

ThinLinc Administrator's Guide for ThinLinc 4.12.1



Cendio®

ThinLinc®

ThinLinc Administrator's Guide for ThinLinc 4.12.1

Copyright © 2021 Cendio AB

Table of Contents

I. Introduction	1
1. Introduction	1
1.1. About the Documentation.....	1
1.2. Finding More Information.....	1
2. ThinLinc Architecture.....	3
2.1. Session Overview	4
II. Installation.....	5
3. Installation.....	5
3.1. Overview	5
3.2. Server Requirements.....	5
3.2.1. ThinLinc System and Software Requirements	5
3.2.2. Server Sizing	6
3.3. Preparing the Network for ThinLinc Installation	6
3.3.1. A Simple ThinLinc Setup	7
3.3.2. ThinLinc in a Novell Network	8
3.3.3. ThinLinc in a Windows Network.....	8
3.3.4. ThinLinc in a NAT/Split-DNS Environment	9
3.3.5. Using ThinLinc Web Access	11
3.3.6. Other Services Required by ThinLinc Servers	11
3.4. Installing the ThinLinc Remote Desktop Server	11
3.4.1. Starting the Installation Program	12
3.5. Upgrading an Old Installation	12
3.5.1. Upgrading a Cluster	13
3.5.2. New Licenses	13
3.5.3. Upgrading the Packages.....	13
3.5.4. Configuration Migration	14
3.6. SELinux enabled distributions.....	15
3.7. VirtualGL.....	15
3.7.1. Overview	15
3.7.2. Installation and configuration	15
4. License Handling	17
4.1. Overview	17
4.2. License Counting.....	17
4.3. Location and format of License Files.....	17
4.4. Log Files and E-mail Messages.....	17
4.5. Checking the Number of Valid Licenses	18
5. Printer Features	19
5.1. Overview of ThinLinc Printer Features	19
5.2. Printer Configuration Overview	19
5.2.1. CUPS Browsing	20
5.2.2. CUPS configuration on the Machine Running VSM Server	20
5.2.3. CUPS configuration on the Machine running VSM Agent	20
5.3. Local printer support	21
5.3.1. Theory of operation.....	21
5.3.2. Device independent mode.....	21
5.3.3. Device dependent mode	22

5.3.4. Installation and Configuration.....	22
5.3.5. Parallel port emulation.....	22
5.4. Nearest printer support	23
5.4.1. Administration of the Nearest Printer Feature in ThinLinc	23
5.4.2. Nearest Printer Selection Algorithm.....	23
5.4.3. Printer Drivers.....	24
5.5. Printer Access Control.....	24
5.5.1. Theory of Operation.....	25
5.5.2. Requirements	25
5.5.3. Activating the Printer Access Control Feature	25
5.5.4. Configuration	26
6. High Availability (HA).....	27
6.1. Overview	27
6.1.1. Background - Reasons For a HA Setup	27
6.1.2. Solution - Elimination of Single Point of Failure	28
6.1.3. Theory of Operation.....	28
6.2. Configuration of ThinLinc for HA Operations.....	29
6.2.1. Installation of a New HA Cluster.....	29
6.2.2. Reconfiguring an existing ThinLinc Installation into HA mode	30
6.3. Recovering from hardware failures	30
6.3.1. Recovering from Minor Failures.....	30
6.3.2. Recovering from Catastrophic Failure	31
7. The ThinLinc Client.....	33
7.1. Client usage	33
7.1.1. The started ThinLinc client.....	33
7.1.2. Logging in to a ThinLinc server	33
7.1.3. Language Settings.....	35
7.1.4. The ThinLinc session life cycle	35
7.1.5. The session menu	36
7.2. Running the ThinLinc client from the command line	36
7.3. Local device export	39
7.3.1. Sound device	39
7.3.2. Serial ports (Windows and Linux only)	40
7.3.3. Drives	40
7.3.4. Printer.....	40
7.3.5. Smart Card Readers	40
7.4. Client configuration	41
7.4.1. Options tab.....	41
7.4.2. Local Devices tab.....	43
7.4.3. Screen tab.....	46
7.4.4. Optimization tab.....	48
7.4.5. Security tab	51
7.5. Client Touch Gestures	55
7.6. The XDM mode (Linux only)	56
7.7. Logfile placement	56
7.7.1. Linux log file.....	57
7.7.2. Windows log file	57
7.7.3. macOS log file.....	57

7.8. Client configuration storage	57
7.8.1. Overview and Parameters	57
7.8.2. Configuration Parameter Storage	65
7.8.3. Adding Custom Branding to the ThinLinc Client Login Window	66
7.9. Client Customizer	66
7.9.1. Introduction	66
7.9.2. Installation	66
7.9.3. Building a Customized Client	66
7.9.4. Adding SSH Host Keys to <code>settings.reg</code>	67
7.10. Launching the Client from a Web Page	67
7.10.1. Requirements	67
7.10.2. Installation	68
7.10.3. Usage	68
7.10.4. The CGI Script <code>tlclient.cgi</code>	69
7.11. Advanced Topics	70
7.11.1. Hardware Address Reporting	70
7.11.2. Client Update Notifications	70
8. Client Platforms	73
8.1. Windows	73
8.1.1. Requirements	73
8.1.2. Installing the Windows Client	73
8.1.3. Running the Windows Client	73
8.2. macOS	73
8.2.1. Requirements	73
8.2.2. Installing the macOS Client	73
8.2.3. Running the macOS Client	74
8.2.4. Command and Alt Keys on macOS	74
8.3. Linux PC	74
8.3.1. Requirements	74
8.3.2. Installing the Linux Client	74
8.3.3. Running the Linux Client	76
8.4. Thin Terminals	76
8.4.1. eLux-based Thin Terminals (Fujitsu Futro et. al.)	76
8.4.2. HP ThinPro Terminals	78
8.4.3. IGEL Universal Desktop	78
8.4.4. Other Thin Terminals	79
8.5. Running ThinLinc on a Thinstation terminal	79
8.5.1. Installing and Building the Package	80
8.5.2. Configuring the ThinLinc client when running on a Thinstation Terminal	80
9. ThinLinc Web Access	83
9.1. Overview	83
9.2. Requirements	83
9.3. Server Configuration	83
9.3.1. Certificates	83
9.4. Usage	84
9.4.1. Logging in to a ThinLinc server	85
9.4.2. The Toolbar	86
9.4.3. Extra Keys	86

9.4.4. Clipboard.....	87
9.4.5. Touch Gestures.....	88
9.4.6. Command and Alt Keys on macOS and iOS	88
10. Authentication in ThinLinc	89
10.1. Pluggable Authentication Modules	89
10.1.1. Configuration files for PAM.....	89
10.2. Limitations.....	89
10.3. Using Public Key Authentication	89
10.3.1. Introduction.....	89
10.3.2. Key Generation	89
10.3.3. Server Configuration	90
10.3.4. Client Configuration	90
10.4. Using Smart Card Public Key Authentication.....	90
10.4.1. Introduction.....	91
10.4.2. General Requirements.....	91
10.4.3. Key Generation	91
10.4.4. Server Configuration	91
10.4.5. Client Configuration	92
10.4.6. Automatic Connection	92
10.4.7. LDAP Automatic Update (tl-ldap-certalias)	92
10.5. Using One Time Passwords.....	95
10.5.1. Introduction.....	96
10.5.2. General Requirements.....	96
10.5.3. Configuration for RSA SecurID.....	96
11. File Access	99
11.1. Accessing Windows File Servers	99
11.1.1. Introduction.....	99
11.1.2. Requirements	99
11.1.3. Mounting and Unmounting Shares	100
11.2. Restricting write access to users home directory	101
11.2.1. Introduction.....	101
11.2.2. Activation.....	102
11.2.3. Configuration	102
11.2.4. Security Considerations and Limitations	102
III. Administration	103
12. Accessing Client Resources from the ThinLinc session	103
12.1. Accessing the Clients Local Drives.....	103
12.1.1. Introduction.....	103
12.1.2. Mounting and Unmounting Local Drives	103
12.1.3. Mounting Drives at Login.....	104
12.1.4. Limitations and additional information	104
12.2. Using Serial Port redirection	104
12.2.1. Introduction.....	104
12.2.2. Requirements	104
12.2.3. Enabling Serial Port Redirection	105
12.2.4. Accessing the redirected port from applications.....	105
12.2.5. Limitations and additional information	105

12.3. Using Sound Device Redirection	105
12.3.1. Introduction.....	106
12.3.2. Requirements	106
12.3.3. PulseAudio applications.....	106
12.3.4. OSS applications.....	106
12.3.5. ALSA applications.....	107
12.3.6. Choosing sound system	107
12.3.7. Limitations and additional information	107
12.4. Using Smart Card Redirection.....	107
12.4.1. Introduction.....	108
12.4.2. Requirements	108
12.4.3. Enabling Smart Card Redirection	108
12.4.4. Limitations and additional information	108
13. Commands on the ThinLinc Server	109
14. Server Configuration	117
14.1. Configuring ThinLinc Servers in a Cluster	117
14.1.1. Cluster Configuration.....	117
14.1.2. Cluster Management	118
14.2. Server Configuration Parameters.....	119
14.2.1. Parameters in /vsmagent/	120
14.2.2. Parameters in /vsmserver/	122
14.2.3. Parameters in /vsmserver/subclusters/	123
14.2.4. Parameters in /vsm/	124
14.2.5. Parameters in /sessionstart/	125
14.2.6. Parameters in /shadowing/	126
14.2.7. Parameters in /tlwebadm/.....	126
14.2.8. Parameters in /webaccess/.....	127
14.3. Configuring Logging on ThinLinc servers	128
14.3.1. ThinLinc server components.....	128
14.3.2. Per-Session Logging	130
14.4. Customizing the User's Session	130
14.4.1. Session startup - the big picture	130
14.4.2. Session startup on VSM Agent	132
14.4.3. Profiles and the standard xstartup.default file.	133
14.4.4. Session Startup with a Client Supplied Start Program	135
14.4.5. Configuring available profiles	135
14.4.6. Configuring different Linux Desktops based on the selected profile.....	137
14.4.7. Speeding up Session Startup.....	137
14.4.8. Configuring the language environment on the server based on the client language	137
14.5. Limiting Lifetime of ThinLinc Sessions	138
15. Shadowing.....	139
15.1. Introduction	139
15.2. Disable shadowing feature.....	139
15.3. Granting shadowing access to users	139
15.4. Shadowing notification	139
15.5. Shadowing a user session	140
16. Hiveconf	141

16.1. Overview	141
16.1.1. Basic Syntax.....	141
16.1.2. Tree Structure.....	141
16.1.3. Mounting Datasources	142
16.1.4. Hostwide Configuration	142
16.1.5. Hiveconf Tools	143
16.2. Hiveconf and ThinLinc.....	143
16.2.1. The ThinLinc Configuration Tool - tl-config	143
17. Administration of ThinLinc using the Web Administration Interface	145
17.1. Introduction	145
17.2. Modules	145
17.2.1. The System Health Module	145
17.2.2. The Status Module	146
17.2.3. The VSM Module	147
17.2.4. The Profiles Module.....	149
17.2.5. The Locations Module	151
17.2.6. The Desktop Customizer Module	155
18. Building Custom Linux Desktops with the ThinLinc Desktop Customizer	157
18.1. Introduction	157
18.2. Using the ThinLinc Desktop Customizer	157
18.2.1. Concepts.....	157
18.2.2. Using the ThinLinc Desktop Customizer	159
18.2.3. Handling Applications	159
18.2.4. Defining a Menu Structure.....	160
18.2.5. Defining Application Groups.....	161
18.2.6. Distribute Configuration to all agent hosts	163
18.3. Enabling the Custom Desktops for users.....	163
18.4. Tips & Tricks with TLDC	163
18.4.1. Unwanted Icons on the Desktop with KDE.....	163
18.4.2. File Associations for Applications Not In the Menu	164
18.4.3. Home Icon not Working in KDE?	164
IV. Appendixes	165
A. TCP Ports Used by ThinLinc	165
A.1. On Machine Running VSM Server.....	165
A.2. On Machine Running VSM Agent	165
B. Troubleshooting ThinLinc.....	169
B.1. General troubleshooting method.....	169
B.2. Troubleshooting Specific Problems	170
B.2.1. Problems Where the Client Reports an Error.....	170
B.2.2. Problems that Occur After Session Start.....	172
C. Restricting access to ThinLinc servers	173
C.1. Disabling SSH access	173
C.2. Disabling shell access	173
C.2.1. Changing the configured shell.....	173
C.2.2. Using ForceCommand.....	173
C.3. Disabling port forwarding.....	174
C.3.1. Disabling remote port forwarding	174

C.4. Disabling clipboard.....	174
C.5. Disabling local drives	174
D. GnuTLS priority strings	177
D.1. Standard configuration.....	177
D.1.1. Cipher suites.....	177
D.1.2. Protocols.....	177
D.1.3. Ciphers	178
D.1.4. MACs	178
D.1.5. Key Exchange Algorithms	178
D.1.6. Groups.....	178
D.1.7. PK-signatures	179
D.2. Available algorithms.....	179
D.2.1. Cipher suites.....	179
D.2.2. Certificate types.....	183
D.2.3. Protocols.....	183
D.2.4. Ciphers	183
D.2.5. MACs	184
D.2.6. Digests.....	184
D.2.7. Key exchange algorithms	185
D.2.8. Compression.....	185
D.2.9. Groups.....	185
D.2.10. Public Key Systems.....	185
D.2.11. PK-signatures	186

Chapter 1. Introduction

1.1. About the Documentation

This document is separated into five parts. This, the first part, is an introduction to the subject with general information about the product. The second part is about how to install different components in ThinLinc and integrate those with other systems, such as user account databases and file servers. Part three discusses the administration of ThinLinc after it is installed. The last part contains appendices with extra information.

Note: Before you start using ThinLinc, please read the release notes supplied in both Server and Client Bundles and online at <http://www.cendio.com/> (<https://www.cendio.com/thinlinc/docs/relnotes>)

1.2. Finding More Information

If you need more information about ThinLinc, contact your supplier and/or visit the ThinLinc homepage, <http://www.cendio.com/>. At the ThinLinc homepage you will find information about courses, upgrades, etc.

If you need more information about Linux, we recommend looking at the Linux Documentation Project homepage (<http://www.tldp.org/>) as well as the homepage for your Linux distribution.

Chapter 2. ThinLinc Architecture

The goal of this chapter is to give a technical overview of how the system works for someone who will install or maintain a ThinLinc installation.

ThinLinc is a product for managing *server based computing*. The system is largely based on open source software, which has led to an expansion of the product to encompass solutions for authentication, availability systems, emulation and conversion between different computer systems. ThinLinc can be used as a gateway between different types of clients and a large number of base systems.

The system architecture allows an existing infrastructure to be maintained while a new architecture is gradually introduced to the organization. The system can be launched alongside the existing systems for a gradual migration to a new platform, and at the same time it acts as a link or gateway between the existing systems.

The architecture is designed to be flexible in order to handle larger organizations with autonomous office applications or functions, whilst maintaining management and security. The system can be supplemented with an automated system for installation, configuration and administration of the client hardware, such as through the use of PXE. It's also possible to create different user groups. In this way departments with special needs are easily administered in the case of adaptations or user-driven application development.

Figure 2-1. The System Architecture of ThinLinc

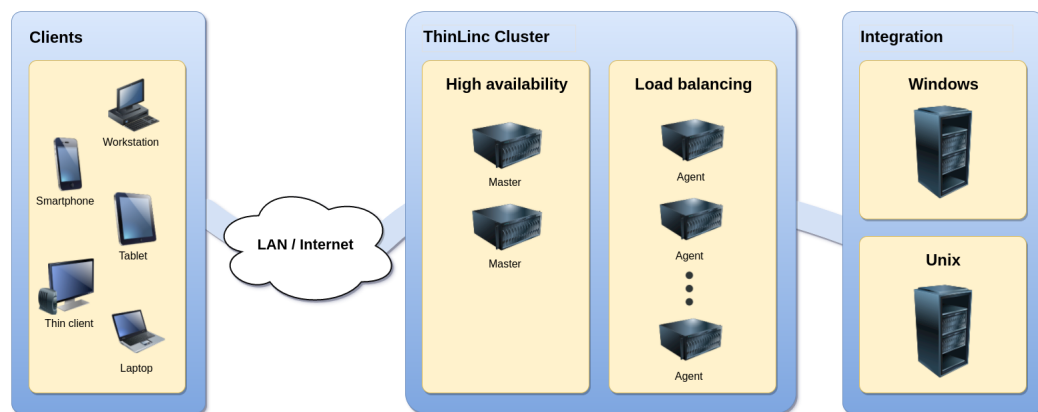


Figure 2-1 gives an overview of the ThinLinc architecture.

Several different devices can be used to connect to a ThinLinc system. ThinLinc client applications are available for Linux, macOS, Windows and selected thin terminals. ThinLinc Web Access is also available, enabling web browsers to act as ThinLinc clients.

The clients connect to a ThinLinc system located on the Local Area Network (LAN) or on a Wide Area Network (WAN) such as the Internet. Depending on the network type and the bandwidth available, several bandwidth-saving algorithms can be used to provide good performance even over narrow-banded links. Encryption is used to secure all information sent between the client and the server.

When a user connects to a ThinLinc server, a *session* is created. This session is the user's starting point for running applications either on the ThinLinc server(s) or on other servers reachable from the ThinLinc server. ThinLinc has a Single Sign-On (SSO) mechanism that enables passwordless but secure logins to (for example) Windows Remote Desktop Servers and other Unix Servers running special applications.

The ThinLinc servers runs on Linux platform. There is support for High Availability and advanced two-level load balancing.

2.1. Session Overview

When a user logs in from a native ThinLinc client, the following will happen:

- The client establishes a SSH tunnel to the server entered in the server field of the client interface. If this fails, then the login process will be interrupted and an error message will be displayed.
- The client tries to authenticate with the VSM server, through the SSH tunnel. The VSM server (VNC Session Manager) is the main process of ThinLinc, responsible for allocating and keeping track of user sessions.
- If the authentication succeeds, the server will check if there already exists a session for the user. If there is a session, then information about it will be returned. If there is no session a new one will be started on an agent server and information about it will be returned. If more than one agent server exists, load balancing will be used to select which server to start a session on.
- The client now disconnects the SSH tunnel to the VSM server and checks the information it received to see which agent server it should connect against.
- The client now establishes a new SSH tunnel to the VSM agent server it received information about from the VSM server. Port forwarding for VNC is always established, as well as other ports depending on which local devices have been enabled. All tunnels are multiplexed over the same SSH connection.
- The client now starts the VNC viewer, which will connect to the remote VNC server via the SSH tunnel.

Chapter 3. Installation

3.1. Overview

This chapter describes how to install the ThinLinc server software. To upgrade an existing installation, see Section 3.5.

1. Read through any platform-specific notes for your distribution. These can be found at <https://www.cendio.com/thinlinc/docs/platforms>.
2. Install the ThinLinc Master machine, following the instructions in Section 3.4.1.
3. Optionally, install an additional ThinLinc Master for a High Availability setup. More information regarding HA can be found in Chapter 6.
4. Optionally, install the ThinLinc Agent machines, following the same instructions as for the ThinLinc Master. Instructions for setting up a cluster can be found in Section 14.1.

3.2. Server Requirements

3.2.1. ThinLinc System and Software Requirements

- A compatible CPU architecture:
 - An i686 (or compatible) CPU with MMX and SSE support
 - An x86_64 (or compatible) CPU
- GLIBC 2.12, or newer
- RPM or dpkg support
- Libraries and commands from LSB 4.1, specifically those listed in the Core and Printing modules (except LSB specific interfaces). Additionally, "libX11" is also required.
- MIT Kerberos runtime libraries (libgssapi_krb5.so.2).
- ss from iproute2
- Python 2.6 or newer 2.X version
- PyGTK 2.16.0 or newer
- python-ldap (required when using ThinLinc LDAP tools.)
- CUPS (Common UNIX Printing System) (only required when using nearest printer or local printers, see Chapter 5)
- An SSH (secure shell) server
- Accurate time synchronization between all ThinLinc servers

As long as your platform fulfills the requirements above, ThinLinc should work as expected.

3.2.2. Server Sizing

The amount of computer resources needed to run a ThinLinc cluster varies greatly with the number of users, the type of hardware used for the servers, the application mix run by the users and the type of users. Trying to estimate the number of servers needed for a specific cluster is not something that can be done using a predefined table of facts. Instead decisions should be made based on benchmarks and experience.

Below, we will try to give some ideas on what kind of resources are needed based on customer experience. With time and experience from your own cluster with your own application set, you will work out your own set of figures.

It is important to remember that the ThinLinc load balancing feature makes it easy to add another server when the need arises. Start out with a number of servers and add more as the load increases.

3.2.2.1. Types of Resources

There are several types of resources needed in a ThinLinc cluster.

- *Disk*

About 100MiB of disk is needed for the software and data being part of ThinLinc. Each active session also requires a very small amount of data (normally less than 100KiB) for storage of session data and the session log. In addition to that, there must be disk available for the operating system, the applications users run and logs.

- *CPU*

The amount of CPU is very hard to estimate as it depends completely on the set of applications run by the users, and also on how active the users are as well as which response times are accepted by the users. A server that without problem copes with 100 users running LibreOffice calc updating a spreadsheet now and then will cope with a considerably lower amount of concurrent users if they are accessing internet sites with streaming video.

For a full desktop (KDE or Gnome) with typical office and internet applications (LibreOffice, Firefox, some graphics program and users visiting multimedia-intensive web pages, the amount of CPU needed is somewhere between 150 and 300MHz per active user.

The CPU figures above are based on experience from customers running Intel Xeon 7140M (Netburst) CPUs. For other types of CPU, the figures should be adjusted accordingly.

- *Memory*

The amount of memory, just as the amount of CPU, is also very dependent on type application set and how active the users are.

For a full desktop (KDE or Gnome), expect the need for 100-200MiB of memory per user, not including the memory required for individual applications.

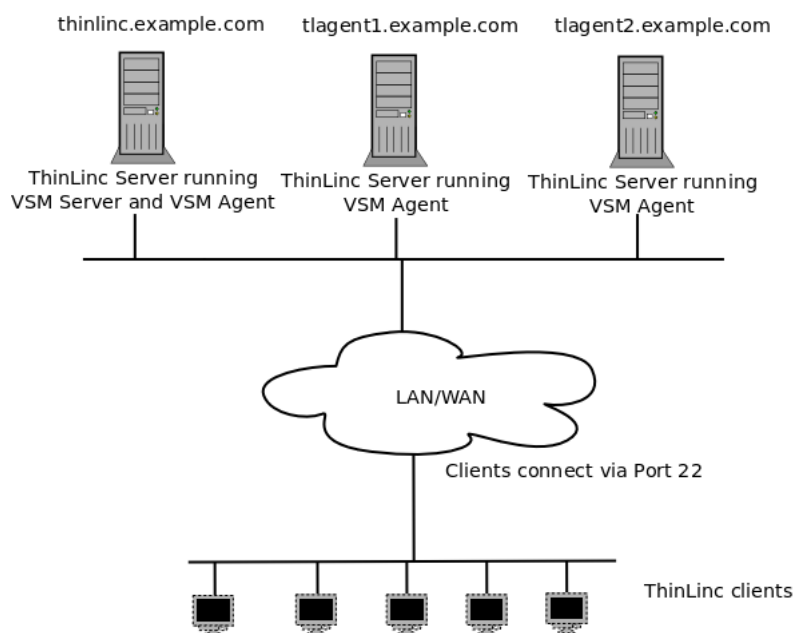
3.3. Preparing the Network for ThinLinc Installation

Naturally, the network at the site where ThinLinc is to be installed needs to be prepared for the installation. This section aims to help in understanding the requirements of the network for a successful ThinLinc installation.

We will explain the most common setups, including a typical Novell site and a typical Microsoft site. Also, we will explain how a site with NAT can use a NAT/Split-DNS setup to access ThinLinc in an efficient way both from the inside network as well as from the Internet.

3.3.1. A Simple ThinLinc Setup

Figure 3-1. A Simple ThinLinc Setup



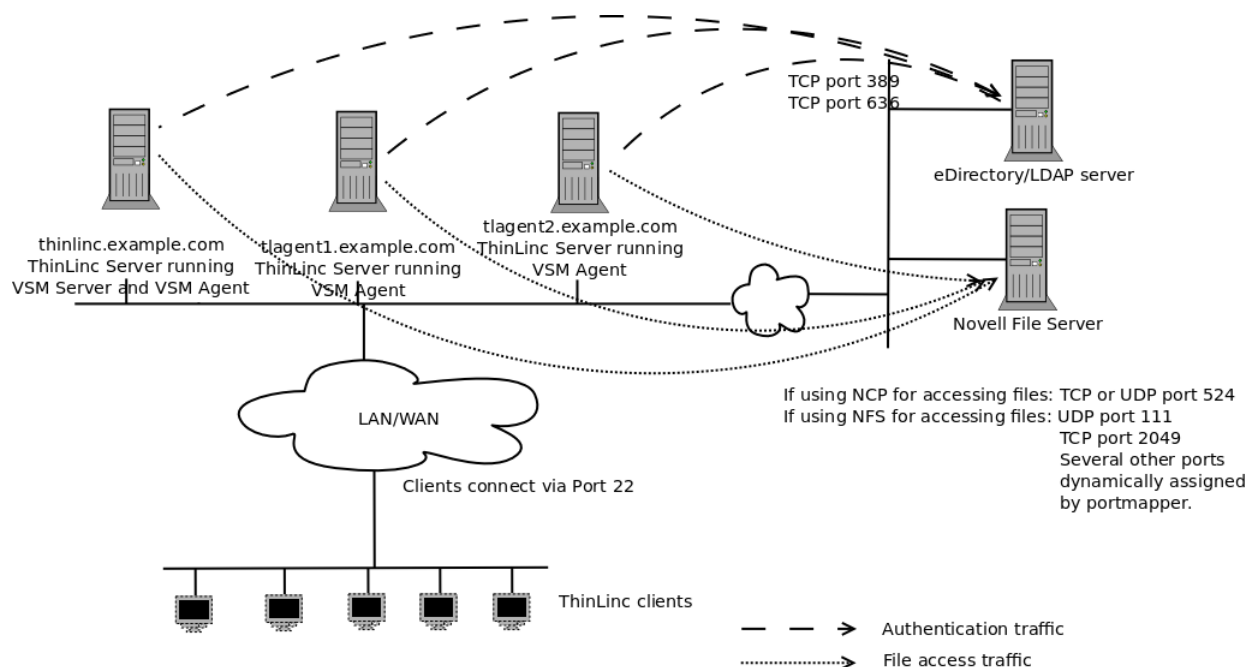
In Figure 3-1, a very simple ThinLinc setup is shown. In this setup, clients are configured to connect to *thinlinc.example.com*, DNS is configured with information about what IP addresses correspond to the hostnames *thinlinc.example.com*, *tlagent1.thinlinc.com* and *tlagent2.thinlinc.com* and no firewalls are in the path between the clients and the servers.

The number of VSM agents will range from 1 (on the same host as the VSM server) to a larger number, based on the number of users that are using the system. In this example, there are one host running both VSM server (the software controlling the whole ThinLinc cluster) and VSM agent, and two dedicated VSM agent hosts running only sessions.

Clients will communicate with the servers solely via SSH (by default port 22).

3.3.2. ThinLinc in a Novell Network

Figure 3-2. ThinLinc in a Novell Network



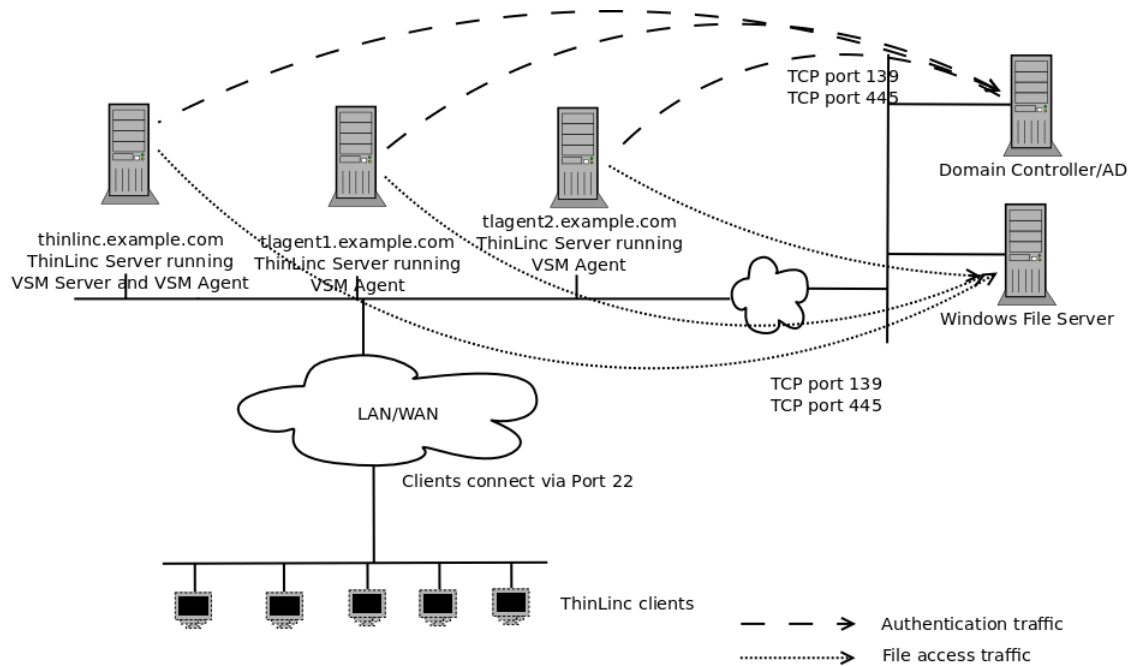
In Figure 3-2, ThinLinc is installed in a Novell environment, and integration with Novell eDirectory and/or Novell Netware filesystems are in use.

The ThinLinc servers will need to communicate with the eDirectory servers on either port 389, if using unencrypted LDAP, or on port 636, if using encrypted LDAP (ldaps).

The ThinLinc servers will also need to communicate with the Novell Netware file servers. In the case where NCP is used to access the files, the ThinLinc servers need to communicate with the Netware servers on TCP or UDP port 524. In the case where NFS is used to access files, UDP port 111, TCP and UDP port 2049 and a range of dynamically allocated UDP ports are used to communicate with the file servers. If there is a firewall between the ThinLinc servers and the Netware file servers, it needs to have support for understanding portmap requests, opening NFS UDP ports on demand, or there can be no restrictions for the traffic between the ThinLinc servers and the Netware file servers.

3.3.3. ThinLinc in a Windows Network

Figure 3-3. ThinLinc in a Windows Network



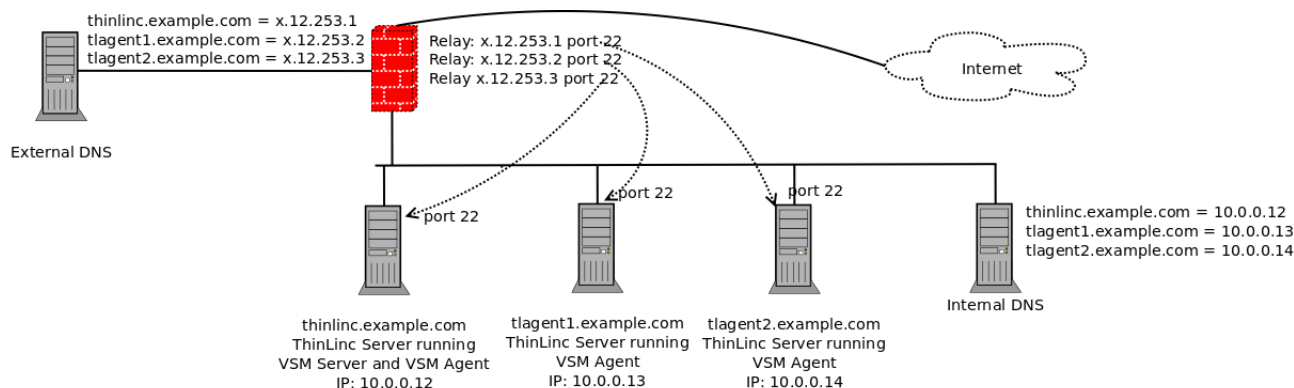
In Figure 3-3, ThinLinc is installed in a Windows environment, and integration with Windows Domain Services and/or Windows Fileservers are in use.

The ThinLinc servers need to communicate with the Windows Domain Controller on TCP port 139.

The ThinLinc servers will need to communicate with the Windows file servers using TCP port 139 and/or TCP port 445.

3.3.4. ThinLinc in a NAT/Split-DNS Environment

Figure 3-4. ThinLinc in a NAT/Split-DNS Environment



At many sites, the internal network is behind a firewall doing Network Address Translation (NAT). This means that the IP addresses on the internal network are allocated from so-called RFC1918 space, i.e., they are within the range 10.0.0.0-10.255.255.255, 172.16.0.0 - 172.31.255.255 or 192.168.0.0 - 192.168.255.255.

As long as ThinLinc servers are only meant to be accessed from the internal network, this is no problem, and the situation will be like the one described in Section 3.3.1. However, if the ThinLinc servers are meant to be accessed from the Internet as well, special arrangements need to be made.

Note: An alternative to using a split DNS configuration is to use a client side translation configured by the `HOST_ALIASES` parameter, but in most cases, a proper DNS setup is recommended. See Section 7.8 for more information.

3.3.4.1. Relays

First, relays must be configured in the firewall. One IP address reachable from the outside network per ThinLinc server needs to be available, and each should be equipped with a relay forwarding traffic from TCP port 22 on the outside to TCP port 22 on one specific ThinLinc server. In our example, as shown in Figure 3-4, there is one relay listening to TCP port 22 on the externally reachable IP address `x.12.253.1` forwarding traffic to the ThinLinc server on the internal network with IP address `10.0.0.12`, one relay listening on TCP port 22 on the externally reachable IP address `x.12.253.2` forwarding traffic to the ThinLinc server on the internal network with IP address `10.0.0.13`, and so on.

3.3.4.2. DNS

After configuring the relays, DNS must be configured so DNS queries for the hostnames of the ThinLinc servers get different answers depending on the origin of the query. DNS queries originating from the internal network should be answered with the real IP addresses of the servers, and DNS queries originating from the outside network should be answered with the IP addresses on the firewall, where the relays are listening.

In our example, if a host on the internal network is asking for the IP address of the hostname *thinlinc.example.com* it should get the IP address *10.0.0.12* as answer. If a outside host is asking for the IP address of the same hostname it should instead get the IP address *x.12.253.1* as answer.

When configured this way, a client connecting from the internal network will communicate directly with the ThinLinc servers, without the need to pass the firewall, while clients connecting from the outside will pass through the firewall and the relays to communicate with the ThinLinc servers. This will ensure optimal performance for clients from the internal network, at the same time lowering the load on the firewall.

3.3.4.3. Configuring the VSM Agents

Finally, after configuring relays and DNS, the VSM agents must be configured to respond with the correct hostname when asked by the VSM server what hostname the clients should connect to. The default behaviour is to respond with the IP address of the host, but that will not work in this case since clients connecting from the external network won't have any route to for example *10.0.0.13*. Instead, the VSM agents should be configured to respond with the hostnames that can be found in both the internal and the external DNS.

This is done by setting the parameter `/vsmagent/agent_hostname` on each of the VSM agents in the ThinLinc cluster. In our example, set `/vsmagent/agent_hostname` to *tlagent1.example.com* on the machine with IP address *10.0.0.13*.

3.3.5. Using ThinLinc Web Access

If users are supposed to be able to connect using a web browser, using ThinLinc Web Access, they must be able to connect to port 300 on both the VSM server and on all VSM agents.

In the NAT/Split-DNS setup, relays must obviously be configured in the firewall for each ThinLinc server and the port 300.

3.3.6. Other Services Required by ThinLinc Servers

In order for ThinLinc to function properly together with the rest of the network, they will need to synchronize time with some internal or external time source. Linux machines use the Network Time Protocol (NTP), so if there is one or several NTP servers on the internal network, the ThinLinc servers will need to communicate with them. Otherwise, the ThinLinc servers should be configured to use some external time source, and should be allowed to communicate with it.

3.4. Installing the ThinLinc Remote Desktop Server

3.4.1. Starting the Installation Program

The installation program is located in the root directory of the Server Bundle. Extract the bundle and start the installation program as follows:

```
sh ./install-server
```

If you prefer, you can also install the ThinLinc packages by hand. These packages are located in subdirectory `packages` of the Server Bundle.

After installing the software packages, ThinLinc must be configured. This is done by ThinLinc Setup, which is started by running `/opt/thinlinc/sbin/tl-setup`. If `install-server` is used, it will ask about starting ThinLinc Setup automatically at the end of the package installation. ThinLinc Setup must be run on all ThinLinc servers that make up the ThinLinc cluster. The role of the server in the cluster can be specified at the start of ThinLinc Setup, it's a choice between agent or master. Instructions for configuring newly installed master and agent machines to create a cluster can be found in Section 14.1.

3.4.1.1. Automating ThinLinc Setup

You can automate ThinLinc Setup by providing it with an answer file. Begin by generating an answer template by running the following command.

```
# /opt/thinlinc/sbin/tl-setup -g OUTPUT-FILE
```

A list of questions which ThinLinc Setup would ask is written to `OUTPUT-FILE`. Edit this file with suitable answers for your system. The file uses the same Hiveconf syntax also used for the ThinLinc configuration files, described in Chapter 16. You can now use the `-a` option for ThinLinc Setup to make it read answers from the given file.

```
# /opt/thinlinc/sbin/tl-setup -a INPUT-FILE
```

3.5. Upgrading an Old Installation

This chapter will detail the process of upgrading ThinLinc servers and clusters. There are several important items that has to be considered regarding ThinLinc upgrades:

- It is required that all servers (including HA nodes) in a cluster are running the same ThinLinc version.
- Users will not be able to reconnect to running sessions when the master service is stopped or when the agent service is stopped on the agent server running the session.
- Users will not be able to create new sessions when the master service is stopped or when no agent servers are available.
- Running sessions will be unaffected by a upgrade. This means that users can continue working. This also means that running sessions will not be getting the benefits from the new version.

3.5.1. Upgrading a Cluster

The recommended workflow for upgrading a ThinLinc cluster is as follows:

1. Review configuration changes in the release notes for the new release. More information regarding configuration migration can be found in Section 3.5.4.
2. Check licenses and install new ones if needed. For details see Section 3.5.2.
3. Schedule the upgrade and if necessary prepare the users on that reconnections or creation will be unavailable during the upgrade process. The command `tl-notify` can be used to send messages to users in running sessions. Documentation for this command can be found in Chapter 13.
4. Stop the agent services on all agent servers. The command `tl-ssh-all` can be used to run commands on all agent servers in the cluster. Documentation for this command can be found in Chapter 13. This step will prevent reconnections and the creation of new sessions.
5. Remove all agent servers from the cluster by clearing the configuration parameter `/vsmserver/subclusters/<name>/agents` on the master. Details on this parameter can be found in Section 14.2.3. Restart the master service to take the change into effect. If HA is used, do this on both master servers.
6. Upgrade the master server. Details for installing an upgrade can be found in Section 3.5.3 and Section 3.5.4. If HA is used, stop the master service on both master servers and then upgrade both servers.
7. Upgrade each agent server and manually add them back into the upgraded cluster. Upgrading agents works the same way as upgrading a master server. Add each upgraded agent to the configuration parameter `/vsmserver/subclusters/<name>/agents` on the master. Restart the master service afterwards. If HA is used, do this on both master servers. Once at least one agent is added users will again be able to create new sessions.

Once all agent servers are upgraded and added back into the cluster all users will be able to reconnect to existing sessions and the upgrade is complete.

3.5.2. New Licenses

Before performing an upgrade the first step is to find out if new license files are required to run the new version. ThinLinc license files delivered with version `x.y.z` will still work for versions with the same `x` and `y` but higher `z`, but not for increased `x` or `y`. For example, license files for ThinLinc 3.1.0 will still work for ThinLinc 3.1.1, but not for ThinLinc 3.2.0 or ThinLinc 4.0.0.

As the new licenses will work with the old (current) version, it's a good idea to install them as the first step in the upgrade process.

3.5.3. Upgrading the Packages

The same installation program that you used to install ThinLinc is also used to upgrade it. It is located in the root directory of the Server Bundle. Extract the bundle and start the installation program as follows:

```
sh ./install-server
```

and answer the questions. If you prefer, you can also upgrade the ThinLinc packages by hand. These packages are located in subdirectory `packages` on the Server Bundle.

If `install-server` was used, it will ask if ThinLinc Setup should be started at the end of the package upgrade. If ThinLinc Setup wasn't started automatically, it should be started manually after the package upgrade by running `/opt/thinlinc/sbin/tl-setup`.

3.5.4. Configuration Migration

Once the packages has been upgraded, a decision will sometimes be required regarding how to migrate the configuration. When a conflict between the saved configuration and the configuration in the new release arises, a choice has to be made.

ThinLinc Setup will present choices regarding migration of Hiveconf files. Conflicting files that aren't Hiveconf files are not affected by ThinLinc Setup. In these cases the package upgrade will have kept your configuration in place and saved the new default values from the new ThinLinc version as `.rpmnew` or `.dpkg-dist` versions of the conflicting files. Potential migration of non-Hiveconf files has to be done manually.

The three options that are presented in ThinLinc Setup are as follows:

- Use the new Hiveconf files and migrate the parameters and values from the old files.

With this option, all configuration changes done in the earlier version are preserved. The configuration will be based on the new files. Values of parameters that have been moved or renamed in the new release will be migrated to the new parameters. Parameters that have been removed will be deleted. Comments will not be migrated. The file structure and file names may also be different. All parameters and values from the listed Hiveconf files are copied over. This means that unchanged parameters in these files will use the default values from the earlier release.

Note that a certain parameter will be defined if it is defined either in the new or old Hiveconf files. This means that if you have removed some parameters, for example one of the example profiles, those parameters will again exist after the migration. For profiles, however, this will not affect the user session, since profiles are only visible if they are also listed in the "order" parameter.

Parameters will be removed from the new Hiveconf files if they are defined elsewhere. For example, if `/vsmagent/agent_hostname` has been moved from `vsmagent.hconf` to `local.hconf`, this change will be preserved.

- Use all old Hiveconf files.

With this option, all the old files are used. Custom comments and the file structure are preserved, but no new parameters or comments from the new release are introduced. Please note that configuration files which are identical in the old and new release are not listed or processed. This means that new default values in such files are introduced even with this option.

- Ignore old Hiveconf files and use the new files.

With this option, the listed configuration files are ignored and the new files are used instead. Please note that configuration files which are identical in the old and new release are not listed or processed. This means that configuration changes to such files are preserved even with this option.

3.6. SELinux enabled distributions

ThinLinc is designed to run with reference SELinux policy and users in the unconfined context. It is possible to use ThinLinc with other policies and more restricted contexts, but will most likely require modifications to your policy to accommodate ThinLinc.

The local system policy will optionally be modified by ThinLinc Setup during installation. The SELinux module and other policy changes performed can be examined in `/opt/thinlinc/share/selinux`. Execute the command `/opt/thinlinc/share/selinux/install` to reapply ThinLinc's policy changes.

Note: The ThinLinc policy module is distributed in source form and therefore requires the reference policy build environment. On Red Hat based systems this is always installed, but other systems might require extra packages.

3.7. VirtualGL

3.7.1. Overview

VirtualGL is used to provide server-side hardware 3D acceleration to applications displayed on a remote client. VirtualGL can be used with ThinLinc to provide accelerated graphics for OpenGL applications running in Linux environment.

Although ThinLinc is designed to work in combination with VirtualGL, VirtualGL is not developed or maintained directly by Cendio AB, and as such is not shipped as a part of the ThinLinc product.

3.7.2. Installation and configuration

Full documentation regarding the installation and configuration of VirtualGL can be found online at <https://virtualgl.org/Documentation/Documentation>.

Note: The following section numbers references the VirtualGL 2.3.3 documentation. Documentation for past or future VirtualGL releases may have different section numbers.

For the general case, it should be sufficient to consult the following sections:

- 5.1 - Installing VirtualGL on Linux

Chapter 3. Installation

- 6.1 - Granting Access to the 3D X Server

And see also:

- 9.1 - Using VirtualGL with an X Proxy on the Same Server

For more advanced configuration, such as using a remote application server with VGL Transport, see the following sections:

- 6.3 - SSH Server Configuration
- 8 - Using VirtualGL with the VGL Transport

Chapter 4. License Handling

4.1. Overview

To run a session against a ThinLinc cluster, the server must be equipped with license files. The license files specify the number of concurrent users the cluster is allowed to run.

If no license files are installed on the cluster, a maximum of five concurrent users are allowed.

Each cluster can have one or several license files. Each file contains licenses for a specific number of concurrent users. When the VSM Server starts up, it reads all license files and creates a sum of the number of concurrent users allowed based on the licenses from all files.

License files have one soft and one hard limit. When the soft limit is reached, new sessions can still be started, but a license violation will be logged and sent to the administrator (see Section 4.4). If however the hard limit has been reached, new sessions cannot be started. The purpose of this system is to allow growing organisations some time to adapt the number of licenses to a growing number of concurrent sessions, avoiding loss of production.

4.2. License Counting

One license is required for each pair of (*username, client hardware*). This means that if a user runs several sessions from the same client, only one license is used. If the same user runs multiple concurrent sessions from different client hardware, multiple licenses are required by the user.

4.3. Location and format of License Files

License files are delivered either in the form of text files (filename extension `.license`) or ZIP files (filename extension `.zip`). Transfer each file to your ThinLinc master server and place it in `/opt/thinlinc/etc/licenses`. Make sure that the transfer of the files uses binary mode, or the license file might not be verifiable. We recommend transferring via `scp` or `sftp`.

After adding new license files, either restart VSM Server by running `/bin/systemctl restart vsmserver` or wait until the VSM Server automatically reads in the new licenses, something that happens once every 12 hours.

Note: When running VSM Server in a High Availability setup (see Chapter 6), license files should be copied to `/opt/thinlinc/etc/licenses` on both nodes.

4.4. Log Files and E-mail Messages

ThinLinc logs user license violations to the file `/var/log/thinlinc-user-licenses`. Other license-related messages are logged to `/var/log/vsmserver.log`.

If license violations occurs, ThinLinc sends email to the person defined as system administrator in the parameter `/vsmserver/admin_email` in `vsmserver.hconf`. E-mail messages warning about license violations are sent every 12 hours if any license violations have occurred.

4.5. Checking the Number of Valid Licenses

You can use the program `/opt/thinlinc/sbin/tl-show-licenses` to verify the number of valid user licenses. There is also a graph available in the administrative interface. See Chapter 17 for more information.

Chapter 5. Printer Features

5.1. Overview of ThinLinc Printer Features

ThinLinc has several printer-related features that aims to provide the user with maximum flexibility while making the administrator's work easier. A ThinLinc system normally uses CUPS (Common Unix Printing System) to provide normal printing services. By integrating with CUPS, ThinLinc also provides the following features:

- *Local Printer support* allows users to print documents on a printer that is connected to their terminal from applications running on the ThinLinc server.

See Section 5.3 for documentation on this feature.

- *Nearest Printer* is a feature that simplifies the printing process for the user by automatically printing to a printer that is located at the terminal the user is currently using. Users only need to know that they should always print to the *nearest* printer - the system will figure out the rest based on a database of terminals, printers and locations, eliminating the need to learn the names of printers at different locations. This decreases the need for support.

See Section 5.4 for documentation on this feature.

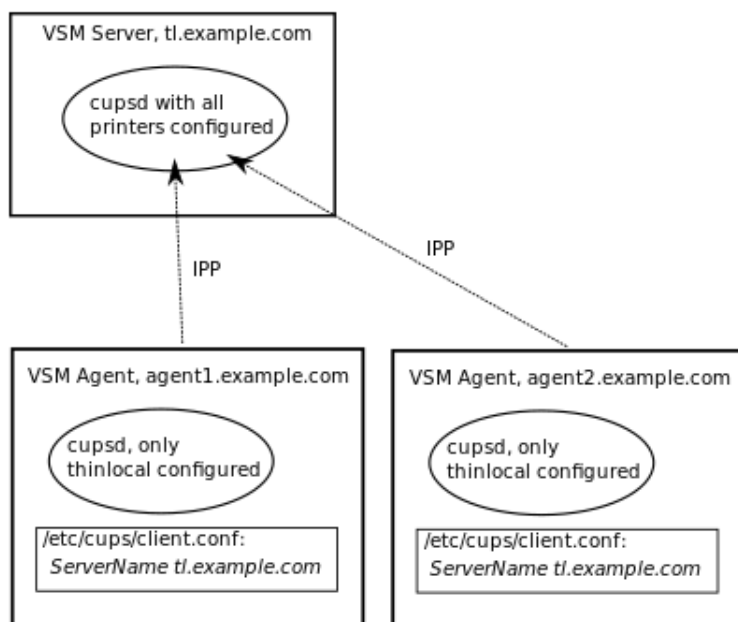
- *Printer Access Control* uses the same database of terminals, locations and printers as the *Nearest Printer* feature to dynamically limit which printers a user may print to based on the terminal the user is currently using. This feature also limits the list of printers seen by each user to the printers the user are allowed to use, simplifying choice of printer for the user by only showing the printers that are relevant at the current location.

See Section 5.5 for documentation on this feature.

5.2. Printer Configuration Overview

This section provides an overview of how printing is configured in a ThinLinc cluster.

Figure 5-1. Printer Configuration Overview



5.2.1. CUPS Browsing

It is important that the CUPS Browsing feature is turned *off* on all machines in the cluster, or problems with duplicate thinlocal printers will occur.

5.2.2. CUPS configuration on the Machine Running VSM Server

Configure all printers that need to be available in the CUPS configuration on the machine running VSM Server. Either use distribution-specific tools, or the built-in administration interface in CUPS which can usually be reached by using a web browser, connecting to port 631 on the machine, i.e. <http://tl.example.com:631/>.

The nearest and thinlocal queues, used by the nearest printer and the local printer features respectively, are added by ThinLinc Setup.

Printers, with one exception (see below) only needs to be configured on the machine running VSM Server. Agent nodes will use the CUPS daemon (cupsd) on the VSM Server machine for printing.

5.2.3. CUPS configuration on the Machine running VSM Agent

The machines in the cluster that run VSM Agent, i.e., the machines that host user sessions, need a running CUPS daemon (cupsd), but this cupsd only needs one printer defined - the *thinlocal* queue. The reason for this is that the local printer backend needs to run on the same machine as the session of the user printing to local printer to be able to access the endpoint of the SSH tunnel used to transport the printer job to the client.

The thinlocal queue is added by ThinLinc Setup when installing the agent.

Note: The CUPS daemon on each agent must listen to requests on the network interface, and allow printer jobs from the machine running VSM Server to be submitted to the thinlocal queue.

When a user submits a job to the local printer, i.e. to the thinlocal queue, the printer job will be submitted to the CUPS daemon running on the VSM Server host. It will then be respooled to the cupsd on the agent server hosting the session. This is to make central configuration of all other printers possible.

5.3. Local printer support

5.3.1. Theory of operation

With ThinLinc, it is possible to print to a printer attached to the client computer. Two primary modes of operation are available: device independent and device dependent. Both modes can be used at the same time. See below for details about the two modes.

The thinlocal printer is cluster-aware. If a user submits a print job on a node in a ThinLinc cluster which does not host the user's session, the print job will automatically be respooled to the correct node. This is used in the recommended setup (see Section 5.2).

If a user has more than one session, print jobs submitted to the local printer will be redirected to the client that made the last connection.

The local printer features are implemented as a backend to CUPS (Common Unix Printing System).

Note: When using local printers, we recommend that you activate the parameter `/vsmserver/unbind_ports_at_login`.

5.3.2. Device independent mode

The device independent mode is designed to provide universal access to any local printer without having to install drivers on the ThinLinc server. This is achieved by converting the print job to the Adobe Portable Document Format (PDF) on the remote desktop server, and then sending it through an encrypted tunnel to the client. The client subsequently prints the job on the local printer using a built-in PDF renderer.

Because the driver on the ThinLinc server is device independent, it has no way to know what capabilities (duplex ability, trays, paper size, etc.) the printer connected to the client has. At the same time, applications that want to print need to know about these capabilities to print correctly.

As a compromise, the universal printer is configured with a PPD (Postscript Printer Definition) that covers a broad range of printer capabilities - it's a *Generic Postscript Printer* driver. This makes it possible for CUPS to convert input formats to the correct format before sending them to the local printer. It also means that default values can be set for some of the configuration parameters, for example paper size, using the CUPS configuration interface.

5.3.3. Device dependent mode

The device dependent mode is to be used when it is necessary to access all options on the printer, or when the communication with the printer cannot be expressed in terms of normal pages (e.g. a label printer). In this mode the printer driver is installed on the ThinLinc server and the data is sent unmodified to the local printer.

Note: ThinLinc has no way of verifying that the connected printer is the correct one, so it is up to the user to make sure that a device dependent queue is not used with a different printer.

5.3.4. Installation and Configuration

Use ThinLinc Setup to install the PDF conversion filter, the backend and queue in CUPS on all machines running VSM Agent. This adds a new queue named *thinlocal* to CUPS and makes it available to your users. This queue is the one to use for device independent mode described above.

After installation, the local printer is ready for use. Make sure your ThinLinc client is configured to allow redirection of printers, then print to the *thinlocal* queue, and the job will be rerouted to the default printer of the client you're currently using.

Device dependent queues are installed as if installing the printer locally on the ThinLinc server. The only difference is that the URI shall be specified as `thinlocal:/.` Example:

```
# lpadmin -p thinlocal-label -v 'thinlocal:/' -P /media/cd/label-printer.ppd
```

5.3.5. Parallel port emulation

ThinLinc also includes a very basic form of parallel port emulation that gives legacy application access to the local printer. It is built on top of the *thinlocal* queue, which means it only works if certain requirements are satisfied:

- The application must only write to the port. Reading is not supported, neither is monitoring or altering the port status pins.

- After a print job is completed, the application must close the port. As the emulation is unaware of the printer protocol, closing the port is the only way it can determine where one job ends and another begins.

To access the emulated parallel port, configure the application to use the port

```
$TLSESSIONDATA/dev/lp0.
```

5.4. Nearest printer support

With the ThinLinc *nearest printer* feature, printer jobs are sent to a printer selected based on the physical address of the users terminal. This is typically used to implement printer queues based on physical proximity.

The *Nearest Printer* is implemented as an extra printer queue, on top of the real printers. Printer jobs sent to the *nearest* queue will be sent to the Nearest Printer backend. The backend is a program which is called by CUPS together with all needed information. The backend will look at the user name requesting the printout and ask the ThinLinc VSM server for more information about this user. The information includes which terminal the user is currently using. The backend then queries the information stored in Hiveconf for a list of printers that are considered near the terminal used by the printing user. When a printer is known the backend will place the job in that printer queue.

The *nearest* queue is added to the VSM master server by ThinLinc Setup. The recommended setup is to configure one *nearest* printer queue in the CUPS daemon on the VSM Server host, and then let all agents use this CUPS daemon. See Section 5.2 for an overview of printer setup in a ThinLinc cluster.

5.4.1. Administration of the Nearest Printer Feature in ThinLinc

The Nearest Printer system needs information about groups of terminals, known as Locations, which typically represents some physical layout. The information connects Terminals to Locations and also links printers to the Locations. Available printers are automatically fetched from the underlying printing system and are available for assignment to Locations and/or Terminals.

Information about Terminals, Locations and their associated printers can be administrated using the ThinLinc Web Administration, see Chapter 17.

Each Location should be entered with a name, and may have an optional description. A Location can for example represent a classroom, a department, a house, and so on. Each Location can be associated with one or more printers, including the special 'nearest' and 'thinlocal' printers. Typically it will include all printers available near that physical location the Location represents. If the location is so big that different printers are close to different parts of the location, then you should probably divide the Location into smaller parts, each represented by a separate Location.

A Location can be set to handle clients which are not defined using a Terminal definition ("unknown terminals").

Each Terminal in the ThinLinc Web Administration represents one physical terminal in the installation and is defined by its terminal network interface hardware (MAC) address. The hardware address can be entered in many formats, but will be converted to all uppercase hexadecimal form separated by colon, i.e. "01:23:45:67:89:AB".

A Terminal must be associated with a Location.

5.4.2. Nearest Printer Selection Algorithm

If a terminal has a printer directly assigned to it in the terminals module in `tlwebadm`, that printer will be the nearest printer for that terminal. For Terminals without a printer directly assigned (the normal situation), the first printer in the list of printers for the terminal's Location is selected when the user submits a printer job to the *nearest* queue.

If the client is not a known Terminal, i.e. its hardware address was not found, it will use the printer for the Location marked as handling "unknown terminals". If not, there will be no printer available.

If a user is using multiple sessions, print jobs submitted via *nearest* printer will be redirected to the printer that is found starting from the client that made the last connection.

5.4.3. Printer Drivers

When printing via the *nearest* printer, the CUPS client can't get hold of all information about the real printer where the job will actually be printed, because it doesn't know that the printer job will be rerouted by the nearest driver. Therefore, the printing application has no way to know about the number of trays, the paper sizes available etc.). This is a problem for some applications, and it also adds to the number of applications that will be misconfigured, for example selecting the wrong paper size.

As a compromise, the nearest printer is configured with a PPD (Postscript Printer Definition) that covers a broad range of printer capabilities - it's a *Generic Postscript Printer* driver. This makes it possible to configure default values for some of the settings, for example paper size, using the CUPS configuration interface.

If all the printers in your organisation are of the same type, it may be a good idea to replace the Generic Postscript PPD installed for the nearest queue with a PPD for the specific printer in use. That will let CUPS-aware applications select between the specific set of features available for the specific printer model.

5.5. Printer Access Control

In a ThinLinc cluster, all printers that any user of the cluster needs to be able to print to must be defined centrally, or the user will not be able to print from applications that run in a ThinLinc session. For large installations, this leads to a very long list of available printers.

A long list of printers leads to usability problems - having to select printer from a long list can be troublesome. Also, it opens for problems with printer jobs being printed at remote locations by mistake (or on purpose, by users finding it amusing to send "messages" to other locations).

The solution to this problem is the Printer Access Control feature of ThinLinc. By integrating with CUPS (the Common Unix Printing System), the list of printers a user is presented with and allowed to print to is limited to the printers that should be available to a specific terminal, based on information in a database of printers, terminals and locations.

Note: The Printer Access Control feature will affect all users on the ThinLinc cluster. The only user excepted from limitations of the printer list is the superuser (root) - all other users will only see and be able to use printers based on the location of their terminals, when the Printer Access Control feature is enabled.

5.5.1. Theory of Operation

Each time a user requests a new session or reconnects to an existing session, the hardware (MAC) address of the terminal is sent along with the request from the ThinLinc client. Using the same database as the *nearest printer* feature used to find which printer is closest to the user, the printer access control feature calculates which printers the user is allowed to use, and then configures the access control of the printing system (CUPS).

This way, the user is presented with a list of printers that only contains the printers relevant for the location where the terminal the user is currently using is located. In a situation where a user has multiple sessions running from multiple clients, all printers associated with the different terminals will be made available.

5.5.2. Requirements

- CUPS v1.2 or higher.

5.5.3. Activating the Printer Access Control Feature

First, make sure you have configured the printers in your ThinLinc cluster as documented in Section 5.2. For the Printer Access Control Feature, a central CUPS daemon on the VSM Server host is required, and all agent hosts must have a correctly configured `/etc/cups/client.conf`.

To activate the printer access feature, create two symlinks on the host running VSM Server, as follows:

```
ln -s /opt/thinlinc/sbin/tl-limit-printers /opt/thinlinc/etc/sessionstartup.d
ln -s /opt/thinlinc/sbin/tl-limit-printers /opt/thinlinc/etc/sessionreconnect.d
```

The first symlink makes sure **tl-limit-printers** is run when sessions are started. The second makes sure it is run at reconnects to existing sessions. More details about the session startup can be found in Section 14.4.

Note: With the above configuration (symlinking `tl-limit-printers` into `sessionstartup.d` and `sessionreconnect.d`), the client will not get an answer back from the server until `tl-limit-printers` has finished its execution. This is the desired behaviour if it is strictly necessary that printer access rights are correct when the user connects to the session. In environments where it is acceptable that the final list of printers shown to the user may not be finished when the user connects to the session, place the execution of `tl-limit-printers` in the background, as detailed in Section 14.4.1.1, as that will decrease the time the user has to wait for the session to appear on his client.

After creating the symlinks, try connecting to your ThinLinc cluster with a ThinLinc client and bring up an application that lists the available printers. The list of printers should now be limited according to configuration.

Note: The printer list limitation doesn't work for applications that use the deprecated *cupsGetPrinters* library call. This means that older applications might show the whole list of printers. The access control are still enforced, which means that even if a disallowed printer is shown in the list of printers, users can't submit jobs to it.

Most applications in a modern Linux distribution doesn't have this problem.

5.5.4. Configuration

Configuration of the printer access control feature is mostly a matter of using *tlwebadm* (see Chapter 17 for details) to add the hardware address of all terminals as well as information about where they are located and which printers are to be available for each location.

5.5.4.1. Unknown Terminals / Terminals Without Hardware Address

When a client reports a hardware address that is not present in the database of terminals, or when no hardware address is reported, the default behaviour is to disallow access to all printers, rendering an empty printer list for the user.

There is however a way to give even unknown terminals access to one or more printers - define a special location and check the *Use for unknown terminals and terminals without hardware address* checkbox. Then add the printers that should be available for the unknown terminals.

One common configuration is to add such a location and then add the *thinlocal* printer to this location. This way, unknown terminals, for example people working from their home computers, will be able to use their locally connected printer, but no other printer will be available.

Chapter 6. High Availability (HA)

6.1. Overview

This chapter describes how to setup ThinLinc with High Availability (from now on referred to as "HA") for the VSM server. Since the VSM server service handles load-balancing and the session database, it can be problematic if the machine fails. ThinLinc HA provides protection for this service against the single point of failure that the hardware running the VSM server normally is.


The basic principle behind this setup is to have two equal machines, both capable of running VSM server. If one of the machines goes down for some reason, the other machine will take over and serve VSM server requests with no or short interruption of service.

Note: The HA functionality provided by ThinLinc provides synchronization of the ThinLinc session database across two VSM servers. The software used by these machines to implement failover is not part of ThinLinc, and must be installed and configured according to your requirements. The industry standard for doing so on Linux is provided by the Linux-HA project; see <http://linux-ha.org> for more information.

6.1.1. Background - Reasons For a HA Setup

In a standard ThinLinc setup, there is a single point of failure - the machine running the VSM server. If the VSM server is down, no new ThinLinc connections can be made, and reconnections to existing sessions can't be established. Existing connections to VSM agent machines still running will however continue to work. A ThinLinc cluster of medium size with one machine running as VSM server and three VSM agent machines is illustrated in Figure 6-1

Figure 6-1. A non-HA ThinLinc cluster setup



Here the incoming connections are handled by the VSM server which distributes the connections to the three VSM agent machines. If the VSM server goes down, no new connections can occur. The VSM server is a single point of failure in your ThinLinc setup.


6.1.2. Solution - Elimination of Single Point of Failure

In order to eliminate the single point of failure, we configure the VSM server in a HA configuration where two machines share the responsibility for keeping the service running. Note that ThinLinc's HA functionality only handles the parts of your HA setup that keeps the ThinLinc session database synchronized between the two machines. Supplementary software is required, read more about this in Section 6.1.3.

When ThinLinc as well as your systems are configured this way, the two machines are in constant contact with each other, each checking if the other one is up and running. If one of the machines goes down for some reason, for example hardware failure, the other machine detects the failure and automatically takes over the service with only a short interruption for the users. No action is needed from the system administrator.

6.1.3. Theory of Operation

Figure 6-2. A ThinLinc HA cluster setup



In a HA setup, as illustrated in Figure 6-2 two equal machines are used to keep the VSM server running. One of the machines is primary, the other one is secondary. The primary machine is normally handling VSM server requests, but if it fails, the secondary machine kicks in. When the primary machine comes online again, it takes over again. That is, in normal operation, it's always the primary machine that's

working, the secondary is just standby, receiving information from the primary about new and deleted sessions, maintaining its own copy of the session database.

Both machines have an unique hostname and an unique IP address, but there is also a third IP address that is active only on the node currently responsible for the VSM server service. This is usually referred to as a resource IP address, which the clients are connecting to. ThinLinc does not move this resource IP address between servers, supplementary software is required for this purpose.

6.2. Configuration of ThinLinc for HA Operations

In this section, we describe how ThinLinc is configured for High Availability.

6.2.1. Installation of a New HA Cluster

In this section, we will describe how to setup a new HA cluster. In the examples we will use a primary node with the hostname *tlha-primary* and IP address 10.0.0.2, a secondary node with the hostname *tlha-secondary* and IP address 10.0.0.3, and a resource IP address of 10.0.0.4 with the DNS name *tlha*.

1. Begin by installing ThinLinc as described in Chapter 3 on both nodes.
2. Both nodes in the HA cluster must have the same SSH host key. Copy `/etc/ssh/ssh_host_*` from the primary host to the secondary host, and restart ssh on the secondary host.
3. Install and configure the system-level high-availability software, for example the software provided by the Linux-HA project, which can be found at <http://linux-ha.org>. This and other high-availability software may also be provided as part of your distribution, so check for the solution which best fits your requirements before proceeding.
4. Configure the system's high-availability software to watch the status of the other machine via the network, and to enable the resource IP address *10.0.0.4* on the active node. The machine with the hostname *tlha-primary* should normally be active.
5. Configure each VSM agent to allow privileged operations both from *tlha-primary* and *tlha-secondary*:

```
[root@agent root] tl-config '/vsmagent/allowed_clients=tlha-primary tlha-secondary'
```

Also, set the master_hostname to the DNS name of the HA interface:

```
[root@agent root] tl-config /vsmagent/master_hostname=tlha
```

Restart all VSM agents after changing the configuration values.

If the `tl-config` command is not found, logout and login again in order to let the login scripts add `/opt/thinlinc/bin` and `/opt/thinlinc/sbin` to the PATH.

6. Verify operations of VSM Server on both nodes. Make sure you can start the VSM server properly on both hosts, and connect to the respective hosts when VSM server is running (i.e., it should be possible to connect, using `tlclient`, to both *tlha-primary* and to *tlha-secondary*).

Both nodes should be configured with the same subcluster configuration.

Warning

It is **VERY IMPORTANT** that *127.0.0.1* is not in the list of agent servers of any subcluster. If the machines running VSM server are also VSM agents, their unique hostnames or IP addresses must be added to `/vsmserver/subclusters/<name>/agents` instead of *127.0.0.1*. The reason for this is that *127.0.0.1* will be a different server based on which VSM server is currently active.

7. After verifying that normal ThinLinc connections work as intended when using both the primary and the secondary VSM server's hostname, it is time to enable HA in the VSM servers. This is done by setting `/vsmserver/HA/enabled` to 1, and by specifying the nodes in the cluster in `/vsmserver/HA/nodes`. For example:

```
[root@tlha-primary root] tl-config /vsmserver/HA/enabled=1
[root@tlha-primary root] tl-config '/vsmserver/HA/nodes=tlha-primary.example.com tlha-secondary.example.com'
```

Configuration should be identical on both nodes. Restart the VSM server on both nodes after configuration.

8. If `vsmserver` can't safely determine which of the two nodes in `/vsmserver/HA/nodes` is the remote node, and which is the local node, it will start without HA enabled, and log a message. If this happens, validate your hostname and DNS setup. One of the entries of `/vsmserver/HA/nodes` must match the local machine. Either the resolved IP of one of the entries in `/vsmserver/HA/nodes` must match the local IP, or one entry must exactly match the local hostname as returned by `uname -n`.
9. Once HA has been configured, tests should be performed in order to confirm that the failover works as expected. This can normally be done by simply removing the network cable from the primary node, and ensuring that the secondary node then takes over. Check also that any active ThinLinc sessions have been synchronized from the primary to the secondary node, and that logging in to such a session on the secondary node succeeds once the primary node has been disabled.

Your ThinLinc HA cluster is now configured! When sessions are created, changed or deleted on the currently active node, the information about them will be transferred to the other node using a inter-VSM server protocol. If the other node has to take over service, its copy of the session data will be up to date, and it can start serving new requests instantly. When the primary node comes up again, the secondary node will resynchronise with the master.

6.2.2. Reconfiguring an existing ThinLinc Installation into HA mode

If you have an existing ThinLinc installation and want to eliminate the single point of failure (the VSM server), the procedure is very much like the procedure for installing a new HA cluster.

6.3. Recovering from hardware failures

If situations occur where the secondary node has been forced to take over service because the primary node failed for some reason, it's important to know how to recover.

6.3.1. Recovering from Minor Failures

If the primary went down because of a minor failure (overheating trouble, faulty processor, faulty memory etc.) and the contents of the files in `/var/lib/vsm` are untouched, recovery is very simple and fully automatic. Simply start the server and let the two VSM servers resynchronize with each other.

6.3.2. Recovering from Catastrophic Failure

If a catastrophic failure has occurred, and no data on the disks of the primary can be recovered, ThinLinc needs to be reinstalled and HA must be reinitialized.

Install ThinLinc as described in Section 6.2.1, but before starting the VSM server after enabling HA in the configuration file, copy the file `/var/lib/vsm/sessions` from the secondary to the primary. That will preload the database of active sessions with more current values on the primary.

Chapter 7. The ThinLinc Client

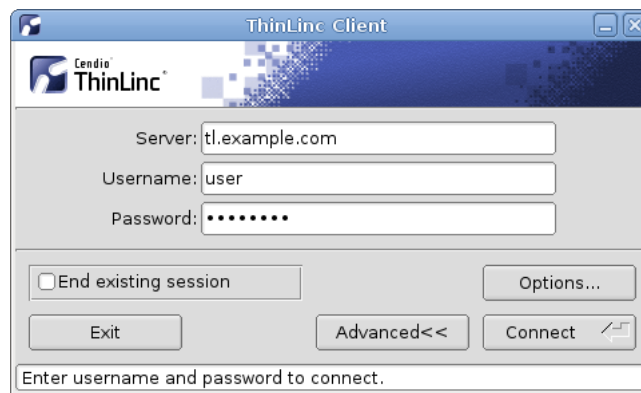
7.1. Client usage

Starting the ThinLinc client is normally easy, but the method can differ somewhat between the available operating systems. See Chapter 8 for instructions on how to start the client on different platforms.

7.1.1. The started ThinLinc client

When the ThinLinc client is started it will show the login window. This window contains a ThinLinc logo, text fields where needed information can be entered, buttons for control and at the very bottom a status field that gives information about the login procedure.

Figure 7-1. The ThinLinc client login window



7.1.2. Logging in to a ThinLinc server

To login into a ThinLinc server the client needs to do a successful user authentication. This means that it needs to tell the ThinLinc server a user name and a corresponding authentication information (a password or an encryption key). The ThinLinc server controls that the information is valid and accepts or denies the login attempt.

The things the client needs to know to successfully login the user against a ThinLinc server is a server address, a user name and the corresponding authentication information. When the client is normally started it will display two text fields labeled "Server" and "Name", and one text field labeled "Password", "Key" or "Certificate". This can differ some depending on command line arguments, but this is the normal behavior.

Accepted values for the server field is the hostname or the IP address of the server. The name field should be filled in with the ThinLinc username. The authentication information needed depends on the type of authentication used:

- For password authentication, a plain text password should be entered. The password won't be shown as clear text when entered.
- For public key authentication, the path to an encryption key must be entered or browsed to using the "..." button.
- For smart card authentication, a certificate must be chosen using the drop down menu next to the certificate name field.


The server name, username, key path and certificate name are saved when the user tries to start the session, so they don't have to be entered again each time a new session is wanted.

When the user has entered server address, username and authentication information, it becomes possible to login. This is done by pressing the *Connect* button or the *enter key* on the keyboard. The client will then try to establish a connection with the ThinLinc server. If any of the fields has a bad value that prevents the client from successfully logging in, for example if the username or password is incorrect, there will be response message shown as a message box with the suiting information.

Note: By default, usernames are case-sensitive when logging in via the ThinLinc client. This behaviour may be changed using an option in the client configuration file - see `LOWERCASE_LOGIN_NAME` in Section 7.8 for details.

If the login attempt is successful a ThinLinc session will start, an old one will be reused or a session selection box might be presented, all depending on the client's settings and how many sessions the user has running, a . See Section 7.4.1 for more information on how the choice is made.

Figure 7-2. The ThinLinc client session selection window



The session selection window presents the user with a list of relevant sessions and several buttons to act on those sessions:

Connect

Connect to the selected session, or create a new session if the current selection is "Create new session...".

End session

Forcefully terminate the selected session and restart the connection procedure.

Cancel

Abort the connection and return to the main window.

The server will then prepare a graphical session on a ThinLinc server. The client then connects to this session and displays it. Normally the user now sees a dialog with different session options. The user can there select for example to run a Linux session or a Windows session. Depending on the choice the server at the other end will start that kind of session.

7.1.3. Language Settings

The ThinLinc client gets all its strings from a database. This way it can be easily translated, by just providing a new database for a new language.

On Linux based systems, the client picks up which language to use by reading the standard POSIX *locale* environment variables. A somewhat simplified description of these follow here:

- `LC_ALL` : If this environment variable is set, it takes precedence over all other locale variables. It will affect all locale settings, including message strings, sorting order, money representation, decimal numbers, etc.
- `LC_MESSAGES` : If `LC_ALL` is not set but this one is, it will make the messages of the client adhere to the language in question, in effect making the client use that language. There are several other variables of this kind, but they do not affect the ThinLinc client.
- `LANG` : If `LC_ALL` is not set then the value of this variable will be used for all locale categories that are not explicitly set, e.g. `LC_MESSAGES` .

There is also a variable called `LANGUAGE` on some systems, but it is non-standard, and we do not recommend the use of it.

If none of these variables are set, the locale defaults to C, which in practice means American English. The value of the variables should be of the form *language_country*, where language and country are 2 letter codes. Currently, the languages delivered with the client are Brazilian Portuguese (`pt_BR`), English (`en_US`), Dutch (`nl_NL`), French (`fr_FR`), German (`de_DE`), Italian (`it_IT`), Russian (`ru_RU`), Spanish (`es_ES`), Swedish (`sv_SE`), and Turkish (`tr_TR`).

On Windows, the same environment variables can be set in a script that also starts the ThinLinc client. An example script called `altlang.cmd` is installed with the ThinLinc client for Windows. If nothing is set, the Windows client will use the language setting that was given with the control panel.

7.1.4. The ThinLinc session life cycle

When the user has started a ThinLinc session the client login interface disappears from the desktop. The client program continues to run in the background as long as the ThinLinc session is running. The client enters a service mode where it handles services needed to fulfill the requested features. For example the client handles the export of local printers, serial ports, and so on. When the ThinLinc session quits the client service engine quits as well.

There are several ways a session can end. The most common one is that the user chooses to logout from the session. That causes the session to finish on the server side. The ThinLinc server finds out that the session has finished and disconnects the client. Another possibility is to intentionally disconnect the client, without finishing the session on the server. This can be done by using the session menu. See Section 7.1.5 below for information about how to do this. When the client disconnects before the session running on the server is told to end, then the session will continue to run on the server. The next time the user logs in the server will reconnect the user to the very same session. This way it's possible to, for example, disconnect a session at work, go home, reconnect to that session and continue to work.

If the user knows that there already is a session running on the server, but still wants to start a new fresh session, then it's possible to check the *End existing session* check box that exists in the client login interface (advanced mode only). The client will then tell the server that it wants to end the existing session (if it exists) and get a new one.

Network issues or problems with ThinLinc services can sometimes prevent the servers from checking the status of a session. Such a session will be considered unreachable and the client will not be able to reconnect to it. The user can choose to abandon the session or wait for the problem to be resolved. However abandoning the session causes the ThinLinc server to stop tracking it and can leave applications running without any way of reaching them.

7.1.5. The session menu

When the ThinLinc session is authenticated and the ThinLinc session is running it's possible to control the session. For example it's possible to change between full screen mode and window mode, and to disconnect the ThinLinc client from the server.

To switch to windowed mode there is a session menu that pops up when the user press a predefined key. The default key for this is F8, but the key is configurable from the client options. See Section 7.4.1 for more information about how to change this key. In the session menu you should select *Full screen* to toggle full screen mode.

7.2. Running the ThinLinc client from the command line

To run the ThinLinc client from the command line you run the program `tlclient`, optionally followed by options and a server name. The correct program syntax is as follows.

```
tlclient [options] [server][:port]
```

The optional *server* field can be used to specify a ThinLinc server that should be predefined in the server field when client is started. The optional *port* parameter causes the client to try to connect another TCP/IP port number than the normal SSH port when establishing it's secure connection to the ThinLinc server. More information about custom SSH settings is available at Section 7.4.5.

The ThinLinc client is highly controllable from the command line by the use of command line arguments. Many parts of the client can be controlled this way. The more simple things to control is the server or user name. It is possible to force settings and lock tabs and fields in the config interface to prevent them from being changed.

All arguments written on the command line overrides the settings saved from previous sessions. The options window will show the current settings, including the settings from the command line. The client settings is only stored to file when the user press the *OK button* in the options window. This means that options from the command line normally don't affect the saved settings. But if the user opens the options window and accepts the settings by pressing the *OK button* then the settings, including the one from the command line, will be saved.

For a complete list of arguments supported by your client you can run the client with the argument `-?`.

Description of available command line arguments

Here follows a description for all available command line arguments.

`-?, --help`

Display a help summary.

`--version`

Display client version information and exit.

`-a, --advanced`

Start client in advanced mode. Advanced mode means that the client will show the *Server field*, *Options...* button and the *End existing session* checkbox. The advanced mode is the normal mode used when you start the ThinLinc client. A simpler mode, where those interface components are hidden, is used automatically when you enter a server name as a command line argument. By adding this argument you override that and always use the advanced mode.

`-C, --configfile FILE`

Specifies an additional configuration file. Parameter values in this configuration file overrides the values specified by the system wide and user configuration file. Settings changed from the GUI will be stored in this configuration file, instead of the user's configuration file.

`-d, --debug LEVEL`

The ThinLinc client logs information about the current session to the file `~/thinlinc/tlclient.log` on Linux systems and `%TMP%\tlclient.log` on Windows systems. When the client is started, any existing log file is renamed as `tlclient.old.log`. The amount of information to log can be configured with this option followed by a number from 1 to 5. A low number gives less logging than a higher number. The default is a log level of 3. For more information about log file placement, see Section 7.7 below.

`-u, --user USER`

This option sets the user name that should be filled in into the *Name* field. This can be used to override the name that is automatically saved from last session. If you for example, in a school

classroom, want it to always start with an empty Name field, then you can use this parameter with the empty string "".

-p, --password PASSWORD

This option sets the password that should be filled in into the *Password* field. When this option is used and a user name exists (either saved from previous session or entered with the *-u* parameter) the client will automatically try to login, directly after start. If the login try fails it will return focus to the client interface, making it possible to adjust the values. Note that the command line of *tlclient*, and therefore the password, will be visible to other processes running on the client operating system. If this is a problem in your environment, consider using the *-P* option documented below.

-P, --askpass PROGRAM

This option makes it possible to specify an askpass program that should be used to achieve the password. This program should in some way ask the user for a password and then return that password together with an exit code. This triggers the auto login (see argument *-p* above).

-e, --encodings ENCODING, ...

This option makes it possible to select which VNC encoding you want to use (see Section 7.4.4 for more information about VNC encodings). Valid encodings for this option are: *Tight*, *ZRLE*, *Hexile* and *Raw*.

-l, --lock ITEM, ...

This option makes it possible to lock different parts of the client interface. This can be used to prevent things from being changed. Locked parts will still be shown, but will be "grayed out", which means that they can't be made active for change. The items that should be locked should follow this option as a comma separated list. The following items are possible to lock.

- *server*: Server entry field
- *user*: Username entry field
- *options*: Options tab
- *localdevices*: Local Devices tab
- *screen*: Screen tab
- *optimization*: Optimization tab
- *security*: Security tab

-h, --hide ITEM, ...

This option makes it possible to hide different parts of the client interface. This can be used to remove parts of the interface that can confuse novice users, or to prevent them from reaching parts of the interface. The following items are possible to hide.

- *options*: options button

-f, --force SETTING, ...

This option makes it possible to force a setting to a value. This can be used to preset a client with values and to force them to reset to those values each time, even if the users make changes. When an option is forced it is turned on. The following items are possible to force.

- *terminate*: terminate session
- *fullscreen*: fullscreen mode
- *sound*: sound mode
- *sshcomp*: ssh compression

-M, --minimize

This option causes all other applications to be minimized when the ThinLinc client starts.

-s, --startprogram

Specifies the program to start in the session. Overrides the `START_PROGRAM_ENABLED` and `START_PROGRAM_COMMAND` configuration parameters.

--loop

This option causes the client to run forever. The exit button is removed, and when a session has ended, a new client process is automatically started.

Note: The only way to stop the client from restarting is to terminate the `tlclient` process.

7.3. Local device export

ThinLinc supports export of different local devices. This means that a device that exists on your client computer or terminal can be reached from the ThinLinc session that runs on the server. The type of devices that can be exported varies depending on which operating system the ThinLinc client runs on. The export is, very generalized, done by establishing secure tunnels for the data transmission and services that connect both ends. Here follows more information about each type of possible export; for detailed information about how to enable each type of export in the client, see Section 7.4.2 below.

7.3.1. Sound device

This feature makes it possible to hear sound from applications that runs on the ThinLinc server. Sound will be sent from the ThinLinc server to your local client through a secure connection. A small local sound daemon will be automatically started by the ThinLinc client. A secure tunnel for sound will be established during the ThinLinc session setup.

All programs that support PulseAudio should automatically be aware of this tunnel and send their sound to the client. See also Section 12.3 for information about supporting other applications.

The sound data that is sent from the server session to the local client is uncompressed audio data. This means that it can be relatively large and may use relatively much network bandwidth. This feature should not be used if you plan to use ThinLinc over low bandwidth connections such as modems or ISDN connections.

7.3.2. Serial ports (Windows and Linux only)

This feature makes it possible to export two local serial ports to the ThinLinc session. When serial port redirection is enabled, a small redirection daemon will be automatically started by the ThinLinc client during session startup. A secure tunnel for serial port data will be established.

Warning

When activating serial port redirection, all users on the terminal server can access the serial port of the client machine.

7.3.3. Drives

This feature makes it possible to, in a secure way, export one or many local drives from the client machine to the server session. This can be local hard disk volumes, local CD-ROM drives, and so on. The local drive will be made available on the ThinLinc server session.

Each exported device can have individual permission settings. All export settings are made in the ThinLinc client options interface.

7.3.4. Printer

This feature makes it possible to export a local printer to make it available from the ThinLinc session. When enabled, the client will setup a secure tunnel for printer jobs. The client will also activate a small built-in print server that listens for printer jobs on this tunnel.

When you print to the special printer queue *thinlocal* in your ThinLinc session, then the job will be sent through this tunnel and then printed on the client machine. On Linux platforms, the print job will always be sent to the default printer. On Windows and macOS, it is possible to select whether the print job should be sent to the default printer or if the printer selection dialog should be used every print. Note that device dependent print jobs will always go to the default printer.

For more information about printer redirection in ThinLinc, see Section 5.3.

7.3.5. Smart Card Readers

This feature makes it possible to export all local smart cards and smart card readers to make them available from the ThinLinc session. All smart card readers available to the system will be exported to the session so there is nothing to configure except an activation switch.

The ThinLinc client relies on the PC/SC interface present on the system to communicate with the smart card readers. If you have a reader that uses another system, then that reader will not be exported.

7.4. Client configuration

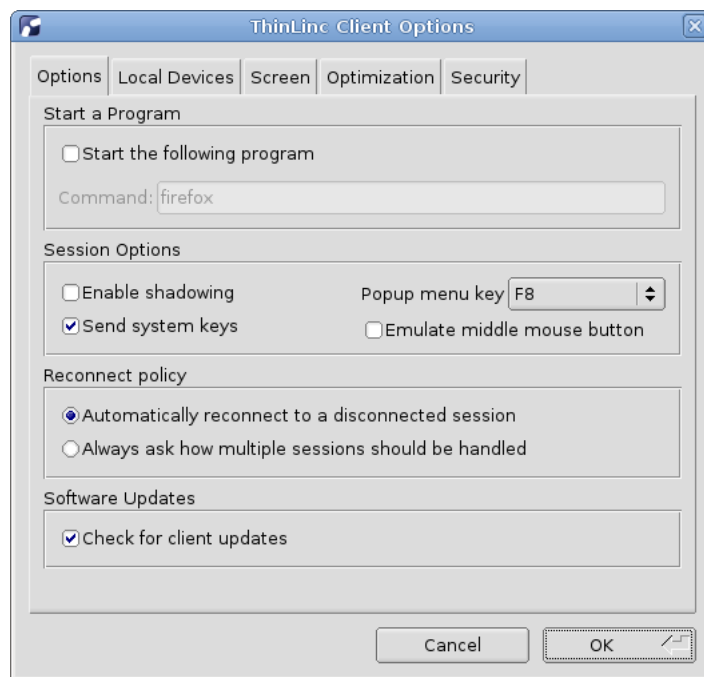
To configure the ThinLinc client you press the button labeled *"Options..."* in the client window. That brings up the client options window. This window contains several pages of settings, ordered in tab sets. The following sections will describe each of these pages and all individual settings.

When a user press the *OK button* all the current settings in the options window is saved. For more information about the config file format, see Section 7.8.

7.4.1. Options tab

The Options tab contains general options for the ThinLinc session. This includes settings for which program to execute in the session, shadowing another users session, reassignment of session pop-up key and how reconnections are handled.

Figure 7-3. Client settings Options tab



Description of options tab settings

Here follows detailed description of the settings available in the options tab.

Start a Program

If enabled, the client requests that the server should start the session with the command supplied by the client. Otherwise, the session command is determined by the server configuration.

Enable shadowing

When enabled, an extra text field will be present in the client main window. This field is used to enter the user name of the user whose session you want to shadow. For more information, see Chapter 15.

Send system keys

When this setting is enabled and the client is in full screen mode, key combinations such Alt+Tab will be sent to the remote system instead of being handled locally. To regain access to the local system without ending the session, the menu key must be used.

Emulate middle mouse button

When enabled, middle mouse button can be emulated by pressing left and right mouse button simultaneously.

Popup menu key

During a ThinLinc session you can press a specific key to bring up the session control pop-up window. This window can for example be used to toggle to and from full screen mode and to disconnect the session. The default key for this is *F8*, but other keys can be configured here. The feature can also be disabled by selecting *None*.

Reconnect policy

When the client connects to a ThinLinc server, there might already be multiple sessions running on it. Some of these sessions might be connected to another client, and some might be disconnected. The client can be configured to automatically handle some of these cases, or always ask the user what to do.

Note: Sessions that have been started with a command different from the one currently used will be ignored.

Automatically reconnect to a disconnected session

1. If there is no disconnected session and additional sessions are allowed, create a new session.
2. If there is a single disconnected session, or if server allows only one session, reconnect to existing session.
3. Otherwise, ask how to proceed.

Always ask how multiple sessions should be handled

1. If there is no running session, create a new session.
2. If server allows only one session, reconnect to existing session.
3. Otherwise, ask how to proceed.

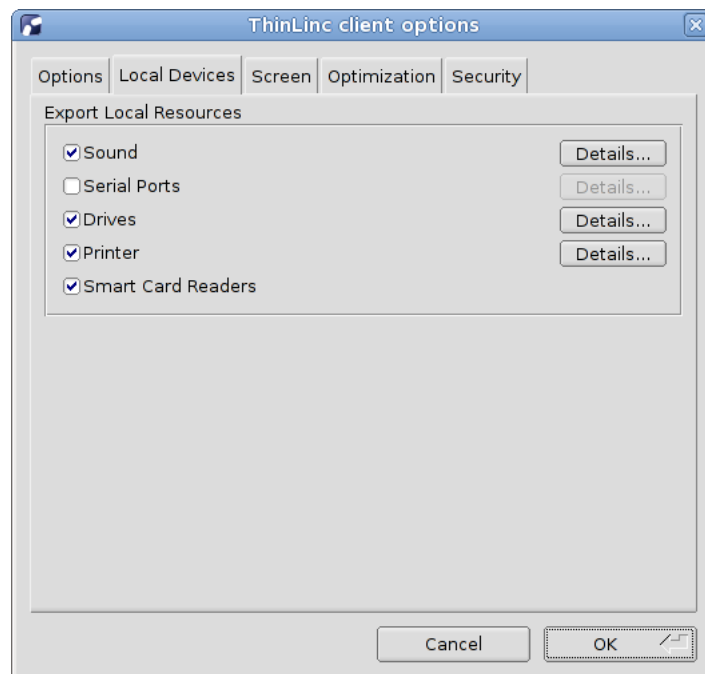
Software Updates

If enabled, the client will periodically query the `UPDATE_URL` value specified in `tlclient.conf` for updates. If a newer version is available, the user will be asked if they want to install it.

7.4.2. Local Devices tab

The Local Devices tab contains options for which local devices should be exported to the server and in what manner.

Figure 7-4. Client settings Local Devices tab



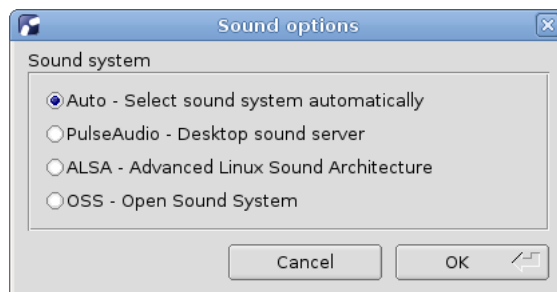
Description of local devices tab settings

Here follows detailed description of the settings available in the local devices tab.

Export - Sound Device

When enabled, sound will be sent from the ThinLinc server to your local client. A small local sound daemon will be started by the client, which connects to a secure tunnel to the server. See Section 12.3 for more information about this topic.

Figure 7-5. Sound system selection interface

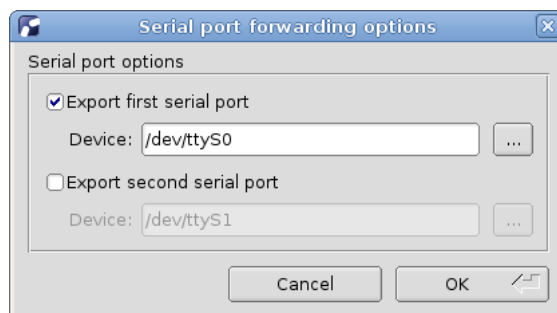


On Linux there is a "Details..." button next to the Sound check box that will allow you to choose between PulseAudio, ALSA and OSS for the local sound system. You can also let the ThinLinc client select the correct system automatically.

Export - Serial Ports

It is possible to forward two serial ports from the client to make it available to programs you run on the server. To select which of your local serial devices to export you can press the "Details..." button next to the Serial Port check box. This will bring up a dialog where you can select which two serial ports should be exported.

Figure 7-6. Serial port selection interface




The *Device* should be a path to a Linux serial device (such as `/dev/ttyS0`) or a Windows COM port name (such as `COM1`). Enter the device to export in the text field or press the "... " button to browse to the wanted device.

Export - Drives

This check box turns on export of local devices from your terminal to the ThinLinc server. This makes your local drives available from your ThinLinc session. To select which drives to export you press the "Details..." button next to Drives check box. That presents a dialog where you can build a list of drives to export and set export permissions.

Figure 7-7. Local drive export selection interface



The *Export Local Drives* window consists of two parts. At the top there is a list containing exported paths, with two control buttons below. The lower half contains settings fields for the currently selected path. When you select a path listed in the upper list you will see its corresponding settings in the Drive Options field below. You can then change the selected path by changing the values on the options field.

To add a new path to the list you press the *Add* button. That creates a new empty land in the path list. The new path will be automatically selected. you can then modify the settings in the lower half. Set the path and export permission for the new export. To set the export path you can either write it manually in the path text field or press the "... " button to bring up a file navigation window.

To remove a path you simply select a path and press the *Remove* button.

The Windows client features a mechanism that makes it easy to export the "My Documents" folder. This feature is activated by pressing the "My Documents" button. Regardless of the local folder


name, this folder will be mounted as "MyDocuments" on the server.

The export permissions can be one of the following three options, *Not Exported*, *Read Only* and *Read and Write*. The *Not Exported* option can be used to temporarily turn off an export without having to delete it. The *Read Only* option means that you from the ThinLinc session will be able to read from the export, but not write. The *Read and Write* option means that you from the ThinLinc session will be able to both read and write.

Export - Printer

By checking this check box the client will export your local printer to make it available from the ThinLinc session. For more information about this feature, see Section 7.3.4 and Section 5.3.

Figure 7-8. Printer options dialog




On Windows and macOS there is a "Details..." button next to the Printer check box that will allow you to select if the print job should be sent to the default printer or if the printer selection dialog should be used on every print.

Export - Smart Card Readers

This check box makes all local smart card readers and smart cards available to applications on the ThinLinc server. It is not necessary to check this box to authenticate using smart cards, but it is needed if you also wish to authenticate using smart cards to a Windows Remote Desktop Server.

7.4.3. Screen tab

The "Screen" tab contains options regarding the session screen. This includes initial screen size, resize behaviour and full screen mode.

Figure 7-9. Client settings Screen tab

Description of screen tab settings

Here follows detailed description of the settings available in the screen tab.

Size of session

In this box of radio buttons you can select the screen size you want for your ThinLinc session. The first five options are static with four very common screen sizes (800*600, 1024*768, 1280*1024 and 1600*1200).

The option *Current monitor* makes the ThinLinc session just as large as the monitor that the main client window is currently on. This can be used to run ThinLinc in full screen mode on one monitor, whilst retaining access to the local desktop on the other monitors.

The option *All monitors* makes the session large enough to cover all available monitors. This is a good choice when using *full screen mode*.

The options *All monitors* and *Current monitor* are identical if there is only a single monitor connected.

The option *Work area (maximized)* makes the ThinLinc session size suitable for a maximized window.

The final possible size option is to manually enter the wanted width and height. The two text boxes close to the last radio button is supposed to contain the width and the height of the wanted session as numbers. These numbers must be larger than 128 and not larger than 16384.

Resize remote session to the local window

This option makes the remote session follow the size of the local ThinLinc Client window. If the local window is resized, the remote session will be adjusted to match. If this option is disabled, or if the server is too old, padding or scroll bars will be added as needed when the remote session does not match the size of the local window.

Full screen mode

This option enables *full screen mode* during sessions. That means that the ThinLinc session will cover all of the screen area. If the session is smaller than your screen resolution, there will be black borders around your session which will be centered on the screen. If you run in full screen mode and want to reach the native session that is hidden by the ThinLinc session you can switch out from full screen mode. To do this you press the key assigned to bring up the session pop-up menu. Normally this menu is bound to the F8 key, but can be manually changed. See the *Popup menu key* setting on the *Options tab* above for more information on this. In the session menu you should select *Full screen* to toggle full screen mode.

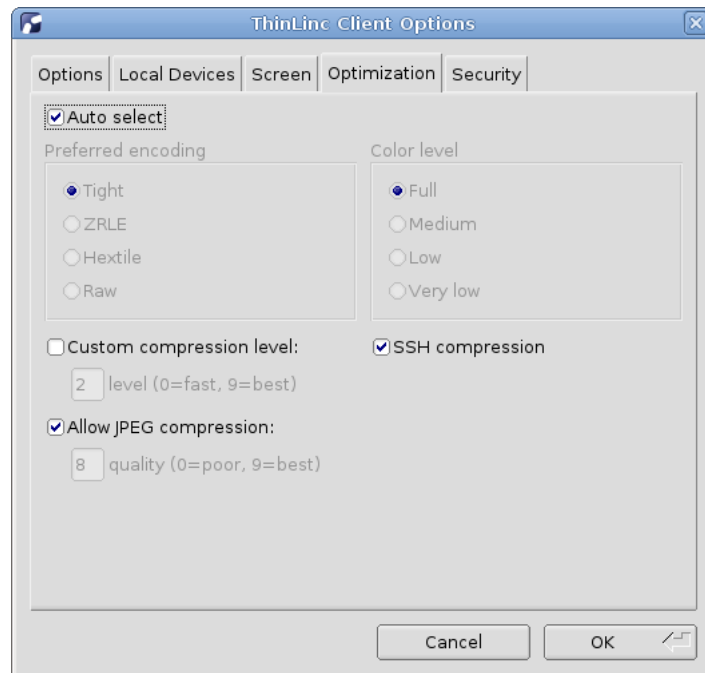
Enable full screen mode over all monitors

This option controls if *full screen mode* should use just the current monitor, or all monitors connected to the local client.

7.4.4. Optimization tab

The "Optimization" tab contains various settings that affects the protocols used to transfer the graphic information. This includes algorithm used for the graphic encoding. The best choices may differ from case to case. Factors that affects the algorithm choices can for example be network bandwidth, network latency, and client computer performance.

The default setting is to use the *Auto select* mode, to automatically select the best suited algorithms.

Figure 7-10. Client settings Optimization tab

Description of optimization tab settings

Here follows detailed description of the settings available in the optimization tab.

Auto select

This option makes the ThinLinc system try to automatically select the best suited encoding algorithm. The network performance is measured during the life of the connection and the encoding options are adjusted based on the results. This means that the encoding options can be changed automatically during the connection, if the network performance changes. Activating this option will "gray out" the *Preferred encoding* and *Color level* options, to show that they aren't manually controlled.

Preferred encoding

This block of settings affects the VNC protocol encoding. There are several different ways to compress and encode the graphic data that is sent from the server to your client. In this box you select one of four possible encodings. The methods differ much: Some try to use smart algorithms to select and compress the areas to send, which means slightly higher CPU usage, but most likely less bandwidth usage and faster sessions where the bandwidth is limited. Other methods use less CPU capacity but more network bandwidth. The best choice can vary much depending on place and

situation. A safe choice is to let the system automatically select the best encoding by checking the *Auto select* checkbox above.

Encoding: Tight

This choice selects the Tight encoding method. With this encoding the *zlib* compression library is used to compress the pixel data. It pre-processes the data to maximize compression ratios, and to minimize CPU usage on compression. Also, JPEG compression may be used to encode color-rich screen areas. The *zlib* compression level and the JPEG compression ratio can be manually changed. See *Custom compression level* and *Allow JPEG compression* below. *Tight encoding is usually the best choice for low-bandwidth network environments (e.g. slow modem connections).*

Encoding: ZRLE

This choice selects the ZRLE encoding method.

Encoding: Hextile

This choice selects the Hextile encoding method. With Hextile the screen is divided into rectangles, split up in to tiles of 16x16 pixels and sent in a predetermined order. *Hextile encoding is often the best choice for using in high-speed network environments (e.g. Ethernet local-area networks).*

Encoding: Raw

This choice selects the Raw encoding method. This is the simplest of the encoding methods. It simply sends all the graphic data of the screen, raw and uncompressed. *Since this method use the least processing power among the possible methods this is normally the best choice if the server and client runs on the same machine.*

Custom compression level

By selecting this option you choose to override the standard compression level used when compressing data with the Tight encoding. You can manually select the wanted compression level by entering a number between 0 and 9. Level 0 means no compression. Level 1 uses a minimum of CPU performance and achieves weak compression ratios, while level 9 offers best compression but is slow in terms of CPU consumption on the server side. Use high levels with very slow network connections, and low levels when working over high-speed network connections. *This applies to the Tight encoding only!*

Allow JPEG compression

By selecting this option you choose to override the standard JPEG compression quality of color-rich parts of the screen. JPEG is a "lossy" compression method for images that helps the Tight encoding to significantly reduce the size of the image data. The drawback is that the resulting image, depending of selected compression ratio, can be blurred and grainy. You can manually select the wanted image quality by entering a number between 0 and 9. Quality level 0 gives bad image quality but very impressive compression ratios, while level 9 offers very good image quality at lower compression ratios. Note that the Tight encoder uses JPEG to encode only those screen areas that look suitable for lossy compression, so quality level 0 does not always mean unacceptable image quality.

Color level

This block of choices selects the number of colors to be used for the graphic data sent from the server to the client. The setting has four levels, *Full*, *Medium*, *Low* and *Very low*. The default and normal is to use the *Full* setting. Selecting a lower number of colors will highly affect the resulting image to the worse, but may also speed up the transfer significantly when using slow network connections.

In this context, *Full* means the number of colors supported by the clients graphics hardware.

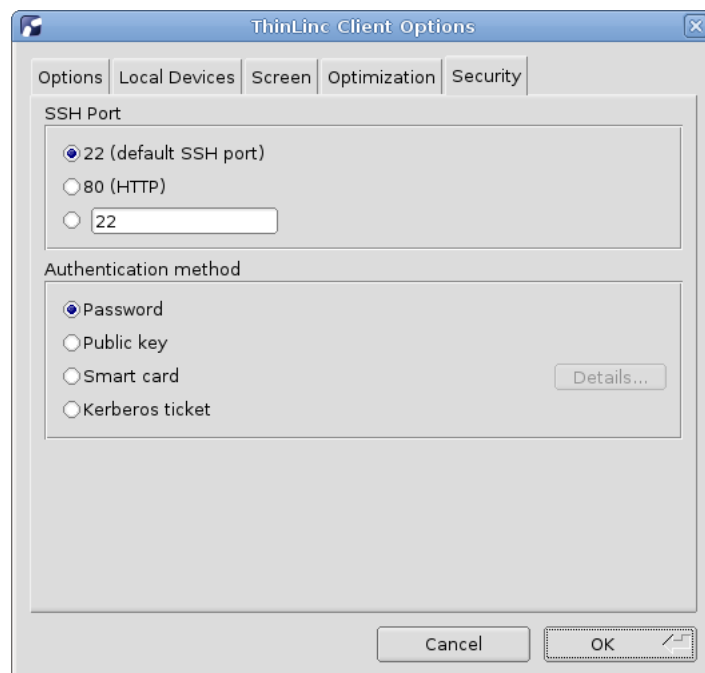
SSH compression

This choice selects whether or not to use SSH compression for all the data sent between ThinLinc server and client. This is normally not used since an extra compression step, above a compressing graphic encoding normally doesn't help making it smaller, only use more CPU performance. There can still be occasions where this is worth trying though. It is possible that this can help speeding up printing or other exports over slow connections.

7.4.5. Security tab

The "Security" tab controls how the client authenticates against the ThinLinc server. The main interface of the client will be different depending on the choices made here.

Figure 7-11. Client settings Security tab



Description of security tab settings

Here follows detailed description of the settings available in the security tab.

SSH Port

This option selects the TCP/IP port to use when the client tries to establish an SSH connection with the ThinLinc server. The normal SSH port is 22, which also is the default selection for this option. There can be reasons to use another port on some occasions. If you for example need to use ThinLinc over the Internet, from a location where port 22 is blocked by a firewall. Then you can select a port that is let open. Port 80 which is used for HTTP, the protocol used for transport when surfing the WWW is one port that often is open. To be able to use a port other than 22 you need to make sure that the SSH daemon (sshd), which runs on the ThinLinc server, listens to the port you want to use. The SSH daemon can be told to listen to any wanted ports. In the client interface you can select between the default port 22, port 80 and an arbitrary port number which you can enter by yourself.

Note about SSH host key updates: If the SSH host key on the server changes, e.g. due to an upgrade of the OS or SSH server software, the client will note this fact. It will then, at the next login, open a dialog and let the user confirm that the new host key is valid. If the user clicks OK , then the host key on the client for this particular server is updated on disk.

The administrator can disallow this by manually setting the parameter `ALLOW_HOSTKEY_UPDATE` to 0. See Section 7.8 for more information.

Password

This option makes the client try to authenticate using a regular password.

Public key


This option makes the client try to authenticate using public key encryption. The user will be asked to provide a private encryption key instead of a text password.

Smart card

This option makes the client try to authenticate using public key encryption, but with the private key securely stored on a smart card. The user will be asked to select a certificate on the smart card and to provide the passphrase for it.

Note: Smart card authentication requires that the smart card is readable by your PKCS#11 library. The library included by default supports PKCS#15 compliant smart cards and relies on the PC/SC interface. This is always present on Windows systems and is usually installed by default on Linux systems.

The "Details..." button lets you change the options for smart card usage and managing the certificate filters which are used to match accepted certificates for authentication.

Figure 7-12. Smart card authentication settings**Use certificate subject as login name**

Enable this options if you want to enable automatic login, this will also hide the input box for login name from user.

Connect when smart card is inserted

This options will try to automatic connect and is dependent on certificate filters, automatic connect will only occur if only one certificate is available after the filtration.

Read more about automatic connection below where certificate filters is discussed.

See Section 10.4.6 for information on how to configure the server for automatic smart card connection.

Disconnect when smart card is removed

Enabling this options makes the client automatically disconnect when the smart card used to authenticate is removed.

Send smart card passphrase (PIN) to server

This option makes the client transmit the smart card passphrase, as entered by the user, over to the ThinLinc server. It is required to enable smart card single sign-on.

Warning

Enabling this option reduces the security of the smart card as the passphrase would otherwise never leave the client system. The option should be left disabled if smart card single sign-on is not used.

Smart card - certificate filter

A certificate filter is used to present only allowed certificates for authentication, certificates that does not match any filter will be hidden from the user.

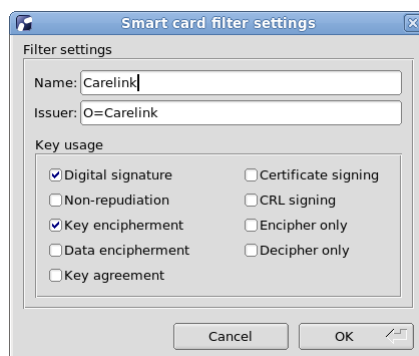
When no certificate filters are configured, all available certificates on the smart card will be available for authentication and therefore the autoconnect feature will not work.

If the resulting filtered list of certificate evaluates only one certificate for authentication and the autoconnect feature is enabled, it will be used for authentication.

When the login dialog is displayed and the key shortcut **control-shift-F8** is pressed, the certificate filtering functionality is bypassed and gives you access to all certificates available on the smart card for authentication.

To add a new filter just press the add button as shown in dialog Figure 7-12 or select an available filter item in the list and press edit to change the settings for specific filter. Either way, the certificate filter settings dialog Figure 7-13 will be shown where you can modify the settings of the specific filter.

Figure 7-13. Certificate filter settings



Name

Enter name of the filter that will be seen in the list of filters.

Issuer

The certificate issuer field consists of a comma separated list of attribute-value pairs that all must be present in the certificate issuer field. Commonly the "common name" of the issuer is used, e.g. "cn=My CA". It is also possible to allow any issuer that are part of the same organisation, e.g. "o=My Company Ltd.". Any registered object identifier descriptor can be used as an attribute name (see IANA (<https://www.iana.org/assignments/ldap-parameters/ldap-parameters.xhtml>) for a full list).

Key usage

The certificate must have all the key usage bits selected in this window. Having more bits than those selected does not exclude a certificate.

Kerberos ticket

This option makes the client try to authenticate using an existing kerberos ticket.

Note: This requires that a valid kerberos ticket is available on the client, and that the SSH daemon on the ThinLinc server is configured to accept this ticket during authentication. For information about how to configure kerberos authentication on your particular platform(s), please see the relevant vendor documentation.

7.5. Client Touch Gestures

The ThinLinc client has support for a number of touch gestures when used on a touch capable monitor. These gestures allow the user to simulate certain mouse operations that would otherwise not be possible.

Note: Touch gestures are not available on macOS as it currently lacks native support for touch capable monitors.



Click

Tap a single finger to simulate a click of the left mouse button.



Drag

Press a single finger and drag it across the screen to simulate pressing the left mouse button and moving the cursor.



Right click

Press and hold a single finger to simulate a press of the right mouse button.



Right click, alternative

Tap two fingers to simulate a press of the right mouse button.



Middle click (not available on Windows)

Tap three fingers to simulate a press of the middle mouse button (the mouse wheel).



Pan / Scroll

Press two fingers and drag them across the screen to simulate rotating the vertical or horizontal mouse wheel.



Zoom

Press two fingers and move them close or further away from each other to simulate pressing the Control key and rotating the mouse wheel. Many applications interpret this combination as a request to zoom the open document in or out.

7.6. The XDM mode (Linux only)

When installing dedicated clients, for example old PCs or thin terminal boxes, it's common to install the client to run in XDM mode. XDM is an acronym for X Display Manager and is the name of a small graphical program used for graphical logins in many Linux systems. By using the ThinLinc client in XDM mode you can make sure that the client appears automatically when the client hardware is started and that it reappears directly after a user logs out.

To run the client in XDM mode you need to start it with the `-x` option. When running in XDM mode the following changes will be made to the client interface.

- The Exit button is removed.
- Mouse acceleration and keyboard repetition settings are loaded from the client configuration file and applied at startup.
- Always use fullscreen
- More limited F8 menu

7.7. Logfile placement

The ThinLinc client logs its activities to the human readable file, `tlclient.log`. The locations of this file differs between Linux and Windows systems and are explained below.

The log file is truncated every time the ThinLinc client is started. This means that the log file always contains information about the latest session only.

7.7.1. Linux log file

On Linux systems the logfile is located in the home directory for the user that runs the ThinLinc client. The path is `~/.thinlinc/tlclient.log`.

7.7.2. Windows log file

On Windows systems the logfile is located in the directory referenced from the `%TMP%` variable. The value of this variable can be achieved by running any of the following commands in a Windows command window.

```
C:\> echo %TMP%
```

or

```
C:\> set
```

Observe that some directories in the Windows `%TMP%` path may be flagged as hidden by the Windows system. This means that you need to change directory options to display hidden files and directories to navigate to the log file.

7.7.3. macOS log file

On macOS systems the logfile is located in the home directory for the user that runs the ThinLinc client. The path is `~/.thinlinc/tlclient.log`.

7.8. Client configuration storage

7.8.1. Overview and Parameters

The ThinLinc client does not use Hiveconf for its configuration. Instead, the Linux and macOS clients use a plain text format with key/value pairs and the Windows client stores the values in the Windows registry.

Note: The configuration parameters should seldom be edited by hand. For a system wide configuration, create a parameter set using the client and copy it to the system wide file.

Configuration Parameters Used by the ThinLinc Client

Both the Windows and the Linux version of the ThinLinc client use the same names for their configuration parameters, although the storage technique used is different (text files vs. registry keys). In this section we will list the parameters and explain their possible values.

ALLOW_HOSTKEY_UPDATE

Set to 1 if SSH host key updates should be allowed. This parameter cannot be changed from the GUI. The result of setting `ALLOW_HOSTKEY_UPDATE` to 0 is that the client cannot connect to the server if the host key is wrong. This enhances security if there is a risk for a man in the middle attack.

AUTHENTICATION_METHOD

This parameter can be set to "password", "publickey", "scpupublickey" or "kerberos" to select the authentication mode used by the client.

AUTOLOGIN

If this parameter is set to 1, the client will automatically login at start, using the server name, user name and password specified in the configuration storage.

CERTIFICATE

Specifies the smart card certificate to use when authenticating.

CERTIFICATE_NAMING

Controls how the client presents a certificate to the user. The parameter consists of a comma separated list of naming tokens that represent bits of information from each card or certificate. Possible tokens:

`card_label`

The label specified on the smart card.

`pin_label`

The label associated with the PIN protecting this certificate.

`subject_*`

A field from the subject in the certificate. Can for example be the common name by specifying *subject_cn* or *subject_commonName*. Any registered object identifier descriptor can be used (see IANA (<https://www.iana.org/assignments/ldap-parameters/ldap-parameters.xhtml>) for a full list).

`issuer_*`

A field from the issuer in the certificate, in the same manner as for *subject_**.

The client will use as many of the tokens as necessary to give each certificate a unique name. That means that certificates on two different cards can be presented with a different number of tokens depending on how much the information between the certificates overlap. An index number will be added to the name if the names are still not unique when all tokens are used.

CUSTOM_COMPRESSION

Set to 1 if a custom compression method is selected.

CUSTOM_COMPRESSION_LEVEL

The selected compression level. An integer between 1 and 9.

DISPLAY_MODE

The display mode. Can be set to values "SIMPLE" and "ADVANCED", or be left empty. In the latter case, the default behaviour is to use simple mode if a server name is given as a parameter and advanced mode otherwise.

EMULATE_MIDDLE_BUTTON

Set to 1 if the client should emulate middle mouse button when pressing left and right mouse buttons simultaneously.

FULL_SCREEN_MODE

Set to 1 if the client should run in fullscreen mode.

FULL_SCREEN_ALL_MONITORS

Set to 1 if the client should use all monitors in full screen mode, instead of just the current monitor.

HOST_ALIASES

This parameter specifies a list of hostname and port translations. This translation list is consulted whenever the client is about to initiate a network connection. This includes the SSH connection to the ThinLinc agent machine. The syntax for this parameter is:

```
[fromhost][:fromport]=[tohost][:toport] ...
```

If `fromhost` is omitted, the translation will apply to all hosts. The same principle is used for ports.

If `tohost` or `toport` is omitted, the original host or port will be used. Multiple translations are separated with whitespace. The translation stops as soon as one match is found.

JPEG_COMPRESSION

Set to 1 if JPEG compression is wanted.

JPEG_COMPRESSION_LEVEL

The wanted compression level.

KILL_EXISTING_SESSIONS

Set to 1 if existing sessions should be ended.

Note: It makes little sense to change this value. The client never saves this setting.

LOGIN_NAME

The username.

LOWERCASE_LOGIN_NAME

Set to 1 if the client should convert the entered username to lowercase before logging into the server. This affects both the login user name and the name of the user to shadow (if applicable).

NEW_PASSWORD_REGEX

This parameter specifies a regular expression. If an interactive SSH prompt matches this expression, the response is taken as a new password. The new password will be used for the SSH connection to the agent machine. It will also be sent to the server to enable Single Sign-On.

NFS_EXPORTS

A list of local drive paths and permissions. The syntax for this parameter is:

[path1], [permissions1], [path2], [permissions2] ...

As seen above, each path should be followed by the desired permissions *disabled* (not exported), *ro* (read only) or *rw* (read and write). See Section 7.4.1 for their meaning. This list specifies local drives to be exported.

NFS_ROOT_WARNING

Set to 1 to enable a warning if running as root and exporting local drives.

NFS_SERVER_ENABLED

Set to 1 if local drives should be exported.

OPTIONS_POPUP_KEY

Key code for key to activate option pop-up menu.

PASSWORD

This parameter allows you to specify a password in the configuration file. It must be specified using a hexadecimal ASCII notation, which means that every character is specified by its hexadecimal value.

Warning

The password value is not encrypted. It should be treated as a clear text password. Avoid storing configuration files with a PASSWORD parameter on disk or transmit such files over networks without encryption.

PKCS11_MODULE

Specifies the PKCS#11 module that will be used to communicate with the smart card. The path can be relative the base prefix of the ThinLinc client or an absolute path.

PRINTER_ENABLED

Set to 1 if local printers should be enabled.

PRINTER_SELECTION

Set to 1 if the local printer selection dialog should be displayed on every print on Windows and macOS clients. Otherwise printing jobs will be sent to the default local printer.

PRIVATE_KEY

This parameter specifies the path to the private key to be used to authenticate the user.

RECONNECT_POLICY

This parameter can be set to "single-disconnected" or "ask" to control the client's reconnect policy. See Section 7.4.1 for their meaning.

REMOTE_RESIZE

Set to 1 if the client should resize the remote session when the local window changes.

REMOVE_CONFIGURATION

If 1, the user configuration file (or the file specified by -C) will be removed after the client has started. Settings changed in the GUI will not be stored to disk. If the client fails to remove the file, it will try to truncate it instead.

SCREEN_SIZE_SELECTION

The default size of the ThinLinc session. Possible values:

- 0 for 640x480
- 1 for 800x600
- 2 for 1024x768
- 3 for 1280x1024
- 4 for 1600x1200
- 5 for Current monitor
- 6 for Work area (maximized)
- 7 for Custom screen size, set using the SCREEN_X_SIZE and SCREEN_Y_SIZE parameters.
- 8 for All available monitors

SCREEN_X_SIZE

Custom width of session, if SCREEN_SIZE_SELECTION is set to 7.

SCREEN_Y_SIZE

Custom height of session, if SCREEN_SIZE_SELECTION is set to 7.

SEND_SYSKEYS

Set to 1 if the client should send system keys (like Alt+Tab) to the remote system when in full screen mode.

SERIAL1_DAEMON_DEVICE

The path to the first local serial port device to be exported.

SERIAL1_PORT_ENABLED

Set to 1 if the first local serial port should be exported.

SERIAL2_DAEMON_DEVICE

The path to the second local serial port device to be exported.

SERIAL2_PORT_ENABLED

Set to 1 if the second local serial port should be exported.

SERIAL_PORTS_ENABLED

Set to 1 if local serial ports should be exported.

SERVER_NAME

The hostname or IP of the ThinLinc server. When using ThinLinc in a cluster setup this should be the hostname or IP of the Master server machine.

SHADOWING_ENABLED

Set to 1 if shadowing should be enabled.

SHADOW_NAME

The username of the user who's session should be shadowed.

SMARTCARD_AUTOCONNECT

Set to 1 if the client should automatically attempt a connection when a smart card with a suitable certificate is found, this will only work if SMARTCARD_SUBJECT_AS_NAME also is set to 1.

SMARTCARD_DISCONNECT

Set to 1 if the client should disconnect automatically when the smart card used for authentication is removed.

SMARTCARD_EXPORT_ENABLED

Set to 1 if local smartcard readers should be exported.

SMARTCARD_FILTER_n

This is a item list of certificate filters replace *n* with a sequence number that defined the order of the filter in the list.

The filter string consists of three fields where each field is sperated using a | (pipe), the defined three fields are: name, attributes and key usage which are documented below. Here follows an example of a filter string showing its format:

SMARTCARD_FILTER_1=Telia|o=TeliaSonera|5

name

The name of the filter which will be displayed in the list of filters defined in the user interface.

attributes

This field holds a comma separated list of certificate attributes that is used when matching against available certificates, for example *O=TeliaSonera*.

key usage

Key usage is a bitmask value used to match against a certificate's key usage flags. It indicates the intended usage of the certificate, such as identification, signing etc.

Use this to match certificates that is intended to be used for logon. For example, identification certificates will be matched using a value of 5, *digital signature + key encipherment = 5*. The values are described in the following table:

1	digital signature
2	non-repudiation
4	key encipherment
8	data encipherment
16	key agreement
32	certificate signing
64	CRL signing
128	encipher only
256	decipher only

SMARTCARD_PASSPHRASE_SSO

Set to 1 if the client should transmit the smart card passphrase to the ThinLinc server to enable smart card single sign-on. See Section 7.4.5 for security implications.

SMARTCARD_SUBJECT_AS_NAME

Set to 1 if the certificate subject should be used as logon name, this will hide the name field from login window.

SOUND_ENABLED

Set to 1 if sound redirection should be enabled.

SOUND_SYSTEM

Which local sound system to use. Only used on platforms that have multiple sound systems to choose from. Possible values:

AUTO

Automatically choose the most appropriate sound system of those available.

PULSE

Use the local PulseAudio server as determined by X11 properties or environment variables.

ALSA

Use the default ALSA device.

OSS

Use the default OSS device.

SSH_ARBITRARY

Custom port number for ThinLinc connection.

SSH_COMPRESSION

Set to 1 to use the compression built into SSH.

SSH_PORT_SELECTION

Port selection for ThinLinc connection. Possible values:

- 0 for port 22 (standard ssh port).
- 1 for port 80.
- 2 for custom port set in the SSH_ARBITRARY parameter.

START_PROGRAM_COMMAND

Specifies the command to use when starting the session, if START_PROGRAM_ENABLED is active.

START_PROGRAM_ENABLED

Specifies if the client should request that the server starts the session with the command supplied by the client.

UPDATE_ENABLED

Set to 1 to enable periodic checks for new versions.

UPDATE_INTERVAL

This parameter specifies the time interval, in seconds, between client update checks.

UPDATE_LASTCHECK

This parameter specifies the time that the last update check was performed.

UPDATE_MANDATORY

If set to 1, updating to new client versions is mandatory.

UPDATE_URL

The HTTP URL to client update configuration file.

VNC_AUTOSELECT

Set to 1 to dynamically autoselect the compression algorithm during the session.

VNC_COLOR_LEVEL

The color level used for the session.

VNC_ENCODING_SELECTION

The encoding to use for VNC. Possible values:

- 0 for Raw
- 5 for Hextile
- 7 for Tight
- 16 for ZRLE

YESNO_PROMPT_REGEX

This parameter specifies a regular expression. If an interactive SSH prompt matches this expression, a graphical yes/no dialog will be presented, instead of a dialog for text input. Additionally, if the prompt is known to the client, an alternate text will be used. The dialog buttons Yes and No will send "yes" and "no" to the server, respectively.

7.8.2. Configuration Parameter Storage

Configuration parameters are typically stored in text based configuration files. The format is simple: Each parameter is written on one line, followed by an equal sign (=) and the value of the parameter, as in the following example:

```
SOUND_ENABLED = 0
SERVER_NAME = tl.example.com
```

By using the -C option, additional configuration files can be specified. Any name is accepted, but the file extension `.tlclient` is recommended. The Windows, Linux, and macOS packages configures the system to automatically recognize such files as configuration files for the ThinLinc Client. Additionally, the Internet Media Type "application-vnd.cendio.thinlinc.clientconf" is linked to such configuration files.

7.8.2.1. Linux Client Configuration Files

The Linux client first reads the file `/opt/thinlinc/etc/tlclient.conf`, if it exists. It then reads the file `.thinlinc/tlclient.conf` in the user's home directory, and the values there override the values from `/opt/thinlinc/etc/tlclient.conf`. This way, a system administrator can set global defaults for client operations, while each user can still customize the client to wanted behavior.

7.8.2.2. macOS Client Configuration Files

The macOS client first reads the file `/Library/Application Support/ThinLinc Client/tlclient.conf` if it exists. It then reads the `.thinlinc/tlclient.conf` in the user's home directory, and the values there overrides the values from `/Library/Application Support/ThinLinc Client/tlclient.conf`. This way, a system administrator can set global defaults for client operations, while each user can still customize the client to wanted behavior.

7.8.2.3. Windows Client Configuration

On Windows, the ThinLinc client reads its configuration from the registry. All ThinLinc client data is stored under `Software\Cendio\ThinLinc\tlclient` in the HKLM and HKCU hives. The parameter names are the same as for the Linux client.

The behaviour of global and user-specific settings are identical to that of the Linux client, where settings in HKLM correspond to `/opt/thinlinc/etc/tlclient.conf` and those in HKCU correspond to `.thinlinc/tlclient.conf`.

7.8.3. Adding Custom Branding to the ThinLinc Client Login Window

It is possible to add a custom logo to the main ThinLinc client window, making it easily distinguishable from a generic client. The custom logo will be placed to the right of the input fields.

Adding the logo is easy. The new logo must be a PNG file with maximum width and height of 50 pixels. On Windows, just add the file `branding.png` in the same directory as the executable with the custom logo. On Linux, the file name is `/opt/thinlinc/lib/tlclient/branding.png`.

7.9. Client Customizer

7.9.1. Introduction

This software lets you create customized ThinLinc client installation programs. This means that when users install the customized version, they will automatically get the default settings you have configured.

One advantage with this is that you can provide, for example, a default server name. A custom client can also be used to enhance security: You can distribute SSH host keys with the client itself, so that users don't need to be concerned with SSH host key fingerprint verification.

Note: The Client Customizer only works for the Windows client.

7.9.2. Installation

Before you can start, you have to install the build environment. This is done by running the command **tl-4.12.1-client-customizer.exe** located in the Client Bundle. This will also create a shortcut to the build directory in the Start menu.

7.9.3. Building a Customized Client

To create a customized client, do the following:

1. Edit `settings.reg`. This file contains all the parameter names and default values that are used in **tlclient**. To customize the client, edit any of these values, and they will be installed in the registry when you install the customized client itself. You can also add your servers SSH host keys (see below).
2. Custom branding can be added to the client by simply dropping a file called `branding.png` in the root directory of the Client Customizer. See Section 7.8.3 for more details.
3. Run **build.bat** in the same directory. The file `setup.exe` will now be created. This is the installation program for the customized client.

7.9.4. Adding SSH Host Keys to `settings.reg`

To add your servers SSH host key to `settings.reg`, do the following:

1. Use **tlclient** to connect to your server, if you haven't already done so. Confirm the servers host key, if necessary.
2. Run the registry editor, and select
`HKEY_CURRENT_USER\Software\Cendio\ThinLinc\tlclient\KnownHosts`
3. Export this key to an external text file.
4. Open the text file from the previous step in an editor.
5. Copy the line corresponding to your ThinLinc server. Paste this line into `settings.reg`, section
`HKEY_LOCAL_MACHINE\Software\Cendio\ThinLinc\tlclient\KnownHosts`
6. Save `settings.reg`, and proceed to create the customized client as described above.

7.10. Launching the Client from a Web Page

This feature allows a web server, such as an intranet or a web portal, to initiate a ThinLinc Client connection with a given configuration on behalf of a user.

7.10.1. Requirements

Web Integration requires an Apache HTTP Server, configured for TLS, with the ability to run CGI scripts.

Note: Some Linux distributions distributes their Apache HTTP server with the `mod_cgi` module disabled. This module is required for Web Integration to work. On Ubuntu systems, it can be enabled by running the **a2enmod cgi** command and restarting the `httpd` service.

Note: If Web Integration is used over HTTP, an attacker with access to the network may be able to intercept user passwords. To protect from this happening, Web Integration automatically redirects to a HTTPS connection when HTTP is used.

7.10.2. Installation

The Web Integration feature is not enabled by default in a ThinLinc installation. An installation script, `/opt/thinlinc/share/web_integration/install-web-integration`, is provided for ease of installation.

Example 7-1. Installing Web Integration configuration

```
# /opt/thinlinc/share/web_integration/install-web-integration
```

Note: After installing the Apache HTTP configuration file, make sure to restart the `httpd` service to load the new configuration.

While the installation script works as-is on most supported platforms, two environment variables grants you more control over where the configuration file is installed. This can be useful if you have an `httpd` installation at a custom location.

APACHE_CONF_DIR

If your Apache HTTP server has been installed to a non-standard location, set this environment variable to tell the installation script where the configuration directory is located.

If this parameter is unset, the installation script will attempt to find the configuration directory from a list of known locations.

Example 7-2. Installing Web Integration configuration to a custom `httpd` directory

```
# env APACHE_CONF_DIR=/usr/local/etc/httpd/conf.d \
/opt/thinlinc/share/web_integration/install-web-integration
```

APACHE_CONF_NAME

The default behavior of the installation script is to install the configuration file in the configuration directory with the name `thinlinc.conf`. If you already have a file with that name in the configuration directory that you wish to keep, set this environment variable to an different name.

Example 7-3. Installing Web Integration configuration with a custom file name

```
# env APACHE_CONF_NAME=web-integration.conf \
/opt/thinlinc/share/web_integration/install-web-integration
```

7.10.3. Usage

The process works like this:

1. The CGI script is called with the desired parameters.
2. The CGI script generates a "launch file", which is a normal client configuration file. When the browser receives this file, it launches the locally installed ThinLinc client.

The launch file delivered to the client is generated from the template `/opt/thinlinc/etc/tlclient.conf.webtemplate`. The CGI script performs some substitutions on this file, before sending it to the client. Currently, the following variables are substituted:

`$server_name$`

The server name where the CGI script resides.

`$login_name$`

The user name, specified by the `username` CGI parameter.

`$password$`

The password in hexadecimal ASCII notation, specified by the `password` or `hexpassword` CGI parameters.

`$autologin$`

The value of the `autologin` CGI parameter.

7.10.4. The CGI Script `tlclient.cgi`

The CGI script `tlclient.cgi` is used to start the native client, when launched from a web page. It accepts many parameters which affects its operation. These are described below:

`server_name`

The desired server name.

`username`

The desired user name. No default.

`password`

The desired password, in plain text. No default.

`hexpassword`

The desired password, in hexadecimal ASCII notation. This parameter overrides the `password` parameter. No default.

`redirto`

After launching the native client, the browser will redirect to the web page specified by this parameter. Default value: the empty string.

`loginsubmit`

This boolean parameter specifies if a login should be directly executed, instead of showing a login form. Default value: 0

`autologin`

This boolean parameter specifies if the native client should automatically connect to the specified server at startup. Default value: 1

`start_program_enabled`

This boolean parameter specifies if the native client should request that the server starts the session with the command supplied by the client, as indicated by the `start_program_command` parameter. Default value: 0.

`start_program_command`

This parameter specifies the command to use when starting the session. Default value: "tl-single-app firefox".

`displayurl`

This boolean parameter can be used for debugging and development purposes. It will display and URL with all submitted parameters, and do nothing else. Default value: 0

`shadowing_enabled`

This boolean parameter specifies if the native client should activate shadowing. Default value: 0

`shadow_name`

This parameter specifies the user to shadow. Default value is the empty string.

To make it easier to test various parameters, the HTML file `cgitest.html` is included, in the same location as `tlclient.cgi`. It also demonstrates how to create icons on web pages, which launches ThinLinc sessions.

7.11. Advanced Topics

7.11.1. Hardware Address Reporting

When the client connects to server, it reports its hardware address. On Linux, the active interface with the smallest MAC address is used. On Windows, the address of the first interface (as listed in the Control Panel) is used.

7.11.2. Client Update Notifications

The client includes a feature which can periodically check for new versions. This functionality is affected by the configuration parameters `UPDATE_ENABLED`, `UPDATE_INTERVAL`, `UPDATE_LASTCHECK`, `UPDATE_MANDATORY`, and, `UPDATE_URL`. These are described in Section 7.8. During an update check, the client retrieves the file specified by `UPDATE_URL`. An example follows:

```
WINDOWSINSTALLER = https://www.cendio.com/downloads/clients/tl-latest-client-windows.exe
LINUXINSTALLER = https://www.cendio.com/downloads/clients/thinlinc-client-latest.i686.rpm
DEFAULTINSTALLER = https://www.cendio.com/thinlinc/download
OKVERSIONS = 3.2.0 3.3.0
```

The `OKVERSIONS` parameter specifies a list of valid client versions. If the running client version is different, the client will notify the user. The `WINDOWSINSTALLER`, `LINUXINSTALLER`, and `DEFAULTINSTALLER` parameters specifies the updated client packages for Windows, Linux, and other platforms, respectively.

Chapter 8. Client Platforms

There are several ways to run the ThinLinc client, and also some ways to access ThinLinc servers without running the client.

In this chapter we will document how to install, configure and run the ThinLinc client on all different platforms including dedicated thin terminals.

8.1. Windows

8.1.1. Requirements

The supported Windows versions are 7, 2008 R2, 8, 2012, 8.1, 2012 R2, 10 and 2016. Windows CE is currently not supported.

8.1.2. Installing the Windows Client

To install the client on a Windows machine, unpack the Client Bundle and enter the `client-windows` directory. Then click on the file `tl-4.12.1-client-windows.exe` and follow the instructions.

You will also find unpacked versions of the ThinLinc Windows client, both 32bit and 64bit under `tl-4.12.1-client-windows-x86` and `tl-4.12.1-client-windows-x64` directories. It makes it possible to run the client directly from the bundle or other media, like a portable application, without the requirement of installing the client.

For more information about how to configure the client, read Section 7.8.

8.1.3. Running the Windows Client

During installation the ThinLinc client will be added to the Start menu. To start the client you select it from the Start menu.

8.2. macOS

8.2.1. Requirements

- macOS (formerly OS X) version newer than 10.6 running on 64-bit Intel hardware

Note: macOS (formerly OS X) versions newer than 10.9 installs with a default setting that breaks the multi monitor functionality of the ThinLinc client. A workaround to this problem is to disable setting "Displays have separate Spaces" in settings for "Mission Control" found in "System Preferences" .

8.2.2. Installing the macOS Client

The client for macOS can be found in the directory `client-macos` in the Client Bundle. To install the client, follow these steps:

1. Double-click on the file `tl-4.12.1_6733-client-macos.iso`.
2. Drag the "ThinLinc Client" application to an application folder of your choice.
3. Eject "ThinLinc Client".

8.2.3. Running the macOS Client

To start the ThinLinc Client, double click on the client application. The client can also be added to and started from the Dock.

8.2.4. Command and Alt Keys on macOS

The Alt key (also known as the Option key) behaves very differently on macOS compared to its behaviour on other platforms. It closely resembles the PC AltGr key, found on international keyboards. ThinLinc therefore treats these keys in a special manner on macOS in order to provide a good integration between the client and the remote ThinLinc system.

Whenever either of the Alt keys are pressed, ThinLinc will behave as if AltGr was pressed. The left Command key is used as a replacement for Alt in order to use shortcuts like Alt+F. The right Command key retains its behaviour of acting like the Super/Windows key.

8.3. Linux PC

8.3.1. Requirements

- A compatible CPU architecture:
 - An i686 (or compatible) CPU with MMX and SSE support
 - An x86_64 (or compatible) CPU
 - An ARMv7 (or compatible) CPU with Thumb-2 and VFP3D16
- GLIBC 2.12, or newer
- A working Fontconfig configuration, or basic fonts available in `/usr/share/fonts` or `/usr/X11R6/lib/X11/fonts`.
- 32 MiB RAM

8.3.2. Installing the Linux Client

The Linux client is distributed in three different kinds of packages. One that can be installed using the RPM package system, one in the DEB package format, and one in compressed tar archive form for any Linux distribution.

If you need more information than mentioned here, read Section 7.8.

In the instructions below, we will assume that you have unpacked your Client Bundle to `~/tl-4.12.1-clients`.

8.3.2.1. RPM-based Installation on RPM-based distributions

The RPM-based client can be found in the directory `client-linux-rpm` in the Client Bundle. It is suitable for systems such as Red Hat, Fedora, SuSE, and Mandrake. Perform the following steps to install it on a 32-bit system:

```
$ cd ~/tl-4.12.1-clients/client-linux-rpm
$ sudo rpm -Uvh thinlinc-client-4.12.1-6733.i686.rpm
```

or the following steps on a 64-bit system:

```
$ cd ~/tl-4.12.1-clients/client-linux-rpm
$ sudo rpm -Uvh thinlinc-client-4.12.1-6733.x86_64.rpm
```

or the following steps on a 32-bit ARM hard-float system:

```
$ cd ~/tl-4.12.1-clients/client-linux-rpm
$ sudo rpm -Uvh thinlinc-client-4.12.1-6733.armv7hl.rpm
```

8.3.2.2. DEB-based Installation on Debian and Ubuntu based distributions

The DEB-based client can be found in the directory `client-linux-deb` in the Client Bundle. It is suitable for systems such as Debian and Ubuntu. Perform the following step to install it on a 32-bit system:

```
$ cd ~/tl-4.12.1-clients/client-linux-deb
$ sudo dpkg -i thinlinc-client_4.12.1-6733_i386.deb
```

or the following steps on a 64-bit system:

```
$ cd ~/tl-4.12.1-clients/client-linux-deb
$ sudo dpkg -i thinlinc-client_4.12.1-6733_amd64.deb
```

or the following steps on a 32-bit ARM hard-float system:

```
$ cd ~/tl-4.12.1-clients/client-linux-deb
$ sudo dpkg -i thinlinc-client_4.12.1-6733_armhf.deb
```

8.3.2.3. Installation on Other Linux Distributions

A client without any specific package format can be found in the directory `client-linux-dynamic` in the Client Bundle. It is possible to run this client from any directory, even from the unpacked Client Bundle. We generally recommend installing it in `/opt/thinlinc`. Perform the following steps to install the client to `/opt/thinlinc` on a 32-bit system:

```
$ cd ~/tl-4.12.1-clients/client-linux-dynamic
$ sudo mkdir -p /opt/thinlinc
$ sudo cp -a tl-4.12.1-6733-client-linux-dynamic-i686/* /opt/thinlinc/
```

or the following steps on a 64-bit system:

```
$ cd ~/tl-4.12.1-clients/client-linux-dynamic
$ sudo mkdir -p /opt/thinlinc
$ sudo cp -a tl-4.12.1-6733-client-linux-dynamic-x86_64/* /opt/thinlinc/
```

or the following steps on a 32-bit ARM hard-float system:

```
$ cd ~/tl-4.12.1-clients/client-linux-dynamic
$ sudo mkdir -p /opt/thinlinc
$ sudo cp -a tl-4.12.1-6733-client-linux-dynamic-armhf/* /opt/thinlinc/
```

The client is also available as tar archives for easy transfer to other systems without having to copy the entire Client Bundle.

8.3.3. Running the Linux Client

On Linux systems the client will be installed as `/opt/thinlinc/bin/tlclient`. The client package contains settings that adds `/opt/thinlinc/bin` to the path of the users.

To run the client, click on the "ThinLinc Client" icon in your desktop environment. Typically, the icon is found in the Internet category. You can also run the client by executing `/opt/thinlinc/bin/tlclient`.

8.4. Thin Terminals

ThinLinc has support for several thin terminals, i.e. hardware built with the task of providing a thin client as primary design goal.

8.4.1. eLux-based Thin Terminals (Fujitsu Futro et. al.)

ThinLinc has support for running the client on eLux-based thin terminals. This includes the Fujitsu Futro as well as NEXTerminal terminals such as the NEXceed. eLux is a modern embedded operating system which works well.

ThinLinc supports models that uses the eLux RP 5.x operating system. Microphone devices are known to work at least on the Futro models. However, make sure to disable "Sound in XDMCP sessions".

8.4.1.1. Installing/Upgrading the ThinLinc Client on eLux

Below we will describe how to install and configure the ThinLinc client on terminals running eLux. For details on how to use the administrative tools (Scout and ELIAS) as well as on how to manually configure terminals without use of Scout, please refer to the eLux documentation available at <http://www.mylux.com>.

1. Create a new firmware image including the ThinLinc client. This is done by copying the files `thinlinc-client-*` from the appropriate container directory in the Client Bundle to your ELIAS/Scout container directory.

Note: Some eLux distributions already include the ThinLinc client, rendering this step unnecessary.

2. Add the ThinLinc client to your firmware image using ELIAS.
3. Install the image on your clients using either Scout (for centrally-managed clients), or via http (for single clients).
4. Configure your clients by adding an application of type *Custom* under the **Local** tab under **Configuration** on the client (or create a new application if using Scout). Set the application name to "ThinLinc" and the command line to `tlclient`. Check the "Application restart" and the "Start automatically after 0 s" checkboxes to make sure the ThinLinc client is restarted after end of session, and that it is started automatically at boot.

Note: All commandline parameters accepted by the Native Linux client is also accepted by the client run on eLux (it's the same client). This means that server name, username, lockdown options etc. can be used on the commandline specified when configuring the client. An example commandline instructing the client to start with an empty username and to connect to `<server name>` is provided below.

```
tlclient -u " <server name>
```

This configuration can be done either on the client, under the configuration tab in the starter, or centrally, using Scout, by adding an application in the Applications container in an appropriate organizational unit.

5. Configure the starter not to start automatically by unchecking the checkbox in "Autostart Starter after 0 s" under **Setup Desktop Advanced**. Unchecking the "Enable" button for the Taskbar in the same location may also be a good idea to give users a cleaner environment.

This configuration can be done either on the client, under the Setup tab in the starter, or centrally, using Scout.

6. If custom configuration of the client, for example to support local drives, is needed, transfer a custom client configuration file (`tlclient.conf`) to `/setup/thinlinc/tlclient.conf` on the eLux clients using the "Transfer files" functionality available in the properties of devices or organisation units in Scout (under the Files tab). The files are transferred when the client is restarted.

8.4.2. HP ThinPro Terminals

HP ThinPro terminals are based on Ubuntu, and therefore one can use the DEB package provided in our ThinLinc Client bundle for this terminal.

8.4.2.1. Manual installation/upgrade of ThinLinc Client

Below we will describe the process of manually installing the ThinLinc Client on Ubuntu based HP ThinPro Linux terminals.

1. Use the tool *"Administrator/User mode switch"* to authenticate as administrator.

2. Start an X terminal from the advanced tab in the control panel.

3. Unlock the filesystem:

```
# fsunlock
```

4. Copy the ThinLinc Client .deb package from ThinLinc Client bundle onto a USB memory stick, and connect it to the terminal. Go into the directory which represents your connected USB device with command:

```
# cd /tmp/tmpfs/media/my_usb_storage
```

As an alternative, it is also possible to download the client package from a web server using the "wget" command.

5. Install the ThinLinc Client package using Debian package manager command:

```
# dpkg -i thinlinc-client*.deb
```

6. Lock down the filesystem before closing the X terminal window:

```
# fslock
```

7. Reboot.

8. Add a ThinLinc connection in the connection manager.

The HP "Connection Wizard" does not include an entry for ThinLinc. Press "Skip", then add a ThinLinc Connection in the "Connection Manager".

The default user and administrator share the same home directory, and it is therefore important to NOT start the ThinLinc Client as administrator the first time. This will make the ThinLinc Client configuration only accessible by administrator and not the default user.

On "zero" clients, the default server name is set when the ThinLinc connection type is selected. To change server name, temporarily switch to another connection type, then switch back to ThinLinc. Also, to configure the ThinLinc Client, enter an invalid username/password combination in the HP login dialog. Acknowledge the error. It is then possible to access the full ThinLinc Client interface.

8.4.3. IGEL Universal Desktop

A client package for IGEL Universal Desktop terminals is provided. It is included in the directory `client-igel` in the Client Bundle. IGEL Universal Desktop is a modern embedded operating system which works well. Some editions includes a bundled ThinLinc client. We do not recommend this client. Instead, install the client as described below.

Note: Installation of our client package is only possible on IGEL terminals with the "Custom Partition" feature. Please ask your IGEL representative for more information.

8.4.3.1. Installing/Upgrading the ThinLinc Client on IGEL terminals

Below we will describe how to install and configure the ThinLinc client on IGEL terminals, using the "Custom partition". You can use either the Universal Management Suite software running on a separate workstation, or the setup software installed on the terminal. You will need access to a web server which allows you to publish the client files.

1. Edit the configuration of the terminal. Select System, Firmware Customization, Custom Partition.
2. Under the Partition option, make sure that "Enable Partition" is checked. Enter a size, such as "100M". The partition must be at least 25 MiB. The upper limit depends on the hardware used. Make sure that the mount point is /custom.
3. Under the Download option, press the star to create a new data source. Enter the URL to the web server where the ThinLinc client package definition is located. Example:
http://www.example.com/client-igel/thinlinc-i686.inf
4. Under Custom Application, press the star to create a new application entry. Use a Session Name such as "ThinLinc".
5. Click on Settings. Enter the Icon name:

`/custom/thinlinc/icon.png`

To setup the client to use the terminals normal language, enter this Command:

`/custom/thinlinc/bin/tlclient`

To setup the client to use Swedish, use this Command:

`env LC_ALL=sv_SE.UTF-8 /custom/thinlinc/bin/tlclient`

6. Press OK to save the configuration.

8.4.4. Other Thin Terminals

The ThinLinc client can be made to run on almost any Linux-based Thin Terminal as well as on some Windows-based appliances. Contact Cendio if you need help on a consultancy basis.

8.5. Running ThinLinc on a Thinstation terminal

The Thinstation project (<http://thinstation.github.io/thinstation/>) is an opensource thin client Linux distribution that can be booted in many different ways, including entirely over the network on diskless machines and via a LiveCD.

A client package for ThinStation is shipped as part of the ThinLinc client distribution. In this section, we will document how to use and configure this package with Thinstation.

Note: To run the ThinLinc client in ThinStation, you need ThinStation version 2.2 or later. However, on ThinStation version 2.4 and later, the ThinLinc client will be downloaded automatically during the build process. Thus the download procedure is only needed for ThinStation 2.2. However, the configuration step is needed for both 2.2 and 2.4.

There are two alternative methods of getting a Thinstation image with the ThinLinc client included. The first one is to use one of the TS-O-Matic servers available. They allow you to build Thinstation images online, and should generally have the ThinLinc client available as an option. The TS-O-Matic servers are available from the Thinstation webpages (<http://thinstation.github.io/thinstation/>). The second is to download the Thinstation distribution, add the ThinLinc client package and then configure and build a Thinstation image manually. This requires access to a Linux box.

Below, we will describe the second method

8.5.1. Installing and Building the Package

Begin by downloading and unpacking the Thinstation main distribution available from the Thinstation webpages (<http://thinstation.github.io/thinstation/>).

Enter the Thinstation directory created while unpacking, and unpack the ThinLinc Thinstation client package, found in the `client-thinstation` directory in the Client Bundle, into this directory:

```
$ tar zxvf tl-4.12.1-6733-client-thinstation.tar.gz
```

This will unpack files into the `packages/thinlinc` directory, as well as the file `README.thinlinc` which contains some summarized information on the package.

Edit the `build.conf` and add a line 'package thinlinc' in the Applications section.

Run the `build` script and wait for its completion.

If everything went well, there will now be Thinstation images available in the `boot-images` directory. Use the appropriate boot image for your preferred boot method.

8.5.2. Configuring the ThinLinc client when running on a Thinstation Terminal

When running on a network-booted Thinstation terminal, the client is configured by adding statements to the configuration file that is downloaded at boot by Thinstation. The default name of this file is `thinstation.conf.network`, located in your tftpboot. There can also be other filenames that configures specific terminals based on their IP or hardware (MAC) addresses.

8.5.2.1. Basic Configuration

For the ThinLinc client to appear at all, a Thinstation "session" must be created. This is done by adding a few lines to the `thinstation.conf.network` file. Here's an example:

```
SESSION_0_TYPE=thinlinc
SESSION_0_THINLINC_SERVER=tl.example.com
SESSION_0_THINLINC_OPTIONS="-u johndoe"
```

```
SESSION_0_THINLINC_CONFIG_NFS_SERVER_ENABLED=0
```

The above example will make the ThinLinc appear on the display of the client after boot. It will set the servername to *tl.example.com*, and it will reset the username field. It will also disable export of local drives. See below for information on enabling local drives on Thinstation.

All standard client options can be added to the `SESSION_0_THINLINC_OPTIONS` variable. For example, to lock down the server field, add *-l server*.

8.5.2.2. Configuration using the Client Configuration File

Some of the features of the ThinLinc client can't be configured via command line options. Instead, the configuration file must be altered. To allow features such as local drive and sound redirection to work when running on Thinstation, the ThinLinc client package for Thinstation has features for altering the configuration file on the client.

To alter the configuration file, add statements on the form

SESSION_0_THINLINC_CONFIG_<configuration file variable name> = <value> to `thinstation.conf.network`. An example follows:

```
SESSION_0_THINLINC_CONFIG_NFS_SERVER_ENABLED=1
SESSION_0_THINLINC_CONFIG_SOUND_ENABLED=1
```

The above example will set the `NFS_SERVER_ENABLED` to *1* and the `CONFIG_SOUND_ENABLED` to *1*, and so on.

8.5.2.3. Enabling Sound and Local Drive Redirection

If the hardware running Thinstation has support for it and the correct sound and disk device modules has been loaded, the ThinLinc client will be able to support sound and local drive redirection. The following configuration lines in `thinstation.conf.network` will enable sound redirection and local drive redirection for USB storage devices:

```
SESSION_0_THINLINC_CONFIG_NFS_EXPORTS=/mnt/usbdevice,rw,/mnt/cdrom,ro
SESSION_0_THINLINC_CONFIG_NFS_SERVER_ENABLED=1
SESSION_0_THINLINC_CONFIG_SOUND_ENABLED=1
SESSION_0_THINLINC_CONFIG_NFS_ROOT_WARNING=0
```

8.5.2.4. Avoiding Question about Server Host Key

When running on a device with non-volatile storage, such as a hard disk, the ThinLinc client stores the public part of the ssh host key of the ThinLinc client the first time it connects to the server after asking the user to verify the fingerprint of the key. At subsequent connects, this copy is used to verify that the client is connecting to the correct server.

When running on a diskless Thinstation host, the key can be stored only in volatile memory (on a RAM disk), so the client will ask the user to verify the fingerprint once each time the client has been rebooted. Since its normal behaviour to reboot a Thinstation terminal once a day, this will lead to a confusing situation for users, not to mention that it will decrease security.

To solve this problem, the ThinLinc client package for Thinstation tries to download a file name `ssh_known_hosts` from the tftpboot. If it exists it will be used as database of known hostkeys on the client.

To create this file, login with the client to the ThinLinc server, using the same servername as the one that will be configured on the clients. Then copy the file `~/.thinlinc/known_hosts` to `<tftpboot>/ssh_known_hosts`.

Chapter 9. ThinLinc Web Access

9.1. Overview

ThinLinc Web Access is a ThinLinc client that runs in modern browsers and allows access to a ThinLinc server or cluster without installing any extra software on a client device.

9.2. Requirements

ThinLinc Web Access requires a web browser that supports modern web technologies such as WebSockets and Canvas. It is verified to work correctly on the latest versions of the major web browsers:

- Internet Explorer
- Microsoft Edge
- Firefox
- Google Chrome
- Safari

9.3. Server Configuration

ThinLinc Web Access is served by the service `tlwebaccess`. The default TCP port number for this HTTP service is 300. It can be changed to some other port such as 443, assuming this port is free. The configured port must be allowed in any firewalls.

In a cluster setup, the `tlwebaccess` service must run on all machines. The same service port should be used, and all machines must be accessible from the clients.

The setting `/webaccess/login_page` will also need to be configured in a cluster setup. The client first authenticates with the master. Once the master server has chosen an agent server for the session, the client will authenticate with that agent server. The browser will thus present pages from two different servers. First a page from the master, and then from the agent, unless the agent is on the same server of course. This parameter is a means for the agent to know the public hostname of the master server. Thus when it's properly set, the user can, when the session has ended, click a button to return from the agent to the master to login again. The default value, which is `/`, will not redirect to another server and is only usable if you are running a stand alone ThinLinc server, i.e. not a cluster. If you have more than one server or are using Network Address Translation (NAT), you must set this parameter to an address that all clients can connect to. Example:

```
login_page = https://thinlinc-master.example.com:300/
```

Please see Section 14.2.8 for details on all possible settings for ThinLinc Web Access.

9.3.1. Certificates

For best security and user experience, we strongly recommend that you use valid TLS certificates. The certificates should match the server host names. For correct behavior, you should set the parameter `/vsmagent/agent_hostname` on each of the agents in the ThinLinc cluster.

If you can't obtain a valid TLS certificate but still want to test ThinLinc Web Access you can use a self-signed certificate. Such a certificate, created for "localhost", is bundled with Web Access. Any use of self-signed certificates is insecure and most browsers will display warnings when they are used. Self signed certificates must be manually approved.

Note: In Safari, the certificates MUST match the server hostname, while other browsers might be content with a warning. Firstly, this means that you cannot connect through an IP address. Secondly, you can not use the bundled self-signed certificate. You can create a new self-signed certificate using our shipped helper script `make-dummy-cert`. OpenSSL is required to be installed for this script. Use it like this:

```
$ sudo /opt/thinlinc/etc/tlwebaccess/make-dummy-cert `hostname --fqdn`
```

Manually approving the self-signed certificate requires some additional steps in Safari compared to other browsers. On macOS the user must expand the browser dialog that complains about the certificate and choose to always accept that certificate. If the user already dismissed that dialog, then Safari has to be restarted. A self signed certificate must be manually approved for all machines in a cluster.

If you must test a browser on iOS with a self-signed certificate you have to use Safari to manually type the hostname and port of your server and add `/server.crt` to the URL in order to download the server certificate (e.g. `https://thinlinc-master.example.com:300/server.crt`). After downloading the certificate, on iOS version 12 and later you have to install it in "Settings/General/Profile". Then you also have to enable the full trust of that root certificate in the "Certificate Trust Settings" which can be found at the bottom of the "Settings/General/About" page. See Apple's instructions here (<https://support.apple.com/HT204477>). After using Safari to install the certificate, you can use Web Access in any browser on iOS.

Warning

The above steps for iOS are very insecure and are not recommended for production systems. iOS does not have a mechanism for ignoring bad certificates for a single site. This means that following the method above will result in that your device considers the certificate as a generally trusted authority. This can in turn allow whoever has access to that certificate's private key to generate a certificate that falsely appears valid for any site. For example, an evil website could appear to have a valid certificate for your bank.

9.4. Usage

ThinLinc Web Access is accessed with your web browser by browsing to the master machine, for example `https://thinlinc-master.example.com:300`. If you have configured the service to run on port 443, ":300" can be omitted.

Note: On iOS and Android devices, you can add an icon to the home screen. When the ThinLinc Web Access is launched from the home screen, it will run in full screen mode.

9.4.1. Logging in to a ThinLinc server

The first thing presented to the user when browsing to ThinLinc Web Access is a login form where the user's username and password can be specified.

Figure 9-1. ThinLinc Web Access Login



The login form features the Cendio ThinLinc logo at the top. Below the logo are two input fields: 'Username:' and 'Password:'. A 'Login' button is positioned to the right of the password field. At the bottom, the text 'Version 4.11.0post (build 6416) on localhost' and 'Copyright © Cendio AB 2019' are displayed.

Cendio®
ThinLinc®

Username:

Password:

Login

Version 4.11.0post (build 6416) on localhost
Copyright © [Cendio AB](#) 2019

To login into a ThinLinc server Web Access needs to do a successful user authentication. For most systems the password will be sufficient. If more information is needed, e.g. when using One Time Passwords or when a password change is needed, then Web Access will present a series of prompts for the user until the user has been fully authenticated.

If the login attempt is successful a ThinLinc session will start or an old one will be reused, depending on if the user already has a session running or not.

Note: ThinLinc Web Access does not fully support multiple sessions for a the same user. If the user has multiple session then a random session will be chosen

9.4.2. The Toolbar

Once connected ThinLinc Web Access will display a toolbar on one side of the browser window for various functions. This toolbar can be hidden by clicking the small tab on the side of it. Clicking the tab again will make the toolbar reappear. The toolbar can also be moved to either side by grabbing the tab and dragging it to the desired side.

Figure 9-2. ThinLinc Web Access Toolbar



The ThinLinc Web Access toolbar has the following functions:

Move/Drag Viewport

Toggle between sending mouse events to the ThinLinc session or panning a session that is larger than the current browser window. This button will only be shown on devices that do not have visible scrollbars.

Show Keyboard

Toggle the on screen keyboard for the device. This button will only be shown if a touch device has been detected.

Show Extra Keys

Displays a secondary toolbar with virtual keys for devices with limited or no physical keyboard. See Section 9.4.3 for details.

Clipboard

Opens the clipboard dialog. See Section 9.4.4 for details.

Disconnect

Disconnects ThinLinc Web Access from the current session.

9.4.3. Extra Keys

Some physical keyboards and most on screen keyboards lack a number of keys that are commonly used in applications and desktop environments. To simplify use of these an extra toolbar is available that can simulate these keys.

Figure 9-3. ThinLinc Web Access Extra Keys



Control

Simulates pressing or releasing the left Control key.

Alt

Simulates pressing or releasing the left Alt key.

Windows

Simulates pressing or releasing the left Windows key.

Tab

Simulates pressing and releasing the Tab key in sequence.

Escape

Simulates pressing and releasing the Escape key in sequence.

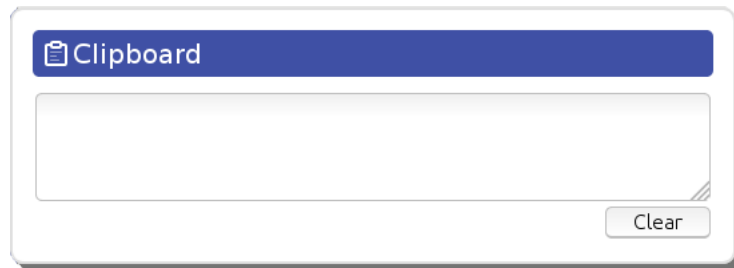
Ctrl+Alt+Delete

Simulates pressing and releasing the Control, Alt and Delete keys in sequence.

9.4.4. Clipboard

For security reasons the browsers prevent ThinLinc Web Access from directly integrating with the local clipboard. Copying text to or from the ThinLinc session must therefore be handled manually via the clipboard dialog.

Figure 9-4. ThinLinc Web Access Clipboard Dialog



The contents of the clipboard dialog will automatically be updated whenever the contents of the clipboard in the ThinLinc session changes. Correspondingly, if the contents of the clipboard dialog is changed by the user then the clipboard in the session will be updated to match.

9.4.5. Touch Gestures

ThinLinc Web Access has support for the same touch gestures as the ThinLinc client when used on a touch capable monitor. These gestures allow the user to simulate certain mouse operations that would otherwise not be possible. Please see Section 7.5 for details on what gestures are available.

9.4.6. Command and Alt Keys on macOS and iOS

The Alt key (also know as the Option key) behaves very differently on macOS and iOS compared to its behaviour on other platforms. It closely resembles the PC AltGr key, found on international keyboards. ThinLinc therefore treats these keys in a special manner on macOS and iOS in order to provide a good integration between the client and the remote ThinLinc system.

For more information on how ThinLinc treats these keys see Section 8.2.4.

Chapter 10. Authentication in ThinLinc

In this chapter we will describe how authentication of users is performed in ThinLinc

10.1. Pluggable Authentication Modules

Authentication of users in ThinLinc is performed using the *Pluggable Authentication Modules* (PAM). This means ThinLinc can authenticate users using any system for which there is a PAM module. Examples of PAM modules are *pam_ldap* for accessing LDAP directories (including Novell NDS/eDirectory) and *pam_winbind* for authenticating against a Windows Domain. Of course, authentication using the standard plaintext password files of Linux is also possible using the PAM module *pam_unix*.

10.1.1. Configuration files for PAM

PAM is configured by editing the files located in the directory `/etc/pam.d/`.

Different Linux distributions have slightly different ways of configuring PAM. The ThinLinc installation program will setup ThinLinc to authenticate using the same PAM setup as the Secure Shell Daemon, by creating a symbolic link from `/etc/pam.d/thinlinc` to either `/etc/pam.d/sshd` or `/etc/pam.d/ssh`, depending on which of the latter files that exists at installation. This seems to work on most distributions. Be aware that the PAM settings for the Secure Shell Daemon might really be somewhere else. For example, on Red Hat distributions, the file `/etc/pam.d/system-auth` is included by all other pam-files, so in most cases, that is the file that should be modified instead of the file used by sshd.

10.2. Limitations

Some PAM modules and authentication mechanisms are case sensitive, while others are not. Usernames in the ThinLinc client are case sensitive by default, however this behaviour can be changed. See `LOWERCASE_LOGIN_NAME` in Section 7.8 for details.

The SSH server should be configured to allow "keyboard-interactive" authentication for full functionality. The "password" authentication method does not allow multiple interactive prompts which is required for things such as password changes during login.

10.3. Using Public Key Authentication

10.3.1. Introduction

Public key authentication is a more secure alternative to passwords. It uses a challenge/response mechanism that prevents even the server from gaining access to the authentication information. This section will describe how to configure ThinLinc to use it.

10.3.2. Key Generation

In order to use public key authentication, a pair of encryption keys must be generated. Tools to accomplish this should be included with the SSH server. On Linux, that server is normally OpenSSH and the tool to generate keys is called **ssh-keygen**.

Example:

```
# ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key (/home/johndoe/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/johndoe/.ssh/id_rsa.
Your public key has been saved in /home/johndoe/.ssh/id_rsa.pub.
The key fingerprint is:
e1:87:0d:82:71:df:e9:4a:b0:a8:e3:cd:e8:79:58:32 johndoe@example.com
```

Remember that the private key (`id_rsa` in the example) is a password equivalent and should be handled with care. The public key (`id_rsa.pub` in the example) does not need to be kept secret.

Note: The SSH key format is not standardised, so it may be required to convert the keys depending on which servers you will be using.

10.3.3. Server Configuration

All SSH servers must support public key authentication, so any SSH server will work. It is important, however, to verify that public key authentication is not disabled. Refer to the documentation for your SSH server for instructions on how to do this.

Next, the public keys need to be associated with the correct users. For OpenSSH, you must store a copy of the public key in the users' home directory, specifically in the file `~/.ssh/authorized_keys`. This file can contain multiple keys, each on a separate line.

10.3.4. Client Configuration

The client must have a copy of the private key associated with the public key stored on the server. The key can be stored anywhere, although on Linux it is commonly stored as `~/.ssh/id_rsa`. The user will be able to specify where the key is located in the ThinLinc Client interface.

Note: The client currently only supports the OpenSSH key format. If your keys are in another format, e.g. for PuTTY, then they must first be converted before they can be used with ThinLinc.

10.4. Using Smart Card Public Key Authentication

10.4.1. Introduction

Smart card public key authentication is an advanced version of the method described in Section 10.3. It uses the same basic principle but stores the private key on a smart card, where it can never be extracted. This section will describe how to configure ThinLinc to use it.

10.4.2. General Requirements

- Smart cards with an appropriate PKCS#11 library. The library included with ThinLinc requires PKCS#15 compliant smart cards and PC/SC libraries on the client system.

10.4.3. Key Generation

The keys on the smart card are generated when the smart card is issued. How this is done is not covered by this guide.

10.4.4. Server Configuration

To use a smart card with ThinLinc, the public key must be extracted off the card and associated with a user on the ThinLinc server. The method for doing this depends on your smart card and your SSH server.

On Linux, with the OpenSSH server and an PKCS#15 compliant smart card, the tool **pkcs15-tool** (part of the OpenSC suite) is able to extract the public key.

The first step is identifying the certificate on the card:

```
# pkcs15-tool --list-certificates
X.509 Certificate [identification]
    Flags      : 0
    Authority: no
    Path       : 3f0050154331
    ID         : 45
```

The second step is to extract the key, based on the ID number:

```
# pkcs15-tool --read-ssh-key 45
1024 65537 918282501237151981353694684191630174855276113858858644490084487922635
27407657612671471887563630990149686313179737831148878256058532261207121307761545
37226554073750496652425001832055579758510787971892507619849564722087378266977930
98757520821634533353538210518946058646748977963861144645730357512544251473818679
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQGCCxIx/xtVoDR2qwY4Pym7F6yKmdJsB26MUbbTiGT7o
6M6G4A215Go1kdQRNjOWDJE9HZWToaApSkVprNeiQLeOkbELz2yND2Te/Oyl3u44YeIUImT1v4t7q9jC
MTpfZ+TpxLf0sd3DAk2So8EBAAtUkhib/ugKqfTCqB7WN0Hf0Nw==
```

The second line, starting with "ssh-rsa", is the one needed for SSH version 2 authentication. For instructions on how to associate this key with a user, see Section 10.3.

10.4.5. Client Configuration

The ThinLinc client requires no special configuration to use the smart card.

10.4.6. Automatic Connection

The client is able to automatically connect to the server when a smart card is inserted (see Section 7.4.5). It does, however, require that the user is able to log in using the subject name on the card. As that is rarely a valid user name, ThinLinc ships with a special NSS module, called *nss-passwdaliases*, that enables alternate names for users.

The module is configured by editing the file `/etc/passwdaliases`. The file is a colon-delimited table of alternate names and their corresponding user ids. Example:

```
givenname=John,sn=Doe,c=us:572
```

To activate the *nss-passwdaliases* module, it must be added to the list of NSS modules for the *passwd* database. This is specified in the file `/etc/nsswitch.conf`. For example, replace the following line:

```
passwd: files ldap
```

with this line:

```
passwd: files ldap passwdaliases
```

10.4.7. LDAP Automatic Update (tl-ldap-certalias)

ThinLinc includes the tool **tl-ldap-certalias** that can automatically update the local databases needed for smart card public key authentication, provided the system uses the OpenSSH server (or any SSH server that uses a compatible format and location for authorized public keys) and standards compliant LDAP servers where users and certificates are stored.

The **tl-ldap-certalias** command can also perform validation of certificates it finds in LDAP databases. Read more about this in Section 10.4.7.3.

- On invocation, a list of all users and matching certificates is gathered. How is determined by the `certificate_user_match` configuration variable. If `allow_invalid_certificates` is no, only matching valid certificates will be gathered.
- The user's home directory, as well as the `.ssh` directory, are created if they are required and do not already exist. **tl-ldap-certalias** reuses the `/vsmagent/make_homedir_mode` configuration variable from `vsmagent` for determining the default permissions of newly created home directories.

- Any old public keys added by **tl-ldap-certalias** are removed from the `authorized_keys` file and the keys from the current set of certificates are added.
- The file `/etc/passwdaliases` is updated with a list of subject names and user id:s, to allow for login without usernames. See Section 10.4.6 for more information.

Note: It should be noted that any custom entries in `~/.ssh/authorized_keys` will be retained, but custom changes to `/etc/passwdaliases` will be overwritten each time **tl-ldap-certalias** is run.

After deployment, **tl-ldap-certalias** is meant to be run from cron at regular intervals, for example every 15 minutes. This makes sure that the ThinLinc system automatically keeps all user certificates updated. However, please note that if you're using certificate validation, downloading and parsing certificate revocation lists may take a long time (up to 5 minutes each). This is mitigated by caching the data from the CRL:s, but the first run and whenever the CRL needs to be updated may take a long time. Thus, if you have certificates from a lot of different certificate authorities, don't run **tl-ldap-certalias** too often.

Since the default use of this tool is to be run from cron, the default behaviour is to produce no output other than error messages. If you want more output from **tl-ldap-certalias**, see options in Section 10.4.7.1.

Note: The root user must be able to write to the users' home directories for **tl-ldap-certalias** to be able to update the `~/.ssh/authorized_keys` files.

10.4.7.1. Command line options

tl-ldap-certalias accepts a number of different command line options that affects how the program interacts with its environment.

`-v`

`--verbose`

Turn on program status output to standard output. This is off by default.

`-d`

`--debug`

Turn on extra debugging output to standard output. This is off by default.

`-s`

`--simulate`

Dry run mode. Specifying this option tells **tl-ldap-certalias** to avoid writing any changes to disk. This is off by default.

`-h`

`--help`

Show usage information and exit.

10.4.7.2. Configuration

tl-ldap-certalias uses the `/utils/tl-ldap-certalias` hiveconf folder for configuration purposes. On a standard ThinLinc installation, it's located in `/opt/thinlinc/etc/conf.d/tl-ldap-certalias.hconf`.

Configuration parameters

`/utils/tl-ldap-certalias/ldap_schema`

Specify the schema type that is used on the target LDAP server. Valid options are `rfc2307` and `AD`. `rfc2307` is default and should be used with standard LDAP servers that complies to `rfc2307`. `AD` should be used if target LDAP server is an Active Directory.

`/utils/tl-ldap-certalias/allow_invalid_certificates`

This parameter controls whether to perform validation on certificates found in the LDAP database. Possible values are `yes` and `no`.

Please note that by setting this to `yes`, you will allow users with expired, invalid, revoked or untrusted certificates to log in as if their certificates are valid.

Note: If you want **tl-ldap-certalias** to match the behaviour of **tl-ldap-certalias** versions earlier than 3.2.0, set this to `yes`.

`/utils/tl-ldap-certalias/certificate_user_match`

The method to use for finding certificates assigned to the user. Valid options are `sameobject` and `novell_certificate_subjectname`.

`sameobject` makes **tl-ldap-certalias** search for certificates in the `userCertificate` attribute on user objects it finds.

`novell_certificate_subjectname` allows for certificates to be stored in another LDAP tree. User objects are associated to certificates by storing subject names of certificates in a multivalued attribute called `sasAllowableSubjectName` (OID 2.16.840.1.113719.1.39.42.1.0.69) on the user object. **tl-ldap-certalias** can handle both DN:s as written by Novell iManager (`DC=com.DC=example.OU=test.CN=foo`) and as returned by **tl-certtool --subject** (`cn=foo,ou=test,dc=example,dc=com`).

`/utils/tl-ldap-certalias/users/uri`

A LDAP server URI for finding users on the form `ldap [s] :// name [:port]`

`/utils/tl-ldap-certalias/users/base`

The LDAP search base for finding users.

`/utils/tl-ldap-certalias/users/binddn`

The username to bind as for searching for users. If left blank, no bind is performed.


```
/utils/tl-ldap-certalias/users/bindpw
```

The password to use in combination with binddn for bind operations. If binddn is left empty, this can also be left empty.

```
/utils/tl-ldap-certalias/certs/uri
/utils/tl-ldap-certalias/certs/base
/utils/tl-ldap-certalias/certs/binddn
/utils/tl-ldap-certalias/certs/bindpw
```

If certificate_user_match is not sameobject, these settings will be used to determine where to look for certificates. They follow the same rules as the settings for users.

10.4.7.3. Certificate validation

tl-ldap-certalias can perform validation of certificates found in LDAP databases by the following methods if allow_invalid_certificates is set to yes:

Certificate validity and expiry dates

tl-ldap-certalias now checks the certificate validity and expiry dates and rejects certificates that are not valid yet or have expired.

Matching certificate to certificate issuers

Place the CA certificates you wish to trust certificates from in `/opt/thinlinc/etc/ca/`. The CA certificates must be in DER form. If **tl-ldap-certalias** finds a certificate with an issuer that does not match any of the certificates in `/opt/thinlinc/etc/ca/`, the certificate will be considered invalid and ignored.

Certificate revocation lists

tl-ldap-certalias searches the certificates it encounter for certificate revocation lists (CRL), to make sure that the certificate has not been revoked by its issuer. Once downloaded, the CRL will be cached until the time for the next scheduled update found in the CRL list has passed.

Note: **tl-ldap-certalias** can only handle CRL lists distributed with HTTP.

Validation of certificate signatures.

tl-ldap-certalias can verify that the certificate signature is valid and thus assures that the certificate has not been tampered with.

10.5. Using One Time Passwords

10.5.1. Introduction

In this section, we will describe how to configure ThinLinc for authentication against One Time Password (OTP) servers, such as the RSA SecurID. By using OTPs, you can increase the system security.

10.5.2. General Requirements

- An OTP server which accepts the OTP twice. This is due to the ThinLinc architecture: The client first contacts the master machine, and then the agent host. When using RSA SecurID, we recommend using the Steel-Belted Radius server as a "Token Caching Server".
- An user database (directory server) that is supported both by the operating system on the ThinLinc servers, as well as the OTP server. We recommend using a LDAP directory server, such as Novell eDirectory.
- The operating systems on the ThinLinc servers must support the OTP servers authentication protocol. We recommend using the RADIUS protocol, by using the `pam_radius_auth` PAM module from the FreeRADIUS project (<http://freeradius.org>).
- The SSH server on the ThinLinc servers must accept "keyboard-interactive" authentication. It's recommended to disable "password" authentication.

10.5.3. Configuration for RSA SecurID

This section describes how to deploy a OTP solution based on RSA SecurID with ThinLinc. When using this solution, the SecurID PASSCODEs are used instead of normal passwords. The PASSCODE should be entered in the ThinLinc client password input field. Please observe the following limitations:

- When SecurID requests additional information, in addition to the PASSCODE initially entered, a popup dialog will be used. This happens, for example, in Next Token or New PIN mode. After finishing the dialog, the ThinLinc client will display a "Login failed!" error message. This happens since the SBR server clears the token cache when additional information is requested. When this happens, wait until the token changes once more, and login again.
- The ThinLinc Single Sign-On mechanism will store the string entered in the clients password input field. When using SecurID, this is the PASSCODE, which cannot be used for further logins. To use the Single Sign-On mechanism, the user must be prompted for their real password. This can be done with the program **tl-sso-update-password**. To configure ThinLinc so that this program is executed during login, execute this command:

```
# ln -s /opt/thinlinc/bin/tl-sso-update-password /opt/thinlinc/etc/xstartup.d/05-tl-sso-update-pass
```

The configuration example below assumes that you are using LDAP and RADIUS, and the Steel-Belted Radius (SBR) server. Step 8 through 11 should be repeated on all ThinLinc servers.

1. Install and configure RSA Authentication Manager (ACE server). For basic configuration tasks such as creating users and assigning tokens, we refer to the RSA documentation.
2. Create a new Agent Host for the SBR server, with type "Net OS Agent". Select which users should be able to login through ThinLinc. To allow all users, check the "Open to All Locally Known Users" checkbox.
3. Generate a configuration file for the SBR server, by selecting Agent Host->Generate Configuration File. Copy this file to `c:\windows\system32` on the machine running SBR.
4. Open the SBR Administrator. Create clients for all ThinLinc servers, using default settings. Make sure you enter a shared secret.
5. Use SBR Administrator to create a SecurID user. The user should typically be of type <ANY>.
6. Modify the SBR Authentication Policy, so that the only active method is "SecurID User". Exit SBR Administrator.
7. Enable ACE authentication caching by editing the configuration file
`c:\radius\service\radius.ini` and set:

```
[SecurID]
CachePasscodes          = yes
SecondsToCachePasscodes = 60
```

Restart SBR after changing the configuration file. For more information about ACE authentication caching, refer to the Steel-Belted Radius Tech Note RD310.
8. Install `pam_radius_auth`. Some distributions, such as SUSE, includes this module.
9. Configure `pam_radius_auth`, by creating `/etc/raddb/server` . It should contain the SBR server name, port, and a shared secret. Example:
`myotpserver.example.com:1812 mysecret`
10. Configure the ThinLinc servers for RADIUS authentication by modifying its PAM configuration. The exact procedure depends on the system, but typically, this can be done by modifying `/etc/pam.d/system-auth`, by inserting the line

```
auth          sufficient      /lib/security/pam_radius_auth.so use_first_pass
```

after the line containing `pam_unix.so`.
11. Restart the VSM and SSH server.
12. Login to the system with a SSH client, and verify that an OTP is required and accepted.
13. Login to the system with a ThinLinc client, and verify that an OTP is required and accepted.

Chapter 11. File Access

11.1. Accessing Windows File Servers

11.1.1. Introduction

This chapter describes how to setup a ThinLinc server to access Windows file servers via the SMB/CIFS protocol. CIFS is a modern version of SMB. In this document, we use the term CIFS, but the procedure described in this documentation works for SMB servers as well.

CIFS is different from NFS in that CIFS mounts are per user, not per system. For example, with NFS, it's possible to mount all network file systems when the server boots. One NFS mount can be used by all users on the system. With CIFS, each user must have their own mounts. Also, when mounting a CIFS file system, the password of the user is usually required.

ThinLinc and many other Linux applications requires that hard links are supported in the user's home directory. There are often other POSIX file system semantic requirements as well. This means that the user's home directories cannot be a mounted CIFS filesystem. The Linux CIFS client (`smbfs`) does not support all POSIX file operations, such as hard links. The newer CIFS client (`cifsfs`) supports the CIFS UNIX extensions, but few CIFS servers support this and this feature has not been tested with ThinLinc.

ThinLinc includes two utility programs for dealing with CIFS mounts: `tl-mount-cifs` and `tl-umount-all-cifs`. These are described below.

The method described in this chapter mounts all CIFS shares below the directory `~/winshares`. The user's CIFS home directory, if any, is mounted at `~/winshares/home`.

11.1.2. Requirements

11.1.2.1. CIFS Server Requirements

This document assumes that you are using a Windows file server. However, you should be able to use any CIFS file server.

Username and passwords must be synchronized between the file server and the ThinLinc server. Usually, this is accomplished by letting the ThinLinc servers and the CIFS file server use a common directory server. For details, please refer to Chapter 10.

11.1.2.2. ThinLinc Server Requirements

Either of `smbmount/smbumount` or `mount.cifs/umount.cifs` must be installed. On Red Hat distributions, they are available in the package `samba-client`. Refer to your distribution for how to install these applications.

The programs `smbmnt/smbumount` and `mount.cifs/umount.cifs` must be setuid root. This is accomplished by the following commands:

```
# chmod u+s /usr/bin/smbmnt /usr/bin/smbumount
# chmod u+s /sbin/mount.cifs /sbin/umount.cifs
```

11.1.3. Mounting and Unmounting Shares

11.1.3.1. Using tl-mount-cifs

tl-mount-cifs is a small wrapper for smbmount and mount.cifs , which adds:

1. Automatically selects which file system implementation to use. `cifsfs` is used if the command `mount.cifs` is available. Otherwise, `smbfs` is tried.
2. Automatically submits password, using the ThinLinc Single Sign-On mechanism.
3. Automatically creates mount point directory, if it does not exist.
4. Can optionally fetch the service and drive letter corresponding to the users home directory specified in Active Directory.
5. Will automatically use the options specified in Hiveconf (as explained below).

The syntax for tl-mount-cifs resembles smbmount/mount.cifs :

tl-mount-cifs [-r] [--verbose] [-o options] service mount-point

tl-mount-cifs [-r] [--verbose] [-o options] --homedir [mount-point]

The `-r` option removes the mount point if the mount fails. The `--verbose` option executes both `tl-mount-cifs` and the actual mount command with debugging information. Additional mount options can be specified using the `-o options` option. Refer to the `smbmount/mount.cifs` documentation for more information. If the `--homedir` option is specified, it is not necessary to specify the service to mount. Instead, the service corresponding to the users home directory will be fetched automatically from Active Directory. This requires that the Samba `net` command is available. When `--homedir` is used, the `mount-point` argument is optional. If omitted, the service will be mounted on a directory in the users home directory corresponding to the drive letter specified in Active Directory (without the trailing colon).

The Hiveconf parameter `/utils/tl-mount-cifs/cifsmount_args` specifies default arguments for the `tl-mount-cifs` command. This Hiveconf parameter is normally found in `/opt/thinlinc/etc/conf.d/tl-mount-cifs.hconf`. The default value of this parameter is `"-o dir_mode=0700"`, which makes CIFS mounts user-private. This option is however only recognized by `mount.cifs`.

Example 1: User "johndoe" has a home directory on the CIFS file server `\\alabama`, shared as `"johndoe$"`, which should be mounted on `/home/johndoe/winshares/home`. To do this, he runs the following command:

```
$ tl-mount-cifs -r //alabama/johndoe$ ~/winshares/home
```

Example 2: User "johndoe" is part of a workgroup that shares files using a share called "project" on the file server \\alabama. This share can be mounted on /home/johndoe/winshares/project with the following command:

```
$ tl-mount-cifs //alabama/project ~/winshares/project
```

Example 3: If a home directory and home drive is specified in Active Directory, the home directory of user "johndoe" can be executed with the command:

```
$ tl-mount-cifs --homedir
```

11.1.3.2. Using tl-umount-all-cifs

tl-umount-all-cifs is a utility that unmounts the current user's mounted CIFS shares (all CIFS mounts below the user's home directory). It requires no arguments. The optional argument -a will unmount all CIFS filesystems on the host.

11.1.3.3. Mounting Shares at Login

Often, it's convenient to automatically mount CIFS shares for all users upon login. This can be accomplished by creating a script in /opt/thinlinc/etc/xstartup.d. It can be named anything. The script should contain something like:

```
#!/bin/sh
/opt/thinlinc/bin/tl-mount-cifs //alabama/${USER}$ ~/winshares/home
```

You should also make sure that tl-umount-all-cifs runs at logout. This can be done with the following command:

```
# ln -s /opt/thinlinc/bin/tl-umount-all-cifs /opt/thinlinc/etc/xlogout.d
```

11.2. Restricting write access to users home directory

11.2.1. Introduction

When accessing directories from CIFS and NCP servers, these are mounted in subdirectories of the users Linux home directory. It is not possible to place the Linux home directory on a CIFS or NCP server, since these typically does not support the necessary POSIX file system semantics (such as hard links). In a typical setup, applications such as Mozilla uses the Linux home directory for settings (~/.mozilla), while the user saves documents in ~/MyDocuments. In this case, it might be desirable to restrict access

to the Linux home directory: Forbid saving arbitrary files to it. This can be solved by using a feature of ThinLinc called `homecreatefilter`.

11.2.2. Activation

To activate `homecreatefilter`, create a symbolic link in the `xstartup.d` directory:

```
# ln -s /opt/thinlinc/libexec/tl-homecreatefilter.sh /opt/thinlinc/etc/xstartup.d/06-tl-homecreatefilter.sh
```

11.2.3. Configuration

The configuration file `/opt/thinlinc/etc/homecreatefilter.conf` controls which files and directories are allowed. By default, all files starting with a dot are allowed, as well as the files necessary for KDE to start.

The configuration file is line based. A line not starting with a colon specifies a file object pattern that should be allowed. A line starting with a colon specifies a command line pattern. Processes matching this pattern will also be allowed write access, even if no file object pattern allows access.

11.2.4. Security Considerations and Limitations

The `homecreatefilter` feature is based on the `LD_PRELOAD` mechanism, which means it does not support statically linked applications. Since environment variables can be modified by the user, the user can disable the filter at will. `homecreatefilter` should not be regarded as a security mechanism, but rather a mechanism that prevents the user from saving documents to the Linux home directory by mistake.

In addition to the home directory, `homecreatefilter` restricts write access to the `~/Desktop` directory.

Chapter 12. Accessing Client Resources from the ThinLinc session

In this chapter we will describe how to access client resources, such as local drives and serial ports, from the ThinLinc session.

12.1. Accessing the Clients Local Drives

12.1.1. Introduction

Using ThinLinc, it is possible to access the clients' drives and filesystems from the ThinLinc session. With thin terminals, one might want to access a local CD-ROM drive. When running the client on a workstation, applications on the remote desktop server can access all filesystems mounted at the workstation, just like local applications can.

Note: Many Digital Cameras can be accessed as a USB storage device, and can be exported as a local drive.

12.1.2. Mounting and Unmounting Local Drives

The exported local drives can be mounted with the command `tl-mount-localdrives`. The drives will be mounted below `$TLSESSIONDATA/drives`. A symbolic link called "thindrives" will be created in the user's home directory, pointing to this directory. The syntax for `tl-mount-localdrives` is:

tl-mount-localdrives [-h] [-v]

The `-v` option causes the tool to be executed in verbose mode, while `-h` shows the syntax.

The Hiveconf parameter `/utils/tl-mount-localdrives/mount_args` specifies the mount arguments. This Hiveconf parameter is normally found in `/opt/thinlinc/etc/conf.d/tl-mount-localdrives.hconf`. The options `mountport`, `port`, `mountvers`, `nfsvers`, `nolock`, and `tcp` will always be used.

Mounted local drives can be unmounted with the command `tl-umount-localdrives`. If some applications are using a mount at this time, they can continue to access the mount, even though the mount has been removed from the filesystem hierarchy (so called "lazy" unmount). The syntax for `tl-umount-localdrives` is:

tl-umount-localdrives [-a] [-s] [-l]

If `-a` is specified, then all mounted local drives, for all users on this machine, will be unmounted. If `-s` is specified, then all mounted local drives, for all sessions belonging to the current user, will be unmounted. If `-l` is specified, the `thindrives` link will not be updated.

Note: When using multiple sessions per user, the `thindrives` link will point to the newest session that executed **tl-mount-localdrives**. **tl-umount-localdrives** will restore the link to the newest session which is not newer than the current session and which has mounted local drives.

12.1.3. Mounting Drives at Login

Often, it's convenient to automatically mount all local drives for a user when the session starts. This is done by default via a symbolic link in `/opt/thinlinc/etc/xstartup.d`, pointing at `/opt/thinlinc/bin/tl-mount-localdrives`. This link is created for you during installation, as well as its counterpart in `/opt/thinlinc/etc/xlogout.d` which points to `/opt/thinlinc/bin/tl-umount-localdrives`.

12.1.4. Limitations and additional information

- Linux kernel 2.6.23 or later is required.
- A mounted local drive, for example `/var/opt/thinlinc/sessions/johndoe/47/drives/cdrom`, is only usable during the lifetime of the ThinLinc session. If the user ends the session without unmounting and then starts a new session, the mount will not be usable even if the session number happens to be same. In this case, any attempts to access the mount will give the error message "Stale NFS file handle". To be able to use the local drive, the user needs to run `tl-mount-localdrives`.
- The mounted local drive does not fully support POSIX semantics. The usual limitations of NFSv3 applies. Additionally, if the file is moved to another directory while a process has the file open, the process will get a "Stale NFS file handle" error on any subsequent file operation for that file.
- Local files are uniquely identified by their inode number. Some file system implementations, such as the Linux kernel FAT implementation, do not provide persistent inode numbers. Inode numbers will change on each remount, which usually results in "Stale NFS file handle" errors.

12.2. Using Serial Port redirection

12.2.1. Introduction

Using ThinLinc, it is possible to access the serial ports of the client from the ThinLinc session. This means that you can utilize peripheral devices which connect through a serial port, such as digital cameras, PDAs and modems. Up to two serial ports are supported at a time.

12.2.2. Requirements

- The application which communicates with the serial port must be dynamically linked. Statically linked applications are not supported.

12.2.3. Enabling Serial Port Redirection

Serial port redirection is activated (for the current user session) by sourcing the file `tl-serial-redir.sh`. It can be done manually with this command:

```
$ source /opt/thinlinc/libexec/tl-serial-redir.sh
```

It is necessary to source this file, because it sets the environment variables `CYCLADE_DEVICES` and `LD_PRELOAD`. Thus, all applications needing serial port access should be started as a subprocess to this shell. The easiest way to accomplish this is to source `tl-serial-redir.sh` from the session startup scripts. To automatically activate serial port redirection at login for all users, execute this command:

```
# ln -s /opt/thinlinc/libexec/tl-serial-redir.sh /opt/thinlinc/etc/xstartup.d/42-tl-serial-redir.sh
```

12.2.4. Accessing the redirected port from applications

When using redirected serial ports, applications should be configured to use a special, personal device-file, instead of a port such as `/dev/ttyS0`. The two device files are called `$TLSESSIONDATA/dev/ttyS0`. and `$TLSESSIONDATA/dev/ttyS1`.

Best Practice: Since the session number varies, it's often convenient to use the symbolic link `/var/opt/thinlinc/sessions/$USER/last`, which points to the last started session directory. For example, the first serial port can be accessed as `/var/opt/thinlinc/sessions/$USER/last/dev/ttyS0`.

12.2.5. Limitations and additional information

- When reconnecting to an existing session, it might take up to 10 seconds before the serial ports are available.
- A maximum of two serial ports per session can be redirected.
- The redirection is handled by processes called `cyclades-ser-cli`. It writes debugging information to `$TLSESSIONDATA/ttyS0.log`. and `$TLSESSIONDATA/ttyS1.log`. These processes will automatically terminate when the session terminates.
- Applications that uses the `ioctl TIOCMGET` are not supported yet.

12.3. Using Sound Device Redirection

12.3.1. Introduction

With ThinLinc, it is possible to access the client's sound device from the ThinLinc session. This means that you can run sound applications on the remote desktop servers and listen to the sound through the client's sound device and speakers. Input devices such as microphones can also be used.

ThinLinc can support sound redirection for almost all applications, provided that the correct libraries and utilities are installed on the ThinLinc server.

12.3.2. Requirements

- PulseAudio client libraries to support applications with native PulseAudio support and the ALSA plug-in. ThinLinc supports version 0.9 of PulseAudio.
- padsp to support OSS applications via PulseAudio.
- alsa-plugins, version 1.0.12 or later, to support ALSA applications via PulseAudio.

12.3.3. PulseAudio applications

All applications that can communicate using the PulseAudio protocol will also work automatically in ThinLinc. Most current distributions are configured to use PulseAudio by default, but older ones might require some configuration to work properly.

12.3.4. OSS applications

Most applications that use the Open Sound System (OSS) can be made to work with ThinLinc through the padsp application.

padsp redirects OSS applications to the PulseAudio protocol. The following command line should be used:

```
padsp <application>
```

See the padsp manual page for more information.

The application which communicates with the sound device must be dynamically linked to glibc. It is not possible to intercept the accesses to OSS in a statically linked application. Most applications that you find on a Linux system will satisfy this requirement, but a test with ldd can also be done:

```
$ ldd /usr/bin/someapp
        not a dynamic executable
```

When using padsp on 64-bit platforms, make sure that you have both 32- and 64-bit versions of the necessary libraries (libpulsedsp.so and libpulse.so.0). Usually, these are found in /usr/lib and /usr/lib64. Also, the padsp script must not contain absolute references to these libraries. Instead,

the system should automatically select the correct library, depending on if you are executing a 32- or 64-bit application. In this case it's necessary to have both library directories included in `/etc/ld.so.conf`.

Although it is rare, some applications manage to misuse the OSS API in a way that works with local sound cards but not `padsp`. If you encounter problems, try updating the application to the latest version as it might contain fixes for such bugs.

12.3.5. ALSA applications

All applications that use the Advanced Linux Sound Architecture (ALSA) will also work well with ThinLinc provided the correct plug-ins are installed and configured. The plug-ins can be found in the `alsa-plugins` package (called `libasound2-plugins` on Debian-based distributions). The PulseAudio client libraries are also needed to build and use the plug-ins.

To redirect ALSA applications to use the plug-ins, the ALSA configuration must be modified. This can be done on a global level in `/etc/asound.conf` or per user in `~/.asoundrc`. Add the following to the file (creating it if necessary):

```
pcm.!default {
    type pulse
}
ctl.!default {
    type pulse
}
```

Unfortunately, there are some applications that use the ALSA API in an incorrect way. When using local hardware this usually doesn't matter, but when advanced ALSA features, like `dmix` or this plug-in, are used, then problems start to appear. If an application misbehaves, the first step should be to upgrade it to the latest version. With some luck, the API is used more correctly in a later version.

12.3.6. Choosing sound system

Many applications support several sound systems and it can be difficult to know which one to use. Applications should be configured in the following manner, listed from the best solution to the worst:

1. Native PulseAudio application.
2. ALSA application using the PulseAudio plug-in.
3. OSS application using `padsp`.

12.3.7. Limitations and additional information

- Transferring sound over the network requires a lot of bandwidth, so it is only suitable for high-speed networks, such as LANs.

12.4. Using Smart Card Redirection

12.4.1. Introduction

Using ThinLinc, it is possible to access the locally connected smart cards and smart card readers from the ThinLinc session. This means that you can use smart cards for encrypting your email, signing documents and authenticating against remote systems.

12.4.2. Requirements

- The application which communicates with the smart card must be using the PC/SC API and be dynamically linked to pcsc-lite.

12.4.3. Enabling Smart Card Redirection

Smart card redirection is always activated on the server so there is no administration required to enable it.

12.4.4. Limitations and additional information

- When a client disconnects, all smart cards and smart card readers will disappear for the applications. Some applications do not handle hot-plug and must therefore be restarted when this happens.

Chapter 13. Commands on the ThinLinc Server

In this chapter, we will describe the commands shipped as part of the ThinLinc server that are meant for the common user.

Commands in `/opt/thinlinc/bin`

tl-session-param [*options*] *parameter*

The **tl-session-param** command is used to access the session information managed by the VSM server. This includes information sent by the client, such as if the client has exported any local drives, or what language is set on the client side. This command is used by for example **tl-set-clientlang.sh**, documented later in this chapter.

tl-config *options*

The **tl-config** command is used to access configuration parameters used by the ThinLinc system. It is also used to set parameters from scripts, and can be used instead of an editor when some parameter needs to be changed. **tl-config** uses **hivetool**, part of the Hiveconf system. See Chapter 16 for more information about Hiveconf.

tl-desktop-restore

When a user's Gnome or KDE desktop needs to be reset to default, the command **tl-desktop-restore** can be run. This will move the settings directories for KDE and Gnome to a backup directory named `.old-thinlinc-desktop` in the user's home directory, which will make both Gnome and KDE revert to the default settings.

tl-limit-printers

This command is run by VSM Server at session startup and reconnect if the Printer Access Control feature of ThinLinc is activated. See Section 5.5 for details.

tl-mount-cifs

This command is used to mount CIFS/SMB network file systems at login-time. See Section 11.1 for documentation on this subject.

tl-memberof-group *groupname...*

This command can be used to determine if the current user is a member of the specified groups. It returns true (0) if the user is a member of any of the groups, false (1) if the user is not a member and false (2) if any of the specified groups do not exist.

tl-passwd

This command is used to let the user change their password, both in the underlying authentication mechanism and in the ThinLinc Single Sign-On mechanism.

In order for this to work, any user must be able to read the file `/etc/pam.d/sshd` (or, more correct, the file that the symbolic link `/etc/pam.d/thinlinc` points at).

Also, in the case where the underlying authentication mechanism is LDAP or eDirectory, make sure that the parameter `pam_password` in `/etc/ldap.conf` is set to a value that is appropriate for your environment. If you're authenticating against eDirectory servers, it must be set to `nds`. See the comments in `ldap.conf` for more information.

tl-run-xstartup.d

This command is run by the default session startup file (`/opt/thinlinc/etc/xstartup.default`) to execute all start scripts in the directory `/opt/thinlinc/etc/xstartup.d/`. Files with the suffix `.sh` will be sourced. All other files will be executed.

tl-select-profile

This command is run by the session setup file (`/opt/thinlinc/etc/xstartup.default` or `~/.thinlinc/xstartup`) and provides a menu where the user can choose what kind of session to run. See Section 14.4 for more information.

tl-set-clientlang.sh

By creating a symlink from `/opt/thinlinc/etc/xstartup.d` to this command, the user's LANG environment will be set to the language environment reported by the client.

tl-shadow-notify

This command starts the **tl-shadow-notify** command for the lifetime of the session. This will enable notifications when the session is shadowed.

tl-single-app *command* [*arguments*]

The **tl-single-app** command can be used to execute a single application in a ThinLinc session. A window manager with a suitable configuration is automatically started. All top level windows are automatically maximized. Window titles are displayed in the title bar of the ThinLinc Client, not in the ThinLinc session. The client close button will disconnect the session as usual. Inner close buttons closes application windows. The **tl-single-app** command can be specified as a client supplied start program (see Section 14.4.4), or used with the ThinLinc profile selector (see Section 14.4.5).

Switching Between Windows: If the application opens multiple top level windows, you can switch between them by clicking on the application icon in the top left corner.

tl-sso-update-password

This command requests a password from the user, to be used with the Single Sign-On mechanism of ThinLinc. It is useful when the password is not already available, for example, when using One Time Passwords. See Section 10.5.3 for more information.

tl-support [-p *listen-port*] [-u *user*] [*host*]

The **tl-support** command can be used to enable a support technician to login to your ThinLinc server, even though the server is behind a firewall that doesn't allow connections to the ssh port. This is accomplished by opening a ssh connection from the server to an external server on the internet, at the same time setting up a tunnel from the remote host to the local host's ssh port. The default server to connect to is support.thinlinc.com with the default username "support". This command should only be used after contacting your ThinLinc support technician.

tl-umount-all-cifs

This command is used to unmount CIFS/SMB network file systems at logout-time. See Section 11.1 for documentation on this subject.

tl-disconnect

This command is used to disconnect from the current session. This can be used to provide an alternative to the F8 key, such as a disconnect button on the Gnome panel.

tl-sso-password [--check] [--remove]

This command can be used to hook up the Single Sign-on mechanism of ThinLinc with new applications. It can be used to test for the presence of a valid password and to feed that password out on standard output to another application.

To check for the existence of a valid password, invoke the command as **tl-sso-password --check**. A return code of zero indicates a valid password.

If the **--remove** option is specified, the password will be removed, after the retrieval or check.

There are two basic models to connect **tl-sso-password** to an application. The first is to use shell pipes:

```
# tl-sso-password | /usr/bin/application --read-password-on-stdin
```

The second is to have the application invoke **tl-sso-password** as needed:

```
# /usr/bin/application --password-prog tl-sso-password
```

tl-sso-token-passphrase [--check] [--remove]

This command is identical to **tl-sso-password**, except that it uses the smart card token passphrase (PIN) instead of the user's password. For usage, see the **tl-sso-password** section above.

tl-env [-d] [-n *nr*] [*command* [*arg...*]]

tl-env [-s] [-n *nr*]

This command can be used to save and restore the ThinLinc session environment variables. It operates on the file `xstartup.env` in the session directory. During session startup, **tl-env** is called with the `-s` option after everything in `xstartup.d` have been executed. Later, **tl-env** can be used to execute a command in this environment, even outside the ThinLinc session. During restore, the `DISPLAY` environment variable can be excluded by specifying `-d`. By default, this command operates on the "last" session number for the invoking user. An alternative session number can be specified with the `-n` option.

Commands in /opt/thinlinc/sbin**tl-notify** [-u *username*] *message*

This command sends a user-visible message to ThinLinc sessions on the server. The default is to send the message to all sessions, but the -u option can be used to send the message to a single recipient instead.

To send messages to all users in a ThinLinc cluster, you can use this command in combination with the **tl-ssh-all** command described in this section.

tl-rsync-all

This command is used to synchronize files and directories in a ThinLinc cluster. It runs the rsync command over SSH against all agent servers in the cluster. When using this command, it's convenient if password-less SSH login between the servers in the clusters has been setup.

See also **tl-ssh-all** below for some tips regarding password-less running of ssh.

tl-ssh-all

This command is used to perform shell commands on all agents in a ThinLinc cluster. It works by running the ssh command against all agent servers in the cluster. When using this command, it's convenient if password-less SSH login between the servers in the clusters has been set up.

Best Practice: An alternative approach to using password-less login is to use the SSH agent to cache the passphrase of a SSH keypair. This increases the security, since a malicious party that gains access to the server which is configured to login to the other servers with SSH key-pair does not automatically get access to the rest of the servers - a password is needed.

First, setup the SSH key-pair as described below:

```
#
# First time / One time procedure
#
# Generate a private and public key-pair for SSH with SSH keygen.
# When prompted pick a secret password for the key-pair.
#
ssh-keygen -t dsa

# Copy the public key to SSH authorized_keys
cp /root/.ssh/id_dsa.pub /root/.ssh/authorized_keys

# Make sure the authorized key has the right permissions
chmod 600 /root/.ssh/authorized_keys

# Copy the authorized key to all ThinLinc Agents
tl-rsync-all /root/.ssh/authorized_keys
```

Next, before using `tl-ssh-all`, do as follows

```
eval `ssh-agent`  
ssh-add  
  
# Run your commands  
tl-ssh-all rpm -Uvh /root/kdelibs-3.5.1-1.fc4.i386.rpm
```

Commands in `/opt/thinlinc/libexec`

tl-crossover-drives

CodeWeavers CrossOver allows you to configure the mapping between Windows drive letters and paths in the Linux file system. This can be done globally by adding symbolic links to the directory `/opt/cxoffice/support/BOTTLENAME/dosdevices`. However, this does not work if drive letters should correspond to different paths for different users. In this case, a bottle hook script is required. **tl-crossover-drives** is such a script that automatically maps "personal" mounts to separate drive letters in CrossOver. This includes all mounts mounted on subdirectories in the users home directory. The first character of the directory name determines the drive letter. To activate this command for all bottles, execute:

```
# mkdir /opt/cxoffice/support/scripts.d  
# ln -s /opt/thinlinc/libexec/tl-crossover-drives \  
/opt/cxoffice/support/scripts.d/02.tl-crossover-drives
```

tl-has-gnome-2

The **tl-has-gnome-2** command is used to check if Gnome 2 is installed on the system, in a way which works for most distributions. It is used by the default profile configuration.

tl-unity-2d [`--test`]

The **tl-unity-2d** command is used to start the Unity 2D desktop environment, in a way that works on most distributions. It is used by the default profile configuration. The `--test` option can be used to test if this desktop environment is installed.

tl-kinit.sh

The **tl-kinit.sh** command is used to obtain a Kerberos ticket automatically during start of the session, using the single sign-on mechanism.

tl-kdestroy.sh

The **tl-kdestroy.sh** command is used to destroy the Kerberos ticket cache. It calls **kdestroy** during logout.

Chapter 14. Server Configuration

14.1. Configuring ThinLinc Servers in a Cluster

In this section, we will describe how to configure a ThinLinc cluster with multiple agent servers to allow load-balancing and redundancy.

Note: This section does *not* address configuration of high availability (HA). For information on configuring your ThinLinc cluster for high availability, see Chapter 6.

A ThinLinc cluster consists of one master server (or multiple master servers in a HA configuration) with multiple agent servers grouped into subclusters. While ThinLinc in its simplest configuration may be run with both the master and agent installed on the same machine, running ThinLinc in a cluster configuration conveys numerous advantages:

1. A cluster configuration allows automatic load-balancing of sessions across multiple agents
2. Having multiple agents offers redundancy; for example, if one agent goes down or is taken out of service for maintenance, other agents are still available to handle user sessions
3. A cluster configuration is scalable. Since most of the workload is taken up by the agents and not the master, adding more capacity to your ThinLinc installation is generally as easy as installing one or more new agent servers

14.1.1. Cluster Configuration

When configuring ThinLinc servers as a cluster, one needs to configure both the master server and the agents. The master server needs configuration for subclusters (even if there is only one agent) and the agents need to know which server is the master for access control.

The configuration parameter `/vsmagent/master_hostname` for each agent that is included in a ThinLinc cluster should be configured with the address of the master server for the cluster. This gives the master access to communicate with and control the agent server.

Once the master and agents within a cluster are configured, and the `vsmagent` and `vmsserver` services have been restarted, these ThinLinc servers will then function as a cluster.

14.1.1.1. Subclusters

A subcluster is a group of agents. At least one subcluster is always active, even in a single server setup. Each subcluster can serve a specific purpose within the ThinLinc cluster. The dimension for grouping can be chosen at will and could for example be; location, project, application etc.

ThinLinc ships with one default subcluster configuration which is used for creating new sessions by any user. It is allowed to define as many subclusters as needed. Each subcluster can be associated with users and with user groups.

To associate a user with a subcluster, either use the `users` or `groups` configuration parameter for the specific subcluster. See Section 14.2.3 for details on these subcluster configuration parameters.

If a subcluster does not have neither user nor group associations configured, it is used as a default subcluster. Users that does not belong to any subcluster, will have their sessions created on the default subcluster. If a user is not associated with a subcluster and there is no default subcluster configured, the user will not be able to logon to the ThinLinc cluster.

Loadbalancing of new sessions is performed using the list of agents defined in the `agents` parameter within each subcluster.

A subcluster is defined as a folder under the `/vsmserver/subclusters/` configuration folder in `vsmserver.hconf`. The foldername defines a unique subcluster name.

Here follows an example subcluster configuration for a geographic location based grouping of agents:

```
[/vsmserver/subclusters/default]
agents=tla01.eu.cendio.com tla02.eu.cendio.com

[/vsmserver/subclusters/usa]
agents=tla01.usa.cendio.com tla02.usa.cendio.com
groups=ThinLinc_USA

[/vsmserver/subclusters/india]
agents=tla01.india.cendio.com tla02.india.cendio.com
groups=ThinLinc_India
```

During the selection process for which subcluster a new session is created on, the following rules apply:

1. `users` has precedence over `groups`. This means that if a user belongs to a group that is associated with subcluster A and the user also is specified in `users` for subcluster B, the user session will be created on subcluster B.
2. `group` has precedence over the default subcluster. This means that if a user belongs to a group that is associated with subcluster B, the user session will be created subcluster B and not on the default subcluster A.
3. If user does not belong to an associated group nor is explicitly specified by `users` for a subcluster, the new session will be created on the default subcluster.

Note: Try to avoid the following configurations that will result in undefined behaviors for subclusters:

1. Avoid two subclusters without associated `users` and `groups`, eg. default subclusters. It is undefined which of them that will be the default subcluster used for users that are not associated to a specific subcluster.
2. If a user is a member of two user groups which are used for two different subclusters, it is undefined which subcluster the new session will be created on.
3. If a user is specified in `users` of two different subclusters, it is undefined which subcluster the new session will be created on.

14.1.2. Cluster Management

When multiple agents have been configured as part of a cluster, it is usually desirable to keep their configurations synchronised. Instead of making the same change separately on each agent, ThinLinc ships with the tool `tl-rsync-all` to ensure that configuration changes can be synchronised easily across all agents in a cluster. See Chapter 13 for more information on how to use this tool.

The `tl-rsync-all` command should be run on the master server, since it is the master which has the knowledge of which agents there are in the cluster. For this reason, it is often useful to configure the master server as an agent as well, even if it will not be used to host user sessions in general. This allows the master to be used as a "template" agent, where configuration changes can be made and tested by an administrator before pushing them out to the rest of the agents in the cluster. In this way, configuration changes are never made on the agents themselves; rather, the changes are always made on the master server, and then distributed to the agents using `tl-rsync-all`.

An example of how one might implement such a system is to configure the master server as an agent which only accepts sessions for a single administrative user. The steps to do this are as follows:

1. Configure the master as an agent too. On a ThinLinc master, the `vsmagent` service should already have been installed during the ThinLinc installation process; make sure that this service is running.
2. Create an administrative user, for example `tladmin`. Give this user administrative privileges if required, i.e. `sudo` access.
3. Create a subcluster for the master server and associate administrator user `tladmin` to it. See following example:

```
[/vmsserver/subclusters/Template]
agents=127.0.0.1
users=tladmin
```

See Section 14.1.1.1 for details on subcluster configuration.

4. Restart the `vmsserver` service.

In this way, configuration changes are never made on the agents themselves; rather, the changes are always made on the master server, and then tested by logging in as the `tladmin` user. If successful, these changes are then distributed to the agents using `tl-rsync-all`.

14.2. Server Configuration Parameters

The ThinLinc server is configured using a number of configuration parameters stored in Hiveconf. For information about how to access and set the parameters, please refer to Chapter 16. In this chapter, we will describe the different parameters and their meaning.

The parameters used in ThinLinc are divided into a number of folders, each having zero or more subfolders. The following folders exist:

- `/vsm/` contains parameters common to both the VSM agent and the VSM server. This folder normally resides in `/opt/thinlinc/etc/conf.d/vsm.hconf`
- `/vsmagent/` contains parameters specific to the VSM agent. This folder normally resides in `/opt/thinlinc/etc/conf.d/vsmagent.hconf`

- `/vsmserver/` contains parameters specific to the VSM server. This folder normally resides in `/opt/thinlinc/etc/conf.d/vsmserver.hconf`
- `/vsmserver/subclusters/` contains definitions of subclusters within the ThinLinc cluster. This folder normally resides in `/opt/thinlinc/etc/conf.d/vsmserver.hconf`
- `/profiles/` contains parameters for configuring the different session profiles. This folder normally resides in `/opt/thinlinc/etc/conf.d/profiles.hconf`
- `/utils/` contains parameters used by miscellaneous ThinLinc utilities. Each utility has its own configuration file, but all parameters are then merged in under `/utils` when read by the HiveConf framework.
- `/sessionstart/` contains some parameters used during session startup. This folder normally resides in `/opt/thinlinc/etc/conf.d/sessionstart.hconf`
- `/shadowing/` contains parameters used to control access to the shadowing feature. This folder normally resides in `/opt/thinlinc/etc/conf.d/shadowing.hconf`
- `/tlwebadm/` contains parameters for the tlwebadm web configuration interface. This folder normally resides in `/opt/thinlinc/etc/conf.d/tlwebadm.hconf`
- `/webaccess/` contains parameters for the server part of ThinLinc Web Access. This folder normally resides in `/opt/thinlinc/etc/conf.d/webaccess.hconf`

14.2.1. Parameters in `/vsmagent/`

In this section, we will describe all the parameters currently used by the VSM agent.

`/vsmagent/agent_hostname`

Public hostname; the hostname that clients are redirected to. If not defined, the agent will use the computer's IP address. This is the default configuration, and means that ThinLinc does not require DNS to work properly. However, if you are using Network Address Translation (NAT), you must set this parameter to a IP address or DNS name that all clients can connect to. Example:

```
agent_hostname = thinlinc.example.com
```

`/vsmagent/allowed_clients`

This is the space-separated list of VSM servers that should be allowed to connect to this VSM agent and create new sessions. The localhost is always allowed as well as the IP of the hostname the VSM agent runs on, and the host specified in the `/vsmagent/master_hostname/` parameter.

`/vsmagent/default_environment`

This subfolder of `/vsmagent` contains environment variables that should be set in each user's session. Example:

```
[/vsmagent/default_environment]
TOWN=Springfield
LC_CTYPE=sv_SE.UTF-8
FOOBAR=foobar
```

This will set the `TOWN` environment variable to `Springfield`, the `LC_CTYPE` variable to `sv_SE.UTF-8` and the `FOOBAR` variable to `foobar` in each user's session.

Note: `xsession` is executed via a login shell, which may modify the environment and override values in `[/vsmagent/default_environment]`.

`/vsmagent/default_geometry`

The default session size, to be used when clients are not requesting any specific session size.

`/vsmagent/display_max`

The maximum display number to be used for ThinLinc sessions on each specific VSM agent host. Default value is 2000.

The maximum ThinLinc sessions allowed on a specific VSM Agent host is `/vsmagent/display_max - /vsmagent/display_min`.

`/vsmagent/display_min`

The lowest display numbers to use for clients. The default is 10, and unless there are other processes needing display numbers, the recommendation is not to change this number. See Appendix A for an in-depth explanation of port allocation.

`/vsmagent/listen_port`

The TCP port VSM Agent listen to for incoming requests. This should normally be set to the same value as `/vsm/vsm_agent_port`.

`/vsmagent/lowest_user_port`

The lowest port to be used by normal user processes. This may *never* be lower than `/vsmagent/max_session_port`. See Appendix A for an in-depth explanation of port allocation.

`/vsmagent/make_homedir`

If this parameter is true, the users home directory will be automatically created if it doesn't exist.

`/vsmagent/make_homedir_mode`

When a home directory is created (see parameter `/vsmagent/make_homedir` above), the mode for the newly created directory will be determined by this parameter.

`/vsmagent/master_hostname`

This parameter specifies the hostname of the master machine, i.e. the machine that runs the VSM server. In a HA setup, this should be the hostname of the IP address that is on the machine that is currently the active node, to ensure that services on the agents that need to access the VSM Server always connects to the machine that is up and running.

`/vsmagent/max_session_port`

The highest port to use for VNC and tunnel ports on the VSM Agent. See Appendix A for an in-depth explanation of port allocation.

`/vsmagent/single_signon`

This parameter decides whether the passwords of the users should be saved in order to support Single Sign-On when connecting to servers from the ThinLinc session, for example when running a Windows session.

`/vsmagent/xserver_args`

Extra arguments to pass on to the Xserver Xvnc. One common case is to use `-localhost`, which makes Xvnc require connections to originate from localhost, thus forcing applications to either be local or use a tunnel (which often also means that the traffic is encrypted). Other examples include `-IdleTimeout` and `-MaxIdleTime`. For more information, see Section 14.5.

`/vsmagent/xauthority_location`

This parameter controls the location of the `Xauthority` file. Currently, two values are supported: With "homedir", the file will be placed in the users home directory. With "sessiondir", the file will be placed in the session directory below `/var/opt/thinlinc/sessions`. The `XAUTHORITY` environment variable is set accordingly by the VSM agent.

14.2.2. Parameters in `/vsmserver/`

In this section, we will describe all the parameters currently used by the VSM server.

`/vsmserver/admin_email`

The administrator's email address. This is where warnings about overuse of Licenses are sent, among with other administrative messages. Make sure this is a valid address.

`/vsmserver/allowed_clients`

A space-separated list of hosts from which privileged operations are allowed. The default (empty) allows localhost to do this. Privileged operations are for example to deactivate a session, something that should be allowed by the host running the ThinLinc Web Administration service.

`/vsmserver/allowed_groups`

ThinLinc access can be limited to certain groups. If the `allowed_groups` space-separated list is empty, all users are accepted. Otherwise, the user must be a member of the groups listed below, to be able to use ThinLinc. Example:

```
allowed_groups = students teachers
```

`/vsmserver/bogomips_per_user`

Estimated bogomips required for each user.

`/vsmserver/existing_users_weight`

This parameter decides the importance of the amount of logged in users on a VSM agent host when calculating load balance parameters. A host with low load, but a lot of users, is generally more likely to get a higher load within short time when the users get active. For this reason, the load balance calculating code takes the number of users at a certain host into its calculation. The `/vsmserver/existing_users_weight` controls how important this factor is. A higher value of this parameter means the load balancing code will care less about a high number of users on a certain machine.

Note: This parameter should normally not be changed, unless when fine-tuning the load balancing.

`/vsmserver/HA/enabled`

If this parameter is true, the VSM server will try to replicate information about sessions to the other VSM server node. See Chapter 6 for more information about ThinLinc in a High Availability configuration.

`/vsmserver/HA/nodes`

This parameter lists the hostnames of both nodes in a ThinLinc HA setup. The space-separated list should include the hostname of the current node. This means that `vsmserver.hconf` can be identical on both nodes.

`/vsmserver/listen_port`

The TCP port VSM Server listen to for incoming requests. This should normally be set to the same value as `/vsm/vsm_server_port`.

`/vsmserver/load_update_cycle`

The number of seconds allowed for updating the load status in the entire cluster.

`/vsmserver/max_sessions_per_user`

The maximum number of sessions allowed per user. 0 means no limit.

`/vsmserver/ram_per_user`

Integer, number of estimated MiB memory required for each session.

`/vsmserver/unbind_ports_at_login`

If this parameter is true, processes occupying the users' interval of forwarded ports will be killed at login. This means that if a user logs in twice to the same session, the second login will get working tunnel ports, if this parameter is true. The first session's tunnel ports will stop working. If the parameter is false, the first session will keep the ports.

14.2.3. Parameters in `/vsmserver/subclusters/`

In this section, we will describe all the parameters used for defining subclusters. For more information about subclusters see Section 14.1.1.1.

`/vsmserver/subclusters/<name>/agents`

All ThinLinc agents part of this ThinLinc subcluster. This should be a space-separated list of DNS host names. These will be used for communication between the master and the agent. The name reported to the client is fetched from the agent itself; names in `/vsmserver/subclusters/<name>/agents` are not reported directly to the clients.

`/vsmserver/subclusters/<name>/users`

All users that should be associated with this specific ThinLinc subcluster. This should be a space-separated list of usernames.

`/vsmserver/subclusters/<name>/groups`

All user groups that should be associated with this specific ThinLinc subcluster. This should be a space-separated list of groupnames.

14.2.4. Parameters in `/vsm/`

Parameters in the `/vsm/` folder are used by both the VSM agent and the VSM server. Neither of them need to be changed on a normal ThinLinc installation.

`/vsm/tunnel_bind_base`

The tunnels setup by the client to access various resources (audio, serial port, network resources, local printer) need one port number each on the server running the VSM agent the client is connected to. This parameter decides the lowest such port that is allocated by the VSM agent. Each user has a port range defined by the formula `/vsm/tunnel_bind_base + display-ID*10 + service_slot` where the `service_slot` depends on which service will use the tunnel. This port range is however used only for sessions with display numbers less than 100. See Appendix A for an in-depth explanation of port allocation.

Note: This parameter should normally not be changed.

`/vsm/tunnelservices/`

There are several parameters under the `/vsm/tunnelservices` folder. Each one decides which ports are used at serverside termination points for the tunnels used to access client resources. See Appendix A for an in-depth explanation of port allocation.

Note: None of these parameters should normally be changed.

`/vsm/tunnelslots_per_session`

The number of ports to reserve for tunnel port endpoints on the server. The number of ports actually used depends on the number of services defined under `/vsm/tunnelservices/`. We recommend letting this parameter have its default value (10), since that provides a margin for future services. See Appendix A for an in-depth explanation of port allocation.

Note: This parameter should normally not be changed and must not be changed whilst there are any running sessions.

`/vsm/vnc_port_base`

The port base for VNC communication. The VNC protocol runs on one port per active user on the VSM agent host, and this is the base of the numbers used. That is, for the first user, the port will be `/vsm/vnc_port_base + 1`, for the second user `/vsm/vnc_port_base + 2` and so on. This algorithm is used only for display numbers below 100. See Appendix A for an in-depth explanation of port allocation.

Note: This parameter should normally not be changed.

`/vsm/vsm_agent_port`

VSM agent communication. This is the port that the VSM server connects to on VSM Agents. This traffic is not encrypted.

Note: This parameter should normally not be changed

`/vsm/vsm_server_port`

The port that the VSM server listens to.

Note: This parameter should normally not be changed

14.2.5. Parameters in `/sessionstart/`

In this section, we will describe all the parameters currently used by the session startup scripts.

`/sessionstart/background_color`

The initial color of the background that is set early during session startup. By default this is a dark blue color.

`/sessionstart/background_image`

A PNG image used as the initial background. The image will always be scaled to cover the entire screen.

If the image contains transparency then the color set by `background_color` will shine through.

`/sessionstart/keyboard_layout`

The default virtual keyboard layout used by Xvnc. The protocol is not dependent on this being configured, but some applications can misbehave if a different virtual layout is configured compared to the real keyboard layout on the client device.

A list of possible keyboard layouts is given from this command:

```
$ man /opt/thinlinc/share/man/man7/xkeyboard-config.7
```

14.2.6. Parameters in `/shadowing/`

In this section, we will describe all the parameters currently used by the shadowing feature.

`/shadowing/allowed_shadowers`

A space-separated list of users and/or groups that are allowed to shadow other users. Group names are prefixed with `+` sign. Please note that these users will gain full access to other users' sessions. See Chapter 15 for more information.

`/shadowing/shadowing_mode`

A constant string value of; `reject`, `silent`, `notify` or `ask`. This value configures in which way a shadowing request is handled.

`reject`

All shadowing requests are rejected. You should set this if you want to disable the shadowing feature.

`silent`

All shadowing requests are accepted and the user will not be notified about being shadowed.

`notify`

All shadowing requests are accepted and a message box will be shown to notify the user when the shadowing starts and when the shadowing ends.

`ask`

Shows a dialog to the user and gives him the full control of deciding to accept or reject the shadowing request. If the request timeout is reached without the user making a decision then the shadowing request will be rejected. Like for `notify` the user is also informed when the shadowing ends.

See Chapter 15 for more information.

14.2.7. Parameters in /tlwebadm/

In this section, we will describe all the parameters currently used by the ThinLinc Web Administration.

`/tlwebadm/username`

The username to authenticate with when accessing the web interface.

`/tlwebadm/password`

The password for the above user. The tool `/opt/thinlinc/sbin/tl-gen-auth` may be used to create hashes of the format required for use with this parameter.

`/tlwebadm/cert`

The path to the certificate file to be used for TLS encryption.

`/tlwebadm/certkey`

The path to the certificate private key file.

`/tlwebadm/listen_port`

The local port for the web server to listen on.

`/tlwebadm/gnutls_priority`

The GnuTLS priority string is used to select the order and availability of TLS versions, ciphers, key exchange, MAC, compression, signature and elliptic curve algorithms for TLS sessions. See Appendix D for possible values.

`/tlwebadm/logging/logfile`

The file to use for logging tlwebadm messages. By default, this is `/var/log/tlwebadm.log`.

14.2.8. Parameters in /webaccess/

In this section, we will describe all the parameters currently used by the ThinLinc Web Access client.

`/webaccess/cert`

The path to the certificate file to be used for TLS encryption.

Note: This certificate may be downloaded by connecting clients to be installed in their browsers. Make absolutely sure that this file does not contain a private key.

`/webaccess/certkey`

The path to the certificate private key file used for TLS encryption.

`/webaccess/login_page`

The URL which is used to redirect back to the Web Access login page on the master server. The default value is `/`, which redirects to the current server. This parameter needs to be changed when ThinLinc Web Access is used in a cluster setup.

`/webaccess/listen_port`

The local port for this service to listen on. The default port used is 300.

`/webaccess/gnutls_priority`

The GnuTLS priority string is used to select the order and availability of TLS versions, ciphers, key exchange, MAC, compression, signature and elliptic curve algorithms for TLS sessions. See Appendix D for possible values.

`/webaccess/logging/logfile`

The file to use for logging tlwebaccess messages. By default, this is `/var/log/tlwebaccess.log`.

14.3. Configuring Logging on ThinLinc servers

In this section we will describe how ThinLinc logs activities on the server, and how the logging can be configured.

Logs are written by each individual session and by the following ThinLinc server components.

- VSM server
- VSM agent
- Web Integration
- Web Administration Interface
- Web Access (HTML5 client)

14.3.1. ThinLinc server components

Logging is configured by editing parameters in Hiveconf. Each component that uses the logging system has a Hiveconf folder named `logging`. In this folder and its subfolder, the logging system is configured.

14.3.1.1. Log destinations

Logs can be written either to file on the local disk, to syslog or both.

14.3.1.1.1. Writing Logs to File

The file name for the log file written to local disk is configured by editing the parameter `logfile` under the `logging` folder. To turn off logging to file, set the parameter `log_to_file` to 0. Note that the log

file will still be created. If abnormal situations occur because of programming errors, data may appear in the file.

14.3.1.1.2. Writing Logs to Syslog

For large installations, using a central loghost might be very convenient. ThinLinc supports writing logs to syslog, which makes it possible to collect all logs in one place.

By setting the parameter `log_to_syslog` under the `logging` folder to 1, logs will be written to syslog. Specify the syslog facility in the parameter `syslog_facility`. The default behaviour is not to log to syslog.

If the parameter `syslog_host` is set, logs will be sent via UDP to the syslog daemon on the host specified. If not, logs will be sent to syslog by writing to the socket specified in `syslog_socket`. The latter is the default.

14.3.1.2. Subloggers

Each program doing logging uses a number of sub loggers. Sub loggers are a way to distinguish different types of information written by the program. For example, the VSM server uses the subloggers *license*, *session* and *shadow* for logging license-related messages, information about sessions and information about shadowing respectively. Every sublogger can be configured to use a different log level. This allows the system administrator to, for example, increase the information about new sessions without increasing the overall loglevel, easing debugging of specific problems.

14.3.1.3. Log levels

The amount of logging can be configured using log levels. The log levels available are:

Table 14-1. Log Levels

Log Level	Explanation
ERROR	Unrecoverable Errors
WARNING	Warnings - something went wrong, but ThinLinc can recover.
INFO	Messages that are useful in daily maintenance.
DEBUG	Messages that can be of use for system administrators when debugging problems.
DEBUG2	Messages useful to trained ThinLinc personel when doing advanced debugging.

The log level configured can be seen as a barrier. If the log level is set to for example INFO, log messages with a level of INFO or higher are let through. If the log level instead is set to DEBUG2, all log messages are let through, since all log levels are higher than DEBUG2.

There is one default loglevel, and one loglevel per sublogger defined. If the log level for a sub level is set to a lower value than the default loglevel, more information will be written by that specific sublogger.

The default loglevel is configured in the `logging/defaultlevel` parameter. Each sublogger's level can then be configured by setting the parameters under the `logging/levels` folder.

14.3.1.4. Summary

The default logging configuration is summarized in Table 14-2.

Table 14-2. Default Log Behaviour

Component	Default Behaviour	Log Configuration Hive Folder
VSM server	Log to <code>/var/log/vsmserver.log</code>	<code>/vsmserver/logging</code>
VSM agent	Log to <code>/var/log/vsmagent.log</code>	<code>/vsmagent/logging</code>
Web Administration Interface	Log to <code>/var/log/tlwebadm.log</code>	<code>/tlwebadm/logging</code>
Web Access (HTML5 client)	Log to <code>/var/log/tlwebaccess.log</code>	<code>/webaccess/logging</code>

14.3.2. Per-Session Logging

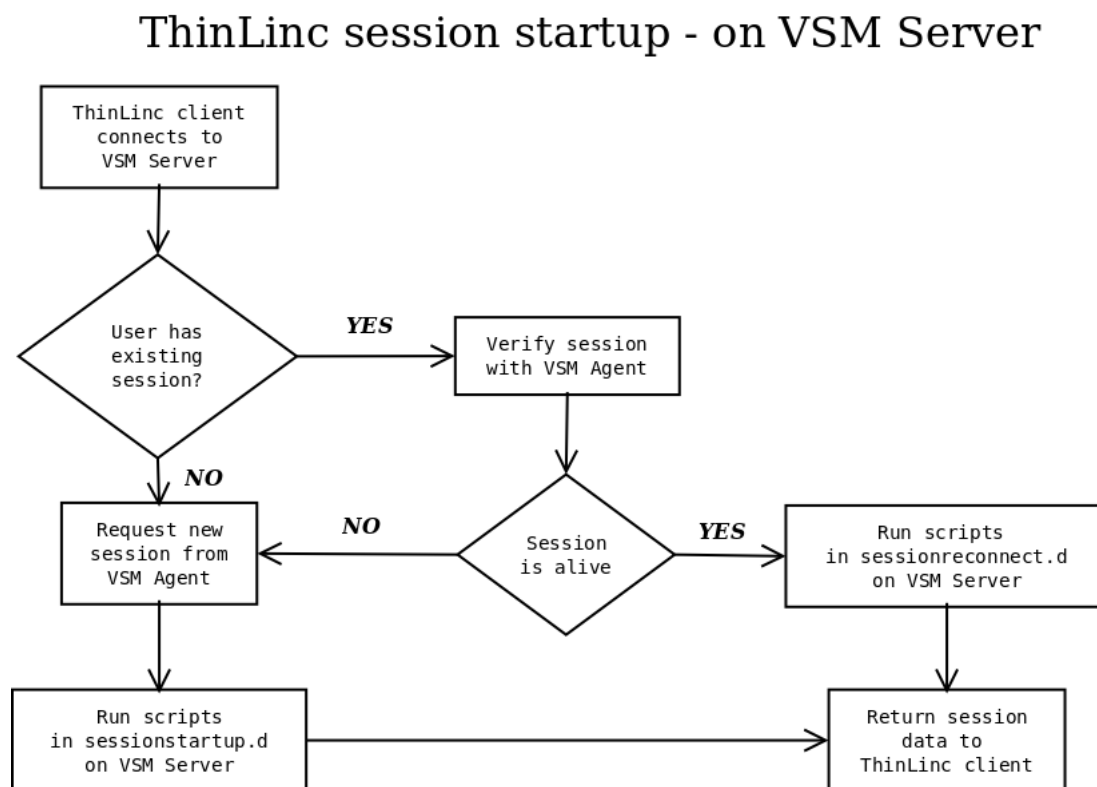
Each Session writes what is written to standard output and standard error output to a file named `xinit.log` which is located in the session directory for a specific session. For example, the log for the last session of the user "johndoe" is located in `/var/opt/thinlinc/sessions/johndoe/last/`. This log contains for example output written by software run in the session, but it also has some output from ThinLinc software that is run by the user.

14.4. Customizing the User's Session

In this section, we will describe how the session startup in ThinLinc can be customized.

14.4.1. Session startup - the big picture

The session setup is constructed to be easy to use and configure yet still easy to customize for advanced use cases.

Figure 14-1. Session Startup Procedure - on VSM Server.

In Figure 14-1, shows a (simplified) description of what happens on the VSM Server when a client connects to login:

- The VSM Server checks if the user has an existing session.
- If a session exists, VSM Server contacts the VSM Agent running on the host where the session is running, and asks it to verify that the session is still alive.
- If the session was alive, VSM Server runs any scripts placed in `/opt/thinlinc/etc/sessionreconnect.d`. When all such scripts are completed, session information is returned to the client. The client proceeds by contacting the agent on which the session is running.
- If the existing session was not alive, or if there were no existing session at all, VSM Server finds out which VSM Agent has the least load, and contacts this agent to request a new session.
- When the agent responds that a new session has been created, VSM Server runs any scripts placed in `/opt/thinlinc/etc/sessionstartup.d`. When all such scripts are completed, session information is sent back to the client. The client proceeds by contacting the agent on which the session was started.

14.4.1.1. Scripts run at session startup/reconnect

Scripts in `/opt/thinlinc/etc/sessionstartup.d` and `/opt/thinlinc/etc/sessionreconnect.d` are run by the `root` user, on the VSM Server. Session information will not be sent back to the client until these scripts have completed. This makes it possible to ensure that commands have been run before the client connects to the VSM Agent.

If background execution is desired, place the command to be run in the background and make sure all file descriptors are closed. Here's an example on how to execute a script in the background.

```
${TLPREFIX}/sbin/tl-limit-printers < /dev/null > /dev/null 2>&1 &
```

14.4.2. Session startup on VSM Agent

Figure 14-2. Session Startup Procedure - on VSM Agent

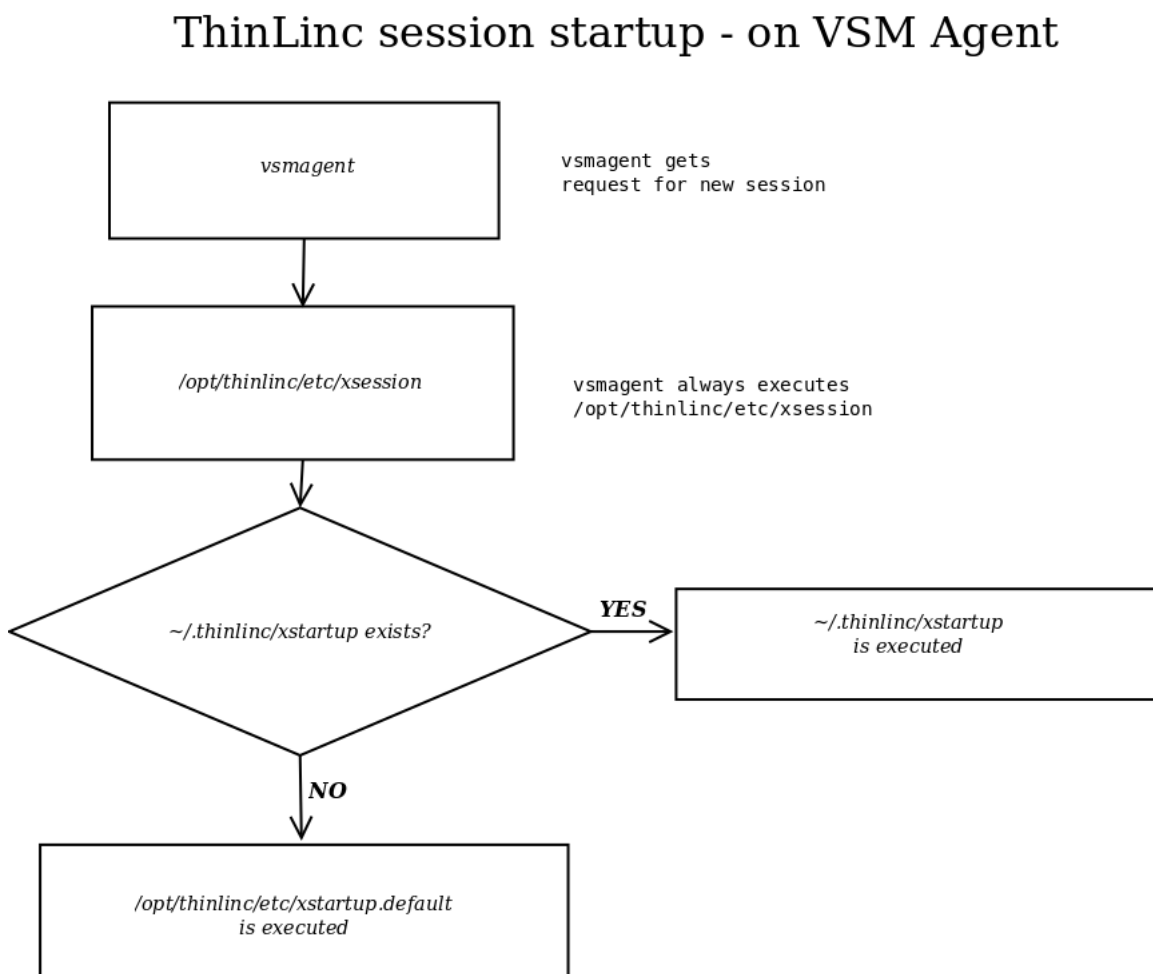


Figure 14-2 outlines what happens when an VSM Agent is contacted by VSM Server to request a new session. In detail, the following happens:

1. The VSM agent on the machine where the session will reside executes the script **/opt/thinlinc/etc/xsession**.
2. The file **/opt/thinlinc/etc/xsession** is a shell script that can be customized by advanced users. In its standard version, as delivered with ThinLinc, it will check if there is a file named **~/.thinlinc/xstartup** in the user's home directory. If there is such a file, it will be executed. If no such file exists, the file **/opt/thinlinc/etc/xstartup.default** is executed instead. See Section 14.4.3 for a description of the standard behaviour of this file.

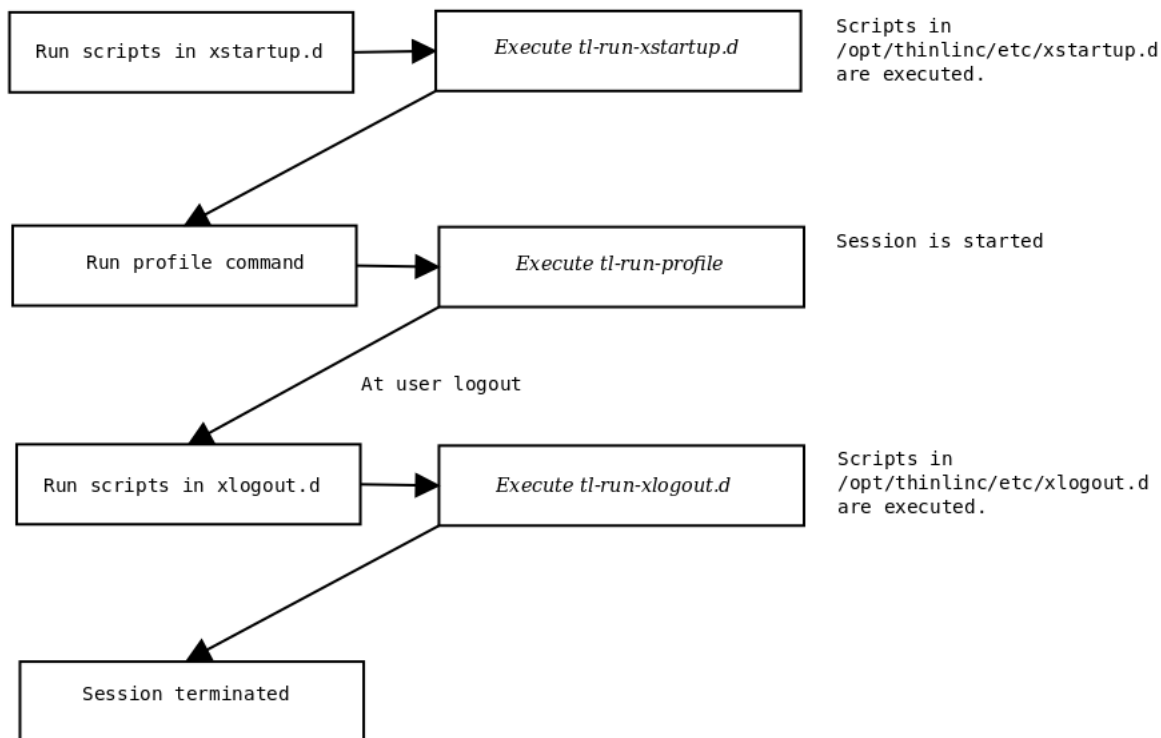
This system allows for experienced users to customize how their session startup should work by editing the file **~/.thinlinc/xstartup**. On the other hand, at sites where users should not be able to customize their system startup, **/opt/thinlinc/etc/xsession** can be modified so that it doesn't try to execute user-specific xstartup-files. The standard setup should however suit the needs of the majority of installations.

14.4.3. Profiles and the standard xstartup.default file.

ThinLinc allows for different "profiles" when starting up a user session. The users will be presented with a menu after logging in, where they can choose for example between a desktop suited for engineering users, a desktop suited for the marketing department or a Windows desktop. The example configuration files that are delivered with ThinLinc have several different alternatives, however only those sessions that are actually available on the system are displayed. This is just an example configuration, meant to be customized for each individual ThinLinc installation.

Figure 14-3. The ThinLinc profiles and xstartup.default

ThinLinc profiles and xstartup.default



As described in Section 14.4.1, **/opt/thinlinc/etc/xstartup.default** is executed if there is no **~/thinlinc/xstartup** for the user. This file, in its unmodified version as delivered with ThinLinc, executes three steps, as outlined in Figure 14-3.

1. The command **tl-run-xstartup.d** is executed, which causes all files in **/opt/thinlinc/etc/xstartup.d/** to be executed. Files that have filenames ending with **.sh** will be *sourced* as shell scripts. Other files are executed normally. This way, environment variables that persist down to the session command can be set in ***.sh** files.

If a specific execution order is needed for the scripts in the **xstartup.d/** directory, let the names of the scripts begin with numbers, where a script with a lower number will be executed before one with a higher number. For example **10setuphomedir** will be executed before **20copyfiles**.

By default, the script **/opt/thinlinc/etc/xstartup.d/20-tl-select-profile.sh** will invoke **tl-select-profile**, to let the user choose among the possible profiles. See Section 14.4.5 for documentation on how to setup profiles. If only one profile is available, **tl-select-profile** will select it without asking the user. The environment variable **TLPROFILE** is set to the name of the selected profile.

Worth noting is that the environment variable `TLPROFILE` is available when running the scripts in `xstartup.d`, for decisions based on what profile will be run.

2. The command **tl-run-profile** is run. This runs the commands associated with the selected profile, for example **startkde** to start a KDE session.
3. When the commands run by **tl-run-profile** exits, **xstartup.default** runs the command **tl-run-xlogout.d** which runs scripts and commands located in the directory `/opt/thinlinc/etc/xlogout.d`. The same information that applies to files in `xstartup.d` (as documented in 1) applies to files in this directory.

Note: Scripts in `/opt/thinlinc/etc/xstartup.d` and `/opt/thinlinc/etc/xlogout.d` are run *on the agent*, with the same rights as the user owning the session.

14.4.4. Session Startup with a Client Supplied Start Program

If the client has requested that the session should be started with a command supplied by the client, VSM agent will set the environment variable `TLCOMMAND` to this command. In this case, the profile selection dialog will be disabled and **tl-run-profile** will execute the command specified by the client, instead of a profile command. To disable client supplied start programs, create the file `/opt/thinlinc/etc/xstartup.d/00-no-startprog.sh`, containing:

```
unset TLCOMMAND
```

14.4.5. Configuring available profiles

The profiles choosable via the **tl-select-profile** command are configured via Hiveconf, under the `/profiles` path. The default configuration includes a number of examples.

If the `default` parameter is present, it specifies the default profile. The profile chooser will have this entry selected when it starts, and it may also be used automatically for some error conditions.

The `order` parameter selects which profiles should be available for selection, and the order in which they are displayed. This is a space-separated list.

If the `show_intro` parameter is true, a configurable introduction text will be displayed and requires user input to proceed with the logon process. The `introduction` parameter is a text that will be displayed if introduction is shown, this text block does also supports Pango Markup format styling for a fancier text layout.

Each profile is defined under a section named `/profiles/<profile key>`. It has the following fields:

`xdg_session`

Connects this ThinLinc profile with a system desktop session configuration. The directories `/etc/X11/sessions` and `/usr/share/xsessions` will be searched for a file matching `<xdg_session>.desktop`. It is recommended that this field is used for all modern desktop environments as it sets up important environment variables.

The fields *name*, *description*, *icon*, *cmdline* and *testcmd* will all be implicitly filled in by the system configuration. You can override those values individually by specifying a different value in the ThinLinc configuration.

Multiple values can be specified in this field, separated by spaces. The first matching configuration will be used. If no matching configuration can be found then the profile will not be shown.

Note: If the configuration is listed in `/etc/upstart-xsessions` then the specified command is ignored and an Upstart user session will be started instead. A manually specified *cmdline* can still be used to override the command.

name

A short description of the profile, shown in the profile list.

description

A longer description, shown under the screen shot when the profile is selected.

icon

A 22x22 image shown next to the name in the profile list. Paths can be absolute or relative
`/opt/thinlinc/share/tl-select-profile.`

screenshot

A 200x150 image shown when the profile is selected. Paths can be absolute or relative
`/opt/thinlinc/share/tl-select-profile.`

cmdline

The command to execute if this profile has been chosen.

If *xdg_session* is set then the environment variable `XDG_EXEC` will be set to the original command line from the system desktop session configuration.

testcmd

A shell expression or command that is executed to determine if this profile should be visible or not. A non-zero return code causes the entry to be hidden. If this field is empty or missing then the entry will always be shown.

If *xdg_session* is set then the environment variable `XDG_TRY_EXEC` will be set to the expected binary from the system desktop session configuration. Note that this value differs in behaviour from *testcmd*. `XDG_TRY_EXEC` should only name a executable binary in `PATH`, whilst *testcmd* will be executed and its return code inspected.

The *name* and *description* variables mentioned above also support Pango Markup (<https://developer.gnome.org/pango/stable/PangoMarkupFormat.html>) format styling which provides a simple way to display formatted text (for example, bold or italicized words).

The following example changes the name to be displayed using a blue foreground color and a large font, and the description with a boldface tag to emphasise the desktop environment name:

```
/profiles/gnome
name=<span foreground="blue" size="x-large">GNOME Desktop</span>
description=This is a standard <b>GNOME</b> desktop.
xdg_session=gnome
```

14.4.6. Configuring different Linux Desktops based on the selected profile

Please read Chapter 18 for documentation on how to configure different desktops with for example different menu and desktop icons depending on what profile were selected.

14.4.7. Speeding up Session Startup

If a user has a complicated session startup with many time-consuming operations, it can take quite a while before the user's desktop environment (for example KDE or Gnome) begins to start. Prime examples of when this happens is when mounting local drives, or when mounting some shared directories from a Netware server.

One way of speeding up this process is to execute some of the operations in the background. Most often, there is no need to mount the local drives before starting KDE, because it takes longer time to start KDE than it takes to mount the local drives. The two operations can easily run in parallel. The same goes for the example of mounting shared directories.

The easiest way to accomplish this is to add an & sign after commands run by scripts in `/opt/thinlinc/etc/xstartup.d`.

Make sure that commands that must be run before starting the window environment are run sequentially. For example, configuring desktops via TLDC must be done before starting KDE.

14.4.8. Configuring the language environment on the server based on the client language

The ThinLinc client reports the language settings on the client side when requesting a session. This can be used to configure the language on the server side. The idea is that in an environment where several languages are in use, a user could automatically get their preferred language based on what their client computer is configured for.

To activate this, a symlink needs to be created:

```
# ln -s /opt/thinlinc/libexec/tl-set-clientlang.sh /opt/thinlinc/etc/xstartup.d/00-tl-set-clientlang.sh
```

Also, make sure no other parts of the startup environment are trying to set the LANG variable. For example, on Fedora, the files `/etc/profile.d/lang.sh` and `/etc/profile/lang.csh` will override the LANG variable set by `tl-set-clientlang.sh`.

14.5. Limiting Lifetime of ThinLinc Sessions

The Xserver has three options which controls the maximum lifetime of ThinLinc sessions. These are described below, and can be added to the parameter `/vsmagent/xserver_args`.

- `-MaxDisconnectionTime s`

Terminate when no client has been connected for `s` seconds. Note: Never use a value smaller than 60.

- `-MaxConnectionTime s`

Terminate when a client has been connected for `s` seconds

- `-MaxIdleTime s`

Terminate after `s` seconds of user inactivity. Note: Never use a value smaller than 60.

In addition to the options above which control the lifetime of the ThinLinc session, the option `-IdleTimeout` can be used to configure how long an idle session should remain connected. The `-IdleTimeout` option takes a number of seconds as an argument, and can be added to the parameter `/vsmagent/xserver_args` as per the options described above.

Note: Setting `-IdleTimeout s` will simply disconnect the client from the session after `s` seconds; it will not terminate the ThinLinc session itself.

Chapter 15. Shadowing

15.1. Introduction

Shadowing is a feature that lets a user connect to, view, and interact with ThinLinc sessions of other users. This can be useful in remote assistance and support scenarios, where trusted support personnel can connect to a user session and aid with for example application problems.

Because shadowing gives the shadowing user full control over the shadowed session, this feature should be used with caution.

The shadowing feature is enabled by default and is configured to ask the user to accept or reject a shadowing request.

15.2. Disable shadowing feature

The shadowing feature is enabled by default when installing ThinLinc. You can disable this feature if required, using the following command.

```
# sudo tl-config /shadowing/shadowing_mode=reject
```

When the shadowing feature is disabled, all requests to shadow a user session is actively rejected. Details about this configuration parameter is described in */shadowing/shadowing_mode*.

Note: The above command should be run on all of the ThinLinc servers in your cluster.

15.3. Granting shadowing access to users

Because of the security implications of this feature, the system administrator needs to grant this permission to named users and/or groups before it can be used.

The `vsmserver` service controls whether a user requesting to shadow another user is authorized to do so. The configuration parameter `/shadowing/allowed_shadowers` from the `/opt/thinlinc/etc/conf.d/shadowing.hconf` file is read by the `vsmserver` service on startup. This parameter is described in detail in */shadowing/allowed_shadowers*.

Note: After the configuration variable has been set, the `vsmserver` service needs to be restarted before the change is made active.

15.4. Shadowing notification

Notification behaviour of the shadowing feature is configured by the system administrator. The notification mechanism can be configured in four different modes as described here.

- Shadow requests are silently rejected
- Shadow requests are silently accepted
- Shadow requests are accepted and the user is notified
- Shadow requests are interactively accepted or rejected by the user

To configure the shadowing mode use the following command and select a value of; reject, silent, notify or ask. Details about this configuration parameter is described in `/shadowing/shadowing_mode`.

```
# sudo tl-config /shadowing/shadowing_mode=ask
```

Note: The above command should be run on all ThinLinc servers in your cluster.

Note: Only newly started session are affected by the above change.

15.5. Shadowing a user session

The ThinLinc client must be configured for shadowing. See Section 7.4.1 for more information.

Once the client has been configured for shadowing, enter the username of the user you wish to shadow in the User to shadow field and connect.

Chapter 16. Hiveconf

16.1. Overview

Hiveconf is the name of the configuration system used in ThinLinc. It is however not a ThinLinc-specific configuration system, but instead a generic configuration framework for storing key/value pairs in a human readable way, although still in a format that's easy to read and modify from a computer program.

Hiveconf stores data using a "backend", meaning configuration data can be stored in different ways. The default backend which is also used in ThinLinc is using a text file format similar to Windows `.INI`-files, or the format used in `smb.conf` from Samba.

In this section, we will describe Hiveconf from a general point of view and also describe ThinLinc-specific details.

16.1.1. Basic Syntax

Basically, a Hiveconf file consists of key/value pairs with a equalsign(=) between them, as in the following example:

```
vsm_server_port = 9000
vnc_port_base = 5900
```

The values after the equal sign can be of the following types:

- String
- Boolean
- Integer
- Float
- Binary data as hexadecimal ASCII

Data can also be lists of the above types, these lists are space-separated.

16.1.2. Tree Structure

Parameters in Hiveconf all reside in folders. Folders are just like a directory or folder in a normal file system. By adding folder directives to Hiveconf files, the parameters will be split up in a tree structure, meaning each parameter will be addressed using a path. This way, two folders can have two parameters with the same name without collision.

The benefits of this is that a software suite (for instance ThinLinc) can have one common configuration namespace, without having to name all configuration parameters uniquely, since every component in the suite can have its own namespace. In ThinLinc, the VSM server has its parameters in the `vsmserver/` folder, the VSM agent has its parameters in the `vsmagent/` folder and so on.

Looking from a system global point of view, every software package has its own folder, meaning *all* configuration parameters of the system can be accessed using a common tool.

Folders are put into the configuration files by adding a path inside square brackets to the file as in the following example:

```
[root@tlha-master conf.d]# cat vsmserver.hconf
#
# Hiveconf configuration file - VSM server
#
[/vsmserver]
unbind_ports_at_login=true

# Administrators email
admin_email = root@localhost
```

In this example, both parameters (*unbind_ports_at_login* and *admin_email*) reside in the */vsmserver* folder. This means that they should be addressed as */vsmserver/unbind_ports_at_login*, */vsmserver/admin_email* respectively if used from inside a program using the Hivetool libraries. This is of course not that important from the system administrator's point of view, but it's important to understand the principle.

16.1.3. Mounting Datasources

One Hiveconf file can use another Hiveconf file by mounting the other file using a mount command as in the following example:

```
%mount HA.hconf
```

The mount should be compared to a mount on a Linux. That is, the mount adds the tree structure of the file mounted at exactly the place in the current tree structure where the mount command was found.

Mounts can also use wildcards, as in the following example

```
%mount conf.d/*.hconf
```

The above is exactly what you'll find if you look into the file */opt/thinlinc/etc/thinlinc.hconf*. Hiveconf will mount all files in */opt/thinlinc/etc/conf.d* and add them to the current folder. This is a very convenient way to add all configuration files for a specific software suite to the Hiveconf namespace.

16.1.4. Hostwide Configuration

As we hinted in Section 16.1.2, Hiveconf lays the foundation for a hostwide configuration system where all applications on a host can be configured using a single system with a common API. This can be accomplished because each application will get its own subfolder in the hostwide configuration folder, so that two applications parameters won't collide even if they have the same name. Using the mount command, every application can have its own configuration file, while still exporting its parameters to the hostwide folder system.

There is a hostwide Hiveconf "root", implemented by the file `/etc/root.hconf`. This file mounts all files in `/etc/hiveconf.d/` where an application can drop its own hiveconf file at install-time, just like it can drop a file in for example `/etc/logrotate.d` or `/etc/profile.d`.

16.1.5. Hiveconf Tools

In addition to the system libraries used by applications to read and write configuration parameters that reside in Hiveconf files, there is a command line utility named **hivetool** for inspecting and setting parameters from the command line. This can be very convenient, for example when scripting setup scripts that need to set some parameter.

Hivetool without parameters will do nothing. To see all parameters on the system, run:

```
# hivetool -Ra /
```

which instructs **hivetool** to print all parameters, beginning from the root (/) and recursing downwards. With a standard Hiveconf installation this will list Samba and KDE configuration parameters. If ThinLinc is installed, it will list ThinLinc parameters as well.

To print a specific parameter, run **hivetool** with the name of the parameter as parameter. For example:

```
[root@tlha-primary etc]# hivetool /thinlinc/vmsserver/admin_email
root@localhost
```

Setting a parameter is equally easy. To set the `admin_email` parameter above, execute the following:

```
# hivetool /thinlinc/vmsserver/admin_email=johndoe@example.com
```

16.2. Hiveconf and ThinLinc

ThinLinc uses Hiveconf as its primary configuration system on the serverside. In this section, we will describe the convenience utility shipped with ThinLinc. For descriptions of the folders and parameters used by ThinLinc, please refer to Chapter 14

16.2.1. The ThinLinc Configuration Tool - tl-config

In order to access the ThinLinc part of the Hiveconf configuration namespace without having to address it using the hostwide path (i.e. to avoid having to add `/thinlinc/` to all parameters, a tool named **tl-config** is shipped with ThinLinc.

tl-config takes the same parameters as **hivetool** and works the same way. Refer to Section 16.1.5 for information about **hivetool**. Try for example

```
# tl-config -Ra
/
```

a command that will print all ThinLinc-related parameters.

Chapter 17. Administration of ThinLinc using the Web Administration Interface

17.1. Introduction

This chapter describes the web-based ThinLinc administration interface called `tlwebadm`. This administration interface is installed automatically by the ThinLinc installation program, and may be accessed by pointing your web browser to `https://<hostname>:1010`.

For information on configuring `tlwebadm`, for example setting a password or changing the default port, see Section 14.2.7.

Note: The password must be set and the `tlwebadm` service restarted before use.

17.2. Modules

The `tlwebadm` interface consists of several modules which address different aspects of ThinLinc configuration:

- *System Health*, for viewing information about ThinLinc master and agent services, and testing user or group lookup performance. See Section 17.2.1 below.
- *Status*, for viewing information such as license usage, server load and sessions. See Section 17.2.2 below.
- *VSM*, for viewing information and managing ThinLinc subclusters, the master service and the agent service. See Section 17.2.3 below.
- *Profiles*, for viewing and configuring profiles. See Section 17.2.4 below.
- *Locations*, for viewing and configuring printers and terminal locations. See Section 17.2.5 below.
- *Desktop Customizer*, for configuring desktops. See Section 17.2.6 below.
- *Documentation Center*, a module containing documentation and other useful information.

These modules are described in more detail in the following sections.

17.2.1. The System Health Module

The System Health module allows you to check the running state of the ThinLinc services VSM Master and VSM Agent. There is also a tool to perform a user or group lookup which reports the time for the lookup.

- *VSM Master* reports current running state of VSM Master service.
- *VSM Agent* reports current running state of VSM Agent service.

- *Users and Group Lookup* allows you to test performance of user and group lookup on the system. Fill in values for **username** and/or **group** and click the **Test user and group lookup** button to perform a lookup.

17.2.2. The Status Module

The Status module allows you to view or manipulate the following aspects of ThinLinc, by selecting the relevant submenu:

- *Licenses* allows you to view current and historic license usage, as well as the current number of licenses.
- *Load* allows you to check the current server load on the ThinLinc agents.
- *Sessions* allows you to terminate, shadow or view details of sessions. This feature is described in more detail in the next section, Section 17.2.2.1

17.2.2.1. The Sessions Menu

When you select the sessions menu, a table with all currently active users is displayed. To perform additional tasks, click on the corresponding user name. This will bring up the session details page, which displays all the session parameters for each session the user has running. The information table is described below.

- **Agent Server** : The DNS host name of the server that is hosting this session. If you only have one ThinLinc server, this server will host all sessions. If you have several ThinLinc servers in a cluster, new sessions will be created on the server with the lightest load.
- **Display Number** : Each session on a certain server has a unique number, the X Window System display number. Display zero is reserved, and never used for ThinLinc sessions.
- **Terminal ID** : An identification of the thin terminal. This is the terminal's ethernet hardware address (MAC address).
- **Framebuffer Size** : The size (resolution) of the active session.
- **Local Screen Size** : The size (resolution) of the terminal's screen.
- **Session process ID** : The PID (process identification number) of the tl-session process, which is the parent for all processes belonging to a certain session.
- **Command** : The command line that was specified when starting this session. This is usually empty for full desktop sessions.

Below each table, there are two buttons:

- **Terminate Session** : By clicking this button, you can terminate a session immediately. Caution: This can lead to data loss, since applications running on the ThinLinc servers may not hold unsaved data.

- **Shadow Session** : This button will generate a ThinLinc "launch file" (see Section 7.10) that starts the native ThinLinc client, preconfigured to shadow the current user.

Note: The user will not be informed that shadowing is in progress, unless `tl-shadow-notify` is enabled.

17.2.3. The VSM Module

VSM contains information about subclusters and VSM Master and Agent services. This module also allows managing of subclusters and configuration of Master and Agent.

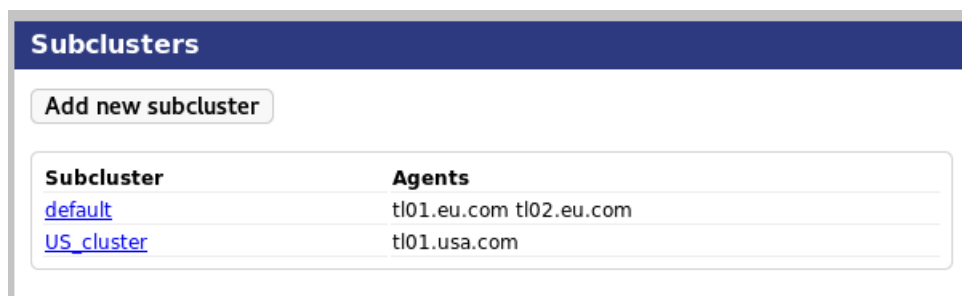
- *Subclusters* allows viewing, adding, modifying or deleting subclusters.
- *VSM Master* allows starting or stopping the service, and modifying a subset of the configuration options.
- *VSM Agent* allows starting or stopping the service, and modifying a subset of the configuration options.

17.2.3.1. Subclusters

On this page, the configured subclusters of the ThinLinc Master can be found. Subclusters are groups of agents that can be associated with users or groups. Adding, modifying or deleting existing subclusters is possible. A restart of the `vsmserver` service is required after changes to subcluster configuration in order for the changes to take effect. The `vsmserver` service can be restarted from the *VSM Master*-page described in Section 17.2.3.2.

The page will present a list of currently configured subclusters. This should be something like the example in Figure 17-1. To edit a subcluster, click on its name in the list.

Figure 17-1. Subclusters



Subcluster	Agents
default	tl01.eu.com tl02.eu.com
US_cluster	tl01.usa.com

Figure 17-1 shows a system with a total of two subclusters. The subcluster called *default* is configured with two agent servers and *US_cluster* is configured with one agent.

To add a new subcluster to the list, press the *Add new subcluster* button. This will bring up an empty subcluster edit form. See figure Figure 17-2 for an example.

Figure 17-2. New subcluster form

The screenshot shows a web interface titled "Subclusters". At the top, there is a button labeled "Add new subcluster". Below this is a form for editing a subcluster. The form has two main sections: "Subcluster" and "Agents". The "Subcluster" section contains a text input field with the value "NewSubcluster". The "Agents" section contains three text input fields, labeled "Agents", "Users", and "Groups". Below these fields are two buttons: "Delete" and "Save". There is also a checkbox labeled "Yes, really delete NewSubcluster". At the bottom of the form, there is a table listing existing subclusters.

Subcluster	Agents
default	tl01.eu.com tl02.eu.com
US_cluster	tl01.usa.com

There are four editable fields in this view; Subcluster, Agents, Users and Groups. These are described in Section 14.2.3. To save changes, press the *Save* button. Remember to restart the *vsmserver* service afterwards.

17.2.3.2. VSM Master

On this page you can start or stop VSM Master service. There are also a subset of configuration options available for configuration of the service.

- *Service Status* gives you the ability to start or stop the VSM Master service.
- *Sessions per user* allows you to configure how many session are allowed per user.

A value of zero means no limit and will give unlimited sessions per user.

- *Allowed Groups* allows you to configure which groups should be given access to connect to ThinLinc. If no groups are specified, all users will have access to connect to ThinLinc
- *Allowed Shadows* allows you to configure which users should be able to shadow other ThinLinc sessions.

Click the **Save** button when you want to save your changes to the configuration files.

Note: You need to restart the service to apply your changes.

17.2.3.3. VSM Agent

On this page you can start or stop the VSM Master service. There are also a subset of configuration options available for configuration of the service.

- *Service Status* gives you the ability to start or stop the VSM Agent service.
- *Agent Hostname* allows you to configure the hostname that clients are redirected to.

Note: This configuration is needed if running ThinLinc in a NAT environment. See Section 3.3.4 for more information.

- *Extra Arguments to X Server* allows you to configure additional arguments to the Xserver (Xvnc) for new sessions that are started.

Click the **Save** button when you want to save your changes to the configuration files.

Note: You need to restart the service to apply your changes.

17.2.4. The Profiles Module

On this page you can modify text shown in the profile chooser, and manage profiles. You can create or delete a profile and configure the profile order.

- *Introduction Texts* allows you to modify and manage translation of texts used in the profile chooser.
- *Profile List* allows you to configure the available profiles and their order.

17.2.4.1. Introduction Texts

Introduction texts contains translation tables for greetings and introduction texts. There is also a configuration option to enable or disable the use of introduction texts.

- *Greeting Text* the text to show at the top of the profile chooser.
- *Show Introduction* disable or enable the introduction text which is shown to the user before the profile selection dialog.
- *Introduction Text* the text to show before presenting the list of profiles.

To add a new translation, fill in language code and the translated string. Click the **Save** button to save the new translation and add it to the translation table.

To delete a translation select the row using the checkbox in *Delete* column of the translation table. Click the **Save** button to carry out the actual deletion of selected rows.

17.2.4.2. Profile List

The Profile List module contains functionality to manage your profiles. You can change the default profile, or create new and edit existing profiles. You can also change the order of profiles.

- *Default Profile* allows you to specify the default profile to be selected in the profile chooser.
- *Profile List* allows you to modify profiles and their order, or create new profiles.

Create a new profile by clicking the **Add new profile** button. If you want to edit an existing profile, click the profile name in the table of available profiles.

When creating a new or editing an existing profile a form is displayed. This form is shared between both modes and each field details are described below.

- **Identification**

An unique string identifier for the profile which is used when referencing this profile.

- **XDG Session Desktop**

The system desktop session configuration that this profile should be connected to.

- **Default Name**

A name for the profile which is displayed in the profile chooser.

- **Availability**

This will make the profile available to be selected and used. If you uncheck it will not be shown to the user in the profile chooser.

- **Take Description From**

The description is shown in the profile chooser when a profile is selected. This field can be a static text which is defined in the input field **Default Description** .

- **Test command** : This will take and use the output of defined **Test Command** as description for the profile.

- **Manually defined text below** : This will use the text defined in the Default Description field below.
- **Default Description**
A text used as description for the profile. This is text is used if Take Description From above is selected to use the manually defined text.
- **Icon Path**
A filename of the icon to use in the profile chooser.
- **Screenshot Path**
A filename of the screenshot to use in the profile chooser.
- **Command Line**
This command is used to start up a session. In most cases this is something simple like **xfce4-session**, but in some cases there might also be arguments like **gnome-session --session gnome-classic**.
- **Test Command**
This command is evaluated and if it returns true, the profile is shown to the user. If the command evaluates as false, the profile will not be shown in the list of available profiles for the user.
ThinLinc includes the tool "tl-memberof-group" which may be used to test membership of groups. You can use this tool as test command, such as **\${TLPREFIX}/bin/tl-memberof-group my_profile_access_group**. This example will give members of group my_profile_access_group access to the profile.
If you only want to give a specific user access to the profile you may specify **test \${USER} = user**.

When you have filled out the form, or changed any fields, click **Save** button at the bottom of the form to save your changes into the configuration file.

To delete a profile click the profile name in table of available profiles. Then click the checkbox at the bottom of the form to verify your intention about deletion of the profile. Complete the deletion by clicking the **Delete** button.

17.2.5. The Locations Module

Locations contains information about locations where terminals and printers reside. A location can be a room, a floor, a house or some other type of geographical delimitation.

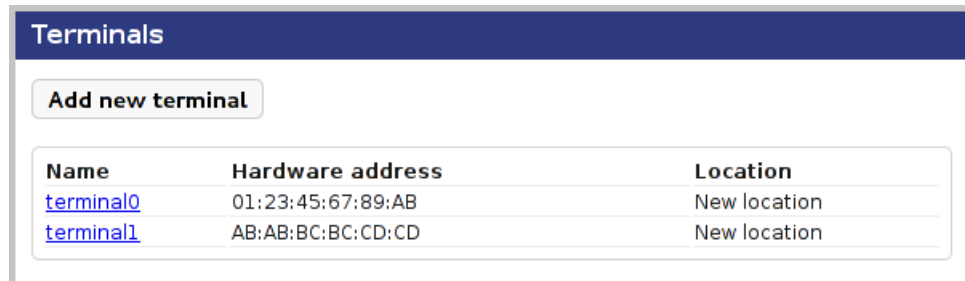
Every terminal should be assigned as a member of a location. In addition to terminals, printers may also be assigned to locations.

17.2.5.1. Terminals

Terminals contains necessary information about all terminals. The most important information here is every terminal's interface hardware (MAC) address.

Each terminal should be entered as described in this section. Enter the terminals module by clicking on the menu item. You will be presented with a list of currently entered terminals. This could be something like the example in Figure 17-3.

Figure 17-3. Terminals



Name	Hardware address	Location
terminal0	01:23:45:67:89:AB	New location
terminal1	AB:AB:BC:BC:CD:CD	New location

Figure 17-3 shows a system with a total of two terminals.

To edit a terminal, click on its name in the list.

To add a new terminal to the list you press the **Add new terminal** button. This will bring up an empty terminal edit form. See Figure 17-4 for an example.

Figure 17-4. New terminal form

The screenshot shows a web interface titled "Terminals". At the top, there is a button labeled "Add new terminal". Below this is a table with three columns: "Name", "Hardware address", and "Location". The table contains two entries: "terminal0" with hardware address "01:23:45:67:89:AB" and location "New location", and "terminal1" with hardware address "AB:AB:BC:BC:CD:CD" and location "New location". Below the table is a form for adding a new terminal. The form has four fields: "Terminal name (required)" with a text input containing "New terminal", "Hardware (MAC) address (required)" with an empty text input, "Location" with a dropdown menu showing "New location", and "Printers" with an "Add printer" button. At the bottom of the form are two buttons: "Delete" and "Save". Below the "Delete" button is a checkbox labeled "Yes, really delete New terminal".

There are three editable fields in this view, Hardware (MAC) address , Terminal name and Location . They are described in Table 17-1 below.

To save changes, press the **Save** button. When you have pressed the **Save** button you will see that the Hardware (MAC) address field will change from being an editable field to become a static text label. To assure data integrity between the modules you aren't allowed to change this field after it's added.

When a new terminal is saved or when clicking an existing in the terminals list, there will be three buttons inside the form. The **Save** button is used to save changes made to the terminal. The **Delete** button deletes the currently viewed terminal. The **Add Printer** button will add a new printer field to the form.

Table 17-1. Terminal properties

Name	Description
Hardware (MAC) address	hardware (MAC) address of the main interface of the terminal. This field is important! Without a correct value the <i>nearest printer</i> won't work!
Terminal name	Name of the terminal. Could for example be the terminal's DNS name or a name following a naming scheme that identifies the terminal.

Name	Description
Location	Which of the locations, entered in the Locations module, this terminal belongs to.

It is also possible to add a printer to a terminal in the terminal module. This can be used if a terminal has its own printer or is closer to another printer than the ones assigned to this terminal's location. You should use this feature moderately since it may cause more administration.

To add a printer you do exactly as in the Locations menu. Click the **Add printer** button, select the printer in the pop-up menu and then press **Save** to make sure that the settings are stored. To delete it, check the relevant Delete checkbox(es) for the printer(s) you wish to remove, and click **Save**.

17.2.5.2. Locations

To edit a location, click on its name in the list.

To add a new location to the list you press the **Add location** button. This will bring up an empty location edit form. See Figure 17-5 for an example.

Figure 17-5. New Location Form

The screenshot shows a web interface titled "Locations". At the top, there is a button labeled "Add new location". Below this is a table with three columns: "Name", "Description", and "Use for Unknown Terminals". The table contains one row with the text "New location" under "Name" and "No" under "Use for Unknown Terminals". Below the table is a form for editing the selected location. This form has several sections: "Location name (required)" with a text input field containing "New Location"; "Description" with an empty text input field; "Unknown Terminals" with a checkbox labeled "Use this location for unknown terminals"; "Printers" with an "Add printer" button; and a "Delete" section with a "Delete" button and a checkbox labeled "Yes, really delete this location". A "Save" button is also present at the bottom right of the form.

Fill the Name and Description fields with relevant information. Check the checkbox if this location is to be used for unknown terminals when using the printer access control feature (see Section 5.5 for details).

To save changes, press the **Save** button. When you have pressed the **Save** button you will see that the **Name** field will change from being an editable field to become a static text label. To assure data integrity between the modules you aren't allowed to change the name of an item after it's added.

The **Delete** button deletes the currently viewed location, but only if the confirmation checkbox is also checked. The **Add Printer** button will add a new field to the form, a drop-down menu with all possible printers. An example of this can be seen in Figure 17-6.

Figure 17-6. Location Details With Printer

The screenshot shows a web interface titled "Locations". At the top, there is a button labeled "Add new location". Below this is a table with three columns: "Name", "Description", and "Use for Unknown Terminals". The first row in the table is labeled "New location" and has "No" in the "Use for Unknown Terminals" column. Below the table is a form for adding a new location. The form has three main sections: "Location name (required)" with a text input field containing "New location"; "Description" with a text input field; and "Unknown Terminals" with a checkbox labeled "Use this location for unknown terminals". Below these is a "Printers" section with an "Add printer" button and a dropdown menu showing "MX-2700N". To the right of the dropdown is a checkbox labeled "Delete". At the bottom of the form are two buttons: "Delete" and "Save". Below the "Delete" button is a checkbox labeled "Yes, really delete this location".

The **Printer** field above has the printer *europa* selected. Remember to save the changes if you change printer! You can assign more printers to this location by clicking **Add printer** again to bring up another printer line. To remove a printer you select the **Delete** checkbox corresponding to the printer(s) you want to delete, and click **Save** to apply the changes. The printer(s) will disappear.

Note: Printers contains entries for all available printers. These entries are just shadows of the real CUPS (Common Unix Printing System) printer system entries. This means that you first need the printers to be installed in the CUPS printer system and then added to this list.

17.2.6. The Desktop Customizer Module

The ThinLinc Desktop Customizer is described more fully in its own chapter, Chapter 18. Links to sections of this chapter pertaining to the respective menus of the Desktop Customizer Module are provided below for convenience.

17.2.6.1. Application Groups

For information on configuring Application Groups using TLDC, see Section 18.2.5

17.2.6.2. Applications (Manual)

For information on configuring Manual Applications using TLDC, see Section 18.2.3

17.2.6.3. Applications (System)

For information on configuring System Applications using TLDC, see Section 18.2.1.1

17.2.6.4. Menu Structure

For information on configuring Menu Structures using TLDC, see Section 18.2.4

Chapter 18. Building Custom Linux Desktops with the ThinLinc Desktop Customizer

In this chapter, we will document how to create custom desktops for ThinLinc users using either the K Desktop Environment (<https://www.kde.org>) or the Gnome Desktop Environment (<https://www.gnome.org>), in combination with the ThinLinc Desktop Customizer (TLDC).

The TLDC's core functionality is to build the menu of ThinLinc users based on factors such as group membership, user name and ThinLinc profile. It can also add icons to the desktop of each user, based on the same premises.

18.1. Introduction

The ThinLinc Desktop Customizer is a combination of a web-based administration tool and a command that is run at session startup for all users. It enables the administrator to decide what menu entries should be presented for specific users, and what icons should be made available on the desktop. Which menu entries and/or desktop entries are given to a specific user is decided based on the Linux group memberships of the user, the username and what ThinLinc profile was chosen (if any).

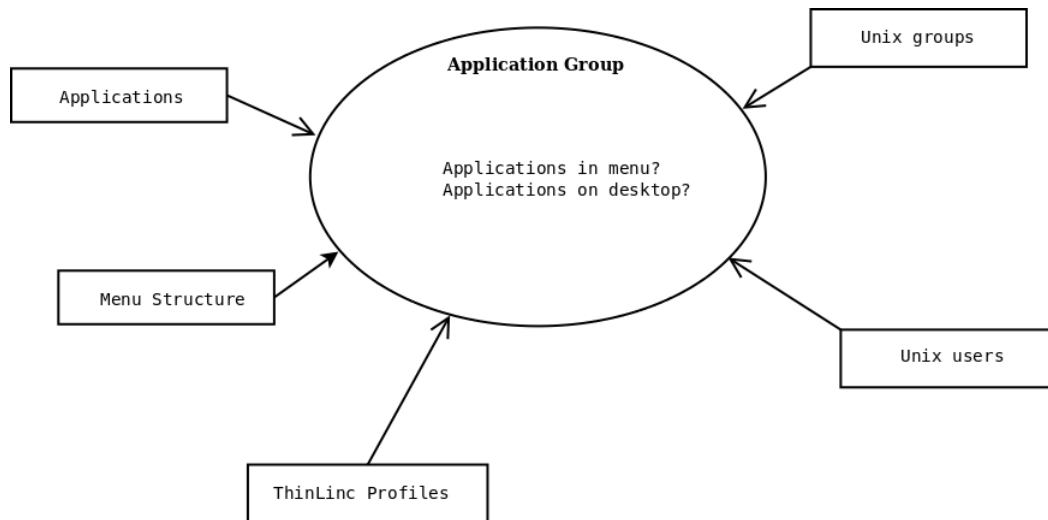
Note: Since KDE4 uses a different desktop configuration mechanism to previous versions of KDE, icons added to the desktop using the ThinLinc Desktop Customizer will not be shown in KDE4. This can be solved by changing the "Desktop layout" setting of your KDE4 desktop to "Folder view".

18.2. Using the ThinLinc Desktop Customizer

Using the ThinLinc Desktop Customizer, the system administrator can decide what applications should be available in the menu and/or on the desktop for specific users or for users that are members of some Linux group. The ThinLinc Desktop Customizer is configured via a web interface, part of the ThinLinc Web Administration. Chapter 17 describes the interface in general, this section will describe the Desktop Customizer part of it.

18.2.1. Concepts

Figure 18-1. ThinLinc Desktop Customizer Concepts



The main concept in the ThinLinc Desktop Customizer is the *Application Group*. The Application Group combines data about applications, a menu structure, Linux groups and users, and some other settings.

18.2.1.1. Applications

The Applications referred to in the Application Groups are found by scanning the space-separated list of directories defined in the Hiveconf parameter `/utils/tl-desktop-customizer/xdg_data_dirs` for files named `*.desktop`. The files are read according to the Freedesktop.org (<https://www.freedesktop.org/>) Desktop Menu Specification. The TLDC scans the directories in the same way that KDE will do when building the menu.

Some applications are marked by the system to be shown only for root, or only in either Gnome or KDE. On SuSE, there is also a `X-SuSE-Unimportant` parameter in some `*.desktop` files, which will make the KDE packaged with SUSE hide the application. TLDC handles this by adding a comment to the application in the applications listing, and in the selectboxes used when creating application groups.

In addition to the desktop files automatically found, it is also possible to manually define applications. This is needed for example when an application without a `*.desktop`-file has been installed or when an application has been installed in a non-standard location.

18.2.1.2. Menu Structure

Each Application Group can add applications to a specific place in the menu structure. The available menu structure is edited in the "Menu structure" part of the web based administration interface.

18.2.1.3. Linux Groups and Users

An Application Group is used by zero or more Linux groups and by zero or more specific users. An example would be an educational environment. Let's say that all pupils attending the class "biology 4" are members of the Linux group "bio4". By creating an Application Group named "Biology 4" with all applications that are specific to the biology class, and then adding the "bio4" Linux group as one of the groups that should be assigned the "Biology 4" Application Group, all students attending the class will automatically get the applications specific to the biology class in their menu. By adding the teacher of the class as a specific user, he/she as well will also get access to the applications.

18.2.2. Using the ThinLinc Desktop Customizer

The Daily use of the TLDC consists of one or several of the following steps:

- Create an Application
- Create a folder in the Menu Structure
- Bind one or several applications to a folder in the menu structure, using an Application Group

In the following sections, we will more thoroughly describe the different actions that may be needed.

18.2.3. Handling Applications

The handling of applications is normally the first step in using the TLDC. Click on the "Applications (Manual)" link in the TLDC, and you will enter a view where the applications you've defined manually are listed. Several example applications are included with ThinLinc at installation. By clicking on the text "Applications defined by system", you can also see what applications are found automatically by scanning, as described in Section 18.2.1.1.

If the application you want to add to a menu or to the desktop is not found among "Applications defined by system", you need to define it manually. This is the case for applications installed without adding a .desktop file in the correct location.

Defining applications manually is done by clicking on the button "Add new application" (located at the top of the list of applications). This leads to a page where you can define the following properties for the new application:

- *Default Application Name*

This is the name of the application. It's the name that is written next to the icon (if any), in the menu, and under the icon if the application is to be added to the desktop.

The Default Application Name is used if there is no name defined for the language in use when the application is shown, or if the language is english.

- *Application Name (<language-code>)*

This is the name of the application in the language with the RFC1766 language code <language-code>. This name is shown if the locale is set to that language when the menu or desktop is shown.

The languages that should be configurable are set by editing the space-separated list in the parameter `/utils/tl-desktop-customizer/desktop_languages`. The default value of this parameter is `sv`, which means that the TLDC will allow you to set the default name and the name in Swedish.

- *Command*

This specifies the command to run to start the application. Enter the path to the command followed by any arguments in the *Command* box. The input box follows Bourne shell syntax rules.

Example:

```
"/usr/bin/my program" --fullscreen --title "My title"
```

- *Path to Icon file*

The filename of the icon for the application. If the icon is available in one of the directories where KDE automatically looks for icons, just the filename without the extension can be given. Otherwise, the complete path must be specified.

- *Command Startup Feedback*

Check the box to instruct the Window Manager to show a special icon while the command is starting. Note that not all applications support this functionality.

Press save when done filling the fields. The application will now show up among the other manually defined applications.

If you want to, you can add the application directly to an existing application group by checking the checkbox in front of the application name, then selecting the application group and if the application should be added to the menu or desktop of this application group, in the form at the top of the page. This can be done for both manually defined, and automatically found applications.

18.2.4. Defining a Menu Structure

With TLDC, the normal menu structure as defined by the Linux Distributor or by the KDE team is not used. Instead, a new menu structure is defined. This gives more flexibility in designing menus. The TLDC administrator can fully decide where in the menu structure a certain application is placed.

To define the menu structure, click on the "Menu structure" submenu in the left pane of the TLDC administration interface. This leads to a view where a menu structure can be defined. The Root menu folder is always available and can't be removed.

Note: A menu called "*Hidden Menu*" is shipped with the default ThinLinc configuration. See Section 18.4.3 for an explanation of its functionality. Please don't remove it if you are planning to use KDE.

The following properties can be edited for a menu:

- *Default Menu Name*

This is the name of the menu, as it will be shown in the menu.

- *Menu Name (<language-code>)*

This is the name of the menu in the language with the RFC1766 language code <language-code>. This name is shown if the locale is set to the language at runtime.

- *Path to Icon File*

The filename of the Icon for the menu, shown to the left of the menu name in the KDE menu. If the icon is available in one of the directories where KDE automatically looks for icons, just the filename without the extension can be given. Otherwise, the complete path must be specified.

- *Hide This Menu*

If this radio button is set to *Yes*, the menu will be a hidden menu. It will not be shown in the menu, but any applications that are added to this menu via an application group will be available in the KDE File Associations.

Just as for Applications, the name of the menu can be defined in several languages. The *Default Menu Name* is used if no language-specific name is defined, or if the locale specifies that the language is english. The list of languages that can be defined using the TLDC is found in `/utils/tl-desktop-customizer/desktop_languages`.

18.2.5. Defining Application Groups

Enter the "Applications Groups" part of the "Desktop Customizer". This will present you with a list of existing application groups and their settings.

Note: An application group called "*Hidden*" is shipped with the default ThinLinc configuration. See Section 18.4.3 for an explanation of its functionality. Please don't remove it if you are planning to use KDE.

Press the button "Add new group" (located at the top in the table of existing application groups) to create a new application group. This will open a rather large form, where you can define the following properties:

- *Name of the Application Group*

This is the name of the Application Group. This is not displayed to the users, but only to the System Administrator using the ThinLinc Desktop Customizer. Set to something that reflects the contents of the Application Group.

- *Applications Added to Menu*

First, define in the dropdown box what location in the menu structure applications chosen in the boxes below it should be added to.

Add to the left selectbox the applications that should appear in the menu folder selected above, for the users that are assigned this Application Group. The right selectbox lists the applications defined or

found installed on the system. If there are no applications available, you've forgotten to define Applications, as documented in Section 18.2.5.

- *Applications Added to Desktop*

Add to the left selectbox the applications that should appear as icons on the desktop of the users that are assigned this Application Group. Just as for Application added to the menu, only applications earlier defined, or automatically found, will show up as selectable.

- *Linux Groups with this Application Group*

This is where you connect Linux groups to Application Groups. If for example a specific school should be assigned this Application Group, and all the pupils of that school are members of the Linux group "school-1", add the Linux group "school-1" to the left selectbox. When logging in, the group memberships of each user is inspected to determine which Application Groups to assign to the user.

Note: If the mapping between the numerical group id and the group name doesn't work, the group is shown as `#<gid>`. This might be because the group has been removed from the system, or because the operating system has problems in the connection to the directory service used.

- *Specific Users with this Application Group*

This parameter allows you to decide that specific users should be assigned this Application Group as well, even if they are not a member of one of the groups that were added above. This way, for very specialized applications, no Linux group needs to be created. Another way of using this field would be that the teachers of a specific class could be added to the Application Group for that class, if the teachers are not part of the Linux group that is associated with the class.

Note: If the mapping between the numerical user id and the user name doesn't work, the user is shown as `#<uid>`. This might be because the user has been removed from the system, or because the operating system has problems in the connection to the directory service used.

- *ThinLinc profiles with this Application Group*

This setting allows you to connect the Application Group to ThinLinc Profiles as documented in Section 14.4.3. This allows for different Application Groups to be selected based on user input after login.

- *Shell Command Activating this Application Group*

This setting allows you to activate application groups based on the return value of an arbitrary command. If the command returns 0 (which is the standard return code for success for shell commands), the application group will be activated.

This can be used for example to activate application groups based on group membership by using the **tl-memberof-group** command. It can also be used to activate an application group for all users by running **/bin/true** as activation command.

The command is run via the shell in the current user's environment when running **tl-desktop-activate.sh**. The environment variable `TLDCGROUP` is set to the application group currently under consideration for activation.

- *Save!*

Don't forget to press the Save-button, or none of the changes will be written to the database.

18.2.6. Distribute Configuration to all agent hosts

After doing changes to the Desktop Configuration, the new configuration must be copied to all VSM agent hosts. The files/directories to be copied are

`/opt/thinlinc/etc/conf.d/tl-desktop-customizer.hconf` and all subdirectories of `/opt/thinlinc/desktops`.

Best Practice: Use the **tl-rsync-all** command as described in Chapter 13 to copy the files.

18.3. Enabling the Custom Desktops for users

Enabling the Custom Desktops for users is easy. Simply create a symbolic link in `/opt/thinlinc/etc/xstartup.d`:

```
# ln -s
    /opt/thinlinc/bin/tl-desktop-activate.sh
    /opt/thinlinc/etc/xstartup.d/35-tl-desktop-activate.sh
```

The ThinLinc session startup will read this file and make sure the environment variable `KDEDIRS` is set correctly. It will also write a `~/.config/menu/applications.menu` and possibly create symbolic links under `~/.kde/share/apps/kdesktop/Desktop`. Your profile should then execute the command **startkde**.

Note: The TLDC activation script only runs TLDC for non-root users. Test your TLDC configuration using a normal user.

18.4. Tips & Tricks with TLDC

18.4.1. Unwanted Icons on the Desktop with KDE

At first login for each user, KDE copies files from `/usr/share/apps/kdesktop/DesktopLinks` to the Desktop directory of the user. This means that if there is a Home Icon in `DesktopLinks` and you add a Home Icon via TLDC, there will be two Home Icons.

Remove the contents of `/usr/share/apps/kdesktop/DesktopLinks` to solve the problem, and let TLDC be the sole provider of icons on the desktop.

Note: If you KDE is based somewhere else than under `/usr`, the `DesktopLinks` directory will be situated elsewhere. For example, on SuSE, KDE is based at `/opt/kde3`.

18.4.2. File Associations for Applications Not In the Menu

When KDE tries to determine what application to use for opening a specific file, it is only looking for applications that are available in the menu. There are cases where not all applications that may be used for opening files are meant to be available in the menu.

In this case, create a hidden menu by setting *"Hide this Menu"* to *Yes* in the Menu Structure Editor, and then create an Application Group that adds the applications that should be available for file associations in to this menu.

18.4.3. Home Icon not Working in KDE?

This is a case of the problem above where File Associations are not working. Create an Application Group that includes the Konqueror (`kde-kfmclient_dir`) application in a hidden menu, and make sure this application group is added for all relevant users, and the home icon will work again.

Note: A menu named *"Hidden Menu"* is created by the application group *"Hidden"* which is by default activated for the profile *kde*. This menu contains the `kde-kcmclient_dir` to make sure the home icon is working. Make sure this application group is activated for all users with a desktop based on KDE.

Appendix A. TCP Ports Used by ThinLinc

A.1. On Machine Running VSM Server

22: SSH Daemon

Port 22 is not used by ThinLinc *per se*, but since no ThinLinc installation can work without a running SSH daemon, we list port 22 here. Port 22 is the normal SSH port, but basically any port can be used - the client has support for connecting to any port. Note however that if the SSH daemon on the VSM server is listening on port *x*, all VSM agents must also have their SSH daemons configured to listen on port *x*.

300: ThinLinc Web Access

By default, ThinLinc's Web Access service `tlwebaccess` is available on TCP port 300. Traffic to this port is encrypted (TLS).

Note: The port on which `tlwebaccess` runs is configurable via the parameter `/webaccess/listen_port`. See Section 14.2.8 for details.

1010: ThinLinc Administration Interface (`tlwebadm`)

By default, ThinLinc's web-based administration interface is available on TCP port 1010. In order to access this interface remotely, port 1010 will need to be reachable. Traffic to this port is encrypted (TLS).

Note: The port on which `tlwebadm` runs is configurable via the parameter `/tlwebadm/listen_port`. See Section 14.2.7 for details.

9000: VSM server

The VSM server listens on port 9000. The traffic is not encrypted, but sensitive information will only be shared with root or connections originating from a port lower than 1024, from a list of known IP addresses. The protocol used is XML-RPC through HTTP (using a minimal internal HTTP server in the VSM server).

A.2. On Machine Running VSM Agent

22: SSH Daemon

Just as for the VSM server, there must be a SSH Daemon running on all VSM agent machines. This daemon is normally listening to port 22, although it can listen to other ports as well. See the entry about port 22 on Section A.1.

300: ThinLinc Web Access

By default, ThinLinc's Web Access service `tlwebaccess` is available on TCP port 300. Traffic to this port is encrypted (TLS).

Note: The port on which `tlwebaccess` runs is configurable via the parameter `/webaccess/listen_port`. See Section 14.2.8 for details.

904: VSM Agent

The VSM agent listens on port 904 for incoming requests from the VSM server host. The traffic is not encrypted, but the VSM agent only allows connections originating from a port lower than 1024, from a list of known IP addresses. The protocol in use is XMLRPC through HTTP.

1010: ThinLinc Administration Interface (`tlwebadm`)

By default, ThinLinc's web-based administration interface is available on TCP port 1010. See the entry about port 1010 at Section A.1.

5901-5999: VNC servers for first 99 sessions

Ports 5901-5999 are used by `Xvnc` processes serving display numbers strictly below 100.

4900-5899: Tunnels to clients

The ports in this interval is used as serverside endpoints for the SSH tunnels used to access local resources of the client, for example local drives, serial ports and sound.

This interval is used for sessions with display number strictly below 100.

The algorithm used for calculating which ports to use for a specific display number in this interval is $4900 + display-id * 10 + SERVICE_SLOT$ where `SERVICE_SLOT` is a number defined under `/vsm/tunnelservices`.

6001-8000: X Display ports

If `Xvnc` is configured to listen for incoming TCP requests from X Window System clients on other hosts, ports 6001-8000 are used by display numbers 1-2000. The default is not to listen for incoming TCP requests, so in the default configuration, the ports in this interval are not open.

Port 32767 downwards to 11857

The algorithm described below is used to allocate ports for the `Xvnc` process and for the serverside endpoints for SSH tunnels to access local resources of the client. This algorithm is used for sessions with display numbers strictly higher than 99.

Each session is allocated `/vsm/tunnelslots_per_session` (default value 10) + 1 ports, leading to an allocation of 11 ports per session with the default configuration. The ports are allocated starting with the port given as `/vsmagent/max_session_port` (default 32767), and then downwards. This means that the ports are aligned upwards as closely as possible to the upper limit defined. This is a good practice to avoid collisions with other services running on the machine.

Some examples follow

Display number 50

The VNC port will be 5950 which is $5900 + display$.

The tunnel ports allocated for this display are 5400-5409, which is $4900 + (10 * display) + SERVICE_SLOT$ where SERVICE_SLOT is 0-9.

Display number 100, `/vsmagent/display_min = 10` (the default),
`/vsmagent/max_session_port = 32767`.

The VNC port will be 32757, which is $32767 - ((display - 100) * (/vsm/tunnelslots_per_session + 1) + /vsm/tunnelslots_per_session)$.

Ports 32758-32767 (inclusive) will be used for tunnel ports.

Display number 300, `/vsmagent/display_min = 100`, `/vsmagent/max_session_port = 32767` (the default).

The VNC port will be 30557 which is $32767 - ((display - 100) * (/vsm/tunnelslots_per_session + 1) + /vsm/tunnelslots_per_session)$.

Ports 30558-30567 (inclusive) will be used for tunnel ports.

Display number 600, `/vsmagent/display_min = 300`, `/vsmagent/max_session_port = 32767` (the default).

The VNC port will be 29457, which is $32767 - ((display - 300) * (/vsm/tunnelslots_per_session + 1) + /vsm/tunnelslots_per_session)$.

Ports 29458-29467 (inclusive) will be used for tunnel ports.

Appendix B. Troubleshooting ThinLinc

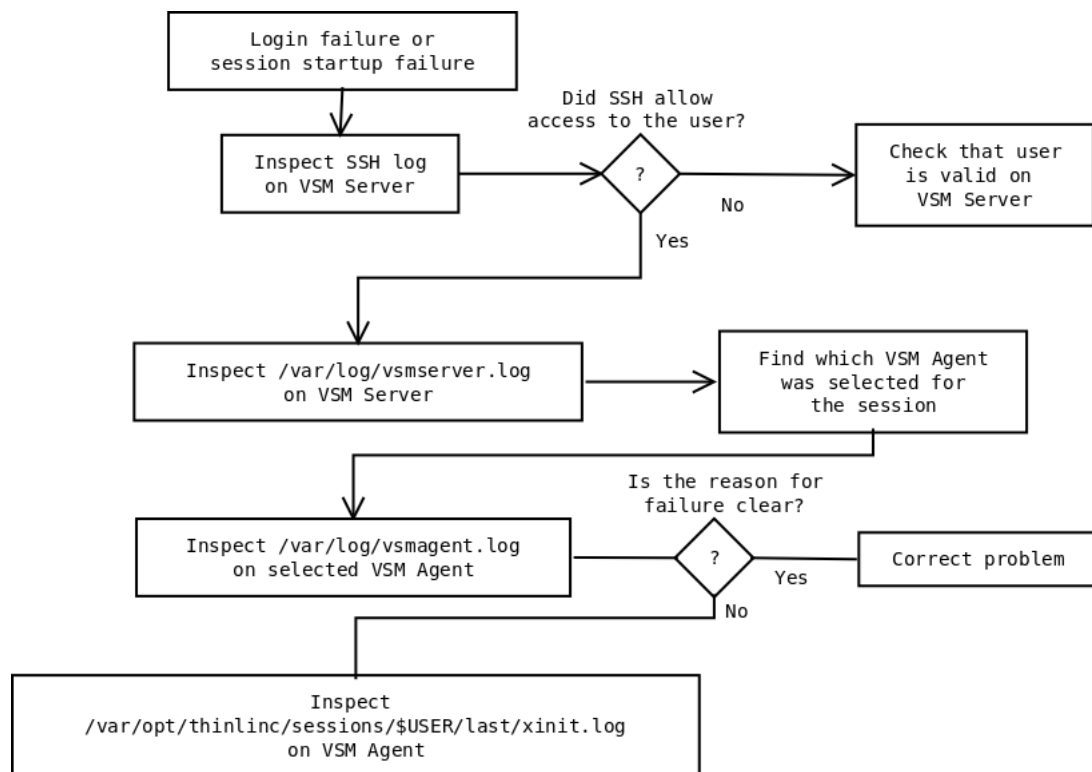
In this appendix, we will describe how to troubleshoot common problems in a ThinLinc installation.

We will begin by giving a general view of the recommended troubleshooting method, and then continue with more detailed instructions for troubleshooting specific problems.

B.1. General troubleshooting method

In most cases, troubleshooting a ThinLinc session problem should follow the method outlined in Figure B-1.

Figure B-1. The General Troubleshooting Method



The method is to first check that the user was let in by SSH on the VSM server. This information is found on different places on different distributions. Common log filenames for SSH information are `/var/log/secure`, `/var/log/auth.log` or `/var/log/daemon.log`. If the user was let in by SSH, the VSM server log (`/var/log/vsmserver.log`) is inspected. In some cases, the reason for session failure can be found there, but most of the times, it's necessary to find out which VSM agent was selected for the session, and inspect the VSM agent log (`/var/log/vsmagent.log`) on the server in question.

If inspecting `/var/log/vsmagent.log` on the server that was selected for the session does not reveal the reason for the failure, there is a per-session log in `/var/opt/thinlinc/sessions/<username>/last/xinit.log` where the output of commands run during session startup is stored.

In very rare cases, it might also be necessary to inspect the SSH log on the VSM agent.

B.2. Troubleshooting Specific Problems

The first step should be to check if your specific problem is mentioned in the Platform Specific Notes available at <https://www.cendio.com/thinlinc/docs/platforms>. If your problem isn't mentioned in the Platform Specific Notes you should read the sections below.

B.2.1. Problems Where the Client Reports an Error

In the following sections, we will describe how to cope with problems where the ThinLinc client is reporting an error.

B.2.1.1. Couldn't set up secure tunnel to ThinLinc server. (Couldn't establish SSH tunnel, SSH terminated.)

This error is caused by failure to connect to the SSH daemon on the ThinLinc server (the server running the VSM server). This could be caused by the fact that the SSH daemon is simply not running, or that it is not letting the user in for some reason.

Another possible reason is that there is a firewall between the user and the ThinLinc server, that forbids communication.

B.2.1.2. "Login Failed! Wrong username or password."

This error is very often caused simply by the user entering an incorrect password. Begin by verifying that the user is actually entering the correct username and password.

If the username and password are correct and this is the first time the user tries to login, check the SSH logs of the server. If SSH says that the user is invalid, that means that something is incorrect in the user's user information database entry. For example, this may happen if a user stored in eDirectory has two *cn* attributes, one of them different than the other.

The **getent** command can be a valuable tool to dissect problems of this type. If the output of **getent passwd <username>** doesn't produce any output, that is a sign that the user is in fact invalid.

Note: Usernames beginning with numbers (for example *96aabbcc*) are parsed as numeric uids by **getent**, rendering **getent** rather useless for debugging purposes in environments with username schemes beginning with numbers.

If usernames with numbers in the beginning are in use, the following python code can be used to verify a username

```
python-thinlinc -c 'import pwd, sys; print pwd.getpwnam(sys.argv[1])' <username>
```

B.2.1.3. The SSH connection succeeded, but the ThinLinc server connection failed. Perhaps this server doesn't run a ThinLinc server?

This error is most often caused by the fact that the VSM server is not running on the server. Start the VSM server and try again.

A user entering the wrong hostname, for example the hostname of one of the VSM agents, would also get this error message. Check that the user has entered the correct hostname. In very rare cases, this could also be caused by incorrect DNS data.

B.2.1.4. ThinLinc login failed. (No agent server was available)

This error is reported if there were no working VSM agents available according to the load balance information in the VSM server.

In a system with few VSM agent servers, restoring a VSM agent that has been down for some reason doesn't take effect immediately - the load balance information is only updated once every 40 seconds by default. Either wait for the load balance cycle to complete, or restart the VSM server. In a small cluster it might be a good idea to lower the load balance cycle, by setting the parameter `/vsmserver/load_balance_cycle`.

The load balance information can be inspected in the ThinLinc Web Administration, see Chapter 17.

B.2.1.5. ThinLinc login failed. (Couldn't create your session)

When this error occurs, the user was valid on the VSM server, but for some reason, the session couldn't be created on the VSM agent.

One very common reason for this problem is that the VSM agent has lost its connection to the user database backend (LDAP, Windows domain or other database), so the user exists on the VSM server, but not on the VSM agent. If this is the case, the VSM agent log on the selected server will clearly state that the user doesn't exist on the system.

Another very common reason is home directory trouble on the VSM agent. Verify that the home directory exists on the selected server, and that it is owned by the correct uidNumber/gidNumber. Of course, the user must have write permissions on his/her home directory.

To verify that the home directory works, the following command can be used:

```
ssh <username>@<agenthost> touch .
```

If the home directory is correctly mounted and writable by the user, the above command will not produce any output except the password question.

B.2.1.6. Couldn't set up secure tunnel to VNC! (Couldn't establish SSH tunnel, SSH terminated.)

This error is caused by failure to connect to the SSH daemon on the selected VSM agent. This could be caused by the fact that the SSH daemon is simply not running, or that it is not letting the user in for some reason.

Another possible reason is that there is a firewall between the user and the selected VSM agent that disallows communication.

B.2.2. Problems that Occur After Session Start

In this section we will discuss some problems that can occur after the successful login, that is, after the ThinLinc login window has closed. In this phase, a number of session startup problems can occur

B.2.2.1. Session starts, but closes down immediately

If the ThinLinc login window closes, and the session starts up but then immediately shuts down, inspect `xinit.log` found in `/var/opt/thinlinc/sessions/<username>/last/` on the selected VSM agent. Some of the commands run during session startup will probably have written an error message that will be stored in that file.

It may also be of value to inspect the VSM agent log on the selected server.

B.2.2.2. At login, user is reconnected to previous, faulty, session

If a previous session still exists and is faulty, for example because of desktop environment failures, the user is reconnected to the same session when logging in. Disconnect from the session, enable the *"End existing session"* checkbox and log in again. That will terminate the current session and start a new one.

B.2.2.3. Login Succeeds, but the ThinLinc Desktop Configuration fails

When using the ThinLinc Desktop Customizer, as documented in Chapter 18, the KDE or Gnome menu and the entries on the desktop are customized at each login. If this fails, quota problems are very often the problem. Check the quota of the user in question.

B.2.2.4. Login Succeeds, but KDE Fails to Start

If KDE fails to start, complaining about being unable to create symlinks and similar, quota problems are very often the real problem. Check the quota of the user in question.

Appendix C. Restricting access to ThinLinc servers

In some cases it might be desirable or required to restrict the users' access to the ThinLinc servers and their ability to move data in and out of the system. This chapter describes some ways this can be achieved, as well as the consequences of such restrictions.

C.1. Disabling SSH access

The system's SSH server often includes a lot of functionality for accessing the system. Completely disabling this service is a quick way to restrict most of the external access to the system. However the native ThinLinc client requires SSH to function so users will be limited to only using the HTML based Web Access client.

Many SSH servers also support limiting access to just certain users. OpenSSH has settings such as `AllowGroups` and `Match` that can limit functionality without completely disabling the SSH server.

C.2. Disabling shell access

User sessions are normally started via the user's configured shell, so restricting the shell is a good method to restrict what kind of sessions the user can start. Primarily this is useful to prevent users from running custom commands via SSH.

C.2.1. Changing the configured shell

Commonly the user's shell is configured to `/bin/false` in order to disable shell access. Unfortunately this also prevents access to ThinLinc as it needs to run the commands `thinlinc-login` and `/opt/thinlinc/etc/xsession` via the user's shell. As an alternative it is possible to configure `/usr/bin/thinlinc-login` as the shell. This will allow ThinLinc to function whilst preventing any other type of session.

Note that this method prevents any terminals inside the session from functioning as well. In most cases it also does not prevent users from running custom scripts and shell commands as they can use a text editor to construct such scripts.

C.2.2. Using ForceCommand

OpenSSH has the ability to ignore the user's configured shell and run a different command instead. This makes it possible to keep a normal shell for the user and only restrict access when connecting via SSH. However this prevents the native ThinLinc client from connecting as it needs to be able to run the command `thinlinc-login` with specific arguments. The following script can be specified as `ForceCommand` to allow only ThinLinc access via SSH:

```
#!/bin/bash
thinlinc-login -c "${SSH_ORIGINAL_COMMAND}"
```

It is also possible to apply this restriction only to a subset of users by using the `Match` setting. Please see OpenSSH's documentation for how to configure this mechanism.

C.3. Disabling port forwarding

ThinLinc relies on SSH port forwarding in order to function. However it is possible to limit that port forwarding in order to avoid unwanted network access. ThinLinc only requires forwarding via the loopback interface, so the SSH server can always be configured to only allow this without limiting ThinLinc in any way. For OpenSSH this is configured by specifying the following in `sshd_config`:

```
GatewayPorts no
PermitOpen 127.0.0.1:*
```

Note that it is also necessary to disable shell access in order to completely prevent users from forwarding ports as otherwise they could run their own forwarding software over the shell channel.

C.3.1. Disabling remote port forwarding

It is possible to use ThinLinc with remote port forwarding completely disabled. However this will prevent local devices such as sound, drives and printers from functioning. In OpenSSH this is configured by adding the following to `/etc/ssh/sshd_config`:

```
AllowTcpForwarding local
```

It is also possible to apply this restriction only to a subset of users by using the `Match` setting. Please see OpenSSH's documentation for how to configure this mechanism.

Note: Local port forwarding cannot be disabled as it is required for even the basic ThinLinc functionality.

C.4. Disabling clipboard

It is possible to disable clipboard transfers in either direction in order to avoid easily moving data to and from the server. The first step is adding `-noclipboard` to the ThinLinc setting `/vsmagent/xserver_args`. This prevents the user from later changing the clipboard settings. Next add `-AcceptCutText=0` to disable clipboard transfers going to the server, and `-SendCutText=0` to prevent transfers going from the server.

C.5. Disabling local drives

ThinLinc local drives redirection relies on being able to ask the kernel to mount a NFS share. This is a privileged operation that only root is permitted to perform, and as such this feature requires a setuid helper binary. This helper is called `/opt/thinlinc/libexec/tl-mount-personal` and the setuid permission can be removed by running the follow:

```
$ sudo chmod u-s /opt/thinlinc/libexec/tl-mount-personal
```


Note that this only disables the ability to use the kernel NFS client. If a user can start some other NFS client then they can still access the local drive redirection. The setuid permission is also restored each time ThinLinc is upgraded.

Appendix D. GnuTLS priority strings

ThinLinc uses priority strings to allow the administrator to select their own preferred availability and order of algorithms used by GnuTLS for services that uses `tlstunnel`. The priority string is a colon-delimited list of strings being either keywords (groups of algorithms) or algorithms which can be individually enabled or disabled.

For more information, see the GnuTLS documentation about priority strings.

D.1. Standard configuration

ThinLinc comes configured with the priority string "NORMAL", which means the standard, secure GnuTLS algorithms. This is the order and availability of algorithms for that priority string.

D.1.1. Cipher suites

```
TLS_AES_256_GCM_SHA384
TLS_CHACHA20_POLY1305_SHA256
TLS_AES_128_GCM_SHA256
TLS_AES_128_CCM_SHA256
TLS_ECDHE_ECDSA_AES_256_GCM_SHA384
TLS_ECDHE_ECDSA_CHACHA20_POLY1305
TLS_ECDHE_ECDSA_AES_256_CCM
TLS_ECDHE_ECDSA_AES_256_CBC_SHA1
TLS_ECDHE_ECDSA_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_AES_128_CCM
TLS_ECDHE_ECDSA_AES_128_CBC_SHA1
TLS_ECDHE_RSA_AES_256_GCM_SHA384
TLS_ECDHE_RSA_CHACHA20_POLY1305
TLS_ECDHE_RSA_AES_256_CBC_SHA1
TLS_ECDHE_RSA_AES_128_GCM_SHA256
TLS_ECDHE_RSA_AES_128_CBC_SHA1
TLS_RSA_AES_256_GCM_SHA384
TLS_RSA_AES_256_CCM
TLS_RSA_AES_256_CBC_SHA1
TLS_RSA_AES_128_GCM_SHA256
TLS_RSA_AES_128_CCM
TLS_RSA_AES_128_CBC_SHA1
TLS_DHE_RSA_AES_256_GCM_SHA384
TLS_DHE_RSA_CHACHA20_POLY1305
TLS_DHE_RSA_AES_256_CCM
TLS_DHE_RSA_AES_256_CBC_SHA1
TLS_DHE_RSA_AES_128_GCM_SHA256
TLS_DHE_RSA_AES_128_CCM
TLS_DHE_RSA_AES_128_CBC_SHA1
```

D.1.2. Protocols

VERS-TLS1.3
VERS-TLS1.2
VERS-TLS1.1
VERS-TLS1.0
VERS-DTLS1.2
VERS-DTLS1.0

D.1.3. Ciphers

AES-256-GCM
CHACHA20-POLY1305
AES-256-CCM
AES-256-CBC
AES-128-GCM
AES-128-CCM
AES-128-CBC

D.1.4. MACs

SHA1
AEAD

D.1.5. Key Exchange Algorithms

ECDHE-ECDSA
ECDHE-RSA
RSA
DHE-RSA

D.1.6. Groups

GROUP-SECP256R1
GROUP-SECP384R1
GROUP-SECP521R1
GROUP-X25519
GROUP-FFDHE2048
GROUP-FFDHE3072
GROUP-FFDHE4096
GROUP-FFDHE6144
GROUP-FFDHE8192

D.1.7. PK-signatures

```
SIGN-RSA-SHA256
SIGN-RSA-PSS-SHA256
SIGN-RSA-PSS-RSAE-SHA256
SIGN-ECDSA-SHA256
SIGN-ECDSA-SECP256R1-SHA256
SIGN-EdDSA-Ed25519
SIGN-RSA-SHA384
SIGN-RSA-PSS-SHA384
SIGN-RSA-PSS-RSAE-SHA384
SIGN-ECDSA-SHA384
SIGN-ECDSA-SECP384R1-SHA384
SIGN-RSA-SHA512
SIGN-RSA-PSS-SHA512
SIGN-RSA-PSS-RSAE-SHA512
SIGN-ECDSA-SHA512
SIGN-ECDSA-SECP521R1-SHA512
SIGN-RSA-SHA1
SIGN-ECDSA-SHA1
```

D.2. Available algorithms

Here are all the available algorithms for use in a priority string in ThinLinc.

D.2.1. Cipher suites

```
TLS_AES_128_GCM_SHA256
TLS_AES_256_GCM_SHA384
TLS_CHACHA20_POLY1305_SHA256
TLS_AES_128_CCM_SHA256
TLS_AES_128_CCM_8_SHA256
TLS_RSA_NULL_MD5
TLS_RSA_NULL_SHA1
TLS_RSA_NULL_SHA256
TLS_RSA_ARCFOUR_128_SHA1
TLS_RSA_ARCFOUR_128_MD5
TLS_RSA_3DES_EDE_CBC_SHA1
TLS_RSA_AES_128_CBC_SHA1
TLS_RSA_AES_256_CBC_SHA1
TLS_RSA_CAMELLIA_128_CBC_SHA256
TLS_RSA_CAMELLIA_256_CBC_SHA256
TLS_RSA_CAMELLIA_128_CBC_SHA1
TLS_RSA_CAMELLIA_256_CBC_SHA1
TLS_RSA_AES_128_CBC_SHA256
```

Appendix D. GnuTLS priority strings

```
TLS_RSA_AES_256_CBC_SHA256
TLS_RSA_AES_128_GCM_SHA256
TLS_RSA_AES_256_GCM_SHA384
TLS_RSA_CAMELLIA_128_GCM_SHA256
TLS_RSA_CAMELLIA_256_GCM_SHA384
TLS_RSA_AES_128_CCM
TLS_RSA_AES_256_CCM
TLS_RSA_AES_128_CCM_8
TLS_RSA_AES_256_CCM_8
TLS_DHE_DSS_ARCFOUR_128_SHA1
TLS_DHE_DSS_3DES_EDE_CBC_SHA1
TLS_DHE_DSS_AES_128_CBC_SHA1
TLS_DHE_DSS_AES_256_CBC_SHA1
TLS_DHE_DSS_CAMELLIA_128_CBC_SHA256
TLS_DHE_DSS_CAMELLIA_256_CBC_SHA256
TLS_DHE_DSS_CAMELLIA_128_CBC_SHA1
TLS_DHE_DSS_CAMELLIA_256_CBC_SHA1
TLS_DHE_DSS_AES_128_CBC_SHA256
TLS_DHE_DSS_AES_256_CBC_SHA256
TLS_DHE_DSS_AES_128_GCM_SHA256
TLS_DHE_DSS_AES_256_GCM_SHA384
TLS_DHE_DSS_CAMELLIA_128_GCM_SHA256
TLS_DHE_DSS_CAMELLIA_256_GCM_SHA384
TLS_DHE_RSA_3DES_EDE_CBC_SHA1
TLS_DHE_RSA_AES_128_CBC_SHA1
TLS_DHE_RSA_AES_256_CBC_SHA1
TLS_DHE_RSA_CAMELLIA_128_CBC_SHA256
TLS_DHE_RSA_CAMELLIA_256_CBC_SHA256
TLS_DHE_RSA_CAMELLIA_128_CBC_SHA1
TLS_DHE_RSA_CAMELLIA_256_CBC_SHA1
TLS_DHE_RSA_AES_128_CBC_SHA256
TLS_DHE_RSA_AES_256_CBC_SHA256
TLS_DHE_RSA_AES_128_GCM_SHA256
TLS_DHE_RSA_AES_256_GCM_SHA384
TLS_DHE_RSA_CAMELLIA_128_GCM_SHA256
TLS_DHE_RSA_CAMELLIA_256_GCM_SHA384
TLS_DHE_RSA_CHACHA20_POLY1305
TLS_DHE_RSA_AES_128_CCM
TLS_DHE_RSA_AES_256_CCM
TLS_DHE_RSA_AES_128_CCM_8
TLS_DHE_RSA_AES_256_CCM_8
TLS_ECDHE_RSA_NULL_SHA1
TLS_ECDHE_RSA_3DES_EDE_CBC_SHA1
TLS_ECDHE_RSA_AES_128_CBC_SHA1
TLS_ECDHE_RSA_AES_256_CBC_SHA1
TLS_ECDHE_RSA_AES_256_CBC_SHA384
TLS_ECDHE_RSA_ARCFOUR_128_SHA1
TLS_ECDHE_RSA_CAMELLIA_128_CBC_SHA256
TLS_ECDHE_RSA_CAMELLIA_256_CBC_SHA384
TLS_ECDHE_ECDSA_NULL_SHA1
TLS_ECDHE_ECDSA_3DES_EDE_CBC_SHA1
TLS_ECDHE_ECDSA_AES_128_CBC_SHA1
TLS_ECDHE_ECDSA_AES_256_CBC_SHA1
```

TLS_ECDHE_ECDSA_ARCFOUR_128_SHA1
TLS_ECDHE_ECDSA_CAMELLIA_128_CBC_SHA256
TLS_ECDHE_ECDSA_CAMELLIA_256_CBC_SHA384
TLS_ECDHE_ECDSA_AES_128_CBC_SHA256
TLS_ECDHE_RSA_AES_128_CBC_SHA256
TLS_ECDHE_ECDSA_CAMELLIA_128_GCM_SHA256
TLS_ECDHE_ECDSA_CAMELLIA_256_GCM_SHA384
TLS_ECDHE_ECDSA_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_AES_256_GCM_SHA384
TLS_ECDHE_RSA_AES_128_GCM_SHA256
TLS_ECDHE_RSA_AES_256_GCM_SHA384
TLS_ECDHE_ECDSA_AES_256_CBC_SHA384
TLS_ECDHE_RSA_CAMELLIA_128_GCM_SHA256
TLS_ECDHE_RSA_CAMELLIA_256_GCM_SHA384
TLS_ECDHE_RSA_CHACHA20_POLY1305
TLS_ECDHE_ECDSA_CHACHA20_POLY1305
TLS_ECDHE_ECDSA_AES_128_CCM
TLS_ECDHE_ECDSA_AES_256_CCM
TLS_ECDHE_ECDSA_AES_128_CCM_8
TLS_ECDHE_ECDSA_AES_256_CCM_8
TLS_ECDHE_PSK_3DES_EDE_CBC_SHA1
TLS_ECDHE_PSK_AES_128_CBC_SHA1
TLS_ECDHE_PSK_AES_256_CBC_SHA1
TLS_ECDHE_PSK_AES_128_CBC_SHA256
TLS_ECDHE_PSK_AES_256_CBC_SHA384
TLS_ECDHE_PSK_ARCFOUR_128_SHA1
TLS_ECDHE_PSK_NULL_SHA1
TLS_ECDHE_PSK_NULL_SHA256
TLS_ECDHE_PSK_NULL_SHA384
TLS_ECDHE_PSK_CAMELLIA_128_CBC_SHA256
TLS_ECDHE_PSK_CAMELLIA_256_CBC_SHA384
TLS_PSK_ARCFOUR_128_SHA1
TLS_PSK_3DES_EDE_CBC_SHA1
TLS_PSK_AES_128_CBC_SHA1
TLS_PSK_AES_256_CBC_SHA1
TLS_PSK_AES_128_CBC_SHA256
TLS_PSK_AES_256_GCM_SHA384
TLS_PSK_CAMELLIA_128_GCM_SHA256
TLS_PSK_CAMELLIA_256_GCM_SHA384
TLS_PSK_AES_128_GCM_SHA256
TLS_PSK_NULL_SHA1
TLS_PSK_NULL_SHA256
TLS_PSK_CAMELLIA_128_CBC_SHA256
TLS_PSK_CAMELLIA_256_CBC_SHA384
TLS_PSK_AES_256_CBC_SHA384
TLS_PSK_NULL_SHA384
TLS_RSA_PSK_ARCFOUR_128_SHA1
TLS_RSA_PSK_3DES_EDE_CBC_SHA1
TLS_RSA_PSK_AES_128_CBC_SHA1
TLS_RSA_PSK_AES_256_CBC_SHA1
TLS_RSA_PSK_CAMELLIA_128_GCM_SHA256
TLS_RSA_PSK_CAMELLIA_256_GCM_SHA384
TLS_RSA_PSK_AES_128_GCM_SHA256

Appendix D. GnuTLS priority strings

```
TLS_RSA_PSK_AES_128_CBC_SHA256
TLS_RSA_PSK_NULL_SHA1
TLS_RSA_PSK_NULL_SHA256
TLS_RSA_PSK_AES_256_GCM_SHA384
TLS_RSA_PSK_AES_256_CBC_SHA384
TLS_RSA_PSK_NULL_SHA384
TLS_RSA_PSK_CAMELLIA_128_CBC_SHA256
TLS_RSA_PSK_CAMELLIA_256_CBC_SHA384
TLS_DHE_PSK_ARCFOUR_128_SHA1
TLS_DHE_PSK_3DES_EDE_CBC_SHA1
TLS_DHE_PSK_AES_128_CBC_SHA1
TLS_DHE_PSK_AES_256_CBC_SHA1
TLS_DHE_PSK_AES_128_CBC_SHA256
TLS_DHE_PSK_AES_128_GCM_SHA256
TLS_DHE_PSK_NULL_SHA1
TLS_DHE_PSK_NULL_SHA256
TLS_DHE_PSK_NULL_SHA384
TLS_DHE_PSK_AES_256_CBC_SHA384
TLS_DHE_PSK_AES_256_GCM_SHA384
TLS_DHE_PSK_CAMELLIA_128_CBC_SHA256
TLS_DHE_PSK_CAMELLIA_256_CBC_SHA384
TLS_DHE_PSK_CAMELLIA_128_GCM_SHA256
TLS_DHE_PSK_CAMELLIA_256_GCM_SHA384
TLS_PSK_AES_128_CCM
TLS_PSK_AES_256_CCM
TLS_DHE_PSK_AES_128_CCM
TLS_DHE_PSK_AES_256_CCM
TLS_PSK_AES_128_CCM_8
TLS_PSK_AES_256_CCM_8
TLS_DHE_PSK_AES_128_CCM_8
TLS_DHE_PSK_AES_256_CCM_8
TLS_DHE_PSK_CHACHA20_POLY1305
TLS_ECDHE_PSK_CHACHA20_POLY1305
TLS_RSA_PSK_CHACHA20_POLY1305
TLS_PSK_CHACHA20_POLY1305
TLS_DH_ANON_ARCFOUR_128_MD5
TLS_DH_ANON_3DES_EDE_CBC_SHA1
TLS_DH_ANON_AES_128_CBC_SHA1
TLS_DH_ANON_AES_256_CBC_SHA1
TLS_DH_ANON_CAMELLIA_128_CBC_SHA256
TLS_DH_ANON_CAMELLIA_256_CBC_SHA256
TLS_DH_ANON_CAMELLIA_128_CBC_SHA1
TLS_DH_ANON_CAMELLIA_256_CBC_SHA1
TLS_DH_ANON_AES_128_CBC_SHA256
TLS_DH_ANON_AES_256_CBC_SHA256
TLS_DH_ANON_AES_128_GCM_SHA256
TLS_DH_ANON_AES_256_GCM_SHA384
TLS_DH_ANON_CAMELLIA_128_GCM_SHA256
TLS_DH_ANON_CAMELLIA_256_GCM_SHA384
TLS_ECDH_ANON_NULL_SHA1
TLS_ECDH_ANON_3DES_EDE_CBC_SHA1
TLS_ECDH_ANON_AES_128_CBC_SHA1
TLS_ECDH_ANON_AES_256_CBC_SHA1
```



```
TLS_ECDH_ANON_ARCFOUR_128_SHA1
TLS_SRP_SHA_3DES_EDE_CBC_SHA1
TLS_SRP_SHA_AES_128_CBC_SHA1
TLS_SRP_SHA_AES_256_CBC_SHA1
TLS_SRP_SHA_DSS_3DES_EDE_CBC_SHA1
TLS_SRP_SHA_RSA_3DES_EDE_CBC_SHA1
TLS_SRP_SHA_DSS_AES_128_CBC_SHA1
TLS_SRP_SHA_RSA_AES_128_CBC_SHA1
TLS_SRP_SHA_DSS_AES_256_CBC_SHA1
TLS_SRP_SHA_RSA_AES_256_CBC_SHA1
```

D.2.2. Certificate types

```
CTYPE-X.509
CTYPE-Raw Public Key
```

D.2.3. Protocols

```
VERS-TLS1.0
VERS-TLS1.1
VERS-TLS1.2
VERS-TLS1.3
VERS-DTLS0.9
VERS-DTLS1.0
VERS-DTLS1.2
```

D.2.4. Ciphers

```
AES-256-CBC
AES-192-CBC
AES-128-CBC
AES-128-GCM
AES-256-GCM
AES-128-CCM
AES-256-CCM
AES-128-CCM-8
AES-256-CCM-8
ARCFOUR-128
ESTREAM-SALSA20-256
SALSA20-256
CAMELLIA-256-CBC
CAMELLIA-192-CBC
CAMELLIA-128-CBC
CHACHA20-POLY1305
CAMELLIA-128-GCM
CAMELLIA-256-GCM
```

Appendix D. GnuTLS priority strings

GOST28147-TC26Z-CFB
GOST28147-CPA-CFB
GOST28147-CPB-CFB
GOST28147-CPC-CFB
GOST28147-CPD-CFB
AES-128-CFB8
AES-192-CFB8
AES-256-CFB8
AES-128-XTS
AES-256-XTS
GOST28147-TC26Z-CNT
3DES-CBC
DES-CBC
RC2-40
NULL

D.2.5. MACs

SHA1
SHA256
SHA384
SHA512
SHA224
UMAC-96
UMAC-128
AEAD
MD5
GOSTR341194
STREEBOG-256
STREEBOG-512
AES-CMAC-128
AES-CMAC-256
AES-GMAC-128
AES-GMAC-192
AES-GMAC-256
GOST28147-TC26Z-IMIT

D.2.6. Digests

SHA1
SHA256
SHA384
SHA512
SHA224
MD5
GOSTR341194
STREEBOG-256

STREEBOG-512

D.2.7. Key exchange algorithms

ECDHE-RSA
ECDHE-ECDSA
RSA
DHE-RSA
DHE-DSS
PSK
RSA-PSK
DHE-PSK
ECDHE-PSK
SRP-DSS
SRP-RSA
SRP
ANON-DH
ANON-ECDH
RSA-EXPORT

D.2.8. Compression

COMP-NULL

D.2.9. Groups

GROUP-SECP192R1
GROUP-SECP224R1
GROUP-SECP256R1
GROUP-SECP384R1
GROUP-SECP521R1
GROUP-X25519
GROUP-FFDHE2048
GROUP-FFDHE3072
GROUP-FFDHE4096
GROUP-FFDHE6144
GROUP-FFDHE8192

D.2.10. Public Key Systems

RSA
RSA-PSS
RSA
DSA
GOST R 34.10-2012-512
GOST R 34.10-2012-256
GOST R 34.10-2001
EC/ECDSA
EdDSA (Ed25519)
DH
ECDH (X25519)

D.2.11. PK-signatures

SIGN-RSA-SHA256
SIGN-RSA-SHA384
SIGN-RSA-SHA512
SIGN-RSA-PSS-SHA256
SIGN-RSA-PSS-RSAE-SHA256
SIGN-RSA-PSS-SHA384
SIGN-RSA-PSS-RSAE-SHA384
SIGN-RSA-PSS-SHA512
SIGN-RSA-PSS-RSAE-SHA512
SIGN-EdDSA-Ed25519
SIGN-ECDSA-SHA256
SIGN-ECDSA-SHA384
SIGN-ECDSA-SHA512
SIGN-ECDSA-SECP256R1-SHA256
SIGN-ECDSA-SECP384R1-SHA384
SIGN-ECDSA-SECP521R1-SHA512
SIGN-ECDSA-SHA3-224
SIGN-ECDSA-SHA3-256
SIGN-ECDSA-SHA3-384
SIGN-ECDSA-SHA3-512
SIGN-RSA-SHA3-224
SIGN-RSA-SHA3-256
SIGN-RSA-SHA3-384
SIGN-RSA-SHA3-512
SIGN-DSA-SHA3-224
SIGN-DSA-SHA3-256
SIGN-DSA-SHA3-384
SIGN-DSA-SHA3-512
SIGN-RSA-RAW
SIGN-RSA-SHA1
SIGN-RSA-SHA1
SIGN-RSA-SHA224
SIGN-RSA-RMD160
SIGN-DSA-SHA1

SIGN-DSA-SHA1
SIGN-DSA-SHA224
SIGN-DSA-SHA256
SIGN-RSA-MD5
SIGN-RSA-MD5
SIGN-RSA-MD2
SIGN-ECDSA-SHA1
SIGN-ECDSA-SHA224
SIGN-GOSTR341012-512
SIGN-GOSTR341012-256
SIGN-GOSTR341001
SIGN-DSA-SHA384
SIGN-DSA-SHA512

