



OpenVPN Access Server v1.6 System Administrator Guide

Rev 1.0

COPYRIGHT NOTICE

Copyright OpenVPN Technologies ©2010

TABLE OF CONTENTS

1	Introduction	2
1.1	Access Server Deployment Topology	2
1.2	Access Server Deployment Terminology	3
1.3	Deployment Overview (Quick Start)	4
2	OpenVPN Access Server Operation	5
2.1	Services and TCP/UDP Ports	5
2.2	Typical Network Configurations	5
2.2.1	One Network Interface on Private Network Behind the Firewall	6
2.2.2	Two Network Interfaces, One on Public and One on Private Network	6
2.2.3	One Network Interface on Public Network	7
2.3	User Authentication and Management	8
2.4	Client Configuration Generation and Management	8
2.5	Virtual VPN Subnet Configuration	9
3	Installation	10
3.1	Prepare the Server	10
3.2	Obtain License Key	10
3.3	Install OpenVPN Access Server RPM/DEB Package	10
3.4	Run ovpn-init	11
3.4.1	Configure Initial Admin Web UI Network Settings	12
3.4.2	Finalize the Initial Configuration	13
3.5	Configure Access Server with the Admin Web UI	14
4	Admin Web UI Reference	17
4.1	Status Pages	17
4.1.1	Status Overview	17
4.1.2	Log Reports	18
4.2	Configuration Pages	20
4.2.1	License	20
4.2.2	Server Network Settings	21
4.2.3	VPN Mode	24
4.2.4	VPN Settings	25
4.2.5	Advanced VPN	28
4.2.6	User Permissions	32
4.2.7	Group Permissions	34
4.3	Authentication Pages	35
4.3.1	General	35
4.3.2	PAM	36
4.3.3	RADIUS	37
4.3.4	LDAP	38
4.4	Tools Pages	39
4.4.1	Profiles	39
4.4.2	Connectivity Test	41
4.4.3	Support	43
5	Connect Client	44
5.1	Connect	45
5.2	Login	46
5.3	Rebranding the Admin UI	48
5.4	Certificates	49
5.5	Server-locked Profile	51

6	Additional Information on RADIUS Support	51
6.1	RADIUS Authentication Attributes.....	51
6.2	RADIUS Accounting Attributes	51
7	How to authenticate users with Active Directory	52
7.1.1	Configuring Access Server LDAP Authentication.....	52
7.1.2	Specifying Additional Requirements for LDAP Authentication.....	53
8	Failover	54

1 Introduction

The OpenVPN Access Server consists of a set of installation and configuration tools which allow for simple and rapid deployment of VPN remote access solutions using the OpenVPN open source project. The Access Server software builds upon the usability and popularity of OpenVPN, while easing VPN configuration and deployment by providing the following features:

1. Simplified server configuration

Access Server presents the administrator with only the most useful of the many configuration options supported by the sophisticated OpenVPN server and clients. An easy-to-use, Web-based configuration interface makes setting up and maintaining the Access Server deployment straight-forward and efficient.

2. Support for external user authentication database

Rather than requiring you to create and manage credentials for each valid VPN user, OpenVPN Access Server offers the ability to integrate with existing user authentication systems using one of the following:

1. PAM¹: the system for authenticating user accounts on the Unix server
2. an external LDAP or Active Directory server
3. one or more external RADIUS servers

3. Easy intuitive Web-Based client access

Once a user fires up a Web browser they can then enter their credentials and connect to the VPN. In addition a user can download a pre-configured Windows installer for their Windows Operating System. Since the installer file was dynamically generated specifically for the user in question, that user can instantly connect to the VPN without need for additional client-side configuration.

4. Compatibility with a large base of OpenVPN clients

An authenticated user can also download an OpenVPN client configuration file (also generated specifically for the user) from the Connect Client and use it with an OpenVPN v2.1+ client other than the Windows GUI client. In this way, OpenVPN Access Server is immediately compatible with OpenVPN clients running on non-Windows platforms, such as the Tunnelblick client on MacOSX and the Community Projects OpenVPN client on Unix/Linux.

Of course, none of these benefits would matter without the robust security of client-server communication provided by OpenVPN's use of SSL/TLS.

1.1 Access Server Deployment Topology

An OpenVPN Access Server deployment consists of one server, many clients and many users, as depicted in Figure 1. Each client machine in this topology uses the public IP network (the Internet) to communicate with the OpenVPN Access Server and thereby gains VPN-protected access to the private IP Network connected (if present).

¹ PAM stands for "Pluggable Authentication Modules," the common system for authenticating users on a Unix system.




Figure 1: OpenVPN Access Server Topology

1.2 Access Server Deployment Terminology

The following terminology is used when referring to an OpenVPN Access Server deployment:

Term	Definition
OpenVPN Access Server	The OpenVPN server daemon along with the Access Server's configuration and maintenance software running on a server computer.
User	An individual attempting remote access to private network resources via the public Internet.
Client	A computer (operated by a user) running OpenVPN client software in order to gain access to private network services via the OpenVPN Access Server.
User Credentials	A username and password used to authenticate a user.
OpenVPN Desktop Client for Windows	The OpenVPN Desktop client for Windows is a legacy client which has now been replaced by the Connect Client; however it is still available for users who need it.
Client Configuration File	A file which contains all of the information required for an OpenVPN client to securely connect to the OpenVPN server. User credentials are not included in the client configuration.
Connect Client	A client running on the Access Server which delivers client configuration files and/or pre-configured Windows client installer files to authenticated users. The Connect Client also allows for a user to login and connect through the browser.
Admin Web UI	A Web server running on the Access Server which is used by the administrator to configure the settings of the Access Server.

Table 1 Access Server Deployment Terminology

1.3 Deployment Overview (Quick Start)

Setting up the OpenVPN Access Server involves taking the following basic steps:

1. Determine the network configuration and IP addresses to use for server

See Section 2.1 for descriptions of typical network configurations. In short, you need to ensure that clients on the Internet can connect to the Access Server (either via a public IP address on the Access Server or via forwarding from a border firewall) and that the Access Server is connected to the private network, if one is to be used.

2. Obtain a license key

Register and sign in to www.openvpn.net to obtain an Access Server license key. If you are evaluating this product, we have already allocated a two-user test key to the Access Server.

3. Download and install the OpenVPN Access Server package file

Also from www.openvpn.net, download the appropriate binary package file for your server's particular version of Linux. Then (as root) install the package. For example, on Fedora/CentOS/RHEL:

```
rpm -i openvpn-as-1.6.0-Fedora9.x86_64.rpm
```

and on Ubuntu:

```
dpkg -i openvpn-as-1.6.0-Ubuntu8.amd_64.deb
```

4. Run `ovpn-init` to set initial configuration settings

Post 1.5.6: By default the `ovpn-init` tool is already run after the package install. If you still feel the need to run the tool again (to configure more advanced settings) you can run the tool again.

Run `ovpn-init` (without command-line arguments) using the `bash` shell:

```
/usr/local/openvpn_as/bin/ovpn-init --force
```

The `ovpn-init` utility asks a few questions regarding what IP address and port should be used for the Access Server Admin Web UI, and what user credentials should be used to login to the Admin Web UI to administer the Access Server, information about licensing and whether you are setting this up as a primary or secondary node (you will usually select primary unless using a failover setup).

5. Administrator uses Admin Web UI to complete configuration

The administrator uses a Web browser to open the URL of the OpenVPN Access Server, such as <https://vpn-gw.example.net/admin> or <https://x.x.x.x:943/admin>. The administrator logs in with the root username and password of the machine, and adjusts settings on the pages of the Admin Web UI. At a minimum, the administrator enters the license key on the **License** page and then starts the VPN Server.

6. User authenticates to the Connect Client

The user's Web browser opens a URL such as <https://vpn-gw.example.net> and the user signs on with a username and password. Once the user is authenticated, the Connect Client generates an OpenVPN client configuration file and a pre-configured OpenVPN-AS Windows Client GUI installer file specifically for that user and then allows that user to either connect through the interface or download the necessary certificates.

7. User connects to VPN

After the user has authenticated against the VPN Server the client software will initiate a connection. The user will see the connection status in their browser window. After the

connection has been established, the browser window will show the connection status and list the address of the server the user is connected to along with the amount of data that has been transferred between the users client and the vpn server. The systray icon will also show the connection status and will display a status message informing the user they are connected after the connection has been established.

2 OpenVPN Access Server Operation

This section elaborates on some of the characteristics of OpenVPN Access Server deployments and further describes the operation of several components of the Access Server.

2.1 Services and TCP/UDP Ports

The OpenVPN Access Server provides three network services:

Network Service	TCP/UDP	Default
VPN Server	TCP or UDP	TCP port 443 , if forwarding service for Connect Client UDP port 1194
Connect Client (HTTPS)	TCP	port 443 (via service forwarding) port 943 (direct)
Admin Web UI (HTTPS)	TCP	port 443 (via service forwarding) port 943 (direct)

Table 2 Access Server Services and Ports

The VPN Server is the daemon that creates the VPN tunnels with VPN clients. If TCP is configured as the protocol for VPN Server communication, the VPN Server can also forward services to the Connect Client and/or Admin Web UI

The Client Web Service is a secure Web service handling SSL-protected HTTP from Web browsers. Users log in to the Connect Client in order to download a pre-configured OpenVPN Windows client installer file or a client configuration file. The normal port for such traffic is TCP port 443.

The VPN Tunnel service can be configured to use either TCP or UDP. In the TCP case, it can also be configured to forward the Connect Client and/or Admin Web UI services. If service forwarding is used, only one TCP port needs to be made available to Internet clients. If applications requiring UDP communication (such as VoIP) are to be used over the VPN, configuring OpenVPN Access Server to use UDP for VPN Tunneling will result in the VPN tunnel communication being more efficient. In this case, the UDP port (number 1193, by default) on the server must also be made available to Internet clients.

2.2 Typical Network Configurations

The following sections describe the three most common supported network configurations used with OpenVPN Access Server deployments.

2.2.1 One Network Interface on Private Network Behind the Firewall

This configuration is most commonly seen when the Access Server resides in an internal corporate network, providing VPN access to users outside the corporate network. In this configuration the Access Server has one network interface connected to the private network (note that other interfaces may be present on the system but will not be utilized in by the Access Server). This scenario is illustrated in Figure 2.




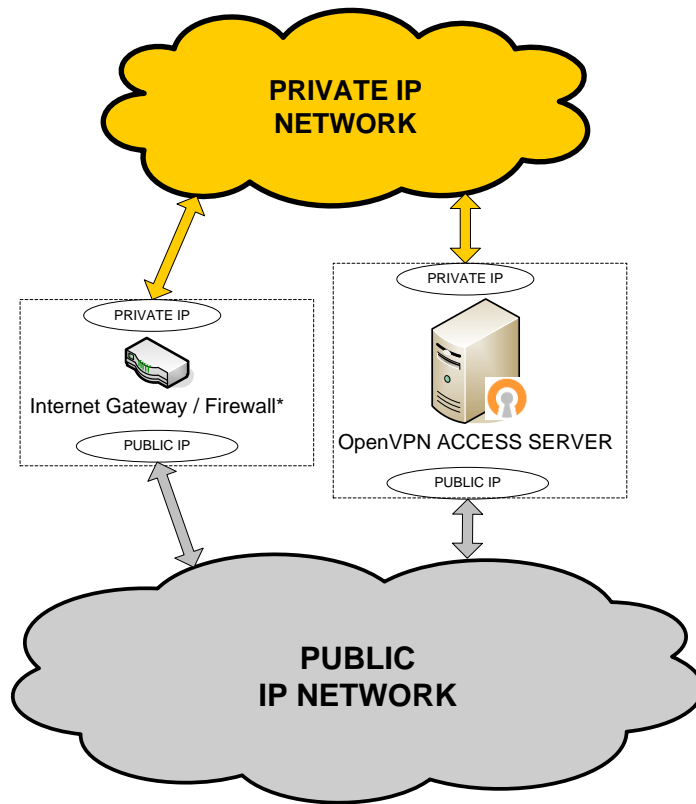
Figure 2: Access Server on Private Network Behind a Firewall

In this configuration, it is required that the Internet Gateway be set up to forward desired TCP/UDP port traffic from the public IP to the Access Server's private IP address. At a minimum, one TCP port (typically port 443) needs to be forwarded. That can carry both the VPN tunnel traffic and the Web Client Server/Connect Client traffic. Optionally, the VPN tunneling can be separated from the Web Client Server traffic, in which case an additional TCP or UDP port (e.g., UDP port 1193) must be forwarded for the VPN tunnel purposes.

A variation on this network configuration has the Access Server with one interface attached to a DMZ network provided by the firewall. The same forwarding of client traffic is required (as above); additionally, the firewall may need to be configured to allow traffic between the Access Server and the private network behind the firewall.

2.2.2 Two Network Interfaces, One on Public and One on Private Network

This configuration is most commonly seen when the Access Server resides in an internal corporate network but it also has its own public IP address (see Figure 3). The Access Server communicates with clients outside the corporate network via its public IP interface. It uses another network interface to communicate with hosts on the private IP network and to propagate packets between VPN tunnels and the private network.



* Internet Gateway is Optional

Figure 3: Access Server with Two Network Interfaces

2.2.3 One Network Interface on Public Network

This configuration is most commonly seen when the Access Server is located in a data center and its purpose is to create a virtual IP network to which all VPN clients can connect in order to communicate with services deployed on the server itself.




Figure 4: Access Server with One Network Interface on Public Internet

2.3 User Authentication and Management

OpenVPN Access Server can manage its own internal database and can also work with a variety of popular authentication methods. The currently supported systems are:

1. Local: Internal Database authentication
2. PAM: the system for authenticating users with accounts on the Access Server Linux host
3. Active Directory/LDAP Server
4. RADIUS Server(s)

The user authentication service may reside on the same server as the Access Server (as is always the case when PAM is chosen); or it can reside on a completely separate server, as long as the server is reachable by the Access Server via either the private or public network. A typical deployment with an external user database is shown in Figure 5 below.




Figure 5: Access Server Deployment with External User Database

2.4 Client Configuration Generation and Management

The client files (OpenVPN client configuration file and Windows software installer) for a particular user are automatically created when the user successfully logs in to the Connect Client. This process takes place without any need for interaction from the administrator, as long as the user can authenticate against the user database chosen by the administrator during installation of the Access Server.

If a user is disabled or deleted from the user authentication database, the user's VPN client becomes implicitly disabled due to the fact that the user can no longer authenticate successfully when the VPN client connects to the OpenVPN server. Thus, there is no need for the administrator to delete a user's configuration files on the Access server.

Also note that each generated client configuration is user-locked – it can only be used by that particular user. So a user that successfully signs on to the Connect Client cannot enable a different user to access the VPN simply by giving away the client configuration and/or Windows installer file (since the different user will not have the required user credentials).²

2.5 Virtual VPN Subnet Configuration

When deployed, the Access Server creates an independent, virtual VPN IP subnet on which each of the connected VPN clients is assigned an IP address³. If access to private networks is enabled by the administrator, the Access Server will also set up a NAT or internal routing system to allow VPN clients from the VPN subnet to reach the private network via the server's private IP address. An illustration of this system is shown in Figure 6.




Figure 6: OpenVPN Access Server Virtual VPN Subnet Configuration

² As of version 1.2.0, particular users can have the “auto-login” permission enabled (see the User Permissions page) which allows users to connect to the VPN without entering a password.

³ As of version 1.2.0, the Access Server may configure two virtual networks: one for “static” VPN IP addresses (i.e., the admin assigns specific VPN IP addresses to particular users) and one for “dynamic” VPN IP addresses.

3 Installation

This section describes in detail the steps for installing the OpenVPN Access Server.

Important Note:

The administrator should execute the commands listed below while running as root, and the use of the `bash` command shell is strongly recommended (instead of `csch`, `tcsh`, etc.).

3.1 Prepare the Server

Before performing the installation of OpenVPN Access Server, the following steps should be taken to prepare the server platform:

1. Ensure that SELinux is disabled (disabling SELinux requires a system reboot to take effect).
2. Configure the server with the interface IP address(es) and domain name desired. Ensure that the network settings will permit OpenVPN clients to access the Access Server, and that the server's domain name resolves properly to the desired interface address.

Completing the second step usually involves configuring the server in one of the following ways:

- a) The server has a static IP address that is reachable from clients on the Internet, at least for the TCP ports used by Access Server (see Section 2.1). Preferably, the server has a Fully Qualified Domain Name (FQDN) as its host name.
- b) The server has a dynamic IP address that is reachable by clients on the Internet and a dynamic DNS host name which tracks the changing IP address (this service is offered for free by various providers).

In either case, having the server located on a private network behind a corporate firewall implies that the firewall must be configured to forward client traffic (on the ports used by Access Server) between the public IP address and the server's private IP address.

3.2 Obtain License Key

A two concurrent user key is allocated to the OpenVPN Access Server by default for trial purposes.

3.3 Install OpenVPN Access Server RPM/DEB Package

Download the Access Server package from www.openvpn.net that is appropriate for your server operating system. Assuming the server runs a RedHat flavor of Linux, the package is an RPM file. For example, `openvpn-as-1.6.0-Fedora9.x86_64.rpm` would be appropriate for a 64-bit installation of Fedora 9.

Run one of the following commands, substituting the filename of the downloaded RPM or DEB file:

For Fedora/CentOS/RHEL hosts:

```
rpm -i openvpn_as_rpm_filename
```

For Ubuntu hosts:

```
dpkg -i openvpn_as_deb_filename
```

Once the package installation completes, you should see this message:

```
Please configure OpenVPN-AS by running /usr/local/openvpn_as/bin/ovpn-init
```

Note that if you ever have to remove the OpenVPN Access Server package, the command to use is one of the following:

For Fedora/CentOS/RHEL hosts:

```
rpm -e openvpn-as
```

For Ubuntu hosts:

```
dpkg -r openvpn-as
```

3.4 Run `ovpn-init`

OpenVPN Access Server is initially configured using an interactive configuration utility called `ovpn-init`. This utility prompts the administrator with a few questions in order to construct the desired Access Server initial configuration. Versions 1.6+ Already run this tool to the default settings, if you have a need to run it again you can use the following command:

```
/usr/local/openvpn_as/bin/ovpn-init --force
```

```
ovpn-init

      OpenVPN Access Server
      Initial Configuration Tool
-----
OpenVPN Access Server End User License Agreement (OpenVPN-AS EULA)

  1. Copyright Notice: OpenVPN Access Server License;
     Copyright (c) 2010 OpenVPN Technologies, Inc.. All rights reserved.
  2. Redistribution of OpenVPN Access Server binary forms and documents,
     are permitted provided that redistributions of OpenVPN Access Server
     binary forms and documents must reproduce the above copyright notice.
  3. You agree not to reverse engineer, decompile, disassemble, modify,
     translate, make any attempt to discover the source code of this software,
     or create derivative works from this software.

  4. THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESSED OR IMPLIED
     WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF
     MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO
     EVENT SHALL OPENVPN TECHNOLOGIES, INC BE LIABLE FOR ANY DIRECT, INDIRECT,
     INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
     LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA,
     OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF
     LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING
     NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS
     SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Please enter 'yes' to indicate your agreement [no]:
```

If you have already run `ovpn-init` previously, you will instead see an error message and `ovpn-init` will exit immediately:

```
Error: ovpn-init has already been run on this system. Use --force option.
```

You can force `ovpn-init` to re-initialize the Access Server configuration by running “`ovpn-init -force`”. Note that this will re-generate all keys and certificates used by Access Server and restore the configuration to the initial defaults.

```
Will this be the primary Access Server node?
(enter 'no' to configure as a backup or standby node)
> Press ENTER for default[yes]:
```

You will need to list whether this Access Server will be a Primary node or a Standby Node. Most users will choose Primary, you will only choose standby if you plan on using our Failover feature and have already configured a primary node for this feature.

3.4.1 Configure Initial Admin Web UI Network Settings

The main way to configure the Access Server is using its Admin Web UI. You must specify the network address(es) and port number to be used by the Web server that provides the Access Server's Admin Web UI. First, `ovpn-init` asks for the IP address:

```
Please specify the network interface and IP address to be
used by the Admin Web UI:
(1) all interfaces: 0.0.0.0
(2) eth0: 192.168.1.112
(3) eth1: 10.55.20.33
Please enter the option number from the list above (1-3).
> Press Enter for default [2]:
```

If you select “1” at the prompt, the Access Server Admin Web UI will listen on all available IP addresses (on the port specified in the next step). Otherwise, the Admin Web UI will be available on the IP address that you select; e.g., it will listen on 10.55.20.33 if option “3” is selected in the example above. You may wish to choose the IP address for the Admin Web UI so that it is only available on a private network.

Next, `ovpn-init` prompts for a TCP port number for the Admin Web UI:

```
Please specify the port number for the Admin Web UI.
> Press ENTER for default [943]:
```

You should use a port number that is not already used by a network service running on the Access Server host. Note that the Admin Web UI port number need not be given out to VPN client users, so it does not need to be a well-known port number.

Next, `ovpn-init` prompts for TCP port number for VPN connections. Make sure you use a port that is not in use by other services.

```
Please specify the TCP port number for the OpenVPN Daemon.
> Press ENTER for default [443]:
```

Next, `ovpn-init` prompts for a decision on whether or not you want all VPN traffic from clients routed (including internet traffic).

```
Should client traffic be routed by default through the VPN?
> Press ENTER for default [yes]:
```

Next, `ovpn-init` prompts for a decision on whether or not you want to use the RFC1918 private subnets which are:

```
10.0.0.0      - 10.255.255.255 (10/8 prefix)
172.16.0.0   - 172.31.255.255 (172.16/12 prefix)
192.168.0.0  - 192.168.255.255 (192.168/16 prefix)
```

```
Should RFC1918 private subnets be accessible to clients by default?
> Press ENTER for default [yes]:
```

Configure Initial Authentication Settings for Administrator

Once you have Access Server initially configured, you can use the Admin Web UI to set up User Authentication using RADIUS or LDAP. However, PAM is used initially to give an administrator access to the Admin Web UI. Thus, you must choose a username for a Unix account on the Access Server host to use when logging in initially to the Admin Web UI. You can either use the existing “root” account credentials or specify a different user account for this purpose. In the latter case, you can choose an existing account or create one within `ovpn-init`.

```
To initially login to the Admin Web UI, you must use a
username and password that successfully authenticate you
with the host UNIX system (you can later modify the settings
so that RADIUS or LDAP is used for authentication instead).
```

```
You can login to the Admin Web UI as 'root' with your existing
root password or specify a different user account to use for this
purpose. If you choose to use a non-root account, you can create
a new user account or specify an existing user account.
```

```
Do you wish to login as 'root'?
> Press ENTER for default [yes]:
```

Answering “yes” to this question indicates that you wish to use username “root” and the root account’s password when you initially log in to the Admin Web UI; `ovpn-init` will then be done asking questions and it will finalize the initial configuration.

Answering “no” to the above question will lead to the following prompt:

```
> Specify the username for an existing user or for the new user account:
```

At this prompt, you may type the username of an existing account on the Access Server host, or specify a username for an account you wish to create. If the username you type does not exist on the host, `ovpn-init` will prompt you for the password for the new user account (the username is “admin” in the example below):

```
Type the password for the 'admin' account:
Confirm the password for the 'admin' account:
```

The password characters you type are not echoed to the console during this step.

```
>Please specify your OpenVPN-AS license key (or leave blank to specify
later):
```

If you would like to activate a license through the `ovpn-init` you can do so at this option.

3.4.2 Finalize the Initial Configuration

Once you have supplied the necessary input to `ovpn-init`, it generates the initial Access Server configuration. The output seen during this step should be similar the following text:

```
Initializing OpenVPN...
Adding new user login...
useradd "admin"
Writing as configuration file...
Writing config.json...
Perform sa init...
Wiping any previous userdb...
Creating default profile...
Modifying default profile...
Adding new user to userdb...
Modifying new user as superuser in userdb...
```



```

Getting hostname...
Hostname: vpn-gw.example.net
Preparing web certificates...
Getting web user account...
Adding web group account...
Adding web group...
Adjusting license directory ownership...
Initializing userdb...
Generating init scripts...
Generating PAM config...
Generating init scripts auto command...
Starting openvpn-as...

NOTE: Your system clock must be correct for OpenVPN Access Server
to perform correctly. Please ensure that your time and date
are correct on this system.

Initial Configuration Complete!

You can now continue configuring OpenVPN Access Server by
directing your Web browser to this URL:

https://10.55.20.33:943/admin

Login as "admin" with the same password used to authenticate
to this UNIX host.

See the Release Notes for this release at:
http://www.openvpn.net/access-server/rn/openvpn\_as\_1\_6\_0.html

```

Note that near the end of the output text, the URL for the Admin Web UI is displayed (in the above example the URL is **https://10.55.20.33:943/admin**).

3.5 Configure Access Server with the Admin Web UI

Once **ovpn-init** completes, you can access the Admin Web UI by entering the its URL into your Web browser. The URL to use will end with “/admin” and it is shown near the end of the output of **ovpn-init**.

When you initially connect to the Admin Web UI (and/or the Connect Client), your browser will display a warning regarding the server certificate. This warning is to be expected, and it is due to the (automatically-generated) certificate for the Access Server ‘s Web components *not* being issued from a Certificate Authority (CA) that is already trusted by your Web browser. See Section 5.4 for information on preventing this browser security warning). After you instruct your Web browser to go ahead and connect to the secure server, you should see the login prompt shown in Figure 7.



Figure 7: Login page for Admin Web UI

Enter the credentials for the administrative user you specified during the `ovpn-init` step. Note that since PAM is initially selected for Access Server user authentication, the administrator's password is the same one as used to authenticate to the Access Server's Unix host. Once you have successfully authenticated, you see the **Status Overview** page of the Admin Web UI (see Figure 8). The “Welcome to the Access Server Admin UI” message box is seen only the first time that you use the Admin Web UI.

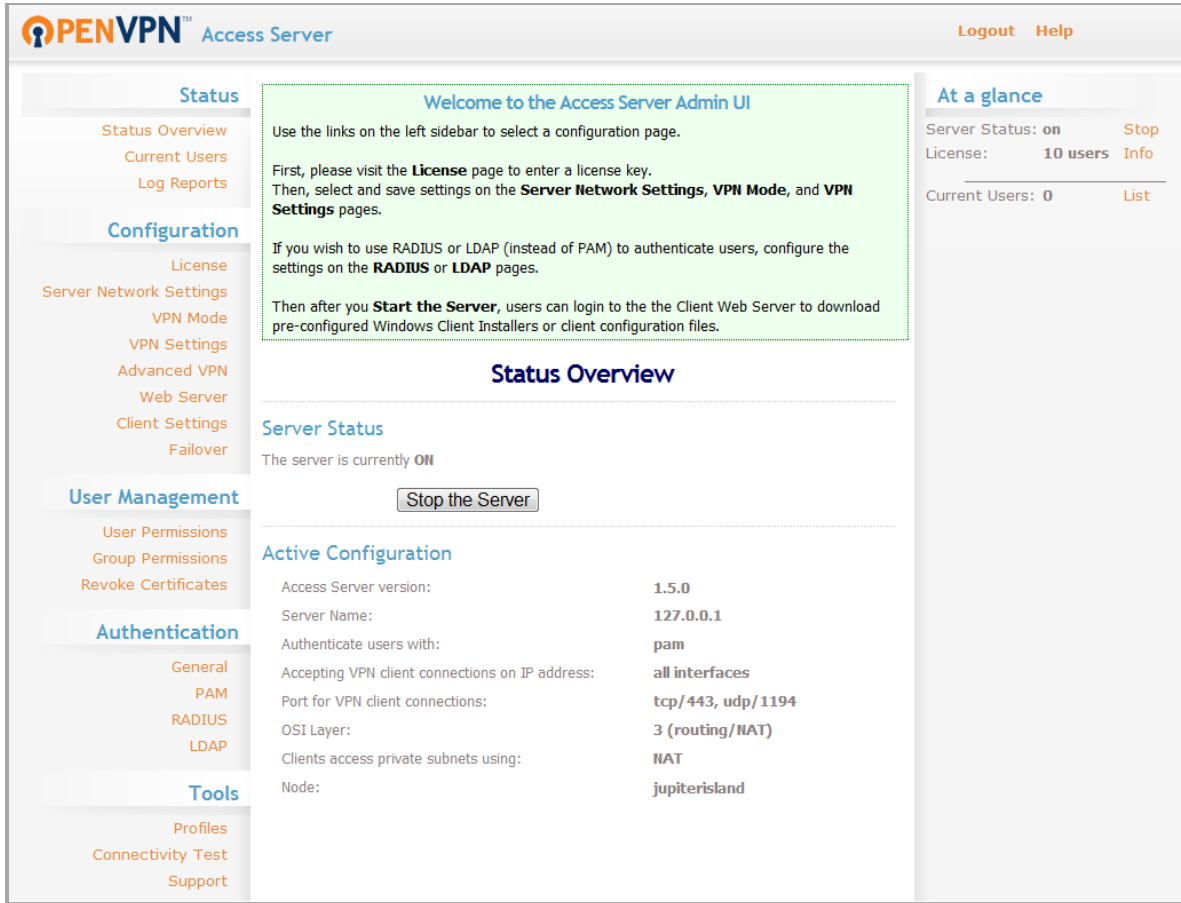


Figure 8: Landing page (Status Overview) for Admin Web UI

You may configure the Access Server using its Admin UI pages in any order, though the “Welcome” message suggests a minimal set of configuration steps.

The two sidebars are seen in all Admin UI pages:

The left sidebar contains links (orange text) to each Admin UI page, with the page links grouped under (blue text) headings. See Figure 9.

The right sidebar (see Figure 10) contains an “At a glance” display of the some of the Access Server properties, along with links (orange text) to the page in the Admin UI for performing an action related to the property:

VPN Server status (on or off), with a link to Start or Stop the VPN Server, depending on its status
the number of concurrent VPN users allowed by the installed license key(s), with a link to the License page for more information

the count of currently-connected VPN users, with a link to the Status Overview page (which contains a table listing all currently-connected VPN users)

Note that the “At a glance” right sidebar is not seen on some Admin UI pages with wide output (specifically the **Log Reports** page).

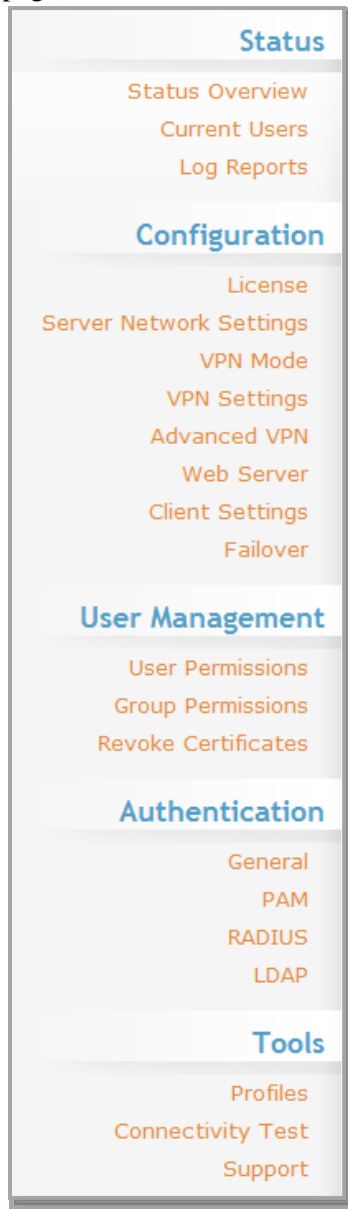


Figure 9: Left Sidebar (Page Menu List) of Admin Web UI




Figure 10: Right Sidebar (“At a glance” panel) of Admin Web UI

4 Admin Web UI Reference

This section describes each page of the Admin Web UI. Since most pages contain several “panels” with grouped settings, the descriptions cover each page’s panels individually.

4.1 Status Pages

4.1.1 Status Overview



The screenshot shows the 'Status Overview' page. At the top, the title 'Status Overview' is centered. Below it, the 'Server Status' section indicates 'The server is currently ON' and features a 'Stop the Server' button. The 'Active Configuration' section lists various settings in a key-value format:

Access Server version:	1.6.0
Server Name:	67.228.212.68
Authenticate users with:	local
Accepting VPN client connections on IP address:	all interfaces
Port for VPN client connections:	tcp/443
OSI Layer:	3 (routing/NAT)
Clients access private subnets using:	NAT
Node:	jupiterisland

Figure 11: Status Overview page

The Status Overview page has three main components:

The Server Status section shows you whether the VPN Server is currently ON or OFF. Based on the current status, you can either Start the Server or Stop the Server with a button near the top of the page.

The Configuration section displays a few of the more important Access Server configuration settings. These settings are not modifiable on the Status Overview page.

The Current VPN Users section displays a list of users that are currently connected to the VPN Server, or '(none)' if either the VPN Server is not running or if it is running but no users are connected. A column of buttons under the “Block” heading allows the administrator to immediately disconnect and disable access for a user that is currently connected to the VPN (provided that the user is not the administrator viewing the page). When this button is pushed for a given user, the user is added to the User Permissions configuration (if the username is not already present in that configuration) with the “Deny” property enabled. Thus, you can re-enable access to such a user by unchecking the user’s “Deny” checkbox on the User Permissions page.

4.1.2 Log Reports

The **Log Reports** page lets you query the Access Server's log database to see auditing information that meets the constraints you specify with query inputs.

The screenshot shows the 'Log Reports' interface. At the top, there are three input fields: 'Username:', 'Real IP Address:', and 'VPN IP Address:'. Below these are instructions: 'May contain multiple usernames separated by commas, such as mark,joe' and 'May contain '%' wildcard character, for example: 192.168.%'. The 'Start Time Range:' section has three radio buttons: 'Do not filter log on Start Time' (selected), 'Within the last [] hours', and 'From [] to []'. Below this is a note: 'Specify start and/or end times in MM/DD/YYYY format or MM/DD/YYYY HH:MM format such as 08/03/2009 14:00'. The 'Services:' section has three radio buttons: 'All' (selected), 'VPN', and 'WEB_ADMIN'. The 'Limit output to the' section has two radio buttons: 'first' and 'last' (selected), with a text input field containing '30' and the text 'log entries.' below it. A 'Query Log' button is centered below the filters. To the right, there is a link: 'Download current log report (below) in CSV format'. Below the link is a table with the following data:

Username	Start Time	Duration	Service	Real IP	VPN IP	Bytes In	Bytes Out	Error
root	09/22/09 . 16:16:40		WEB_ADMIN	192.168.232.1				
root	09/22/09 . 16:19:27		WEB_ADMIN	192.168.232.1				
root	09/23/09 . 03:31:34		WEB_ADMIN	192.168.232.1				
root	09/29/09 . 00:02:55		WEB_ADMIN	192.168.232.1				
root	09/29/09 . 14:22:04		WEB_ADMIN	192.168.232.1				

Figure 12: Log Reports page

You can customize your query by specifying input constraints:

Username: a single username or multiple usernames separated by commas (blank for 'any username').

Real IP address: a public IP address, possibly including a % character as a wildcard (blank for 'any Real IP address').

VPN IP address: a public IP address, possibly including a % character as a wildcard (blank for 'any VPN IP address').

The time range to use for the Start Time (also relative to the host local time).

The service of interest (All, VPN, WEB_CLIENT or WEB_ADMIN).

The number of log entries to display, and whether to display the beginning or end of the matching entries from the log database.

Username	Start Time	Duration	Service	Real IP	VPN IP	Proto	Port	Bytes In	Bytes Out	Error
root	11/11/09 . 02:00:15		WEB_ADMIN	10.7.31.3						
root	11/11/09 . 02:19:38		WEB_CLIENT	10.7.31.160						
root	11/11/09 . 02:21:16		XML_API							
test	11/11/09 . 02:22:15		XML_API							
test	11/11/09 . 02:22:17	00:07	VPN	10.7.31.160	10.8.0.99	UDP	1194	5.84 KB	7.94 KB	
test	11/11/09 . 02:27:08		XML_API							
test	11/11/09 . 02:27:11	02:36	VPN	10.7.31.160	10.8.0.130	UDP	1194	68.33 KB	73.39 KB	
root	11/11/09 . 05:03:15		XML_API							
root	11/11/09 . 05:03:17	00:03	VPN	10.7.31.160	10.8.0.100	UDP	1194	3.61 KB	5.31 KB	
test	11/11/09 . 05:03:27		XML_API							
test	11/11/09 . 05:03:29	00:10	VPN	10.7.31.160	10.8.0.162	UDP	1194	7.71 KB	8.72 KB	
test	11/11/09 . 05:18:59	00:00	VPN	10.7.31.160	10.8.0.2	TCP	443	3.74 KB	4.26 KB	

Figure 13: Sample Log Report on Log Reports page

A given log report displays the following information for each output log entry:

Username of the user.

Start time for the VPN connection or Web session. Time is measured as local time on the Access Server host, not GMT/UTC.

The duration of the VPN connection (empty when the log entry is for a Web session).

The service type: VPN is for VPN connections, WEB_CLIENT is for logins to the Connect Client, and WEB_ADMIN is for logins to the Admin Web UI. XML_API is for connections to the server using REST (i.e; Server-locked Profile, Third Party Plug-ins).

The real IP address of the client.

The VPN address set for the client, when the service is VPN for the log entry.

Proto displays the Protocol used by the VPN Client to connect to the VPN Server.

The Port section displays which port was used to connect to the VPN.

The Bytes In and Bytes Out for the VPN connection, when the service is VPN.

Any error message that occurred for the service access.

You can also export the displayed log report in CSV (comma-separated value) format, for use with spreadsheet software such as Excel.

Note: The `logdba` command line utility (in `/usr/local/openvpn_as/scripts`) has the same functionality as, and even more flexibility than, the Log Reports page in the Admin Web UI. Run `logdba --help` to see the detailed usage information for that CLI utility.

4.2 Configuration Pages

4.2.1 License

License Manager

This page shows licenses activated for your Access Server, and allows you to add new licenses.

Installed License Keys

The license keys that are activated for this Access Server installation are shown below.

License Key	Type	Concurrent Users
B034-Q0TU-P7LH-JFZ4	purchased	30

Concurrent User Limit:

Licensed for 30 concurrent users.

Add A New License Key

License keys can be obtained by registering and logging in at the [Access Server Downloads](#) page.

New License Key:

Figure 14: License Manager page

Before you can successfully Start the VPN Server, a license must be obtained and activated. Licenses are conveyed in the form of **License Keys** a string that looks like:

```
TTAK-3BSH-V9DU-JS9J
```

Once you register and login at the [Access Server Downloads](#), you can obtain a license key for your Access Server installation.

A license key is activated for a particular server host, and that host must be able to access the Internet for activation when the license key is added using the **Add A New License Key** button.

Each license key carries a maximum number of concurrent VPN users. E.g., if your Access Server has one activated license key for 5 concurrent users, then no more than 5 users are allowed to use the VPN at any given time. If you have multiple license keys activated for a given Access Server installation, the combined maximum concurrent user limit will depend on the type of license keys installed.

Purchased Licenses: The concurrent user count for purchased licenses are additive. So if you purchase multiple licenses, the combined concurrent user limit will be the **sum** of the individual user counts for each purchased license.

Free Licenses: The concurrent user count for free licenses are not additive. So if you activate multiple free licenses, the combined concurrent user limit will be the **maximum** of the individual user counts for each free license.

The **License Manager** page lets you view the license keys that are activated, view the **Concurrent User Limit**, and also **Add A New License Key**.

4.2.2 Server Network Settings

The Server Network Settings page contains networking settings for the three network servers comprising the Access Server: the VPN Server, the Admin Web UI, and the Client Web Server.

4.2.2.1 VPN Server

Server Network Settings

VPN Server

Warning: Changing the Hostname, Protocol or Port Number after VPN clients are deployed will cause the existing clients to be unusable (until a new client configuration or VPN installer is downloaded from the Client Web Server)

Hostname or IP Address:

Interface and IP Address

Listen on all interfaces

eth0: 10.7.31.144

Protocol

TCP

UDP

Both (Multi-daemon mode)

Port number:

Service Forwarding

When TCP or Multi-daemon mode is chosen for the VPN Server protocol, the VPN Server can optionally provide access to these services through its IP address and port:

Admin Web Server

Client Web Server

Note: Services are only forwarded when the VPN Server is running.

Figure 15: VPN Server panel of the Server Network Settings page

The settings in the VPN Server section directly affect the contents of the client configuration files issued to VPN Clients (via the Connect Client). Thus, these settings should not be altered once users have downloaded VPN Client installers and/or configuration files. Any modification to the **Hostname or IP Address**, **Protocol**, or **Port Number** will invalidate existing client configurations. If this is done the user will need to redownload their VPN profile.

The **Hostname or IP Address** is the name or IP address that VPN Clients will use to access the VPN Server. Thus, it must be a public IP address or Fully-Qualified Domain Name (FQDN). It is strongly recommended that a FQDN be used for this setting.

The **Protocol** specifies whether TCP or UDP is used for VPN client-to-server communication. In general, Multi-Daemon mode is preferred as this gives you the option to use https:// without the specification of the port when accessing the Admin or Connect Client as well as the use of UDP when connecting to the VPN. By default, when multi-daemon mode is enabled, the client will first try to connect via UDP and if it is unable to connect via UDP it will attempt to connect via TCP. The reason UDP is first tried is because you can run into TCP meltdown when accessing certain TCP traffic over the VPN tunnel.

Multi-Daemon Mode: On the Server Network Settings page, you will see a new protocol option "Both" along with TCP and UDP. When "Both" is checked, multi-daemon mode is enabled. In this mode, you will see some new fields appear on the Server Network Settings page that allow you to define the TCP and UDP ports, and how many daemons to run concurrently for each. In addition, when "Both" is checked, the Connect Client will return OpenVPN client configurations that will adaptively try UDP first, then fail over to TCP if no response is received within 4 seconds. The connectivity test has also been extended to handle multi-daemon mode -- both TCP and UDP ports will be tested for connectivity when multi-daemon mode is enabled.

Interface and IP Address

Listen on all interfaces

eth0: 10.7.31.189

Protocol

TCP

UDP

Both (Multi-daemon mode)

Multi-Daemon Mode

In Multi-Daemon mode, the Access Server will load-balance connecting VPN clients across multiple OpenVPN daemons to fully leverage the capability of multi-core servers. NOTE: It is not recommended to set the number of TCP and UDP daemons to a higher value than the number of processor cores on the machine. Doing so may result in resource exhaustion and system instability.

Number of TCP daemons:

TCP Port number:

Number of UDP daemons:

UDP Port number:

Service Forwarding

When TCP or Multi-daemon mode is chosen for the VPN Server protocol, the VPN Server can optionally provide access to these services through its IP address and port:

Admin Web Server

Client Web Server

Note: Services are only forwarded when the VPN Server is running.

Service Forwarding is selectable when the VPN Server Protocol is chosen to be TCP. The Admin Web UI and/or the Client Web Server can be made available through the VPN Server. Of course,

when the VPN Server is OFF, these Web servers are not available through Service Forwarding (though they can still be accessed via their configured IP address and port number).

4.2.2.2 Admin Web UI

Admin Web UI

The IP address and port number for the Admin Web Server may be the same as for the Client Web Server.

Interface and IP Address

- Listen on all interfaces
- eth0: 10.7.31.144
- localhost: 127.0.0.1

Port number:

Figure 16: Admin Web UI panel of the Server Network Settings page

The Admin Web UI can listen on the same address and port as the Client Web Server, or it can listen on a separate address and port. In either case, the IP address and port of the Admin Web UI Server must not be the same as the address and port of the VPN Server (use **Service Forwarding** to access the Admin Web UI through the address and TCP port of the VPN Server).

4.2.2.3 Client Web Server

Client Web Server

Users login to the Client Web Server to obtain an auto-generated VPN config or customized VPN Client Installer.

- Use the same address and port as the Admin Web Server
- Use a different IP address or port:

Interface and IP Address

- Listen on all interfaces
- eth0: 10.7.31.144
- localhost: 127.0.0.1

Port number:

Figure 17: Client Web Server panel of the Server Network Settings page

The Client Web Server can be combined with the Admin Web UI, using the same IP address and port number. Alternatively, you can specify an IP address and port number for the Client Web Server that are different than those of the Admin Web UI.

When the Admin Web UI and Client Web Server listen on the same IP address and port number (as in the case that both services are forwarded through the VPN Server), the Client Web Server is accessed by the base URL ('/') while the Admin Web UI is accessed by the URL '/admin'. E.g., if the two Web Servers are configured to use IP address 192.168.1.20 and port 443, the Client Web Server is accessed by the URL:

`https://192.168.1.20/`
and the Admin Web UI is accessed by the URL:
`https://192.168.1.20/admin`

4.2.2.4 SSL Options for Web Services

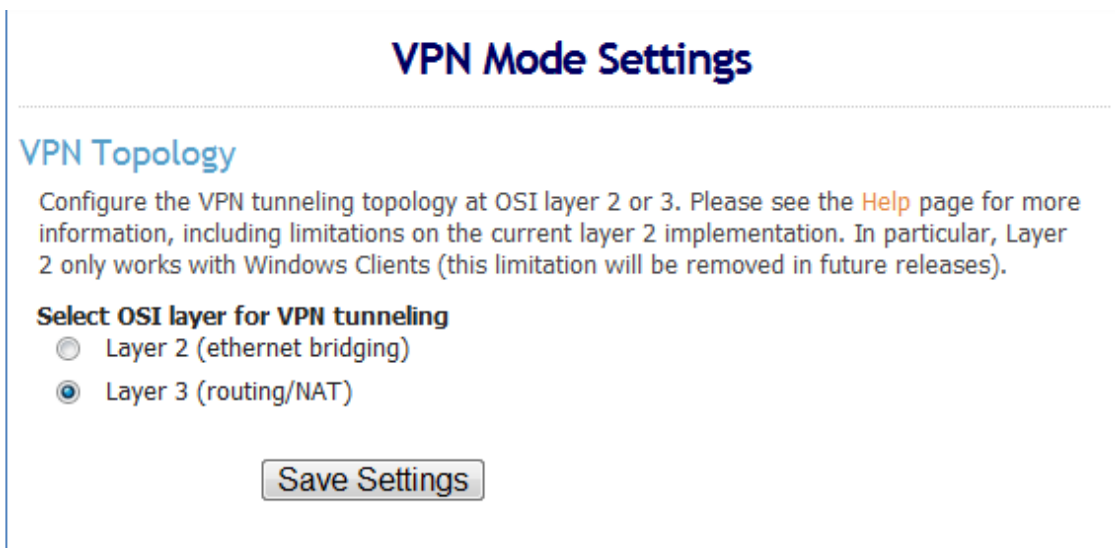


Figure 18: SSL Options panel of the Server Network Settings page

By default, the Web services (Admin Web UI and Client Web Server) components of the Access Server support the broadest set of SSL/TLS options, which includes SSLv2 cipher suites. Recently, SSLv2 has fallen into disfavor as a less secure version of the SSL/TLS protocol. Thus, to maximize communications security for your deployment you may wish to disable SSLv2 along with its corresponding weak cipher suites. Note, however, that interoperation with certain older Web browsers (specifically, Internet Explorer v6) may suffer as a result of disabling SSLv2.

4.2.3 VPN Mode

VPN Mode allows you to use either Layer 2 Ethernet Bridging or Layer 3 (Routing/NAT)




If you select Layer 2 Bridging your VPN Settings Page will only allow you to set whether you want Internet Traffic routed through the VPN since everything else would be configured from your external router.

***NOTE:** At the time of this writing, only Windows Clients support Layer 2 Ethernet Bridging Mode

4.2.4 VPN Settings

4.2.4.1 VPN IP Network



The screenshot shows the 'VPN Settings' page with a sub-section titled 'VPN IP Network'. Below this title is a descriptive sentence: 'Specify the addresses and netmasks for the virtual networks created for VPN clients'. There are three main sections for configuration:

- Dynamic IP Address Network:** Includes a descriptive sentence: 'When a user does not have a specific VPN IP address configured on the User Permissions page, the user's VPN client is assigned an address from this network.' It has two input fields: 'Network Address' with the value '5.5.5.0' and 'Number of Bits in Netmask' with the value '24'.
- Static IP Address Network (Optional):** Includes a descriptive sentence: 'Any static VPN IP addresses specified for particular users on the User Permissions page must be within this network. It must be different than the Dynamic IP Address Network above.' It has two empty input fields: 'Network Address' and 'Number of Bits in Netmask'.
- Group Default IP Address Network (Optional):** Includes a descriptive sentence: 'When a group does not have a specific Dynamic IP Address pool setting, the dynamic IP address pool for the group will be allocated from this list of subnets.' It has a large empty text area for input.


Figure 19: VPN IP Network panel of the VPN Settings page

When a VPN Client successfully connects to the Access Server, it is assigned a unique IP Address on the virtual VPN IP Network. If the user associated with the VPN Client has a specific, 'static' VPN IP address specified (on the User Permissions page), then that IP address is assigned to the connecting VPN Client. Otherwise, the VPN Client is automatically assigned a VPN IP address from the **Dynamic IP Address Network**.

For both the **Dynamic IP Address Network** and the **Static IP Address Network** you specify the network by defining the **Network Address** and **Number of Bits in Netmask**. Note that the number of bits in the netmask determines an upper bound on the maximum number of VPN Clients that may concurrently use the subnet. E.g., a 24-bit netmask yields a maximum of 254 simultaneous

VPN Clients (provided that the Access Server License also allows this number of concurrent users).

4.2.4.2 Routing



Routing

Should VPN clients have access to private subnets (non-public networks on the server side)?

No

Yes, using NAT

Yes, using routing (advanced)

Specify the private subnets to which all clients should be given access (as '*network/netmask_bits*', one per line):

```
10.33.5.0/24
10.44.0.0/16
```

Should client Internet traffic be routed through the VPN?

No

Yes

Should clients be allowed to access network services on the VPN gateway IP address?

No

Yes

Figure 20: Routing panel of the VPN Settings page

The Access Server can enable VPN clients to access private subnets available on the Access Server host. The options for the **Should VPN clients have access to private subnets?** setting are as follows:

No: VPN Clients are not allowed to access any private subnet.


Yes using NAT: VPN Clients can access private subnets, and each VPN Client's virtual address is transformed via NAT so that the Access Server host's IP address is used as the source address on client packets destined for private subnets.

Yes, using routing (advanced): VPN Clients can access private subnets, and it is the virtual address of each VPN Client that is used as the source address on client packets destined for private subnets. Routing must be configured on hosts on the private subnet(s) so that response packets can be routed back to the VPN Clients via the Access Server host's IP address on the private subnet.

Note: NAT is usually preferred for allowing VPN Clients access to private subnets. Routing is more complicated to configure, as it requires routing changes on the network infrastructure. Routing is offered to accommodate applications that do not function properly through NAT. When one of the **Yes** options above is selected, the private subnets must be specified. You can enter multiple subnets, each specified as a network/netmask_bits pair such as **10.33.4.0/24** on a separate line in the textbox.

When **Yes** is selected for the **Should clients' Internet traffic be routed through the VPN?** setting, the default route on a newly-connected VPN Client host is set to point to the VPN gateway's virtual IP address. This setting prevents 'split tunneling'. All network traffic on the VPN Client host flows through the Access Server (with the client's Internet traffic going through the Access Server's public IP address)

4.2.4.3 DNS Settings



DNS Settings

Pushing DNS servers to clients is optional, unless clients' Internet traffic is to be routed through the VPN

- Do not alter clients' DNS server settings
- Have clients use the same DNS servers as the Access Server host
- Have clients use these DNS servers:

Primary DNS Server:


Secondary DNS Server:

Figure 21: DNS Settings panel of the VPN Settings page

When a client connects to the VPN, the its DNS settings may be altered so the client can resolve names of hosts on the private network. When **Have clients use the same DNS servers as the Access Server host** is selected, the VPN clients' DNS settings are altered so that the client resolves names using the DNS servers configured for the Unix host running Access Server (typically specified in the host's `/etc/resolv.conf` file).

You can also specify particular DNS servers for the VPN clients to use. You must then configure the IP address of the primary DNS server, and optionally the IP address of a secondary DNS server. Note that when **127.0.0.1** is specified, the VPN address of the Access Server host (e.g., '10.8.0.1') is pushed to VPN clients as the DNS server.

4.2.4.4 Default Domain Suffix



Default Domain Suffix (optional)


Default domain suffix:

Figure 22: Default Domain Suffix panel of the VPN Settings page

Optionally, you can specify the default domain suffix that clients use when they are connected to the VPN. For example, a default domain suffix of 'example.net' would cause a VPN client to resolve the host name 'foo' as if it were 'foo.example.net'.

4.2.5 Advanced VPN

4.2.5.1 Inter-Client Communication



The screenshot shows a settings panel titled "Advanced VPN Settings" with a sub-section "Inter-Client Communication". The main question is "Should clients be able to communicate with each other on the VPN IP Network?". There are three options: "Yes" (radio button), "No (unless explicitly allowed by User Permissions)" (radio button, selected), and "Allow VPN users with Administrator privilege to access all VPN client IP addresses" (checkbox).

Figure 23: Advanced VPN Settings panel of the Advanced VPN Settings page

When **Yes** is selected for the **Should clients be able to communicate with each other?** setting, packets can be exchanged between individual VPN clients on the VPN virtual subnet. When this option is set to **No** there is the additional option of only allowing users with Administrator privileges (as configured on the User Permissions page) to access all client VPN IP addresses.

4.2.5.2 Multiple Sessions per User




The screenshot shows a settings panel titled "Multiple Sessions per User". There is one option: "Allow multiple concurrent VPN connections for a user (automatically disabled when static VPN IP addresses are configured for users)" with a checked checkbox.

Figure 24: Multiple Sessions per User panel of the Advanced VPN Settings page

When this option is enabled, a single VPN user can establish multiple, concurrent VPN connections. In terms of licensed user limits, each such connection is counted as a separate concurrent user. Note that when a user is configured to have a static IP address (on the User Permissions page), that user cannot have multiple concurrent VPN connections, even if this setting is enabled.

4.2.5.3 Default LZO Settings


Figure 25: Default LZO Settings



This feature will enable LZO for all users on the OpenVPN Access Server. LZO is a lossless data compression library. There are many benefits when using LZO compression with OpenVPN Access Server. The results of using realtime compression with OpenVPN Access Server are less bandwidth usage, quicker transmission times (files reach the destination faster since there is less data being transferred). LZO is a real-time compression.

4.2.5.4 Private Routed Subnets

Figure 26: Private Routed Subnets



This option allows administrators the ability to set certain subnets as routed subnets versus natted subnets. Please keep in mind that you cannot set a subnet to use both natting and routing (you cannot list a private routed subnet as a natted subnet in the Routing box on the VPN Settings page).

4.2.5.5 Connection Security Refresh



Figure 27: Connection Security Refresh panel of the Advanced VPN page

At regular intervals, the Access Server renegotiates a TLS session with a given VPN client. This is to maintain the security of the TLS connection. The refresh interval is specified as a number of minutes. 60 minutes is a reasonable default; less than 10 minutes is not recommended.

During the security refresh, the VPN Client user is re-authenticated, however VPN Clients may cache the user's credentials and make this re-authentication go unnoticed by the user.

4.2.5.6 Windows Networking

Windows Networking (optional)

Don't alter Windows networking settings on clients

Disable NetBIOS over TCP/IP on clients

Enable NetBIOS over TCP/IP and use these Windows networking settings on clients:

Primary WINS Server:

Secondary WINS Server:

NetBIOS over TCP/IP Node type:

b-node (broadcasts)

p-node (point-to-point name queries to a WINS server)

m-node (broadcast, then query name server)

h-node (query name server, then broadcast)

NBDD Server:

NBS Scope-ID:

Figure 28: Windows Networking panel of the Advanced VPN page

To allow Windows VPN clients to convert NetBIOS host names into IP addresses, you can configure the IP address of a primary (and optionally, a secondary) WINS server. You can also specify the Node type for the NetBIOS over TCP/IP communication. Additionally, you can specify the IP address of a NBDD (NetBIOS over TCP/IP Datagram Distribution server).

Optionally, the NetBIOS over TCP/IP Scope ID can be specified as a character string (which is appended to a NetBIOS name). The use of NetBIOS Scope IDs allow computers to use the same (NetBIOS) computer name, as long as have different Scope IDs.

4.2.5.7 Additional OpenVPN Config Directives

Additional OpenVPN Config Directives (Advanced)

You can specify OpenVPN directives to be added to client and/or server configuration files used with the Access Server. Enter the desired directives, one per line, in one or both of the textboxes below.

Note: If a line begins with the minus sign, then the directive is **removed** from the configuration. An asterisk can be used as a wildcard to match 0 or more characters. For example, this config directive removes the 'duplicate-cn' directive:

-duplicate-cn

Warning: Minimal checking is performed on the supplied directives. Be sure you are using valid OpenVPN directives, as defined in the [OpenVPN 2.1 Manual](#)

Server Config Directives

```
push "route-delay 6"
```

Client Config Directives

```
keepalive 10 60
```

Figure 29: Additional OpenVPN Config Directives panel on Advanced VPN page

If you are already familiar with OpenVPN 2.1 configuration directives, and you need to use OpenVPN options that are not handled by the Access Server Admin Web UI, you can specify the configuration directives in these textboxes. The specified config directives are then added to the VPN configuration files generated by Access Server.

Beware that there is no meaningful checking of the supplied options until the VPN server or client attempt to use the configuration file. Also note that some OpenVPN options, such as bridging, may not be usable with Access Server (as they require special handling due to interaction with Access Server functionality).

4.2.6 User Permissions

User Permissions

Search By User Name or users in group by Group Name

Username	Group	More Settings	Admin	Allow Auto-login	Deny Access	Delete
root	No Default Group	Show	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	
test1	No Default Group	Hide	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Select IP Addressing : Use Dynamic Use Static

VPN Static IP Address:

VPN Gateway

Configure VPN Gateway: No Yes

Allow client to act as VPN gateway for these client-side subnets:

List subnets in *network/nbits* form

DMZ settings

Configure DMZ IP address: No Yes

DMZ IP Address:

List of IP address in *network.tcp/port or udp/port*

test2	No Default Group	Show	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
test3	No Default Group	Show	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

New Username:

	No Default Group	Show	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	
--	------------------	------	--------------------------	--------------------------	--------------------------	--

Figure 30: User Permissions page

Though the Access Server does not store user credentials, it does maintain permission information for particular users, as defined by their usernames. The following permissions may be specified for particular users:

Administrator: The user is allowed to use the Access Server Admin Web UI.

Allow Auto-login: The user can choose a client configuration that enables connecting to the VPN without authenticating with a password (useful for unattended VPN clients residing on network server hosts. When an autologin profile is used this will allow

Deny Access: The user is not permitted to use the VPN.

Additionally, you can use the Show link to show the drop-down list of other settings that can be configured for the user, including:

VPN Static IP Address: The IP address to be given to the user when the VPN connection is made. This address must be within the Static IP Address Network configured on the VPN Settings page.

VPN Gateway: You can configure a user to act as a gateway for the VPN server. The Gateway option allows users from the VPN to access local machines and services on the LAN that the VPN Gateway sits on.

DMZ Settings: This option allows you to make a users full machine or single service available over the VPN's Public IP

You can add new usernames to the User Permissions table with the New Username: textbox. Modify the user permissions by checking or unchecking the permissions box for the desired user in the table, or delete the user's reference with the Delete checkbox. Note that when a user's reference is deleted from the User Permissions page, the user's account may still be active in the authentication system (PAM, RADIUS or LDAP). In particular, the user (whose reference was removed from the User Permissions page) may still be able to successfully authenticate to the authentication system and thus also be able to connect to the Access Server's VPN.

By default, the Access Server allows VPN access to any successfully authenticated user. However, the Deny access to all users not listed above setting allows you to prevent VPN access to any user that is not listed in the User Permissions table.

4.2.7 Group Permissions

Group Permissions

Group	More Settings	Admin	Allow Auto-login	Deny Access	Delete
Sales	Show	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Engineers	Hide	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

VPN IP Addresses

Subnets assigned to this group:

List subnets in *network/nbits* form

Dynamic IP Address Pool: From: To:

Example: From: 10.6.0.10 To: 10.6.0.100

Access Control

Use Access Control? No Yes

Client Scripting

Use Client Scripting? No Yes

Environment Variables [\[Windows\]](#) [\[Mac\]](#) [\[Linux\]](#)

WINDOWS SCRIPT [\[user-connect\]](#) [\[user-disconnect\]](#) [\[admin-connect\]](#) [\[admin-disconnect\]](#)

MAC SCRIPT [\[user-connect\]](#) [\[user-disconnect\]](#) [\[admin-connect\]](#) [\[admin-disconnect\]](#)

LINUX SCRIPT [\[user-connect\]](#) [\[user-disconnect\]](#) [\[admin-connect\]](#) [\[admin-disconnect\]](#)

Default Group Permissions to use for any User not in any Group:

Figure 31: Group Permissions page

Subnets assigned to this group: You will need to assign a subnet for this group to use for IP Addressing.


Dynamic IP Address Pool (VPN IP Addresses): You can list a range in the assigned subnet for OpenVPN Access Server to use for clients who are acquiring Dynamic IP Addresses, be sure not to use any of the IP's in this range for Static Addressing.

Access Control: Access Control will allow you to give the group access to certain subnets, services, groups and users.

Client Scripting: Access Server can run certain client scripts for different clients on different Operating Systems. This could be as simple as mounting a network drive to writing python scripts.

4.3 Authentication Pages

4.3.1 General



User Authentication

User credentials are validated using one of the three (external) user databases below or using the locally configured users on 'Users Permissions' page.

Authenticate users using:

- Local
- PAM
- RADIUS
- LDAP

Figure 32: General User Authentication page

Access Server does not store and manage user credentials. Instead, it interfaces with one of the following systems for user authentication:

Local: This is a user authentication system that is managed by OpenVPN Access Server. You can set the vpn users password on the User Permissions page when Local Authentication is enabled.

PAM: Pluggable Authentication Modules - The system used to authenticate users to the Unix host running Access Server.

RADIUS: Between one and five RADIUS servers can be contacted for user authentication and (optionally) also user accounting.

LDAP: An Active Directory domain controller or other LDAP server is used to validate user credentials.

On the General User Authentication page, you can choose between the three methods of authenticating Access Server users.

This setting can also be changed on the configuration pages for PAM, RADIUS and LDAP (e.g., on the RADIUS page, you can press Use RADIUS if RADIUS isn't already chosen). Note, however, that only one authentication method type can be chosen for Access Server user authentication.

4.3.2 PAM

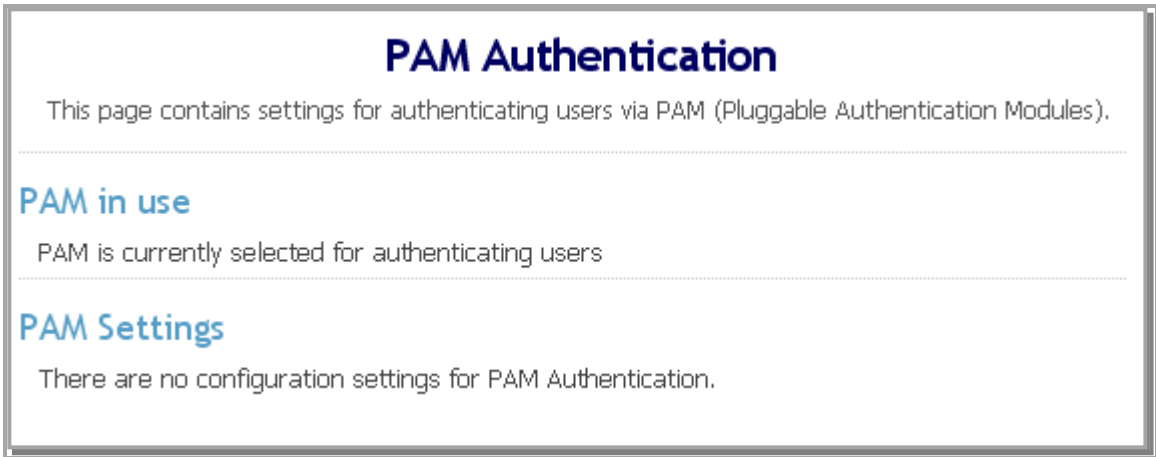


Figure 33: PAM Authentication page

PAM stands for Pluggable Authentication Modules and is the standard method for authenticating users on a Unix system. Selecting PAM for authenticating OpenVPN Access Server users means that users must provide the same username and password credentials to the Access Server as they would when authenticating to the Unix host that runs the Access Server.

When PAM is not already selected to be used to authenticate users, the Use PAM button selects PAM (instead of RADIUS or LDAP) for authentication. When PAM is selected, there are no configuration settings to adjust in the Admin Web UI.

Internally, the Access Server authenticates using a PAM service named `openvpn_as`, which corresponds to the file `/etc/pam.d/openvpn_as` (added during initial configuration of Access Server).

4.3.3 RADIUS

RADIUS Authentication

This page contains settings for authenticating users via RADIUS.

RADIUS is NOT in use

PAM is currently selected for authenticating users

RADIUS Settings

Hostname or IP Address	Shared Secret	Authentication Port	Accounting Port
<input type="text" value="rad1.example.net"/>	<input type="text" value="••••••"/>	<input type="text" value="1812"/>	<input type="text" value="1813"/>
<input type="text" value="rad2.example.net"/>	<input type="text" value="••••••"/>	<input type="text" value="1812"/>	<input type="text" value="1813"/>
<input type="text"/>	<input type="text"/>	<input type="text" value="1812"/>	<input type="text" value="1813"/>
<input type="text"/>	<input type="text"/>	<input type="text" value="1812"/>	<input type="text" value="1813"/>
<input type="text"/>	<input type="text"/>	<input type="text" value="1812"/>	<input type="text" value="1813"/>

Enable RADIUS Accounting

Figure 34: RADIUS Authentication page

When RADIUS is not already selected to be used to authenticate users, the Use RADIUS button selects RADIUS (instead of PAM or LDAP) for authentication.

Up to five redundant RADIUS servers may be configured. For each server, the Hostname or IP Address, Shared Secret, and Authentication Port must be specified. The Accounting Port is only needed when RADIUS Accounting is enabled (see below).

To authenticate users via RADIUS, the Access Server attempts to communicate with one of the configured RADIUS servers (chosen randomly). If the communication times out (after 5 seconds), the Access Server will retry the same server once more. The Access Server will attempt communication with up to three RADIUS servers in this way.

Accounting information is also conveyed to the RADIUS server when the Enable RADIUS Accounting checkbox is enabled. The user's accounting information includes the time length of the user's VPN session, as well as the input and output bytecounts for the user's VPN traffic. More specifically, the supported Accounting Attributes (as prescribed by RFC2865 and RFC2866) are listed in Section 6.2.

4.3.4 LDAP

LDAP Authentication

This page contains settings for authenticating users via LDAP. Please click the [Help](#) for more information on LDAP Authentication, or see the [Howto Page](#) for authenticating with Active Directory.

LDAP is NOT in use

PAM is currently selected for authenticating users

LDAP Settings

Primary server:

Secondary server:

Use SSL to connect to LDAP servers

Credentials for Initial Bind:

Bind anonymously

Use these credentials:

Bind DN:

Password:

Base DN for User Entries:

Username Attribute:

The **Username Attribute** is often **uid** for generic LDAP servers and **sAMAccountName** for Active Directory LDAP servers.

Additional LDAP Requirement: (Advanced)

This additional requirement uses LDAP query syntax. E.g., to require that the user be a member of a particular LDAP group (specified by DN) use this filter:

```
memberOf=CN=VPN Users, CN=Users, DC=example, DC=net
```

Figure 35: LDAP Authentication page

When LDAP is not already selected to be used to authenticate users, the Use LDAP button selects LDAP (instead of RADIUS or PAM) for authentication.

Note: In the context of LDAP, DN refers to a Distinguished Name, such as

```
cn="Jane Smith", cn="Users", dc="example", dc="com"
```

To authenticate users via LDAP, the Access Server performs these steps:

Bind to the LDAP server initially (either anonymously or with the specified Credentials for Initial Bind).

Perform an LDAP query to find the user's entry, using the Base DN for User Entries. A user's entry is the one whose Username Attribute value matches the username entered by the user at the login page.

Obtain the user's DN from the user entry, if found.

Re-bind to the LDAP server with the user's DN and the password entered by the user at the login page.

A Primary Server must be specified, either as a hostname or IP Address. Specifying a Secondary Server is optional; if present, Access Server attempts to communicate with the Secondary Server when attempts to contact the Primary Server fail. When Use SSL to connect to LDAP servers is enabled, Access Server establishes a secure, SSL-protected connection to the LDAP server(s) for all LDAP operations.

The optional **Additional LDAP Requirement** setting specifies a restriction (specified in LDAP query form) on a user's LDAP entry that must be true for the authentication to succeed. This can be used, for instance, to require membership in a particular LDAP group (specified by its group DN) for all users permitted to authenticate to the Access Server.

For more information on configuring LDAP authentication for interoperability with Active Directory, see Section 7.

4.4 Tools Pages

4.4.1 Profiles

A Configuration Profile contains all settings used by Access Server, with exception of the User Permissions database and the keys and certificates used by the SSL server components. Using multiple profiles may be considered a feature for “advanced users” of OpenVPN Access Server.

4.4.1.1 Active Profile and Edit Profile

Configuration Profiles

Configuration profiles are used to select between different sets of settings for the Access Server. You can edit one profile while a different one is being used for the running server (so current users aren't affected by editing changes in progress).

Active Profile

The settings of the **active profile** are used by the running Access Server.

Select Profile to Activate:

Default

Select Profile to Edit

The **edit profile** is the profile whose settings you edit in this Web UI.

Edit Profile:

Default

Figure 36: Active Profile and Edit Profile selection panels on Configuration Profiles page

The **Active Profile** is the profile that is currently selected for use with the VPN Server. Select the desired Active Profile using the **Select Profile to Activate** drop-down list.

The **Edit Profile** is the profile whose settings are currently being viewed and modified in the Admin Web UI. Select the desired Edit Profile using the **Edit Profile** drop-down list. By selecting an Edit Profile that is different from the Active Profile, you can edit the Access Server configuration without altering the behavior of the VPN Server (and thus, any current VPN client users will not be affected by editing changes). If the Edit Profile is the same as the Active Profile, then changes saved in the Admin UI can affect current VPN client users, once the running server is updated.

Note: If the VPN Server is running, and the Edit Profile is the same as the Active Profile, then changes made in the Admin Web UI are first saved to the profile and then, optionally, propagated to the running server. I.e., until you press the **Update Running Server** button, the settings in use by the running VPN Server may differ from those stored in the Active Profile.

4.4.1.2 Creating and Deleting a Profile

The screenshot displays two panels on the Configuration Profiles page. The top panel, titled "Create a New Configuration Profile", includes the instruction "Make a copy of an existing profile to a new profile." It features a "Select Profile to Copy:" dropdown menu with "Default" selected, a "Name for new profile:" text input field, a checkbox labeled "Allow overwrite of existing profile" which is currently unchecked, and a "Create" button. The bottom panel, titled "Delete a Configuration Profile", includes a "Note" stating that the current active or edit profile cannot be deleted. It features a "Select Profile to Delete:" dropdown menu with "Default" selected and a "Delete" button.

Figure 37: Profile Creation and Deletion panels on Configuration Profiles page

To create a new profile, make a copy of an existing profile and give it a new profile name. Select the source profile from the **Select Profile to Copy** drop-down list. Specify the name for the profile to be created using the **Name for new profile** box. The **Allow overwrite of existing profile** lets you delete an existing profile if its name is the same as the new profile name.

You can delete an unwanted profile by choosing its name in the **Select Profile to Delete** drop-down list and pressing **Delete**.

4.4.2 Connectivity Test

The Connectivity Test helps determine if VPN clients on the Internet will be able to connect to the VPN Server, given its current network configuration settings.

During the test, the Access Server communicates with a test host on the Internet. The test host reports the public IP address of the connection with the Access Server as well as the hostname obtained through a reverse DNS lookup on that public IP address.

The test server then attempts to establish a test connection to the Access Server, to simulate the connectivity that Internet VPN clients will encounter.

Connectivity Test

This test checks whether VPN clients on the Internet will be able to reach your VPN Server.

Current Server Network Settings

The VPN Server is currently configured with these network settings:

Host Name	IP Address	Port Number	Protocol
vpn-gw.example.net	all	443	tcp

The Connectivity Test attempts to determine the public IP address and FQDN (Fully-Qualified Domain Name) corresponding to the IP address: **all**, and also whether or not clients on the Internet will be able to connect to the VPN Server

Perform Test

The connectivity test may take several seconds to complete.

Figure 38: Connectivity Test page

Note: When the connectivity test runs, the Access Server dynamically adjusts the iptables rules so that the test traffic can be sent and received. These iptables rule changes are temporary and are removed when the connectivity test completes.

When the administrator presses the “Test Connectivity” button, several seconds may elapse before the test results are seen (see example results in Figure). If the test is successful but the detected public IP address or FQDN does not match the “Hostname or IP address” configured on the **Server Network Settings** page, a warning will be displayed to this effect.

Test Results

Result:	Success
Explanation:	Connectivity between an Internet client and your Access Server was confirmed.
Public IP Address:	174.33.177.207
Reverse DNS result:	174.33.177.207-static.reverse.softlayer.com

Warning: Your configured **Hostname or IP Address** is '**vpn-gw.example.net**' which is not the detected public IP address or FQDN.

To ensure that VPN clients on the Internet can connect to your VPN Server, it is suggested that you set the **Hostname or IP Address** setting on the **Server Network Settings** page to either '**174.33.177.207-static.reverse.softlayer.com**' or '**174.33.177.207**'.

Perform Test Again

Note: For maximally useful test results, the connectivity test should be performed when the VPN Server is not running. This allows the connectivity test to use the port number that the VPN Server is configured to use.

The connectivity test may take several seconds to complete.

Figure 39: Connectivity Test Results page

4.4.3 Support

The Support link takes the administrator to the online Support site for the Access Server software. This website is the main vehicle for communications with OpenVPN Technologies regarding the Access Server. Once you have logged in with your registered account, you can view and submit support tickets on this site or initiate a Live Support Session during OpenVPN Technologies business hours.

5 Connect Client

The Connect Client’s role is to create and distribute client configuration files and/or pre-configured OpenVPN Connect Client installers to authenticated users. This is the only way that VPN client installations are deployed with OpenVPN Access Server.

The client configuration and installer files generated by the Connect Client for a particular user are locked to that user. No other user can connect to the VPN with those files. Note that more than one connection profile may be installed on a client machine, for those situations where multiple users share the same machine.

The user accesses the Connect Client by entering the appropriate **https** URL into his or her Web browser. The URL to use is described in Section 4.2.2.3. Typically, the Connect Client URL is simply the server’s FQDN preceded by “**https://**”. When the browser connects, the user will likely see a warning or error displayed due to the untrusted server certificate (see Section 5.4 for information on preventing users from seeing such warnings). Once the user confirms that the server should be accessed, the user is presented with a simple login page, as shown in Figure below.




Figure 40: Connect Client login page

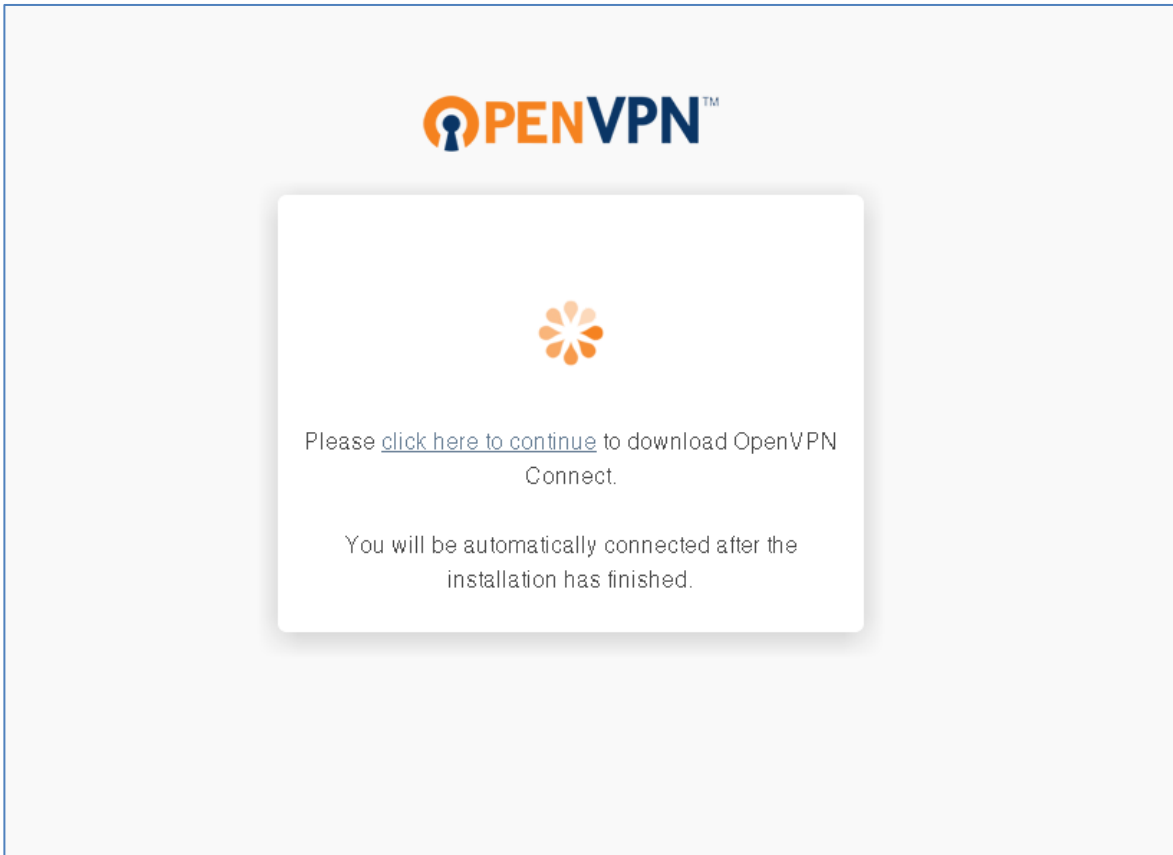
Users are authenticated against the authentication scheme configured by the Access Server administrator (see Section 4.3.1). When authenticating a user has two options which are seen in the drop-down menu next to the “Go” button. The user has a choice to connect or login.

5.1 Connect

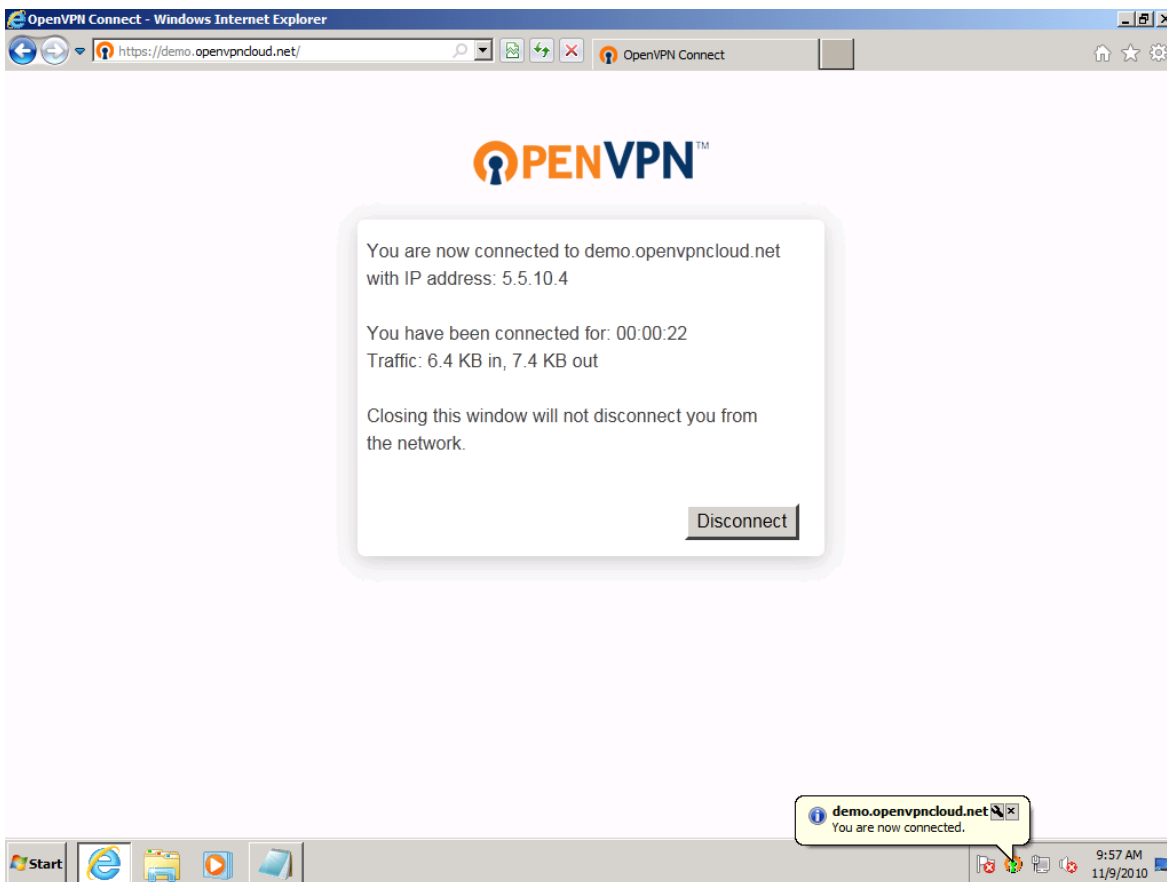
Connect: When the user chooses to connect for the first time they will be asked to download the OpenVPN Connect installer. The installer will include the user's bundled profile. Once installed, the connect client will automatically connect to the VPN Server.

Figure 41: Connect Page

This figure represents the page which users are brought to upon first login. They will be asked to download and run the installer:



Once the installer has completed the browser will continue to connect to the VPN Server. After the connection is successful they will be shown the status page that lists the server they are connected to as well as the amount of data that has passed to and from the VPN server via the client machine.



Tray Icon: The tray icon is a feature in the Connect Client that gives the user the ability to connect and disconnect from the Access Server directly through the tray. If the user is using an autologin profile they have the ability to do this without ever needing to communicate with the web browser. When the user is required to enter a username and password they have the ability to go to the Connect Client interface by selecting “Go to vpn.example.domain” which will then launch their default browser and bring them to the Connect Client interface. The user will also have the ability to disconnect from any active profile from the tray icon.



5.2 Login

Login: When accessing the Connect Client the user also has another choice aside from connect: Login. When logging into the Connect Client the user can download various different profiles assigned to them, different Windows client downloads, and tutorials for connecting to the VPN server from other Operating Systems.

Client downloads: The user will be given the ability to download multiple clients. There are currently two windows clients unique to the Access Server; the Connect client and the Desktop client. The connect client is browser based and works directly with the Connect Client to connect the user to the Access Server, this client is the easiest to use, the smallest in size and the recommended client for use with the Access Server. The Desktop client is a standalone client that is not integrated with the browser. You can import multiple profiles from different servers. You can also connect to the server by entering the Access Servers hostname or IP address in the “Server Address” field. For this option to work, you need to make sure you have the limited or complete API enabled from the client settings page in the Admin UI.

We also include guides for connecting to the Access server from both a Linux and Mac client which can be accessed by clicking the “OpenVPN for MAC OS X” or “OpenVPN for Linux” urls.

Profile Downloads: We also offer download links to Server-locked profile, user-locked profiles and autologin profiles for that user. Some of these profile may not be accessible to the user depending on what you allowed them to have permissions to via the User Permissions page and Client Settings page via the Admin UI.

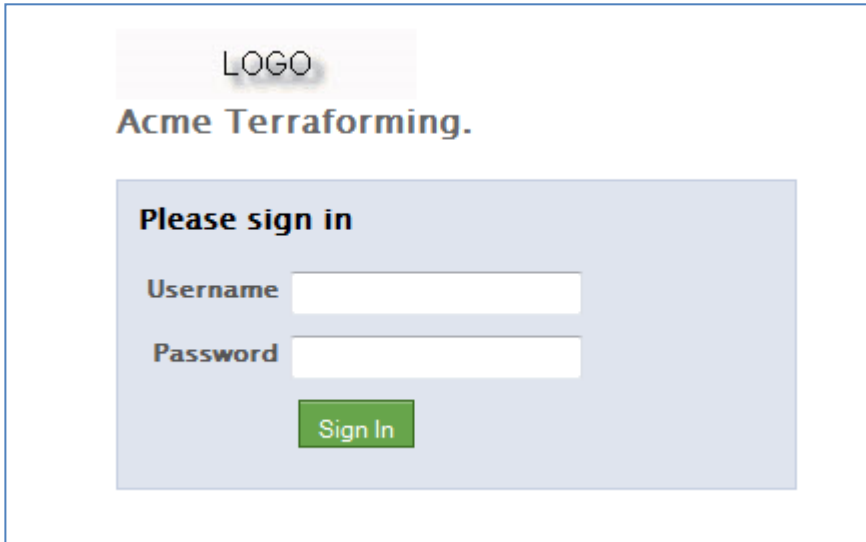
Connect: We also give the option for the user to connect to the Access Server through the login page, this can be done by clicking the “Connect” button on the top of the Login Page.

Admin: There is an Admin button at the top of the Login page that will take an Admin user to the Admin UI. This button will only show for users who are designated Admins for this Access Server.



5.3 Rebranding the Admin UI

OpenVPN Access Server now offers the option of rebranding the Admin UI with your Company Name and Company Logo.



In `as.conf`, set the following var to point to a `.jpg`, `.gif`, or `.png` image:

```
sa.logo_image_file=/my/dir/logo.jpg
```

To rebrand the company name, edit the following var in `as.conf`:

```
sa.company_name=Acme Terraforming
```

(the `as.conf` file can be found in: `/usr/local/openvpn_as/etc`)

```
system_users_local.0=root
system_users_local.1=openvpn_as

# The user/group that the web server will run as
cs.user=openvpn_as
cs.group=openvpn_as

# socket directory
general.sock_dir=~/.sock

# path to linux openvpn executable
# if undefined, find openvpn on the PATH
#general.openvpn_exe_path=

# source directory for OpenVPN Windows executable
# (Must have been built with MultiFileExtract)
sa.win_exe_dir=~/.exe

# The company name will be shown in the UI
sa.company_name=Acme Terraforming.
sa.logo_image_file=/usr/local/openvpn_as/logo.jpg

# server agent socket
sa.sock=~/.sock/sagent
```

5.4 Certificates

During the Access Server configuration process (specifically, during `ovpn-init`), server certificates for the Connect Client and OpenVPN server are created using a newly-generated Certificate Authority (“CA”). The CA’s self-signed certificate (with a Common Name of “OpenVPN Access Server Self CA”) is not among the trusted CA certificates pre-loaded into Web browsers. Thus, when a user connects to the Client Web Server, the Web browser will display a security warning that the server certificate is untrusted.

To eliminate the browser security warning, you must either:

1. Make the Web browsers in question add the generated CA certificate (for the “OpenVPN Access Server Self CA”) to the set of trusted CA’s.
2. Obtain a new server certificate for the Web Client Server using an external CA that is trusted by Web browsers.

Adding a new CA certificate to a browser’s set of trusted CA’s can typically be done by placing the CA certificate on a Web server and having the browser user open the appropriate URL (such as `https://corp.example.net/openvpn_as_ca.crt`). The browser prompts the user with confirmation dialog boxes, to verify that the new CA certificate should be trusted.

Of course, adding the new CA certificate to the browsers of all relevant users may be infeasible for a given deployment. Therefore, it is generally recommended that the Client Web Server certificate be replaced with one from a trusted CA. The steps for accomplishing this depend upon the choice of CA. Typically one must purchase these server certificates and provide proof of identity, along with submitting the required public key material in the form of a CSR (Certificate Signing Request). Below is an example of how to obtain an external certificate using the `openssl` utility is shown below:

1. Make sure you know the desired hostname for your server. This name will be the public name used by VPN clients to connect to your Access Server, and it should also be specified as the "Hostname or IP Address:" on the "Server Network Settings" page in the Access Server Admin Web UI. The hostname will be encoded in your certificate from the CA, so it will not be changeable.
2. Make a copy of the files in `/usr/local/openvpn_as/etc/web-ssl/` into a backup directory, just in case.

```
mkdir /root/keyfiles_bak
cp /usr/local/openvpn_as/etc/web-ssl/* /root/keyfiles_bak
```

3. Generate the new keypair and CSR (Certificate Signing Request) using these commands on your Access Server host machine:

```
cd /usr/local/openvpn_as/etc/web-ssl
openssl genrsa -out new.key 1024
openssl req -new -key new.key -out new.csr
```

In the last step, you will be prompted for input. Your CA may have certain requirements on the fields you specify. Often it is desirable to have the Common Name on the CSR match the hostname of your server. An example run of the above commands is shown below. Note that several fields are left blank by just hitting Return at the input prompt.

```
# openssl genrsa -out new.key 1024
Generating RSA private key, 1024 bit long modulus
.....+++++
.....+
+++++
e is 65537 (0x10001)
# openssl req -new -key new.key -out new.csr
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:US
State or Province Name (full name) [Some-State]:California
Locality Name (eg, city) []:Anytown
Organization Name (eg, company) [Internet Widgits Pty Ltd]:Exampletronix,
Inc.
Organizational Unit Name (eg, section) []:
Common Name (eg, YOUR name) []:vpn.example.net
Email Address []:
Please enter the following 'extra' attributes
to be sent with your certificate request
A challenge password []:
An optional company name []:
```

4. Give the contents of the "new.csr" file to your CA (via a Web upload or email or whatever method is preferred).
5. The CA may perform additional verification of your identity and/or your rights to use the names you specified. You may also have to pay for the certification service. In the end, the CA will provide a certificate and probably also a bundle with one or more CA certificates. All of these certificates should be PEM-encoded text strings, including BEGIN/END lines:

```
-----BEGIN CERTIFICATE-----
-----END CERTIFICATE-----
```

6. Save the server certificate (issued by the CA) as the file `server.crt` in `/usr/local/openvpn_as/etc/web-ssl` (overwriting the existing file).
7. Copy the `new.key` file as `server.key` in `/usr/local/openvpn_as/etc/web-ssl`.
8. Save the CA certificate bundle as `ca.crt` in the `/usr/local/openvpn_as/etc/web-ssl` directory. The CA certificates should appear in order, with the first certificate being that of the CA that issued the server certificate, and the last certificate being that of the "trusted root CA". The certificates can be concatenated, with the BEGIN and END lines included (so that the BEGIN line of one certificate follows the END line of the previous one).
9. Restart the Access Server using this command:

```
/etc/init.d/openvpnas restart
```

The new key and certificate should now be in use.

Note that to avoid security warnings with Web browsers, the server certificate must have a Subject Name with a “Common Name” field equal to the FQDN or IP address that clients will use to access the server. That is the purpose of specifying the FQDN in the “openssl req” step above.

5.5 Server-locked Profile

The server-locked profile allows any VPN User the ability to connect with the profile. This was created for a one size fits all solution. This profile is now offered to all users by default.

6 Additional Information on RADIUS Support

As of OpenVPN Access Server version 1.1 the RADIUS support includes support for RFC2865 and RFC2866. Please note that extensions beyond the previous mentioned RFC’s, such as Microsoft extension MS-CHAP V2 are **not supported** at this time. *This should be kept in mind when configuring a RADIUS server to interoperate with OpenVPN Access Sever.*

6.1 RADIUS Authentication Attributes

As of OpenVPN Access Server version 1.1 the RADIUS support includes the following Authentication Attributes as prescribed by RFC2865 and RFC2866:

1. **User-Name**
2. **User-Password**
3. **NAS-Identifier**
4. **NAS-Port-Type**
5. **NAS-Port**
6. **NAS-IP-Address**
7. **Service-Type**
8. **Framed-Protocol**
9. **Framed-IP-Address**
10. **Framed-IP-Netmask**

6.2 RADIUS Accounting Attributes

As of OpenVPN Access Server version 1.1 the RADIUS support includes the following Accounting Attributes as prescribed by the RFC2865 and RFC2866:

1. **Acct-Status-Type**
2. **Acct-Session-Id**
3. **Acct-Session-Time**
4. **Acct-Terminate-Cause**
5. **Acct-Input-Octets**
6. **Acct-Output-Octets**

7 How to authenticate users with Active Directory

OpenVPN Access Server's LDAP authentication feature is *general* in that it interoperates with various LDAP servers. A popular specific case is configuring Access Server to authenticate users with a Windows Active Directory server. You will need to know a few details about your Active Directory configuration to perform this configuration with Access Server.

Note: "DN" means Distinguished Name, a name encoding with multiple attribute=value pairs, such as `CN=Joan Smith, CN=Users, OU=Finance Group, DC=example, DC=com`

What you need to know:

- **The "Base DN" for User Entries of all users to be authenticated by Access Server.**
For an AD domain of "example.net", a typical Base DN for User Entries would be:
`CN=Users, DC=example, DC=net`
- **The Full DN and password of a user with administrative privileges in Active Directory.**
This user's credentials are used by Access Server to bind to the Active Directory server so that it can perform a search for a given VPN user's entry in the LDAP database.

7.1.1 Configuring Access Server LDAP Authentication

On the **LDAP** page in the Access Server Admin Web UI,

- Enter the hostname or IP address of the Active Directory server (typically also the Domain Controller) for the domain in the **Primary Server** field. If there is a secondary/backup Active Directory server, enter its hostname or IP address in the **Secondary Server** field.
- Configure the **Base DN for User Entries** setting with the Base DN described above.

Note that all users to be authenticated by Access Server must have full DNs that end with the specified Base DN. For example, with a Base DN of
`CN=Users, DC=example, DC=net`

these user DNs are valid:

```
CN=David Jones, CN=Users, DC=example, DC=net
CN=Users, DC=example, DC=net
```

However, these user DNs are *not* valid:

```
CN=Fred Murtok, DC=example, DC=net
CN=Alice Barnes, CN=Users, OU=Eng Group, DC=example, DC=net
```

- For the **Credentials for Initial Bind:** setting, choose **Using these credentials:**

Then enter the **Full DN** and password of the administrative user (see above). Note that you cannot simply enter "Administrator". The Full DN must be used, such as
`CN=Administrator, CN=Users, DC=example, DC=net`

- Be sure that the **Username Attribute** setting is set to "**sAMAccountName**". This is the attribute name that Active Directory uses to store a user's username (e.g., "abarnes")

7.1.2 Specifying Additional Requirements for LDAP Authentication

Starting with Access Server v1.2.0, the LDAP authentication capability can impose additional requirements on a user's LDAP entry in Active Directory. Any user that does not meet the requirements will not be successfully authenticated by Access Server (and thus, cannot use the VPN or Client Web Server).

The **Additional LDAP Requirement** setting specifies one or more requirements in the form of LDAP query syntax. If you are not fluent in LDAP query syntax, the examples below may still be useful.

Examples:

- Requiring membership in an Active Directory group

If you want to require that all VPN users be members of a particular group with group Full DN of
`CN=VPN Users, CN=Users, DC=example, DC=net`

then use this text as the **Additional LDAP Requirement**:
`memberOf=CN=VPN Users, CN=Users, DC=example, DC=net`

- Requiring that user accounts **not** be **disabled** in Active Directory

A user account in Active Directory that is marked as "disabled" may still have valid authentication results, from Access Server's perspective. To require that disabled user accounts be rejected in the context of Access Server authentication, use this text as the **Additional LDAP Requirement**:
`!(userAccountControl:1.2.840.113556.1.4.803:=2)`

- Combining multiple requirements

Multiple requirements can be required by surrounding each requirement with parentheses and then appending them together, and then preceding the combined string with either an ampersand ("&") for Logical AND, or with a vertical bar ("|") for Logical OR. For example, you can require that a user **both** be a member of a particular group **AND** not have a disabled account using this text as the **Additional LDAP Requirement**:

```
&(!(userAccountControl:1.2.840.113556.1.4.803:=2))
(memberOf=CN=VPN Users, CN=Users, DC=example, DC=net)
```

(the above text should be pasted as one single line into the textbox for **Additional LDAP Requirement**)

For more information on forming LDAP queries, see this [Microsoft Article](#).

8 Failover

Failover Settings

Redundancy Model

The Access Server supports a LAN-based redundancy model using Active/Standby failover. Please see the [Help](#) page for more information.

Select Redundancy Model

- LAN model (UCARP-based failover)
- No redundancy

LAN model

In the LAN redundancy model, a virtual IP address on the LAN is shared between the primary and secondary failover nodes, which must reside on the same LAN. When LAN model redundancy is enabled, all Access Server services will be provided via this IP address.

Shared virtual IP address

Primary Node

Configure the Primary Node in the failover pair.

Hostname/IP

SSH Username

SSH Password (optional)

SSH Port

Secondary Node

Configure the Secondary Node in the failover pair. Remember to install a fresh Access Server package on the Secondary Node, and when running `ovpn-init`, make sure to explicitly designate the node as a Secondary Node.

Hostname/IP

SSH Username

SSH Password (optional)

SSH Port

OpenVPN Access Server has a built-in failover mechanism which utilizes UCARP. With this failover system you can have a Primary Node and a Secondary Node which share a virtual IP. If the Primary Node goes down the secondary node will take over. This is an active-standby model.

In order to link the two servers together you will need to enter in the correct root password and ssh port for both primary and secondary nodes. You will also need to have an extra IP free to use as the shared virtual IP. Once failover is enabled you can access the admin ui through the shared IP.

****NOTE: Rsync is required on both the primary and secondary node for failover to work properly.***