nessquik 2.5

User Guide

Tim Rupp

Last Modified
07/06/2007

Table of Contents

Introduction....4
History....4
Requirements....6
Software....6
nessquik web....6
Nessus 2.x....7
Nessus 3.x....7
Hardware....7
nessquik web....7
Nessus Server(s)....8
Installation Options....9
The “All eggs in one basket” Approach....9
The “Little over here, little over there” Approach....9
The “One client to rule them all” Approach....10
Installation....11
Path Default Locations....11
nessquik....12
SELinux....12
Database....12
Website installation....13
nessquik Client....15
Requirements....15
Create a Scanner User....15
Installation via nessquik....16
Installation via standalone tarball....16
HTTP Authentication Modifications....17
scan-me-now....18
Website installation....18
Assign Elevated Privileges....18
portscan-me-now....20
Website Installation....20
Assign Elevated Privileges....20
Post Install Configuration....22
nessquik....22
Apache ACLs....22
Limit IP access....22
Limit Access with a Username and Password....23
Limit Viewing of Potentially Sensitive Directories....23
The templates Directory....24

The templates/(js | metrics | styles) Directories....24
Nessus 2.x....25
The Interface....26
create....26
settings....28
scans....29
help....31
Administration....32
home....32
admin....32
whitelist....32
settings....33
metrics....35
scans....36
help....37
Removing....38
nessquik....38
nessquik client....38
scan-me-now....38
portscan-me-now....38
Appendix A – Upgrading From nessquik 2.0....39
Backup Your Database....39
Backup Your Config Files....40
Download nessquik....40
Upgrade nessquik....40
Install nessquik clients....42
Appendix B – nessquik 2.5 Complete Feature List....43
Appendix C – Troubleshooting nessquik....48
Is PHP in your path?....49
Is your client key correct?....49
Is your RPC path correct?....49
Have you bumped up your RAM sufficiently?....50
Using HTTP authentication? Does the client know that?....50

Introduction

nessquik is a web front end for the Nessus vulnerability scanner. It provides a fast interface for scheduling vulnerability scans that allow an administrator to schedule scans quickly and easily.

History

FNAL's Computer Security Team encourages the system administrators on site to regularly scan their applications, devices, appliances, etc with the Nessus vulnerability scanner to make sure that any new vulnerabilities are found before they become security breaches.

In the past, the Nessus client was used for this purpose. Due to policies that we have on site however, the job of maintaining the Nessus accounts, and who is allowed to scan what, became impossible. This project was started to address the shortcomings of the Nessus client in our particular environment. Namely that,

- Authentication and authorization maintenance were a nightmare
- Creating scans was too overwhelming for many users due to the myriad of options available

nessquik solves 50% to 80% of the most common needs of our system administrators. It leverages the Lab's use of our Kerberized CA and several separate databases on site to control authentication and authorization automatically. The interface, by design, hides many of the configuration options that Nessus provides. In general, our user community has no interest in using 90% of the features that the original Nessus client provides.

Before undertaking the task of writing a new client, we researched the available clients. Our key requirements were.

1. Must be open source. We **will** be modifying it
2. Must support certificate authentication
3. Must be quick and easy to use
4. Must be web based since we have no interest in distributing clients to everyone

We were aware of only a single open source web based Nessus front end at the time. We

evaluated it to see if it could meet our needs, but decided to forgo using it because of our concerns about speed, certificate authentication, and the community supporting it. No other web based front ends were known at the time. That being the case, nessquik was started to fill this gap and meet our needs.

Requirements

Requirements are broken up into several categories. nessquik can be run on several different software configurations, so requirements for each of those configurations is mentioned below. On our production systems we use RedHat Enterprise Linux 4. Several individuals from the community have reported that nessquik also works on Ubuntu. nessquik should work on any *nix based operating system. If you find that it doesn't, please contact me so I can work out the problem.

Software

nessquik web

The following is a list of required software that needs to be installed to run the web portion of nessquik. I've listed version numbers only because that is what I test with on our development machines and it is what is used on the production machines.

Software	Version
PHP 5.x	5.2.1
Apache 2.x	2.0.52-28.ent
MySQL 4.x	4.1.20-1.RHEL4.1

The nessquik web installation also requires that you substantially increase the amount of RAM that PHP is allowed to allocate. I can't say definitively what that amount is, but I would start at 48 meg. You should run the nessquik-client code with debugging enabled to check for any out of memory errors. Increase the `memory_limit` in your php.ini file to make these errors go away. Be sure you **restart your web server** if you change the memory limit.

Nessus 2.x

Originally, nessquik was developed using the 2.x branch of the Nessus vulnerability scanner as it's backend. Since then, nessquik has moved to 3.x as it's backend. I expect that 2.x will continue to work as long as the output generated by Nessus during the scan doesn't change.

The following is needed by Nessus when installing using the universal installer. I'm basing these requirements off of my standard RHEL 4 installation. They can be installed via yum.

1. sharutils
2. gtk2-devel

Nessus 3.x

Nesuss 3.x is the currently supported branch of Nessus with regard to nessquik. It is the platform that is used for the production version of Nessus on-site. Installation of Nessus 3.x is dead simple on RHEL 4. You will need to fetch the RPM from Tenable's site though.

It is also suggested that you register your scanner with them so that you can get up-to-date plugins. At this point in time there are two types of feeds available. These are discussed in detail at Tenable's website.

Hardware

Overall, the hardware requirements for the nessquik part of your installation can be pretty meager. You'll want to devote the majority of your hardware resources to your Nessus servers.

nessquik web

The required hardware for the nessquik web software is. A typical server these days will easily satisfy these requirements.

- Pentium 3 733 Mhz

- 512 Meg of RAM
- 80 gigabyte hard drive
- 10/100 megabit NIC

Nessus Server(s)

The development Nessus server which nessquik relies on has the following hardware specifications.

- 2, Dual Core AMD Opterons at 2.4 Ghz
- 8 gigabytes of RAM
- 2, 250 gigabyte hard drives
- Gigabit NIC

There is no requirement that you match the Nessus Server specs. For your own Nessus servers, I recommend you follow Tenable's requirements as specified in their *Nessus 3.0 Installation Guide*. As of the time of this writing, these specifications are.

- A minimum of 256 MB of RAM for a local “Class C” network
- 1 GB to 4 GB of RAM for larger scans of multiple networks
- A Pentium 3, 733 MHz CPU or higher is recommended

Installation Options

There are several ways that you can approach the installation of nessquik and it's components. If you're creative, there are likely many more. If you set up nessquik in an interesting way and would like to share what you've done, please contact me.

The "All eggs in one basket" Approach

If you intend to set up all the available services on one machine, you're likely going to be using this approach. These services include

- Nessus
- nessquik
- MySQL
- Apache
- scan-me-now
- portscan-me-now

The "Little over here, little over there" Approach

This approach is the recommended approach. With this approach, you separate your nessquik installation from your Nessus scanners. This provides the best approach, in my opinion, because (among other reasons) it lifts the majority of load that you'll see, from the front end.

I've observed that when the Nessus scanner is under load, the web interface can become crippled due to the high CPU/RAM/network usage that the scanner requires. I strongly encourage you to separate **at least** the Nessus scanners from the nessquik web server.

The “One client to rule them all” Approach

This approach involves using a single nessquik-client to schedule scans on any number of Nessus servers. This setup is an extension of the ones above. The big difference with this approach is that you leverage the client-server architecture available in Nessus to connect to remote Nessus servers instead of having to run a nessquik client on every Nessus scanner you use.

The problem with this approach is that nessquik provides no facilities to support it. You're on your own at this point in time if you attempt to try this. There are several architectural issues that must be addressed in nessquik before something like this will be feasible. It is something I plan on addressing in a future release.

Installation

Installation of nessquik can be tricky for those unfamiliar with the system. I've tried my best to wrap as much of the installation in automated scripts as possible. Having said that, there is still some manual work that needs to be done before and after installation. You should **always** be able to install nessquik. If you run into a problem when following the steps below, please contact me so that we can resolve it.

Path Default Locations

You'll need to determine several directory locations before you continue. On a base RHEL 4.x system, the default directories are listed below. For the installation steps I will be using the paths listed in the following table. Please change these paths if needed for your particular installation.

Application	Default Directory
Nessus 2.x	/usr/local/
Nessus 3.x	/opt/nessus/
Apache	/var/www/html/
nessquik	/var/www/html/nessquik/
nessquik-client	/opt/nessquik-client/
scan-me-now	/var/www/html/scan-me-now/
portscan-me-now	/var/www/html/portscan-me-now/

nessquik

SELinux

Disable it or at least let Apache run commands from scripts. If you see messages in syslog that resemble the following,

```
Apr 20 18:21:10 unixc42 kernel: audit(1177111269.385:804): avc:
denied  { getattr } for  pid=10678 comm="ps" name="10157" dev=proc
ino=665649154 scontext=root:system_r:httpd_sys_script_t:s0
tcontext=root:system_r:unconfined_t:s0-s0:c0.c255 tclass=dir
```

Then that means SELinux is blocking the nessquik scripts that need to run.

Database

nessquik requires a MySQL database. The database should be created before you run any scripts to install nessquik. Make sure that you modify your configuration file (in the next section) with the values that you decide on here.

In the following steps, I'm making the assumption that you're at the MySQL command line prompt.

1. Create the MySQL database

   ```
   mysql> CREATE DATABASE `nessquik`
   ```

2. Create the MySQL user that will be used by nessquik.

   ```
   mysql> GRANT USAGE ON * . * TO 'nessquik'@'%'
   > IDENTIFIED BY '***';

   mysql> GRANT SELECT , INSERT , UPDATE , DELETE ,
   > CREATE , DROP , INDEX , ALTER ON `nessquik2` . *
   > TO 'nessquik'@'%';
   ```

The user should have the following privileges on the nessquik database that you created above.

- SELECT
- INSERT
- UPDATE
- DELETE
- CREATE
- DROP
- INDEX
- ALTER

nessquik should have it's own database user. This user only needs to have permission to access the nessquik database.

Website installation

1. Extract the nessquik tarball

   ```
   tar zxf nessquik-2.5.tar.gz
   ```

2. Move the contents of the package to your web space

   ```
   mv nessquik /var/www/html/
   ```

3. Change into the nessquik directory

   ```
   cd /var/www/html/nessquik
   ```

4. Change the ownership of the templates_c directory to be writable by your webserver

   ```
   chown apache:apache templates_c
   ```

5. Change the contents of the config-inc.php file to reflect your setup
6. Create the SQL tables

   ```
   /path/to/php /var/www/html/nessquik/setup/setup.php
   ```

7. If there are no errors, then run the update-plugins.php script from the command line

   ```
   sudo /path/to/php /var/www/html/nessquik/scripts/update-plugins.php
   ```

8. If there are no errors, then run the nasl_name_updater.php file from the command line

```
sudo /path/to/php /var/www/html/nessquik/scripts/nasl_name_updater.php
```

9. Add the entries in the provided crontab file into your own crontab. These are necessary because they control the nightly updates to the database tables and the rescheduling of scans.

nessquik Client

nessquik requires at least one nessquik-client be installed to schedule scans. This idea does not deviate from previous nessquik releases; a client has always been used. Installation of the nessquik client is simple.

Requirements

The current requirements for the nessquik client code are.

- PHP 5
- cURL support compiled into PHP 5

Create a Scanner User

nessquik client requires special privileges(because of Nessus) to run the scans that you schedule. These privileges can be granted via sudo.

In the past, I recommended you give the web user the sudo privileges. In nessquik 2.5 however, an API exists that allows the nessquik-client code to be separated from the server. With this in mind, my current recommendation is to NOT grant the necessary sudo privileges to the web user. Instead, create a special account specifically for scanning.

1. Create an account for the scanner user on your system

   ```
   [root@localhost ~]# adduser scanner
   ```

2. Using visudo, grant the right to run nessus without a password, to the web user. Here's an example line from a sudoers file using Nessus 2.x

   ```
   scanner ALL=NOPASSWD: /usr/local/bin/nessus
   ```

   where 'scanner' is the special scanner user.

   An example line using Nessus 3.x is shown on the next page

   ```
   scanner ALL=NOPASSWD: /opt/nessus/bin/nessus
   ```

Installation via nessquik

1. Change to your nessquik installation directory

   ```
   cd /var/www/html/nessquik
   ```

2. Move the folder nessquik-client/ in the deps/ directory into /opt or some place similar. Do not place it in any web accessible area.

   ```
   mv deps/nessquik-client /opt
   ```

3. Copy the contents of the client crontab file to the crontab for the scanner user. Make sure to change the directory paths for these crontab entries to reflect your installation.
4. In nessquik's web interface, under the Admin Settings section, create a new scanner and assign 'All Groups' to it.
5. Copy the client key that is generated (you can view it under the 'change' link) into the config-inc.php file in the nessquik client's confs/ directory.
6. Your client should start running jobs as they are created.

Installation via standalone tarball

1. Download the nessquik-client tarball from my home area

   ```
   http://home.fnal.gov/~tarupp
   ```

2. Move the tarball to /opt or some place similar. Do not place it in any web accessible area.

   ```
   mv nessquik-client.tar.gz /opt
   ```

3. Extract the tarball

   ```
   tar zxf nessquik-client.tar.gz
   ```

4. Copy the contents of the client crontab file to the crontab for the scanner user. Make sure to change the directory paths for these crontab entries to reflect your installation.
5. In nessquik's web interface, under the Admin Settings section, create a new scanner and assign 'All Groups' to it.
6. Copy the client key that is generated (you can view it under the 'change' link) into the config-

inc.php file in the nessquik client's confs/ directory.

7. Your client should start running jobs as they are created.

HTTP Authentication Modifications

Since nessquik does not currently support an authentication system for use in the general release, I have advised everyone to use htaccess files and HTTP based authentication to restrict access to the application until a supported authentication system is available.

During testing, a major problem was discovered that involved this advice and the nessquik client's extensive use of the XML-RPC API. A supported workaround is discussed below. In the next release of nessquik, this problem will become moot because a supported authentication system will be made available.

If you use the htaccess restrictions that I have recommended, this means you'll likely have a username/password combination for the nessquik area. In the nessquik client, there are three configuration options that you need to change so that the client will be able to talk back to the server using the XML-RPC API. These configuration options are

- _HTTP_AUTH
- _HTTP_AUTH_USER
- _HTTP_AUTH_PASS

Start by changing the value of _HTTP_AUTH to true. After doing this, you must also change the values of the _HTTP_AUTH_USER and _HTTP_AUTH_PASS configuration values. These values should be a username and password that can access the restricted nessquik area of your webserver. Note that even if you are only running the client on the localhost, you'll still need to make this change.

scan-me-now

Installation of scan-me-now is quick and generally easy. I've only ever run scan-me-now on the same server that runs the Nessus scanner, but it's certainly possible to talk to remote Nessus servers. Change the login credentials and server name in the config-inc.php file to use remote servers.

Website installation

1. Extract the scan-me-now tarball

```
tar zxf nessquik-2.5.tar.gz
```

2. Move the contents of the package to your web space

```
mv scan-me-now /var/www/html/
```

3. Change into the scan-me-now directory

```
cd /var/www/html/scan-me-now
```

4. Change the ownership of the template_c directory to be writable by your webserver

```
chown apache:apache templates_c
```

5. Change the ownership of the logs directory to be writable by your webserver

```
chown apache:apache logs
```

6. Change the ownership of the scans directory to be writable by your webserver

```
chown apache:apache scans
```

7. Change the contents of the config-inc.php file in the confs directory to reflect your setup

Assign Elevated Privileges

scan-me-now requires special privileges(because of Nessus) to run the scans that will be started when you surf to the scan-me-now webpage. These privileges can be granted via sudo.

Unlike nessquik, scan-me-now is not tied to the system which uses the nessquik API. As a

result, you must grant the web user access to run the Nessus command with elevated privileges. Creating a separate user who can run the scanner will not give you the desired results.

1. Using visudo, grant the right to run nessus without a password, to the web user. Here's an example line from a sudoers file using Nessus 2.x

   ```
   apache ALL=NOPASSWD: /usr/local/bin/nessus
   ```

   where 'apache' is the web user.

   And here's an example line using Nessus 3.x

   ```
   apache ALL=NOPASSWD: /opt/nessus/bin/nessus
   ```

portscan-me-now

Installation of portscan-me-now is also generally easy. Unlike scan-me-now, there's no need to fill out credentials in a config file. Just make sure that your path to nmap is correct in the config file. You'll also need to add elevated privileges on the webserver that is hosting portscan-me-now so that the web user can run nmap.

Website Installation

1. Extract the portscan-me-now tarball

```
tar zxf portscan-me-now.tar.gz
```

2. Move the contents of the package to your web space

```
mv portscan-me-now /var/www/html/
```

3. Change into the portscan-me-now directory

```
cd /var/www/html/portscan-me-now
```

4. Change the ownership of the templates_c directory to be writable by your webserver

```
chown apache:apache templates_c
```

5. Change the ownership of the logs directory to be writable by your webserver

```
chown apache:apache logs
```

6. Change the contents of the config-inc.php file in the confs directory to reflect your setup

Assign Elevated Privileges

portscan-me-now requires special privileges(because of nmap) to run the port scans that will be started when you surf to the portscan-me-now webpage. These privileges can be granted via sudo.

1. Using visudo, grant the right to run nmap without a password, to the web user. Here's an example line from a sudoers file

```
apache ALL=NOPASSWD: /usr/bin/nmap
```

where 'apache' is the web user.

Post Install Configuration

nessquik

1. Change the ownership of all the directories to root:root except for the `templates_c` directory.
2. ACL off the entire nessquik directory with a `.htaccess` to only allow connections from your network. Don't allow nessquik to face the internet directly without these htaccess controls.

Apache ACLs

Since nessquik, at this time, does not support an authentication system, I recommend that you use Apache .htaccess and .htuser files to restrict access to the application. Note that for any of these recommendations to take effect, you must change your Apache httpd.conf file to allow overriding of several directives. These directives are,

- Limit
- AuthConfig

Limit IP access

At the bare minimum, I'd suggest using a .htaccess to limit the IP addresses that are allowed to view the nessquik installation. An example htaccess file to accomplish this is shown below.

```
<limit GET>
     order deny,allow
     deny from all
     allow from 172.16
</limit>
```

In the above example, the class B IP block 172.16.xx.xx is allowed to reach nessquik. Everyone else is denied.

Limit Access with a Username and Password

In addition to the address restrictions, if you're feeling adventurous you can add HTTP authentication to the mix. This means that upon reaching the nessquik installation, you'll be prompted by a username and a password before you can proceed. This will further restrict who can access the nessquik install. The example below shows a .htaccess file with this configuration.

```
AuthType Basic
AuthName "Restricted Files"
AuthUserFile /etc/http/conf/password
Require user myuser
```

In the password file specified above, you would add username/encrypted password combinations using the htpasswd utility. More information can be found at the link below.

`http://httpd.apache.org/docs/2.0/howto/auth.html`

Note that if you decide to go this route, you will need to modify the nessquik client's config-inc.php file to accommodate this. The configuration options that need to be changed are

- _HTTP_AUTH
- _HTTP_AUTH_USER
- _HTTP_AUTH_PASS

Set these as per your installation or else the nessquik client will not be able to run scans that have been scheduled for it.

Limit Viewing of Potentially Sensitive Directories

In addition to the two suggestions above, I also suggest that you further limit the viewing of several directories to the local Apache server only. The `templates` and `templates_c` directory for example, do not (with a minor exception), need to be reachable from the web. The exceptions to this rule are the `templates/js/` and `templates/styles/` directories as the generated web pages will refer to these files via URL.

The templates Directory

When generating the web content, only apache needs to be able to read from the `templates` directory. A couple subdirectories need to be viewable by everyone, but those will be addressed in another section. To limit viewing of the templates directory only to the local host, you can create a .htaccess file in the `templates` directory and include the following in it.

```
<Limit GET POST>
     order allow,deny
     deny from all
     allow from 127.0.0.1
</Limit>
```

If you do this, **be sure** to add the necessary htaccess files in the templates subdirectories described in the next section.

The templates/(js | metrics | styles) Directories

To allow your browser to view the content in these directories, it is necessary that you role back the restrictive permissions that you applied earlier. The more lax permissions will only apply to these subdirectories. This can be accomplished by placing an htaccess file in the above directories that contains the following rules.

```
<Limit GET>
     order deny,allow
     allow from all
</Limit>
```

The `templates_c` Directory

Since this is a caching directory used by nessquik, there is no need for outsiders to be able to view this directory. Therefore all access to it can be restricted in much the same way as the `templates` directory.

```
<Limit GET POST>
     order allow,deny
```

```
        deny from all
        allow from 127.0.0.1
    </Limit>
```

All these restrictions are simply suggestions that I am offering so as to help you secure your nessquik installation further. None of them are required, however I urge you consider them. nessquik is still able to operate properly with them in place, and by implementing them, you are being proactive in securing nessquik from outsiders.

Nessus 2.x

If you're running Nessus 2.x, you'll want to update your plugins on a regular basis. As far as I know, Nessus 2 requires that the server be restarted to read in new plugins and any plugins that have changed. We've observed that this restarting of the server will stop nessquik scans that are taking place when the server restarts. To prevent this, I had to modify the script that came with Nessus to prevent Nessus from restarting if there are scans running.

Adding the following lines to the `nessus-update-plugins` script right before the big warning line about "DO NOT EDIT..."

```
count=`ps -ef | grep "nessusd: testing" | wc -l`
count=$((count - 1))
if [ $count -gt 0 ]; then
      exit
fi
```

This will cause the update script to exit if it sees any more than 0 scans running. Obviously this will mean that your plugins could be outdated if your server routinely has scans running during the times when it would update it's plugins. Keep an eye on the times that scans usually run, and try to tweak the update script run time to work around the busiest hours of the day.

The Interface

As powerful as Nessus is, I think it exposes too much configuration for an un-learned user to handle. At the lab we encourage individuals to scan their machines, but we've found that the myriad of options available discourages scanning. nessquik's interface is designed so that any user can schedule a scan targeting any machine under their control in as short a time as possible.

In nessquik there are four main sections that users need to care about.

1. create
2. settings
3. scans
4. help

create

When you first reach the nessquik app, you're immediately presented with the create page. On this page there are four sections to the left.

1. Scan Choices
2. Plugins
3. Configure
4. Finish

They're put there in the order that you'd use them. Scheduling a scan is simple. Under the first section, you may have one of several options. If you're at Fermi Lab, your screen has several options such as *registered machines* or *a cluster of computers*. Otherwise, your screen probably has less options for the Scan Choices. The general release of nessquik is a fine example of those “less options” at work. In either interface, by just clicking in the text box, you can begin adding targets to the scan profile.

But Tim! I have all my machines stored in a separate file! What will I do?

Usually this isn't as big of a problem as you may think it is. nessquik separates targets by either comma or new line. Therefore if you have all your targets in a comma separated file, you just need to copy and

paste into the box. Done. If CSV's aren't your forte though, you can also separate targets by new line. So enter a target, press the Enter key, and type in another target. Great!

Next in line is the selection of plugins. Nessus categorizes its plugins into Families and Severities, so, logically, nessquik let's you choose from those categories. In addition to that, there is also the much appreciated *all plugins* item. As you've probably guessed by now, clicking this item will tell Nessus to use all of it's plugins in the scan. Moving on, there is the search box that lets you pick individual plugins. The search box is semi smart, and will search for plugins like this

- If you type only a number, nessquik assumes you meant a plugin id
- If you type *word.nasl*, nessquik assumes you want to search for the exact nasl script name
- If you type *word*.nasl*, nessquik assumes you want to search for all nasl scripts whose file name begins with *word*.
- If you type **word*.nasl*, nessquik assumes you want to search for all nasl scripts that have *word* in their file name.
- If you type *word*, nessquik will return results like so

  No more than 15 matches from plugins where *word* is found in the short description

  No more than 10 matches from plugins where *word* is found in the long description
- If you type more than one word, nessquik will return results like so

  No more than 15 matches from plugins where **all words** exit in the short description

  No more than 10 matches from plugins where **all words** exist in the long description

Plugins that are returned will be placed in the *available devices* area. You can add them to the scan profile by clicking on the “plus” graphic next to the plugin.

On down the line, the section that allows you to configure the scan is next. In previous versions of nessquik, it was not possible to change all aspects of the scan **before** it was scheduled; now, it is. You'll notice that nessquik has a dearth of configuration options. By design, nessquik hides most of the configuration from the end user and only exposes what we consider necessary for our environment. I won't go through all the different options here because there is help text available for all of them. If you click on the option name, the help text will appear.

One thing that I will mention is the way the calendar feature works here. In nessquik 2, there was a calendar that could be used for selecting when you want the scan to be run. It was also used for rescheduling your scan. In 2.5, the calendar was removed. In it's place are clickable fields. If you hover your mouse over the month, day, year, etc, you'll notice it turns to a pointing finger. By **left** clicking on the field, you will **increase** the value. By **shift left** clicking on the field, you will **decrease** the value. Remember that.

Finally, there is one last section of the *create* page that users must be concerned with. Luckily, it's a really really simple section; Finish. When you're done tweaking the scan profile, just click the link under this section to schedule the scan. Note that if you haven't created any scanners in the admin section yet, the link here will be gray.

settings

From here you can configure all the settings for your scans. In nessquik 2.5 there are two setting choices that you're presented with when you first reach the page.

1. General
2. Per-scan

General scan settings affect all the scans that you will create from now on into the future. These settings are also where you can change certain behaviors of the interface. Not to sound like a broken record, but you need it drilled into your head: General settings affect all **future** scans. Think of them like the defaults for your scan profiles.

Now on the other side of the fence are *per-scan* settings. Here's where you'll find the current values for all your scan profiles. When you click on the per-scan link, you'll be presented with all of *your* profiles that nessquik knows about. There are several handy buttons here you'll notice. First, depending on the current status of the scan, you have the ability to either stop the scan, or delete the scan and it's results. Profiles cannot be changed while the scan is running. Clicking the “stop” button however will stop the scan so that you can make changes. The delete button will do just that; delete the profile. A word of caution here, deleting the scan profile will delete all the associated results too!

If you click on a scan profile, you'll be presented with a page that looks eerily similar to the *scan settings* area of the *create* page. Eerie is a good thing though because you'll tend to remember how the bells and whistles work. This page is basically a duplicate of the aforementioned *scan settings* page. For further explaining, see the section above concerning creating new scans and changing their settings.

scans

The scans page displays all your scans; past, present and future. There's a handy summarization of all your scans that you'll see to the left. The overview, as it's called, will update its self automatically as more scans come in or change their status. Since you're likely to want to *see* which scans are which, you can click on any of the overview links to display all the scans with that particular status.

There are several features added to the page when you click a particular type of status. Most of them have just been enhanced from 2.0. You have the ability to schedule and delete scans if they are *not ready to run*, or *finished*. Your most recent scan results are also available for any profile that has scan results saved. Notice that a said **most recent** scan results. nessquik by default will save all your scan results. If you want to view past results though, you need to choose the *view scan history* link on the left.

Results can be viewed in one of three formats

1. HTML
2. Text
3. NBE

Hovering your mouse over any of the *view*, *scan* or *email* links will allow you to get the results in those formats. Unlike 2.0, there is no longer a need to change the scan settings to change the viewable format. Now it's a simple and hover and click.

So what the heck is NBE you ask? Well, if you have to know, then you probably shouldn't care. NBE is one of the many formats that Nessus can generate results in. Generally it's used by developers and geeky people to create fancy schmancy programs, mine included, that create pretty reports out of the vulnerabilities found. nessquik provides the ability to fetch the raw NBE just as a courtesy to third

party users or developers who may want to integrate nessquik into their network. It's one of those formats that typical users don't need to worry about though.

Real quick I'll touch on the email link that nessquik provides. If you've taken a look at the scan settings, you may have noticed that you can send the scan report to multiple people. Well, you may be thinking now that this email link will send the email to all of those people again; you'd be wrong though. nessquik will only send email to multiple users when the scan initially finishes. That email link that you see there is to resend the email to **you**. After you've received it, you can forward it on to others.

As far as the scanning overview features are concerned, that pretty much covers it. Everything on the page should refresh without user intervention. One last thing that I guess I should mention is that you can also cancel scans when they're running from this page. When a scan is running, you'll see a progress bar too, as shown in the picture on the previous page.

Below the scanning overview is a section devoted to past scans. It's name is, well...*past scans*. In 2.0, I conveniently forgot to give you the ability to view past scans. Everyone thought this meant I was deleting the scan results. On the contrary. Your scans were never deleted, they were just hidden from view. With nessquik 2.5 you can see all your previous scan results. Also, as a semi apology for missing that obvious bug, I've gone out of my way to add the *compare results* link that you might have noticed.

Compare results should be considered alpha quality code. It's brand new and has only been tested in a couple ways. It's not destructive though, so I'm soliciting feedback from the community on what needs to be improved with it. If you find a problem with it, you really need to tell me. I may ask you for a copy of the scan results that are causing the problem, so you should just assume you'll need to send me the results. With the compare results, you're limited to comparing results from the same scan profile. This seems to me to be logical though because I'd assume that separate profiles have completely separate devices or plugins in them.

It's brain dead simple how the compare feature works, but first you need to have two (2) or more scan results to compare for a profile. If your profile doesn't have two or more sets of results, you simply wont see anything on the compare page. After you've accumulated some results, you can surf to the compare page where you'll see your scan profiles in a list. Pick one, and you'll be asked to pick a set of

results to **compare from** against a set of results to **compare to**. Select the bullets and click the button, that's about it. The scan results comparison is output as HTML. There are several new fields in the report, they're listed below.

- New hosts found
- Hosts no longer found
- Hosts with changed issues
- Difference in scan time
- List of new hosts
- List of new security issues and fixes for each host

At this point in time the comparison report is only viewable, you cannot download it. Also, the report is available in HTML only. Text output is likely to be included in a future release, and I haven't decided if NBE output is even logical since it would likely deviate from the format generated by Nessus. We'll see though, if the community wants it, the community is likely to get it.

help

Help is provided in nessquik via an FAQ. In the FAQ I tried to cover most of the gotchas and problems that people are likely to run into. The help is split up into categories. These categories are named according to where you're likely to run into the problem. Each category has topics in it that you can view by clicking on the category name. Topics are just like they were in nessquik 2 where you click on them to view the content of the topic. Since nessquik 2, several more topics have been added to address the most common problems I've heard from people. New topics can easily be added in the admin area to accommodate your own installation. If you get stuck with nessquik, I'd first suggest you check out the help pages since they very likely have an answer. If you still come up empty handed, email me.

Topics are more or less just plain text, but HTML can be embedded in them if necessary. Topics are sorted alphabetically, as are the categories, so keep that in mind when searching through them.

Administration

In addition to the normal sections that nessquik provides to any user, nessquik also provides an administration area. For the general release of 2.5 there is only a single account. These administration areas will become more relevant in 2.6 when the user accounts system is made available to the public. Not all the sections described below are available in the general release of nessquik 2.5. Mainly the problem stems from a lack of user accounts. If a feature does not exist, it will be noted as such in the section.

home

As a failsafe, nessquik includes this link on the admin page to link you back to the general user pages of nessquik. This link is not the default page that you surf after clicking on the admin link. If you're interested in returning to the default admin page, click the *admin* link.

admin

The admin link is the first page that you'll arrive at after clicking your username at the top left corner of the main nessquik page. There's really not much to explain on this page. At this point in time it basically says “Welcome” and tells you the version of nessquik that you're running.

whitelist

This feature is not available in nessquik General

The whitelist is a feature specific to Fermi for allowing administrators to scan machines that are not specifically run by them. It is mainly used to allow people with special privileges over a group of machines to scan them. Since there is a policy restricting scanning of machines that you are not the registered admin for, this feature is necessary.

The whitelist contains a list of usernames. Entries in the whitelist pair up a username to an IP

address, host, or range of IP addresses. New entries can be added to the whitelist at any time by clicking the *add entry* link to the left. Entries can also be copied if one user has privileges that you'd like another user to have, or renamed, if you enter a username wrong.

Copying entries has the effect of **appending** all entries from one user, to the list of the second user. If an entry already exists, a duplicate will not be made. Copying entries **does not** overwrite the entries of the user receiving the copy.

Renaming is included as a convenience for the administrator. In the event that a username is entered incorrectly, and entries have already been added, the username can be renamed so that it is correct again.

Individual whitelist entries for a user can be deleted from that user by clicking the *delete* button for the particular entry. Alternatively, all the entries for a user can be removed by clicking the *delete* button next to the username you want to remove.

settings

The admin settings area works slightly differently from how the user setting's area works. First of all, you cannot view the settings of any user scan from this page, so don't try to. There are two sections in the admin settings area, however only one of them is absolutely essential to using nessquik.

1. Scanners
2. Special Plugin Profiles

The “absolutely essential” one is the *scanners* section because that's where you tell nessquik which Nessus scanners exist and who's allowed to use them. At first glance you may say to yourself,

Hey, where do I type in my Nessus username and password?

No you're not seeing things, nessquik does not currently store these credentials in the database. During development, there was a general lack of agreement on how to accomplish this safely, so it was just deferred to a future release. nessquik only asks you for a server name, to make remembering easier, and

a list of groups that are allowed to use the scanner.

The server name does **not** have to be the hostname of the Nessus server. All of that configuration stuff is done in the nessquik client's config file. What you're entering here is just a “friendly name” if you will, to describe what the scanner is.

The groups that are allowed to use the scanner refers to functionality that exists in Fermi's release, but is currently irrelevant in the general release. For all practical purposes, if you use the general release, just get in the habit of choosing *all* groups for now.

Upon “creating” a scanner, there's one more important piece of information that you need and that is the client key for the nessquik client. To get the client key, click the *change* link under the Scanners section. You need to take that client key that you see and copy it into the configuration file for the nessquik client. For more information about how to configure the nessquik client, see the installation instructions earlier in this document.

The only other relevant section in the admin settings is the Special Plugin Profiles section. Some people have asked what this feature actually is. In a nutshell its the equivalent of a plugin composed of another set of plugins. On site it's used for such things as creating a “critical vulnerabilities” plugin which contains all the plugins that check for vulnerabilities that you're absolutely not allowed to have when connecting to our network. It makes the sys admin's life much easier because they can simply schedule a scan with a plugin profile, and if that profile ever has more plugins added to it, the sys admin's scan profile is automatically updated.

To create a new plugin profile you need to know three (3) things.

1. Name for the new profile
2. Plugins you want to include in it
3. Who's allowed to use it

For the general release of nessquik, the first and second bullets above should be filled out along the same lines as a new scanner; descriptive name and *all groups*. The plugins that you choose are the only

different piece. The search box is identical to the one you saw on the main scan creation page. Likewise, clicking the *by family* or *by severity* will accomplish the same thing it does on the scan creation page. Add plugins to the list on the right, click the *Create* button, and you're done.

metrics

This feature is not available in nessquik General

Metrics were added in 2.5 Fermi as I began to think about data mining the scanner results. There's only a single reason they're not included in the general release and that's because there is no user account system. Having said that, metrics will be made available to the general public in 2.6. There are two types of metrics that can be made.

1. Graphs
2. Reports

As of nessquik 2.5 Fermi, there are two (2) graphs and one (1) report available. The available graphs are

1. Number of Scans
2. Scan Trends

The available report is

1. System Statistics

To get an idea of what the interface looks like, check out the metrics page. The one main configurable option that all metrics should support is the *Time Range* option. It may look strikingly similar to the one you've seen on the settings page. Well, fancy a guess at where I copied it from. The controls to use the time range are the same as the settings page, but to get anything out of the metrics page, you have to click a metric.

For example if you click the *Number of Scans* metric you'll see there are two configurable options for this graph which are shown above it. Other graphs also can have configurable options as

they are contained in the actual metric file. After you choose to view, for instance, a graph, now is when the time range really comes into play. By default, the time range is one year ago from the time you view the page. It's used mainly as a factor to limit the amount of data that may be returned from the database. I don't suggest increasing it to cover a massive time range because the time required to generate the graph may be huge.

More metrics can be added to the system but it takes a bit of work to develop them. Refer to the *nessquik Metrics Developer Guide* if you're interested.

scans

The admin scans page is almost identical to the normal user scans page. In fact, in the general release you're likely not to notice much of any difference. From this page, the admin can get an overview of all the scans and their current state. The admin can also perform the same actions that normal users can perform on their scans from this page. These include canceling, rescheduling, and viewing the most recent scan results.

Unlike the normal user scans page however, the admin has a bit more control. First of all as I said earlier, they can cancel, delete, reschedule, etc, any user's scans. If an admin does this, the scan will be rescheduled in the users' name, not the admin's name. The admin can also narrow down the list of scans by username. This functionality is located in the *list by user* item on the left. In this view, only 25 usernames are displayed at any given time. You can narrow the list down by typing in the box. Unlike other lists in nessquik, this list is not automatically refreshed as the status of a scan changes or more scans are scheduled. This is a known limitation and it will be addressed in the future.

In addition to this new area that is available, you can also view the scan history of all the users by clicking on the link *view history and download* located on your left. In this area you're first presented with a list of users and their current scan counts. You have the same refining functionality as you did in the list by user. This time around however, you can click on their username to view all of their scan profiles. This page is important primarily in that it gives you the ability to download the

nessusrc file and machine list file that nessquik is going to use, or has used, to run the scan. It's very helpful in debugging and if you're asked for it by me, you're expected to provide it.

Clicking on a profile name that is blue will show the scan history for that profile. If a scan profile has no results (for instance if the user has told nessquik to not save the scan results) then the profile will be plain black text. Scan profiles can also be deleted from this page deleting a scan profile will, of course, remove all the associated scan results too. If you click on the scan profile to view the results, you also have the ability to remove individual scan results from the profile. As with the normal user history feature, you can view the results in any of the three (3) supported formats, save, and email them to the user.

help

Much like the help area for the normal user, nessquik provides a help area for the admin that includes topics that are specific to features in the admin interface. The concept behind categories and topics is identical to how they're used in the general user help area.

New to the admin help section is the ability to add new help topics and categories, and change the existing ones. I'm not going to try to play this up like it's a useful feature for the typical admin; it's not. But I use it for adding new help topics to the nessquik release so that's why it was included.

Help topics and categories can be edited after they are created. HTML can be embedded in the topic to provide semi-fancy text. Topics and categories can also be removed using the delete button next to each of them. If you delete a category, all the help topics in that category will also be deleted.

Removing

nessquik

The web GUI portion of nessquik is self contained in the folder that you extracted from the tarball. To remove the web portion of nessquik, simply remove the nessquik directory. nessquik's client code is typically found in /opt/nessquik-client if this directory exists, it can be removed. Finally the nessquik database is the last portion of nessquik that can be removed.

nessquik client

The nessquik client is self contained in the directory that is extracted from the tarball. To remove a nessquik client, simply delete the nessquik client directory.

scan-me-now

scan-me-now is self contained in the directory that is extracted when the tarball is uncompressed. To remove scan-me-now, simply delete the scan-me-now directory.

portscan-me-now

portscan-me-now is self contained in the directory that is extracted when the tarball is uncompressed. To remove portscan-me-now, simply delete the portscan-me-now directory.

Appendix A – Upgrading From nessquik 2.0

A lot has changed from nessquik 2.0, so this upgrade should be considered major. As such, you'll want to take every precaution to prevent a disaster from occurring during the upgrade. I've tried my best to cover all the bases with the upgrade. Scripts that are used, have been tested extensively with the datasets we generate in our production environment. If you run into any problems with the upgrade

Dont Panic

After you've stopped panicking, contact me via email and I'll do the best I can to get you back up and running.

In the steps below, I've made several assumptions about directory names on your system. The system being used for these instructions is Red Hat Enterprise 4. You should change the directory names as per your system's configuration. With this change in mind, all the steps should work as written.

Backup Your Database

This cannot be stressed enough. Before you start the upgrade, please backup your existing database. The upgrade scripts **will** modify your database and the information in it. You **will** need this backup in case something goes wrong.

1. To backup your database, you can use the following command.

```
mysqldump -u USER -p DATABASE > nessquik-backup.sql
```

Replace USER with (usually) the root username for the MySQL database.

Replace DATABASE with the name of the nessquik database as stored in MySQL.

example:

```
mysqldump -u root -p nessquik2 > nessquik-backup.sql
```

You will be prompted for the user's password after you enter the command. After the command has run,

you will have a file called nessquik-backup.sql in your current directory. Verify that this file does in fact have data in it. Run 'tail' on the file to make sure there are no errors at the end of it, and to make sure you do in fact have all the data.

Backup Your Config Files

Your configuration file is really the only file that you're ever likely to change in nessquik. You should backup this file to a separate directory (for example your home directory) during the upgrade, so that you can copy it to the new installation after upgrading. The file can be found, relative to your nessquik installation directory, in the `confs/` directory. It is named `config-inc.php`.

Download nessquik

nessquik can be downloaded from my homepage at FNAL.

`http://home.fnal.gov/~tarupp`

Downloads are listed under the nessquik section. For this installation step, you're only interested in the full nessquik package, not scan-me-now, portscan-me-now or the nessquik-client.

Upgrade nessquik

Upgrading nessquik can be accomplished by following the directions below. I'm making the assumption that you have unfettered access to the database. This includes root level privileges on the nessquik database. I'm also assuming that you have benevolent dictator powers over the nessquik installation directory and the systems where your Nessus scanner software is installed.

1. Disable your nessquik crontabs to prevent scans from being scheduled during the upgrade.

   ```
   10 00 * * * /../update-plugins.php && /../scripts/nasl_name_updater.php 2>&1
   */5 * * * * /../scan-runner.php >/dev/null 2>&1
   */5 * * * * /../cron.php >/dev/null 2>&1
   ```

2. Move your own installation to a backup location

```
mv /var/www/html/nessquik /var/www/html/nessquik.backup
```

3. Change to your web root and move the downloaded tarball there

```
cd /var/www/html
mv ~/nessquik-2.5.tar.gz .
```

4. Extract the new nessquik installation to your web area

```
tar zxf nessquik-2.5.tar.gz
```

5. Update the new configuration file. This can be accomplished by running the update-config-inc.php file in the nessquik 2.5 upgrade directory. Make sure you are using an account that has write permissions to the new nessquik confs directory.

```
cd upgrade/2.0_to_2.5/
/path/to/php update-config-inc.php oldconfig
```

The value for *oldconfig* should be the full path to the version 2.0 config file on your nessquik installation. The new configuration file will be written to the version 2.5 confs directory.

6. Update the database using the upgrade.php script

```
/path/to/php upgrade.php
```

7. Change back to the nessquik installation directory

```
cd ../../
```

8. Change the ownership of the templates_c directory to be writable by the web server user

```
chown apache:apache templates_c
```

9. Change the ownership of the logs directory to be writable by the web server

```
chown apache:apache logs
```

10. You're encouraged at this time to restrict access to the nessquik application using .htaccess files with Apache. This is to prevent unwanted visitors from access the tool. An accounts system will be available in the future.

Install nessquik clients

For instructions on how to do this, see the installation steps for nessquik-client earlier in this document. Be sure to read the section concerning HTTP based authentication and the effects it can have on the client. If you are restricting access to the nessquik install using HTTP based authentication, you'll need to make extra client tweaks, or else the client will continually receive a “permission denied” error, and no scans will be run.

Appendix B – nessquik 2.5 Complete Feature List

- Schedule scans from the web providing a single Nessus client for an arbitrary number of people to use.
- Specify an arbitrary number of machines in the scan profile
- Specify machines in any of the following ways

    hostname

    IP address

    CIDR block

    IP range

    Virtual Host (as-in Apache web server virtual hosts)
- Specify an arbitrary number of plugins for a scan profile

    by family

    by severity

    all plugins

    individual plugins
- Quickly search for a plugin

    by nasl name

    by plugin ID

    by words contained in the short description

    by words contained in the long description
- Name your scans so you can remember what they do
- Schedule your scans to run in the future
- Choose whether to ping the host first when scanning or not
- Ability to choose which Nessus scanner to run the scan from
- Specify a port range for the scan
- Specify an arbitrary number of cgi-bin directories for scanning web applications

- Reschedule your scans
    - weekly
    - daily
    - monthly
- For daily rescheduling, specify a specific time of day
- For weekly rescheduling, specify a specific time of day and the days that you want to reschedule the scan
- For monthly rescheduling, specify a time of day, and then choose from two different means of picking the day to schedule on
    - absolute - 1st, 2nd, 3rd etc day of the month
    - relative - 1st Monday, last Friday, etc.
- Save scan reports for viewing at any time
- Scan reports are emailed to you automatically
- Choose the format of the report that you want emailed to you
    - text
    - html
- Specify a custom subject line for the email that is sent to you
- Use macros in the email subject line to customize it further
    - %m is replaced by the machine being scanned
    - %D is replaced with the date the scan started
    - %d is replace with the date the scan finished
    - %T is replaced with the time the scan started
    - %t is replaced with the time the scan finished
- Send the report to an arbitrary number of other users
- Change settings on a global or per-scan basis
- Enable short or full plugin listings when searching for plugins to eliminate clutter
- Most of the scan settings can be configured globally for all scans that will be created in the

future

    ping the host first

    scanner to use

    port scan range

    alternate cgi-bin directories

    saving the scan report

    scan report format

    custom subject line

    sending the report to other people

- Add and remove machines from existing scan profiles
- Add and remove plugins from existing scan profiles
- View the status of all your scans from one dashboard

    scans not ready to run

    scans pending

    scans running

    scans finished

- No need to reload the page to update the scan status
- View the history for a scan profile
- Compare results from different runs of a scan profile
- View your scan results in different formats from the scans page

    html

    text

    nbe

- Save your scan results in different formats from the scans page

    html

    text

    nbe

- Email your scan results in different formats from the scans page
    - html
    - text
    - nbe
- Remove scan profiles completely from the scans page
- Stop running scans from the scans page
- Reschedule “not ready” and “finished” scans from the scans page
- Help FAQs are provided for many categories
- Ability to whitelist machines for users (Fermi version only)
- Ability to copy and rename whitelist entries (Fermi version only)
- Create special plugin profiles that contain other plugins
- Special plugin profiles can contain any type of plugin including families and severities
- Ability to add an arbitrary number of scanners to the system so that users can pick which scanner to run their scan on
- Ability to assign scanners to specific groups so that only members of those groups can view and schedule scans on those scanners (Fermi version only)
- Several metrics made available
    - Number of Scans
    - Scan Trends
    - System Statistics
- Metrics can be narrowed down via a time range
- Individual metrics allow for runtime customization
- Administrators can view the status of every scan
- Ability to stop, delete, or reschedule any scan for any user via the administration pages
- Ability to search for scans by a user and view the current status of it
- View the scan history of any user including the number of scans that they have
    - not ready to run

pending

running

finished

- View scan results for any user in any of the supported formats from the admin page
- Email the scan results for any user in any of the supported formats from the admin page
- Save the currently configured nessusrc file for any scan profile from the admin area
- Save the currently configured machine list for any scan profile from the admin area
- Administrative help topics that are relevant to options in the admin area
- Ability to add new help topics to either the end user help area or the admin help area
- Ability to edit and delete any help topic

Appendix C – Troubleshooting nessquik

If you are having trouble installing or using nessquik and need to contact me, there are several things that I need to know. Make sure you include the following in your email.

1. PHP version
2. MySQL version
3. Full path to your PHP binary

```
[root@localhost]# which php
```

4. Version of Nessus that you are using

If you are specifically having trouble with the nessquik-client, **make sure** that you attach a copy of the XML-RPC debug output with your mail. This can be done like so.

- Change to your nessquik-client installation directory

```
[root@localhost]# cd /opt/nessquik-client/
```

- Modify the `config-inc.php` in the nessquik-client `confs/` directory and change the following value to the one shown below.

```
// For troubleshooting, you may be asked to set...
define("_DEBUG", true);
```

- Reschedule the scan
- All XML-RPC debug output will be written to a file called

```
xml-rpc-debug-XXX
```

  where `XXX` is replaced with the timestamp of when the scan ran. **Attach this output with your help request**.

Is PHP in your path?

If you're able to schedule scans, but none of those scans ever run, you may have an incorrect path to PHP specified in several of the scripts that nessquik uses on the backend. First of all, log in as a normal user on your server and use the which command to locate PHP.

```
[user@localhost]$ which php
```

If that returns **any** error, your first problem is that PHP cannot be found. Find the location of the PHP binary. With that location in mind, you need to update the following scripts or else nessquik will not work. Each of the scripts below can be located in the `scripts/` folder of the respective applications.

1. nessquik webserver
   - `cron.php`
   - `nasl_name_updater.php`
   - `update-plugins.php`
2. nessquik clients
   - `scan-maker.php`
   - `scan-runner.php`

Is your client key correct?

If you're able to schedule scans but those scans aren't running, another problem may be that your client key is incorrect. Verify that your client key as shown in the Admin > Settings > Change Scanners matches the client key that you entered in the nessquik client's config-inc.php file.

Is your RPC path correct?

The `_NESSQUIK_PATH` variable defines where on the server your installation of nessquik is located. The nessquik client uses this path to query the XML-RPC server during scan scheduling. If this path is not correct, no scans will ever take place and you will see errors like the following in your

client's debug output.

```
HTTP/1.1 404 Not Found
```

associated with this file, or something similar

```
POST /nessquik/xmlrpc/jobs.php HTTP/1.1
```

Have you bumped up your RAM sufficiently?

If your nessquik client is running scans but sending you (literally) empty output, you may be running PHP with too little RAM.

Since my nessquik environment is very a-typical, I can't say for sure what the needed amount of RAM should be for PHP. Start with 48 meg and bump it up as necessary. I have heard users say that 60 meg is adequate. If you're not sure if you're encountering a “lack of RAM” problem, the nessquik client debug output will be able to tell you. You will likely see output in the XML file that looks like this.

```
Allowed memory size of 8388608 bytes exhausted (tried to
allocate 524277 bytes)
```

If that's the case, start bumping up the memory_limit specified in your php.ini file. On our production system we give PHP 256 meg of RAM. This may sound like a lot, but the system has 8 gig available. Make sure you restart the webserver if you boost the amount of memory in the php.ini file.

Using HTTP authentication? Does the client know that?

With nessquik general, there is no form of access control built into the software. Instead, other means of access control must be used. For instance, the installation document suggests using Apache's built in HTTP authentication. If this method is used, the nessquik client will need to know what credentials to use to reach the nessquik installation.

In the nessquik client there are three (3) configuration options that you can set.

1. _HTTP_AUTH

2. _HTTP_AUTH_USER
3. _HTTP_AUTH_PASS

If you are using HTTP authentication, you must set _HTTP_AUTH to true and then change the username and password in the next two settings. After making this change, the nessquik client will be able to talk to the nessquik installation.

ip_conntrack too small (1)

If you begin to see messages in your syslog that are similar to the one shown below

```
ip_conntrack: table full, dropping packet.
```

Then it's likely that your IPTables network connection table is too small. To check how many sessions you currently have open right now, use the following command

```
cat /proc/net/ip_conntrack | wc -l
```

To find out the max number of conntrack entries, use this command

```
cat /proc/sys/net/ipv4/ip_conntrack_max
```

If the first number is very close, or matching the second number, you'll need to bump up the maximum number of connections. On linux, this can be done like so.

```
echo "<some_bigger_number>" > /proc/sys/net/ipv4/ip_conntrack_max
```

And if you want the changes to survive a reboot, place the following in /etc/sysctl.conf

```
sys.net.ipv4.ip_conntrack_max = <some_big_number>
```

or

```
net.ipv4.ip_conntrack_max = <some_big_number>
```

Appendix D – References

1. Chemko, Daniel, et all. “ip_conntrack: table full, dropping packet. SAY WHA????”. Netfilter mailing list. 18 June 2007. http://lists.netfilter.org/pipermail/netfilter/2003-September/046641.html