

---

# **snapintime**

***Release 1.1.1***

**Eric Mesa**

**Apr 01, 2020**



**CONTENTS:**

<b>1</b>	<b>Usage</b>	<b>3</b>
1.1	Creating Local Snapshots . . . . .	3
1.2	Backing Up to Remote Location . . . . .	3
1.3	Culling Local Snapshots . . . . .	4
1.4	Putting it All Together . . . . .	4
<b>2</b>	<b>config.json</b>	<b>7</b>
<b>3</b>	<b>Origins of my culling algorithm</b>	<b>9</b>
<b>4</b>	<b>create_local_snapshots</b>	<b>11</b>
<b>5</b>	<b>Culling</b>	<b>13</b>
<b>6</b>	<b>remote_backup</b>	<b>17</b>
<b>7</b>	<b>config</b>	<b>19</b>
<b>8</b>	<b>date</b>	<b>21</b>
<b>9</b>	<b>Indices and tables</b>	<b>23</b>
	<b>Python Module Index</b>	<b>25</b>
	<b>Index</b>	<b>27</b>



snapintime is meant to manage the creation, culling, and send to a remote location of btrfs snapshots.

At this point in time it creates snapshots, culls local snapshots three days old, and can btrfs send/receive to a remote btrfs subvol.



## USAGE

Grab config.json from the Github repo (<https://github.com/djotaku/Snap-in-Time>), edit it, and place it in \$HOME/.config/snapintime (or /root/.config/snapintime/ if you're going to run as root)

### 1.1 Creating Local Snapshots

If running from a git clone:

```
pip -r requirements.txt
cd snapintime
python create_local_snapshots.py
```

If running from PyPi, run: `python -m snapintime.create_local_snapshots`

If you want to run it from cron in a virtual environment, you can adapt the following shell script to your situation:

```
#!/bin/bash
cd "/home/ermesa/Programming Projects/python/cronpip"
source ./bin/activate
python -m snapintime.create_local_snapshots
```

Make it executable and have cron run that script as often as you like.

For a more involved script, useful for logging, see *Putting it All Together*.

### 1.2 Backing Up to Remote Location

This code makes the assumption that you have setup ssh keys to allow you to ssh to the remote machine without inputting a password. It is recommended to run the remote backup code BEFORE the culling code to increase the chances that the last snapshot on the remote system is still on the local system. (This will minimize the amount of data that has to be transferred to the remote system.)

```
pip -r requirements.txt
cd snapintime
python remote_backup.py
```

If running from PyPi, run: `python -m snapintime.remote_backup`

## 1.3 Culling Local Snapshots

The culling follows the following specification:

- Three days ago: Leave at most 4 snapshots behind - closest snapshots to 0000, 0600, 1200, and 1800. (implemented)
- Seven days ago: Leave at most 1 snapshot behind - the last one that day. In a perfect situation, it would be the one taken at 1800. (implemented)
- 90 days ago: Go from that date up another 90 days and leave at most 1 snapshot per week. (implemented)
- 365 days ago: Go from that date up another 365 days and leave at most 1 snapshot per quarter (implemented)

(Not going to care about leap years, eventually it'll fix itself if this is run regularly)

I recommend running culling submodule AFTER remote backup (if you're doing the remote backups). This is to prevent the removal of the subvol you'd use for the btrfs send/receive. If your computer is constantly on without interruption, it shouldn't be an issue if you're doing your remote backups daily. And why wouldn't you? The smaller the diff between the last backup and this one, the less data you have to send over the network. So it's more of a precaution in case you turn it off for a while on vacation or the computer breaks for a while and can't do the backups.

```
pip -r requirements.txt
cd snapintime
python culling.py
```

If running from PyPi, run: `python -m snapintime.culling`

## 1.4 Putting it All Together

Here is my crontab output:

```
0 * * * * /root/bin/snapshots.sh
@daily /root/bin/remote_snapshots.sh
0 4 * * * /root/bin/snapshot_culling.sh
```

`remote_snapshots.sh`:

```
#!/bin/bash

cd "/home/ermesa/Programming Projects/python/cronpip"
source ./bin/activate
echo "#####" >> snapintime_remote.log
echo "Starting remote backups" >> snapintime_remote.log
python -m snapintime.remote_backup >> snapintime_remote.log
echo "#####" >> snapintime_remote.log
#!/bin/bash
```

`snapshot_culling.sh`:

```
#!/bin/bash

cd "/home/ermesa/Programming Projects/python/cronpip"
source ./bin/activate
echo "#####" >> snapintime_culling.log
echo "Starting culling" >> snapintime_culling.log
```

(continues on next page)



(continued from previous page)

```
python -m snapintime.culling >> snapintime_culling.log
echo "#####" >> snapintime_culling.log
```

snapshots.sh:

```
#!/bin/bash

cd "/home/ermesa/Programming Projects/python/cronpip"
source ./bin/activate
echo "#####" >> snapintime.log
echo "Starting snapshots" >> snapintime.log
python -m snapintime.create_local_snapshots >> snapintime.log
echo "#####" >> snapintime.log
```



## CONFIG.JSON

An example of the config.json file:

```
{  "0":
  { "subvol": "/home",
    "backuplocation": "/home/.snapshot",
    "remote": "True",
    "remote_location": "user@server",
    "remote_subvol_dir": "/media/backups"
  },
  "1":
  { "subvol": "/media/Photos",
    "backuplocation": "/media/Photos/.Snapshots"
  },
  "2":
  { "subvol": "/media/Archive",
    "backuplocation": "/media/NotHome/Snapshots/Archive"
  }
}
```

- For the 0, 1, 2, 3, etc - there is currently (as of 0.8.1) not any inherent meaning to the fact that they are numbers. They just need to be distinct alpha-numeric sequences.
- subvol: this should be the subvolume you want to create a snapshot of.
- backuplocation: the subvolume that holds your backup subvolumes.
- remote: If set to True, an attempt will be made to backup to the remote location. Any other value or lack of this field means it will not try and backup to the remote location.
- remote\_location: The `username@theserver` where the backup subvolumes will be sent to.
- remote\_subvol\_dir: Just like backuplocation, but on the remote machine.



## ORIGINS OF MY CULLING ALGORITHM

I'm basing it on a conversation I had in the btrfs mailing list. Here's how the guy who inspired me, Duncan, explained it to me:

"However, best snapshot management practice does progressive snapshot thinning, so you never have more than a few hundred snapshots to manage at once. Think of it this way. If you realize you deleted something you needed yesterday, you might well remember about when you deleted it and can thus pick the correct snapshot to mount and copy it back from. But if you don't realize you need it until a year later, say when you're doing your taxes, how likely are you to remember the specific hour, or even the specific day, you deleted it? A year later, getting a copy from the correct week, or perhaps the correct month, will probably suffice, and even if you DID still have every single hour's snapshots a year later, how would you ever know which one to pick? So while a day out, hourly snapshots are nice, a year out, they're just noise.

As a result, a typical automated snapshot thinning script, working with snapshots each hour to begin with, might look like this:

Keep two days of hourly snapshots: 48 hourly snapshots

After two days, delete five of six snapshots, leaving a snapshot every 6 hours, four snapshots a day, for another 5 days:  $4 \times 5 = 20$  6-hourly,  $20 + 48 = 68$  total.

After a week, delete three of the four 6-hour snapshots, leaving daily snapshots, for 12 weeks (plus the week of more frequent snapshots above, 13 weeks total):  $7 \times 12 = 84$  daily snaps,  $68 + 84 = 152$  total.

After a quarter (13 weeks), delete six of seven daily snapshots, leaving weekly snapshots, for 3 more quarters plus the one above of more frequent snapshots, totaling a year:  $3 \times 13 = 39$  weekly snaps,  $152 + 39 = 191$  total.

After a year, delete 12 of the 13 weekly snapshots, leaving one a quarter. At 191 for the latest year plus one a quarter you can have several years worth of snapshots (well beyond the normal life of the storage media) and still be in the low 200s snapshots total, while keeping them reasonably easy to manage. =:^^"



## CREATE\_LOCAL\_SNAPSHOTS

Read in configuration file and create local snapshots.

```
snapintime.create_local_snapshots.create_snapshot (date_suffix: str, subvol: str,  
                                                    backup_location: str)
```

Create a btrfs snapshot.

### Parameters

- **date\_suffix** – a datetime object formatted to be the name of the snapshot
- **subvol** – The subvolume to be snapshotted
- **backup\_location** – The folder in which to create the snapshot

```
snapintime.create_local_snapshots.get_date_time () → str
```

Return the current time, uses system time zone.

```
snapintime.create_local_snapshots.iterate_configs (date_time: str, config: dict) → list
```

Iterate over all the subvolumes in the config file, then call create\_snapshot.

### Parameters

- **date\_time** – The date time that will end up as the btrfs snapshot name
- **config** – The config file, parsed by import\_config.

**Returns** A list containing return values from create\_snapshot

```
snapintime.create_local_snapshots.main ()
```





## CULLING

Thin out the snapshots on disk.

`snapintime.culling.btrfs_del` (*directory: str, subvols: list*) → list  
Delete subvolumes in a given directory.

**Parameters**

- **directory** – The directory holding the subvolumes.
- **subvols** – A list of subvolumes to delete

**Returns** A list with the commands run and the results or, if there weren't any subvolumes to delete, returns a message with that information.

`snapintime.culling.cull_last_quarter` (*config: dict*) → list  
Cull the btrfs snapshots from quarter.

Should leave 1 snapshot per week for 13 weeks.

**Parameters** **config** – The configuration file.

**Returns** A list containing the results of running the commands.

`snapintime.culling.cull_last_year` (*config: dict*) → list  
Cull the btrfs snapshots from quarter.

Should leave 1 snapshot per week for 13 weeks.

**Parameters** **config** – The configuration file.

**Returns** A list containing the results of running the commands.

`snapintime.culling.cull_seven_days_ago` (*config: dict*) → list  
Cull the btrfs snapshots from 7 days ago.

**Parameters** **config** – The configuration file.

**Returns** A list containing the results of running the commands.

`snapintime.culling.cull_three_days_ago` (*config: dict*) → list  
Cull the btrfs snapshots from 3 days ago.

**Parameters** **config** – The configuration file.

**Returns** A list containing the results of running the commands.

`snapintime.culling.daily_cull` (*dir\_to\_cull: list*) → list  
Take a list of snapshots from a directory (already reduced to one day) and cull.

This culling will produce the closest it can to 4 snapshots for that day.

For a perfect set of 24 snapshots, it should leave behind (remove from list):

- day1-0000
- day1-0600
- day1-1200
- day1-1800

**Parameters** **dir\_to\_cull** – A list containing snapshots. Assumes another function has already reduced this list to a list containing only one day’s worth of snapshots.

**Returns** A list containing all the subvolumes to cull.

`snapintime.culling.get_subvols_by_date(directory: str, reg_ex) → list`  
Return a list based on matching regular expression.

This is meant to produce the list that will be the input for one of the culling functions.

**Parameters**

- **directory** – The directory we want to grab subvols from.
- **reg\_ex** – A regular expression to apply to the directory contents.

**Returns** A list of subvolumes for culling.

`snapintime.culling.main()`

`snapintime.culling.print_output(list_of_lists: list)`

`snapintime.culling.quarterly_yearly_cull(dir_to_cull: list) → list`  
Take a list of snapshots from a directory (already reduced to one week or quarter) and cull.

This culling will return a list with the snapshots to remove for the week or quarter.

For a perfect set of snapshots, (where the user has been doing one snapshot per hour and doing the daily culling) it should leave behind (remove from list):

- day7-1800

---

**Note:** May end up combining with weekly cull as they essentially do the same thing.

---

**Parameters** **dir\_to\_cull** – A list containing snapshots. Assumes another function has already reduced this list to a list containing only one week or quarter’s worth of snapshots.

**Returns** A list containing all the subvolumes to cull.

`snapintime.culling.split_dir_hours(subvols: list, reg_ex) → list`  
Return a list based on matching regular expression.

**Parameters**

- **subvols** – A list of subvolumes.
- **reg\_ex** – A re object defining the regular expression to evaluate against.

**Returns** A list that has only the items that passed the regular expression.

`snapintime.culling.weekly_cull(dir_to_cull: list) → list`  
Take a list of snapshots from a directory (already reduced to one day) and cull.

This culling will return a list with the snapshots to remove for the given day.

For a perfect set of snapshots, (where the user has been doing one snapshot per hour and doing the daily culling) it should leave behind (remove from list):

- day1-1800

**Parameters** `dir_to_cull` – A list containing snapshots. Assumes another function has already reduced this list to a list containing only one day's worth of snapshots.

**Returns** A list containing all the subvolumes to cull.



## REMOTE\_BACKUP

Use `btrfs send/receive` to send snapshot to a remote computer.

```
snapintime.remote_backup.btrfs_send_receive(local_subvols: list, remote_subvol: str,
                                              backup_location: str, remote_location: str,
                                              remote_subvol_dir: str)
```

Run command to send/receive `btrfs` snapshot.

### Parameters

- **local\_subvols** – A list of the local subvolumes to choose from.
- **remote\_subvol** – The latest subvolume that is present on both the remote and local systems.
- **backup\_location** – The folder prefix for the local\_subvol
- **remote\_location** – This should be a string like `user@computer` or `user@IPaddress`
- **remote\_subvol\_dir** – This is the directory we will put the backup into on the remote system.

**Returns** A dictionary with the result of the command.

```
snapintime.remote_backup.get_local_subvols(local_subvol_dir: str) → list
```

Grab the subvolumes from the local directory.

**Parameters** **local\_subvol\_dir** – The directory containing the local subvolumes.

**Returns** A list of all the local subvolumes

```
snapintime.remote_backup.get_remote_subvols(remote_location: str, remote_subvol_dir:
                                              str) → list
```

Retrieve the remote subvolumes.

This function assumes user has set up `ssh` keys for passwordless login.

### Parameters

- **remote\_location** – This should be a string like `user@computer` or `user@IPaddress`
- **remote\_subvol\_dir** – This is the directory we will search to get the latest subvolume.

**Returns** A list of the remote subvolumes.

```
snapintime.remote_backup.iterate_configs(config: dict) → list
```

Iterate over all the subvolumes in the config file, then call `btrfs_send_receive` if the value of `remote` is “True”.

**Parameters** **config** – The config file, parsed by `import_config`.

**Returns** A list containing return values from `btrfs_send_receive`

```
snapintime.remote_backup.main()
```

`snapintime.remote_backup.match_subvols` (*local\_subvols: list, remote\_subvols: list*) → str  
Return the latest remote subvol that also exists on the local system.

**Parameters**

- **local\_subvols** – A list of the local subvolumes.
- **remote\_subvols** – A list of the remote subvolumes.

**Returns** The subvolume.

## CONFIG

Load the config file.

`snapintime.utils.config.import_config()` → dict  
Import config file.

**Returns** A dictionary containing configs

**Raises** FileNotFoundError





## DATE

Provide date and Time Operations needed by snapintime.

`snapintime.utils.date.many_dates` (*start\_date: datetime.datetime, interval\_start: int, interval\_end: int*) → list

Provide a list of dates within a certain range.

Used by quarterly culling and yearly culling to determine date range to cull.

### Parameters

- **start\_date** – The reference point for the intervals
- **interval\_start** – How many days ago you want to start getting dates from.
- **interval\_end** – How many days ago you want to stop getting dates from.

`snapintime.utils.date.prior_date` (*start\_date: datetime.datetime, day: int = 0*) → datetime.datetime

Provide a prior date offset by the variable given in day.

Unintuitively, positive numbers subtract days.

### Parameters

- **start\_date** – The date from which you want to count back or forward.
- **day** – The number of days you want to go back.

**Returns** A datetime object day amount of days in the past.

`snapintime.utils.date.quarterly_weeks` (*start\_date: datetime.datetime*) → list

Provide a list of 13 weekly date lists.

**Parameters** **start\_date** – Date from which to go back a quarter.

**Returns** A list of lists containing datetime objects. Each sublist represents a week.

`snapintime.utils.date.yearly_quarters` (*start\_date: datetime.datetime*) → list

Provide a list of 4 quarterly date lists.

**Parameters** **start\_date** – Date from which to go back a year.

**Returns** A list of lists containing datetime objects. Each sublist represents a quarter.



## INDICES AND TABLES

- `genindex`
- `modindex`
- `search`



## PYTHON MODULE INDEX

### S

`snapintime.create_local_snapshots`, [11](#)  
`snapintime.culling`, [13](#)  
`snapintime.remote_backup`, [17](#)  
`snapintime.utils.config`, [19](#)  
`snapintime.utils.date`, [21](#)



## B

btrfs\_del() (in module *snapintime.culling*), 13  
 btrfs\_send\_receive() (in module *snapintime.remote\_backup*), 17

## C

create\_snapshot() (in module *snapintime.create\_local\_snapshots*), 11  
 cull\_last\_quarter() (in module *snapintime.culling*), 13  
 cull\_last\_year() (in module *snapintime.culling*), 13  
 cull\_seven\_days\_ago() (in module *snapintime.culling*), 13  
 cull\_three\_days\_ago() (in module *snapintime.culling*), 13

## D

daily\_cull() (in module *snapintime.culling*), 13

## G

get\_date\_time() (in module *snapintime.create\_local\_snapshots*), 11  
 get\_local\_subvols() (in module *snapintime.remote\_backup*), 17  
 get\_remote\_subvols() (in module *snapintime.remote\_backup*), 17  
 get\_subvols\_by\_date() (in module *snapintime.culling*), 14

## I

import\_config() (in module *snapintime.utils.config*), 19  
 iterate\_configs() (in module *snapintime.create\_local\_snapshots*), 11  
 iterate\_configs() (in module *snapintime.remote\_backup*), 17

## M

main() (in module *snapintime.create\_local\_snapshots*), 11

main() (in module *snapintime.culling*), 14  
 main() (in module *snapintime.remote\_backup*), 17  
 many\_dates() (in module *snapintime.utils.date*), 21  
 match\_subvols() (in module *snapintime.remote\_backup*), 17

## P

print\_output() (in module *snapintime.culling*), 14  
 prior\_date() (in module *snapintime.utils.date*), 21

## Q

quarterly\_weeks() (in module *snapintime.utils.date*), 21  
 quarterly\_yearly\_cull() (in module *snapintime.culling*), 14

## S

*snapintime.create\_local\_snapshots* (module), 11  
*snapintime.culling* (module), 13  
*snapintime.remote\_backup* (module), 17  
*snapintime.utils.config* (module), 19  
*snapintime.utils.date* (module), 21  
 split\_dir\_hours() (in module *snapintime.culling*), 14

## W

weekly\_cull() (in module *snapintime.culling*), 14

## Y

yearly\_quarters() (in module *snapintime.utils.date*), 21