



LogMeIn Security

An In-Depth Look

LogMeIn®

Table of Contents

Remote Access Axioms	3
LogMeIn Architecture	5
LogMeIn Security Mechanisms.....	7
Intrusion Detection.....	13
Data Forwarding.....	18
UDP NAT Traversal	19
Software Updates and Gateway Security	20
Conclusion	20

Abstract	This paper provides an in-depth look at the security features of LogMeln’s remote access and management products, LogMeln Pro and LogMeln Central. At LogMeln we believe in security through transparency. We do not expect our customers to blindly accept our claims. By publishing details on how security mechanisms work and inter-operate in our products, we are also inviting the public to scrutinize our efforts.
Audience	This document is technical in nature and is aimed at network engineers or network designers. Reading this paper can help the reader perform the necessary threat analysis before deploying our product.
Terminology	In the LogMeln architecture, there are three entities that take part in every remote access session. The “client” or the “user” is the person or software (browser, native app, mobile app) accessing a remote resource. The “host” or the “server” is the computer being accessed, or the LogMeln host software on this computer. The “gateway” is the LogMeln service that mediates traffic between the client and the host.
Design Fundamentals	LogMeln is designed to allow secure remote access to critical resources over an untrusted network. During development, security considerations always prevail over usability concerns.

Remote Access Axioms

Everything Is a Target

More and more computers are online 24/7. Most of these computers are operated by home users and have gaping security holes, such as unpatched vulnerabilities and a lack of proper passwords.

The greatest weakness is, however, the user himself. The extremely quick penetration of so-called email viruses illustrates the lack of security-consciousness and the gullible nature of most Internet users. Email viruses, of course, are email attachments that are better classified as Trojan horses. They spread so quickly because users are surprisingly willing to violate fundamental rules when handling untrusted content. If the users themselves are responsible for infecting their computers with Trojans, how can you trust them to properly secure their systems against direct attacks?

Even competent network administrators can slip up and forget to install a patch or two, which, in the worst-case scenario, can allow attackers to run arbitrary code on the affected systems.

Cyber-attacks are nothing new, but recently there have been significant changes in who is capable of conducting them. What was once an illicit profession restricted to a highly skilled, knowledgeable and well-connected few has morphed into an endeavor that nearly anyone can undertake, owing in large part to automated exploit kits that can abuse thousands of known vulnerabilities out of the box.

Remote Access and Security

It is easy to see that many computers connected to the Internet are extremely vulnerable, even without installing a remote access product. Remote access products are perceived as high risk factors, but mainly for psychological reasons. When a user first sees a remote access solution in action, their first negative reaction is usually with regard to the security implications. This is perfectly normal, and, in fact, desirable. The real problem is that users

do not immediately see the threat inherent in other network-enabled applications, such as an email client, a web server or the operating system itself.

All modern operating systems include some sort of remote access solution by default. Windows, for example, ships with Microsoft's Remote Desktop as a simple remote administration interface. Even OpenBSD, the Unix variant which is usually regarded as the most secure operating system available, includes SSH, which, again, is a simple and secure application that allows command-line access over a network connection to the remote computer.

In essence, a well-chosen and well-configured remote access solution does not carry any additional risks. If a network manager can keep a network secure using a reliable remote access software package, such as LogMeIn, productivity can be increased and costs reduced without any adverse effects on network security.

LogMeln Architecture

Before explaining the exact security mechanisms employed by LogMeln, it is necessary to give a quick introduction to the solution architecture.

There are three key components to any remote access session. The roles of the client and the host should be straightforward – the third component is the LogMeln gateway.

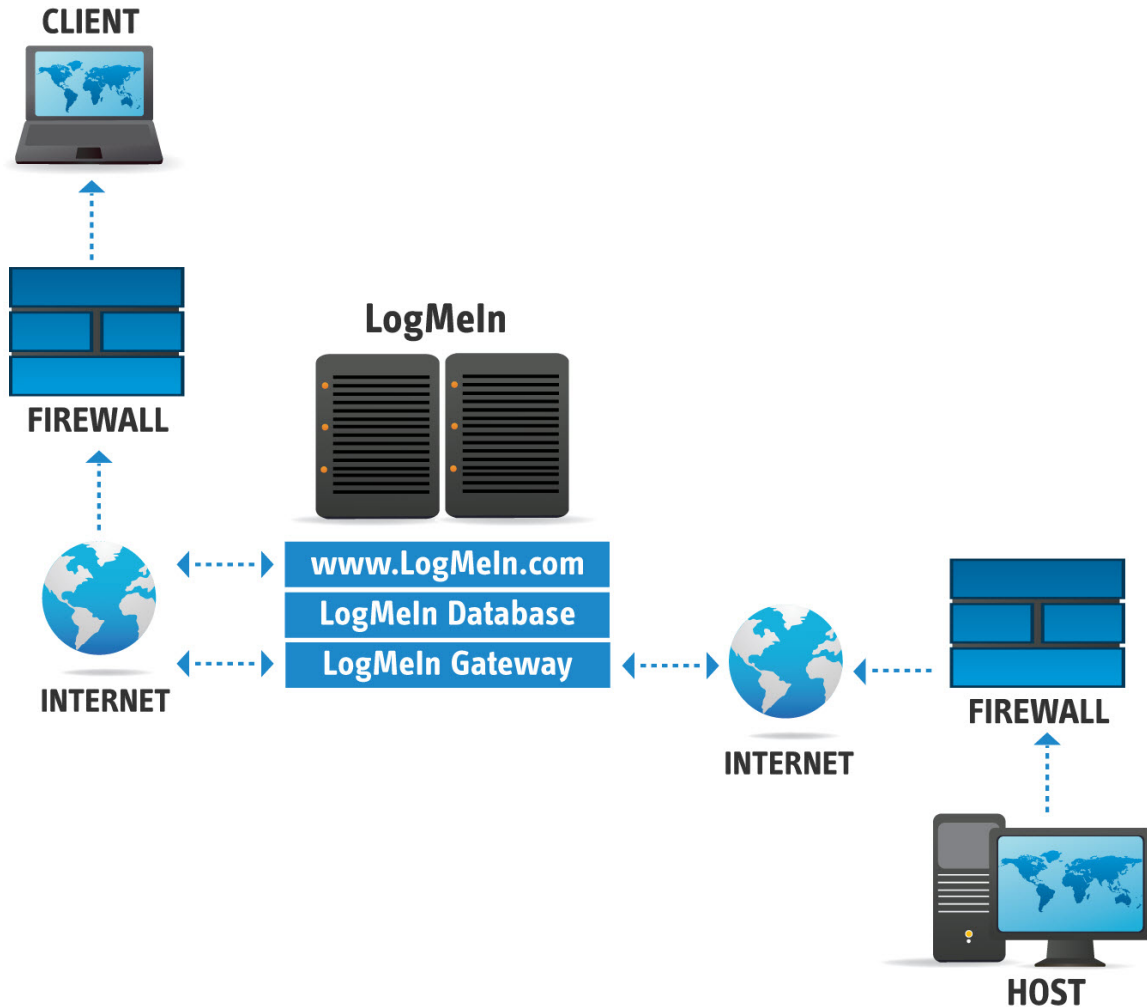


Figure 1. LogMeln Architecture

The LogMeln host in Figure 1 maintains a constant TLS-secured connection with one of the LogMeln gateway servers in one of our physically secure datacenters. This link is initiated by the host and firewalls treat it as an outgoing connection, like secure web-browsing traffic.

The client establishes a connection to LogMeIn and authenticates itself. Based on the client's identity, it is authorized to exchange data with one or more hosts (the hosts belonging to the user's account).The gateway then forwards the subsequent encrypted traffic between the client and the host. It is worth noting that the client will also need to authenticate itself to the host – the gateway mediates the traffic between the two entities, but it does not require that the host implicitly trust the client. Once the host has verified the client's identity and authorized the client to access the computer the actual remote access session begins.

The benefit of using the gateway is that either the client or host (or both) can be firewalled. The LogMeIn gateway ensures that users do not need to configure firewalls.

LogMeln Security Mechanisms

When users think of Internet data security, they are usually concerned about data encryption – to the point where security is measured in the length of the encryption key used. However, encryption and decryption, while being very important, are fairly trivial tasks compared to the other challenges faced by designers of secure systems. As you will see, data encryption is just one of the main goals set forth by the designers of LogMeln.

Authentication of the Gateway to the Client

First and foremost, when a user connects to a LogMeln installation via a gateway – the “server” – they need to be 100% positive that the computer they are about to exchange data with is really the one to which they intended to connect.

Suppose that an attacker poses as the server towards the user, and it poses as the user towards the server. The attacker, in this case, can sit between the two parties while reading, or possibly modifying, the data in transit. This is known as a “Man in the Middle”, or MITM attack and is especially hard to protect against.

LogMeln utilizes TLS 1.x certificate based authentication to verify server identities and thus protect against MITM attacks. When a connection is made, the server’s certificate is verified. A warning is presented if an untrusted certifying authority issued the certificate. A different warning is presented if the hostname in the URL does not match the hostname in the certificate, even if issued by a trusted authority.

If the server passes these verifications, then the user’s client generates a “Pre-Master Secret” or PMS, encrypts it with the server’s public key contained within its certificate, and sends it to the server. As ensured by the use of public key cryptography, only the server that holds the corresponding private key can decrypt the PMS. The PMS is then used to derive the Master Secret by both the user and the server, which, in turn, will be used to derive initialization vectors and session keys for the duration of the secure session.

In short, the above ensures that the user is establishing the connection with the server, and not with a third entity. Should a MITM attack be attempted, either one of the security

warnings will be triggered or the PMS will be unknown to the MITM, effectively rendering the attack impossible.

Authentication of Users to the Gateway

LogMeIn users must be authenticated by both the gateway and the host. An email address and password verification is performed whenever a user logs on to the LogMeIn website. Users are also advised to enable one or more of LogMeIn’s extra security features to strengthen this authentication step.



LogMeIn Central subscribers can enforce a strong password policy. Visit help.logmein.com for details.

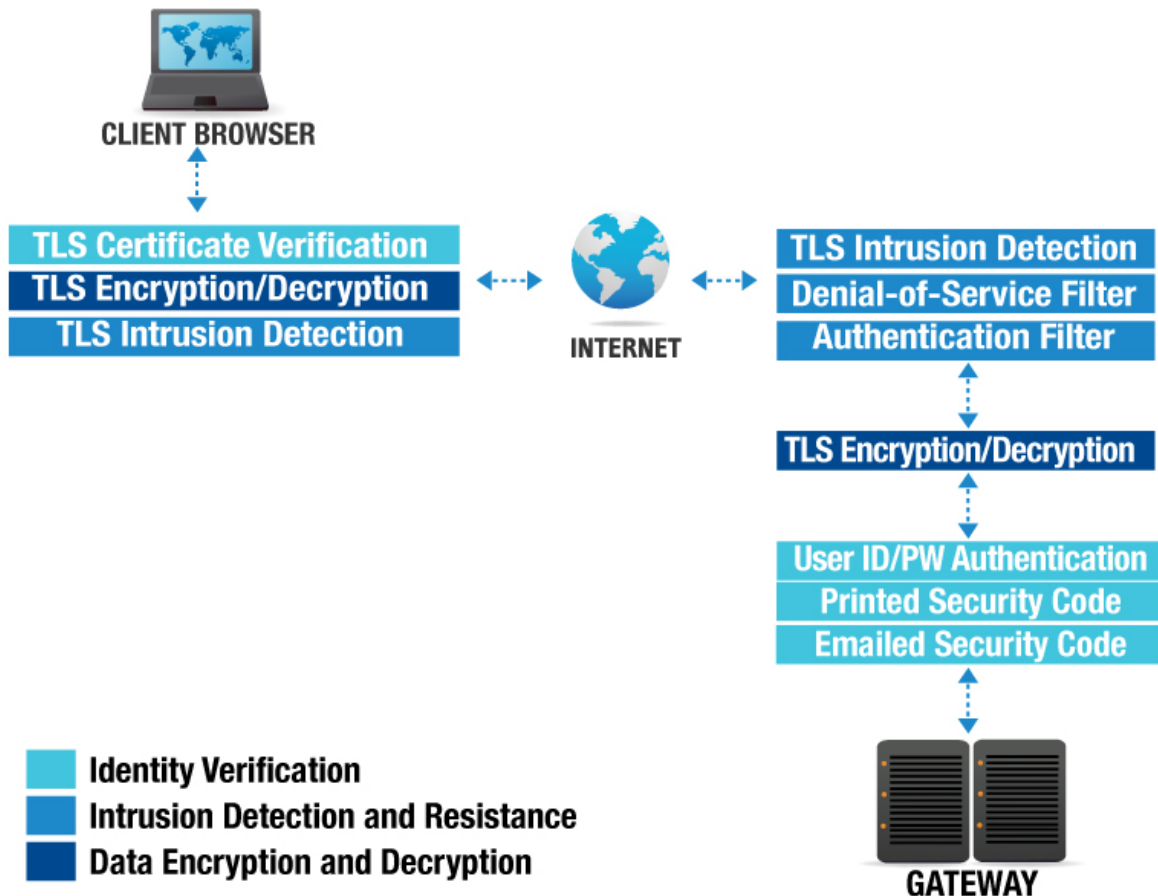


Figure 2. Authentication between Users and the Gateway

Printed Security Codes

One extra security feature is a sheet of printed security codes. When the user enables this feature, he is required to print out a list of nine-character random passwords generated by the gateway. Each time a LogMeIn user logs in to his account at www.LogMeIn.com, he will be prompted to enter one of the security codes from the list in order to gain access to his account. If you have multiple profiles, you must enter a code when you access the currently active profile. Each code can be used only once. Before the user runs out of printed security codes, he is required to print another sheet. This invalidates any previously unused security codes.



How to enable printed security codes.

- 1 Login to your LogMeIn account.
- 2 Click **Settings > Security**.
- 3 Select the **Printed Security Code** option.
- 4 Generate and print a list of security codes.
- 5 Click **Save**.

Emailed Security Codes

Another way to secure your LogMeIn account is to use the Emailed Security Code feature. Each time you log in to your account at LogMeIn.com you will be sent an email containing a security code that you must then enter in the appropriate dialog before you can access your account. Each code can be used only once.

When this feature is turned on and the user authenticates successfully with his email address and password to the LogMeIn gateway, a pass code is generated and sent to the email address. The user receives this pass code in an email and enters the code into the form provided by the gateway. The password expires either upon use or within a few minutes of generation, whichever comes first.



How to enable emailed security codes

- 1 Login to your LogMeIn account.
- 2 Click **Settings** > **Security**.
- 3 On the **Security** tab, select the **Emailed Security Code** option.
- 4 Enter your email address in the field provided.
- 5 Click **Save**.

Account Audit

Keep track of activity in your LogMeIn account. Select events for which you want to receive automatic email notification, such as login attempt failure or password changes. Notifications will be sent to the specified email addresses (for multiple recipients, separate email addresses with a semi-colon). Note that some account events are turned on by default and cannot be disabled.



How to enable the account audit feature (email notifications)

- 1 Login to your LogMeIn account.
- 2 Click **Settings** > **Account Settings**.
- 3 Under **Email notifications**, click **change** and select events for which you want to receive automatic email notification.
Tip: You can also edit the recipient list.
- 4 Click **Save**.

LogMeIn Central subscribers (account holders) can audit additional items at **Settings** > **Security** > **Audit Settings**:

- Changes to security-related settings
- User events (invited users, accepted or deleted invitations, login events)
- Computer-related events (added, deleted, shortcut generated/invalidated, installation package generated)

Two-Step Verification (Two-Factor Authentication, 2FA)

Two-step verification adds a second layer of protection to your account. Just like a cash machine that protects your money by requiring both a bank card and a PIN. Without two-step verification, anyone who knows your password can access your data. Once you set up

two-step verification, your log in procedure will change: After entering your LogMeIn ID and password, you will be required to verify your identity.

LogMeIn Central subscribers can enforce a login policy that forces all users in their account to use two-step verification. For step-by-step instructions, visit help.LogMeIn.com

Authentication of the Gateway to the Host

The gateway must prove its identity to the host before it is trusted with access codes. The host, when making a connection to the gateway, checks the certificate transported during the TLS handshake to make sure it is connecting to one of the LogMeIn gateway servers. This process is very similar to how a gateway authenticates a client.

Authentication of the Host to the Gateway

The gateway verifies the host's identity when it accepts an incoming connection using a long unique identifier string. This string is a shared secret between the two entities and is issued by the gateway when the host is installed. This unique identifier is communicated over a TLS-secured channel only after the host has verified the gateway's identity. Figure 3 illustrates how the host and the gateway authenticate each other before a host is made accessible to the client.

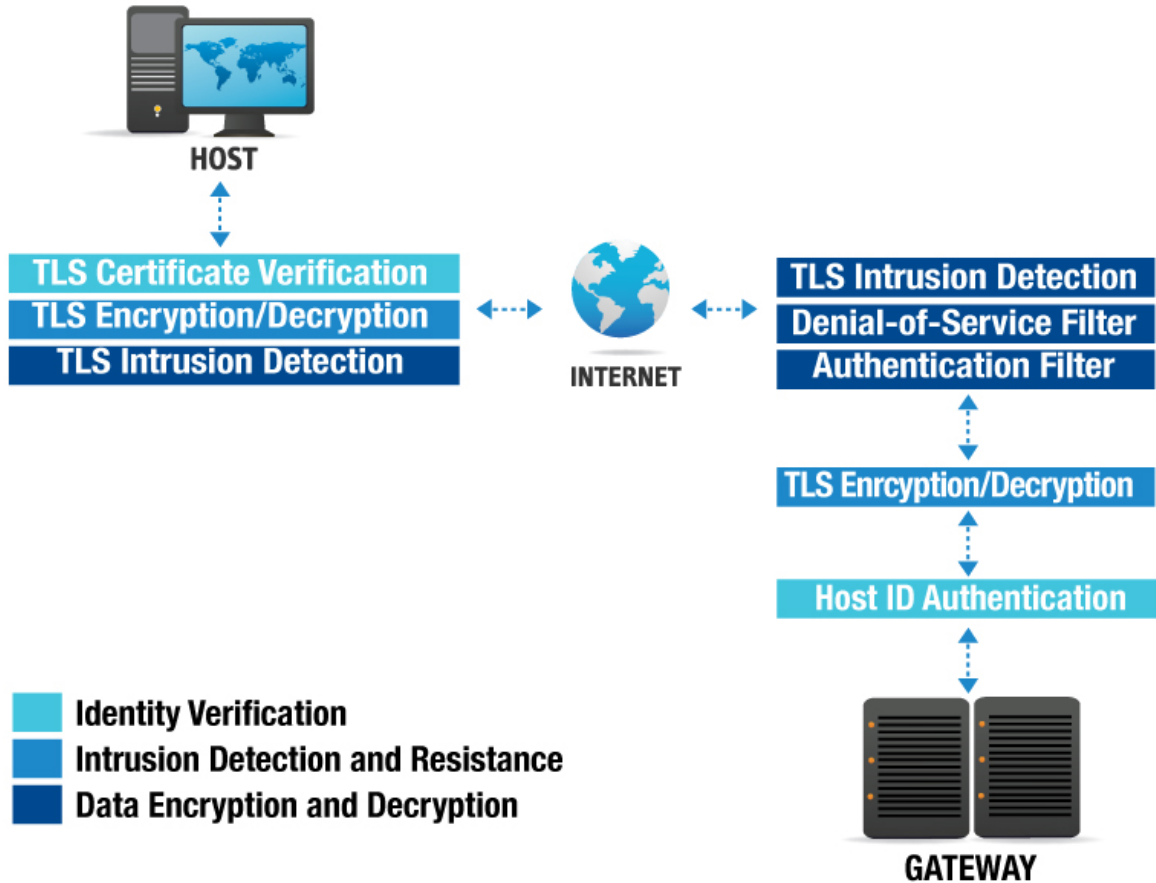


Figure 3. Host and Gateway Authentication

Data Encryption

The TLS standard defines a wide choice of cipher suites, mostly based on AES-based encryption or 3DES for compatibility reasons. AES can utilize 128 or 256 bit keys, while 3DES uses 168 bit keys. The client and the host agree on the strongest cipher possible. The client sends the host a list of ciphers it is willing to use, and the host chooses the one it prefers.

The TLS standard does not define how the host should choose the final cipher. In LogMeIn, the host simply selects the strongest shared cipher suite that the client has offered.

This method allows both the client and the host to decline the use of specific data-protection algorithms without the need of updating both components, should an algorithm be deemed broken or insecure.

Intrusion Detection

LogMeIn provides two layers to detect intrusion attempts: [TLS](#) and [LogMeIn Intrusion Filters](#).

TLS

The first layer of intrusion detection is provided by TLS to ensure that the data have not changed in transit. This is achieved by the following techniques:

Record Sequence Numbering Record Sequence Numbering means that TLS records are numbered by the sender and the order is checked by the receiver. This ensures that an attacker cannot remove or insert arbitrary records into the data stream.

Message Authentication Codes Message Authentication Codes (MACs) are appended to every TLS record. This is derived from the session key (known only to the two communicating parties) and the data contained within the record. If MAC verification fails it is assumed that the data were modified in transit.

LogMeIn Intrusion Filters

The second layer is provided by LogMeIn itself, and is comprised of three intrusion filters

IP Address Filter

When LogMeIn receives a connection request from a client, it first checks its list of trusted and untrusted IP addresses and possibly denies the connection. An administrator can set up a list of IP addresses within LogMeIn that are either allowed or denied to establish a connection to the selected host (for example, designate the internal network and another administrator's home IP address as allowed).

Denial of Service Filter

A Denial of Service Filter rejects connections if the IP address the request is coming from has made an excessive number of requests without authentication within the observation time window. This is done to protect against someone overloading the host computer by, for example, automatically and very quickly requesting the login page over and over again.

Authentication Filter

If the user made an excessive number of failed login attempts, the Authentication Filter rejects the connection. The Authentication Filter is in place to prevent a potential intruder from guessing an account name and password.



How to set filters on a LogMeIn host

- 1 Access the host preferences from either the host or the client:
 - a If you are at the host, open the LogMeIn Control Panel and follow this path:
Options > Preferences > Security
 - b If you are at the client, connect to the host **Main Menu** and follow this path: **Preferences > Security**
- 2 Under **Intrusion Control**, click **Edit Profiles** to begin creating a filter profile.
For details, see the [LogMeIn Pro User Guide](#) or [Central User Guide](#).

Authentication and Authorization of Users to the Host

After being granted access by the previous layers, the user must prove his identity to the host. This is achieved by a mandatory OS-level authentication step.

The user must authenticate himself to the host using his standard Windows or Mac username and password. The host will usually pass this request on to the relevant domain controller. This step not only validates the user's identity, but also ensures that network administrators can control who is able to log in to a specific host.

Personal Password

Personal Password is another optional security measure that can be set up on the LogMeIn host. The user can assign a Personal Password to the host, which, like the OS-level password, is not stored or verified by the gateway. A difference between the operating system password and the Personal Password is that the host never asks for the complete Personal Password so the user never enters it in its entirety in any single authentication session. The user is usually prompted for three random digits of the Personal Password by the host after OS-level authentication has succeeded. If the user enters the correct characters (for example, the first, the fourth and the seventh) he is granted access.



How to set up a Personal Password

- 1 Access the host preferences from either the host or the client:
 - a If you are at the host, open the LogMeIn Control Panel and follow this path:
Options > Preferences > Security
 - b If you are at the client, connect to the host **Main Menu** and follow this path: **Preferences > Security**
- 2 Under **Personal Password**, enter your personal password and then enter it again to confirm.
- 3 Click **Apply**.

LogMeIn and RSA SecurID

To add an extra layer of security over the simple username/password authentication, you can configure LogMeIn to require RSA SecurID authentication. RemotelyAnywhere, the product that pioneered the technology in use by LogMeIn, was officially certified by RSA Security as SecureID Ready in 2003. Since that time, LogMeIn has continued to maintain the high level of security consistent with RSA technology.



For information on the RSA SecurID product, visit the [RSA website](#).

For information on setting up this feature on a LogMeIn host, visit help.logmein.com.

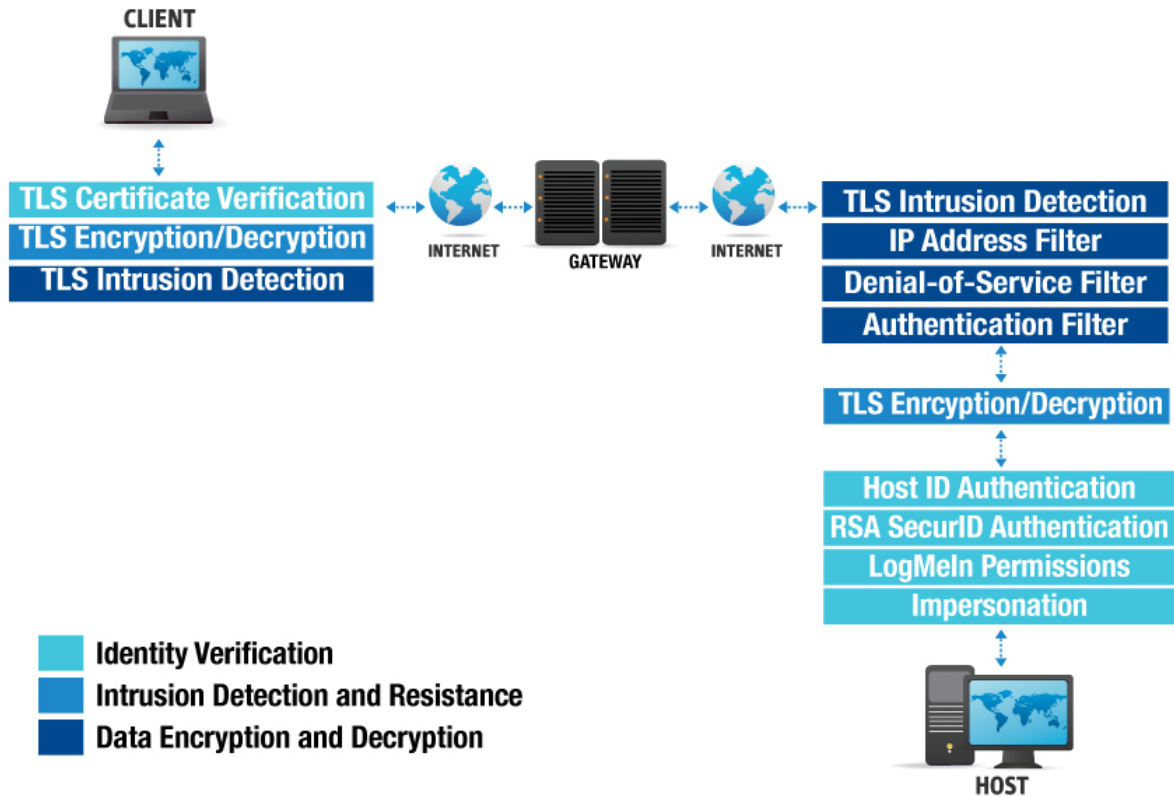


Figure 4: Authentication between Users and the Host

Authentication and Authorization of Users within the Host

Once LogMeIn has verified the user's identity using the above methods, it checks its own internal user database to see which internal modules the user is allowed to access.

System administrators can configure LogMeIn so that users with certain roles have access only to a subset of tools offered by LogMeIn; for example, the Helpdesk department can be configured to only view a computer's screen and performance data, but not actually take over the mouse and the keyboard or make any changes to the system configuration. Alternatively, the Sales department might be given full remote control access to their respective computers, but features such as performance monitoring and remote administration would be made unavailable to them.

Using the operating system access token obtained when the user was authenticated, LogMeIn impersonates the user towards the operating system while performing actions on their behalf. This ensures that LogMeIn adheres to the operating system's security model, and users have access to the same files and network resources as if they were sitting in front of their computer. Resources unavailable to users in Windows or OS X also remain unavailable via LogMeIn.



See “Controlling Who Can Access Your Host Computers” in the [LogMeIn Pro User Guide](#) or [Central User Guide](#).

Auditing and Logging

LogMeIn provides extensive logging capabilities. A very detailed log of the events that occur within the software is kept in the installation directory. The most important events are also placed in the Windows application event log – these events include, for example, logon and logoff actions. The detailed log can also be sent to a central SYSLOG server.



See “How to View LogMeIn Event Log Files” in the [LogMeIn Pro User Guide](#) for details.

For SYSLOG, see “Deployable Host Preferences for Logs and Session Recording” in the [LogMeIn Central User Guide](#).

Data Forwarding

The gateway provides end-to-end encryption by forwarding encrypted data between the host and the client. If you are familiar with how TLS works, this might sound impossible; after all, the assumption is that only the gateway can decrypt the data sent by the client since the client is confident that it is communicating with the gateway. This is a valid point, but LogMeIn made a few important changes to how TLS sessions are handled between the host and the gateway.

The first part of the TLS negotiation is performed between the gateway and the client. The gateway then passes the exchange on to the host, which re-negotiates the TLS session and agrees on a new session key with the client, thereby providing true end-to-end encryption.

When the traffic is relayed through the gateway, the client establishes a TLS session with the gateway using the gateway's certificate. The gateway transfers this TLS session's state (including the pre-master secret) to the host. After agreeing on a new session key, the host uses this session state to handle the rest of the TLS session directly with the client. The gateway's certificate secures the session, leaving the client talking directly with the host without the need for the gateway to decrypt and re-encrypt traffic.

A MITM attack is rendered impossible since both the host and the client verify the gateway's certificate and the client uses its RSA public key to authenticate the encrypted information that is used to derive the TLS Pre-Master Secret.

UDP NAT Traversal

It is important to explain how UDP NAT Traversal is used, especially since UDP is regarded as notoriously insecure. This is not entirely a misconception: If UDP is used as a communications medium, then security can be a serious problem, as UDP datagrams are easy to forge and the sender's IP address can be spoofed.

To counter this, LogMeIn.com does not use UDP as the communications medium itself with UDP NAT Traversal connections. UDP is relegated to the network layer, as defined by the ISO/OSI Network Model, with a TCP-like transport layer built on top of it, complete with flow control, dynamic bandwidth scaling and packet sequence numbering.

LogMeIn.com uses UDP instead of TCP packets (thereby effectively re-implementing a TCP-like transport layer) because most firewalls and NAT devices allow seamless two-way communication over a UDP transport as long as it is initiated from within the security perimeter, but they require significant reconfiguration for TCP and IP packets. After a reliable TCP-like stream is constructed from unreliable UDP packets, the stream is further protected by a TLS layer, providing full encryption, integrity protection and endpoint verification capabilities.

To set up a UDP NAT Traversal connection, both the client and the host send several encrypted UDP packets to the gateway. These packets are encrypted using a secret key shared by the gateway and the respective peer, and communicated over the pre-existing TLS connection. They are impossible to spoof.

The gateway uses these packets to determine the external (Internet) IP addresses of the two entities. It also tries to predict which firewall port will be used for communication when a new UDP packet is sent. It passes its findings down to the peers which then attempt to set up a direct connection. If the gateway can determine the port in use, the connection succeeds. The peers verify each other using another shared secret obtained from the gateway. A TLS session is established. The peers then communicate directly.

If a direct connection cannot be set up, the peers will connect back to the gateway over TCP and request that a forwarded, end-to-end encrypted session be used. This process

takes only a few seconds and is transparent to the user. The only noticeable difference is the improved performance and low latency when a direct connection is in use. For further details see [US Patent no. 7,558,862](#).

Software Updates and Gateway Security

The LogMeln host, based on user preferences, can semi-automatically or automatically update itself on the user's computer. The host software periodically checks the LogMeln.com website for newer versions of the software. If a new version is found, it is automatically downloaded and a message is displayed to the user who can allow the update to take place. The download process uses at most 50% of the available bandwidth, therefore keeping interference with other networking applications to a minimum.

These software updates are digitally signed by LogMeln.com with a private key that is not found on any of our Internet-connected systems. Therefore, even if the LogMeln datacenters were compromised by attackers who then gain complete control over our servers, they would not be able to upload a rogue update and run arbitrary code on our users' computers. The most such a highly unlikely attack could accomplish is access to the LogMeln logon screen on the customer's computer, which, even though it effectively bypasses the gateway security mechanisms, would still require that they enter valid operating system credentials to gain access to the computer. Brute-forcing the password is unfeasible, as the Authentication filter, by default, blocks the user's IP address after a few incorrect passwords.

For cases when a user has the same password for both LogMeln and their computer, note that we do not store actual LogMeln passwords in our database. Instead, we use a one-way cryptographic key derivation function and a per-account salt value to ensure that it is unfeasible to brute force the password even when in possession of the derived value.

Conclusion

A well-designed remote access solution can greatly increase productivity and provide a rapid return on investment. When deployed with care and LogMeln's optional security features are utilized, the benefits greatly outweigh the risks.