

The Cacti Manual

Ian Berry

Tony Roman

Larry Adams

The Cacti Manual

by Ian Berry, Tony Roman, and Larry Adams

Published 2007

Copyright © 2007 The Cacti Group

This project is licensed under the terms of the GPL.

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

All product names are property of their respective owners. Such names are used for identification purposes only and are not indicative of endorsement by or of any company, organization, product, or platform.

Table of Contents

| | |
|---|-----------|
| I. Installation | 1 |
| 1. Requirements | 1 |
| 2. Installing Under Unix | 2 |
| 3. Installing Under Windows | 3 |
| 4. Upgrading Cacti | 8 |
| II. Basics | 9 |
| 5. Graph Overview | 9 |
| 6. How to Graph Your Network | 10 |
| 6.1. Creating a Device | 10 |
| 6.2. Creating the Graphs | 11 |
| 7. Viewing Graphs | 12 |
| 7.1. Graph Trees | 12 |
| 8. User Management | 13 |
| 8.1. Creating a New User | 13 |
| 8.2. Realm Permissions | 14 |
| 8.3. Graph Permissions | 14 |
| 8.4. Graph Settings | 14 |
| III. Advanced Topics | 15 |
| 9. Data Input Methods | 15 |
| 9.1. Creating a Data Input Method | 15 |
| 9.2. Making Your Scripts Work With Cacti | 16 |
| 10. Data Queries | 17 |
| 10.1. Creating a Data Query | 17 |
| 10.2. SNMP Query XML Syntax | 18 |
| 10.3. Script Query XML Syntax | 20 |
| 11. Templates | 23 |
| 11.1. Graph Templates | 23 |
| 11.2. Data Templates | 26 |
| 11.3. Host Templates | 28 |
| 12. PHP Script Server | 29 |
| 12.1. Using the Script Server | 29 |
| 12.2. Upgrade Steps for the Example HostMib Data Queries | 29 |
| 12.3. Migration of Existing PHP Scripts to Script Server | 29 |
| 12.4. Testing Your Script in the Script Server | 31 |
| IV. Using Cacti | 32 |
| 13. How To | 32 |
| 13.1. Graph a Single SNMP OID | 32 |
| 13.2. Simplest Method of Going from Script to Graph (Walkthrough) | 32 |
| 14. Frequently Asked Questions | 37 |
| V. Reference | 40 |
| 15. Variables | 40 |
| 15.1. Graph Variables | 40 |
| 15.2. Host Variables | 43 |
| 16. RRDTool Specific Features | 44 |
| 16.1. GPRINT Presets | 44 |
| 16.2. CDEFs | 44 |

List of Tables

6-1. SNMP Retrieval Types 10

8-1. Field Description: User Management 13

9-1. Field Description: Data Input Methods..... 15

9-2. Field Description: Data Input Fields..... 15

10-1. Field Description: Data Queries 17

10-2. Field Description: Associated Graph Templates..... 17

10-3. SNMP Query XML Field Reference 18

10-4. Script Query XML Field Reference 20

11-1. Field Description: Graph Templates 23

11-2. Field Description: Graph Template Items 24

11-3. Field Description: Graph Template Items 25

11-4. Field Description: Data Templates 26

11-5. Field Description: Data Source Items..... 27

13-1. Field Description: SNMP - Generic OID Template 32

15-1. Nth Percentile Argument Description..... 40

15-2. Nth Percentile Types 41

15-3. Bandwidth Summation Argument Description..... 42

15-4. Host Variables 43

16-1. CDEF Item Types 44

Chapter 1. Requirements

Cacti requires that the following software is installed on your system.

- `RRDTool` 1.0.49 or 1.2.x or greater
- `MySQL` 3.23 or greater, 4.1.x or 5.x or greater
- `PHP` 4.3.6 or greater, 5.x greater highly recommended for advanced features

Packages for RPM-based Operating Systems

- `httpd`
- `php`
- `php-mysql`
- `php-snmp`
- `mysql`
- `mysql-server`
- `net-snmp`

Ports for FreeBSD

- `www/apache2`
- `net/rrdtool`
- `net/net-snmp`
- `www/php4-cgi`
- `lang/php4` (With `MySQL` and `SNMP` Support)
- `databases/mysql323-server`

Chapter 2. Installing Under Unix

1. Extract the distribution tarball.

```
shell> tar xzvf cacti-version.tar.gz
```

2. Create the MySQL database:

```
shell> mysqladmin --user=root create cacti
```

3. Import the default cacti database:

```
shell> mysql cacti < cacti.sql
```

4. Optional: Create a MySQL username and password for Cacti.

```
shell> mysql --user=root mysql
mysql> GRANT ALL ON cacti.* TO cactiuser@localhost IDENTIFIED BY 'somepassword';
mysql> flush privileges;
```

5. Edit `include/config.php` and specify the MySQL user, password and database for your Cacti configuration.

```
$database_default = "cacti";
$database_hostname = "localhost";
$database_username = "cactiuser";
$database_password = "cacti";
```

6. Set the appropriate permissions on cacti's directories for graph/log generation. You should execute these commands from inside cacti's directory to change the permissions.

```
shell> chown -R cactiuser rra/ log/
```

(Enter a valid username for *cactiuser*, this user will also be used in the next step for data gathering.)

7. Add a line to your `/etc/crontab` file similar to:

```
*/5 * * * * cactiuser php /var/www/html/cacti/poller.php > /dev/null 2>&1
```

Replace *cactiuser* with the valid user specified in the previous step.

Replace `/var/www/html/cacti/` with your full Cacti path.

8. Point your web browser to:

`http://your-server/cacti/`

Log in the with a username/password of *admin*. You will be required to change this password immediately. Make sure to fill in all of the path variables carefully and correctly on the following screen.

Chapter 3. Installing Under Windows

Software Components Required

1. (Optional) Apache> - This software is optional if running Windows Internet Information Server.
2. Cacti> - Install from the zip distribution and install in the web root or your choice. Many choose to install into a "Cacti" sub folder.
3. Cactid - Install from the zip distribution into the `c:\cacti` directory. Make sure your `cactid.conf.dist` is located in that directory as well.
4. RRDTool - Install from the Cacti website. Install it into the `c:\cacti` directory.
5. PHP 4.3.6+ or 5.x - Install into the `c:\php` folder. If you choose to install into `c:\Program Files\php`, you will have to use 8.3 filenames to reference it's binaries in Cacti.
6. MySQL 4.x or MySQL 5.x - Install into the default location. This is typically `c:\Program Files\MySQL\MySQL Server X.XX`.
7. (Optional) Cygwin - Download and execute `setup.exe` from the Cygwin website. Keep the `setup.exe` file for later use.
8. (Optional) Net-SNMP - Install to the `c:\net-snmp` directory. If you choose to use `c:\Program Files\net-snmp` you will have to use 8.3 filenames to reference it's binaries in Cacti.

Configure PHP

1. If using PHP 4, move the files in `c:\php\dlls` to `c:\php`
2. Add the following directory to the existing Windows System **PATH** environment variable: `c:\php`. The Windows path can be accessed via the Control Panel at: System | Advanced | Environment Variables | System Variables.
3. Add the following directory to a new Windows System environment variable called **PHPRC**: `c:\php`.
4. Add a new Windows System environment variable called **MIBDIRS**. If using PHP 4, set it to `c:\php\mibs`. If using PHP 5, set it to `c:\php\extras\mibs`
5. If using PHP 4.3.5 or less, create the following directory `c:\tmp`.
6. Rename the file `c:\php\php.ini.dist` to `php.ini`, and make the following changes to it:

If using PHP 4 add/uncomment the following lines.

```
extension_dir = c:\php\extensions
extension=php_snmp.dll
extension=php_sockets.dll
cgi.force_redirect = 0
```

If using PHP 5 uncomment the following lines.

```
extension_dir = c:\php\ext
extension=php_mysql.dll
extension=php_snmp.dll
extension=php_sockets.dll
cgi.force_redirect = 0
```

If using PHP 4.3.5 or less include the following line. If using 4.3.6 or greater, you should remove this line if present.

```
session.save_path=c:\tmp
```

7. In earlier installation guides to PHP, they recommended moving certain DLL's to the `c:\winnt\system32` directory. If so, you will have to remove those files. Please review the PHP installation documentation for instructions on removing those files.
8. If you want to allow template importing, uncomment the following line:

```
file_uploads = On
```
9. Give the user who will be running the scheduled task, modify rights to the `.index` file in the location pointed to by the **MIBDIRS** Windows System environment variable.

Configure the Webserver (Apache)

1. Make sure you have stopped any IIS web servers before you proceed with Apache installation, or make sure Apache is configured on an alternate port.
2. If you are using *Apache 1.3.x*, installation of PHP 5 is not recommended. If using PHP 4, add the following lines to your `httpd.conf` file. You can edit that file by selecting Start | All Programs | Apache HTTP Server X.XX | Configure Server | Edit the `httpd.conf` file pick from the Taskbar.

If using Apache 1.3.x and PHP 4, then add the following lines:

```
LoadModule php4_module c:\php\sapi\php4apache.dll
AddModule mod_php4.c
AddType application/x-httpd-php .php
DirectoryIndex index.html index.htm index.php
```

If using Apache 2.x and PHP 4, then add the following lines:

```
LoadModule php4_module c:\php\sapi\php4apache2.dll
AddType application/x-httpd-php .php
DirectoryIndex index.html index.htm index.php
```

If using Apache 2.x and PHP 5, then add the following lines.

```
LoadModule php5_module c:\php\php5apache2.dll
AddType application/x-httpd-php .php
DirectoryIndex index.html index.htm index.php
```

Configure the Webserver (IIS)

1. Start the Internet Information Services (IIS) Manager, right click on the *Default Web Site* (in most cases) and select *Properties*.
2. Under the *Home Directory* tab, select *Configuration* and click *Add*. Browse to the path of `php4isapi.dll` or `php5isapi.dll`, and type in `.php` as the extension. Note: if using IIS6, Enable All Verbs and Script Engine.
3. Under the *ISAPI Filters* tab, click *Add* and browse to the `php4isapi.dll` or `php5isapi.dll` file. Name the filter "php" and click OK.
4. Under the *Documents* tab, add `index.php` to the list.
5. If using IIS6, goto *Web Service Extensions* and add a new Web Service Extension. Name the extension "php", and click *Add* and browse to the `php4isapi.dll` or `php5isapi.dll` file, enable *Set Extension* status to *Enable*, and click OK.
6. Give the IUSR_XXXX and IIS_WPG users read & execute permissions to the file `%windir%\system32\cmd.exe`. They will also need read permissions on `cacti_web_root/cacti` and it's subfolders.
7. If using IIS6, give the IIS_WPG user modify permissions to the folders `cacti_web_root/cacti/log` and `cacti_web_root/cacti/rrd`.
8. Completely stop and start the IIS service using the following commands:

```
net stop iisadmin
net start w3svc
```

Install Cygwin (optional)

1. Installing a single instance of Cygwin, and using it for all applications that require it is recommended so you do not have different versions of the Cygwin dlls laying around on your system, which can cause conflicts.
2. Run `setup.exe` you previously download.
3. Once you reach the portion of setup entitled *Select Packages*, install the following:

```
Base (include all items)
Libs
  libart_lgpl
  libfreetype26
  libpng12
  zlib
  openssl
Utils
  patch
Web
  wget
```

4. Add `c:\cygwin\bin` to your Windows System PATH environment variable.
5. Move `setup.exe` to `c:\cygwin` for future use.

Install RRDTool

1. Extract the RRDTool zip file from the Cacti web site to `c:\cacti\rrdtool.exe`.

Install MySQL

1. Extract the MySQL zip file to a temp directory and run `setup.exe`.
2. Install MySQL to the default directory, or for the purposes of this manual to the `c:\mysql` directory.

3. If running an older version of MySQL, start it by running `c:\mysql\bin\winmysqladmin.exe`. In more recent versions, this is not required.

4. Set a password for the root user

```
shell> cd mysql\bin
shell> mysqladmin --user=root password somepassword
shell> mysqladmin --user=root --password reload
```

5. Create the MySQL database:

```
shell> mysqladmin --user=root --password create cacti
```

6. Import the default Cacti database:

```
shell> mysql --user=root --password cacti < c:\apache2\htdocs\cacti\cacti.sql
```

7. Create a MySQL username and password for Cacti.

```
shell> mysql --user=root --password mysql
mysql> GRANT ALL ON cacti.* TO cactiuser@localhost IDENTIFIED BY 'somepassword';
mysql> flush privileges;
```

8. If you are running MySQL 4.1 and above, you will need to apply the old password setting in order to authenticate with Cacti. To make this change, stop the MySQL service and add the following to the Start Parameter field. Start it again once it has been added.

```
--old-password
```

You will also need to update the cactiuser account with the old password style.

```
shell> UPDATE mysql.user SET Password = OLD_PASSWORD('cactipwd') WHERE Host = 'localhost' AND User = 'cactiuser'
mysql> FLUSH PRIVILEGES;
```

Install Net-SNMP

1. If you plan to use any hosts with SNMP v2c support, and are using early versions of PHP, you must download and install the Net-SNMP libraries. Net-SNMP provides installers to install their product. However, caution must be taken if you choose to use long file names as Cacti does not them as long file names. You will have to use 8.3 notation. For example `c:\Program Files\Net-SNMP\bin` becomes `c:\progra~1\net-snmpp\bin`.

Install Cactid

1. Extract the Cactid zip file to `c:\cacti` and modify the `cactid.conf.dist` file to include the following statements.

```
DB_Host 127.0.0.1 or hostname (not localhost)
DB_Database cacti
DB_User cactiuser
DB_Password cacti
DB_Port 3306
```

All other pre 0.8.6 settings are obsolete.

2. Cactid now comes with a binary distribution. However, we strongly suggest that you install Cygwin and then remove all the DLL files and `sh.exe` from the `c:\cacti` directory.

Configure Cacti

1. Edit `cacti_web_root/cacti/include/config.php` and specify the MySQL user, password, database, and database port for your Cacti configuration.

```
$database_default = "cacti";
$database_hostname = "localhost";
$database_username = "cactiuser";
$database_password = "cacti";
$database_port = "3306";
```

2. Point your web browser to:

```
http://your-server/cacti/
```

Log in using the username and password of admin/admin. You will be required to change this password immediately.

3. From Cacti, go to *Settings->Paths* and verify/update your paths to point to the correct locations. Recommended examples are posted below. If you plan on using Cactid, then it is very important that all paths include forward slashes instead of backslashes.

PHP Binary Path:

```
c:/php/php.exe
```

RRDTool Binary Path:

```
c:/cacti/rrdtool.exe
```

SNMPGET, SNMPWALK, SNMPBULKWALK, SNMPGETNEXT Paths:

```
c:/progra~1/net-snmp/bin/snmpget.exe
```

```
c:/progra~1/net-snmp/bin/snmpwalk.exe
```

```
c:/progra~1/net-snmp/bin/snmpbulkwalk.exe
```

```
c:/progra~1/net-snmp/bin/snmpgetnext.exe
```

Cacti Logfile Path:

```
c:/mycacti/website/cacti/log/cacti.log
```

Cactid Path:

```
c:/cacti/cactid.exe
```

4. Click on Devices. Delete the Localhost devices as it intended for Linux environments In the upper right corner, click Add. Fill in the following information and then click Add.

Description: My Windows localhost

Hostname: localhost

Host Template: Windows 2000/XP

5. You should now be looking at the localhost device screen. Right under it's name, there should be some SNMP information listed, if not you should double check the SNMP settings on the server and firewall settings. In the upper right-hand corner, click on Create Graphs for this Host. On the following screen, select a disk partition and network interface. At the bottom of the page, click on Create.
6. Log into the user account you'll be using for the scheduled task and verify starting a Cacti polling cycle works. Do this by running the following from the command prompt:

```
php c:/cacti_web_root/cacti/poller.php
```

The output should look something like the following:

```
C:\>php c:\inetpub\wwwroot\cacti\poller.php
```

```
OK u:0.00 s:0.06 r:1.32
```

```
OK u:0.00 s:0.06 r:1.32
```

```
OK u:0.00 s:0.16 r:2.59
```

```
OK u:0.00 s:0.17 r:2.62
```

```
10/28/2005 04:57:12 PM - SYSTEM STATS: Time:4.7272 Method:cmd.php Processes:1 Threads:N/A Hosts:1 HostsPerProce
```

After this has ran once, you should have `cacti.log` in `/cacti/log/` and `rrd` files in `/cacti/rra/`.

7. You are going to need to schedule a task while logged on as an Administrator. This task is required to you can run `poller.php` every 5 minutes. Make sure the Task Scheduler service is started and follow the steps below to begin.

Note: The following instructions are based on Windows XP and Windows Server 2003. You should be able to follow these instructions close enough for Windows 2000 as well.

- a. Select *Start --> Settings --> Control Panel* and double click on *Scheduled Tasks*.
- b. Double click on *Add Scheduled Task*.
- c. Click *Next* and *Browse* on the following screen. Find `c:\php` and select `php.exe`. Choose *Daily* on and click *Next*.
- d. Click *Next* again without changing the time or date settings.
- e. When entering a username and password make sure the user has read and write access to the following directories:

```
cacti_web_root/cacti/rra
```

```
cacti_web_root/log
```

Make sure the user has read, write, and execute access to the following directories:

```
c:\php
```

```
c:\php\sapi
```

- f. Click *Next* and *Finish* to close the wizard.
- g. Right click on the task you just created, and select *Properties*.
- h. Select the *Schedule* tab.
 - i. Make sure *Daily* is selected and click the *Advanced* button.
 - j. Check the *Repeat* checkbox, set it for 5 minutes and set the duration for 24 hours.

k. Click *Ok*

l. In the *Run* textbox enter the following text making sure to use the appropriate paths.

```
c:\php\php.exe c:\mycacti\website\cacti\poller.php
```

The start in box should say `c:\mycacti\website\cacti.`

Apply Patches

1. There are two methods of applying patches to Cacti:
 - a. If you have Cygwin installed, then the patch instructions which use `wget` and `patch`, will work.
 - b. The other method requires you to visit <http://www.cacti.net/downloads/patches/0.8.6h/pre-patched/> and manually download and replace the patched files.
2. You might need to reapply file/folder security on the files patched. Double check they are correct.

Chapter 4. Upgrading Cacti

1. Backup the old Cacti database.

```
shell> mysqldump -l --add-drop-table cacti > mysql.cacti
```

Note: You will probably have to specify the -u and -p flags for the MySQL username and password. This user must have permission to read from Cacti's database or you will end up with an empty backup.

2. Backup the old Cacti directory.

```
shell> mv cacti cacti_old
```

3. Extract the distribution tarball.

```
shell> tar xzvf cacti-version.tar.gz
```

4. Rename the new Cacti directory to match the old one.

```
shell> mv cacti-version cacti
```

5. Edit include/config.php and specify the MySQL user, password and database for your Cacti configuration.

```
$database_default = "cacti";  
$database_hostname = "localhost";  
$database_username = "cactiuser";  
$database_password = "cacti";
```

6. Copy the *.rrd files from the old Cacti directory.

```
shell> cp cacti_old/rra/* cacti/rra/
```

7. Copy any relevant custom scripts from the old Cacti directory. Some script are updated between versions. Therefore, make sure you only over write if the scripts either don't exist or are newer than the distribution's.

```
shell> cp -u cacti_old/scripts/* cacti/scripts/
```

8. Copy any relevant custom resource XML files from the old Cacti directory. Some resource XML files are updated between versions. Therefore, make sure you only over write if the XML files either don't exist or are newer than the distribution's.

```
shell> cp -u -R cacti_old/resource/* cacti/resource/
```

9. Set the appropriate permissions on Cacti's directories for graph/log generation. You should execute these commands from inside Cacti's directory to change the permissions.

```
shell> chown -R cactiuser rra/ log/
```

(Enter a valid username for cactiuser, this user will also be used in the next step for data gathering.)

10. Point your web browser to:

<http://your-server/cacti/>

Follow the on-screen instructions so your database can be updated to the new version.

Chapter 5. Graph Overview

Almost everything in Cacti is somehow related to a graph. At any time, you can list all available graphs by clicking on the *Graph Management* menu item. While it is possible to manually create graphs through this screen, new users should follow the instructions provided in the next chapter for creating new graphs in Cacti.

For users that are familiar with `RRDTOOL`¹, you will immediately recognize that a graph in Cacti is closely modeled after `RRDTOOL`'s graphs. This makes sense since Cacti provides a user friendly interface to `RRDTOOL` without requiring users to understand how `RRDTOOL` works. With this in mind, every graph in Cacti has certain settings and at least one graph item associated with it. While graph settings define the overall properties of a graph, the graph items define the data that is to be represented on the graph. So the graph items define which data to display and how it should be displayed, and also define what should be displayed on the legend.

Each graph and graph item has a set of parameters which control various aspects of the graph. Fortunately through the use of graph templates, it is not necessary to understand the function of each field to create graphs for your network. When you are ready to take on the task of creating your own graph templates, extensive field descriptions for both graphs and graph items are provided in that section of the manual.

Notes

1. <http://www.rrdtool.org/>

Chapter 6. How to Graph Your Network

At this point, you probably realize that graphing is Cacti's greatest strength. Cacti has many powerful features that provide complex graphing and data acquisition, some which have a slight learning curve. Do not let that stop you however, because graphing your network is incredibly simple.

The next two sections will outline the two basic steps which are typically required to create graphs for most devices.

6.1. Creating a Device

The first step to creating graphs for your network is adding a device for each network device that you want to create graphs for. A device specifies important details such as the network hostname, SNMP parameters, and host type.

To manage devices within Cacti, click on the *Devices* menu item. Clicking *Add* will bring up a new device form. The first two fields, *Description* and *Hostname* are the only two fields that require your input beyond the defaults. If your host type is defined under the host template dropdown, be sure to select it here. You can always choose "Generic SNMP-enabled Host" if you are just graphing traffic or "None" if you are unsure. It is important to remember that the host template you choose will not lock you into any particular configuration, it will just provide more intelligent defaults for that type of host.

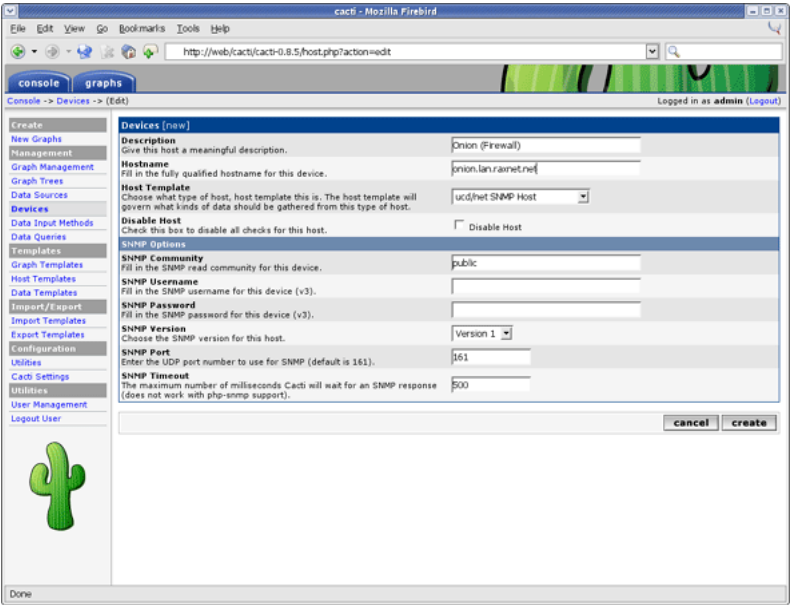


Figure 6-1. Adding a New Device

After saving your new device, you should be redirected back to the same edit form with some additional information. If you configured SNMP for this host by providing a valid community string, you should see various statistics listed at the top of the page. If you see "SNMP error" instead, this indicates an SNMP problem between Cacti and your device.

Towards the bottom of the page there will be two addition boxes, *Associated Data Queries*, and *Associated Graph Templates*. If you selected a host template on the previous page, there will probably be a few items in each box. If there is nothing listed in either box, you will need to associate at least one data query or graph template with your new device or you will not be able to create graphs in the next step. If no available graph template or data query applies to your device, you can check the Cacti templates repository or create your own if nothing currently exists.

6.1.1. A Word About SNMP

The SNMP version that you choose can have a great effect on how SNMP works for you in Cacti. Version 1 should be used for everything unless you have reason to choose otherwise. If you plan on utilizing (and your device supports) high-speed (64-bit) counters, you must select version 2. At the present time, version 3 is not fully implemented and will not work.

The way in which Cacti retrieves SNMP information from a host has an effect on which SNMP-related options are supported. Currently there are three types of SNMP retrieval methods in Cacti and are outlined below.

Table 6-1. SNMP Retrieval Types

| Type | Description | Supported Options | Places Used |
|--------------------------|--|-------------------------------------|---|
| External SNMP | Calls the net-snmp snmpwalk and snmpget binaries that are installed on your system. | All SNMP options | Web interface and PHP poller (poller.php) |
| Internal SNMP (php-snmp) | Uses PHP's SNMP functions which are linked against net-snmp or ucd-snmp at compile time. | Version 1 Only (Community and Port) | Web interface and PHP poller (poller.php) |
| Cactid SNMP | Links directly against net-snmp or ucd-snmp and calls the API directly. | All SNMP options | C-Based Poller (cactid) |

6.2. Creating the Graphs

Now that you have created some devices, it is time to create graphs for these devices. To do this, select the *New Graphs* menu option under the *Create* heading. You will see a screen similar to the image pictured below.

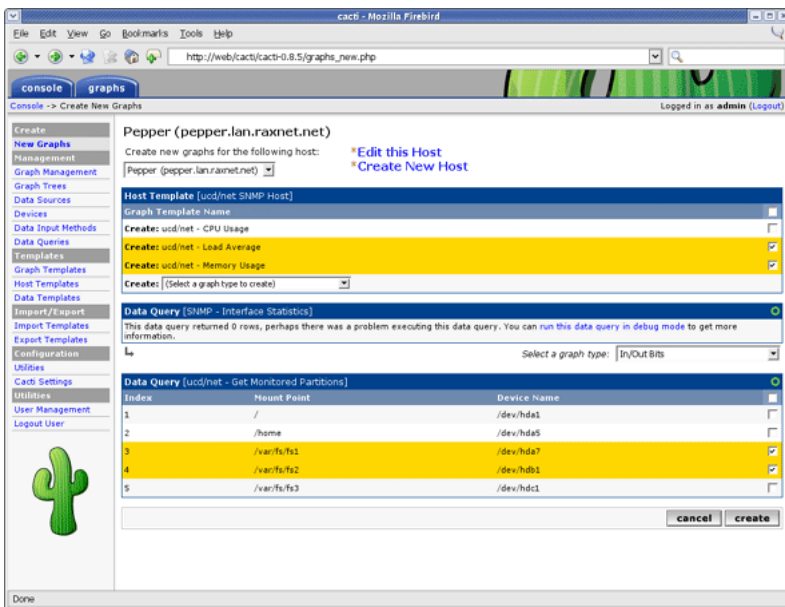


Figure 6-2. Creating New Graphs

The dropdown menu that contains each device should be used to select the host that you want to create new graphs for. The basic concept to this page is simple, place a check in each row that you want to create a graph for and click *Create*.

If you are creating graphs from inside a "Data Query" box, there are a few additional things to keep in mind. First is that you may encounter the situation as pictured above with the "SNMP - Interface Statistics" data query. If this occurs you may want to consult the section on debugging data queries to see why your data query is not returning any results. Also, you may see a "Select a graph type" dropdown box under some data query boxes. Changing the value of this dropdown box affects which type of graph Cacti will make after clicking the *Create* button. Cacti only displays this dropdown box when there is more than one type to choose from, so it may not be displayed in all cases.

Once you have selected the graphs that you want to create, simply click the *Create* button at the bottom of the page. You will be taken to a new page that allows you to specify additional information about the graphs you are about to create. You only see the fields here that are not part of each template, otherwise the value automatically comes from the template. When all of the values on this page look correct, click the *Create* button one last time to actually create your graphs.

If you would like to edit or delete your graphs after they have been created, use the *Graph Management* item on the menu. Likewise, the *Data Source* menu item allows you to manage your data sources in Cacti.

Chapter 7. Viewing Graphs

7.1. Graph Trees

A graph tree can be thought of as a hierarchical way of organizing your graphs. Each graph tree consists of zero or more headers or branch nodes that contain leaf nodes such as graphs or trees. Multiple graph trees or branches within a single tree can be combined to form a very powerful way of organizing your graphs.

7.1.1. Creating a Graph Tree

To create a new graph tree, select the *Graph Trees* menu item under the *Management* header. Select *Add* on this page to create a new tree. The following page will prompt you for a tree name, which will be used to identify the graph tree throughout Cacti. Once you type a name, click the *Create* button to continue. You will be redirected to a page similar to the one below, but without all of the items.

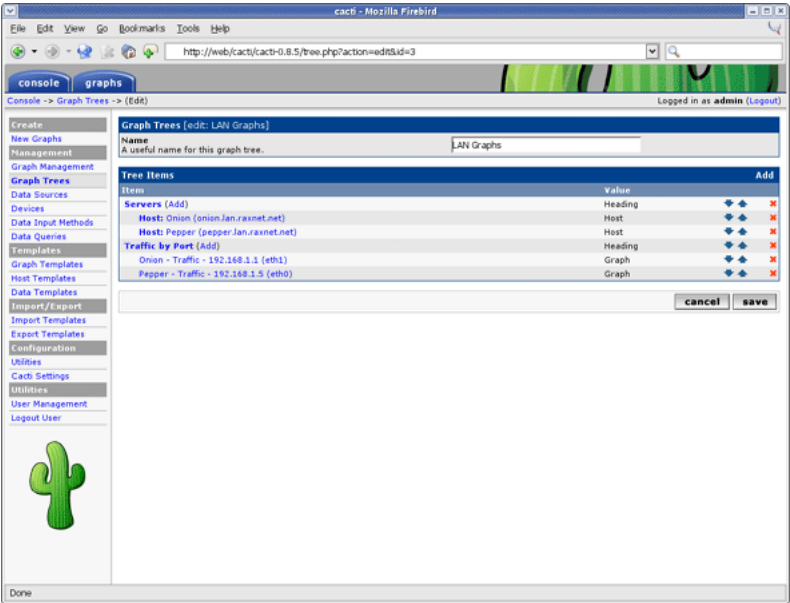


Figure 7-1. Editing a Graph Tree

To start adding items to your tree, click add in the *Tree Items* box. There are currently three different types of tree items you can choose from: header, graph, or host. Simply choose the type you want, fill in the value for that type, and click *Create* to make your new graph tree item. Clicking the *Add* link to the right of any branch will add the new item below that branch, you can change the branch that any item belongs to by changing its *Parent Item* field.

Chapter 8. User Management

In addition to giving you the tools to create sophisticated graphs, Cacti enables you to create users that are tailored specifically to their requirements. Each user has certain settings such as login actions, as well as graph viewing settings. There are also two levels of permissions control, realm permissions and graph permissions which enable you to control what the user can see and change.

Out of the box, there are two users that come with every Cacti installation. The "admin" user, is the main user that by default has access to see and change everything in Cacti. This is the user that you first login with in Cacti, and is probably a good idea to keep around unless you know otherwise. The second user is the "guest" user, which controls which areas/graphs are allowed for unauthenticated users. By default this user only has rights to view, but not change all graphs. This enables any unauthenticated user to visit 'graph_view.php' and view your graphs. This behavior can be changed by either changing the realm permissions for the "guest" user, or disabling the guest user altogether under *Cacti Settings*.

8.1. Creating a New User

To create a new user, select the *User Management* item under the *Utilities* heading on the Cacti menu. Once at the user management screen, click *Add*. You will see a screen that looks similar to the image below.

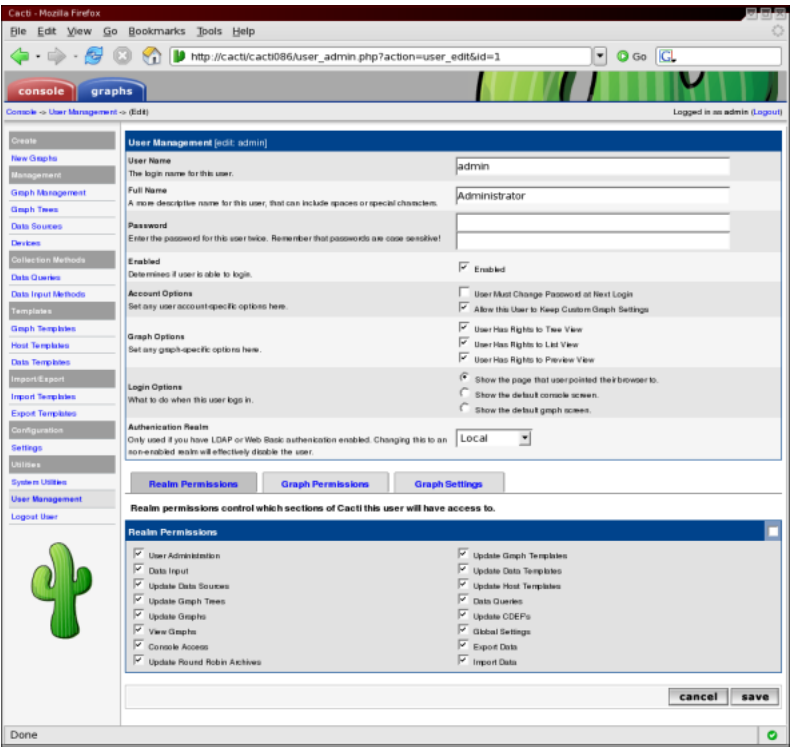


Figure 8-1. Adding a User

At minimum, you must specify a *User Name* and a *Password* for each user. Each user field is described in more detail below. In addition to these fields, each user can have their own realm permissions", graph permissions, and graph settings. Each of these items are described in this section of the manual.

Table 8-1. Field Description: User Management

| Name | Description |
|-----------|---|
| User Name | This field contains the actual login name for the user. There are no character limits this field, but it would make sense to keep it to alphanumeric characters to maintain simplicity. |
| Full Name | (Optional) You can use this field as a more descriptive identifier for the user. It is currently only used for display on the user management page. |

| Name | Description |
|----------------------|---|
| Password | Enter the password for the user twice, once in each text box. Keep in mind that passwords are case sensitive and the password will remain unchanged if the boxes are both left empty. |
| Account Options | There are currently two account related options that can be set for each user. The first, 'User Must Change Password at Next Login', forces a password change immediately after the user logs in. The second option, 'Allow this User to Keep Custom Graph Settings', dictates whether the user can maintain their own custom graph viewing settings which includes expanding and collapsing trees. |
| Graph Options | There are three permission related options that can be set for each user here. The 'User Has Rights to {Tree,List,Preview} View' checkboxes dictate which graph viewing areas the user can see. |
| Login Options | These options dictate what occurs immediately after the user logs in. The first option, 'Show the page that user pointed their browser to.' will point the user to whatever page they were heading to before being interrupted by the login page. The 'Show the default console screen.' option will always point the user to 'index.php' after a successful login. The last option, 'Show the default graph screen.' will point the user to 'graph_view.php' after a successful login. |
| Authentication Realm | Multiple types of users can exist in Cacti. This setting allows you to specify a user for use with Local, Web Basic or LDAP authentication. Template users must be Local. |

8.2. Realm Permissions

Realm permissions control which areas of Cacti a user can access. You can edit a user's realm permissions by selecting *User Management* and choosing the user you want to edit the permissions for. The *Realm Permissions* box will be displayed in the lower part of the screen. Each "realm" is a grouping that represents common tasks in Cacti, making it easier to fine tune each user's access.

If you want to create a user that can only view graphs, you should select the *View Graphs* realm and leave everything else unchecked. See the *Graph Permissions* section for more information about how to fine tune this even more on a per-graph basis. Conversely, if the user needs to access the console, they will need *Console Access* and any additional realms that you see fit.

8.3. Graph Permissions

Graph permissions control which graphs a user is allowed to view, it does not apply to editing graphs. You can edit a user's graph permissions by selecting *User Management* and choosing the user you want to edit the permissions for. Now select the *Graph Permissions* tab to view this user's graph permissions. There are three different ways that you can apply graph permissions, by graph, by host, or by graph template. This works by allowing or denying the user to all graphs associated with the particular group. So you can deny a user to a single graph, all graphs associated with a particular host, or all graphs associated with a particular graph template. Combining these three types of assigning graph permissions results in a very powerful graph policy editor.

Each assignment type contains a *Default Policy* dropdown which can be used to control whether the user should be allowed to view everything or be denied from everything by default in the particular group. It is very important to remember that these policies are evaluated in the order: graph, host, graph template. Therefore, if you set graph's default policy to *Deny*, but kept host and graph template at *Allow*, the user's effective policy would be *Deny* since the graph assignment type is evaluated first. You typically either want to set all default policies to *Deny* for a restrictive user, or *Allow* for a non-restrictive user. If you thoroughly understand Cacti's graph permissions system, these default policies can be combined to provide very complex results.

8.4. Graph Settings

Cacti stores certain graph viewing settings for each user, which enables each user to view graphs in an optimal way. These preferences controls things such as the preview graph size or the default graph viewing mode to use. In addition to managing these settings here under user management, each user can change their own settings by clicking on the *Graphs* tab and selecting the *Settings* tab on the following page. If you want to prevent a user from being able to keep their own graph preferences, uncheck the *Allow this User to Keep Custom Graph Settings* checkbox in user management. The user will be presented with an "Access Denied" error message if they attempt to click the *Settings* tab when this box is checked.

Chapter 9. Data Input Methods

Data input methods allow Cacti to retrieve data to insert into data sources and ultimately put on a graph. There are different ways for Cacti to retrieve data, the most popular being through an external script or from SNMP.

9.1. Creating a Data Input Method

To create a new data input method, select the *Data Input Methods* option under the *Management* heading. Once on that screen, click *Add* on the right. You will be presented with a few fields to populate on the following screen.

Table 9-1. Field Description: Data Input Methods

| Name | Description |
|---------------|--|
| Name | Give the data query a name that you will use to identify it. This name will be used throughout Cacti to identify the data input method. |
| Input Type | Select the type of data input method you are trying to create. Valid options here are 'Script/Command', 'SNMP', 'SNMP Query', and 'Script Query'. |
| Input String | <p>This field is only used when the <i>Input Type</i> is set to 'Script/Command'. It specifies the full path to the script including any per data source variables inside <>'s from the user. For instance, if you are passing an IP address to a script, your input string might look something like:</p> <pre>/path/to/script.pl <ip></pre> <p>When the user creates a data source based on this data input method, they will be prompted for an IP address to pass onto the script.</p> |
| Output String | This field is only here for legacy purposes and will probably go away in a future version of Cacti. |

When you are finished filling in all necessary fields, click the *Create* button to continue. You will be redirected back to the same page, but this time with two new boxes, *Input Fields* and *Output Fields*. The *Input Fields* box is used to define any fields that require information from the user. Any input fields referenced to in the input string must be defined here. The *Output Fields* box is used to define each field that you expect back from the script. *All data input methods must have at least one output field defined*, but may have more for a script.

9.1.1. Data Input Fields

To define a new field, click *Add* next to the input or output field boxes. You will be presented with some or all of the fields below depending on whether you are adding an input or output field.

Table 9-2. Field Description: Data Input Fields

| Name | Description |
|---------------------------------------|---|
| Field/Field Name | Either select or type in the name of the field you want to define. No spaces or other non-alphanumeric characters (except '-' or '_') should be used in this field. |
| Friendly Name | Enter a more descriptive name for this field which will be used for identification throughout Cacti. |
| Regular Expression Match (Input Only) | If you want to enforce a certain regular expression pattern when the user enters a value for this field, enter it here. The regular expression must follow POSIX syntax as it will be passed to PHP's <code>ereg()</code> function. |
| Allow Empty Input (Input Only) | Select whether you want to allow the user to leave the value of this field blank or not. |
| Special Type Code (Input Only) | Sometimes Cacti needs to reference a field internally, but needs to rely on more than just the field name. For instance, if your field requires an IP address from the user, you can enter 'management_ip' here and Cacti will fill this field in with the current IP address of the selected host. Valid values for this field are: 'hostname', 'management_ip', 'snmp_community', 'snmp_username', 'snmp_password', and 'snmp_version'. |

| Name | Description |
|-------------------------------|--|
| Update RRD File (Output Only) | Check this box if you want Cacti to insert the return value from this field into the RRD file. Obviously, this box needs to be checked for at least one output field per data input source, but can be left blank to have Cacti store the value in the database instead. |

When you are finished filling in all necessary fields, click the *Create* button to continue. You will be redirected back to the data input method edit page. From here you can continue to add additional fields, or click *Save* on this screen when finished.

9.2. Making Your Scripts Work With Cacti

The simplest way to extend Cacti's data gathering functionality is through external scripts. Cacti comes with a number of scripts out of the box which are located in the `scripts/` directory. These scripts are used by the data input methods that are present in a new installation of Cacti.

To have Cacti call an external script to gather data you must create a new data input method, making sure to specify *Script/Command* for the *Input Type* field. See the previous section, *Creating a Data Input Method*¹ for more information about how to create a data input method. To gather data using your data input method, Cacti simply executes the shell command specified in the *Input String* field. Because of this, you can have Cacti run any shell command or call any script which can be written in almost any language.

What Cacti is concerned with is the output of the script. When you define your data input method, you are required to define one or more output fields. The number of output fields that you define here is important to your script's output. For a data input method with only one output field, your script should output its value in the following format:

```
<value_1>
```

So if I wrote a script that outputs the number of running processes, its output might look like the following:

Example 9-1. Example script output using 1 field

```
67
```

Data input methods with more than one output field are handled a bit differently when writing scripts. Scripts that output more than one value should be formatted like the following:

```
<fieldname_1>:<value_1> <fieldname_2>:<value_2> ... <fieldname_n>:<value_n>
```

Lets say that I write a script that outputs the 1, 5, and 10 minute load average of a Unix machine. In Cacti, I name the output fields '1min', '5min', and '10min', respectively. Based on these two things, the output of the script should look like the following:

Example 9-2. Example script output using 3 fields

```
1min:0.40 5min:0.32 10min:0.01
```

One last thing to keep in mind when writing scripts for Cacti is that they will be executed as the user the data gatherer runs as. Sometimes a script may work correctly when executed as root, but fails due to permissions problems when executed as a less privileged user.

Notes

1. `data_input_methods.html`

Chapter 10. Data Queries

Data queries are not a replacement for data input methods in Cacti. Instead they provide an easy way to query, or list data based upon an index, making the data easier to graph. The most common use of a data query within Cacti is to retrieve a list of network interfaces via SNMP. If you want to graph the traffic of a network interface, first Cacti must retrieve a list of interfaces on the host. Second, Cacti can use that information to create the necessary graphs and data sources. Data queries are only concerned with the first step of the process, that is obtaining a list of network interfaces and not creating the graphs/data sources for them. While listing network interfaces is a common use for data queries, they also have other uses such as listing partitions, processors, or even cards in a router.

One requirement for any data query in Cacti, is that it has some unique value that defines each row in the list. This concept follows that of a 'primary key' in SQL, and makes sure that each row in the list can be uniquely referenced. Examples of these index values are 'ifIndex' for SNMP network interfaces or the device name for partitions.

There are two types of data queries that you will see referred to throughout Cacti. They are script queries and SNMP queries. Script and SNMP queries are virtually identical in their functionality and only differ in how they obtain their information. A script query will call an external command or script and an SNMP query will make an SNMP call to retrieve a list of data.

All data queries have two parts, the XML file and the definition within Cacti. An XML file must be created for each query, that defines where each piece of information is and how to retrieve it. This could be thought of as the actual query. The second part is a definition within Cacti, which tells Cacti where to find the XML file and associates the data query with one or more graph templates.

10.1. Creating a Data Query

Once you have created the XML file that defines your data query, you must add the data query within Cacti. To do this you must click on *Data Queries* under the *Data Gathering* heading, and select *Add*. You will be prompted for some basic information about the data query, described in more detail below.

Table 10-1. Field Description: Data Queries

| Name | Description |
|-------------------|--|
| Name | Give the data query a name that you will use to identify it. This name will be used throughout Cacti when presented with a list of data queries. |
| Description | (Optional) Enter a more detailed description of the data query including the information it queries or additional requirements. |
| XML Path | Fill in the full path to the XML file that defines this query. You can optionally use the <path_cacti> variable that will be substituted with the full path to Cacti. On the next screen, Cacti will check to make sure that it can find the XML file. |
| Data Input Method | This is how you tell Cacti to handle the data it receives from the data query. Typically, you will select "Get SNMP Data (Indexed)" for an SNMP query and "Get Script Data (Indexed)" for a script query. |

When you are finished filling in all necessary fields, click the *Create* button to continue. You will be redirected back to the same page, but this time with some additional information to fill in. If you receive a red warning that says 'XML File Does Not Exist', correct the value specified in the 'XML Path' field.

10.1.1. Associated Graph Templates

Every data query must have at least one graph template associated with it, and possibly more depending on the number of output fields specified in the XML file. This is where you get to choose what kind of graphs to generate from this query. For instance, the interface data query has multiple graph template associations, used to graph traffic, errors, or packets. To add a new graph template association, simply click *Add* at the right of the *Associated Graph Templates* box. You will be presented with a few fields to fill in:

Table 10-2. Field Description: Associated Graph Templates

| Name | Description |
|------|---|
| Name | Give a name describing what kind of data you are trying to represent or graph. When the user creates a graph using this data query, they will see a list of graph template associations that they will have to choose from. |

| Name | Description |
|----------------|--|
| Graph Template | Choose the actual graph template that you want to make the association with. |

When you are finished filling in these fields, click the *Create* button. You will be redirected back to the same page with some additional information to fill in. Cacti will make a list of each data template referenced to in your selected graph template and display them under the *Associated Data Templates* box. For each data source item listed, you must select the data query output field that corresponds with it. *Do not forget to check the checkbox to the right of each selection, or your settings will not be saved.*

The *Suggested Values* box gives you a way to control field values of data sources and graphs created using this data query. If you specify multiple suggested values for the same field, Cacti will evaluate them in order which you can control using the up or down arrow icons. For more information about valid field names and variables, read the section on suggested values.

When you are finished filling in all necessary fields on this form, click the *Save* button to return to the data queries edit screen. Repeat the steps under this heading as many times as necessary to represent all data in your XML file. When you are finished with this, you should be ready to start adding your data query to hosts.

10.2. SNMP Query XML Syntax

```
<query>
  <name>Get SNMP Interfaces</name>
  <description>Queries a host for a list of monitorable interfaces</description>
  <oid_index>.1.3.6.1.2.1.2.2.1.1</oid_index>
  <oid_index_parse>OID/REGEXP:.*\.[0-9]{1,3}\.[0-9]{1,3}$</oid_index_parse>
  <oid_num_indexes>.1.3.6.1.2.1.2.1.0</oid_num_indexes>
  <index_order>ifDescr;ifName;ifIndex</index_order>
  <index_order_type>numeric</index_order_type>
  <index_title_format>|chosen_order_field|</index_title_format>

  <fields>
    <ifIndex>
      <name>Index</name>
      <method>walk</method>
      <source>value</source>
      <direction>input</direction>
      <oid>.1.3.6.1.2.1.2.2.1.1</oid>
    </ifIndex>
  </fields>
</query>
```

Table 10-3. SNMP Query XML Field Reference

| Field | Description |
|--------------------|--|
| query->name | (Optional) You can enter a "friendly name" for the SNMP query here. It will not be used by Cacti, and is for identification only. |
| query->description | (Optional) You can enter a description for the SNMP query here. It will not be used by Cacti, and is for identification only. |
| query->oid_index | Every SNMP query must have an OID that represents the index values for the query when walked. As described above, any data query in Cacti must contain a field that uniquely identifies each row returned by the query. In the example above, the oid_index points to the OID of ifIndex in the interface MIB. <i>Note:</i> Starting with version 0.8.6c, Cacti is able to parse unique indexes from the OID itself. While the regular expression used for parsing the value from the OID is defined below, you must still specify an OID that can be walked by Cacti in order to obtain the list of OID's. Any OID defined for one of your input fields should work in this case. The values returned from the snmpwalk walk will be completely disregarded. |

| Field | Description |
|--------------------------------|--|
| query->oid_index_parse | This field should only be used if you are trying to parse the unique index from the OID itself. If this field is defined, to obtain a list of indexes, Cacti walks the OID provided in the oid_index field above. It then applies the regular expression provided in this field to the list of OID's that are returned. The matched substrings that remain become the list of indexes for this SNMP query. |
| query->oid_num_indexes | An OID that can be queried to determine the total number of available indexes. If specified, this will be used to determine when to automatically recache this SNMP query when it is attached to a device. |
| query->index_order | As of version 0.8.6, Cacti will attempt to find the best field to index off of based on whether each row in the query is unique and non-null. If specified, Cacti will perform this check on the fields listed here in the order specified. Only input fields can be specified and multiple fields should be delimited with a colon. |
| query->index_order_type | For sorting purposes, specify whether the index is numeric or alphanumeric. <i>numeric</i> : The indexes in this SNMP query are to be sorted numerically (ie. 1,2,3,10,20,31) <i>alphabetic</i> : The indexes in this SNMP query are to be sorted alphabetically (1,10,2,20,3,31). |
| query->index_title_format | Specify the title format to use when representing an index to the user. Any input field name can be used as a variable if enclosed in pipes (). The variable <i> chosen_order_field </i> will be substituted with the field chosen by Cacti to index off of (see index_order above). |
| query->fields | Each field contained within the SNMP query must be defined under this tag. |
| query->fields->ifIndex | Each defined field in the SNMP query must have a unique name given to it. Do not use spaces or any non-alphanumeric characters, this name must be identifiable within Cacti. |
| query->fields->ifIndex->name | Here you can specify a "friendly name" for the field. This name will be used by Cacti to help the user identify this field. |
| query->fields->ifIndex->method | Tell Cacti how you want it to gather SNMP information for this field. <i>get</i> : The 'get' method obtains a list of indexes and does an snmpget for each index of the OID specified for this field. <i>walk</i> : The 'walk' method does a walk of the OID specified for this field. Both methods will return the same values, even though the 'walk' method is typically more efficient. |

| Field | Description |
|-----------------------------------|--|
| query->fields->ifIndex->source | <p>When Cacti obtains a list for this field, you need to tell it how to derive its value for each row.</p> <p><i>value</i>: The 'value' option simply returns the result of the snmpget for each row.</p> <p><i>OID/REGEXP: (regexp_match)</i>: The 'OID/REGEXP:(regexp_match)' can be used when you need to use a POSIX-based regular expression to derive the value from the OID. The most common example of this is to retrieve the IP address of an interface, and can be seen in the 'interface.xml' file.</p> <p><i>VALUE/REGEXP: (regexp_match)</i>: The 'OID/REGEXP:(regexp_match)' option can be used to parse the value based on a regular expression, returning the first match.</p> <p><i>index</i>: Simply use the value of the index for this row as the value. If the index is being parsed from the OID using the oid_index_parse field, you must specify <i>index</i> here.</p> |
| query->fields->ifIndex->direction | <p><i>input</i>: Input values are the "known" values that you will use to derive the output values, this is where the "query" part of SNMP query comes in. When you create a graph based on an SNMP query, Cacti will prompt you to choose the input value to base the graph on.</p> <p><i>output</i>: Output values are "unknown" values that are returned from the script. An SNMP query may return multiple statistics for a single index. For instance, a single interface could return bytes/sec in, errors, packets/sec, etc.</p> <p>A rule of thumb is that input fields contain semi-static data that is not graphable, while the output fields contain the data that will be graphed.</p> |
| query->fields->ifIndex->oid | You must specify the actual OID that corresponds with the field. Each value for this field can be obtained by doing an snmpget on 'oid.(each)snmpindex'. |

10.3. Script Query XML Syntax

```

<query>
  <name>Get Unix Mounted Partitions</name>
  <description>Queries a list of mounted partitions on a unix-based host with the 'df' command.</description>
  <script_path>perl |path_cacti|/scripts/query_unix_partitions.pl</script_path>
  <arg_index>index</arg_index>
  <arg_query>query</arg_query>
  <arg_get>get</arg_get>
  <arg_num_indexes>num_indexes</arg_num_indexes>
  <output_delimiter>:</output_delimiter>
  <index_order>dskDevice:dskMount</index_order>
  <index_order_type>alphabetic</index_order_type>
  <index_title_format>|chosen_order_field|</index_title_format>

  <fields>
    <dskDevice>
      <name>Device Name</name>
      <direction>input</direction>
      <query_name>device</query_name>
    </dskDevice>
  </fields>
</query>

```

Table 10-4. Script Query XML Field Reference

| Field | Description |
|-------|-------------|
|-------|-------------|

| Field | Description |
|-------------------------------------|---|
| query->name | (Optional) You can enter a "friendly name" for the script query here. It will not be used by Cacti, and is for identification only. |
| query->description | (Optional) You can enter a description for the script query here. It will not be used by Cacti, and is for identification only. |
| query->script_path | Enter the complete path to the script or executable that is going to handle your script query. When in doubt, specify the full path to all binaries referenced in this path, the query may not execute otherwise. |
| query->arg_index | Enter the argument that is to be passed to the script to retrieve a list of indexes. |
| query->arg_query | Enter the argument that is to be passed to the script to retrieve a list of values given a field name. |
| query->arg_get | Enter the argument that is to be passed to the script to retrieve a single value given a field name and index value. |
| query->arg_num_indexes | Enter the argument that is to be passed to the script to determine the total number of available indexes. If specified, this will be used to determine when to automatically recache this script query when it is attached to a device. |
| query->output_delimiter | Enter the one character delimiter that will be used to separate output values. This is only used when you "query" the script in which case it outputs 'index(delimiter)value'. |
| query->index_order | As of version 0.8.6, Cacti will attempt to find the best field to index off of based on whether each row in the query is unique and non-null. If specified, Cacti will perform this check on the fields listed here in the order specified. Only input fields can be specified and multiple fields should be delimited with a comma. |
| query->index_order_type | For sorting purposes, specify whether the index is numeric or alphanumeric. <i>numeric</i> : The indexes in this script query are to be sorted numerically (ie. 1,2,3,10,20,31) <i>alphabetic</i> : The indexes in this script query are to be sorted alphabetically (1,10,2,20,3,31). |
| query->index_title_format | Specify the title format to use when representing an index to the user. Any input field name can be used as a variable if enclosed in pipes (). The variable <i> chosen_order_field </i> will be substituted with the field chosen by Cacti to index off of (see <i>index_order</i> above). |
| query->fields | Each field contained within the script query must be defined under this tag. |
| query->fields->dskDevice | Each defined field in the script query must have a unique name given to it. Do not use spaces or any non-alphanumeric characters, this name must be identifiable within Cacti. |
| query->fields->dskDevice->name | Here you can specify a "friendly name" for the field. This name will be used by Cacti to help the user identify this field. |
| query->fields->dskDevice->direction | <i>input</i> : Input values are the "known" values that you will use to derive the output values, this is where the "query" part of script query comes in. When you create a graph based on a script query, Cacti will prompt you to choose the input value to base the graph on. <i>output</i> : Output values are "unknown" values that are returned from the script. A script query may return multiple statistics for a single index. For instance, a single partition could return free disk space, total disk space, fragmentation percentage, etc. A rule of thumb is that input fields contain semi-static data that is not graphable, while the output fields contain the data that will be graphed. |

| Field | Description |
|--------------------------------------|---|
| query->fields->dskDevice->query_name | Enter the name that Cacti must use when asking the script for information about this field. For instance, the following should return values: <code>'(script_name) query (query_name)'</code> . |

Chapter 11. Templates

11.1. Graph Templates

In Cacti, a graph template provides a skeleton for an actual graph. If you have many graphs that share most of their characteristics, using a graph template would probably make sense. After a graph is attached to a particular graph template, all changes made to the graph template will propagate out to all of its graphs.

11.1.1. Creating a Graph Template

To create a new graph template, select *Graph Templates* under the *Templates* heading and click *Add*.

The first thing you must do is give the template a name. This name has nothing to do with the graph title, but is what you will use to identify the template throughout Cacti. Second, you will notice a list of graph field names with Use Per-Graph Value checkboxes next to each one. The nice thing about templates in Cacti is that you can choose whether to template each field on a per-field basis. If you leave the checkbox unchecked, every graph attached to the template will inherit its value from the template. If the checkbox is checked, every graph attached to the template will contain its own value for that particular field.

Table 11-1. Field Description: Graph Templates

| Name | Description |
|--------------------------|---|
| Title | The title of the graph within Cacti and the title that will be printed on the actual graph itself. <i>Note:</i> You can use the keyword <code> host_description </code> in this field, which will be automatically substituted with the current host description. |
| Image Format | Choose whether you would prefer to output your graph images in PNG or GIF. According to the RRDTool graph manual, PNG is preferred because it is both smaller and faster. |
| Height | The height of the graph area in pixels |
| Width | The width of the graph area in pixels |
| Auto Scale | Enable auto scale for the graph. This option must be checked to use the next two options. Upper/Lower limit values are ignored when using autoscale since these boundaries are determined automatically. |
| Auto Scale Options | Choose whether you would like to use <code>--alt-autoscale</code> or <code>--alt-autoscale-max</code> on the graph. The RRDTool graph manual says: "Where <code>--alt-autoscale</code> will modify both the absolute maximum AND minimum values, this option will only affect the maximum value. The minimum value, if not defined on the command line, will be 0. This option can be useful when graphing router traffic when the WAN line uses compression, and thus the throughput may be higher than the WAN line speed". |
| Logarithmic Auto Scaling | Choose if you want logarithmic y-axis scaling. |
| Rigid Boundaries Mode | From the RRDTool manual "Normally <code>rrdgraph</code> will automatically expand the lower and upper limit if the graph contains a value outside the valid range. With this option you can disable this behavior". |
| Auto Padding | If you have ever created RRDTool-based graphs manually, you may have realized how annoying it can be to get text to line up properly. With this option Cacti will do its best to make the columns on your graph legend line up. This option works best when your graph legend has a consistent number of columns. |
| Allow Graph Export | If you choose to use Cacti's static HTML/image export, you can choose whether you want each individual graph to be exported. |
| Upper Limit | The maximum value that will be displayed on the y-axis. This value is ignored if auto-scaling is turned on. |
| Lower Limit | The minimum value that will be displayed on the y-axis. This value is ignored if auto-scaling is turned on. |
| Base Value | Whether you want to base the y-axis labels on 1000 or 1024. This field will typically be set to 1024 for memory and 1000 for traffic measurements. |

| Name | Description |
|---------------------|---|
| Unit Value | Sets the unit value for the y-axis. You should probably keep the unit short as to keep it readable. |
| Unit Exponent Value | Sets the 10^e scaling of the y-axis. Valid values for this field are between -18 and 18. For example, you could use 3 to display everything in 'k' (kilo) or -6 to display everything in 'u' (micro). |
| Vertical Label | The text to print on the left edge of the graph. Usually is the units the data on the graph is measured in. |

Note: For most graph templates, you will want to check the *Use Per-Graph Value* checkbox for the title field so each graph using this template has its own unique title. It also makes sense to enter an initial value in this field that includes the variable `|host_description|` for organizational purposes.

When you are finished filling in values for the graph template, click *Create* and you will be presented with a page similar to the graph edit page.




Figure 11-1. Adding a Graph Template

11.1.1.1. Graph Items

The first thing you should do is create graph items for this graph template, just like for a regular graph. One difference you will notice is that the *Data Sources* dropdown will contain a list of data template items rather than data source items. It is important that Cacti can make this association here, so that Cacti doesn't have to make unnecessary assumptions later.

Table 11-2. Field Description: Graph Template Items

| Name | Description |
|-------------|---|
| Data Source | If this graph item is to represent some sort of data, you must select a your data source here. Keep in mind that not all graph items have a data source. Graph item types such as COMMENT, VRULE, and HRULE typically do not. |
| Color | Depending on the graph item type, you can select a color for the graph item. This field only applies to the graph item types AREA, STACK, LINE1, LINE2, and LINE3. |

| Name | Description |
|------------------------|--|
| Graph Item Type | This field is important because it defines what kind of graph item this is. Types such as AREA, STACK, LINE1, LINE2, and LINE3 are used to represent data on the graph, while COMMENT and GPRINT are used put on the legend. The LEGEND type is specific to Cacti only and can be used if you want to Cacti to automatically create a GPRINT-LAST/GPRINT-AVERAGE/GPRINT-MAXIMUM graph item for you. <i>Note:</i> You must always put an AREA item before using STACK or your graph will not render. |
| Consolidation Function | This tells <code>RRDTOOL</code> which consolidation function to use when representing this data on the graph. You will typically use AVERAGE for most things on the graph area, and LAST/MAXIMUM as well for GPRINT items. |
| CDEF Function | If you want to apply a CDEF function to the graph item, select one here. Check out the CDEF section of the manual for more information. |
| Value | This field is only used with the HRULE/VRULE graph item types. Type any valid integer to draw the line at for HRULE or the time of the day HH:MM for VRULE. |
| GPRINT Type | If this item is a GPRINT, you can choose how you want the number to be formatted. You can add your own in the GPRINT Presets section of Cacti. |
| Text Format | You can enter text to be displayed on the legend here. This field is applicable for all graph item types except for the virtual LEGEND type. |
| Hard Return | Check this box to force graph items onto the next line. |

11.1.1.2. Graph Item Inputs

After creating graph items for your template, you will need to create some graph item inputs. Graph item inputs are unique to graph templates because of the large number of items they sometimes contain. Graph item inputs enable you to take one graph item field, and associate it with multiple graph items.

To create a new graph item input, click *Add* on the right of the *Graph Item Inputs* box. There are various fields that must be filled in for every graph item input:

Table 11-3. Field Description: Graph Template Items

| Name | Description |
|------------------------|--|
| Name | This will be the name used to identify the graph item input on both the graph template and graph edit pages. |
| Description | (Optional) This description will be displayed on the graph edit page of any graph using this template if specified. |
| Field Type | You must choose the field that you are going to associate with one or more graph items. |
| Associated Graph Items | Choose one or more graph items to associate with the field selected for "Field Type". When the user specifies a value for the field, it will be applied to all of the items you select here. |

11.1.2. Applying Graph Templates to Graphs

Applying a graph template to a graph is a very simple process. The first thing you must do is select the graph you want to apply the template to under *Graph Management*. Under the *Graph Template Selection* box, select the graph template that you want to apply to the graph and click *Save*. If this is a new graph or the graph and graph template contains an equal number of graph items, the graph template will be automatically applied. If number of graph items varies from the graph to the target graph template, you will be prompted with a warning that your graph will be changed.

Once the template is applied to the graph, you will notice that you can only change values for the fields that you checked Use Per-Graph Value for. You will also notice a new box, called *Graph Item Inputs*. This is where you can specify values for the graph items inputs that you defined in the graph template. The values specified here will be applied to each graph item tied to the graph item input.

Now any time a change is made to the graph template, it will be automatically propagated to the graphs attached to it.

11.2. Data Templates

In Cacti, a data template provides a skeleton for an actual data source. If you have many data sources that share most of their characteristics, using a data template would probably make sense. After a data source is attached to a particular data template, all changes made to the data template will propagate out to all of its data sources.

11.2.1. Creating a Data Template

To create a new data template, select *Data Templates* under the *Templates* heading and click *Add*.

The first thing you must do is give the template a name. This name has nothing to do with the data source name, but is what you will use to identify the template throughout Cacti. Second, you will notice a list of data source/data source item field names with *Use Per-Data Source Value* checkboxes next to each one. The nice thing about templates in Cacti is that you can choose whether to template each field on a per-field basis. If you leave the checkbox unchecked, every data source attached to the template will inherit its value from the template. If the checkbox is checked, every data source attached to the template will contain its own value for that particular field.

Table 11-4. Field Description: Data Templates

| Name | Description |
|--------------------|--|
| Name | This is the name that cacti uses to identify this data source. You must type something here, but you may include spaces and other non-alphanumeric characters. <i>Note:</i> You can use the keyword <code> host_description </code> in this field, which will be automatically substituted with the current host description. |
| Data Source Path | The full path to the RRD file associated with this data source. When creating a new data source you will typically leave this field blank as Cacti will generate one for you. If you are using an RRD file generated outside of Cacti, you need to specify it here. |
| Data Input Method | Here is where you tell cacti how it is supposed to fetch data for this data source. There are several data input sources that come with cacti, and you can add your own by going to Data Input Methods. If this RRD file is being populated outside of cacti, make sure to leave this field set to "None". |
| Associated RRAs | You need to associate the data source with at least one RRA so <code>RRDTOOL</code> knows how often and for how long to keep its data. You will almost always want to select all of these values however so you can render daily, weekly, monthly, and yearly graphs. |
| Step | This tells <code>RRDTOOL</code> how many seconds there will be between updates. The default is 300 seconds (5 minutes), and is sufficient for most installations. |
| Data Source Active | This is a quick and easy to tell Cacti to stop gathering data for this data source. The data source can still be used on graphs, but no data will be fed to it until it is made active again. |

Note: For most data templates, you will want to check the *Use Per-Graph Value* checkbox for the name field so each data source using this template has its own unique name. It also makes sense to enter an initial value in this field that includes the variable `|host_description|` for organizational purposes.

When you are finished filling in values for the data template, click *Create* and you will be presented with a screen similar to the data source edit screen.

Figure 11-2. Adding a Data Template

11.2.1.1. Data Source Items

Like a graph, a data source can have more than one items. This is useful in situations where a script returns more than piece of data at one time. This also applies to data queries, so you can have a single data template that contains both inbound and outbound traffic, rather than having to create a separate data template for each.

Table 11-5. Field Description: Data Source Items

| Name | Description |
|---------------------------|---|
| Internal Data Source Name | This is the name used by <code>RRDTOOL</code> to identify this particular data source within the RRD file. <code>RRDTOOL</code> places a limit of 19 alphanumeric characters (plus <code>'_'</code> and <code>'-'</code>) on this field. |
| Minimum Value | Here is where you specify the minimum value that is expected for this data source in the RRD file. If a value lower than the minimum is given, it will be stored as Unknown (U). |
| Maximum Value | Here is where you specify the maximum value that is expected for this data source in the RRD file. If a value higher than the maximum is given, it will be stored as Unknown (U). |
| Data Source Type | There are four types of data that <code>RRDTOOL</code> can represent for any given data source: COUNTER, GAUGE, ABSOLUTE, and DERIVE. A gauge is for numbers that are not continuously incrementing. A counter is for continuously incrementing values that might overflow, while a derive does not overflow. The absolute type is used for counters that are reset upon reading. |
| Heartbeat | As defined by <code>RRDTOOL</code> : "The maximum amount of time that can pass before data is entered as "unknown". This field is usually '600' or 2 data gathering intervals". |

11.2.1.2. Custom Data

Assuming you selected a data input source on the previous screen, you should now be presented with a *Custom Data* box. Each custom data field is per-field templatable as all of the other data source fields are. Even if you select the *Use Per-Data Source Value* checkbox, it might be useful to specify a value that will be used as an "initial value" for any data source using this data template.

11.2.2. Applying Data Templates to Data Sources

Applying a data template to a data source is a very simple process. The first thing you must do is select the data source

you want to apply the template to under *Data Sources*. Under the *Data Template Selection* box, select the data template that you want to apply to the data source and click *Save*.

Once the template is applied to the data source, you will notice that you can only change values for the fields that you checked *Use Per-Data Source Value* for.

Now any time a change is made to the data template, it will be automatically propagated to the data sources attached to it.

11.3. Host Templates

Host templates in Cacti serve a different purpose than data and graph templates. Instead of abstracting the fields of a host, a host template allows you to associate graph templates and data queries with a given host type. This way when you assign a host template to a host, all of the relevant graphs to that host type are only one click away from the user.

11.3.1. Adding a Host Template

To create a new host template in Cacti, select the *Host Templates* option under the *Templates* heading. Once on that screen, click *Add* on the right. Type a unique name for the host template and click the *Create* button. You will be redirected back to the edit page with the *Associated Graph Templates* and *Associated Data Queries* boxes. These two boxes allow you to associate certain graph templates or data queries with the host template. Simply select something from the dropdown menu and click *Add* to associate it with your host template.

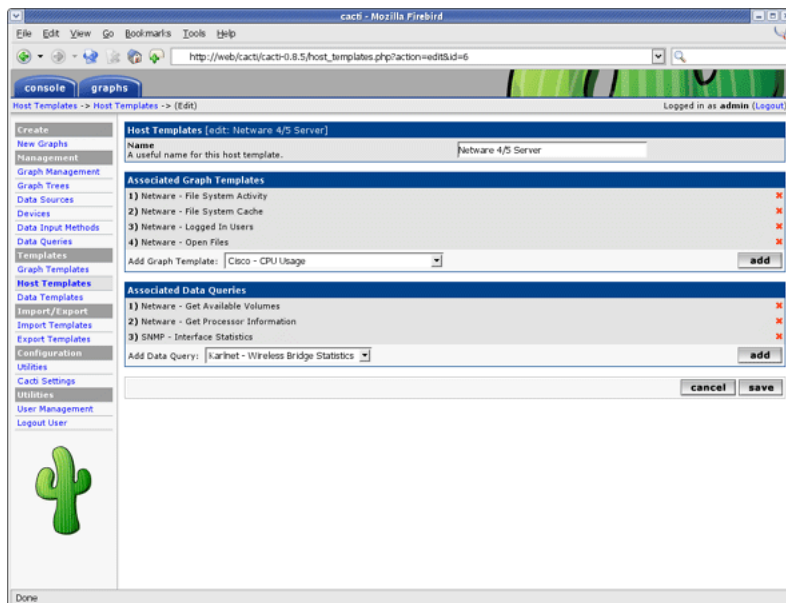


Figure 11-3. Adding a Host Template

Chapter 12. PHP Script Server

The PHP Script Server is a new feature in Cacti 0.8.6. This new feature allows for the rapid execution of PHP based Data Queries in Cacti. The Script Server process is launched by the poller during every polling cycle. It listens for commands from the poller, executes them, and then waits for an quit signal.

The reason that it is so fast is that PHP is started in memory only one time, and for every Data Query called, it's code is interpreted only once. The resulting Data Query binaries are therefore very efficient. Using the Script Server process over the traditional POPEN process nets a 20+ fold speed improvement in Cacti.

Since PHP scripts are so powerful, this new feature in Cacti, makes it an excellent choice for collecting non-SNMP and SNMP based data.

12.1. Using the Script Server

Cacti 0.8.6 contains two sample script server routines. They are for the collection of HostMib CPU and Disk Partition information. These two examples are based off the traditional POPEN version of the HostMib functions found in earlier versions of Cacti.

For new installs, the HostMib functions are defaulted to using the PHP Script Server, therefore, you don't need to do anything to use it.

For upgrades, you must make several changes to start using the PHP Script Server for the HostMib CPU and HostMib Partitions Data Queries. To migrate you must follow the step below.

12.2. Upgrade Steps for the Example HostMib Data Queries

If you are using the two built in script queries, "SNMP - Get Mounted Partitions" and "SNMP - Get Processor Information", you can migrate to the PHP Script Server using the steps below:

1. Verify the existence of New Data Input Method - Go to *Data Input Methods*, verify that you see the "Get Script Server Data (Indexed)" Data Input Method exists and that it is using the "Script Query - Script Server" method.
2. Disable the Poller - Goto *Settings->Poller* and uncheck the *Poller Enabled* checkbox. Press *Save*.
3. Update Your Data Queries - Go to *Data Queries* and Edit the two Data Queries below. Change both their Data Input Method to "Get Script Server Data (Indexed)" and in the XML path replace "script_query" with "script_server" as shown below:
 - a. SNMP - Get Mounted Partitions
`<path_cacti>/resource/script_server/host_disk.xml`
 - b. SNMP - Get Processor Information
`<path_cacti>/resource/script_server/host_cpu.xml`
4. Update Your Data Templates - Goto *Data Templates*, locate the following two data templates and change their data input method to "Get Script Server Data (Indexed)".
 - a. Host MIB - Hard Drive Space
 - b. Host MIB - CPU Utilization
5. Re Enable the Poller - Go to *Settings->Poller* and check the *Poller Enabled* checkbox. Press *Save*.

Following those steps should complete your migration to the new PHP Script Server for the two example HostMIB Data Queries.

12.3. Migration of Existing PHP Scripts to Script Server

If you have other PHP scripts that you wish to migrate, you must follow the steps below to migrate your scripts to the PHP Script Server required format.

12.3.1. Script File Changes

Each PHP Script file must be changed to the new Script Server format. The changes are not dramatic, but required for the proper operation of the PHP Script Server. Follow the steps below to complete.

1. Copy you existing script to a new name. The name must begin "ss_" followed by your script name. The "ss_" identifies the script as being a script server variety of the a PHP script. For example, if you previously had a script called "get_mysql_stats.php", it's new name would be "ss_get_mysql_stats.php".
2. Edit the new PHP script and add the following required lines to the file, where "ss_myfunction" is the same as your filename.

```
<?php
$no_http_headers = true;

/* display No errors */
error_reporting(E_ERROR);

include_once(dirname(__FILE__) . '/../include/config.php');
include_once(dirname(__FILE__) . '/../lib/snmp.php');

if (!isset($called_by_script_server)) {
    array_shift($_SERVER["argv"]);
    print call_user_func_array("ss_myfunction", $_SERVER["argv"]);
}
```

3. What was originally just mainline code, must be replaced with a function name. For example, if your program previously contained the following three lines of code:

```
<?php
$a = 100;
$b = $a / 10;
print $b;
?>
```

Would become:

```
function ss_myfunction() {
    $a = 100;
    $b = $a / 10;
    Print $b;
}
```

4. If you have any additional functions declared within your script file, you must prefix them to make then unique amongst all functions. Our recommendation would be to prefix all functions with the name of the main function. For example if you have a function called "meme" you would rename it to "ss_myfunction_meme". This guarantee's correct Script Server functionality.
5. The last step is to change the function call that could have traditionally returned the value to the Cacti poller using the PRINT function. You must change that line or lines in your code to utilize the RETURN function instead. However, this does not apply to PRINT statements that are not called from the Poller.

12.3.2. XML File Changes

If you are using a "Script Query" type function, then you must also change your XML file. Please reference the XML files in the <path_cacti>/resource/script_server directory for the specifics related to your required modifications. However, you may also follow the instructions below:

1. Modify the <script_path> tag. Change it from:

```
<script_path>|path_php_binary| -q |path_cacti|/scripts/myfucntion.php</script_path>
```

to simply the following:

```
<script_path>|path_cacti|/scripts/ss_myfunction.php</script_path>
```

2. Add the following two XML tags below the <script_path> tag. Replace *ss_myfunction* with your function name:

```
<script_function>ss_myfunction</script_function>
<script_server>php</script_server>
```

3. Save the XML file.

12.3.3. Data Query & Data Template Changes

Your Data Queries and Data Templates must be also changed. Although somewhat self explanatory by now, you must make the following changes:

1. Change it's Input Method to "Get Script Server Data" or "Get Script Server Data (Index)" depending on it's type.

Change the XML file path to point to the new XML file in the `<path_cacti>/resources/script_server/*.xml` path.

For all data templates that use the data query you must change their "Data Input Method" accordingly.

Your final step is to go to the *System Utilities* and *Clear Poller Cache* to apply the new settings. If you script is operating correctly, you should now be migrated to the script server.

12.4. Testing Your Script in the Script Server

To test your script in the script server, simply follow the instructions below. When you have finished you testing, simply type "quit" <cr> at the Script Server command line to exit the script server.

1. Start the script server - You can do this by typing the following command:

```
shell> php <path_cacti>/script_server.php
```

NOTE: Due to a bug in Windows implementation of PHP, you must type the full path name to the `script_server.php` file.

2. Type in your command - Using the example from above, you would type in the following:

```
script server> <path_myfunction> my_function argument1 argument2 ...
```

In the Windows environment, your example could be the following:

```
script server> c:\wwwroot\cacti\scripts\ss_myfunction.php ss_myfunction localhost public 1 get duddle
```

3. If your function is operating properly, you should get a result.
4. To quit the script server, simply type "quit" <cr> at the command line.

NOTE: If there are errors in your script, you must restart the script server before your retest your code.

Chapter 13. How To

13.1. Graph a Single SNMP OID

When dealing with SNMP-enabled devices, there are often times when you want to graph the value of a single OID. This tutorial explains how to do this in Cacti. It also assumes that you have the "SNMP - Generic OID Template" graph template, which is now included in Cacti as of version 0.8.5. If you do not see this template listed under *Graph Templates*, download the template from the Cacti website in XML form and import it using the *Import Templates* menu item.

To start the process of creating a new graph for your OID, click the *New Graphs* menu item and select the host that contains the target OID from the dropdown. Under the *Graph Templates* box, you will see a dropdown on the last line that reads (*Select a graph type to create*). From this dropdown, choose "SNMP - Generic OID Template" and click the *Create* button at the bottom of the page.

You will be presented with several fields that will require input before the new graph can be created. They are described in more detail below.

Table 13-1. Field Description: SNMP - Generic OID Template

| Name | Description |
|---|--|
| (Graph) Title | The title that is to be used for the new graph. It is generally a good idea to keep host_description in the title, as to make the graph easier to identify later. |
| (Graph) Vertical Label | The text that will be printed along the y-axis of the graph. It is generally used to describe units, such as 'bytes' or 'percent'. |
| (Graph Items) Legend Color | The color that will be used to represent the data on the graph. |
| (Graph Items) Legend Text | The text that will be used to describe the data on the graph legend. |
| (Data Source) Name | The title that is to be used for the new data source. It is generally a good idea to keep host_description in the title, as to make the data source easier to identify later. |
| (Data Source) Maximum Value [snmp_oid] | The maximum value that will be accepted from the OID. Make sure you choose a value that is reasonable for the data you are trying to graph because anything larger than the maximum will be ignored. If you are graphing a percentage, you should use '100' as the value should never exceed this. |
| (Data Source) Data Source Type [snmp_oid] | How the data from the OID should be stored by RRDTool and interpreted on the graph. If the value of the OID represents the actual data, you should use <i>GAUGE</i> for this field. If the OID value is a constantly incrementing number, you should use <i>COUNTER</i> for this field. The two remaining field values, <i>DERIVE</i> and <i>ABSOLUTE</i> can be ignored in most situations. |
| (Custom Data) OID | The actual SNMP OID to graph. It is typically a good idea to enter the number OID here as opposed to using MIB names. For instance, to get the number of open files on a Netware server, you would use ".1.3.6.1.4.1.23.2.28.2.7.0" as the OID. |

When finished filling in values for these fields, click the *Create* button. Your new graph can now be accessed through the *Graph Management* page or the *Graphs* tab inside of Cacti.

13.2. Simplest Method of Going from Script to Graph (Walkthrough)

Written by Kevin der Kinderen, <http://kdeuja.com/~kevin/>

This HOWTO walks you through the simplest steps of graphing the output of a single value from a script. As a new user of cacti, I had a difficult time understanding how to graph anything that wasn't canned with the original load. After a lot of playing around, I came up with these procedures which can be built upon for more sophisticated collections and graphs.

I do not use templates in this HOWTO. Templates provide a significant advantage if you are graphing the output for multiple instances or creating graphs for multiple hosts. They help to maintain consistency and simplify setup.

The example here is not realistic, but can be used and expanded upon as a model for creating your own graphs.

I've not put a lot of explanations in the procedures. Refer to the Cacti manual for more details.

Have a tested script ready to go. I used the following script located in `/home/cactiuser/bin` called `random2.pl`:

```
#!/usr/bin/perl -w

print int(rand(10));
```

This script simply prints out a random integer between 0 and 9 every time it's called.

The minimal steps to create a graph from a script are simply:

1. Create a Data Input Method to tell Cacti how to call the script and what to expect from it.
2. Create a Data Source to tell cacti how and where the data will be stored.
3. Create a Graph to tell cacti how the data will be presented in graph form.
4. Add Graph to Graph View so you can view the graph.
5. View the Graph

The details are below:

1. Create Data Input Method

- Click the *Console* tab at the top
- Click *Data Input Methods* under *Management*
- Click *Add*
- Name: Random 2 Input
- Input Type: Script/Command
- Input String: `/home/cactiuser/bin/random2.pl`
- Click *Create*
- Should see "Save Successful." at the top of the page.
- Click *Add* by *Output Fields* (there are no input fields for this example)
- Field [Output]: `random_number`
- Friendly Name: Random Number
- Update RRD File: checked
- Click *Create*
- Should see "Save Successful." and `random_number` listed under *Output Fields*.
- Click *Save*
- Should see "Save Successful." and Random Input listed in *Data Input Methods*.

2. Create a Data Source

- Click *Data Sources* under *Management*
- Click *Add* in the top right
- The *Data Template* Section should be None and None, we're not using templates for this example.
- Click *Create*.
- Name: `random2ds`
- Data Source Path: blank (Cacti will fill this in)
- Data Input Source: Random 2 Input (this is the data input method you created in step 1)
- Highlight each of the Associated RRA's using control-click
- Step: 300 (300 seconds = 5 minutes)
- Data Source Active: checked
- Describe the Data Source Item inside the RRA by...
- Internal Data Source Name: `random_number`
- Minimum Value: 0

- Maximum Value: 0
- Data Source Type: GAUGE
- Heartbeat: 600
- Click *Create*
- Should see "Save Successful." at the top and the Data Source Path should now have a value (my example <path_rra>/random_number_286.rrd)
- Click Turn On Data Source Debugging Mode to see the results of this step.
- Click *Save*
- Should see "Save Successful" and your new Data Source listed.

3. Create Graph

- Select *Graph Management* under *Management*
- Click *Add*
- Selected Graph Template: None
- Host: None
- Click *Create*
- Title: RANDOM NUMBERS
- Image Format: PNG
- Height: 120
- Width: 500
- Auto Scale: checked
- Auto Scale Options: Use --alt-autoscale-max
- Logarithmic Auto Scaling: unchecked
- Rigid Boundaries: unchecked
- Auto Padding: checked
- Allow Graph Export: checked
- Upper Limit: 100
- Lower Limit: 0
- Base Value: 1000
- Unit Value: blank
- Unit Exponent Value: 0
- Vertical Label: Random Number
- Click *Create*
- Should see "Save Successful."
- If you click "Turn On Graph Debug Mode." now you will see: "Error: can't make a graph without contents."
We need to add Graph Items:
- Click *Add* by *Graph Items*
- Data Source: (from list) No Host - random2ds (random_number)
- Color: 0000FF (Blue)
- Graph Item Type: LINE2
- Consolidation Function: AVERAGE
- CDEF Function: None
- Value: blank
- GPRINT Type: Normal
- Text Format: blank
- Insert Hard Return: unchecked

- Click *Create*
- Should see "Save Successful.", Item #1 listed and a graph under debug (probably with nothing in it yet)
- To add a legend, click *Add* by *Graph Items* again
- Data Source: No Host - random2ds (random_number)
- Color: None
- Graph Item Type: GPRINT
- Consolidation Function: LAST
- CDEF Function: None
- Value: blank
- GPRINT Type: Normal
- Text Format: Cur:
- Insert Hard Return: unchecked
- Click *Create*
- Should see "Save Successful." and the graph will have a legend showing the current value. Note: your integer random number has been averaged over the past 5 minutes. At the bottom of the page, click save.
- At the bottom of the page, click *Save*.
- Should see "Save Successful." and your graph listed.

4. Add graph to the graph view

- Click *Graph Trees* under *Management*
- We'll create a tree called "test" to place our graph
- Click *Add*
- Name: test
- Should see "Save Successful"
- Click *Add* beside *Tree Items* (to add our graph to this tree)
- Under Tree Items [graph]...
- Graph: RANDOM NUMBERS (we named this in step 3)
- Round Robin Archive: Daily (5 Minute Average)
- Click *Create*
- Should see "Save Successful." and RANDOM NUMBERS listed under *Tree Items*
- Click *Save*
- Should see "Save Successful." and test listed under *Graph Trees*

5. View Graph

- Select *Graphs* tab
- Select tree view (Tree beside settings tab)
- Select test tree on left
- You may see "Random Graph" but no graph. It takes two or three polls (10 - 15 minutes) to see a graph. I believe poll 1 to create the rrd, poll 2 to get the first data point and poll 3 to have graphable points.
- While waiting, you can click the RANDOM NUMBERS graph. You'll see place holders for 4 graphs. Select [source] under Daily (5 Minute Average). You'll see the source for the call to rrdtool graph. Give it a sanity check.
- Now be patient. Hit your refresh button every few minutes. After the first poll I got a few empty graphs. A few minutes later data started showing up.

Notes

1. <http://kdeuja.com/~kevin/>

Chapter 14. Frequently Asked Questions

1. Using Cacti

How do I create traffic graphs?

Before you can create traffic graphs, Cacti must have working SNMP support. One way to do this is to compile PHP with SNMP support built in, many times you can install the 'php-snmp' package to do this. The second option is to have Cacti call the `snmpget` and `snmpwalk` binaries from your `ucd-snmp` or `net-snmp` installation. You can use the about page to determine whether to are using "built-in" or "external" SNMP support.

Now that you have working SNMP support, follow these steps.

1. Click *Devices* on Cacti's menu.
2. Click *Add* to add a new device.
3. Type a description, hostname, management IP, and SNMP community. Make sure to select "Generic SNMP-enabled Host" under *Host Template* and click *Create*.
4. Click *Create Graphs for this Host* at the top of the page.
5. You should see one or more interfaces listed under the *Data Query [SNMP - Interface Statistics]* box. Place a check next to each interface that you want to graph. Below the box, select type of graph that you want to create (bits, bytes, summation, etc).
6. Click *Create*, and *Create* again on the following screen to create your graphs.

2. General

I get a "Undefined variable: _SERVER" error message from Cacti.

Cacti 0.8.6 and above requires that you have at least PHP 4.1 or greater installed.

I get a "Call to undefined function: mysql_connect()" error message from Cacti.

Your installation of PHP does not have MySQL installed or enabled. On binary-based distributions, make sure you have the 'php-mysql' package installed. Also make sure that 'extension=mysql.so' is uncommented in your `php.ini` file.

I have forgotten my 'admin' password to Cacti, how do I reset it?

To reset the admin account password back to the default of 'admin', connect to your Cacti database at the command line.

```
shell> mysql -u root -p cacti
```

Now execute the following SQL:

```
mysql> update user_auth set password=md5('admin') where username='admin';
```

3. Monitoring

I am polling thousands of items and poller.php takes more than 5 minutes to run.

Give `cactid`, the future replacement for `poller.php` a try. Unlike `poller.php`, `cactid` is written in c and makes use of `pthreads`. On a typical installation, it is not uncommon for `cactid` to poll about 500 items in less then 10 seconds.

I changed x, and now some of my graphs are not updating.

The best thing to do here is to force Cacti to rebuild its poller cache. To do this click *Utilities* on the Cacti menu, and select *Clear Poller Cache*.

I am using Redhat 8.0 and SNMP is not working.

Redhat 8.0 comes with a broken php-snmp package. Updating your 'php-snmp' and 'net-snmp' packages to their latest versions should fix this problem.

How do I configure net-snmp so it will work with Cacti?

Run `snmpconf -g basic_setup` and follow the prompts. Also check that the file being used by snmpd is the correct config file. For Redhat Linux, snmpconf creates the `/etc/snmpd.conf`, but snmpd uses `/etc/snmpd/snmpd.conf`.

If this doesn't work a very simple config file is:

```
# contact
syslocation Something
syscontact "root@someone.com"

# auth
rocommunity public

# disk monitoring
disk /
```

4. Graphs

I just installed Cacti and all of my graphs appear as broken images.

For you to actually get graph images, poller.php must run at least once so it can create .rrd files in Cacti's 'rra/' directory. Double check that you configured your `/etc/crontab` file to execute poller.php every five minutes. Also make sure that the user poller.php runs as has permission to create new files in Cacti's 'rra/' directory.

If all of your settings appear correct, try running poller.php manually by cd'ing to Cacti's directory and typing:

```
shell> php poller.php
```

If you have files in your 'rra/' directory, but your graphs still appear as broken images, you should enable graph debug mode see exactly why the graphs are not rendering. To do this go into Cacti, select *Graph Management* from the menu, select any graph listed, and select *Turn On Graph Debug Mode*.

My graphs render, but they do not contain any data.

More often than not, this problem's cause is permissions. For instance if you run poller.php manually as root for testing, any .rrd files that it creates will be owned by root. Now when poller.php runs from cron, it will not be able to update these .rrd files because they are owned by root. You can double check the owner of your .rrd files by running the following command in your Cacti directory:

```
shell> ls -al rra/
```

If only some of your graphs are not updating correctly, double check the *Maximum Value* field for all data sources used by these graphs. If the value being fed to the .rrd file exceeds its *Maximum Value*, RRDTool will insert an *Unknown* and you will see no data on the graph.

A lot of my graphs contain long ifAlias names, but they are being truncated to 15 characters.

Cacti does this by default to better control graph title formatting. You can change this limit, by going to *Cacti Settings* on the Cacti menu, selecting the *Visual* tab, and changing the value for *Data Queries - Maximum Field Length*.

One of my devices rebooted and now I have a huge spike on my graph!

This occurs because the reboot causes SNMP's counters to reset, which can cause a rather large spike on the graph when RRDTool tries to determine the change between the new small counter value and the large previous value. One way to combat this issue is to specify realistic maximum values for your data sources. RRDTool will ignore any value that is larger than the maximum value.

If you already have a spike on one or more of your graphs, there is a really useful Perl script¹ that will remove them for you.

RRDTool Says: ERROR: unknown option '--slope-mode' or RRDTool Says: ERROR: Garbage ':39:24 To 2005/10/22 16:39:24\c' after command: COMMENT:From 2005/10/21 16:39:24 To 2005/10/22 16:39:24\c

This occurs because the version of RRDTool that you are running does not match the RRDTool version Cacti is configured to use. Double check your Cacti Settings and make sure that the RRDTool version matches what version of RRDTool you are running.

5. Windows Related

I get a "Fatal error: Cannot redeclare title_trim()" error message from Cacti.

Contrary to previous versions of Cacti, you must not put Cacti in your PHP include path. Check your `php.ini` file to make sure there is not something like `include_path = C:/Apache2/htdocs/cacti`.

Notes

1. <http://cricket.sourceforge.net/contrib/files/killspike2>

Chapter 15. Variables

15.1. Graph Variables

The following variables can be used in the *Text Format* and *Value* graph item fields. Below is a description of each of these variables.

15.1.1. Date/Time

|date_time|

This variable will place the date and time of the last poller run on the graph. It can be used to replicate MRTG's "graph last updated" feature.

15.1.2. Data Query Fields

|query_field_name|

You can place the value of any data query field on the by including this variable. Make sure to substitute "field_name" with the actual data query field name.

Example 15-1. Print the IP Address on a traffic graph

|query_ifIP|

Example 15-2. Print the mount point on net-snmp disk graph

|query_dskPath|

15.1.3. Nth Percentile

|[0-9]:(bits|bytes):[0-9]:(current|total|max|total_peak|all_max_current|all_max_peak|aggregate_max|aggregate_sum|agg

Nth percentile is often used by hosting providers to bill customers based on their peak traffic usage while ignoring their top (100 - Nth) percent. This way if a customer has a fairly consistent traffic pattern and decides to download a huge file one day, the large spike will be ignored. Common Nth percentile is 95, which would cut off the top 5% of the traffic.

In Cacti, Nth percentile works just like any other graph variable. To use this variable you must give it five arguments:

Table 15-1. Nth Percentile Argument Description

| Name | Description |
|--------------|---|
| [0-9] | The first argument is a number between 1 and 99. This number represents the Nth percentile that you would like to calculate. |
| (bits bytes) | Choose whether you want to represent your Nth percentile in bits or bytes. The only valid values for this field are 'bits' and 'bytes'. |
| [0-9] | This field determines the power of 10 divisor that will be applied to your Nth percentile number. Enter '0' to leave the number alone, '3' for kilo or '6' for mega, etc. |

| Name | Description |
|---|--|
| (current total max total_peak all_max_current all_max_peak aggregate_max aggregate_sum aggregate_current aggregate) | Nth Percentile Type - Refer to Nth Types |
| [0-9] | Enter the number of digits to use for floating point precision when printing out the final number. The default value for this field is two decimal places. |

Table 15-2. Nth Percentile Types

| Type | Description |
|-----------------|--|
| current | Calculates the Nth percentile based off the selected data source on the graph item which the variable is used. This type requires a selected Data Source Item in the graph item in which it is defined. |
| total | Calculates the Nth percentile based off the selected data source on the graph item which the variable is used. But unlike current, this function totals the Nth percentile results of all same named data source names on the graph. Example, all traffic_in would have their Nth percentile calculated and then totaled. This type requires a selected Data Source Item in the graph item in which it is defined. |
| max | Calculates the Nth percentile based off the selected data source on the graph item which the variable is used, but it is important to note that it selects the higher of the values for each row in data source and then uses the resulting set to calculate the Nth percentile. Example, you have selected a interface data source for traffic_in, this will evaluate traffic_out and traffic_in for that selected data source (rrdtool file) and select the higher of the two values for each row of data. The resulting max values are used to calculate the Nth percentile value. This type requires a selected Data Source Item in the graph item in which it is defined. |
| total_peak | Calculates the Nth percentile based on the max for each data source defined on the graph. Example, the max is taken for each data source defined on the graph, but it's the max of traffic_in or traffic_out for that data source (rrdtool file). The resulting max values are summed and returned. |
| all_max_current | Calculates the Nth percentile based off the selected data source on the graph item which the variable is used, but it only selects the highest (max) value of all the data source items on the graph. Example, there are 3 data sources defined on the graph, the selected data source is traffic_in, all data sources for traffic_in will be compared against each other and the highest Nth value of all graph items will be selected as the result. This type requires a selected Data Source Item in the graph item in which it is defined. |
| all_max_peak | Calculates the Nth percentile by selecting the highest (max) calculated max Nth percentile value for each data source (rrdtool file) selected on the graph. Example, 3 data sources are defined on the graph, for each data source, the max Nth percentile is calculated, and out of the resulting set of values, the highest value is selected and returned. |
| aggregate_max | Calculates the Nth percentile by selecting the highest value for each summed value of like data sources and selecting the maximum value of that set to calculate the Nth percentile value. Example, you have a graph with 5 traffic_in and 18 traffic_out data sources. The traffic_in rows are summed together, then the traffic_out rows are summed together, then for each row, the higher of the 2 values is selected. The Nth percentile is calculated from the resulting maximum values. |
| aggregate_sum | Calculates the Nth percentile by summing all data sources on the graph row for row and calculates the Nth value on the resulting summed data. Example, you have a graph with 4 traffic_in and 3 traffic_out data sources defined. All the defined traffic_in and traffic_out are summed together and then the Nth percentile value is calculated from that set. |

| Type | Description |
|-------------------|--|
| aggregate_current | Calculates the Nth percentile by summing items matching the selected data source name row for row from each data source defined on the graph with the same data source name and then returning the Nth percent value calculated. Example, you have a graph with 3 traffic_out and 5 traffic_in data sources defined. The traffic_in is selected on the graph item where this variable is used, so all graph items for traffic_in are summed and then the Nth value is calculated and returned as the variable. |
| aggregate | Calculates the Nth percentile by summing like data source names row for row from each data source defined on the graph and then returning the highest Nth percent value calculated for like data source names. Example, you have a graph with 3 traffic_out and 5 traffic_in data sources defined. The traffic_in and traffic_out are summed then the Nth value is calculated and the higher of the 2 values is returned as the variable. |

Concerning aggregate Nth percentile functions and rra definitions: All data sources used in a graph MUST have the same rra definitions. If the number of rows do not match when calculations are attempted, the results may produce errors and invalid results. Please also be aware the aggregate functions take more memory and processor power when used. Recommended minimum of 64 megs memory limit for php when querying large ranges or graphs with a large ammount of aggregated data sources.

Example 15-3. Output in the following format, for 95th percentile: 42.58 mbit

```
| 95:bits:6:max:2 |
```

15.1.4. Bandwidth Summation

```
|sum: ([0-9] |auto) : (current|total) : ([0-9]) : ([0-9]+|auto) |
```

Bandwidth summation is useful for summing up all values in an RRD file for a given time range. This is typically useful on traffic graphs where you can see a total of all traffic that has gone through an interface in a given time period.

In Cacti, bandwidth summation works just like any other graph variable. To use this variable you must give it three arguments:

Table 15-3. Bandwidth Summation Argument Description

| Name | Description |
|-----------------|--|
| ([0-9] auto) | This field determines the power of 10 divisor that will be applied to your number. Enter '0' to leave the number alone, '3' for kilo or '6' for mega, etc. You can also enter 'auto' for this field to have Cacti automatically scale the number and insert the appropriate label. |
| (current total) | You can choose to calculate the summation based on the current data source or a total of all data sources used on the graph. The only valid values for this field are 'current' and 'total'. |
| [0-9] | Enter the number of digits to use for floating point precision when printing out the final number. The default value for this field is two decimal places. |
| ([0-9]+ auto) | Enter the number of seconds in the past to perform the summation calculation for. For instance '86400' for 24 hours, '172800' for 48 hours, etc. You can also enter 'auto' for this field to have Cacti use the graph timespan. |

Example 15-4. Output in the following format: 36.47 GB

```
|sum:auto:current:2:auto|
```

15.2. Host Variables

Host variables represent host data and can be placed in graph or data source titles. The following table describes which host variables can be used and which host fields they correspond to.

Table 15-4. Host Variables

| Variable Name | Description |
|---------------------|--------------------|
| host_description | Description |
| host_hostname | Hostname |
| host_snmp_community | SNMP Community |
| host_snmp_version | SNMP Version |
| host_snmp_username | SNMP Username (v3) |
| host_snmp_password | SNMP Password (v3) |

Chapter 16. RRDTool Specific Features

16.1. GPRINT Presets

A GPRINT is a graph item type that enables you to print the values of data sources on a graph. They are typically used to represent legend values on the graph. The output format of these numbers are controlled by a printf-like format string. Cacti enables you to keep a global list of these strings that can be applied to any graph item throughout Cacti.

16.1.1. Creating a GPRINT Preset

To create a new GPRINT preset, select the *Graph Management* menu item under the *Management* heading, and select *GPRINT Presets*. Click *Add* to the right and you will be presented with an edit page containing two fields. Enter a name for your GPRINT preset, and the actual printf-like string in the *GPRINT Text* field. When you are finished, click the *Create* button to create your new GPRINT preset.

16.2. CDEFs

CDEFs allow you to apply mathematical functions to graph data to alter output. The concept of a CDEF comes straight from RRDTool, and are written in reverse polish notation (RPN). For more information regarding the syntax of CDEFs, check out the CDEF tutorial¹.

16.2.1. Creating a CDEF

To create a new CDEF in Cacti, select the *Graph Management* option under the *Management* heading, and select *CDEFs*. Once at this screen, click *Add* to the right. You will be prompted for a CDEF name, for which you can type anything used to describe your CDEF. Click the *Create* button so you are redirected back to the edit page, now with an empty *CDEF Items* box. Construct your CDEF by adding an item for each element in the CDEF string, common types such as operators and functions are enumerated for your convenience. Below is a basic description of each CDEF item type.

Table 16-1. CDEF Item Types

| Type | Description |
|---------------------|--|
| Function | You can choose a CDEF function to use as the item. The RRDTool graph manual ² describes the purpose of each CDEF function. |
| Operator | Just your standard math operators, including modulo (%). |
| Special Data Source | A special data source is basically a flag to tell Cacti to do some special processing when it encounters this CDEF item. The "Current Graph Item Data Source" type basically inserts the name of the data source that is referenced by the graph item that references to this CDEF. Both of the "All Data Sources" types insert a summation of all data sources used on a graph. |
| Another CDEF | You can recursively use another CDEF within this CDEF. |
| Custom String | Sometimes it's just easier to type out the literal CDEF string manually. When referencing to data sources on the graph, remember that Cacti names them 'a', 'b', 'c', '...', starting with the first data source on the graph. |

Notes

1. http://people.ee.ethz.ch/~oetiker/webtools/rrdtool/doc/rrdgraph_data.en.html