use the following search string: "/http:\\\\\\".

Search Example

A common use case is to combine searches in **ToPos** and **FromPos** to extract a substring that is delimited by two other strings. To do so, use search patterns in both fields. An example is as follows: assume a device might generate message in the form "... error XXX occurred..." where "... " represents additional message text and XXX the actual error cause. You would like to extract the phrase "error XXX occurred". To do so, use the following property replacer syntax:

```
%msg:/error/:/ocured/$/%
```

Please note that the FromPos is used without the \$-option, while in ToPos it is used. If it hadn't been used in ToPos, only the part "error XXX " would have been extracted, as the ToPos would point to the last character before the search string.

Similarly, if only " XXX " should be extracted, the following syntax might be used:

```
%msg:/error/$:/ocured/%
```

If you would also like to remove the spaces (resulting in just "XXX"), you must include them into the search strings:

```
%msg:/error /:/ ocured/$/%
```

Default

If not specified, the ending position will be the last character.

10.2.1.4 Options

Options allow you to modify the the contents of the property. Multiple options can be set. They are comma-separated. If conflicting options are specified, always the last option will be in effect (e.g. specifying "uppercase,lowercase" will lead to lowercase conversion of the property value).

The following options are available with this release of the product:

lowercase	All characters in the resulting property extract will be converted to lower case.
uppercase	All characters in the resulting property extract will be converted to upper case.
uxTimeStamp	This is a special switch for date conversions. It only works if the extracted property value is an ISO-like timestamp (YYYY-MM-DD HH:MM:SS). If so, it will be converted to a Unix-like ctime() timestamp. If the extracted property value is not an ISO-like

	timestamp, no conversion happens.
uxLocalTimeStamp	This is the same as <code>uxTimeStamp</code> , but with local time instead of GMT.
date-rfc3339	This option is for replacing the normal date format with the date format from RFC3339.
date-rfc3164	This option is for replacing the normal date format with the date format from RFC3164.
escapecc	Control characters* in property are replaced by the sequence <code>##hex-val##</code> , where <code>hex-val</code> is the hexadecimal value of the control character (at least two digits, may be more).
spacecc	Control characters* in the property are replaced by spaces. This option is most useful when a message contains control characters (e.g. a Windows Event Log Message) and should be written to a log file.
compressspace	Compresses multiple consecutive space characters into a single one. The result is a string where all words are separated by just single spaces. To also compress control characters, use the <code>compressspace</code> and <code>spacecc</code> options together (e.g. <code>%msg:::spacecc,compressspace%</code>).

Please note that space compression happens on the final substring. So if you use the `FromPos` and `ToPos` capabilities, the substring is extracted first and then the space compression applied. For example, you may have the `msg` string "1 2". There are two space between 1 and 2. Thus, the property replacer expression

```
%msg:1:3:compressspace%
```

will lead to "1 " ('1' followed by two spaces). If you intend to receive "1 2" ('1' followed by one space, followed by '2'), you need to use

```
%msg:1:4:compressspace%
```

or

```
%msg:1:/2/$:compressspace%
```

In the second case, the exact length of the uncompressed string is not known, thus a search is used in `"ToPos"` to obtain it. The result is then space-compressed.

compsp	Exactly the same as compressspace , just an abbreviated form for those that like it brief.
csv	For example <code>%variable:::csv%</code> . This option will create a valid CSV string, for example a string like this <code>this is a "test"!</code> becomes <code>this "this is a ""test""!"</code> where quotes are replaced with double quotes.
convgeruml	Converts German Umlaut characters to their official replacement sequence (e.g. "ö" --> "oe")
localtime	Now you can print the Time with <code>localtime</code> format by using <code>%variable:::localtime%</code>

nomatchblank If this is used, the Property Replacer will return an empty string if the FromPos or ToPos is not found.

replacepercent This option replaces all % occurrences with a double %, which is needed for the property replacer engine in case that a string is reprocessed. This is needed because the percent sign is a special character for the property replacer. Once the property is processed, the double % become automatically one %.

* = control characters like e.g. carriage return, line feed, tab, ...

Important: All option values are case-sensitive. So "uxTimeStamp" works while "uxtimestamp" is an invalid option!

10.2.1.5 Examples

Simple Examples

A good example for this is the email subject line, which has severe length constraints. If you would like to have only the first 40 characters of the actual message text in the subject, you could use the replacer: "%msg:1:40%".

If you know the first 10 characters of the message are meaningless for you but you would like to see the full rest of the message (no matter how long it may be), you can use a sequence like "%msg:11%".

If you would just like to see the plain message from beginning to end, you can simply omit FromPos and ToPos: "%msg".

Of course, all of these sample not only work with the "msg" property, but also with all others like "facility" or "priority", or W3C-log header extracted property names.

More complex Examples

If you would like to extract the 50 characters from the message after the word DROP, you would use the following replacer string:
%msg:/DROP/\$:+50%

If you would like to have the first 40 characters in front of the string "-aborted" (including that string):
%msg:/- aborted/\$:-40%

If you would like to receive everything starting from (and including) "Log:":
%msg:/Log/%

If you would like to have everything between the string "FROM" and "TO" including NONE of the both searchstrings:
%msg:/FROM/\$:/TO/%

If you would just like to log lowercase letters in your log messages:
%msg:::lowercase%

And if you would just like to have the first 50 characters (and these in lower case):
%msg:50:::lowercase%

If you need to change a timestamp to a UNIX-like timestamp, you could use this:
`%datereceived:::uxTimeStamp%`

Please see also the focussed sample in the [ToPos description](#).

A real world Sample

We use the following template to generate output suitable as input for MoniLog:

```
%timegenerated:1:10%,%timegenerated:12:19%,%source%,%syslogfacility%,%
syslogpriority%,EvtSlog: %severity% %timereported:::uxTimeStamp%: %source%/
%sourceproc% (%id%) - "%msg%"%$CRLF%
```

Please note: everything is on one line with no line breaks in between. This example is from the "write to file" action (with custom file format).

10.2.2 System Properties

System properties are special sequences that can be helpful. They are available with all event types. They are:

\$CRLF	A Windows newline sequence consisting in the characters CR and LF. If you just need CR, you can use <code>%%\$CRLF:1:1%</code> and if you need use LF you can use <code>%%\$CRLF:2:2%</code>
\$TAB	An US-ASCII horizontal tab (HT, 0x09) character
\$HT	same as \$TAB
\$CR	A single US-ASCII CR character (shortcut for <code>%%\$CRLF:1:1%</code>)
\$LF	A single US-ASCII LF character (shortcut for <code>%%\$CRLF:2:2%</code>)
\$xNN	A single character, whoms value (in hexadecimal) is given by NN. NN must be two hexadecimal digits - a leading zero must be used if a value below 16 is to be represented. The value 0 (<code>%x00</code>) is invalid and - if specified - replaced by the "?" character. As an example, \$CR could also be expressed as <code>%%\$x0d%</code> . Please note that only one character can be represented. If you need to specify multiple characters, you need multiple \$xNN sequences. An example may be \$CRLF which could also be specified as <code>%%\$x0d%%\$x0a%</code> (but not as <code>%%\$x0d0a%</code>).
\$NOW	Contains the current date and time in the format: YYYY-MM-DD HH.MM.SS Please note that the time parts are delimited by '.' instead iof ':'. This makes the generated name directly suitable for file name generation. If you need just parts of the timestamp, please use the property replacer's substring functionality to obtain the desired part. Use <code>%%\$NOW:1:4%</code> to get the year, <code>%%\$NOW:6:7%</code> to get the month,

	... %%\$NOW:1:10% to get the full datestamp , %%\$NOW:12:20% to get the full timestamp
\$NEWUID	Creates a new UUID (Universally Unique Identifiers), a unique 128-bit integer represented as a 32 digit hexadecimal number.

10.2.3 Custom Properties

Users can create an unlimited number of custom properties. These can be created with for example the "PostProcess" action (if the product edition purchased supports this action).

Custom properties can theoretically have any name, but Adiscon highly recommends to prefix them with "u-" (e.g. "u-MyProperty" - "u" like "user"). This ensures that no compatibility problems will arise in current and future versions of the software. Adiscon guarantees that it will never use the "u-" prefix for Adiscon-assigned properties.

Custom properties can be used just like regular properties. Wherever you can specify a property, you can also specify a custom property.

10.2.4 Event-Specific Properties

Each network event is represented by a so-called "Event Record" (sometime also named an "InfoUnit", an "Unit of Information"). Data obtained from all services will end up as an event. For example, Windows Event Log data, syslog data and a file line obtained by the file monitor will all be an event. That kind of generalization make it easy to deal with all of these events in a consistent way.

Each event has a set of properties which in turn have values. For example, there is a property named "source" and it will always contain an indication of which system the event originated on. Obviously, not every event source does support all properties. For example, a syslog message does not contain a Windows NT Event ID - simply because there is no such thing as an event ID in syslog. So, depending on the type of event, it may contain different properties.

In order to make the product really generally useful, some few properties have been defined in a generic way and are guaranteed to be present in every event, no matter what type it may have. Sometimes this is a "natural" common property, like the "fromhost". Sometimes, though, it may look a bit artificial. An example of the later is the "syslogfacility" property. It is guaranteed to be present in every event - but actually this is a syslog-only thing. The non-syslog event sources either emulate this property (in a consistent manner) or allow the user to configure a syslogfacility that should be used for all events generated by that service. At the bottom line, this will ensure that the property is available in all events and - given proper configuration - that can be extremely helpful for the administrators to set up things in a powerful and generic way.

10.2.4.1 Standard Properties

As outlined under [Event Properties](#), these are properties present in all types of events. Some event types have only these standard properties. Others have additional properties. Those with additional properties are documented in the other sections. If there is no specific documentation for a specific event type, this means that it supports the standard properties, only.

msgPropertyDescribed	A human-readable representation of the message text. While this is generally available, the exact contents largely depends on the source of the information. For example, for a file monitor it contains the file line and for a syslog message it contains the parsed part of the syslog message.
source	The source system the message originated from. This can be in various representations (e.g. IP address or DNS name) depending on configuration settings.
resource	A user-assigned numerical value. Does not have any specific meaning. Primarily intended for quick filtering.
CustomerID	A user-assigned numerical value. Does not have any specific meaning. Primarily intended for quick filtering.
SystemID	A user-assigned numerical value. Does not have any specific meaning. Primarily intended for quick filtering.
timereported	The time the originator tells us when this message was reported. For example, for syslog this is the timestamp from the syslog message (if not configured otherwise). Please note that timereported eventually is incorrect or inconsistent with local system time - as it depends on external devices, which may not be properly synchronized. For Windows Event Log events, timereported contains the timestamp from the event log record.
timegenerated	The time the event was recorded by the service. If messages are forwarded via SETP, this timestamp remains intact.
importance	Reserved for future use.
iut	Indicates the type of the event. Possible values are: 1- syslog message 2- heartbeat 3- Windows Event Log Entry 4- SNMP trap message 5- file monitor 8- ping probe 9- port probe 10- Windows service monitor 11- disk space monitor 12- database monitor 13- serial device monitor
iusers	Version of the event record (info unit). This is a monitorware internal version identifier.

10.2.4.2 Windows Event Log Properties

id	Windows Event ID
severity	severity as indicated in the event log. This is represented in string form. Possible values are: [INF] - informational [AUS] - Audit Success [AUF] - Audit failure [WRN] - Warning [ERR] - Error [NON] - Success (called "NON" for historical reasons)
severityid	The severity encoded as a numerical entity (like in Windows API)
sourceproc	The process that wrote the event record (called "source" in Windows event viewer).
category	The category ID from the Windows event log record. This is a numerical value. The actual value is depending on the event source.
catname	The category name from the Windows event log record. This is a string value. The actual value is depending on the event source. This value is a textual representation from the Category ID. This property could contain line feeds, which can be removed by activating the option "Remove Control Characters from String Parameters" in the advanced options of the EventLog Monitor Service.
user	The user name that was recorded in the Windows event log. This is "N\A" if no user was recorded.
NTEventLogType	The name of the Windows event log this event is from (for example "System" or "Security").
bdata	Windows event log records sometimes contain binary data. The event log monitor service can be set to include this binary data into the event, if it is present. If it is configured to do so, the binary data is put into the "bdata" property. Every byte of binary data is represented by two hexadecimal characters. Please note that it is likely for bdata not to be present. This is because the binary data is seldomly used and very performance-intense.

10.2.4.3 Windows Event Log V2 Properties

id	Windows Event ID
severity	severity as indicated in the event log. This is represented in string form. Possible values are: [INF] - informational [AUS] - Audit Success [AUF] - Audit failure [WRN] - Warning [ERR] - Error [NON] - Success (called "NON" for historical reasons)
severityid	The severity encoded as a numerical entity (like in

	Windows API)
sourceproc	The process that wrote the event record (called "source" in Windows event viewer).
category	The category ID from the Windows event log record. This is a numerical value. The actual value is depending on the event source.
catname	The category name from the Windows event log record. This is a string value. The actual value is depending on the event source. This value is a textual representation from the Category ID. This property could contain line feeds, which can be removed by activating the option "Remove Control Characters from String Parameters" in the advanced options of the EventLog Monitor Service.
user	The user name that was recorded in the Windows event log. This is "N\A" if no user was recorded.
nteventlogtype	The name of the Windows event log this event is from (for example "System" or "Security").
channel	The channel property for event log entries, for classic Event logs they match the %nteventlogtype% property, for new event logs, they match the "Event Channel".
sourceraw	This contains the full internal name of the event source for new event logs, for classic event logs it contains the same value as in %sourceproc%.
level	Textual representation of the eventlog level (which is stored as number in %severityid%). This property is automatically localized by the system.
categoryid	Internal category id as number.
keyword	Textual representation of the event keyword. This property is automatically localized by the system.
user_sid	If available, contains the raw SID of the username (%user%) property.
recordnum	Contains the internal event record number. Please note that if the event log has been truncated before, it may not start with 0 or 1 but a higher number.

10.2.4.4 Syslog Message Properties

rawsyslogmsg	The message as it was received from the wire (unparsed).
syslogfacility	The facility of a syslog message. For non-syslog messages, the value is provided based on configuration. In essence, this is simply an integer value that can be used for quick filtering inside your rules.
syslogfacility_text	The facility of a syslog message. This property is automatically created by using the syslogfacility properly and set to these values: "Kernel", "User", "Mail", "Daemons", "Auth", "Syslog", "Lpr", "News", "UUCP", "Cron", "System0", "System1", "System2", "System3", "System4", "System5", "Local0", "Local1", "Local2", "Local3", "Local4", "Local5", "Local6", "Local7"
syslogpriority	The severity of a syslog message. For non-syslog messages, this should be a close approximation to what a syslog

	severity code means.
syslogpriority_text	The severity of a syslog message. This property is automatically created by using the syslogpriority property and set to these values: "Emergency", "Alert", "Critical", "Error", "Warning", "Notice", "Informational", "Debug"
syslogtag	The syslog tag value, a short string. For non-syslog messages, this is provided based on configuration. In most cases, this is used for filtering.
syslogver	Contains the syslog version number which will be one or higher if a RFC 5424 valid message has been received, or 0 otherwise
syslogappname	Contains the appname header field, only available if the Syslog message was in RFC 5424 format. Otherwise, this field will be emulated by the %syslogtag% property
syslogprocid	Contains the procid header field, only set if the Syslog message was in RFC 5424 format.
syslogmsgid	Contains the msgid header field, only set if the Syslog message was in RFC 5424 format.
syslogstructdata	Contains the structdata header field (in raw format), only set if the Syslog message was in RFC 5424 format.
syslogprifac	Contains combined syslog facility and priority useful to build your own custom syslog headers

10.2.4.5 Disk Space Monitor

currusage	The currently used disk space.
maxavailable	The overall capacity of the (logical) disk drive.

10.2.4.6 CPU/Memory Monitor

wmi_type	This variable is a string and can be one of the following variables: cpu_usage, mem_virtual_usage, mem_physical_usage, mem_total usage
cpu_number	Number of the current checked CPU
cpu_load	The workload of the CPU as number, can be 0 to 100
mem_virtual_load	How much virtual memory is used (MB)
mem_virtual_max	How much virtual memory is max available (MB)
mem_virtual_free	How much virtual memory is free (MB)
mem_physical_load	How much physical memory is used (MB)
mem_physical_max	How much physical memory is max available (MB)
mem_physical_free	How much physical memory is free (MB)
mem_total_load	How much total(Virtual+Physical) memory is used (MB)
mem_total_max	How much total(Virtual+Physical) memory is max available (MB)
mem_total_free	How much total(Virtual+Physical) memory is free (MB)

10.2.4.7 File Monitor

genericfilename	The configured generic name of the file being reported.
generatedbasefilename	Contains the generated file name without the full path.

Special IIS LogFile Properties

The Logfile Fields in IIS Logfiles are customizable, so there is no hardcoded command for their use.

The property-name depends on its name in the logfile. For example we take this Logfile:

```
#Software: Microsoft Internet Information Services 5.0
#Version: 1.0
#Date: 2005-10-27 14:15:25
#Fields: date time c-ip cs-username s-ip s-port cs-method cs-uri-stem cs-uri-query
sc-status cs(User-Agent)
2005-10-27 14:15:16 127.0.0.1 - 192.168.0.1 443 POST /eCommerce/asdf.php
2005-10-27 14:15:16 127.0.0.1 - 192.168.0.1 443 POST /eCommerce/asdf.php
2005-10-27 14:15:16 127.0.0.2 - 192.168.0.1 443 POST /eCommerce/asdf.php
2005-10-27 14:15:16 127.0.0.2 - 192.168.0.1 443 POST /eCommerce/asdf.php
```

As you can see, in our sample the fields are named: date, time, c-ip, cs-username, s-ip, ... and so on.

To use them as a Property inside our MonitorWareProducts, just use the names from your Logfile and add a "p-" before it:

p-date	The Date on which the Event occurs
p-time	The Time on which the Event occurs
p-c-ip	The IP Adress of the User which accessed
p-cs-username	The Username of the User which accessed
p-s-ip	The Server IP
p-s-port	The Server Port
p-cs-method	The Client-Server Method (POST,GET)
p-cs-uri-stem	The accessed File including its path

10.2.4.8 Windows Service Monitor

sourceproc	The name of the service whoms status is being reported (from the Windows service registry).
-------------------	---

10.2.4.9 Ping Probe

echostatus	Status returned for the echo request The status value can be one of the following:
-------------------	---

	0 = IP_SUCCESS
	11002 = IP_DEST_NET_UNREACHABLE
	11003 = IP_DEST_HOST_UNREACHABLE
	11010 = IP_REQ_TIMED_OUT
	11013 = IP_TTL_EXPIRED_TRANSIT
	11016 = IP_SOURCE_QUENCH
	11018 = IP_BAD_DESTINATION
roundtriptime	Round trip time for the ping packet (if successful)

10.2.4.10 Port Probe

responsestatus	The status of the probe.
responsemsg	The response message received (if any)

10.2.4.11 Database Monitor

Database-Monitor created events are a bit different than other events. The reason is that the database fields themselves become properties - but obviously these are not fixed but depend on what you monitor.

All queried data fields are available as properties via their database field name **prefixed with "db-"**.

An example to clarify: we assume the following select statement is used for the database monitor:

```
select name, street, zip, city from addresses
```

There is also an ID column named "ID". So the event generated by this database monitor will have the following specific properties:

- db-ID
- db-name
- db-street
- db-zip
- db-city

These properties will contain the field values as they are stored in the database. Please note that NULL values are translated into empty strings (""), so there is no way to differentiate a NULL value from an empty string with this version of the database monitor.

Other than the custom "db-" properties, no specific database monitor properties exist.

10.2.4.12 Serial Monitor

portname	The name of the port that the data originated from (typical examples are COM1, COM2). The actual name is taken from the configuration settings (case is also taken from there).
-----------------	---

10.2.4.13 MonitorWare Echo Request

responsestatus	The status of the echo request. Possible values: 0 - request failed (probed system not alive) 1 - request succeeded If the request failed, additional information can be found in the <i>msg</i> standard property .
-----------------------	---

10.2.4.14 FTP Probe

ftpstatus	The status of the connection.
ftprespmsg	The response of the connection.

10.2.4.15 IMAP Probe

imapstatus	The status of the connection.
imaprespmsg	The response of the connection.

10.2.4.16 NNTP Probe

nntpstatus	The status of the connection.
nntprespmsg	The response of the connection.

10.2.4.17 SMTP Probe

smtpstatus	The status of the connection.
smtprespmsg	The response of the connection.

10.2.4.18 POP3 Probe

pop3status	The status of the connection.
pop3respmsg	The response of the connection.

10.2.4.19 HTTP Probe

httpstatus	The status of the connection.
httprespmsg	The response of the connection.

10.3 Complex Filter Conditions

The rule engine uses complex filter conditions.

Powerful boolean operations can be used to build filters as complex as needed. A boolean expression tree is graphically created. The configuration program is modelled after Microsoft Network Monitor. So thankfully, many administrators are already used

to this type of Interface. If you are not familiar with it, however, it looks a bit confusing at first. In this chapter, we are providing some samples of how boolean expressions can be brought into the tree.

Example 1

In this example, the message text itself shall be checked. If it contains at least one of three given strings, the filter should become true. If none of the string is found, the boolean expression tree evaluates to false, which means the associated action(s) will not be executed.

In pseudo-code, the filter could be written like this:

```
If (msg = "DUPADDRESS") Or (msg = "SPANTREE") Or (msg = "DUPLEX_MISMATCH) then
    execute action(s)
end if
```

Please note: in the example, we have abbreviated "message" to just "msg". Also note that for brevity reasons we use the equals ("=") comparison operator, nicht the contains. The difference between the equals and the contains operator is that with "contains", the string must just be part of the message.

In the filter dialog, this pseudo code looks as follows:

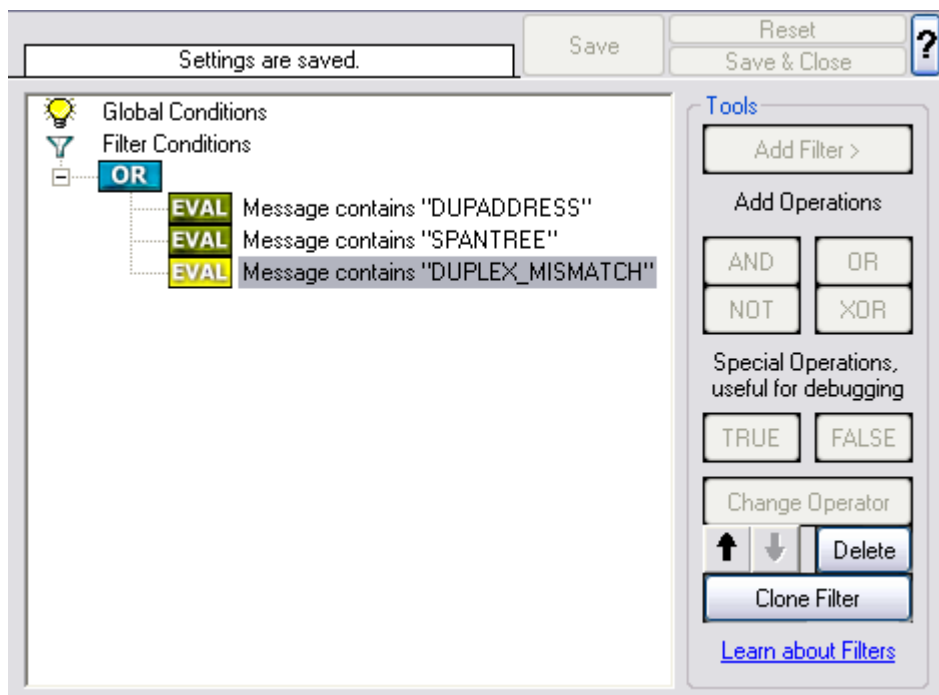


Figure 1 - Example 1

Example 2

Example 2 is very similar to example 1. Again, the message content is to be checked for three string. This time, **all** of these strings must be present in order for the

boolean tree to evaluate to false.

The pseudo code would be as follows (under the same conditions outlined in example 1 above):

```
If (msg = "DUPADDRESS") And (msg = "SPANTREE") And (msg = "DUPLEX_MISMATCH) then
    execute action(s)
end if
```

In the filter dialog, this pseudo code looks as follows:

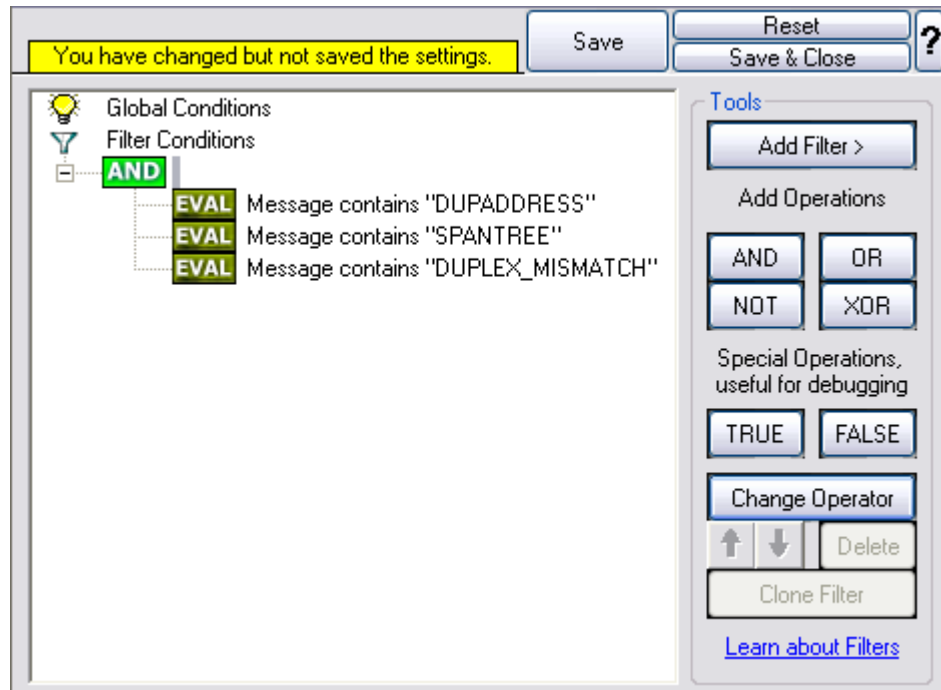


Figure 2 - Example 2

Example 3

This example is a bit more complex version of example 1. Again, the same message text filtering is done, that is if any one of the provided substrings is present, the filter eventually evaluates to true. To do so, the source system must also contain the string "192.0.2", which can be used to filter on a device from a specific subnet.

An example like this can be used for a rule where the administrator of a specific subnet should be emailed when one of the strings indicate a specific event.

The pseudo code would be as follows (under the same conditions outlined in example 1 above):

```
If ((sourceSys = "192.0.2")
    And
    ((msg = "DUPADDRESS") Or (msg = "SPANTREE") Or (msg = "DUPLEX_MISMATCH))
    ) then
    execute action(s)
end if
```

In the filter dialog, this pseudo code looks as follows:

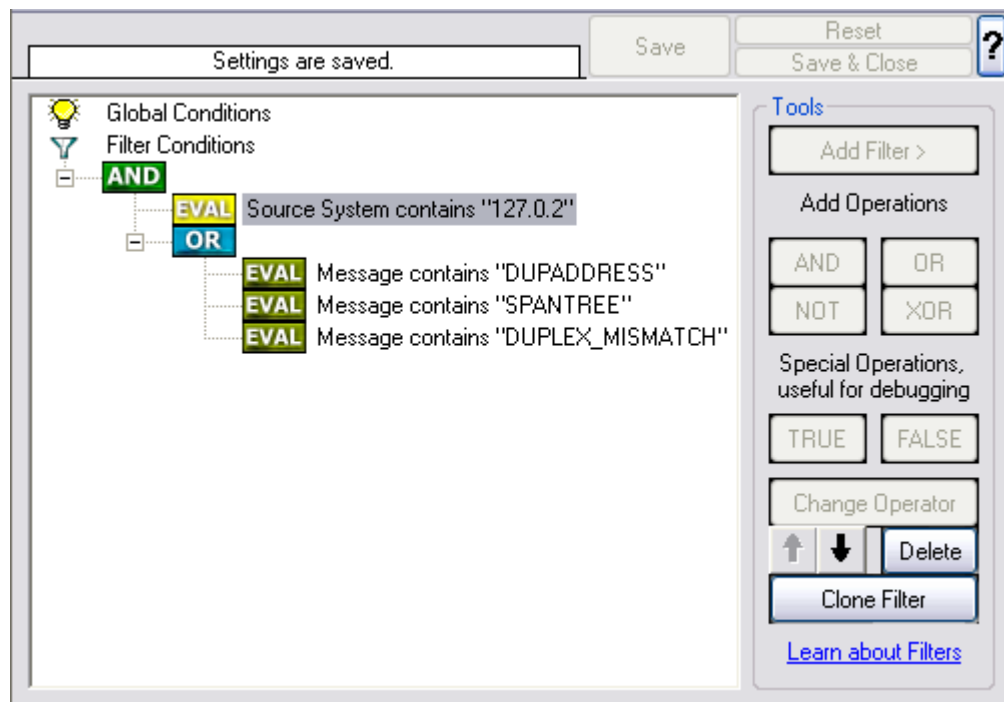


Figure 3 - Example 3

As a side note, you may want to use a range check instead of a simple include for the source system. With a range string check, you can specify that the string must be within a specified column range, in this case obviously at the beginning of the source system IP address.

Real-World Examples

To see some real-world examples of where boolean conditions inside filtering are used, please visit these web links:

- [Detecting Password Attacks under Windows](#)

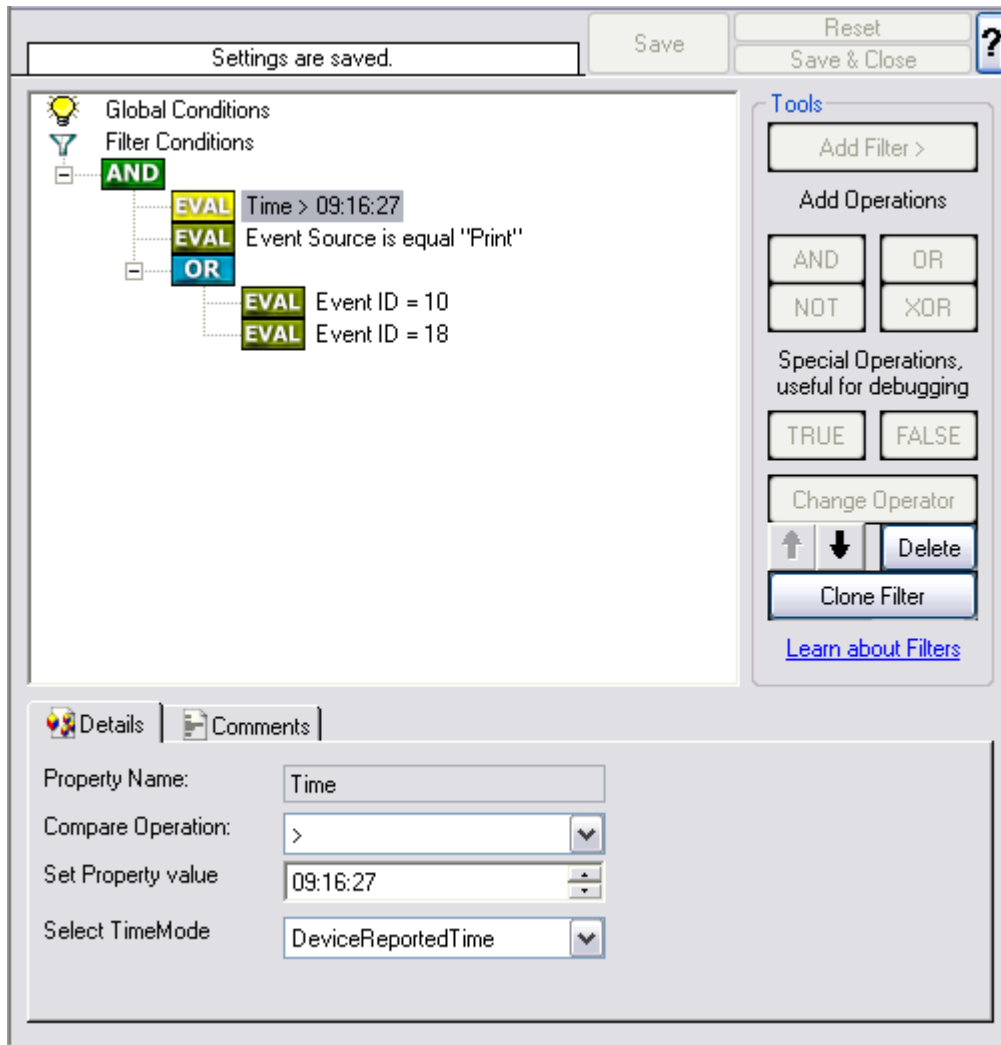
Example 4

In this example, the report is to be filtered in such a way that it shows information only in the case, if the time is greater then certain time with certain event source and one of two event ID's.

In pseudo-code, the filter could be written like this:

```
If (DeviceReportedTime is greater than {9:16:27} AND EventSource is equal to {Print} AND [EventID is equal to {10} OR EventID is equal to {18}])
```

In the filter dialog, this pseudo code looks as follows:



10.4 MonitorWare Agent Shortcut Keys

Use shortcut keys as an alternative to the mouse when working in MonitorWare Agent Client. Keyboard shortcuts may also make it easier for you to interact with MonitorWare Agent. All these shortcuts are usually available in textboxes only. Listed below are the available short keys:

Press	To
CTRL+S	Save
CTRL+X	Cut
CTRL+C	Copy
CTRL+V	Paste
CTRL+Z	Undo

Note: This is in synchronization with most major Windows applications.

10.5 Command Line Switches

There are several command line switches available for using the agent via the command line.

-v	Show version information
-i	Install service
-u	Remove (uninstall) service
-r	Run as console application
-r -o	Run ONCE as console application

If you install the service, you can start and stop the service with the "net start" and "net stop" commands. By using the "-r" switch, you run it only on the command line. When you close the command line, the program will stop working.

The "-v" switch gives you information about the version of the service.

You can install the Service with a custom name by using the command line.

Use

`<servicefile> -i "Service Name"` to install the Service with a custom name, and use `<servicefile> -u "Service Name"` to uninstall the Service.

You can import XML configuration files via the commandline as well. The syntax is quite easy. Simply execute the configuration client and append the name of the configuration file. This could look like this:

```
mwclient.exe example.xml           Sample for MonitorWare Agent
CFGEvntSLog.exe example.xml       Sample for EventReporter
WINSyslogClient.exe example.xml    Sample for WinSyslog
RSyslogConfigClient.exe example.xml Sample for RSyslog Windows Agent
```

or

```
mwclient.exe "example.xml"         Sample for MonitorWare Agent
CFGEvntSLog.exe "example.xml"     Sample for EventReporter
WINSyslogClient.exe               Sample for WinSyslog
"example.xml"
RSyslogConfigClient.exe           Sample for RSyslog Windows Agent
"example.xml"
```

After this is executed, you will see the splash screen of the configuration client and then the import dialogue, which you have to confirm manually.

For doing a silent import, the "/f" (without the quotes) parameter has to be appended. This will look like this:


```
mwclient.exe "example.xml" /f
```

In this case, the filename of the configuration has to be used with the quotes.

10.6 Version Comparison


MonitorWare Agent comes in different versions. Some of them are more feature-rich than others. The manual covers description about the full feature set. In order to remove confusion we have created a Product Comparison Sheet which identifies the differences between different available versions. [Click here](#) to see that which Version provides which services, actions and other features.

10.7 Connect to Computer



Computer Menu

By connecting to another computer, you can remotely configure the machine. Simply go to the computer menu and choose "Connect to Computer". A window will open up.



Select Computer

Here you can enter the name of the machine you want to configure remotely. You can either directly enter the name into the textfield or you use the Browse button to see a list of available machines in the network. If the target machine has a 64Bit Windows operating system, please check the box at **Remote machine is running 64Bit Windows**.

Please Note, for remote configurations, you must ensure, that the remote machine is accessible by network and by the user, that is currently logged on.

11 Copyrights

This documentation as well as the actual MonitorWare Agent product is copyrighted by Adiscon GmbH, Germany. To learn more about other Adiscon products, please visit <http://www.adiscon.com/en/products>. To obtain information on the complete MonitorWare product line, please visit www.monitorware.com.

We acknowledge using these following third party tools. Here are the download links:

Openssl-1.0.1h: <http://www.openssl.org/source/openssl-1.0.1h.tar.gz>
Net-SNMP-5.2.1: <http://www.adiscon.org/3rdparty/net-snmp-5.2.1.tar.gz>
Liblogging: <http://www.adiscon.org/3rdparty/liblogging.zip>
VB6 NeoCaption: http://www.adiscon.org/3rdparty/VB6_NeoCaption_Full_Source.zip

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Other mentioned trademarks are for reference only. They belong to their respective owners.

12 Glossary of Terms

The Glossary of Terms is also available on the Web:

<http://www.monitorware.com/Common/en/glossary/>

The web version most probably has more and more up-to-date content. We highly encourage you to visit the web if in doubt.

12.1 EventReporter

[EventReporter](#) is [Adiscon's](#) solution to forward Windows NT/2000/XP/Vista event log entries to a central system.

These central systems can be either [WinSyslog's](#), other Syslog daemons (e.g. on UNIX) or [MonitorWare Agents](#). EventReporter is part of Adiscon's [MonitorWare line of products](#).

[Click here](#) for more Information about EventReporter.

12.2 Millisecond

A millisecond is a thousandth of a second. It is abbreviated as "ms". As such, 500ms mean half a second.

Inside the [MonitorWare line of products](#), many timers are expressed in milliseconds as a fine control over the services and actions is provided to the administrator.

[Click here](#) for more Information about Milliseconds.

12.3 Monitor Ware Line of Products

[Adiscon's](#) MonitorWare line of products includes monitoring and operations management tools. It consists of several components, each of which can be used either individually or as a complete solution. As of this writing, the following products are available:

- Adiscon Logger (www.monitorware.com/en/logger/)
- ActiveLogger (www.activelogger.com)
- EventReporter (www.eventreporter.com)
- IISLogger (www.iislogger.com)
- MoniLog (www.monilog.com)
- MonitorWare Agent (www.mwagent.com)
- MonitorWare Console (www.mwconsole.com)
- WinSyslog (www.winsyslog.com)

There is also an open source syslog library available for programmers wishing to integrate syslog into their C/C++ programs:

- Liblogging (www.liblogging.org)

New products are continuously being added - please be sure to check www.monitorware.com from time to time for updates.

[Click here](#) for more Information about the MonitorWare Line of Products.

12.4 Resource ID

The Resource ID is an identifier used by the [MonitorWare line of products](#). It is a simple, administrator assigned string value. It can be used to correlate different events - even from different source - to a specific resource.

For example, on a Windows server running Microsoft Exchange, all Exchange events could be assigned to a resource id of "Exchange Server".

In [MonitorWare Agent](#) 1.0 and [WinSyslog](#) 4.0 support for Resource IDs is limited. The field is present and can be persisted to the database or stored in XML files, but besides this there is no value in it.

Later releases of the [MonitorWare Line of Products](#) will much broader support the

Resource ID.

[Click here](#) for more Information about the Resource ID:

12.5 RELP

RELP is the "Reliable Event Logging Protocol". It assures that no message is lost in transit, not even when connections breaks and a peer becomes unavailable. The current version of the RELP protocol has a minimal window of opportunity for message duplication after a session has been broken due to network problems. In this case, a few messages may be duplicated (a problem that also exists with plain tcp syslog).

RELP addresses many shortcomings of the traditional plain tcp syslog protocol. For some insight into that, please have a look at <http://blog.gerhards.net/2008/04/on-unreliability-of-plain-tcp-syslog.html>. Please note that RELP is currently a proprietary protocol. So the number of interoperable implementations is limited.

Note that for reliable operation where messages should be preserved over a service shutdown, queue cache mode must be activated.

12.6 SETP

SETP is the "Simple Event Transfer Protocol". SETP allows reliable delivery of events between SETP supporting systems. [EventReporter](#), [WinSyslog](#) and [MonitorWare Agent](#) support SETP. EventReporter works as SETP Client Only. As such, it can forward events generated and gathered by them to central or intermediary SETP servers. [WinSyslog Enterprise Edition](#) works as SETP client and server, only. The MonitorWare Agent can operate both as a SETP server and client and as such also as a relay. It plays a vital role in a complex, distributed environment.

SETP was developed for MonitorWare. It allows synchronous communication between SETP clients and servers. With SETP, an event can be forwarded exactly as it was on the original event generating system. For example, if a syslog message is received on a remote system, that exact syslog message can be forwarded via as many SETP relays as is configured. During that relaying, no information from the original message is altered or lost. As such, each of the relays as well as the final SETP server will see the original source address, time stamps and message.

Furthermore, SETP guarantees reliable delivery. It is based on [TCP](#), so each of the SETP peers know exactly that the communication partner can successfully receive and process the message. SETP guarantees that new events are only forwarded after the previous ones were successfully received and processed. SETP also checks for on the wire errors. Due to its characteristics, SETP can successfully be used in barely or occasionally connected environments like radio connected systems.

The SETP design is influenced by many industry standard movements, most notably the [BEEP](#) protocol and XML. However, SETP is optimized to have a very lightweight footprint. As such, it can be implemented even in low powered devices with little

overhead.

[Click here](#) for more Information about SETP.

12.7 SMTP

The "Simple Mail Transfer Protocol". This is an Internet standard for sending email messages. Virtually all major email systems are either based on SMTP or at least offer gateways to SMTP capable systems.

SMTP is used for sending email. It can not be used to pick up email messages. For this purpose, protocols like POP3 or IMAP4 are required.

SMTP is highly standardized. As such, a standard email client can work with all SMTP compliant servers. In the public Internet, almost all providers offer SMTP compliant mail servers for their customer's use.

[Click here](#) for more Information about SMTP.

12.8 Syslog Facility

Syslog Facility is one information field associated with a syslog message. It is defined by the [Syslog protocol](#). It is meant to provide a very rough clue from what part of a system the message originated from. Traditionally, under UNIX, there are facilities like KERN (the OS kernel itself), LPD (the line printer daemon) and so on. There are also the LOCAL_0 to LOCAL_7 facilities, which were traditionally reserved for administrator and application use.

However, with the wide adaption of the syslog protocol, the facility field contents has become a little less clear. Most syslog enabled devices nowadays allow configuring any value as the facility. So it is basically left to distinguish different classes of syslog messages.

The facility can be very helpful to define rules that split messages for example to different log files based on the facility level.

[Click here](#) for more Information about Syslog Facility.

12.9 TCP

A reliable IP transport protocol. TCP communication ensures that no packets are lost in transit. As such, it is most useful in low-bandwidth or unreliable environments. Examples are slow WANs or packet radio networks.

12.10 UDP

A non-reliable IP transport protocol. It provides best effort delivery. Typically, in LAN environments UDP packets are never lost. However, in WAN scenarios or with heavily loaded LANs, UDP packets might be lost.

[Click here](#) for more Information about UDP.

12.11 Upgrade Insurance

UpgradeInsurance is [Adiscon's](#) software maintenance plan. It offers free major upgrades as well as priority support. UpgradeInsurance is available for all Adiscon products and can be purchased for a period between 1 and 5 years.

[Click here](#) for more Information about Upgrade Insurance.

12.12 UTC

UTC is the so-called "universal coordinated time". UTC was formerly referred to as "GMT" (Greenwich Mean Time) and is the basis of the international time zone system. For example, New York, USA is 5 hours behind UTC. So if it is 12 noon in New York, the UTC time is 5pm.

The MonitorWare line of products often uses UTC. UTC has the fast advantage of providing one consistent time notation, even if devices are across multiple time zones. This is extremely valuable if a central location is to consolidate events from senders in multiple time zones.

Using UTC might not be appropriate if a whole system is contained within a single time zone. As such, most time parameters inside the MonitorWare line of products can be configured to work with local time instead of UTC.

[Click here](#) for More Information about UTC.

12.13 NNTP

NNTP stands for Network News Transport Protocol. This protocol is used by client and server software to carry USENET postings back and forth over a TCP/IP network.

When you are using any of the common softwares like Netscape, Internet Explorer, etc, you are taking benefit of the NNTP connection to participate in newsgroups.

12.14 SNMP

SNMP stands for Simple Network Management Protocol. A set of standards for communication with devices connected to a TCP/IP network, like routers, hubs and switches. A device is said to be SNMP compatible if it can be monitored and/or controlled using SNMP messages.

SNMP messages are known as PDU's - Protocol Data Units. Devices that are SNMP compatible contain SNMP 'agent' software to receive, send, and act upon SNMP messages. Software for managing devices via SNMP are available for every kind of commonly used computer and are often bundled along with the device they are designed to manage. Some SNMP software is designed to handle a wide variety of devices.

12.15 FTP

FTP stands for File Transfer Protocol. FTP is the best means for moving large files across the Internet. FTP is a client/server protocol that enables a user with an FTP client to log on to a remote machine, navigate the file system of that remote machine, and upload and download files from that machine.

There are two basic types of FTP on the Internet, anonymous ftp and private ftp. With anonymous ftp, one logs in as user anonymous, giving one's email address as a password. With private FTP, one logs in with the username and password one has established on that particular system. You are logged into your home directory, with all the file permissions you would normally have there.

12.16 HTTP

HyperText Transfer Protocol, the underlying protocol used by the World Wide Web. HTTP defines how messages are formatted and transmitted, and what action Web servers and browsers should take in response to various commands. For example, when you enter a URL in your browser, this actually sends an HTTP command to the Web server directing it to fetch and transmit the requested Web page.

12.17 POP3

POP3 (Post Office Protocol 3) is the most recent version of a standard protocol for receiving e-mail. POP3 is a client/server protocol in which e-mail is received and held for you by your Internet server. Periodically, you (or your client e-mail receiver) check your mail-box on the server and download any mail. It's also built into the Netscape and Microsoft Internet Explorer browsers.

12.18 IPv6

Adiscon Products officially support IPv6. The IPv6 support is introduced with the following versions:

MonitorWare Agent 8.0
WinSyslog 11.0

EventReporter 12.0

Support for IPv6 is available in all network related facilities of the engine. All network related actions will automatically detect IPv6 and IPv4 target addresses if configured. You can also use DNS resolution to resolve valid IPv6 addresses. Network related Services can either use IPv4 or IPv6 as internet protocol. In order to support both protocols, you will need to create two services. The only exception is the RELP Listener, which uses IPv4 and IPv6 automatically if available.

12.19 IMAP

Internet Message Access Protocol. A protocol allowing a client to access and manipulate electronic mail messages on a server. It permits manipulation of remote message folders (mailboxes), in a way that is functionally equivalent to local mailboxes.

IMAP includes operations for creating, deleting, and renaming mailboxes; checking for new messages; permanently removing messages; searching; and selective fetching of message attributes, texts, and portions thereof. It does not specify a means of posting mail; this function is handled by a mail transfer protocol such as SMTP.

Index

- I -

IP Address 79
IPv6 285

- M -

MSQueue 228

- N -

NT Services Monitor 180

- S -

Source System (IP) 161
SQL Statement Type 207

Endnotes 2... (after index)

Back Cover