

**MLINKADM-19.4**

**M-Link Administration Guide**

**Isode**

# Table of Contents

<b>Chapter 1</b>	<b>M-Link: Getting Started.....</b>	<b>1</b>
	This chapter discusses M-Link Server creation.	
<b>Chapter 2</b>	<b>M-Link Products.....</b>	<b>6</b>
	This chapter describes the different M-Link products.	
<b>Chapter 3</b>	<b>M-Link MU Gateway.....</b>	<b>8</b>
	This chapter introduces M-Link MU Gateway.	
<b>Chapter 4</b>	<b>Peer Controls and Links.....</b>	<b>12</b>
	This chapter introduces the M-Link Edge Server product, explains what XMPP trunking is, and how M-Link Peer Controls can be used to support it. It also shows the use of <i>links</i> to support operation with XML Guards.	
<b>Chapter 5</b>	<b>TLS Configuration.....</b>	<b>21</b>
	This chapter describes how M-Link is configured to use TLS, covering configuration of certificates/keys used by the local server, and trust of certificates used by other servers.	
<b>Chapter 6</b>	<b>Security Labels, Clearances and Policies.....</b>	<b>25</b>
	This chapter describes how Security Labels, Clearances and Policies are configured and used within M-Link.	
<b>Chapter 7</b>	<b>M-Link Configuration and Management.....</b>	<b>34</b>
	This chapter introduces the M-Link Administration Interface, used for configuration and management of the M-Link server.	
<b>Chapter 8</b>	<b>Instant Messaging Domain and User Configuration.....</b>	<b>43</b>
	This chapter describes how Instant Messaging domains and users are configured within M-Link.	
<b>Chapter 9</b>	<b>HTTP File Upload.....</b>	<b>46</b>
	This chapter describes the HTTP File Upload feature.	
<b>Chapter 10</b>	<b>MUC Room Creation and Administration.....</b>	<b>47</b>
	This chapter introduces the <i>MUC</i> room concept and covers how to create and administer <i>MUC</i> rooms.	
<b>Chapter 11</b>	<b>M-Link Archive Server.....</b>	<b>53</b>
	This chapter describes the M-Link archiving capabilities.	
<b>Chapter 12</b>	<b>Monitoring.....</b>	<b>55</b>
	This chapter introduces the features available for monitoring the state of the server.	
<b>Chapter 13</b>	<b>HTTP API.....</b>	<b>58</b>
	This chapter covers the APIs available for programmatic configuration of M-Link.	
<b>Appendix A</b>	<b>Integration with Isode Icon-Topo.....</b>	<b>75</b>
	This appendix describes how to integrate M-Link MU Gateway with Isode Icon-Topo.	
<b>Appendix B</b>	<b>Pre-defined Transformations.....</b>	<b>77</b>
	This appendix describes the <a href="#">Transformations</a> that are included with M-Link.	

<b>Appendix C Standards Supported by M-Link.....</b>	<b>79</b>
An overview of Open Standards M-Link products conform to.	
<b>Appendix D Glossary.....</b>	<b>82</b>
<b>Appendix E References.....</b>	<b>84</b>

**Isode** and Isode are trade and service marks of Isode Limited.

All products and services mentioned in this document are identified by the trademarks or service marks of their respective companies or organizations, and Isode Limited disclaims any responsibility for specifying which marks are owned by which companies or organizations.

Isode software is © copyright Isode Limited 2002-2023, all rights reserved.

Isode software is a compilation of software of which Isode Limited is either the copyright holder or licensee.

Acquisition and use of this software and related materials for any purpose requires a written licence agreement from Isode Limited, or a written licence from an organization licensed by Isode Limited to grant such a licence.

This manual is © copyright Isode Limited 2023, all rights reserved.

---

## 1 Software version

This guide is published in support of M-Link R19.4. It may also be pertinent to later releases. Please consult the release notes for further details.

---

## 2 Readership

This guide is intended for administrators who plan to configure and manage XMPP services using M-Link R19.4.

---

## 3 Typographical conventions

The text of this manual uses different typefaces to identify different types of objects, such as file names and input to the system. The typeface conventions are shown in the table below.

Object	Example
File and directory names	<i>/var/isode/log</i>
Program and macro names	<code>isode.xmppd</code>
Input to the system	<code>cd newdir</code>
Cross references	see <a href="#">Section 4, “File system place holders”</a>
Additional information to note, or a warning that the system could be damaged by certain actions.	Notes are additional information; cautions are warnings.

---

## 4 File system place holders

Where directory names are given in the text, they are often place holders for the names of actual directories where particular files are stored. The actual directory names used depend on how the software is built and installed. Some of these directories can be changed by configuration.

Certain configuration files are searched for first in  $\$(ETCDIR)$  and then  $\$(SHAREDIR)$ , so local copies can override shared information.

The actual directories vary, depending on whether the platform is Windows or UNIX.

Name	Place holder for the directory used to store...	Windows (default)	UNIX
$\$(ETCDIR)$	System-specific configuration files.	<i>C:\Isode\etc</i>	<i>/etc/isode</i>
$\$(SHAREDIR)$	Configuration files that may be shared between systems.	<i>C:\Program Files\Isode\M-Link\share</i>	<i>/opt/isode/mlink/share</i>

Name	Place holder for the directory used to store...	Windows (default)	UNIX
<code>\$(VARDIR)</code>	Storing local data.	<code>C:\Isode\M-Link</code>	<code>/var/isode/mlink</code>
<code>\$(LOGDIR)</code>	Log files.	<code>C:\Isode\M-Link\log</code>	<code>/var/isode/mlink/log</code>

---

## 5 Support queries and bug reporting

A number of email addresses are available for contacting Isode. Please use the address relevant to the content of your message.

- For all account-related inquiries and issues: [customer-service@isode.com](mailto:customer-service@isode.com). If customers are unsure of which list to use then they should send to this list. The list is monitored daily, and all messages will be responded to.
- For all licensing related issues: [license@isode.com](mailto:license@isode.com).
- For all technical inquiries and problem reports, including documentation issues from customers with support contracts: [support@isode.com](mailto:support@isode.com). Customers should include relevant contact details in initial calls to speed processing. Messages which are continuations of an existing call should include the call ID in the subject line. Customers without support contracts should not use this address.
- For all sales inquiries and similar communication: [sales@isode.com](mailto:sales@isode.com).

Bug reports on software releases are welcomed. These may be sent by any means, but electronic mail to the support address listed above is preferred. Please send proposed fixes with the reports if possible. Any reports will be acknowledged, but further action is not guaranteed. Any changes resulting from bug reports may be included in future releases.

Isode sends release announcements and other information to the Isode News email list, which can be subscribed to from the address: <http://www.isode.com/company/subscribe.html>

---

## 6 Export Controls

M-Link uses TLS (Transport Layer Security) to encrypt data in transit. This means that M-Link is subject to UK Export Controls.

For some countries (at the time of shipping this release, these comprise all EU countries, United States of America, Canada, Australia, New Zealand, Switzerland, Norway, Japan), these Export Controls can be handled by administrative process as part of evaluation or purchase.

For other countries, a special Export License is required. This can be applied for only in context of a purchase order from Isode.

The TLS features of M-Link are enabled by a TLS Product Activation feature. This feature may be turned off, and without this TLS feature M-Link is not export-controlled. This can be helpful to support evaluation in countries that need a special export license.

XMPP servers are generally deployed with TLS for security reasons and Isode strongly recommends that all operational deployments use the export-controlled TLS feature.

You must ensure that you comply with these Export Controls where applicable, i.e. if you are licensing or re-selling M-Link. M-Link is subject to an Isode license agreement and your attention is also called to the export terms of the license agreement.

# Chapter 1 M-Link: Getting Started

This chapter discusses M-Link Server creation.

This chapter contains sections on the following topics:

- [Section 1.1, “Starting and Stopping M-Link Server”](#)
- [Section 1.2, “Creating an Initial Administrative User”](#)
- [Section 1.3, “Installing an Isode Activation Key”](#)
- [Section 1.4, “M-Link Server Runtime User”](#)
- [Section 1.5, “Adding Initial Configuration”](#)
- [Section 1.6, “DNS Configuration”](#)
- [Section 1.7, “Firewall Considerations”](#)

---

## 1.1 Starting and Stopping M-Link Server

Once installed, the M-Link Server will start automatically as a system service, and thereafter the service can be controlled with the system-provided mechanisms.

On first start (and subsequently), M-Link will start with its configuration interface on port 5221, ready to be activated and configured - this is accessible at `https://localhost:5221` on the server, or at `https://hostaddress:5221` from other machines. It will initially generate a self-signed certificate - both the certificate and the default port can be changed later through the configuration interface. The configuration interface can be accessed from a web browser running on a different machine from M-Link by substituting `localhost` with the hostname or IP of the server machine. As the initial certificate is self-signed, it will need to be manually trusted before a browser will accept it - usually by clicking through a security warning.

---

**Note:** Access to the administration interface might be restricted by firewall policies. If the interface is unreachable, please check that there is a rule that allows access to port 5221.

---

---

## 1.2 Creating an Initial Administrative User

M-Link Server can be configured by administrators. Administrators are users distinct from the users of the XMPP service. An initial administrator account is created as the first step through the web configuration interface - this user is stored in M-Link's configuration (the password is stored with a strong one-way hashing function, and cannot be recovered). Currently M-Link is limited to a single administrator account.

---

## 1.3 Installing an Isode Activation Key

In order to use an M-Link Server instance on a host system, an activation key file is required. This file needs to be installed through M-Link's web browser-based configuration interface. The initial configuration process will lead you through generation of an Activation Request and the subsequent installation of the Activation Key supplied by Isode in response to receipt of the Activation Request.

Questions regarding licensing should be directed to [support@isode.com](mailto:support@isode.com).

---

## 1.4 M-Link Server Runtime User

When running on Unix, M-Link Server runs as an unprivileged user, created during package install, and does not need to be started as 'root'.

On Windows, the M-Link Server runs as Windows service(s) under the LocalSystem account.

---

## 1.5 Adding Initial Configuration

Once you're logged into the M-Link Administration Interface and M-Link is activated you're ready to add initial configuration. Depending on the type of M-Link server that is being configured (and has been activated) you will need to either add domains for M-Link to host or add Peer Controls to allow routing between remote servers. Details of the configuration of XMPP Trunking and Peer Routing can be found in the [Chapter 4, Peer Controls and Links](#) chapter.

---

## 1.6 DNS Configuration

XMPP clients and servers typically use the *Domain Name System* (DNS) to determine the IP addresses and TCP ports of XMPP servers they need to connect to. For XMPP Client-to-Server (C2S) the following discovery approach is typically used.

Clients first look for SRV resource records for the XMPP domain they wish to connect to. Each SRV resource record for a domain provides a location (a domain name and a port) for the service (e.g., XMPP C2S). Multiple SRV resource records may be published for each service. In addition to the service location, each SRV resource record has a priority and a weight. Preference of two SRV resource records is given to one with the lower priority value or, in the case the equal priority, to the resource record with the higher weight. The location's domain name is resolved to a set of IP addresses using normal A and AAAA lookup. Depending on the client's configuration, if both IPv4 and IPv6 are supported, preference may be given to one or the other. Otherwise, IP addresses are typically tried in the order provided by the Domain Name System, which may or may not change with each lookup.

If SRV resource records are not published for the domain, clients and servers will resolve the domain name to a set of IP addresses and connect to each using the default port for the protocol. This is accomplished by looking for A and AAAA resource records for IPv4 and IPv6 addresses respectively, possibly with CNAME resolution when necessary. As above, this resolves to a set of IP addresses. Depending on the client's configuration, if both IPv4 and IPv6 are supported, preference may be given to one or the other. Otherwise, IP addresses are typically tried in the order provided by the Domain Name System, which may or may not change with each lookup.

In deployments where DNS is not available, clients will often use other host-to-IP address lookup systems, such as those provided by a hosts file. These typically behave similarly to DNS with only A and AAAA resource records. Configuration of alternative host-to-IP address lookup systems is not detailed in this guide.

To illustrate proper configuration of DNS, consider an XMPP service consisting of a two node cluster providing Instant Messaging service for the domain `example.com`, where the cluster nodes are named `node1.example.com` and `node2.example.com`. The nodes have IPv4 addresses `192.0.2.1`, `192.0.2.2` respectively. The nodes have IPv6 addresses `2001:DB8::1`, `2001:DB8::2` respectively. We assume the XMPP C2S protocol is available on its default port, 5222. Additionally we'll assume this IM service is not co-located with other `example.com` services such as WWW.

Note that these examples utilize the Domain Name System *zone file* format specified in [RFC 1035](#) (§ 5) and [RFC 1034](#) (§ 3.6.1). Furthermore, the examples assume placement in the zone `example.com`:

```
$ORIGIN example.com.
```

Each node should have an appropriate resource records for addressing purposes:

```
node1 IN A 192.0.2.1
node1 IN AAAA 2001:DB8::1
node2 IN A 192.0.2.2
node2 IN AAAA 2001:DB8::2
```

SRV resource records can then be provided. Here we've given equal preference to each node in our cluster.

```
_xmpp-client._tcp IN SRV 1 1 5222 node1
_xmpp-client._tcp IN SRV 1 1 5222 node2
```

While the domain naming each node allow for users to manually configure their client to connect to any one of the two nodes, it is often desirable to provide a domain which they can use to have the client connect to any of the nodes. The following resource records creates `xmpp.example.com` for this purpose as well as other purposes discussed below.

```
xmpp IN A 192.0.2.1
xmpp IN AAAA 2001:DB8::1
xmpp IN A 192.0.2.2
xmpp IN AAAA 2001:DB8::2
```

If the DNS servers providing the `example.com` zone support CNAME round robins, the `xmpp.example.com` A and AAAA resource records can be replaced with:

```
xmpp IN CNAME node1
xmpp IN CNAME node2
```

If a WWW service is available at the URL `https://example.com` or `http://example.com`, host-meta data for BOSH clients can be published using either XML or JSON.

XML host-meta example:

```
<?xml version='1.0' encoding=utf-9'?>
<XRD xmlns='http://docs.oasis-open.org/ns/xri/xrd-1.0'>
  ...
  <Link rel="urn:xmpp:alt-connections:xbosh"
        href="https://xmpp.example.com:5280/http-bind/" />
  ...
</XRD>
```

JSON host-meta example:

```
{
  ...
  "links": [
    ...
    {
      "rel": "urn:xmpp:alt-connections:xbosh",
      "href": "https://xmpp.example.com:5280/http-bind/"
    },
    ...
  ]
}
```

For XMPP S2S the following discovery approach is similar to XMPP C2S. For the IM domain, the following would be published assuming S2S is available on the default port, 5269:

```
_xmpp-server._tcp IN SRV 1 1 5269 node1
_xmpp-server._tcp IN SRV 1 1 5269 node2
```

Additionally, DNS SRV records should be created for each other XMPP domain accessible by other XMPP servers. For instance, assuming a Multi-User Chat service domain of *muc.example.com* and Publish-Subscribe service at *pubsub.example.com*, the following additional resource records should be published:

```
_xmpp-server._tcp.muc IN SRV 1 1 5269 node1
_xmpp-server._tcp.muc IN SRV 1 1 5269 node2

_xmpp-server._tcp.pubsub IN SRV 1 1 5269 node1
_xmpp-server._tcp.pubsub IN SRV 1 1 5269 node2
```

## 1.7 Firewall Considerations

By default, M-Link will listen on a select number of TCP ports. All ports are configurable through the M-Link Administration Interface when the “Customise Ports” is enabled. Depending on the environment, firewall rules will need to be created to block or expose these ports. Isode recommends that ports are only exposed as-needed to provide the service, as in the following table.

**Table 1.1. TCP ports used by M-Link**

Port	Recommended exposure	Description
5001	Users if required	The port used for HTTP File Upload ( <a href="#">XEP-0363</a> ) by M-Link User Server and M-Link MU Server.
5221	Administration only	Port on which the M-Link Administration Interface is reachable. This port should not be exposed generally, but only to machines that will be used to administer the service.
5222	Users	The port used for XMPP Client-to-Server by M-Link User Server and M-Link MU Server.
5224	Localhost only	Port used for communication between various subservices of M-Link, such as the M-Link Archive Server or the BOSH service.

<b>Port</b>	<b>Recommended exposure</b>	<b>Description</b>
5269	Federating servers	The port for used for XMPP Server-to-Server by M-Link Edge Server, M-Link User Server and M-Link MU Gateway .
5280	Users if required	The port for used for BOSH by M-Link User Server and M-Link MU Server.

# Chapter 2 M-Link Products

This chapter describes the different M-Link products.

M-Link is a single set of binaries which provides different products by enabling combinations of functionality (controlled via the Activation Key system). As a result, some chapters of this manual are not applicable to all products, as the features described will not be available.

The M-Link product set is:

---

## 2.1 M-Link User Server

This is the core Instant Messaging server, supporting 1:1 chat, Multi-User Chat (MUC), Personal Eventing and other XMPP services. See chapters:

- [Getting Started](#)
- [TLS Configuration](#)
- [Security Labels, Clearances and Policies](#)
- [M-Link Configuration and Management](#)
- [Instant Messaging Domain and User Configuration](#)
- [MUC Room Creation and Administration](#)
- [M-Link Archive Server](#)
- [Monitoring](#)
- [HTTP API](#)

---

## 2.2 M-Link Edge Server

This is used to provide an XMPP Boundary Guard service to protect organizational boundaries and provide Cross Domain services. See chapters:

- [Getting Started](#)
- [Peer Controls and Links](#)
- [TLS Configuration](#)
- [Security Labels, Clearances and Policies](#)
- [M-Link Configuration and Management](#)
- [M-Link Archive Server](#)
- [Monitoring](#)
- [HTTP API](#)

---

## 2.3 M-Link MU Server

This supports local clients using the standard XMPP Client/Server protocol, and communicates with other M-Link MU Gateway and M-Link MU Server instances using

the XEP-0361: Zero Handshake Server to Server Protocol and additionally supports "XEP-0365: Server to Server communication over STANAG 5066 ARQ", used in conjunction with XEP-0361, to support XMPP communication over HF Radio. See chapters:

- [Getting Started](#)
- [M-Link MU Gateway](#)
- [Peer Controls and Links](#)
- [TLS Configuration](#)
- [Security Labels, Clearances and Policies](#)
- [M-Link Configuration and Management](#)
- [Instant Messaging Domain and User Configuration](#)
- [MUC Room Creation and Administration](#)
- [M-Link Archive Server](#)
- [Monitoring](#)
- [HTTP API](#)

---

## 2.4 M-Link MU Gateway

Used to exchange messages between systems running on high quality links and those running over constrained networks, the M-Link MU Gateway communicates with M-Link MU Gateway and M-Link MU Server instances using XEP-0361 (and XEP-0365 for HF Radio) and with XMPP servers over normal quality links using the standard XMPP Server/Server protocol, allowing units communicating over constrained links to be integrated with existing standard XMPP installations. See chapters:

- [Getting Started](#)
- [M-Link MU Gateway](#)
- [Peer Controls and Links](#)
- [TLS Configuration](#)
- [Security Labels, Clearances and Policies](#)
- [M-Link Configuration and Management](#)
- [M-Link Archive Server](#)
- [Monitoring](#)
- [HTTP API](#)

# Chapter 3 M-Link MU Gateway

This chapter introduces M-Link MU Gateway.

---

## 3.1 Introduction

Many military deployments will use network links with poor and variable quality. High latency is a particular problem for standard XMPP connectivity, especially with handshaking in session establishment. M-Link MU Gateway is specifically designed to overcome the issues of working in such environments.

Specific capabilities for operating over constrained networks include:

- Stream Compression
- Traffic Filtering, including Presence Stripping
- Optimised Server-to-Server Protocol
- Support for HF Radio (STANAG 5066)
- Connection fall-back and fall-forward

### 3.1.1 Stream Compression

Minimising the amount of data transferred is important, and so use of compression is desirable. Compression can be enabled for both X2X and SLEP links. This compression uses the DEFLATE algorithm, which references previous data in the stream and thus becomes more effective for larger data sets or longer-lived streams.

### 3.1.2 Traffic Filtering

Traffic filtering removes data, and so modifies the service provided to the end user.

Filtering options available are:

- Removal of selected types of message (or other stanza)
- Removal of selected elements from messages (message folding)
- Removal of selected elements from presence stanzas (presence folding)

Traffic filtering is implemented via XSLT transformations which are applied to stanzas, and are configurable for individual Peer Controls, as described in [Section 4.5, “Transformations”](#).

### 3.1.3 Optimised Server-to-Server Protocol

M-Link MU Gateway uses the *Zero Handshake Server to Server Protocol (XEP-0361)* operating over TCP to reduce data volumes and remove XMPP-level stream establishment handshaking, improving performance over constrained links. This uses three approaches:

- Use of a single (bi-directional) stream.
- Configuration of options (using peering controls) at both ends of the connection to avoid the need for negotiation during connection establishment.
- Full pipelining of stanzas, so that message stanzas are sent immediately after connection establishment.

The base transport is TCP. There will be a single TCP handshake to establish the TCP connection, and then data can flow without further handshakes. TCP optimizes use of the available link bandwidth.

While Zero Handshake Server to Server Protocol ([XEP-0361](#)) can be used efficiently over wired connections, it is of particular interest over satellite connectivity, where latency can significantly impair other server-to-server connections.

[XEP-0361](#) may be operated over TLS. M-Link supports this to provide data confidentiality with an option to use peer authentication using X.509 strong authentication. With many constrained networks these services will be provided at the data link or network layer, and there is no functional requirement to provide them at the application layer. TLS adds some protocol overhead, and the handshaking may add significant latency.

### 3.1.4 Support for HF Radio

M-Link MU Gateway uses the SIS Layer Extension Protocol protocol with Server to Server communication over STANAG 5066 ARQ ([XEP-0365](#)) to enable communication over STANAG 5066 services, such as HF radio via the STANAG 5066 Subnet Interface Service. The use of SLEP provides a reliable bidirectional stream service which can support compression.

### 3.1.5 Connection fall-back and fall-forward

M-Link MU Gateway allows the configuration of fall-back links in priority order, e.g. falling back to use of STANAG 5066 links when satellite connectivity fails, or between different satellite systems. When a fall-back has occurred, M-Link will attempt to reestablish the primary connectivity and fall-forward to the higher priority link when that one becomes available again.

---

## 3.2 Configuring M-Link MU Gateway

M-Link MU Gateway is an M-Link server with no local users that is used to enable communication between groups of XMPP Servers over slow network links, enabling the optimization of the server-to-server protocol over the slow links.

### 3.2.1 Connectivity

A typical M-Link MU Gateway configuration would prefer use of faster and lower-latency SATCOM links when available, with slower but more reliable HF radio links being used when the SATCOM links are unavailable. This could be achieved by configuring Peer Controls for the appropriate domains with SATCOM (i.e. [XEP-0361](#)) links as the primary links, and the HF (SLEP) links as the fallback for the SATCOM links.

#### 3.2.1.1 Link Fallback

A link may be configured with a fallback link, to be used when the parent link cannot be connected. This allows a chain of links (in order of preference) to be constructed. The appropriate Peer Control is configured to use the primary link, i.e. the link at the top of the chain.

When a connection to the domain handled by the Peer Control is required, M-Link MU Gateway will attempt to connect the Peer Control's primary link. Connection failures are handled by retrying at increasingly long intervals, until a limit is reached at which the link is considered failed. At this point M-Link MU Gateway will start connection attempts for the fallback link, applying the same retry logic. This proceeds until either

a connection is established using one of the links in the chain, or all links have been attempted and failed.

### 3.2.1.2 Link Fall-forward

It may be that a fallback link is in use by a Peer Control, but a more-preferred link for that Peer Control could now be established. For example, a SATCOM link is temporarily unavailable, leading to a less-preferred HF link being established instead, but then the SATCOM link becomes available again. When a less-preferred link is connected for a Peer Control, M-Link MU Gateway will regularly make connection attempts for the more-preferred link. If the more-preferred link is successfully connected, new outbound stanzas will be transferred using it. The less-preferred link will be closed after an idle period.

### 3.2.1.3 Timer Configuration

Connection establishment involves multiple steps. Timeouts and retry limits are configurable for many of these steps. Modification of the default settings may be required to optimize an M-Link MU Gateway installation.

The following options may be modified:

#### Remote Session Connection Timeout

This configures the timeout which is applied to the initial TCP-level connect when establishing a session. This needs to be set to a value greater than the expected connection establishment latency.

#### Remote Session Connection Retry Base Delay

This configures the basic delay applied before a failed connection attempt is retried. A simple algorithm is used which increases the delay by multiplying the previous delay interval by 1.5 each time a failure occurs. This is unlikely to need modifying.

#### Remote Session Connection Retry Delay Limit

This configures the maximum delay interval between connection retries. Once the delay interval reaches this value (through multiplication as described above), it will not increase in length. This is unlikely to need modifying.

#### Remote Session Connection Retry Duration

This configures the maximum length of time for which M-Link MU Gateway will perform connection retries before treating the connection as failed. This affects how quickly a fall-back link will be established, so should be set to be the length of time it is desired for connections to a primary link to be attempted and fail before falling back to a secondary link.

#### Remote Session Authentication Timeout

Once a session has been opened to a Peer, authentication may be required. If this has not completed within this time, the session will be closed. This must be at least as long as authentication handshakes will take for all session types. If using TLS on a Link, it must allow for the latency to perform the several round-trips to establish TLS.

#### Remote Session Dialback Timeout

This configures the maximum time which a remote server may take to verify its own dialback authentication request before we fail it. This is unlikely to need modification as dialback shouldn't be used over constrained bandwidth links.

#### STANAG S5066 Maintenance Timer

SLEP sessions are multiplexed over a single TCP connection between the M-Link MU Gateway and a STANAG 5066 server. This timer controls how often a M-Link MU Gateway will attempt to reestablish this connection. It is unlikely that this will need modification.

#### Remote Session Idle Timer

A session which has been idle (i.e. no stanza has been sent or received) for longer than this timer is closed. For sessions that are expensive to establish, but cheap

to maintain (especially those using TLS over constrained links) this should be moderately high (e.g. the default). For sessions that should be closed quickly when they are unused (such as SLEP links) this should be set lower.

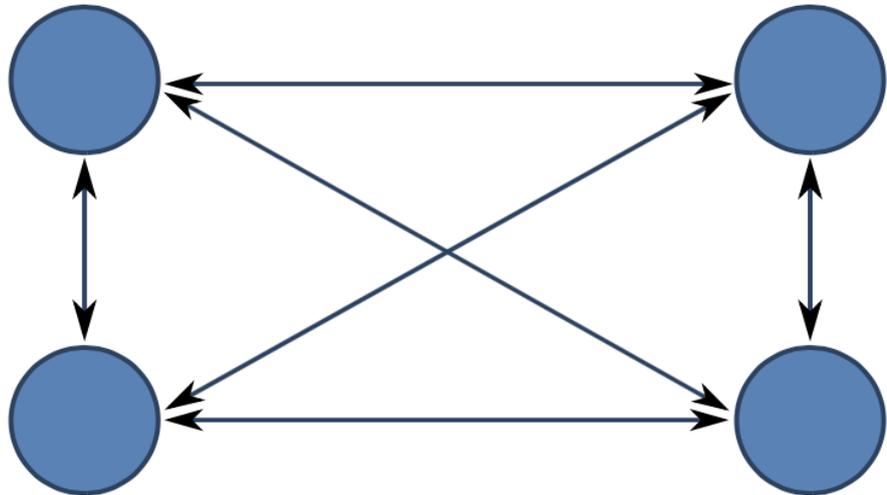
#### Whitespace Ping Time

This configures the length of time for which a session can be idle outbound (i.e. no stanza sent) before a single whitespace character is sent to check that the underlying TCP connection is still active. It is unlikely that this will need modifying.

## Chapter 4 Peer Controls and Links

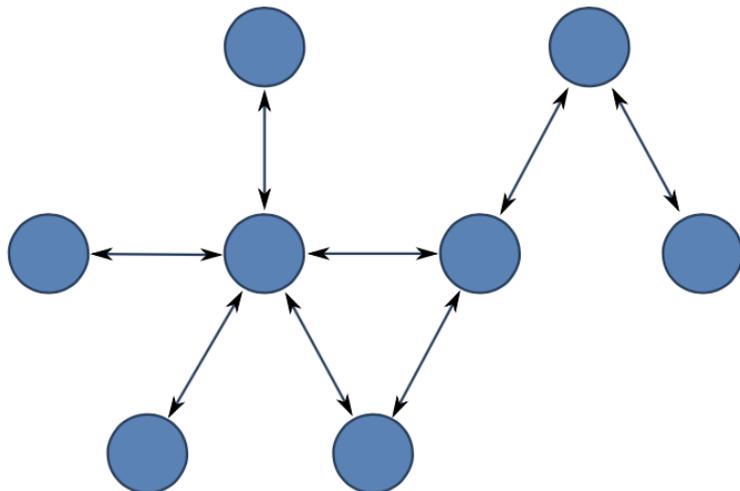
This chapter introduces the M-Link Edge Server product, explains what XMPP trunking is, and how M-Link Peer Controls can be used to support it. It also shows the use of *links* to support operation with XML Guards.

**Figure 4.1. Standard XMPP Server Configuration**



In a standard XMPP configuration, XMPP servers are fully interconnected, as shown in [Figure 4.1, “Standard XMPP Server Configuration”](#). This model works well for open XMPP deployments on the Internet. However, it does not work so well for cross domain, organizational boundary and constrained link scenarios.

**Figure 4.2. XMPP Trunking**



The model of XMPP Trunking, shown in [Figure 4.2, “XMPP Trunking”](#) extends the standard XMPP model to allow for configurations where servers are not fully connected. This structure enables XMPP messages to be switched through intermediate servers.

By default, when an XMPP server initiates a connection to another XMPP server (e.g., to other.organisation.example.com) it will look up the domain in the *Domain Name*

*System* (DNS), first with SRV records and falling back to A/AAAA records, and connect using the IP address determined from this lookup. This leads to the fully connected approach.

Peer Controls are the mechanism used by M-Link to enable configuration of an XMPP Trunking architecture. They are checked prior to the standard DNS lookup, and provide M-Link with information on how to connect to a specific domain. Each peering control relates to a specific remote domain (e.g. `other.organisation.example.com`), to a collection of subdomains (e.g. all subdomains of "isode.com"), or a default catch-all. A peering control can direct the connection to a list of hosts and/or IP address or specialised connection mechanism configured as part of the peering control. Peer controls can be used to reject the traffic to another site entirely, require particular security and authentication, and in some M-Link products restrict it to certain types and control the network connections used for the traffic.

Further information on XMPP Trunking and use of M-Link Peer Controls is provided in the Isode white paper [Providing XMPP Trunking with M-Link Peer Controls](https://www.isode.com/whitepapers/xmpp-trunking.html) [https://www.isode.com/whitepapers/xmpp-trunking.html].

---

## 4.1 M-Link Edge Server

M-Link Edge Server is an M-Link server with no local users that is used in boundary and cross-domain configurations to check and potentially modify traffic. M-Link Edge Server will often be used in conjunction with an XML Guard to provide cross-domain protection and may use [XEP-0361](#) or [GCXP](#) to communicate with the XML Guard.

M-Link and M-Link Edge Server are both provided by the same underlying technical product, but are sold as different products. Note that [XEP-0361](#) may not be used with a standard M-Link deployment that supports local users.

---

## 4.2 Peer Controls

Peer Controls define the rules for communicating with remote domains, and are checked prior to DNS being used. A Peer Control can be defined using one of three Matching Rules:

### Domain

A single specific domain. This is used to force routing for a specific domain.

### Subdomains

All subdomains of the domain, but not the domain itself. If this Matching Rule is selected for a Peer Control for `example.com`, domains that will match include `pubsub.example.com` and `chat.private.example.com`.

### Domains and Subdomains

The domain itself, and all its subdomains. If this Matching Rule is selected for a Peer Control for `example.com`, domains that will match include `example.com` and `pubsub.example.com`.

When multiple Peer Controls match a certain remote domain, the longest match is the one that will take effect. When no match is found for the domain, the *default* Peer Control will be applied.

This combination provides full flexibility to configure arbitrary XMPP trunking configurations.

When a peer control is in place, it is possible to configure specific actions for the peer:

- [Chapter 6, Security Labels, Clearances and Policies](#) checking against a configured security clearance for the peer.
- [Section 4.3, “Peer Authentication”](#) options.
- [Section 4.4, “Relay zones”](#), which allows for relaying stanzas between servers.
- Traffic filtering using [Section 4.5, “Transformations”](#), which enables removal or modification of traffic to and from the peer.

Traffic to a remote domain will by default use S2S connections which are negotiated between the servers. If the option *Use Specialised Connection Mechanism* is enabled, communication will be sent and received using the specified link (see [Section 4.6, “Links”](#)).

SRV lookups for S2S can be overridden by enabling the *Override DNS* option and specifying values for the *Connect Host* and *Connect Port* options. *Connect Host* can either be an IP address, or a hostname. Multiple *Connect Host* and *Connect Port* pairs may be specified, with connection attempts being made in the order in which the pairs are specified.

If the option *Use Specialised Connection Mechanism* is enabled, communication will be sent and received using the specified link.

---

**Note:** Changing options on Peer Controls or Links may result in reestablishment of corresponding sessions and pending connections.

---

## 4.3 Peer Authentication

Authentication of remote servers is usually based on the addressing of the XMPP data that are being sent, not on the DNS name of the server hosting the remote service (although these are sometimes the same). There are two mechanisms for authenticating server to server (S2S) connections: Dialback (*XEP-0220*) and TLS authentication (SASL EXTERNAL).

### 4.3.1 TLS Authentication

TLS Authentication is based on the certificate presented by a remote server during TLS negotiation, and could take four forms.

- Normal. Absent further configuration, TLS Authentication is performed by checking the subjects of the certificate provided by the remote server, and if they match the server address (or Remote Host if a Link is in use) of the XMPP data that are being transmitted, checking to see if the certificate was issued by a trusted root certificate. For a detailed discussion, see [Section 5.2, “Certificate Verification”](#) and [Section 5.3, “Trust Anchors”](#)
- Connect Host. If `Peer Connections` is enabled through "Override DNS" and normal TLS Authentication is attempted and fails, M-Link will perform the same authentication checks:
  - Outbound. Using the hostname from the `Peer Connections` list which has resulted in the successful connection.
  - Inbound. Using each hostname from the `Peer Connections` list in turn.
- Remote Host. If and X2X or GCXP link is being used, M-Link will use the Remote Host value of the Link as the subject to check.
- Pinned Certificates. If a pinned peer or link certificate is set then TLS Authentication is performed only by comparing the certificate provided by the remote server to the one in this option.

### 4.3.2 Dialback

Dialback uses checks against the entries in DNS to authenticate a remote server (connections that use dialback for authentication may also use unauthenticated TLS, and if TLS is set to `Required`, must). This is generally considered a weaker form of authentication than TLS provides. Authentication must happen in both directions (that is: we must authenticate that the remote server is who they claim to be, and the remote server must authenticate that we are who we claim to be), and dialback is controlled independently for each direction. If TLS is set to any value other than `Require Authenticated TLS` we will be willing to authenticate a remote server's identity using dialback. If the `Perform Dialback When Offered` option is enabled then we will allow a remote server to authenticate our identity using dialback.

Where there is a choice of TLS Authentication or Dialback, M-Link will always prefer to use TLS Authentication.

---

## 4.4 Relay zones

Relay zones are the M-Link mechanism for allowing relaying of stanzas between servers in different logical groups (generally by bridging a network boundary of some sort, such that the servers would not be able to establish a connection otherwise). When receiving a stanza from a peer if the 'to' domain is not serviced by the local server a peer control matching the 'to' will be searched for and, if found and if the peer control of the 'to' gives a different relay zone from the peer control of the connection from which the stanza is received, it will be routed out to the target peer. This means that M-Link will not relay stanzas between servers within the same relay zone (because they should be communicating directly, or simply shouldn't be communicating).

---

## 4.5 Transformations

This section describes transformations that may be configured for a Peer Control.

Transformations can be applied to stanzas coming in from or going out to peers. Each Peer Control can have a number of these Transformations, which are executed in the order they are listed. One typical case where this may be useful is when network capacity is constrained: for example, on a slow link it may make sense to set up a transformation to strip all Chat State Notifications.

Transformations are selected from the Transformations object store. Each transformation is an *XSLT 1.0* document operating on a single stanza in the `jabber:server` namespace, wrapped in a minimal XMPP Server-to-Server stream element. The result of a transformation should also be a single stanza wrapped in a stream element.

M-Link ships with transformations for common operations, such as stripping XHTML-IM (*XEP-0071*) elements from messages, or normalizing XML for communications to M-Guard. These transformations are always present, and cannot be removed or renamed, however uploading variants of the built-in rules with modified behaviour is possible. A complete list of all transformations shipped with M-Link can be found in [Appendix B, Pre-defined Transformations](#).

The wizard which allows user-defined *XSLT 1.0* documents to be uploaded to the Transformation object store also allows simple transformations which strip namespace / element combinations to be created and uploaded.

### Example 4.1. A stanza wrapped for transformations

```
<stream:stream id='1' version='1.0' xmlns='jabber:server'
  xmlns:stream='http://etherx.jabber.org/streams'>
<message from='anne@example.com' to='bob@example.net' type='chat'>
  <body>Hello!</body>
  <active xmlns='http://jabber.org/protocol/chatstates' />
</message>
</stream:stream>
```

### Example 4.2. An XSLT that will strip Chat State Notifications from message stanzas

```
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:stream='http://etherx.jabber.org/streams'
  xmlns:xmpp='jabber:server'
  xmlns:xhtml-im='http://jabber.org/protocol/xhtml-im'
  xmlns:csn='http://jabber.org/protocol/chatstates'>

<xsl:output method="xml" omit-xml-declaration="yes" />

<xsl:template match="@*|node()">
  <xsl:copy><xsl:apply-templates select="@*|node()" />
</xsl:copy></xsl:template>

<xsl:template match="/stream:stream/xmpp:message/csn:active" />
<xsl:template match="/stream:stream/xmpp:message/csn:composing" />
<xsl:template match="/stream:stream/xmpp:message/csn:paused" />
<xsl:template match="/stream:stream/xmpp:message/csn:inactive" />
<xsl:template match="/stream:stream/xmpp:message/csn:gone" />
</xsl:stylesheet>
```

If the original stanza had children, but the resulting stanza no longer has any, the option *Allow Childless Message and Presence Stanzas Through* determines whether the stanza is let through. For example, if the stanza in example [Example 4.3, “A stanza with only Chat State Notifications contents”](#) is sent to a peer that has a Transformation rule that removes CSN (as shown in [Example 4.2, “An XSLT that will strip Chat State Notifications from message stanzas”](#) and the option *Allow Childless Message and Presence Stanzas Through* set to *Don't let stanzas through* or *Only allow presence stanzas through*, this stanza will not be sent to the recipient.

### Example 4.3. A stanza with only Chat State Notifications contents

```
<stream:stream id='1' version='1.0' xmlns='jabber:server'
  xmlns:stream='http://etherx.jabber.org/streams'>
<message from='anne@example.com' to='bob@example.net' type='chat'>
  <composing xmlns='http://jabber.org/protocol/chatstates' />
</message>
</stream:stream>
```

If the transformation returns an empty or invalid document, or a document that does not parse into a syntactically valid stanza, or the test described in the previous paragraph blocks the stanza, what happens is determined by the *Action on Blocked Stanza* option associated with the transformation. *Bounce* will send back an error stanza to the source of the stanza if appropriate for the stanza. *Drop* will cause the stanza to be discarded without notifications.

---

## 4.6 Links

A Link defines a specialised connection mechanism for federation. Configuration of connection mechanisms and authentication/security live within the Link configuration for all Peers that use links, rather than in the Peer Control. M-Link currently supports these types of links:

### X2X

*Zero Handshake Server to Server Protocol (XEP-0361)* operating over TCP. This protocol is designed to support constrained bandwidth links with high latency, where standard S2S will lead to performance problems.

### GCXP

*Guard Content Exchange Protocol*, operating over TCP. This protocol is used for communications between M-Link and M-Guard.

By default, GCXP links are configured to send and expect to receive GCXP responses, and M-Guard should be configured to let GCXP responses through. In setups where M-Guard cannot be configured to let GCXP responses through, the option *Send GCXP Responses* must be disabled as well. GCXP responses are used to provide acknowledgement of stanzas delivered between the M-Link Edge Server servers on either side of the Guard, allowing stanza delivery to be re-attempted in the case of a connection or system failure of the Guard or other M-Link Edge Server. They are also used by M-Guard to alert M-Link Edge Server when a stanza has been rejected by the Guard, allowing M-Link Edge Server to 'bounce' the stanza. GCXP responses contain no data other than the ID of the GCXP message to which they refer; connections to the Guard remain unidirectional with respect to the flow of XMPP data.

---

**Note:** Peers using links that communicate using this protocol will generally require the *Normalize XML for M-Guard* and *Strip Legacy Delayed Delivery Transformations*.

---

### SLEP

Connections to other M-Link servers over HF Radio connections using *SIS Layer Extension Protocol (SLEP)* connections to Icon-5066 servers. More details on the configuration of this type of link can be found in [Section 4.6.4, "SLEP Links"](#).

---

**Note:** SLEP links are only available to M-Link MU Gateway and M-Link MU Server servers.

---

### 4.6.1 Disabling a Link

The configuration for each Link includes an `Enabled` item, defaulting to `true`. If set `false`, this Link will be ignored when attempting to open a new outbound connection. If a fallback link is configured, this will be used immediately instead, otherwise a connection failure will occur. The item has no effect on inbound connections or on outbound connections which are already established.

### 4.6.2 Connection parameters

All links can take a fallback link:

#### Fallback Link

This option points to the next link that should be tried when an attempt to establish a connection fails. When this parameter is unset, no further links will be tried. This

can be used to create a list of links that will be tried in order when connecting to a peer.

---

**Note:** Peer Controls using a link with fallbacks must always point to the first link in the chain.

---

X2X and SLEP links support Stream Management ([XEP-0198](#)) acknowledgements (GCXP has its own mechanism for reliability in the form of *GCXP Responses*):

#### Enable XEP-0198

When this option is enabled, the link will use Stream Management ([XEP-0198](#)) to ensure stanzas sent are received by the remote server. This will also enable acknowledgements of stanzas received from the remote server.

#### XEP-0198 Version

This option selects which version of [XEP-0198](#) will be enabled on this link. By default, this will use the most recent version of the standard as supported by M-Link. When interoperating over a Stream Management enabled link with a server running a version of M-Link older than R19.0, this option should be set to *Legacy XEP-0198*.

All X2X and GCXP links support the following parameters:

#### Local Port

TCP port the link will listen on for new connections. The server will try to bind to this port on all local IP addresses.

#### Remote Host

For non-[Listen Only](#) links, this indicates which host to connect to. This might either be an *Fully Qualified Domain Name* (FQDN), or an *IP address*.

#### Remote Port

For non-[Listen Only](#) links, this indicates which port on [Remote Host](#) to connect to.

#### Use TLS

When this option is enabled, and if a certificate is configured, the link will use and require TLS.

X2X links support the following additional parameters:

#### Listen Only

When this option is set, the server will never initiate a connection over this link on its own, but will accept connections coming in. Established *Listen Only* links are always *Bidirectional*.

#### Bidirectional Connections

If set, a single connection can send XML stanzas in both directions. If not set, a connection will be established for each direction of stanza flow.

When [Use TLS](#) is enabled, the following additional options are available:

#### TLS Identity

Key and certificate this link presents to inbound connections, as well as sends to the peer when trying to establish a connection.

#### Pinned Link TLS Certificate

If set, this link will only allow connections with peers that present this exact certificate.

#### Require a Valid Certificate in TLS Session

When set, peers must use a valid certificate that isn't expired, matches the [Remote Host](#), and when [Pinned Link TLS Certificate](#) is set, matches the pinned certificate.

When this option is not set, any TLS certificate will be accepted, and TLS will only be used to prevent eavesdropping.

#### Ignore System Trust Anchors

When set, only trust anchors explicitly configured within M-Link are trusted, and system trust anchors will be ignored.

#### Link Trust Anchors

This is a list of TLS Trust Anchors that will be used for [Section 4.6.3, “Authentication”](#).

## 4.6.3 Authentication

When using Links, authentication is based on the identity of the server to which the Link connects, not to the addressing of the XMPP data that are being transmitted. Unless the `requireValidCertificates` option is enabled on a Link (and TLS is configured) any connection associated with the Link is considered to be authenticated for transmission of XMPP data addressed to any Peer Control configured to use the Link. If the `requireValidCertificates` option is enabled then the certificates used in TLS must be trusted - either by the `pinnedLinkCertificatePEM` option matching the provided certificate, or by the presentation of a certificate whose subject (see [Section 5.2, “Certificate Verification”](#)) matches the `remoteHost` option and was issued by an authority configured as a trust anchor (see [Section 5.3, “Trust Anchors”](#)).

## 4.6.4 SLEP Links

---

**Note:** This section only applies to M-Link MU Gateway and M-Link MU Server.

---

To connect to other servers over HF Radio connections, M-Link MU Gateway supports the *SIS Layer Extension Protocol* (SLEP). This protocol is used to connect to remote sites over HF Radio connections using *Icon-5066* servers. Configuring links of this type takes two steps: creating a definition that will allow connecting to a *STANAG 5066 Subnet Interface Service* (SIS) connection to an Icon-5066 server, and creating a link that uses such a STANAG 5066 definition.

### 4.6.4.1 STANAG 5066 Server Definition

A configuration is required for every local Icon-5066 server. All fields are required, and must reflect the configuration of the Icon-5066 server. The *SAP ID* must be the same for the local and remote M-Link servers, as SLEP connections can only be established between two servers configured for the same *SAP ID*.

M-Link will try to maintain a connection to each defined STANAG 5066 server. If the connection breaks, M-Link will try to reconnect within the amount of time specified in the the *Advanced Connection Management* option *STANAG 5066 Maintenance Timer*.

### 4.6.4.2 Configuring SLEP Links

Once a STANAG 5066 Server Definition has been created, SLEP links can be configured. Multiple links may use the same STANAG 5066 Server Definition as long as the *Remote SIS Address* differs from other links using the same STANAG 5066 server.

Stanza acknowledgement using *XEP-0198* is available on SLEP links. Additionally, SLEP allows compression of the data sent over the radio connections. To configure this, both ends must enable the *Compress* flag in the link configuration.

[Peer Controls](#) using SLEP links will generally need [Transformations](#) to strip out (parts of) stanzas that are not strictly required in order to reduce the bandwidth required.

---

## 4.7 Reconnection Backoff and Retry

S2S and Link connections are subject to reconnection backoff and automatic retry. If a new connection is required (e.g. due to a stanza being queued for delivery to a Peer), a connection attempt will be made. If the connection attempt fails, subsequent attempts for the same connection will be initiated automatically after a delay. The delay increases by 50% on each subsequent connection failure, until the delay length reaches a limit, at which it remains constant.

When a connection succeeds, the associated delay is cleared.

When automatic reconnection attempts for an individual connection instance have been in operation for configurable length of time, further attempts at reconnection are abandoned, and connection failure is recorded.

Three system-wide configuration options allow configuration of the base retry delay, retry delay limit and duration of retry attempts:

**Remote Session Connection Retry Base Delay**

This is the basic delay value, in seconds.

**Remote Session Connection Retry Delay Limit**

This configures the maximum length of the retry delay, in seconds.

**Remote Session Connection Retry Duration**

This configures the duration (in seconds) of automatic reconnection attempts for an individual failed connection.

For example, with a base delay of 10 seconds, a delay limit of 35 seconds and a retry duration of 90 seconds, reconnection attempts would be delayed as follows:

- First attempt: immediate
- Second attempt: delayed by 10 seconds
- Third attempt: delayed by 15 seconds
- Fourth attempt: delayed by 23 seconds
- Fifth attempt: delayed by 35 seconds
- Sixth attempt: delayed by 35 seconds, when this fails no more retries will be attempted, as the total retry duration from the first attempt is greater than 90 seconds

# Chapter 5 TLS Configuration

This chapter describes how M-Link is configured to use TLS, covering configuration of certificates/keys used by the local server, and trust of certificates used by other servers.

---

**Note:** Note: Your M-Link installation must have the TLS feature enabled before TLS can be used.

---

Transport Layer Security (TLS) is used by M-Link to provide privacy, data integrity and authentication for client-to-server and server-to-server connections. Use of TLS requires the configuration of three main classes of information:

## Public Key Certificate Chain

A Public Key Certificate Chain certifies the ownership of a public key (by the named subject of the certificate).

## Private Key

The Private Key corresponding to the Public Key Certificate Chain noted above.

## Trust Anchors

These are a set of certificates which are to be treated as authoritative entities for which trust is assumed.

A Public Key Certificate Chain and corresponding Private Key must be configured before M-Link can act as a TLS server (that is to say, the responding side of a TLS handshake) and provide a TLS-encrypted connection.

The same Public Key Certificate Chain and Private Key will be used by M-Link when it is acting as a TLS client (initiating a TLS exchange with a TLS server).

Whether the Public Key Certificate Chain provided by M-Link in the TLS handshake also provides authentication depends on whether the Certificate Chain has been signed by a Trust Anchor which the other party to the TLS handshake trusts.

The Trust Anchor list identifies a set of trusted signing entities which will be used when verifying the validity of a Certificate Chain. The configured list may be used in addition to Trust Anchors provided externally to M-Link (installed as part of the operating system, or placed in system-wide certificate stores), or can replace these external Trust Anchors completely.

Private Keys are normally stored in encrypted form. In this case, a passphrase which will be used when decrypting the Private Key must also be configured.

A number of secondary TLS controls are available:

## Disable TLSv1

This allows use of TLS protocol version 1 to be disabled. This defaults to false, so that TLSv1 is allowed.

## Cipher Suites

The set of cipher suites which will be proposed or accepted during TLS handshaking can be configured, as a standard OpenSSL cipher string. This is empty by default, allowing any supported cipher suites to be used.

## Diffie-Hellman Parameters

A set of Diffie-Hellman parameters must be provided before some classes of cipher suite can be used by M-Link. A default set of parameters is built in, but can be overridden if necessary.

## Ignore System Trust Anchors

Disables use of built-in Trust Anchors; only those Trust Anchors which have been configured for M-Link will be used. This defaults to false.

Public Key Certificate Chains, Private Keys and Trust Anchors are configured as PEM-encoded strings. Many Certification Authorities provide certificates in this form anyway, and tools are readily available to convert certificates which have been provided in other formats.

---

## 5.1 Detailed Configuration

TLS information can be specified at various points within M-Link's configuration, as described below.

### 5.1.1 Default Settings

Default TLS configuration can be specified for an M-Link server as a whole. In a simple setup this may be all that is required:

Public Key Certificate Chain  
Private Key  
Private Key Passphrase  
Trust Anchors  
Disable TLSv1  
Cipher Suites  
Diffie-Hellman Parameters  
Ignore System Trust Anchors

### 5.1.2 Peer Control Settings

Peer control TLS settings have the effect of overriding the TLS settings for the IM or MUC domain domains to which they apply (and the default settings), when performing server-to-server (S2S) communication:

Public Key Certificate Chain  
Private Key  
Private Key Passphrase  
Trust Anchors  
Ignore System Trust Anchors

A Peer Control may also specify the following TLS-specific settings:

#### TLS

This item controls how TLS is used by the S2S link, and takes the values:

- Never - do not use TLS
- Available - use TLS encryption if available
- Required - require use of TLS encryption
- Required With Auth - require an authenticated TLS connection

#### Pinned Peer Certificate

See below for a detailed description

A Pinned Peer Certificate allows the standard certificate verification (performed during the TLS handshake) to be overridden. Instead of the certificate presented by the Peer being verified against Trust Anchors, checked for revocation etc, a simple comparison with the Pinned Certificate is performed, and if the presented and Pinned certificates match, the authentication is successful.

### 5.1.3 Link Settings

Link TLS settings have the effect of overriding any default TLS settings:

Public Key Certificate Chain

Private Key

Private Key Passphrase

Trust Anchors

Ignore System Trust Anchors

In addition, a Link may specify:

Use TLS

If this is set true, TLS will be used and required for this Link.

Pinned Link Certificate

See below for a detailed description

Require Valid Certificate

If this is set true, a TLS session must present a valid certificate.

A Pinned Link Certificate allows the standard certificate verification (performed during the TLS handshake) to be overridden. Instead of the certificate presented by the Link being verified against Trust Anchors, checked for revocation etc, a simple comparison with the Pinned Certificate is performed, and if the presented and Pinned certificates match, the authentication is successful.

---

## 5.2 Certificate Verification

A certificate which is presented as part of a TLS handshake is verified via a multi-stage process. The first stage takes place during the handshake itself, and checks (among other things) that the certificate has been signed by one of the Trust Anchors which have been configured for the current TLS context. Even if the certificate fails one or more of these checks, the TLS handshake may still complete successfully, with an encrypted TLS session being established.

Once the TLS handshake is complete, secondary checks are performed on the presented certificate, if the configuration of the domain or link requires a valid certificate. These are:

- If a Pinned Certificate is configured, the results of the first-stage verification are ignored and a direct comparison between this and the presented certificate is performed. If they match, the presented certificate is considered valid. If they do not match, the presented certificate is considered invalid. In either case, no further verification of the presented certificate is performed.
- A check of the result of the first-stage verification, described above. If this first-stage verification has failed, no further action is taken, and the certificate is considered invalid.
- For Server-to-Server connections which use a Peer Control, a Subject Alternative Name match is attempted for the XMPP domain to which the connection is being made. If this fails, a Subject Alternative Name match for the Peer Control's connectHost configuration setting is attempted. If both of these matches fail, the certificate is considered invalid.
- For Server-to-Server connections which use a Link, a Subject Alternative Name match against the Link's remoteHost configuration setting is attempted. If this match fails, the certificate is considered invalid.

- If the appropriate checks described above succeed, the certificate is considered valid.

A certificate can contain multiple Subject Alternative Names, of varying types. When attempting to match a domain name, DNS name or hostname against these, a number of different comparisons are performed:

- A match against one of the certificate's DNS Names. This includes wildcard matching, so that a certificate with a DNS Name of `*.isode.com` would match `mary.isode.com`.
- A match against one of the certificate's SRV Names. SRV Names are prefixed with `_xmpp-server` if the certificate belongs to an XMPP server or `_xmpp-client` if presented by an XMPP client. Thus an SRV name of `_xmpp-server.mary.isode.com` would match the domain `mary.isode.com`.
- A match against one of the certificate's XMPP Addresses.
- If the certificate has no other Subject Alternative Names, a match against one of the certificate's Common Name values.

---

## 5.3 Trust Anchors

Trust Anchors are certificates which identify trusted signing entities. These are used by M-Link to verify that a chain of certificates (up to and including an end-entity certificate) received from another XMPP server or client is valid.

Most operating systems provide a built-in set of Trust Anchors which identify commercial Certification Authorities. The location and format of these is system-specific. By default, M-Link will make use of these Trust Anchors. Use of system Trust Anchors can be overridden via configuration either for the whole M-Link installation or per-Peer Control or per-Link.

A set of private Trust Anchors may be specified as part of M-Link's configuration, either across the whole installation, per-Peer Control or per-Link. Use of private Trust Anchors is required when the end-entity certificates being presented have been signed by Certification Authorities whose CA certificates are not configured as part of the operating system.

# Chapter 6 Security Labels, Clearances and Policies

This chapter describes how Security Labels, Clearances and Policies are configured and used within M-Link.

A Security Label is a structured representation of the sensitivity of a piece of information (e.g. a message). This is used in conjunction with a Security Clearance (a structured representation of what information sensitivities an entity is authorized to access) and a Security Policy to control access to each piece of information. For instance, a message stanza could be labelled as "SECRET", hence requiring the sender and the receiver to each have a clearance granting access to "SECRET" information. Labels include a Display Marking which provides a human-readable form of the machine-readable label, as well as information about which colours XMPP clients should use when displaying the label.

The use of Security Labels in XMPP is described in detail in [XEP-0258: Security Labels in XMPP](#). General information on Isode's Security Policy Infrastructure can be found at [Isode Security Policy Infrastructure](#) and its use by Isode within XMPP is described in [Using Security Labels to Control Message Flow in XMPP Services](#).

Determining whether a labelled message stanza is permitted to be transferred to an entity (i.e. a domain or person) is a two stage process:

- The label associated with the destination entity is tested against the clearance associated with the message stanza's sender entity. This is testing "Is the sender of this message cleared to send a message to this entity?".
- The label on the message stanza is tested against the clearance associated with the destination entity. This is testing "Is the recipient of this message cleared to receive a message with this security label?".

Each of the tests described above takes place within the context of a Security Policy. The Security Policy controls how the Label/Clearance comparison is performed, and may provide other facilities such as equivalences to other Security Policies. When Security Label conversion is taking place, two Security Policies will be used: the Policy under which the Label was issued and the Policy to which the Label will be converted.

Each domain which originates or accepts message stanzas may be configured with its own Security Label and Clearance. A system of defaulting within the M-Link configuration means that, in the simplest case, a single server-wide Security Label and Clearance are all that need to be configured. For more complex systems, individual domains can override defaults with their own settings.

A final configuration item which may be associated with each domain entity is a Default Stanza Label. This is the Security Label which will be applied to any unlabelled message stanza originating at that entity.

IM Domain users may be configured with their own Clearance (see [Client Security Clearance Configuration](#)); if not set, this falls back to the [IM Domain User Default Security Clearance](#).

Security Labels and Clearances are maintained within catalogs (a Label Catalog is a collection of associated labels; a Clearance Catalog is a collection of associated clearances). Multiple label and clearance catalogs may be loaded into an M-Link configuration. Specific label and clearance catalogs may then be associated with individual components of the server's configuration, and individual Security Labels and Clearances selected for use from within these catalogs.

---

## 6.1 Object Stores

Security Policies, Security Label catalogs and Clearance catalogs are maintained in type-specific Object Stores within M-Link. All three types of entity are usually maintained and distributed publicly as XML documents. The M-Link Administration Interface allows these XML documents to be loaded into the appropriate Object Stores. Catalogs, Labels and Clearances within catalogs, and Security Policies can then be selected for use in configuration options.

Security Policies are represented within M-Link as SDN.801c Security Policy Information Files in [Open XML SPIF](#) format. Isode provides sample Security Policies and associated Security Label and Clearance catalogs in subdirectories of `$(SHAREDIR)/security_label/example_data`.

---

## 6.2 Server Configuration

The global "Use Security Labelling" option within MLWC controls whether any Security Labelling configuration is exposed and enabled.

---

**Note:** If Security Labelling has been enabled, but no Clearance has been configured for the Server, all messages which contain Security Labels will be blocked.

---

Within the Security Labelling Configuration section of MLWC, the following items may be configured:

### Server Default Security Policy

The default Security Policy to be used by the server. This is a selection from the set of policies which have been loaded into the Security Policy Object Store. A Security Policy will normally define a default Security Label and Clearance; if this is the case, the defaults will be used as described in [Security Policy Default Label and Clearance](#).

### Server Default Security Label

This is the default Security Label to be used by the server. It is selected in a two-stage process; first a Security Label Catalog is chosen from the Security Label Catalog Object Store, and then an individual Security Label is chosen from within that catalog. This allows the transit of messages from originators with unsuitable clearances to be blocked.

### Stanza Default Security Label

This is the default Security Label which is applied to unlabelled stanzas. It is selected in a two-stage process; first a Security Label Catalog is chosen from the Security Label Catalog Object Store, and then an individual Security Label is chosen from within that catalog. The Security Label Catalog defaults to that chosen for the Server Default Security Label.

### Server Default Security Clearance

This is the default Security Clearance to be used by the server. It is selected in a two-stage process; first a Security Clearance Catalog is chosen from the Security Clearance Catalog Object Store, and then an individual Security Clearance is chosen from within that catalog. The Clearance blocks the transit of messages without a suitable label.

**Default XEP-0258 Label Format**

This controls the encoding in which Security Labels will be emitted.

**Require Security Label on Inbound Message**

This mandates that a Security Label is present on all inbound message stanzas.

**Reject Inbound Messages With Unrecognized Security Labels**

This option controls how messages which include valid but unrecognized labels are handled. The default behavior is to reject such messages. Disabling this option causes such messages to have the label discarded and to then be treated as unlabelled.

**User Default Security Clearance**

This is the default clearance which will be applied to Users. It may be overridden by per-IM Domain and per-User configuration.

**Update Display Markings**

Security Labels include a display marking phrase (e.g. "Secret") and optionally foreground and background color specifications. The default behavior of M-Link is to normalize these fields when serializing a label, replacing values present in the label with those defined for the label by the Security Policy. This control allows this functionality to be disabled, so that labels are emitted with the display markings that they arrived with.

---

## 6.3 Security Policy Default Label and Clearance

A security policy may designate one of the classifications within the policy as the default, and derive a clearance and label from this classification. These are used by M-Link as the defaults of last resort when performing access control decisions.

For example, if a message is being delivered to an IM domain, the message's clearance will be tested against a label associated with the IM domain. A security label may be explicitly configured for the IM domain, or it may inherit the server's explicitly configured security label. If neither of these are set, the default label (if any) defined by the security policy will be used. In the same way, if the message does not have an associated clearance, the default clearance defined by the security policy will be used in the access control decision, and if the message is not labelled, the default label configured for the domain will be used, falling back to the security policy default label.

The same defaulting is used when checking that a user has sufficient clearance for an IM domain's label during association establishment to provide defaults for the user clearance and IM domain label.

---

## 6.4 Peer Control Configuration

Individual Peer Controls may have their own Security Labelling configuration. This provides a mechanism for controlling the classification of messages which can flow to and from an external domain. For example, a simple configuration might allow messages labelled as Unclassified, Sensitive or Confidential to be sent between users within the same local IM domain, but only permit Unclassified messages to be passed to external domains.

Individual Peer Controls may also have their own Security Policy configured. When this is the case, messages flowing to the Peer Domain will have their labels converted so that they are consistent with Security Policy which has been configured for the Peer.

If an inbound message includes a label which has not been issued under either the Server Security Policy or the Peer's own Security Policy (if configured), then the label will be discarded and the message will be treated as if unlabelled. An Audit record will be generated to record the label discard. See [Security Label Handling](#) for more information.

The following items can be configured:

#### Peer Security Label

This is the Security Label applied to the Peer: only stanzas with a source clearance which permits access to this Security Label will be transferred to the Peer. The label is selected in a two-stage process; first a Security Label Catalog is chosen from the Security Label Catalog Object Store, and then an individual Security Label is chosen from within that catalog. When setting this option, the catalog choice will default to that selected for the Server as a whole, but the Peer Security Label does not have a default and must be explicitly selected, since it would normally be different from that of the Server.

#### Peer Default Stanza Security Label

This is the default Security Label which is applied to unlabelled stanzas arriving from this Peer. It is selected in a two-stage process; first a Security Label Catalog is chosen from the Security Label Catalog Object Store, and then an individual Security Label is chosen from within that catalog. The Security Label Catalog choice defaults to that selected for the Server as a whole, and the Peer Default Stanza Security Label choice defaults to the Server's Stanza Default Security Label.

#### Peer Outbound Default Stanza Security Label

This allows the configuration of a Security Label which will be applied to all unlabelled outbound messages for this Peer.

#### Peer Security Clearance

This is the Security Clearance applied to the Peer: only stanzas which include a label which the Peer's clearance allows will be transferred to the Peer. The clearance is selected in a two-stage process; first a Security Clearance Catalog is chosen from the Security Clearance Catalog Object Store, and then an individual Security Clearance is chosen from within that catalog. The Clearance Catalog defaults to that selected for the Server as a whole, while the Peer Security Clearance itself must be explicitly selected, since it would normally be different from that of the Server.

#### Peer Relabelling Security Policy

This is the Security Policy which applies to the Peer. If this is set (and different from the Server Default Security Policy), the label on an outbound message (issued under the Server Default Security Policy) will be replaced with an equivalent label issued under the Peer Relabelling Security Policy. For this to be possible, the Server Default Security Policy must have equivalences for the Peer Relabelling Security Policy configured. See [Security Label Handling](#) for more information.

#### Peer Outbound Security Label Format

This controls the format in which outbound Security Labels are emitted. It defaults to the server-wide Default XEP-0258 Label Format setting.

#### Peer Update Display Markings

Security Labels include a display marking phrase (e.g. "Secret") and optionally foreground and background color specifications. The default behavior of M-Link is to normalize these fields when serializing a label, replacing values present in the label with those defined for the label by the Security Policy. This control allows this functionality to be disabled, so that labels are emitted with the display markings that they arrived with.

#### Require Security Label on Inbound Messages

This mandates that a Security Label is present on all inbound message stanzas from a Peer.

#### Reject Inbound Messages With Unrecognized Security Labels

This controls whether an unrecognized Security Label causes an inbound message to be rejected, or whether the Security Label is stripped off and the unlabelled message is allowed to proceed.

#### Reject Inbound Messages With Unrecognized Security Labels

This option controls how messages which include valid but unrecognized labels are handled. Disabling this option causes such messages to have the label discarded and to then be treated as unlabelled.

#### Peer Mandatory Label

This allows configuration of an IC-ISM format label which will be required by default to be present on all inbound message stanzas from the Peer. Messages which do not contain a matching label will be rejected. A label match will cause the label to be removed from the message stanza and the Peer Default Stanza Security Label (if any) to be applied to it. The label may also be applied to outbound messages to the Peer. This is a specialized option and should only be used when suggested by Isode Support. See [Security Label Handling](#) for more information.

#### Require Peer Mandatory Label on Inbound Messages

The default value of this option requires the label configured in the Peer Mandatory Label option to be present in all inbound messages. See [Security Label Handling](#) for more information.

#### Apply Peer Mandatory Label to Outbound Messages

This causes the label configured in the Peer Mandatory Label option to be applied to all outbound messages, replacing any other label present.

#### Add FLOT to Outbound Messages

This configures the conditions under which a First Line Of Text marking is added to the body of messages, and is intended for use when sending messages to Peers which do not support XEP-0258 labeling. The FLOT marking added is derived from the message's Security Label. If the message also includes an HTML payload, this will be stripped. The "Add FLOT if message has non-default Label" option will cause the Security Label on the message to be compared against the Peer Outbound Stanza Default Security Label, if set, otherwise against the label marked as Default in the Security Policy, if any. If a match is detected, no FLOT will be inserted.

#### Peer Fixed Security Label

This configures a fixed Security Label (chosen from a Label Catalog) which will be applied to all outbound messages, replacing any Security Label present on the message.

#### Peer Disable Label Out

If set to true, outbound messages will not include a security label. Access control checks will still be performed if configured.

---

## 6.5 IM Domain Security Configuration

Individual IM Domains may have their own Security Labelling configuration. This provides a mechanism for controlling the classification of messages which can flow to and from an internal domain.

If a message includes a label which has not been issued under the Server Security Policy, then the label will be discarded and the message will be treated as if unlabelled. An Audit record will be generated to record the label discard.

The following items can be configured:

#### IM Domain Security Label

This is the Security Label applied to the IM domain: only stanzas with a source clearance which permits access to this Security Label will be permitted to be sent to or from the domain. The label is selected in a two-stage process; first a Security Label Catalog is chosen from the Security Label Catalog Object Store, and then an individual Security Label is chosen from within that catalog; this defaults to the Server Default Security Label.

**IM Domain Default Stanza Security Label**

This allows the configuration of a Security Label which will be applied to all unlabelled messages sent from this IM domain.

**IM Domain Security Clearance**

This is the Security Clearance applied to the IM domain: only stanzas which include a label which the domain's clearance allows can be sent to or from the domain. The clearance is selected in a two-stage process; first a Security Clearance Catalog is chosen from the Security Clearance Catalog Object Store, and then an individual Security Clearance is chosen from within that catalog; this defaults to the Server Default Clearance.

**IM Domain User Default Security Clearance**

This is the default Security Clearance applied to users within this IM domain. The clearance is selected in a two-stage process; first a Security Clearance Catalog is chosen from the Security Clearance Catalog Object Store, and then an individual Security Clearance is chosen from within that catalog; this defaults to the server's User Default Security Clearance.

**IM Domain Update Display Markings**

Security Labels include a display marking phrase (e.g. "Secret") and optionally foreground and background color specifications. The default behavior of M-Link is to normalize these fields when serializing a label, replacing values present in the label with those defined for the label by the Security Policy. This control allows this functionality to be disabled, so that labels are emitted with the display markings that they arrived with.

**Add FLOT to Messages Before Delivery**

This configures the conditions under which a First Line Of Text marking is added to the body of messages before delivery to IM Domain users. The FLOT marking added is derived from the message's Security Label. If the message also includes an HTML payload, this will be stripped.

**Strip Security Label from Messages Before Delivery**

If set to true, messages will not include a security label when delivered. Access control checks will still be performed if configured.

**IM Domain Outbound Security Label Format**

This controls the encoding in which Security Labels will be emitted for users within the IM domain.

**IM Domain XEP-0258 Label Catalog**

This is the Security Label Catalog which will be used to generate the XEP-0258 label catalog which is served to the IM domain's clients. It is chosen from the Security Label Catalog Object Store. It defaults to the catalog from which the Server Default Security Label has been chosen.

---

## 6.6

## MUC Domain Configuration

Individual MUC Domains may have their own Security Labelling configuration. This provides a mechanism for controlling the classification of messages which can flow into MUC rooms within the domain.

The following items can be configured:

**MUC Domain Security Label**

This is the Security Label applied to the MUC domain: only stanzas with a source clearance which permits access to this Security Label will be permitted to be sent to MUC rooms within the domain. The label is selected in a two-stage process; first a Security Label Catalog is chosen from the Security Label Catalog Object Store, and then an individual Security Label is chosen from within that catalog; this defaults to the Server Default Security Label.

#### MUC Domain Default Stanza Security Label

This allows the configuration of a Security Label which will be applied to all unlabelled messages sent to this MUC domain.

#### MUC Domain Security Clearance

This is the Security Clearance applied to the MUC domain: only stanzas which include a label which the domain's clearance allows can be sent to MUC rooms within domain. The clearance is selected in a two-stage process; first a Security Clearance Catalog is chosen from the Security Clearance Catalog Object Store, and then an individual Security Clearance is chosen from within that catalog; this defaults to the Server Default Clearance.

---

## 6.7 Audit Logging

Access Control Decision Function (ACDF) failures are audited by M-Link. Three different audit types are generated:

#### acdfSourceFailure

This indicates that either the source of the stanza did not have a Clearance, or that the processing entity's Label is not permitted by the Clearance.

#### acdfLabelFailure

This indicates that the stanza's Label is not permitted by the processing entity's Clearance.

#### clientAcdfFail

This audit is logged when a client is binding to an IM Domain, and indicates that the IM Domain's Label is not permitted by the client's Clearance.

In the first two cases, the stanza's To and From address fields are included in the audit record, together with the name of the domain for which the ACDF is being called. In the case of an unexpected ACDF failure, this should enable the identification of configuration point responsible for the failure. A free-form error string indicating the reason for the ACDF failure is also included. For a ClientAcdfFail, the JID of the client is included.

---

## 6.8 Security Label Handling

This section summarizes how an inbound message, which may be labelled or unlabelled, is handled when Security Labelling is enabled. The actions listed below are taken, in order.

- If the message has arrived via a Peer Control, and a Peer Mandatory Label is configured for the domain, and the domain's Require Peer Mandatory Label on Inbound Messages option is set True, messages which are unlabelled or whose label does not match the configured Peer Mandatory Label will be rejected. A message with a matching label will have the label stripped and the unlabelled message will be allowed to continue.
- If the message includes a label, this is extracted. Failure to extract the label (e.g. because it is badly formatted) will cause the message to be rejected.
- If the message has entered via a Peer Control which has a Peer Relabelling Security Policy configured, the label is validated against this alternative policy, and if it passes validation, an Equivalent Label is generated under the Server Security Policy, replacing the message's original label. If the label does not pass validation against the alternative policy, the label is discarded and the unlabelled message is allowed to

continue. If an Equivalent Label cannot be generated, the original label is discarded, and the unlabelled message is allowed to continue.

- The label (or Equivalent Label) is validated against the Server Security Policy. If it passes validation, the message is allowed to continue, with the label attached. If validation fails, the label is discarded and the unlabelled message is allowed to continue.

Following label acquisition and validation, checks for label presence are performed:

- If the message has arrived via a Peer Control or IM Domain which has the Require Security Label on Inbound Messages option set True, and does not now contain a label, it will be rejected.

The final stage of label handling on ingress is to set a default label. If the message does not contain a label, a default label may be applied, using the following rules:

- If the destination domain is a MUC domain, apply the domain's MUC Domain Default Stanza Security Label.
- If the source domain is non-local, apply the appropriate peer control's Peer Default Stanza Security Label.
- If the source domain is a local IM domain, apply the domain's IM Domain Default Stanza Security Label.

## 6.9 XEP-0258 Label Catalogs

M-Link supports XEP-0258 label catalog discovery. The label catalog which is provided to a client in response to a catalog discovery request is generated as follows:

The [IM Domain XEP-0258 Label Catalog](#) configured for the client's IM domain is used as the basis for what will be returned to the client. This is then subject to further filtering which removes any label categories which would not be permitted by the following clearances:

- The clearance configured for the destination IM, MUC or Peer Domain. If the destination does not have a clearance configured, an empty label catalog is returned to the client.
- The clearance configured for the source IM domain: [IM Domain Security Clearance](#).
- The clearance configured for the source client: [Client Security Clearance Configuration](#). If no source client clearance is set, the [IM Domain User Default Security Clearance](#) is used, if set.
- If the target is a client within a local IM domain, the clearance configured for the target client: [Client Security Clearance Configuration](#). If no client clearance is set for the target, the [IM Domain User Default Security Clearance](#) for the target domain is used, if set.

## 6.10 Clearances

A message is assigned a Clearance on ingress to the M-Link server. If the message arrives via a Peer Control, it is assigned the Peer Control's [Peer Security Clearance](#). If ingress is via an IM Domain (i.e. from a client), the message is either assigned the client's Security Clearance (see [Client Security Clearance Configuration](#)) or the IM Domain's [IM Domain User Default Security Clearance](#).

## 6.10.1 Client Security Clearance Configuration

For IM domains which are configured with an Authentication Type of Static Users, a user can be assigned a clearance by setting their User Security Clearance option.

When LDAP authentication is used, the user's clearance will be held as an attribute within the user's Directory entry. M-Link looks for various attribute types, trying the values returned for each attribute type until one which can be successfully interpreted as a valid Clearance issued under the [Server Default Security Policy](#) is found. Three attribute types are attempted, in the order shown below:

- An X.509 certificate, held in a userCertificate attribute. The clearance itself should be held within the X.509 certificate in a Subject Directory Attributes Extension.
- An XML clearance, held in an sioClearance attribute.
- An X.509 clearance, held in a clearance attribute.

---

## 6.11 Client IM Domain Access Control

If an IM Domain has been configured with an [IM Domain Security Label](#), only clients which have been configured with a clearance allowing access to the label will be able to authenticate. See [Client Security Clearance Configuration](#) for more information.

# Chapter 7 M-Link Configuration and Management

This chapter introduces the M-Link Administration Interface, used for configuration and management of the M-Link server.

M-Link supports configuration and management through the M-Link Administration Interface, including management of product activation keys. M-Link has a built-in web server accessible to a browser by giving the server host as URL and specifying the “Web Admin Port”. By default, that is 5221, for example:

```
https://localhost:5221/
```

This example assumes that HTTPS is used for the Admin interface, which is the default unless the TLS feature is not available. Alternatively:

```
http://localhost:5221/
```

It is also possible to configure M-Link programmatically using the interface defined in [Chapter 13, HTTP API](#).

M-Link cannot be used without an activation key. Upon first use, the M-Link Administration Interface takes the user through initial administrator user setup and product activation.

- Initial administrator user credentials are included within the M-Link Administration Interface configuration, to allow set up independently of any external directories. Their input is a one-off set up.
- Product activation key management is addressed later in this chapter (see [Section 7.6, “Product Key Management”](#)).

If the operations listed above have been performed, the M-Link Administration Interface starts with the login page.

Once successfully logged in, the “home page” of the M-Link Administration Interface offers a choice between configuration, monitoring and administration activities. These are briefly described in following sections.

---

## 7.1 UI Navigation

After successful login, the top bar becomes visible and features the following:

- On the left, the application name which varies according to which M-Link features are activated.

This is also a button for returning to this initial location from anywhere.

- A choice of main areas:
  - “Configuration”
  - “Monitoring”
- On the right, a drop-down menu with classic options providing:
  - Access to documentation.
  - Information about M-Link and the activation status.

- Means of logging out or editing credentials.

A sidebar menu appears when entering each of the main areas which adjusts to the current location, providing a choice of sub sections and means of stepping back towards the top level. The “Back” browser action may also be used, but that always returns to the previous page, not necessarily to a page closer to the root.

Every navigation device (button, card or link which, upon mouse click, takes the M-Link Administration Interface to a different page) also supports “Open Link in New Tab” and “Open Link in New Window” normally provided by browsers in their context menus.

When working with multiple servers, it may be useful to identify tabs/windows connected to the same server: see option “Application Title Override” at the bottom of the “Global Options”, the starting point of the “Configuration” area.

---

## 7.2 Configuration Area

The Configuration area consists of a main or root page, displaying “Global Options”, and a number of sections or sub pages, accessible by clicking in a menu bar on the left hand side. These are organized in a tree-like structure, each page having a path from the root.

### 7.2.1 Common Features of Configuration

#### 7.2.1.1 Viewing and Editing Settings

Each option is displayed with a name and followed by a description. Please refer to other chapters of this manual for the full significance of each option. Sometimes, long descriptions appear shortened. These can be viewed in full by clicking on the `MORE . . .` link.

Most options are defined by a simple value which can be edited in various ways as appropriate: mouse clicks, text input or drop-down menu option selection. Changes only take effect, that is, are conveyed to the server, when the Submit button, at the bottom of the page, is clicked. There is also a Cancel button which restores the page to the values currently set in the server.

Some options are only visible and in use when enabled by other settings.

- This serves to remove clutter, such as option “Customise Ports”, which needs to be selected before port number options appear.
- It also helps to manage dependencies: some option combinations would not always make sense. When options are thus controlled from another page, it may be necessary to submit the page with the enabling option before the dependent features appear. For example, the “Security Labelling Configuration” section only appears after submitting the “Global Options” page with “Use Security Labelling” ticked.

A few options are more complex and have a specialized editing mechanism. These are described later.

#### 7.2.1.2 Default Values

A large number of options have defaults. These are denoted by a value displayed in italics, unless it is a boolean option (i.e. selected or not). They will also have a “Use default” tick box.

- This can be ticked to restore the option to its default value. If the tick is showing, then the option is definitely set to its default.

- Changing the value will automatically untick the box. It is, of course, possible to set values which happen to equal the defaults, in which case the server remembers this state explicitly, and the configured values will persist even if the defaults are changed.

Default values are often immutable, possibly defined by a standard. In some cases the default is itself an option which can be modified. Any option used as a default elsewhere is displayed with a message including paths to all the options currently defaulting to that option.

### 7.2.1.3 Validation

Option values are generally constrained by their types. For example, numerical values may be range limited, and domain names must adhere to standard specifications, etc.

- In most cases, validation is applied as you type and errors get highlighted.
- In a few cases, where consistency is required with other options, validation only happens upon submission to the server.

In any case, the M-Link server rejects invalid configuration updates. When that occurs, the page remains unsubmitted and an explanation pops up in a red, persistent box, which can be used to locate the changes required to make a subsequent submission succeed.

An option may be mandatory if no default would be sensible. This is typical of list items that can be added, but not without filling in some essential properties. The mandatory options are marked by a label reading “Required”. Also, the Submit button will remain disabled until all mandatory options are filled.

Due to dependencies between settings, it is possible for a change to enable mandatory settings on other pages. This somewhat rare condition is detected upon submission and produces a pop-up, named “Missing fields”, listing options that will be made visible by the update and which are currently not set, yet mandatory. These are editable in the pop-up and must be configured to complete the submission. It is recommended to make a note of the option paths listed on the “Missing fields” pop-up and visit the containing pages after the current submission completes successfully, as other setting changes may be desirable.

Other unrelated errors may occur, such as caused by a connection failure. They generate red, persistent pop-ups containing an explanatory message.

All error pop-ups can be dismissed by a simple click.

## 7.2.2 Special Options

### 7.2.2.1 Options Referring to Objects

Configuring an M-Link server generally requires importing external data objects which cannot be predetermined. For example, TLS features require keys and certificates which do not ship with the server. This is a two step process:

- Upload of a number of objects into M-Link managed stores.
- Selection of stored objects where needed in the configuration.

For convenience, it is in fact possible to upload an object from the places where an object selection is needed. But logically, the two operations are disconnected, as the same object may be selected in several places of the configuration. An object may also be uploaded but not selected anywhere (i.e. not in use in the configuration) or cease to be selected.

Object stores group objects of the same kind and act like internal catalogs. Setting a configuration option to an object from a store amounts to making a selection from a list of object names. Such options appear with an “Edit” button which offers a drop-

down menu. The options in the drop-down depend on what is currently uploaded in the associated store.

Since some of the uploaded objects are themselves catalogs of items (such as security label catalogs), the internal stores can have two levels. The editing pop-up will then show two drop-down menus. The option is then rendered as a combination of the two choices.

The M-Link Administration Interface has a section for object store management. See [Section 7.2.3.8, “Stores”](#).

### 7.2.2.2 List Options

There are basically two kinds of lists:

- Collections of configurable items.
- Options that happen to be multi valued.

The first case corresponds to M-Link Administration Interface sections (such as “Peer Controls”, see [Section 7.2.3.3, “Peer Controls”](#)) where the side bar menu leads to a page of cards representing the items in existence.

- Such pages may include an “Add item” button, unless displaying a fixed set of items.
- The cards may be clicked to access the features of the corresponding items.
- Items may also be reflected in the side bar menu, where a click will have the same effect.

There are variations with the second case of lists. Depending on complexity, the list items are either

- rendered in full or
- only partially but viewable one at a time in a pop-up.

In all cases, obvious controls are available to add, edit or delete items.

## 7.2.3 Configuration Sub Sections

### 7.2.3.1 Global Options

This page features all the options that do not belong in any of the sub sections. That is, options that apply server wide, excluding the following topics, which are substantial enough to require dedicated sections:

- TLS Configuration, with the exceptions mentioned below.
- Logging Configuration.
- Security Labelling Configuration.

The “Global Options” do include a couple of options that are TLS related:

Use HTTPS for Administration Interface

This controls whether the administration interface (i.e. how the M-Link Administration Interface connects to the server) uses the HTTPS or HTTP protocols. Always use HTTPS, unless encryption is not available.

If switching protocols is required, please note that the current connection will close, so the M-Link Administration Interface will appear to freeze and the browser must be told to visit the new location.

Use Secure Cookie Flag

This controls whether use of the login cookie is restricted to secure connections. When enabled, browsers will be instructed to never use the login cookie over HTTP connections.

The cookie is required for most interactions with the server and the default (which depends on the option above, Use HTTPS for Administration Interface) is generally the correct setting. However, this may need overriding in some circumstances, such as when connecting through a proxy that handles TLS for the administration interface.

---

**Warning:** Improper use of this option can block M-Link Administration Interface updating the configuration.

---

#### TLS Key Pair for Administration Interface

This option is obviously required when HTTPS for the Administration Interface is enabled.

If the TLS Key Pair is misconfigured somehow, M-Link will automatically revert to expecting connections through HTTP.

Please note that the “Security Labelling Configuration” section is only available when option “Use Security Labelling” featured in the “Global Options” is set.

### 7.2.3.2 MUC Domains

This page collects all the MUC domains currently configured, represented as cards. Please refer to [Chapter 10, MUC Room Creation and Administration](#) for an introduction to Multi-User Chat concepts.

- Domains can be added by clicking the “Add item” button.
- Domain settings can be viewed or edited by clicking on the corresponding card.
- Deleting a Domain requires viewing it and using the “Delete...” button on the viewing page.

### 7.2.3.3 Peer Controls

This page collects all the Peer Controls currently configured, represented as cards. The card named “default” is special (it does not represent any real Peer), but is present here so it can be edited in the same way as proper Peer Controls.

- Peer Controls can be added by clicking the “Add item” button.
- Peer Control settings can be viewed or edited by clicking on the corresponding card.
- Deleting a Peer Control requires viewing it and using the “Delete...” button on the viewing page.

The “default” item will be used by M-Link when no Peer Control explicitly matches a domain. Please refer to [Chapter 4, Peer Controls and Links](#) for an explanation of how Peer Controls are used by the M-Link server.

### 7.2.3.4 Links

This page collects all the Links currently configured, represented as cards. Please refer to [Chapter 4, Peer Controls and Links](#) for an explanation of the concept.

- Links can be added by clicking the “Add item” button.
- Link settings can be viewed or edited by clicking on the corresponding card.
- Deleting a Link requires viewing it and using the “Delete...” button on the viewing page.

### 7.2.3.5 TLS Configuration

This page displays defaults and server wide options relating to TLS Configuration. See [Chapter 5, TLS Configuration](#) for details on this topic.

### 7.2.3.6 Logging Configuration

The side bar menu presents a “Logging Streams” entry, which leads to a page of cards.

- Logging Streams can be added by clicking the “Add item” button.
- Stream settings can be viewed or edited by clicking on the corresponding card.
- Deleting a Stream requires viewing it and using the “Delete...” button on the viewing page.

If the “Advanced Logging Level Configuration” is selected in a Logging Stream (and submitted), a “Facility Levels” entry appears in the side bar menu, which leads to a page of cards.

### 7.2.3.7 Security Labelling Configuration

The side bar menu entry for this page only exists if the global option “Use Security Labelling” is selected.

This page displays defaults and server wide options relating to Security Labelling. See [Chapter 6, Security Labels, Clearances and Policies](#) for details on this topic.

### 7.2.3.8 Stores

M-Link comes with a fixed set of object stores, each aggregating objects of a specific type and making them available as potential values for appropriate items in the configuration.

- The stores are represented here as cards.
- The contents of each store can be managed by clicking on the associated card.

Stores may be initially empty. They can be populated by clicking on the “Add item” button.

- In most cases, the operation involves choosing one, or a few, local files which contain the external embodiment of the object, for example, an XML file or a certificate in PEM format. The forms that pop up should be self-explanatory.
- In some cases, additional data needs to be input, such as a password.
- In all cases, a unique name must be given to the object. This defaults to the (main) file name without its extension, if any.

Store contents are presented as a list showing the names and upload dates of each object it contains.

- Each entry has a “Details” button leading to a read-only page showing significant details and all the properties of the object. The properties are specific to each type of object.
- This page also allows object deletion via the “Delete...” button, but only if the object is not in use, that is, has not been selected as the value of an item somewhere in the configuration. Otherwise, the header includes links to the places where it is used.

When objects are themselves a type of catalog, their sub structure is also presented in the form of a sub store. The item list of the store includes “Subitems” buttons, which lead to pages resembling the display of store contents. The differences derive logically from the fact that the sub items are parts of an immutable object:

- The list of sub items has no “Add item” button.
- The sub items cannot be deleted individually. Only the whole containing catalog object can be deleted.
- The sub items have neither upload dates nor source file names: these are attributes of the containing catalog object. They may have individual properties, however, since the containing catalog object may provide these.

### 7.2.3.9 Access Tokens

M-Link allows access to its configuration via application access tokens. These provide an alternative way to authenticate HTTP requests and are carried in `Authorization` headers (see section [Authentication by Access Token](#)).

The side bar includes option “Access Tokens” which leads to a page listing all application access tokens currently known to M-Link. The page includes token management facilities:

- The “Add token” button allows the creation of a new token and requires the input of a unique name.

Please note: the token string for subsequent use in `Authorization` headers is only displayed once: a button is provided to copy the token string to the clipboard.

- “Delete” buttons can be used to permanently remove tokens from M-Link.

---

## 7.3 Monitoring Area

Monitoring activities have their own chapter in this manual. Refer to [Chapter 12, Monitoring](#) for further details.

---

## 7.4 MUC Administration

MUC administration has its own interface, included in the M-Link Administration Interface under tab “Multi-User Chat”

---

**Note:** Please note that MUC domain configuration and administration is only available for M-Link products that support *Multi-User Chat* (MUC). Refer to [Section 1.3, “Installing an Isode Activation Key”](#).

---

See [Chapter 10, MUC Room Creation and Administration](#).

---

## 7.5 Miscellaneous Features

### 7.5.1 Server Version

The bottom of the side bar displays the current version of the server connected to the M-Link Administration Interface. Please include that version number in any support request.

The current version number can also be found with the “Product Activation” status in the “About” side bar (see below).

### 7.5.2 Documentation Side Bar

The “About Side Bar” is brought up by selecting the “Documentation” option in the drop-down menu of the top bar. It provides buttons for:

- Displaying the online manual.

- Displaying the release notes.

### 7.5.3 About Side Bar

The “About Side Bar” is brought up by selecting the “About” option in the drop-down menu of the top bar. It provides the following features:

- Display of product activation details and
- Management of the activation key (see [Section 7.6, “Product Key Management”](#)).

### 7.5.4 Signing Out or Editing Credentials

Normally the M-Link server keeps the M-Link Administration Interface logged in unless inactive for the duration of the “Administrators inactivity logout period” (configurable in the “Global Options”). When logged out at the server's initiative, the M-Link Administration Interface may not react immediately depending on the current activity. Any action requiring valid credentials will automatically redirect the current page to a login page.

It is possible however for the user to log out deliberately by choosing the “Sign Out” option in the top bar drop-down menu. This takes the M-Link Administration Interface straight back to a login page.

Credentials may be changed by clicking on the username in the drop-down menu. This opens the “Edit Credentials” form.

The main purpose of this form is to set a new password for the administrator currently logged in. It is also possible to edit the user name at the same time.

- Current (old) credentials are required input.
- Leaving field “New Username” blank preserves the current user name.

The changes take effect immediately (unless they are rejected). In other words, the user is not logged out, and opening a new connection to the server will require entering the new credentials.

---

## 7.6 Product Key Management

Product Key Management is accessible from the “About” side bar, which slides out when the “About” button is clicked in the drop-down menu. The side bar displays the details of the activation key currently loaded.

A new product key is required for the following purposes:

- To change the details of the activation.
- To extend the activation to a later release.
- To add or remove features.
- To widen or narrow the scope of the activation.

The “Update Key” button initiates the same operation required during initial activation.

- The pop-up may be closed at any time to continue working with the current key.
- The final step (the actual key update) overwrites the existing, but only if the new key is valid.

If an M-Link server is no longer needed, the product key can be removed using the “Deactivate” button. This operation cannot be undone, unless the key was saved

elsewhere. As this returns the M-Link product to a non-activated state, product activation is then the only operation available.

# Chapter 8 Instant Messaging Domain and User Configuration

This chapter describes how Instant Messaging domains and users are configured within M-Link.

---

## 8.1 Instant Messaging Domain Configuration

A single M-Link Server may host multiple Instant Messaging domains. Within the IM Domains section of the M-Link Administration Interface, the following items may be configured:

### Domain

The name of the Instant Messaging domain. This must be a valid JID domain component.

### Authentication Type

Specifies whether users within the domain are configured statically within M-Link Server or are held within an external LDAP Directory.

---

**Note:** The use of Static Users is only appropriate for a demonstration system. Production systems should use an external LDAP Directory to hold user configuration.

---

### LDAP Configuration

If the Instant Messaging domain is configured to use LDAP for authentication, this item specifies the LDAP configuration which is to be used. See [LDAP Configurations](#) for more information.

### TLS Required

When this option is enabled, and if a TLS identity is configured, only TLS-encrypted connections from clients will be accepted

### TLS Identity

The key and certificate this Instant Messaging domain presents to inbound connections. If this is configured but the "TLS Required" item is set false, TLS will be offered but not required.

### Enable HTTP Upload

Support for HTTP File Upload (*XEP-0363*) can be enabled or disabled per domain. See [Chapter 9, HTTP File Upload](#) for more information.

An Instant Messaging domain may also have associated Security Labelling configuration: this is described in [Section 6.5, "IM Domain Security Configuration"](#)

### 8.1.1 Instant Messaging Domain Static User Configuration

If an Instant Messaging Domain is configured with an Authentication Type of "Static Users", a list of Static Users can be configured subordinate to the Domain itself. Each Static User may be configured with the following items:

#### JID

The JID of the user. This must have the same domain component as the Instant Messaging Domain.

#### Password

The password for the user.

A User may also be assigned a Security Clearance. This is described in [Section 6.10.1, “Client Security Clearance Configuration”](#).

## 8.1.2 Static Groups

One or more Static Groups can be configured within the Domain. Each Static Group contains a list of JIDs, which should all be members of the Instant Messaging Domain.

## 8.1.3 Roster Groups

An Instant Messaging Domain may be configured with one or more Roster Groups. A Roster Group provides a list of JIDs to all users who themselves are members of the Roster Group. Each Roster Group may be configured with:

### Name

A name assigned to the Roster Group. This name will be exposed to Roster Group members.

### Group Type

This configures the type of the Roster Group; either "Static", referencing an existing Static Group configured for the Instant Messaging Domain, or "LDAP".

### Static Group

If a Group Type of "Static" is selected, this option specifies the Static Group to be used.

### LDAP Group DN

If a Group Type of "LDAP" is selected, this option specifies the Distinguished Name of the LDAP entry which holds the Roster Group.

## 8.2 LDAP Configurations

An LDAP Configuration defines how M-Link will use an LDAP directory for User and Group lookup. Multiple LDAP Configurations may be configured; the same Configuration may be used by multiple Instant Messaging Domains. An LDAP Configuration may have the following settings:

### Configuration Name

A user-friendly name assigned to the LDAP Configuration.

### LDAP URI

This encodes the address and port of the LDAP Directory server, for example "ldap://ldap.example.com:19389".

### Use TLS

Configures how TLS is used when connecting to the LDAP server.

### LDAP Authentication Identity

The identity used for authentication to the LDAP directory. This may be either a SASL identity or a string-encoded DN as appropriate.

### LDAP Authorization Identity

The identity used for authorization by the LDAP directory. This defaults to the LDAP Authentication Identity value.

### LDAP Password

The password associated with the LDAP Authentication Identity.

### Perform authentication by binding to the DSA

Select how user password validation takes place. If this option is selected, the server will bind to the LDAP server as the user; if using an Active Directory LDAP server you should always select this option. If the option is not selected, the server will read userPassword attributes from the LDAP server and perform validation locally;

the SCRAM-SHA-1 and PLAIN SASL mechanisms will be offered. Note that for this second option to work, the LDAP Authorization Identity must have read permission for userPassword attributes.

#### LDAP User Base

This configures the base for User searches. It is configured as a string-encoded Distinguished Name.

#### LDAP User Attribute

This configures the name of the LDAP attribute which holds a User's JID. If unset, the default used is "mail".

#### LDAP User Search Scope

This configures the scope of the search operation used to locate users beneath the LDAP User Base entry.

#### LDAP User Filter

This enables the configuration of an additional LDAP filter term to be used when searching for users beneath the LDAP User Base entry.

#### LDAP Group Base

This configures the Distinguished Name of the entry in the Directory beneath which LDAP Groups are located. This must be set for LDAP Group lookup to occur. An LDAP Group is any entry which contains "member", "uniqueMember" or "roleOccupant" attribute values.

#### LDAP Group Search Scope

This configures the scope of the search operation used to locate group entries beneath the LDAP Group Base entry.

#### LDAP Group Filter

This enables the configuration of an additional LDAP filter term to be used when searching for groups beneath the LDAP Group Base entry.

#### LDAP UID Rules

This allows the configuration of multiple regex/replacement pairs which operate on user account name values. For example, this could be used to normalize a "mail" attribute value of "alice@development.example.com" into a user account name of "alice@example.com".

---

**Note:** This item is located on the sidebar menu rather than on the page for the LDAP configuration itself.

---

---

## 8.3 User Provisioning Using Cobalt

Refer to the Cobalt User Manual for information on how to provision users.

# Chapter 9 HTTP File Upload

This chapter describes the HTTP File Upload feature.

---

## 9.1 HTTP File Upload

M-Link MU Gateway and M-Link MU Server support HTTP File Upload ([XEP-0363](#)), which provides a standardised way for users to share files using a [HTTP](#) service that is advertised by XMPP servers. Using this protocol, clients supporting the protocol can request a location to upload a file to, and share a (public) link to the uploaded file.

The HTTP File Upload service shares configuration settings with the M-Link Server running on the host system, and will automatically detect and apply changes relevant to its operation.

---

## 9.2 Configuration

By default, HTTP File Upload is disabled. It can be enabled through the M-Link Administration Interface on the main HTTP File Upload configuration page. A hostname must be provided, this is the [DNS](#) name that will be used for both uploading and downloading shared files. If [TLS](#) is enabled, a `TLS identity` must be selected as well. This identity should cover the configured `Default HTTP Hostname` to ensure clients will be able to use the service. Once enabled, this page also allows configuration of file size limits.

Once support for HTTP File Upload has been enabled, every configured IM Domain will support the feature, although this can be disabled per domain on the corresponding IM Domain section. Uploading files is by default limited to local users of the domain only, but can be exposed to users of any local or remote domain by disabling the `Limit to local users` option.

# Chapter 10 MUC Room Creation and Administration

This chapter introduces the *MUC* room concept and covers how to create and administer *MUC* rooms.

If the activated M-Link product supports MUC the M-Link Administration Interface will feature a "Multi-User Chat" tab. Under that, an administrator may create new *MUC* rooms and manage all existing *MUC* rooms.

---

## 10.1 MUC Room Concepts

---

**Note:** The Multi-User Chat protocol is defined in [XEP-0045](#).

---

A *Multi-User Chat* (MUC) room is an addressable chat for multiple users (those users currently “in” the room are known as “occupants”). Rooms belong to a domain hosted by an XMPP server and have a unique address combining a room name with the domain name: <room-name>@<domain-name>.

Once a hosted *MUC* domain has been created (see [Section 7.2.3.2, “MUC Domains”](#)), XMPP users may “join” rooms in that domain from their XMPP clients, possibly creating the rooms on the fly if they didn't exist depending on configuration. Rooms can be persistent or transient, with transient rooms being destroyed when the last “occupant” leaves them. Rooms created through the M-Link Administration Interface are persistent.

Rooms have “affiliation” lists, defining a relationship between the listed users and the room, used by the room to determine their initial roles and privileges upon joining. Please refer to [XEP-0045](#) for an exact definition of roles and privileges. In decreasing privilege order, the affiliation types are:

### Owner

Room owners have full control over the room configuration and existence. A user who creates a room in their XMPP client is automatically an owner of the room.

### Administrator

Users who are administrators can moderate other occupants, but not configure the room.

### Member

For “members only” rooms, a user may only enter if they have a member (or higher) affiliation.

### Outcast

Users affiliated as outcasts are banned from the room. This means they can neither see nor send any message exchanged in the room, and they get expelled from the room should they become outcasts while in it.

Affiliations are determined by two levels of configuration, domain and room levels. The lists refer to JIDs and can contain bare user JIDs or domain names (domain affiliations are equivalent to an affiliation for every user on that domain).

### Domain-Wide

A MUC domain administrator can configure affiliations that automatically apply to all rooms in the domain. For example, to ban a specific user from all rooms in the domain the user's JID can be added to the outcast affiliation list of the MUC domain.

### Per Room

Room specific affiliations will apply unless overridden by domain-wide affiliations in the following cases:

- A domain-wide affiliation exists, either for the same JID or for a domain or group it belongs to, and confers higher privileges.
- An outcast affiliation at the domain-wide applies to the user JID, also explicitly or by inclusion.

---

## 10.2 Room Administration

---

**Note:** The facilities for creation, configuration and removal of *MUC* domains can be found in the Configuration area under [Section 7.2.3.2, “MUC Domains”](#).

---

An administrator may use the M-Link Administration Interface to perform the following operations:

- View the configuration of all existing rooms in the MUC domains hosted by the server.
- Create rooms.
- Change default settings for new rooms.
- Change current settings of any room (room settings are not affected by changes in the defaults after their creation).
- Manage the list of affiliations associated with a room.
- Manage domain-wide affiliations that apply to all rooms within it, potentially overriding their own room specific affiliations.
- Monitor the status of rooms (e.g. number of occupants, *FMUC* status).
- Destroy rooms.

### 10.2.1 Available MUC Domains

The landing page under tab "Multi-User Chat" displays the MUC domains hosted by this M-Link server. The creation and removal of domains is a matter for server configuration. Refer to [Section 7.2.3.2, “MUC Domains”](#).

Each domain is represented by a card showing the domain name and the number of rooms it currently has. Clicking on a card leads to the Domain Management page for that domain. See the next section.

### 10.2.2 Domain Management

#### Listing Rooms in a Domain

A MUC domain management page displays the name of the domain and represents all its rooms as cards. A click on a card opens the associated room management pop-up. See [Section 10.2.4, “Change Room Settings”](#).

---

**Note:** This is a dynamic page and rooms can be transient. Please note the auto-refresh button and setting.

---

The page also provides access to domain level settings with two buttons beside the domain name:

#### Manage Domain...

This button opens a pop-up where domain-wide settings can be changed. Tabs in its header allow switching between the features:

#### Room Defaults

For details on domain-wide room defaults, refer to [Section 10.2.3, “Room Creation”](#).

#### Domain Affiliations

For details, refer to [Section 10.2.5, “Affiliation Management”](#).

Changes to the defaults or affiliations take effect upon clicking the “Update” button in the footer of the pop-up.

#### Create New Room...

This button opens the room creation pop-up.

For details on room creation, refer to [Section 10.2.3, “Room Creation”](#).

Clicking the “Add” button creates the room and dismisses the pop-up.

Rooms can also be destroyed from the room configuration pop-up, [Section 10.2.6, “Room Destruction”](#).

## 10.2.3 Room Creation

### 10.2.3.1 Domain-Wide Room Defaults

Room defaults are settings that will be given to new rooms upon their creation. They are defaults in the sense of an initial room configuration: Subsequent erasure of a room setting does not amount to specifying a value from the defaults and subsequent changes to the defaults do not affect existing rooms.

### 10.2.3.2 Room Creation by a MUC Administrator

The “Create New Room...” on a Domain Management page opens the room creation pop-up.

The first field is the mandatory name for the room, which must be unique within the domain. It must be valid as the node part of a JID.

Other fields will have been filled according to the room defaults described above. They may be modified before clicking the “Add” button.

### 10.2.3.3 Room Creation by an XMPP User

Users may attempt to join rooms that do not exist. A successful join (the facility depends on the XMPP client being used) means the room was created on the fly using the room defaults set up by a MUC administrator for the domain. It will be owned by the creating user.

XMPP clients may also provide means for users with owner affiliation (and administrator affiliation to a lesser extent) to change room settings as well.

## 10.2.4 Change Room Settings

From a domain management page, clicking on the card for a room opens a pop-up offering a number of tabs:

#### Room Configuration

This tab displays all the configurable settings of the room. Changes take effect when clicking the “Update” button.

#### Room Affiliations

This form works in exactly the same way as for domain-wide affiliations. Please refer to [Section 10.2.5, “Affiliation Management”](#).

#### FMUC Status

This tab is for monitoring purposes and only appears if FMUC is enabled in the “Global Options”. Please refer to [Section 10.2.7, “Room Status”](#).

## 10.2.5 Affiliation Management

Domain-wide and room-specific affiliations are maintained in the same way, so this section covers both. As summarized in [Section 10.1, “MUC Room Concepts”](#), domain-wide affiliations may take precedence over room-specific affiliations.

- All affiliation types are displayed by default (if any), but the list can be filtered according to the check-boxes in the “Filter by” box.
- The “Add Affiliation” button can be used to add new entries.
- Entries have a “Remove” button and an affiliation selector that allows subsequent affiliation changes without having to remove and re-insert the entry.
- Entries are sorted by JID and the sort order can be changed by clicking on the column header.

## 10.2.6 Room Destruction

The footer of the Room Configuration pop-up includes a “Destroy Room...” button.

---

**Note:** Possible path to get there:

- Click on the “Main Menu” link in the left side-bar (while in the “Multi-User Chat” area).
  - Click on the domain card.
  - Click on the room card.
- 

Destruction is effective when confirming with a click on the “Yes I’m sure” button.

## 10.2.7 Room Status

The cards display some status information about the rooms:

- Number of occupants.
- Number of joined *FMUC* rooms, or “N/A” if the room is not federated.
- If the room is federated, any remote nodes failing to join.

Clicking on the card for a room opens a pop-up which includes an “FMUC Status” tab. If applicable, this tab provides federation details which cannot be shown on the card.

---

## 10.3 MUC Federation

---

**Note:** The Federated Multi-User Chat protocol is defined in [XEP-0289](#).

---

Although XMPP provides a service federated over multiple servers, standard MUC rooms operate on a single server. This centralized MUC model works well in many situations, but is not ideal in all environments. In particular:

- Where there is a constrained link between a pair of servers, operating a MUC room on one server with clients on both servers leads to inefficient operation. It also prevents operation of remote clients when the link is not working.
- In a cross-domain environment, it is undesirable to have clients directly accessing a MUC room in the remote domain

Federated MUC (FMUC) addresses these issues. FMUC enables MUC rooms on multiple servers to federate and provide a single MUC room. This provides optimized performance for constrained links. More information on FMUC is provided in the Isode whitepaper "Federated Multi-User Chat: Efficient and Resilient Operation over Slow and Unreliable Networks", available at <https://www.isode.com/whitepapers/federated-muc.html>.

FMUC is enabled server-wide (see [Section 10.3.1, "Enabling FMUC for the server"](#)); individual MUC rooms can be configured to participate in a federated FMUC room, simply by including the names of the federated rooms to which messages are sent. FMUC is simply configured using a set of standard MUC rooms that will work together.

FMUC configuration is per-room, but first requires FMUC to be enabled for the server.

### 10.3.1 Enabling FMUC for the server

A control in the M-Link Administration Interface "Global Options" allows FMUC to be enabled for the M-Link server. Doing this makes a number of additional controls visible:

#### FMUC Connection Retry Timer

The interval (in seconds) at which the M-Link server will attempt to establish federation between MUC rooms which are configured to be federated but have become detached. Configuring a value of zero will disable background reconnection attempts.

#### FMUC Ping Timer

The interval (in seconds) at which XMPP "ping" stanzas will be sent to each MUC room's federated nodes.

#### FMUC Stanza Timeout

The period (in seconds) after which an FMUC stanza (including a "ping" stanza) will time out. Stanza timeout results in federated room detachment.

#### FMUC Join Timeout

The period (in seconds) after which an FMUC join attempt times out.

### 10.3.2 Configuring FMUC for a room

Once FMUC has been enabled for the server, a set of FMUC options are made available on each MUC room configuration form. These are:

- **"Allow FMUC from arbitrary sources"** will, if enabled, accept requests to start federating from any other MUC, irrespective of whether it is in the configured node list or not.
- **"Federated MUC nodes"** is a list of the MUCs that this MUC is configured to federate with - note that while it is possible to federate with many nodes, there must be no 'loops' in the config (that is: it must be an acyclic graph). So to have the rooms *demo@conference.example.com* and *sample@rooms.server.local* federate you would add *sample@rooms.server.local* to the **Federated MUC nodes** list of *demo@conference.example.com*, and add *demo@conference.example.com* to the **Federated MUC nodes** list of *sample@rooms.server.local*.
- **"Add remote FMUC domain name to nicks"** controls whether the remote FMUC domain name should be appended to the nickname for those MUC occupants who are present via federation.
- **"History requested from FMUC"** determines how many historic messages should be requested when MUC federation occurs.

### 10.3.3 Things to note

- When a federated MUC becomes unavailable it may take some time before the link is considered 'dead' and remote occupants in the MUC are marked as having left the room. It is possible for occupants of different nodes of an FMUC room to see

messages arrive in a slightly different order when either messages are sent at the same time or the network link between servers is slow.

- You can only federate between local and remote MUC rooms, you cannot federate between two MUC rooms hosted on the same M-Link instance.
- If FMUC is enabled for the server, anyone able to create a MUC room can then configure it to federate with a remote host - this provides a mechanism for potential abuse if your users are not trusted.

# Chapter 11 M-Link Archive Server

This chapter describes the M-Link archiving capabilities.

---

## 11.1 M-Link Archive Server

The M-Link Archive Server is used to archive the XMPP traffic on behalf of M-Link Server. Communication between the M-Link Archive Server and the M-Link Server uses a high performance protocol.

Most stanzas that pass through the M-Link Server using XMPP, covering all types of traffic, are sent to the configured M-Link Archive Server. Key state changes are also sent.

The archiving mechanism is asynchronous, and so normal processing of XMPP traffic is not affected by database writes (or reads). The M-Link Server writes queued archive data to disk until they have been transferred to the M-Link Archive Server and acknowledged. This makes archiving resilient to outages and restarts of both M-Link Archive Server and M-Link Server processes.

The M-Link Archive Server normally shares configuration settings with the M-Link Server running on the host system. The M-Link Archive Server server needs to be restarted if any configuration settings which affect it are modified.

---

## 11.2 User Access to Archives

End user access to the archives is provided using the XMPP extension protocol called Message Archive Management (MAM). XMPP clients that implement MAM issue queries to the M-Link Server which in turn queries the M-Link Archive Server. Standard XMPP authentication and access control applies, so users can only see traffic that they were party to or could have been party to.

Isode provides the Message Archive Browser, an XMPP over BOSH client that uses MAM to provide user access to the archives.

---

## 11.3 Archive Data and Configuration

M-Link Administration Interface allows you to configure archiving for all IM and MUC domains within a server. The following configuration options are available:

### Use Message Archiving

This option turns on message archiving.

---

**Note:** All chat and groupchat messages are currently archived if this option is enabled.

---

### Archive Queue Path

This is the location in which messages which have not yet been acknowledged as archived are stored.

---

**Note:** Changing this path will not automatically move any existing unarchived items.

---

#### Queue Timeout

The minimum amount of time (in milliseconds) the server waits for acknowledgement from the archiving service until it persists messages to an on-disk queue.

---

## 11.4 Message Archive Browser Configuration

The configuration file for Message Archive Browser is located at  $$(VARDIR)/webapps/config.js$ . A sample config file is provided at  $$(SHAREDIR)/webapps/config_sample.js$ , copying this file and renaming should be sufficient for most deployments.

The configuration is stored in a JavaScript file in an Object variable named `cfg`; it is important that you do not change the name of this variable.

#### boshURL

This will be used to prefill the BOSH URL on the login form.

#### modules

This is a JavaScript Object using the module name as the key and a Boolean value to indicate whether the module should be shown or not. There is currently one module available: `mam`.

# Chapter 12 Monitoring

This chapter introduces the features available for monitoring the state of the server.

---

## 12.1 Sessions Listing and Monitoring

Sessions connecting the local M-Link server to remote servers or to clients have different properties and are monitored in a similar way but on separate pages. Hence the following two sub options or cards in the `Monitoring` area for displaying a list of current sessions of the corresponding kind:

- `Server Sessions`
- `Client Sessions` (only available for M-Link products that support *Instant Messaging* (IM))

The lists are snapshots: use the refresh button provided on the page to get an updated list, or enable `Auto Refresh` and choose a refresh rate from the drop-down menu beside it.

The second half of the page displays traffic data from the sessions that have been selected using the check boxes in the left hand side column. It displays stream fragments (XML) carried by the sessions since they were selected. This is transient data which is lost when leaving the page. Deselecting a session stops data collection for that session and also removes any already collected data from the display. All traffic data can also be reset by clicking its `Clear` button (with a trash icon).

To monitor traffic data from new sessions as they are created, use the checkbox at the top. Checking it has the effect of automatically checking the monitoring checkbox for any new session created while this top checkbox has a check mark.

---

**Warning:** Traffic data may accumulate while automatically monitoring new sessions.

---

Unchecking the top checkbox will only disable the automatic selection facility; it will not suspend monitoring for sessions being monitored.

The session list may be empty if there are no sessions of the selected kind.

### 12.1.1 Server Session Properties

The list of server sessions presents the following session properties:

**A monitoring checkbox**

This checkbox controls whether or not traffic from the session is included in the traffic data monitor window in the second half of the page.

**SID**

A unique identifier for the session.

**Type**

A code for the session type. There are three types of sessions supported at the moment: `S2S`, `X2X`, `GCXP`. For details about these types see [Section 4.2, “Peer Controls”](#) and [Section 4.6, “Links”](#).

**Dialback**

When an incoming session authenticates using the dialback mechanism, a separate session is needed to interrogate the calling server according to DNS. If no session to

the relevant server already exists, a new outgoing session is initiated which will be indicated in this column.

#### TLS

This column indicates whether the session is protected with TLS.

#### Authentications

This lists the XMPP domain endpoint pairs (local and remote) for which this session has been authenticated (allowing stanza routing). Specialised connection mechanisms authenticate based on Peer Control configuration and for those endpoint pairs will appear here once stanzas for that pairing have been exchanged. When Relaying is in use the 'local' domain of the pair will be the domain on to which M-Link will be forwarding the received stanzas, or on behalf of which it is forwarding.

#### Initiated

This column indicates whether the local server (Outbound) or the remote server (Inbound) initiated the session.

#### Direction

This column indicates whether the session is able to send (Outgoing) or receive (Incoming) stanzas, or both (Bidirectional).

#### Local Port

Sessions are underpinned by a TCP/IP socket uniquely identified by IP address and port number. This column uses the standard IP notation (IPv4 or IPv6 as applicable) with port separated by a colon.

#### Remote Port

As the Local Port, but indicating the IP/port pair used by the remote server.

#### Created On

The moment the session was created in the web browser's locale.

#### End

The last column includes a button that allows termination of the session. Clicking on the button only displays a confirmation dialogue from which it is still possible to step back.

## 12.1.2 Client Session Properties

The list of client sessions presents the following session properties:

#### A monitoring checkbox

This checkbox controls whether or not traffic from the session is included in the traffic data monitor window in the second half of the page.

#### SID

A unique identifier for the session.

#### Type

Indicates whether the session uses the BOSH protocol or is transported by TCP only.

#### JID

The Jabber identifier used by the connected client. This will include a resource part.

#### TLS

This column indicates whether the session is protected with TLS.

#### Local Port

Sessions are underpinned by a TCP/IP socket uniquely identified by IP address and port number. This column uses the standard IP notation (IPv4 or IPv6 as applicable) with port separated by a colon. This field is blank for BOSH sessions.

#### Remote Port

As the Local Port, but indicating the IP/port pair used by the client.

#### Created On

The moment the session was created in the web browser's locale.

End

The last column includes a button that allows termination of the session. Clicking on the button only displays a confirmation dialogue from which it is still possible to step back.

# Chapter 13 HTTP API

This chapter covers the APIs available for programmatic configuration of M-Link.

## 13.1 Authentication HTTP API

M-Link supports two authentication methods:

- Access by login, which is used by the M-Link Administration Interface, and
- access by application access tokens, which are meant for external application programmatic access.

### 13.1.1 Logging in as an Administrator

URL	/api/login
Method	POST
Request Content-Type	application/json

#### 13.1.1.1 Description

The method allows user authentication based on a cleartext password.

#### 13.1.1.2 Request

The request body contains values for username and password in the following JSON format:

```
{
  "username": "admin",
  "password": "adminpass"
}
```

#### 13.1.1.3 Response

The response has no content. If successful, it includes a Set-Cookie header setting the “M-LinkLoginToken<admin port>” cookie.

This cookie must be included in the Cookie header of all subsequent requests that require it, as specified in their related sections. Failure to include the loginToken cookie will fail these requests with a status code of 403 (FORBIDDEN) and no content in the response body.

#### 13.1.1.4 Status Codes

204 (NO CONTENT)	Correct credentials were provided and cookie “M-LinkLoginToken<admin port>” is set as HTTP only (so the browser environment does not have access to it) and possibly secure as well (cookie forbidden on HTTP) depending on option secureCookie.
403 (FORBIDDEN)	The credentials provided were incorrect.

## 13.1.2 Acquiring the Security Token

URL	/api/token
Method	GET
Response Content-Type	application/json

### 13.1.2.1 Description

Provides a CSRF token to authenticate users.

### 13.1.2.2 Request

The request has no content.

### 13.1.2.3 Response

If a valid login token is set, responds with 200 (OK) and a body containing the token as in the example below.

```
{
  "token": "..."}

```

The client should store this token for later use in the X-Auth-Token header of requests that require it (please refer to their own sections). Failure to include this header with a valid token will fail these requests with a status code of 403 (FORBIDDEN) and no content in the response body.

Clients SHOULD fetch a new CSRF token after an hour, as they expire after two hours.

### 13.1.2.4 Authentication

The request must have a cookie obtained from the /api/login end point. Failure to provide a valid loginToken cookie will cause a return with an error status of 403 (FORBIDDEN) and an empty response body.

### 13.1.2.5 Status Codes

200 (OK)	A valid login token was set. The body contains: { "token": "..."}.
403 (FORBIDDEN)	No login token was present in the request or it was invalid.

## 13.1.3 Who Is Logged in?

URL	/api/me
Method	GET
Response Content-Type	application/json

### 13.1.3.1 Description

This end point returns the user name of the currently logged in administrator. The request must have a cookie obtained from /api/login.

### 13.1.3.2 Request

The request has no content.

### 13.1.3.3 Response

The response body will be a JSON object. If no initial administrator has been set up, the "username" will not be present. The response object may additionally hold values for internal use by the M-Link Administration Interface, the format of which may vary between releases, and should therefore not be relied upon.

```
{
  "username": "$username"
}
```

### 13.1.3.4 Authentication

The request must have a cookie obtained from the /api/login end point, if the initial admin user has been set up. Failure to provide a valid loginToken cookie then causes a return with an error status of 403 (FORBIDDEN) and an empty response body.

### 13.1.3.5 Status Codes

200 (OK)	If the current session is fully authenticated, the response will contain a "username".  Servers that do not have any administrative users configured yet will also respond with a 200 (OK), but the response object will not contain a "username".
403 (FORBIDDEN)	The initial admin user is set up, but the login token cookie was missing or invalid. The body contains no "username".

## 13.1.4 Changing Administrator username and/or password

URL	/api/change_credentials
Method	POST
Request Content-Type	application/json

### 13.1.4.1 Description

The method allows the administrative user's name or cleartext password to be changed. The current administrative username and password must be provided in addition to the desired new password and optionally new username.

### 13.1.4.2 Request

The request body contains values for existing and new username and password in the following JSON format:

```
{
  "username": "admin",
  "password": "adminpass",
  "newUsername": "administrator",
  "newPassword": "secret"
}
```

### 13.1.4.3 Response

The response has no content.

### 13.1.4.4 Authentication

The request must have a cookie obtained from the `/api/login` end point and a token obtained from the `/api/token` end point. Alternatively, the request may include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.1.4.5 Status Codes

204 (NO CONTENT)	Correct current credentials were provided and the new username and password have been stored.
422 (UNPROCESSABLE ENTITY)	The request body was invalid, or is missing required information, or the current credentials provided were incorrect.

## 13.1.5 Authentication by Access Token

As an alternative to the login cookie and, where applicable, the security token, external applications may access M-Link end points programmatically by including an Authorization header in their HTTP requests. See [RFC-7235](#).

---

**Tip:** Syntax reminder:

```
Authorization: <scheme-name> <credentials>
```

Please note: The scheme name and credentials string are separated by a single space.

---

M-Link uses a proprietary authentication scheme named `Isode-API-Token` (scheme names are case insensitive), which carries a token string as credentials. For example:

```
Authorization: Isode-API-Token xLSNPnA+eJIImp2TKxt7qLHz0791SNaT
```

Application access tokens can only be obtained manually using the M-Link Administration Interface (see section [Access Tokens](#)).

---

## 13.2 Configuration HTTP API

This section describes the HTTP REST interface which allows third party clients to configure options in the option database.

The options can be thought of as held in a [JSON](#) document, described by a [JSON Schema](#), which can be obtained from the `/api/config/schema` end point. The standard Schema has been extended with the following elements (please note: these are not always valid in combination):

```
{
  "properties": {
    "choice": {
      "isode_label": "someName",
      "isode_required": true,
      "isode_validator_type": "aValidatorName",

```

```

        "isode_multi_line": true,
        "isode_underlying_type": "NotNegativeExample",
        "isode_enum_titles": [
            "Display This Instead of the First Enum Value",
            "Display This Instead of the Second Enum Value"
        ],
        "isode_default_comes_from": "JSON Pointer",
        "isode_only_if": [
            {
                "check": "is",
                "ref": "1/enabled",
                "value": true
            }
        ],
        "isode_store_uri": "/api/store/someObjectStore",
        "isode_substore_item": true,
        "isode_icon": "address-book"
    }
}
}
}

```

#### isode\_label

This property is appropriate for lists of complex options (i.e. with nested sub options). It refers by name to one of the sub options, singling it out as distinctive of the list item. For example, the label might designate a sub option meant to be unique to each item and thus suitable as an item identifier. Lists of simple options do not require any label as they only carry one value per item.

#### isode\_multi\_line

Indicates new-line characters in the option value are to be preserved.

#### isode\_validator\_type

Denotes a validator which constrains string values to certain formatting rules. When taking input for an option with a validator, the GUI should validate it using end point `/api/config/validator`.

#### isode\_underlying\_type

An informative property naming the underlying type of the option. It might be used to identify options that the end point consumer can validate by itself. However, it should only appear along side the "isode\_validator\_type" property.

#### isode\_enum\_titles

Short definitions of the choices identified by an enum option. The option should also have the standard "enum" property, listing the enum values in the same order.

#### isode\_only\_if

In the absence of any isode\_only\_if, the values of options fetched from the `/api/config` end point are either their default if unset, or the value they were last set to. An isode\_only\_if property is a boolean value computed as the logical AND of all the conditions it lists. The clauses in the isode\_only\_if connect the value of an option, referred to by a [JSON Pointer](#) (the "ref" element) to a constant value (the "value" element) with an operator (the "check" element) which is either "is" or "isNot". If the isode\_only\_if property is true, the behaviour is the same as when there is no isode\_only\_if property. If it is false, fetching the option value returns its default, whether or not the option was explicitly set. Please note that there are several ways to generate a default. If the option is complex (that is, it has nested sub options), a false isode\_only\_if property is logically applied to all the sub options, overriding any explicit "isode\_only\_if" on sub options. If the option is a list, a false "isode\_only\_if" makes it appear empty. While the "isode\_only\_if" is false, set values are not lost and options may still be set, modified or unset.

#### isode\_default\_comes\_from

Extends schema property default: instead of a value, refers to another option, via a [JSON Pointer](#), from which to collect a default value. In cases where a default value is returned (e.g. fetching an unset option or fetching an option with false "isode\_only\_if"), "isode\_default\_comes\_from" takes precedence over "default".

Multiple options may be chained with "isode\_default\_comes\_from". In other words, the target of an "isode\_default\_comes\_from" may be unset or have a false "isode\_only\_if", in which case its own default is retrieved, wherever that comes from. Please note that some types have a natural default, meaning that even without any "isode\_default\_comes\_from" or "default" a fetch will succeed. Other types must have some form of explicit default.

#### isode\_required

Indicates that a value must be provided when patching the configuration. Defaults to false if missing on an option. This is similar to the standard "required" property. However, this extension accommodates the "isode\_only\_if" extension: false "isode\_only\_if" conditions effectively disable the requirement to be set, as indeed only the default can then be fetched. In other words, unset options with "isode\_required" but a false "isode\_only\_if" do conform to the schema (which would not be the case with a standard "required" property).

#### isode\_store\_uri

Indicates that the option value is the unique identifier of an object held in an object store. Object stores are persistent containers of objects of a type specific to the store. These can be manipulated through a separate API (see section [Object Stores HTTP API](#)). The value of the isode\_store\_uri property is the path to the object store according to that API, relative to the server URI.

#### isode\_substore\_item

If true, indicates that the option value is the unique identifier of an object within a substore. Substores are object stores which are themselves items of an object store. Please refer to section [Object Stores HTTP API](#). The identifier only makes sense when knowing which substore it belongs to: Options with this property are expected to exist within a complex option also including a sub option with the same name plus "Substore" and providing the required information (see property isode\_store\_uri). Property isode\_substore\_item is incompatible with isode\_underlying\_type: Types with the latter are expected to support a validator (implicitly or explicitly), however substore items require validators which cannot be triggered via the /api/config/validator endpoint. They may still have validators applied by the server after calling PATCH /api/config, but those are not declared in the [JSON Schema](#).

#### isode\_icon

The value is only a hint to the UI client, drawing meaning by convention: If it is appropriate for the option, the UI may interpret the value as the class name of an icon/glyph within some version of FontAwesome (tm) and use it where appropriate in relation to the rendering of that option.

Available operations are described below. Some end points allow multiple methods. They allow configuration retrieval, option update through [JSON patches](#) and validation. A status code of 200 (OK) implies that the request was successful.

## 13.2.1 Obtaining the JSON Schema

URL	/api/config/schema
Method	GET
Response Content-Type	application/json

### 13.2.1.1 Description

As mentioned in the overview of [Section 13.2, "Configuration HTTP API"](#), the set of configuration options presents as a [JSON](#) document. In order to set values properly and derive defaults, it is necessary to retrieve a [JSON Schema](#) describing the whole document.

### 13.2.1.2 Request

Fetches the [JSON Schema](#) for the application options. The request body is empty.

### 13.2.1.3 Response

Returns the [JSON Schema](#) as a [JSON](#) document in the body.

Example:

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "description": "Server configuration",
  "id": "http://isode.com/api/config/schema.json",
  "properties": {
    "useTls": {
      "default": true,
      "description": "Boolean value",
      "type": "boolean"
    },
    "max_stanza_size": {
      "default": 10024,
      "description": "Maximum size of a stanza",
      "maximum": 9223372036854775807,
      "minimum": -9223372036854775808,
      "type": "integer"
    }
  },
  "propertyOrder": [
    "max_stanza_size",
    "useTls"
  ],
  "title": "Config",
  "type": "object"
}
```

### 13.2.1.4 Authentication

The request must have a cookie obtained from the `/api/login` end point or include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.2.1.5 Status Codes

Status codes relating specifically to this method:

200 (OK)	Schema successfully retrieved.
	Failure codes are possible, but only due to higher layer processing (such as authentication issues).

## 13.2.2 Retrieving the Current Configuration

URL	<code>/api/config</code>
Method	GET
Response Content-Type	<code>application/json</code>

### 13.2.2.1 Description

Retrieves the current configuration in [JSON](#) format. This configuration format is described in the [JSON Schema](#) provided by the `/api/config/schema` end point. Please note that the document may be incomplete: where options have not been explicitly set they are omitted, as their values can be derived either from the default included in the

schema, or from an `isode_default_comes_from`, or the type has a natural default, such as strings which can be empty.

### 13.2.2.2 Request

Fetches the current configuration in [JSON](#) format. The request body is empty.

### 13.2.2.3 Response

The [JSON](#) document returned in the body conforms to the [JSON Schema](#). It only includes explicitly set options.

Example:

```
{
  "max_stanza_size": 10000
}
```

### 13.2.2.4 Authentication

The request must have a cookie obtained from the `/api/login` end point or include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.2.2.5 Status Codes

Status codes relating specifically to this method:

200 (OK)	Successfully retrieved the configuration.
	Failure codes are possible, but only due to higher layer processing (such as authentication issues).

## 13.2.3 Applying Changes to the Configuration

URL	<code>/api/config</code>
Method	PATCH
Request Content-Type	<code>application/json-patch+json</code>
Response Content-Type	<code>application/json</code>

### 13.2.3.1 Description

Applies changes, described in a [JSON Patch](#) document, to the current configuration.

### 13.2.3.2 Request

The request body should contain a valid [JSON Patch](#) document describing the changes to be applied to the configuration, seen as a [JSON](#) document conforming to the [JSON Schema](#) provided by the `/api/config/schema` end point. A [JSON Patch](#) is an array of objects, each describing a single operation to be performed on a specific option referred to by its path (a [JSON Pointer](#)).

Only a sub set of [JSON Patch](#) operations are supported: replace, add, remove.

replace

This operation sets an existing single option to a value. It fails if the path leads to a complex or list option.

Example:

```
[
  {
    "op": "replace",
    "path": "/max_stanza_size",
    "value": 99000
  }
]
```

**add**

This operation adds an item to a list. The path must lead to a list option with '/' appended to it (suggesting the index number beyond the current last). If the item is complex, the value is a compound **JSON** object matching the sub tree. Only list items can be added as the option tree is fixed.

Example: N.B. the list assumed below does not feature in the simple schema example given out above:

```
[
  {
    "op": "add",
    "path": "/imDomains/0/users/-",
    "value": {
      "jid": "alice@wonderland.lit",
      "password": "secret"
    }
  }
]
```

**remove**

This operation removes an item from a list. The path must lead to an existing list item. Other options are not removable.

Example: N.B. the list assumed below does not feature in the simple schema example given out above:

```
[
  {
    "op": "remove",
    "path": "/imDomains/0/users/5"
  }
]
```

**13.2.3.3 Response**

If the patch was successfully applied, the new configuration is returned in the response body in the same way as in response to a GET /api/config.

Example:

```
{
  "max_stanza_size": 99000,
  "useTls": true
}
```

If the patch failed, the body is empty.

**13.2.3.4 Authentication**

The request must have a cookie obtained from the /api/login end point and a token obtained from the /api/token end point. Alternatively, the request may include a valid access token delivered in an Authorization header. Failure to comply will cause a return

with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.2.3.5 Status Codes

Status codes relating specifically to this method:

200 (OK)	The patch was successfully applied.
415 (UNSUPPORTED MEDIA TYPE)	The content type or the <a href="#">JSON</a> syntax was invalid.
422 (UNPROCESSABLE ENTITY)	The <a href="#">JSON Patch</a> provided was invalid or could not be applied.
	Other failure codes are possible, but only due to higher layer processing (such as authentication issues).

## 13.2.4 Validating Option Values

URL	/api/config/validator
Method	POST
Request Content-Type	application/json
Response Content-Type	application/json

### 13.2.4.1 Description

This end point provides a facility to validate a particular option value before submitting it to end point /api/config with PATCH. Most value types are fully defined in the [JSON Schema](#). For example, integers might have a maximum and minimum. Other types are defined as strings but are actually more elaborate and/or have specific formatting restrictions that cannot be easily expressed in the schema. These will have a "validator\_type" property with which to query the option database through this end point.

### 13.2.4.2 Request

The body is a [JSON](#) document defining a type and value to validate. The type is the value of a "validator\_type" property and should be recognized by the option database.

Example:

```
{
  "type": "typeName",
  "value": "valueToTest"
}
```

### 13.2.4.3 Response

If the validator ("type" in the request) is supported and the value valid, the body contains:

```
{
  "valid": true
}
```

If the validator is supported but the value invalid, the body content resembles the following, the value of userMessage being supplied by the validator.

```
{
```

```

    "valid": false,
    "userMessage": "a hint"
  }

```

#### 13.2.4.4 Authentication

The request must have a cookie obtained from the `/api/login` end point or include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

#### 13.2.4.5 Status Codes

Status codes relating specifically to this method:

200 (OK)	The validator is supported. If the value is valid, the body contains: { "valid": true }, otherwise: { "valid": false, "userMessage": "a hint" }
400 (BAD REQUEST)	Unsupported content type or unsupported validator. In the latter case, the body contains: { "error": "message" }
	Other failure codes are possible, but only due to higher layer processing (such as authentication issues).

---

## 13.3 Object Stores HTTP API

This section describes the HTTP end points for the management of Object Stores. All Object Store end points require a valid login token. The calls that modify the store or its contents will additionally require a CSRF token.

The API provides access to two kinds of stores: Object Stores and Object Substores. Object Substores provide another hierarchy layer to Object Stores where each item in the Object Store is another Object Store containing items - e.g. every Security Label Catalog in the Security Label Catalog Store can be addressed as a substore, providing access to all the labels in the catalog. Substores are immutable, so any operation attempting to modify a substore will fail.

Each available Store has a base path associated with it. Store paths will be represented by `/storepath` in this section. The base paths for most stores can be found from the `GET /api/stores` endpoint.

The base paths for substores are generated by taking the path of the base store, `/storepath`, appending the id of the item in the base store that you want to use as a substore, and appending `/sub_store`. For example, if a Security Label Catalog store was located at `/api/stores/label_catalogs`, and contained a catalog with id 'a64ef', `/api/stores/label_catalogs/a64ef/sub_store` would be the base path for a substore containing the labels provided by that catalog.

### 13.3.1 List available Object Stores

URL	<code>/api/stores</code>
Method	GET
Response Content-Type	application/json

### 13.3.1.1 Description

This method can be used to discover the object stores that are available.

### 13.3.1.2 Request

The request has no content.

### 13.3.1.3 Response

Returns an array of objects. Each object has at least the members `name`, `path`, `description`, and `upload_fields` as shown. The `path` member represents the base path of each store. If the `has_substores` member is present and set to true, this store has substores. The `upload_fields` member is an array of specifications of the fields which the store requires when creating a new entry. In the simple case, this is a single file containing the data of the entry, but more complex stores may require multiple components to be supplied (e.g. a certificate chain, private key and passphrase for a TLS identity).

```
[
  {
    "name": "Example Store",
    "path": "/api/stores/example",
    "description": "All available examples are stored here.",
    "upload_fields": [
      {
        "name": "object",
        "type": "file",
        "title": "Filename",
        "description": "Choose a file to upload",
        "accept": ".pem,.crt",
        "required": true
      }
    ]
  },
  {
    "name": "Example Store Containing Substores",
    "path": "/api/stores/example_catalogs",
    "description": "Example catalogs are stored here.",
    "has_substores": true
  }
]
```

### 13.3.1.4 Authentication

The request must have a cookie obtained from the `/api/login` end point or include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.3.1.5 Status Codes

200 (OK)	The request was successful.
403 (FORBIDDEN)	The login token cookie was missing or invalid.

## 13.3.2 List all objects in a Store or Substore

URL	<code>/storepath</code>
Method	GET
Response Content-Type	application/json

### 13.3.2.1 Description

This call returns a list of all items in the specified Object Store or Object Substore.

### 13.3.2.2 Request

The request has no content.

### 13.3.2.3 Response

Returns an array of objects. Each object consists of the members `id`, `name` and an optional `creationDate`, in seconds since 1970-01-01T00:00:00Z. `name` can be used for display purposes. The `id` of a given object will not change, and is used in methods that operate on a specific stored item.

```
[
  {
    "id": "7b5237b4-c3ea-4f9f-8d16-20eb96486b0e",
    "name": "An example object",
    "creationDate": 1582798994
  },
  {
    "id": "174edbfaf-e2b3-4e5c-b5bb-d3783a364e3d",
    "name": "Another object",
    "creationDate": 1585582903
  }
]
```

### 13.3.2.4 Authentication

The request must have a cookie obtained from the `/api/login` end point or include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.3.2.5 Status Codes

200 (OK)	The request was successful.
403 (FORBIDDEN)	The login token cookie was missing or invalid.
404 (NOT FOUND)	The requested store does not exist.

## 13.3.3 Add an object to the Store

URL	<code>/storepath</code>
Method	POST
Request Content-Type	multipart/form-data
Response Content-Type	javascript/json

### 13.3.3.1 Description

This call can be used to add an object to the specified store.

The object content is wrapped within a [multipart/form-data](#), as explained in the request section below. Its `filename` parameter is meant to record the name of the file the object is uploaded from. The name of the new object in the store may also be set using a `name` field, but defaults to the filename. It is returned in the response as the value of property `"name"`, along with property `"id"` which is generated by the POST operation as a unique identifier. Other attributes may also exist, but those are object type specific.

### 13.3.3.2 Request

The call expects a valid [multipart/form-data](#) payload.

You can provide form values for each field mentioned in the `upload_fields` entry from the HTTP GET which discovered the Object Store. Fields marked as required must always be sent.

Fields of type `file` must have their `filename` attribute set. When multiple fields of type `file` are submitted, the object will be stored using the first filename encountered.

This is compatible with regular HTTP file upload mechanisms found in most web clients. Multiple objects may thus be stored in a single request. Objects should a Content-Type values of `application/octet-stream`, or something more specific.

The [multipart/form-data](#) payload may also include a field named `name`. Its value is an arbitrary name to be associated with the object. It must have a Content-Type of `text/plain` with charset of UTF-8, but these are default values. If omitted, the `filename` parameter of the first field which contains one will be used.

### 13.3.3.3 Response

If the call succeeds, the same object information as returned by 'Get object information'. In case of an error, the body will be an object with a `message` member with a user-presentable error.

### 13.3.3.4 Authentication

The request must have a cookie obtained from the `/api/login` end point and a token obtained from the `/api/token` end point. Alternatively, the request may include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.3.3.5 Status Codes

200 (OK)	The request was successful.
403 (FORBIDDEN)	The login token cookie was missing or invalid.
404 (NOT FOUND)	The requested store does not exist.
422 (UNPROCESSABLE ENTITY)	The request body was invalid, or is missing required information.

## 13.3.4 Retrieve an object from the Store

URL	<code>/storepath/id</code>
Method	GET
Response Content-Type	<i>Content specific</i>

### 13.3.4.1 Description

This call returns the object data for an item in an Object Store. Unless the store holds composite objects which are not convertible to a single file, in which case this feature is disabled and a retrieval attempt yields an error.

### 13.3.4.2 Request

The request has no content.

### 13.3.4.3 Response

Returns the data associated with object *id* in the specified store. The Content-Type will depend on the object stored. In case of an error, the Content-Type will be `application/json`, and the body will be an object with a `message` member with a user-presentable error.

### 13.3.4.4 Authentication

The request must have a cookie obtained from the `/api/login` end point or include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.3.4.5 Status Codes

200 (OK)	The request was successful.
403 (FORBIDDEN)	The login token cookie was missing or invalid.
404 (NOT FOUND)	The requested store or item does not exist or the store does not support content retrieval.

## 13.3.5 Rename an object

URL	<code>/storepath/id</code>
Method	POST
Request Content-Type	<code>javascript/json</code>
Response Content-Type	<code>javascript/json</code>

### 13.3.5.1 Description

This call updates the name of the specified object in the Object Store.

### 13.3.5.2 Request

The request body must be an object with a single `name` member that holds the desired new name of the object.

```
{
  "name": "New object name"
}
```

### 13.3.5.3 Response

If the call succeeds, the object identified by *id* will have its `name` updated. In case of an error, will return an object with a `message` member with a user-presentable error.

Please note that the call will fail if the object is read-only (see section [Get object information](#)).

### 13.3.5.4 Authentication

The request must have a cookie obtained from the `/api/login` end point and a token obtained from the `/api/token` end point. Alternatively, the request may include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.3.5.5 Status Codes

204 (NO CONTENT)	The request was successful.
403 (FORBIDDEN)	The login token cookie was missing or invalid.
404 (NOT FOUND)	The requested store or item does not exist.
422 (UNPROCESSABLE ENTITY)	The request body was invalid, is missing the required information, or refers to a read-only object.

## 13.3.6 Remove an object from the Store

URL	<i>/storepath/id</i>
Method	DELETE
Response Content-Type	javascript/json

### 13.3.6.1 Description

This call removes the specified object from the Object Store.

### 13.3.6.2 Request

The request has no content.

### 13.3.6.3 Response

If the call succeeds, the object identified by *id* will no longer be present in the Object Store. In case of an error, will return an object with a `message` member with a user-presentable error.

Please note that the call will fail if the object is read-only (see section [Get object information](#)).

### 13.3.6.4 Authentication

The request must have a cookie obtained from the `/api/login` end point and a token obtained from the `/api/token` end point. Alternatively, the request may include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.3.6.5 Status Codes

204 (NO CONTENT)	The request was successful.
403 (FORBIDDEN)	The login token cookie was missing or invalid.
404 (NOT FOUND)	The requested store or item does not exist.
422 (UNPROCESSABLE ENTITY)	The requested item is still in use in an active configuration, or is read-only.

## 13.3.7 Get object information

URL	<i>/storepath/id/info</i>
Method	GET
Response Content-Type	application/json

### 13.3.7.1 Description

This call returns all available additional information on an object in the specified store.

### 13.3.7.2 Request

The request has no content.

### 13.3.7.3 Response

Returns the information as an object with at least the following members: `id`, `name`.

Further there may be the following members: `creationDate`, `filename`, `contentType`, `readOnly`.

`id`

The object's unique identifier generated when it was added to the store.

`name`

The name chosen for the object. This may have defaulted to the filename of origin, or may have been set when adding the object or using the rename procedure.

`filename`

The filename given when adding the object.

`contentType`

The MIME content type given when adding the object.

`creationDate`

The date and time the object was added. This is conveyed as a number of seconds since the Epoch (1970-01-01T00:00:00Z).

`readOnly`

If this member is present and set to `true`, the object is read-only and cannot be removed, renamed or overwritten.

Additionally, there may be an `attributes` member with an object as value. This object will hold additional information; which attributes are available depends on the store. Attributes which have empty values will not be returned, and this may mean that the `attributes` member is empty.

```
{
  "id": "7b5237b4-c3ea-4f9f-8d16-20eb96486b0e",
  "name": "An example object",
  "filename": "example.xml",
  "contentType": "text/xml",
  "creationDate": 1574091301,
  "attributes": {
    "valid_from": "2020-01-01Z00:00:00"
  }
}
```

### 13.3.7.4 Authentication

The request must have a cookie obtained from the `/api/login` end point or include a valid access token delivered in an Authorization header. Failure to comply will cause a return with an error status (such as 403 (FORBIDDEN)); please refer to the [Authentication HTTP API](#) for information on how to get the tokens.

### 13.3.7.5 Status Codes

200 (OK)	The request was successful.
403 (FORBIDDEN)	The login token cookie was missing or invalid.
404 (NOT FOUND)	The requested store or item does not exist.

# Appendix A Integration with Isode Icon-Topo

This appendix describes how to integrate M-Link MU Gateway with Isode Icon-Topo.

---

## A.1 Enabling support for Icon-Topo

Support for integration with Icon-Topo is enabled using the *Enable Icon-Topo Support* option in *Global Options*.

Isode Icon-Topo will need an *Access Token* to interface with M-Link MU Gateway. This token can be generated using the part of the Management Interface described in [Section 7.2.3.9, “Access Tokens”](#). Details on how to configure M-Link integration on Icon-Topo servers can be found in its manual.

---

## A.2 Configuration managed by Icon-Topo

Icon-Topo will create and maintain the routing information of an M-Link MU Gateway server. More specifically, it will maintain the configuration of Links, STANAG 5066 Servers and Peer Controls. As long as the *Enable Icon-Topo Support* option is set, routing information configured through Icon-Topo will always override locally made changes.

### A.2.1 Links

When Icon-Topo defines a link, M-Link will first check whether there's already an existing link that was previously configured by Icon-Topo with the same name.

If it cannot find an existing Icon-Topo-managed link, it will try to find an existing link that matches. For [XEP-0361](#) links, an existing link with the same *Local Port* will match. For STANAG 5066 links, the STANAG 5066 Server *Host* and *Port* options, as well as the *Remote SIS Address* will be used. If it finds a matching link, it will automatically become associated with the Icon-Topo-provided name. Taking over existing links can be useful to pre-configure certain aspects (such as TLS details) of links before they are actually used for routing.

Once a suitable link is found, or a new one has been created if there was no existing matching link, M-Link will ensure that the details provided by Icon-Topo will remain configured.

The following aspects of Links are managed, maintained according to the latest configuration provided by Icon-Topo, and should not be changed manually:

- Link Name
- Link Type
- Local Port ([XEP-0361](#) only)
- Remote Host ([XEP-0361](#) only)
- Remote Port ([XEP-0361](#) only)
- STANAG 5066 Server (STANAG 5066 only)
- Remote SIS Address (STANAG 5066 only)

## A.2.2 STANAG 5066 Server Definitions

For each STANAG 5066 link defined by Icon-Topo, M-Link will ensure there is a corresponding STANAG 5066 Server Definition. The *Local SIS Address* will be updated if it changes in a new Icon-Topo configuration.

## A.2.3 Peer Controls

Icon-Topo has the ability to maintain both the *default* Peer Control, as well as any additionally configured Peer Controls (also known as *routes*.)

Peer Controls that were previously configured by Icon-Topo and are no longer present in the newest configuration will have the domain extended with `.disabled.invalid`. When a Peer Control is added for which there is an existing control with a domain ending in `.disabled.invalid`, the existing route will be updated and made active again by having the suffix stripped. These two mechanisms will allow for persistence of local changes made to routes even when a Peer is temporarily unavailable.

The following aspects of Peer Controls are managed, maintained according to the latest configuration provided by Icon-Topo, and should not be changed manually:

- Domain
- Matching Rule
- Use Link
- Link
- Inbound-Only Links
- Relay Zone

---

**Note:** Peer Controls are matched on both their *Domain* and *Matching Rules*. Icon-Topo supports the *Domain* or *Domain and Subdomains* matching rules. Use of manually configured Peer Controls using the *Subdomains* matching rule on a system managed by Icon-Topo is not advised.

---

# Appendix B Pre-defined Transformations

This appendix describes the [Transformations](#) that are included with M-Link.

---

## B.1 Block common File Transfer and VOIP mechanisms

This transformation will ensure that common mechanisms for file transfer and VOIP will be removed.

The protocols covered are In-Band Bytestreams ([XEP-0047](#)), Stream Initiation ([XEP-0095](#)), and Jingle ([XEP-0166](#)) and their associated standards.

When using this transformation, it is recommended to set the *Action on Blocked Stanza* option on the rule to *Bounce* to ensure the blocked IQ requests are replied to with an error.

---

## B.2 Normalize XML for M-Guard

This transformation should be used to normalize the XML sent to M-Guard servers. It should generally be the last one configured on a given peer.

This will ensure `xml:lang` attributes are all lowercase, and that whitespace is normalized.

---

## B.3 Strip Chat State Notifications

This transformation will remove Chat State Notifications ([XEP-0085](#)) elements from message stanzas. Chat State Notifications are used to inform users of the state of their conversation partner in a chat session ("inactive", "composing", "paused", etc.)

Disabling these notifications may be desirable in cases where bandwidth is restricted.

---

## B.4 Strip Legacy Delayed Delivery

This transformation removes Legacy Delayed Delivery ([XEP-0091](#)) elements from message or presence stanzas. The Legacy Delayed Delivery specification has been deprecated in favour of Delayed Delivery ([XEP-0203](#)).

Disabling Legacy Delayed Delivery and allowing Delayed Delivery ensures that only one representation of timestamps is emitted, e.g. when crossing a security guard appliance.

---

## B.5 Strip Message Delivery Receipts

This transformation removes Message Delivery Receipts ([XEP-0184](#)) elements from message stanzas. Message Delivery Receipts are a mechanism that will allow a sender to ask the recipient to acknowledge receipt of a content message by returning an ack message.

Disabling Message Delivery Receipts may be desirable in cases where bandwidth is restricted.

---

## B.6 Strip XHTML-IM parts from messages

This transformation will remove XHTML-IM ([XEP-0071](#)) markup from message stanzas. XHTML-IM is a method to include lightweight text markup.

Disabling XHTML-IM parts may be desirable in cases where bandwidth is restricted.

# Appendix C Standards Supported by M-Link

An overview of Open Standards M-Link products conform to.

---

## C.1 Supported by M-Link Edge

[RFC 0793](#) Transmission Control Protocol

[RFC 1034](#) Domain names - concepts and facilities

[RFC 1035](#) Domain names - implementation and specification

[RFC 6120](#) Extensible Messaging and Presence Protocol (XMPP): Core

[RFC 6121](#) Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence

[RFC 6122](#) Extensible Messaging and Presence Protocol (XMPP): Address Format

[XEP-0004](#) Data Forms

[XEP-0030](#) Service Discovery

[XEP-0082](#) XMPP Date and Time Profiles

[XEP-0198](#) Stream Management

M-Link only supports Stanza Acknowledgements. Stream Resumption is not supported.

[XEP-0199](#) XMPP Ping

[XEP-0220](#) Server Dialback

[XEP-0258](#) Security Labels in XMPP

[XEP-0361](#) Zero Handshake Server to Server Protocol

---

## C.2 Supported by M-Link MU Gateway

[RFC 0793](#) Transmission Control Protocol

[RFC 1034](#) Domain names - concepts and facilities

[RFC 1035](#) Domain names - implementation and specification

[RFC 6120](#) Extensible Messaging and Presence Protocol (XMPP): Core

[RFC 6121](#) Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence

[RFC 6122](#) Extensible Messaging and Presence Protocol (XMPP): Address Format

[XEP-0004](#) Data Forms

[XEP-0030](#) Service Discovery

[XEP-0082](#) XMPP Date and Time Profiles

[XEP-0198](#) Stream Management

M-Link only supports Stanza Acknowledgements. Stream Resumption is not supported.

[XEP-0199](#) XMPP Ping

[XEP-0220](#) Server Dialback  
[XEP-0258](#) Security Labels in XMPP  
[XEP-0361](#) Zero Handshake Server to Server Protocol  
[XEP-0365](#) Server to Server communication over STANAG 5066 ARQ

---

## C.3 Supported by M-Link MU Server

[RFC 0793](#) Transmission Control Protocol  
[RFC 1034](#) Domain names - concepts and facilities  
[RFC 1035](#) Domain names - implementation and specification  
[RFC 6120](#) Extensible Messaging and Presence Protocol (XMPP): Core  
[RFC 6121](#) Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence  
[RFC 6122](#) Extensible Messaging and Presence Protocol (XMPP): Address Format  
[XEP-0004](#) Data Forms  
[XEP-0030](#) Service Discovery  
[XEP-0045](#) Multi-User Chat  
[XEP-0049](#) Private XML Storage  
[XEP-0054](#) vcard-temp  
[XEP-0060](#) Publish-Subscribe  
[XEP-0082](#) XMPP Date and Time Profiles  
[XEP-0092](#) Software Version  
[XEP-0124](#) Bidirectional-streams Over Synchronous HTTP (BOSH)  
[XEP-0191](#) Blocking Command  
[XEP-0198](#) Stream Management  
M-Link only supports Stanza Acknowledgements. Stream Resumption is not supported.  
[XEP-0199](#) XMPP Ping  
[XEP-0203](#) Delayed Delivery  
[XEP-0206](#) XMPP Over BOSH  
[XEP-0220](#) Server Dialback  
[XEP-0227](#) Portable Import/Export Format for XMPP-IM Servers  
M-Link only supports importing data from XEP-0227 sources.  
[XEP-0258](#) Security Labels in XMPP  
[XEP-0280](#) Message Carbons  
[XEP-0289](#) Federated MUC for Constrained Environments  
[XEP-0313](#) Message Archive Management  
M-Link currently supports versions 0.3 (MAMv0) and 0.5.1 (MAMv1) of the specification.  
[XEP-0361](#) Zero Handshake Server to Server Protocol  
[XEP-0363](#) HTTP File Upload  
[XEP-0365](#) Server to Server communication over STANAG 5066 ARQ

---

## C.4 Supported by M-Link User Server

[RFC 0793](#) Transmission Control Protocol

[RFC 1034](#) Domain names - concepts and facilities

[RFC 1035](#) Domain names - implementation and specification

[RFC 6120](#) Extensible Messaging and Presence Protocol (XMPP): Core

[RFC 6121](#) Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence

[RFC 6122](#) Extensible Messaging and Presence Protocol (XMPP): Address Format

[XEP-0004](#) Data Forms

[XEP-0030](#) Service Discovery

[XEP-0045](#) Multi-User Chat

[XEP-0049](#) Private XML Storage

[XEP-0054](#) vcard-temp

[XEP-0060](#) Publish-Subscribe

[XEP-0082](#) XMPP Date and Time Profiles

[XEP-0092](#) Software Version

[XEP-0124](#) Bidirectional-streams Over Synchronous HTTP (BOSH)

[XEP-0191](#) Blocking Command

[XEP-0198](#) Stream Management

M-Link only supports Stanza Acknowledgements. Stream Resumption is not supported.

[XEP-0199](#) XMPP Ping

[XEP-0203](#) Delayed Delivery

[XEP-0206](#) XMPP Over BOSH

[XEP-0220](#) Server Dialback

[XEP-0227](#) Portable Import/Export Format for XMPP-IM Servers

M-Link only supports importing data from XEP-0227 sources.

[XEP-0258](#) Security Labels in XMPP

[XEP-0280](#) Message Carbons

[XEP-0289](#) Federated MUC for Constrained Environments

[XEP-0313](#) Message Archive Management

M-Link currently supports versions 0.3 (MAMv0) and 0.5.1 (MAMv1) of the specification.

[XEP-0363](#) HTTP File Upload

# Appendix D Glossary

This appendix provides a glossary of terms.

## Technical Terms

### Domain

See [Domain Name](#).

### Domain Name

A name within the [Domain Name System](#). See [RFC1035].

### Domain Name System (DNS)

A service for providing a mapping between [domain names](#) (for example, `example.com`) and [IP addresses](#). See [RFC 1035](#).

### Extensible Messaging and Presence Protocol (XMPP)

A collection of open standards for real-time communication, including those for [instant messaging](#), presence, and [multi-user chat](#). See [RFC6120].

### Extensible Markup Language (XML)

A markup language used to represent structured information which is designed to be readable by both humans and machines.

### Federated Multi-User Chat (FMUC)

An [instant messaging](#) service allowing multiple multi-user chat rooms to be connected together. See [XEP-0289].

### Fully Qualified Domain Name (FQDN)

A complete [domain name](#) identifying a host system or other entity on the Internet. See Also [Domain Name System \(DNS\)](#).

### Guard Content Exchange Protocol (GCXP)

Protocol used for communications between M-Link and M-Guard.

### Hypertext Transfer Protocol (HTTP)

An application layer network protocol for distributed, collaborative, hypermedia information systems.

### Icon-5066

A modem-independent STANAG 5066 server. It enables applications to work efficiently over HF Modems/Radios and allows multiple applications to work simultaneously.

### Instant Messaging (IM)

Real-time text-based chatting between two or more people. XMPP IM is described in [RFC6121].

### IP address

An address which identifies a host machine on an Internet network. For IPv4, it is 32-bit number commonly written in dotted number notation of the form `192.0.1.100`. For IPv6, it is a 128-bit number commonly written in a notation of the form `2001:db8::100`.

### Multi-User Chat (MUC)

An [instant messaging](#) service allowing multiple users to chat with each other; a group chat service. See [XEP-0045].

### STANAG 5066 Subnet Interface Service (SIS)

Protocol that enables an application to connect to an HF modem through a [STANAG 5066](#) server over TCP/IP. Also known as “Subnet Interface Sublayer”.

**SIS Node Address**

An HF subnetwork node address. Usually represented in a form that resembles an IPv4 address in the range 0 . 0 . 0 . 0 — 31 . 255 . 255 . 255.

**SIS Layer Extension Protocol (SLEP)**

Layer protocol that can be used over *SIS* to provide a set of core services that can be used by different applications. This protocol is designed as a replacement for RCOP and UDOP.

See Also [ISODE-S5066-APP3](#).

**STANAG 5066**

STANAG 5066 "Profile for High Frequency (HF) Radio Data Communication" is a NATO specification to enable applications to communicate efficiently over HF Radio.

**Transmission Control Protocol (TCP)**

A stream-oriented protocol for providing reliable data communications over the Internet. TCP is the primary transport protocol for *XMPP*. See [RFC 0793](#).

**Transport Layer Security (TLS)**

A cryptographic protocol designed to provide secure communications, widely used in applications such as email and instant messaging, and most visibly used to secure HTTP transfers.

# Appendix E References

The documents listed in this appendix provide references to the appropriate standards and other sources of information.

If documents can be obtained electronically, the location is stated as part of the reference. For other documents, please see [Section E.6, “Obtaining documents”](#).

---

## E.1 RFCs

RFC 0793

*Transmission Control Protocol* [<https://tools.ietf.org/html/rfc0793>]. J. Postel, September 1981

RFC 1034

*Domain names - concepts and facilities* [<https://tools.ietf.org/html/rfc1034>]. P. Mockapetris, November 1987

RFC 1035

*Domain names - implementation and specification* [<https://tools.ietf.org/html/rfc1035>]. P. Mockapetris, November 1987

RFC 6120

*Extensible Messaging and Presence Protocol (XMPP): Core* [<https://tools.ietf.org/html/rfc6120>]. P. Saint-Andre, March 2011

RFC 6121

*Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence* [<https://tools.ietf.org/html/rfc6121>]. P. Saint-Andre, March 2011

RFC 6122

*Extensible Messaging and Presence Protocol (XMPP): Address Format* [<https://tools.ietf.org/html/rfc6122>]. P. Saint-Andre, March 2011

RFC 6901

*JavaScript Object Notation (JSON) Pointer* [<https://tools.ietf.org/html/rfc6901>]. P. Bryan, K. Zyp, M. Nottingham, April 2013

RFC 6902

*JavaScript Object Notation (JSON) Patch* [<https://tools.ietf.org/html/rfc6902>]. P. Bryan, M. Nottingham, April 2013

RFC 7235

*Hypertext Transfer Protocol (HTTP/1.1): Authentication* [<https://tools.ietf.org/html/rfc7235>]. R. Fielding, J. Reschke, June 2014

RFC 7578

*Returning Values from Forms: multipart/form-data* [<https://tools.ietf.org/html/rfc7578>]. L. Masinter, July 2015

RFC 8259

*The JavaScript Object Notation (JSON) Data Interchange Format* [<https://tools.ietf.org/html/rfc8259>]. T. Bray, December 2017

---

## E.2 XMPP Extension Protocols

XEP-0004

*Data Forms* [<https://xmpp.org/extensions/xep-0004.html>]. Ryan Eatmon, Joe Hildebrand, Jeremie Miller, Thomas Muldowney, Peter Saint-Andre, July 2022

- XEP-0030  
*Service Discovery* [<https://xmpp.org/extensions/xep-0030.html>]. Joe Hildebrand, Peter Millard, Ryan Eatmon, Peter Saint-Andre, October 2017
- XEP-0045  
*Multi-User Chat* [<https://xmpp.org/extensions/xep-0045.html>]. Peter Saint-Andre, December 2022
- XEP-0047  
*In-Band Bytestreams* [<https://xmpp.org/extensions/xep-0047.html>]. Justin Karneges, Peter Saint-Andre, January 2021
- XEP-0049  
*Private XML Storage* [<https://xmpp.org/extensions/xep-0049.html>]. Peter Saint-Andre, Russell Davis, March 2004
- XEP-0054  
*vcard-temp* [<https://xmpp.org/extensions/xep-0054.html>]. Peter Saint-Andre, July 2008
- XEP-0060  
*Publish-Subscribe* [<https://xmpp.org/extensions/xep-0060.html>]. Peter Millard, Peter Saint-Andre, Ralph Meijer, March 2023
- XEP-0071  
*XHTML-IM* [<https://xmpp.org/extensions/xep-0071.html>]. Peter Saint-Andre, March 2018
- XEP-0082  
*XMPP Date and Time Profiles* [<https://xmpp.org/extensions/xep-0082.html>]. Peter Saint-Andre, Tobias Markmann, August 2021
- XEP-0085  
*Chat State Notifications* [<https://xmpp.org/extensions/xep-0085.html>]. Peter Saint-Andre, Dave Smith, September 2009
- XEP-0091  
*Legacy Delayed Delivery* [<https://xmpp.org/extensions/xep-0091.html>]. Peter Saint-Andre, May 2009
- XEP-0092  
*Software Version* [<https://xmpp.org/extensions/xep-0092.html>]. Peter Saint-Andre, February 2007
- XEP-0095  
*Stream Initiation* [<https://xmpp.org/extensions/xep-0095.html>]. Thomas Muldowney, Matthew Miller, Ryan Eatmon, November 2017
- XEP-0124  
*Bidirectional-streams Over Synchronous HTTP (BOSH)* [<https://xmpp.org/extensions/xep-0124.html>]. Ian Paterson, Dave Smith, Peter Saint-Andre, Jack Moffitt, Lance Stout, Winfried Tilanus, May 2021
- XEP-0166  
*Jingle* [<https://xmpp.org/extensions/xep-0166.html>]. Scott Ludwig, Joe Beda, Peter Saint-Andre, Robert McQueen, Sean Egan, Joe Hildebrand, September 2018
- XEP-0184  
*Message Delivery Receipts* [<https://xmpp.org/extensions/xep-0184.html>]. Peter Saint-Andre, Joe Hildebrand, August 2018
- XEP-0191  
*Blocking Command* [<https://xmpp.org/extensions/xep-0191.html>]. Peter Saint-Andre, March 2015
- XEP-0198  
*Stream Management* [<https://xmpp.org/extensions/xep-0198.html>]. Justin Karneges, Peter Saint-Andre, Joe Hildebrand, Fabio Forno, Dave Cridland, Matthew Wild, July 2018

- XEP-0199  
[XMPP Ping](https://xmpp.org/extensions/xep-0199.html) [https://xmpp.org/extensions/xep-0199.html]. Peter Saint-Andre, March 2019
- XEP-0203  
[Delayed Delivery](https://xmpp.org/extensions/xep-0203.html) [https://xmpp.org/extensions/xep-0203.html]. Peter Saint-Andre, September 2009
- XEP-0206  
[XMPP Over BOSH](https://xmpp.org/extensions/xep-0206.html) [https://xmpp.org/extensions/xep-0206.html]. Ian Paterson, Peter Saint-Andre, Lance Stout, Winfried Tilanus, April 2014
- XEP-0220  
[Server Dialback](https://xmpp.org/extensions/xep-0220.html) [https://xmpp.org/extensions/xep-0220.html]. Jeremie Miller, Peter Saint-Andre, Philipp Hancke, March 2015
- XEP-0227  
[Portable Import/Export Format for XMPP-IM Servers](https://xmpp.org/extensions/xep-0227.html) [https://xmpp.org/extensions/xep-0227.html]. Magnus Henoch, Waqas Hussain, Matthew Wild, June 2021
- XEP-0258  
[Security Labels in XMPP](https://xmpp.org/extensions/xep-0258.html) [https://xmpp.org/extensions/xep-0258.html]. Kurt Zeilenga, November 2018
- XEP-0280  
[Message Carbons](https://xmpp.org/extensions/xep-0280.html) [https://xmpp.org/extensions/xep-0280.html]. Joe Hildebrand, Matthew Miller, Georg Lukas, December 2021
- XEP-0289  
[Federated MUC for Constrained Environments](https://xmpp.org/extensions/xep-0289.html) [https://xmpp.org/extensions/xep-0289.html]. Kevin Smith, March 2021
- XEP-0313  
[Message Archive Management](https://xmpp.org/extensions/xep-0313.html) [https://xmpp.org/extensions/xep-0313.html]. Matthew Wild, Kevin Smith, March 2023
- XEP-0361  
[Zero Handshake Server to Server Protocol](https://xmpp.org/extensions/xep-0361.html) [https://xmpp.org/extensions/xep-0361.html]. Steve Kille, September 2017
- XEP-0363  
[HTTP File Upload](https://xmpp.org/extensions/xep-0363.html) [https://xmpp.org/extensions/xep-0363.html]. Daniel Gultsch, January 2022
- XEP-0365  
[Server to Server communication over STANAG 5066 ARQ](https://xmpp.org/extensions/xep-0365.html) [https://xmpp.org/extensions/xep-0365.html]. Steve Kille, March 2022

---

## E.3 W3C Recommendations

- XSLT 1.0  
[XSL Transformations \(XSLT\) Version 1.0](https://www.w3.org/TR/xslt-10/) [https://www.w3.org/TR/xslt-10/]. W3C Recommendation, 16 November 1999

---

## E.4 Other Publications

- [JSON Schema](https://json-schema.org/) [https://json-schema.org/].  
[Open XML SPIF](http://www.xmlspif.org/) [http://www.xmlspif.org/].

---

## E.5 Isode White Papers

ISODE-MSG-FLOW

*Security Policy Based Access Controls* [<https://www.isode.com/products/security-policy-infrastructure.html>].

ISODE-S5066

*STANAG 5066 (The Standard for Data Applications over HF Radio)* [<https://www.isode.com/whitepapers/stanag-5066.html>].

ISODE-S5066-APP3

*SIS Layer Extension Protocol (SLEP) (S5066-APP3)* [<https://www.isode.com/whitepapers/S5066-APP3.html>].

ISODE-XMPP-LABELS

*Using Security Labels to Control Message Flow in XMPP Services* [<https://www.isode.com/whitepapers/controlling-message-flow.html>].

---

## E.6 Obtaining documents

### E.6.1 ISO/IEC documents

ISO/IEC standards and draft documents may be obtained from:

ISO Central Secretariat  
International Organization for Standardization (ISO)  
1, rue de Varembe  
Case postale 56  
CH-1211 Geneva 20  
Switzerland

Telephone: +41 22 749 01 11

Fax: +41 22 733 34 30

Web: <http://www.iso.org/>

### E.6.2 ITU-T (CCITT) documents

International Telecommunications Union  
Place des Nations  
CH-1211 Geneva 20  
Switzerland

Telephone: +41 22 730 61 41 Fax: +41 22 730 51 94

Email: [sales@itu.int](mailto:sales@itu.int)

Web: <http://www.itu.int/>

### E.6.3 RFCs

Electronic copies of RFCs are available from the following servers:

- <http://ftp.isi.edu/in-notes/>
- <http://www.rfc-editor.org/>