
awslimitchecker Documentation

*Release 0.1.1@0.1.1**

Jason Antman

December 02, 2015

1	Status	3
2	What It Does	5
3	Requirements	7
4	Installation and Usage	9
5	Getting Help and Asking Questions	11
6	Changelog	13
7	Contents	15
7.1	Getting Started	15
7.2	Command Line Usage	17
7.3	Python Usage	21
7.4	Required IAM Permissions	24
7.5	Supported Limits	25
7.6	Development	29
7.7	Internals	32
7.8	awslimitchecker	33
7.9	Getting Help	55
8	Indices and tables	57
8.1	License	57
	Python Module Index	59

A script and python module to check your AWS service limits and usage using `boto`.

Users building out scalable services in Amazon AWS often run into AWS' [service limits](#) - often at the least convenient time (i.e. mid-deploy or when autoscaling fails). Amazon's [Trusted Advisor](#) can help this, but even the version that comes with Business and Enterprise support only monitors a small subset of AWS limits and only alerts *weekly*. `awslimitchecker` provides a command line script and reusable package that queries your current usage of AWS resources and compares it to limits (hard-coded AWS defaults that you can override, or data from Trusted Advisor where available), notifying you when you are approaching or at your limits.

Status

This project is currently in very early development. At this time please consider it alpha code and not reliable; furthermore its API may be changing rapidly. I hope to have this stabilized soon. I wouldn't call it ready for use, but contributions are certainly welcome.

What It Does

- Check current AWS resource usage against AWS Service Limits
- Show and inspect current usage
- Override default Service Limits (for accounts with increased limits)
- Compare current usage to limits; return information about limits that exceed thresholds, and (CLI wrapper) exit non-0 if thresholds are exceeded
- Define custom thresholds per-limit
- where possible, pull current limits from Trusted Advisor API

Requirements

- Python 2.6 through 3.4. Python 2.x is recommended, as `boto` (the AWS client library) currently has incomplete Python3 support. See the [boto documentation](#) for a list of AWS services that are Python3-compatible.
- Python `VirtualEnv` and `pip` (recommended installation method; your OS/distribution should have packages for these)
- `boto` \geq 2.32.0

Installation and Usage

See *Getting Started*.

Getting Help and Asking Questions

See *Getting Help*.

Changelog

See <https://github.com/jantman/awslimitchecker/blob/develop/CHANGES.rst>

7.1 Getting Started

7.1.1 What It Does

- Check current AWS resource usage against AWS Service Limits
- Show and inspect current usage
- Override default Service Limits (for accounts with increased limits)
- Compare current usage to limits; return information about limits that exceed thresholds, and (CLI wrapper) exit non-0 if thresholds are exceeded
- Define custom thresholds per-limit
- Where possible, pull current limits from Trusted Advisor API

7.1.2 Nomenclature

Service An AWS Service or Product, such as EC2, VPC, RDS or ElastiCache. More specifically, Services in `AwsLimitChecker` correspond to distinct APIs for [AWS Services](#).

Limit An AWS-imposed maximum usage for a certain resource type in AWS. See [AWS Service Limits](#). Limits are generally either account-wide or per-region. They have AWS global default values, but can be increased by AWS Support. “Limit” is also the term used within this documentation to describe `AwsLimit` objects, which describe a specific AWS Limit within this program.

Usage “Usage” refers to your current usage of a specific resource that has a limit. Usage values/amounts (some integer or floating point number, such as number of VPCs or GB of IOPS-provisioned storage) are represented by instances of the `AwsLimitUsage` class. Limits that are measured as a subset of some “parent” resource, such as “Subnets per VPC” or “Read Replicas per Master” have their usage tracked per parent resource, so you can easily determine which ones are problematic.

Threshold The point at which `AwsLimitChecker` will consider the current usage for a limit to be problematic. Global thresholds default to usage $\geq 80\%$ of limit for “warning” severity, and usage $\geq 99\%$ of limit for “critical” severity. Limits which have reached or exceeded their threshold will be reported separately for warning and critical (we generally consider “warning” to be something that will require human intervention in the near future, and “critical” something that is an immediate problem, i.e. should block automated processes). The `awslimitchecker` command line wrapper can override the default global thresholds. The `AwsLimitChecker` class can both override global percentage thresholds, as well as specify per-limit thresholds as a percentage, a fixed usage value, or both.

7.1.3 Requirements

- Python 2.6 or 2.7 (`boto` currently has incomplete python3 support)
- Python `VirtualEnv` and `pip` (recommended installation method; your OS/distribution should have packages for these)
- `boto`

7.1.4 Installing

It's recommended that you install into a virtual environment (`virtualenv` / `venv`). See the [virtualenv usage documentation](#) for more details, but the gist is as follows (the `virtualenv` name, “`limitchecker`” here, can be whatever you want):

```
virtualenv limitchecker
source limitchecker/bin/activate
pip install awslimitchecker
```

7.1.5 Credentials

`awslimitchecker` does nothing with AWS credentials, it leaves that to `boto` itself. You must either have your credentials configured in one of `boto`'s supported config files, or set as environment variables. See the [boto configuration documentation](#) for further information.

The recommended way of handling multiple accounts is to use one of the [credential configuration files](#) (`~/.aws/credentials` is recommended, as it should be supported by all AWS SDKs and tools), define a [section per account](#), and then export `AWS_PROFILE=section_name` to tell `boto` which section to use.

7.1.6 Regions

At this time, AWS Limit Checker has no knowledge of AWS regions. As most (all?) limits are calculated on a per-region basis, this isn't a major issue. To check multiple regions, simply run `awslimitchecker` multiple times, once for each region, using a different `AWS_PROFILE` environment variable setting, and entries in `~/.aws/credentials` like:

```
[myuser-us-east-1] region = us-east-1 aws_access_key_id = <your access key> aws_secret_access_key =
<your secret key>
```

```
[myuser-us-west-2] region = us-west-2 aws_access_key_id = <your access key> aws_secret_access_key
= <your secret key>
```

Support for setting the region, or multiple regions, directly through `awslimitchecker` will be implemented in the future depending on demand; it shouldn't be too complicated to retrofit into the existing code.

7.1.7 Required Permissions

You can view a sample IAM policy listing the permissions required for `awslimitchecker` to function properly either via the CLI client:

```
awslimitchecker --iam-policy
```

Or as a python dict:

```

from awslimitchecker.checker import AwsLimitChecker
c = AwsLimitChecker()
iam_policy = c.get_required_iam_policy()

```

You can also view the required permissions for the current version of awslimitchecker at [Required IAM Permissions](#).

7.2 Command Line Usage

awslimitchecker ships with a command line script for use outside of Python environments. awslimitchecker is installed as a [setuptools entry point](#), and will be available wherever you install the package (if you install in a virtual environment as recommended, it will be in the venv's bin/ directory).

The command line script provides simple access to the most common features, though not full access to all configuration possibilities. In addition, when checking usage, the script will exit 0 if everything is OK, 1 if there are warnings, and 2 if there are critical thresholds exceeded (though the output is not currently suitable for direct use as a Nagios-compatible plugin).

```

(venv)$ awslimitchecker --help
usage: awslimitchecker [-h] [-S SERVICE] [-s] [-l] [--list-defaults]
                    [-L LIMIT] [-u] [--iam-policy] [-W WARNING_THRESHOLD]
                    [-C CRITICAL_THRESHOLD] [--skip-ta] [--no-color] [-v]
                    [-V]
Report on AWS service limits and usage via boto, optionally warn about any
services with usage nearing or exceeding their limits. For further help, see
<http://awslimitchecker.readthedocs.org/>
optional arguments:
  -h, --help                show this help message and exit
  -S SERVICE, --service SERVICE
                           perform action for only the specified service name;
                           see -s|--list-services for valid names
  -s, --list-services       print a list of all AWS service types that
                           awslimitchecker knows how to check
  -l, --list-limits         print all AWS effective limits in
                           "service_name/limit_name" format
  --list-defaults          print all AWS default limits in
                           "service_name/limit_name" format
  -L LIMIT, --limit LIMIT
                           override a single AWS limit, specified in
                           "service_name/limit_name=value" format; can be
                           specified multiple times.
  -u, --show-usage         find and print the current usage of all AWS services
                           with known limits
  --iam-policy              output a JSON serialized IAM Policy listing the
                           required permissions for awslimitchecker to run
                           correctly.
  -W WARNING_THRESHOLD, --warning-threshold WARNING_THRESHOLD
                           default warning threshold (percentage of limit);
                           default: 80
  -C CRITICAL_THRESHOLD, --critical-threshold CRITICAL_THRESHOLD
                           default critical threshold (percentage of limit);
                           default: 99
  --skip-ta                do not attempt to pull *any* information on limits
                           from Trusted Advisor
  --no-color               do not colorize output
  -v, --verbose            verbose output. specify twice for debug-level output.
  -V, --version            print version number and exit.

```

```
awslimitchecker is AGPLv3-licensed Free Software. Anyone using this program,
even remotely over a network, is entitled to a copy of the source code. Use
`--version` for information on the source code location.
```

7.2.1 Examples

In the following examples, **output has been truncated** to simplify documentation. When running with all services enabled, awslimitchecker will provide *many* lines of output. (...) has been inserted in the output below to denote removed or truncated lines.

Listing Supported Services

View the AWS services currently supported by awslimitchecker with the `-s` or `--list-services` option.

```
(venv) $ awslimitchecker -s
AutoScaling
EBS
EC2
ELB
ElastiCache
RDS
VPC
```

Listing Default Limits

To show the hard-coded default limits, ignoring any limit overrides or Trusted Advisor data, run with `--list-defaults`:

```
(venv) $ awslimitchecker --list-defaults
AutoScaling/Auto Scaling groups           20
AutoScaling/Launch configurations        100
EBS/Active snapshots                     10000
EBS/Active volumes                       5000
EBS/General Purpose (SSD) volume storage (GiB) 20
(...)
VPC/Rules per network ACL                 20
VPC/Subnets per VPC                     200
VPC/VPCs                                  5
```

Viewing Limits

View the limits that awslimitchecker currently knows how to check, and what the limit value is set as (if you specify limit overrides, they will be used instead of the default limit) by specifying the `-l` or `--list-limits` option. Limits followed by (TA) have been obtained from Trusted Advisor.

```
(venv) $ awslimitchecker -l
AutoScaling/Auto Scaling groups           100 (TA)
AutoScaling/Launch configurations        200 (TA)
EBS/Active snapshots                     10000 (TA)
EBS/Active volumes                       5000 (TA)
EBS/General Purpose (SSD) volume storage (GiB) 30720 (TA)
(...)
VPC/Rules per network ACL                 20
```

VPC/Subnets per VPC	200
VPC/VPCs	5 (TA)

Disabling Trusted Advisor Checks

Using the `--skip-ta` option will disable attempting to query limit information from Trusted Advisor for all commands.

```
(venv) $ awslimitchecker -l --skip-ta
AutoScaling/Auto Scaling groups          20
AutoScaling/Launch configurations        100
EBS/Active snapshots                     10000
EBS/Active volumes                       5000
EBS/General Purpose (SSD) volume storage (GiB) 20
(...)
VPC/Rules per network ACL                20
VPC/Subnets per VPC                     200
VPC/VPCs                                 5
```

Checking Usage

The `-u` or `--show-usage` options to `awslimitchecker` show the current usage for each limit that `awslimitchecker` knows about. It will connect to the AWS API and determine the current usage for each limit. In cases where limits are per-resource instead of account-wide (i.e. “Rules per VPC security group” or “Security groups per VPC”), the usage will be reported for each possible resource in `resource_id=value` format (i.e. for each VPC security group and each VPC, respectively, using their IDs).

```
(venv) $ awslimitchecker -u
AutoScaling/Auto Scaling groups          50
AutoScaling/Launch configurations        50
EBS/Active snapshots                     1065
EBS/Active volumes                       922
EBS/General Purpose (SSD) volume storage (GiB) 5582
(...)
VPC/Rules per network ACL                max: acl-0c279569=4 (acl-0c279569=4, acl-c6d7
VPC/Subnets per VPC                     max: vpc-73ec9716=11 (vpc-a926c2cc=4, vpc-c30
VPC/VPCs                                 4
```

Overriding Limits

In cases where you’ve been given a limit increase by AWS Support, you can override the default limits with custom ones. Currently, to do this from the command line, you must specify each limit that you want to override separately (the `set_limit_overrides()` Python method accepts a dict for easy bulk overrides of limits) using the `-L` or `--limit` options. Limits are specified in a `service_name/limit_name=value` format, and must be quoted if the limit name contains spaces.

For example, to override the limits of EC2’s “EC2-Classic Elastic IPs” and “EC2-VPC Elastic IPs” from their defaults of 5, to 10 and 20, respectively:

```
(venv) $ awslimitchecker -L "AutoScaling/Auto Scaling groups"=321 --limit="AutoScaling/Launch configurations"=456
AutoScaling/Auto Scaling groups          321
AutoScaling/Launch configurations        456
EBS/Active snapshots                     10000 (TA)
EBS/Active volumes                       5000 (TA)
```

EBS/General Purpose (SSD) volume storage (GiB)	30720 (TA)
(...)	
VPC/Rules per network ACL	20
VPC/Subnets per VPC	200
VPC/VPCs	5 (TA)

This example simply sets the overrides, and then prints the limits for confirmation.

Check Limits Against Thresholds

The default mode of operation for `awslimitchecker` (when no other action-specific options are specified) is to check the usage of all known limits, compare them against the configured limit values, and then output a message and set an exit code depending on thresholds. The limit values used will be (in order of precedence) explicitly-set overrides, Trusted Advisor data, and hard-coded defaults.

Currently, the `awslimitchecker` command line script only supports global warning and critical thresholds, which default to 80% and 99% respectively. If any limit's usage is greater than or equal to 80% of its limit value, this will be included in the output and the program will exit with return code 1. If any limit's usage is greater than or equal to 99%, it will include that in the output and exit 2. When determining exit codes, critical takes priority over warning. The output will include the specifics of which limits exceeded the threshold, and for limits that are per-resource, the resource IDs.

The Python class allows setting thresholds per-limit as either a percentage, or an integer usage value, or both; this functionality is not currently present in the command line wrapper.

To check all limits against their thresholds (in this example, one limit has crossed the warning threshold only, and another has crossed the critical threshold):

```
(venv) $ awslimitchecker --no-color
EC2/Running On-Demand EC2 instances (limit 20) CRITICAL: 111
EC2/Running On-Demand m3.medium instances (limit 20) CRITICAL: 75
EC2/Security groups per VPC (limit 100) CRITICAL: vpc-c300b9a6=138
ElastiCache/Clusters (limit 50) CRITICAL: 61
ElastiCache/Nodes (limit 50) CRITICAL: 61
(...)
RDS/VPC Security Groups (limit 5) CRITICAL: 39
VPC/Internet gateways (limit 5) WARNING: 4
VPC/VPCs (limit 5) WARNING: 4
```

Set Custom Thresholds

To set the warning threshold of 50% and a critical threshold of 75% when checking limits:

```
(venv) $ awslimitchecker -W 97 --critical=98 --no-color
EC2/Running On-Demand EC2 instances (limit 20) CRITICAL: 111
EC2/Running On-Demand m3.medium instances (limit 20) CRITICAL: 75
EC2/Security groups per VPC (limit 100) CRITICAL: vpc-c300b9a6=138
ElastiCache/Clusters (limit 50) CRITICAL: 61
ElastiCache/Nodes (limit 50) CRITICAL: 61
ElastiCache/Subnet Groups (limit 50) CRITICAL: 62
RDS/DB snapshots per user (limit 50) CRITICAL: 69
RDS/Subnet Groups (limit 20) CRITICAL: 39
RDS/VPC Security Groups (limit 5) CRITICAL: 39
```


Required IAM Policy

awslimitchecker can also provide the user with an IAM Policy listing the minimum permissions for it to perform all limit checks. This can be viewed with the `--iam-policy` option:

```
(venv) $ awslimitchecker --iam-policy
{
  "Statement": [
    {
      "Action": [
        "autoscaling:DescribeAutoScalingGroups",
        (...)
      ]
    },
    ...
  ],
  "Version": "2012-10-17"
}
```

For the current IAM Policy required by this version of awslimitchecker, see [IAM Policy](#).

7.3 Python Usage

The full feature set of awslimitchecker is available through the Python API. This page attempts to document some examples of usage, but the best resources are [runner](#), the command line wrapper, and the API documentation.

7.3.1 Simple Examples

Many of these examples use `pprint` to make output a bit nicer.

Instantiating the Class

Here we import and instantiate the `AwsLimitChecker` class; note that we also setup Python's `logging` module, which is used by awslimitchecker. We also import `pprint` to make the output nicer.

```
>>> import pprint
>>> import logging
>>> logging.basicConfig()
>>> logger = logging.getLogger()
>>>
>>> from awslimitchecker.checker import AwsLimitChecker
>>> c = AwsLimitChecker()
```

Setting a Limit Override

Override EC2's "EC2-Classic Elastic IPs" limit from its default to 20, using `set_limit_override()`.

```
>>> c.set_limit_override('EC2', 'EC2-Classic Elastic IPs', 20)
```

Checking Thresholds

To check the current usage against limits, use `check_thresholds()`. The return value is a nested dict of all limits with current usage meeting or exceeding the configured thresholds. Keys are the AWS Service names (string), values are dicts of limit name (string) to `AwsLimit` instances representing the limit and its current usage.

```
>>> result = c.check_thresholds()
>>> pprint.pprint(result)
{'EC2': {'Magnetic volume storage (TiB)': <awslimitchecker.limit.AwsLimit object at 0x7f398db62750>,
        'Running On-Demand EC2 instances': <awslimitchecker.limit.AwsLimit object at 0x7f398db55910>,
        'Running On-Demand m3.medium instances': <awslimitchecker.limit.AwsLimit object at 0x7f398db62790>,
        'Security groups per VPC': <awslimitchecker.limit.AwsLimit object at 0x7f398db62790>}}
```

Looking at one of the entries, its `get_warnings()` method tells us that the usage did not exceed its warning threshold:

```
>>> result['EC2']['Magnetic volume storage (TiB)'].get_warnings()
[]
```

But its `get_criticals()` method tells us that it did meet or exceed the critical threshold:

```
>>> result['EC2']['Magnetic volume storage (TiB)'].get_criticals()
[<awslimitchecker.limit.AwsLimitUsage object at 0x7f2074dfeed0>]
```

We can then inspect the `AwsLimitUsage` instance for more information about current usage that crossed the threshold:

In this particular case, there is no resource ID associated with the usage, because it is an aggregate (type-, rather than resource-specific) limit:

```
>>> result['EC2']['Magnetic volume storage (TiB)'].get_criticals()[0].resource_id
>>>
```

The usage is of the EC2 Volume resource type (where one exists, we use the [CloudFormation Resource Type](#) strings to identify resource types).

```
>>> result['EC2']['Magnetic volume storage (TiB)'].get_criticals()[0].aws_type
'AWS::EC2::Volume'
```

We can query the actual numeric usage value:

```
>>> pprint.pprint(result['EC2']['Magnetic volume storage (TiB)'].get_criticals()[0].get_value())
23.337
```

Or a string description of it:

```
>>> print(str(result['EC2']['Magnetic volume storage (TiB)'].get_criticals()[0]))
23.337
```

The “Security groups per VPC” limit also crossed thresholds, and we can see that it has one critical usage value:

```
>>> len(result['EC2']['Security groups per VPC'].get_warnings())
0
>>> len(result['EC2']['Security groups per VPC'].get_criticals())
1
```

As this limit is per-VPC, our string representation of the current usage includes the VPC ID that crossed the critical threshold:

```
>>> for usage in result['EC2']['Security groups per VPC'].get_criticals():
...     print(str(usage))
...
vpc-c300b9a6=100
```

Disabling Trusted Advisor

To disable querying Trusted Advisor for limit information, simply call `get_limits()` or `check_thresholds()` with `use_ta=False`:

```
>>> result = c.check_thresholds(use_ta=False)
```

7.3.2 Logging

awslimitchecker uses the python `logging` library for logging, with module-level loggers defined in each file. If you already have a root-level logger defined in your program and are using a simple configuration (i.e. `logging.basicConfig()`), awslimitchecker logs will be emitted at the same level as that which the root logger is configured.

Assuming you have a root-level logger defined and configured, and you only want to see awslimitchecker log messages of WARNING level and above, you can set the level of awslimitchecker's logger before instantiating the class:

```
alc_log = logging.getLogger('awslimitchecker')
alc_log.setLevel(logging.WARNING)
checker = AwsLimitChecker()
```

It's *highly* recommended that you do not suppress log messages of WARNING or above, as these indicate situations where the checker may not present accurate or complete results.

If your application does not define a root-level logger, this becomes a bit more complicated. Assuming your application has a more complex configuration that uses a top-level logger 'myapp' with its own handlers defined, you can do something like the following. Note that this is highly specific to your logging setup:

```
# setup logging for awslimitchecker
alc_log = logging.getLogger('awslimitchecker')
# WARNING or higher should pass through
alc_log.setLevel(logging.WARNING)
# use myapp's handler(s)
for h in logging.getLogger('cm').handlers:
    alc_log.addHandler(h)
# instantiate the class
checker = AwsLimitChecker()
```

7.3.3 Advanced Examples

CI / Deployment Checks

This example checks usage, logs a message at WARNING level for any warning thresholds surpassed, and logs a message at CRITICAL level for any critical thresholds passed. If any critical thresholds were passed, it exits the script non-zero, i.e. to fail a CI or build job. In this example, we have multiple critical thresholds crossed.

```
>>> import logging
>>> logging.basicConfig()
>>> logger = logging.getLogger()
>>>
>>> from awslimitchecker.checker import AwsLimitChecker
>>> c = AwsLimitChecker()
>>> result = c.check_thresholds()
>>>
>>> have_critical = False
>>> for service, svc_limits in result.items():
```

```

...     for limit_name, limit in svc_limits.items():
...         for warn in limit.get_warnings():
...             logger.warning("{service} '{limit_name}' usage ({u}) exceeds "
...                             "warning threshold (limit={l})".format(
...                                 service=service,
...                                 limit_name=limit_name,
...                                 u=str(warn),
...                                 l=limit.get_limit(),
...                             )
...         )
...     for crit in limit.get_criticals():
...         have_critical = True
...         logger.critical("{service} '{limit_name}' usage ({u}) exceeds "
...                         "critical threshold (limit={l})".format(
...                             service=service,
...                             limit_name=limit_name,
...                             u=str(crit),
...                             l=limit.get_limit(),
...                         )
...         )
...     )
...
CRITICAL:root:EC2 'Magnetic volume storage (TiB)' usage (23.417) exceeds critical threshold (limit=20)
CRITICAL:root:EC2 'Running On-Demand EC2 instances' usage (97) exceeds critical threshold (limit=20)
WARNING:root:EC2 'Security groups per VPC' usage (vpc-c300b9a6=96) exceeds warning threshold (limit=100)
CRITICAL:root:EC2 'Running On-Demand m3.medium instances' usage (53) exceeds critical threshold (limit=5)
CRITICAL:root:EC2 'EC2-Classic Elastic IPs' usage (5) exceeds critical threshold (limit=5)
>>> if have_critical:
...     raise SystemExit(1)
...
(awslimitchecker)$ echo $?
1

```

7.4 Required IAM Permissions

Below is the sample IAM policy from this version of awslimitchecker, listing the IAM permissions required for it to function correctly:

```

{
  "Statement": [
    {
      "Action": [
        "autoscaling:DescribeAutoScalingGroups",
        "autoscaling:DescribeLaunchConfigurations",
        "ec2:DescribeAddresses",
        "ec2:DescribeInstances",
        "ec2:DescribeInternetGatewaysec2:DescribeNetworkAcls",
        "ec2:DescribeNetworkAcls",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeReservedInstances",
        "ec2:DescribeRouteTables",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSubnets",
        "ec2:DescribeSubnets",

```

```

    "ec2:DescribeVolumes",
    "ec2:DescribeVolumes",
    "ec2:DescribeVpcs",
    "ec2:DescribeVpcs",
    "elasticache:DescribeCacheClusters",
    "elasticache:DescribeCacheParameterGroups",
    "elasticache:DescribeCacheSecurityGroups",
    "elasticache:DescribeCacheSubnetGroups",
    "elasticloadbalancing:DescribeLoadBalancers",
    "rds:DescribeDBInstances",
    "rds:DescribeDBParameterGroups",
    "rds:DescribeDBSecurityGroups",
    "rds:DescribeDBSnapshots",
    "rds:DescribeDBSubnetGroups",
    "rds:DescribeEventSubscriptions",
    "rds:DescribeOptionGroups",
    "rds:DescribeReservedDBInstances",
    "support:*",
    "trustedadvisor:Describe*"
  ],
  "Effect": "Allow",
  "Resource": "*"
}
],
"Version": "2012-10-17"
}

```

7.5 Supported Limits

7.5.1 Trusted Advisor Data

So long as the Service and Limit names used by Trusted Advisor (and returned in its API responses) exactly match those shown below, all limits listed in Trusted Advisor “Service Limit” checks should be automatically used by awslimitchecker. The following service limits have been confirmed as being updated from Trusted Advisor:

- AutoScaling
 - Auto Scaling groups
 - Launch configurations
- EBS
 - Active snapshots
 - Active volumes
 - General Purpose (SSD) volume storage (GiB)
 - Magnetic volume storage (GiB)
 - Provisioned IOPS
 - Provisioned IOPS (SSD) storage (GiB)
- EC2
 - Elastic IP addresses (EIPs)
- ELB

- Active load balancers
- RDS
 - DB instances
 - DB security groups
 - DB snapshots per user
 - Max auths per security group
 - Storage quota (GB)
- VPC
 - Internet gateways
 - VPCs

7.5.2 Current Checks

The section below lists every limit that this version of awslimitchecker knows how to check, and its hard-coded default value (per AWS documentation). Limits marked with ^(TA) are confirmed as being updated by Trusted Advisor.

AutoScaling

Limit	Default
Auto Scaling groups ^(TA)	20
Launch configurations ^(TA)	100

EBS

Limit	Default
Active snapshots ^(TA)	10000
Active volumes ^(TA)	5000
General Purpose (SSD) volume storage (GiB) ^(TA)	20
Magnetic volume storage (GiB) ^(TA)	20
Provisioned IOPS (SSD) storage (GiB) ^(TA)	20
Provisioned IOPS ^(TA)	40000

EC2

Limit	Default
EC2-VPC Elastic IPs	5
Elastic IP addresses (EIPs) ^(TA)	5
Rules per VPC security group	50
Running On-Demand EC2 instances	20
Running On-Demand c1.medium instances	20
Running On-Demand c1.xlarge instances	20
Running On-Demand c3.2xlarge instances	20
Running On-Demand c3.4xlarge instances	20
Running On-Demand c3.8xlarge instances	20
Continued on next page	

Table 7.1 – continued from previous page

Limit	Default
Running On-Demand c3.large instances	20
Running On-Demand c3.xlarge instances	20
Running On-Demand c4.2xlarge instances	20
Running On-Demand c4.4xlarge instances	10
Running On-Demand c4.8xlarge instances	5
Running On-Demand c4.large instances	20
Running On-Demand c4.xlarge instances	20
Running On-Demand cc2.8xlarge instances	20
Running On-Demand cg1.4xlarge instances	2
Running On-Demand cr1.8xlarge instances	2
Running On-Demand d2.2xlarge instances	20
Running On-Demand d2.4xlarge instances	10
Running On-Demand d2.8xlarge instances	5
Running On-Demand d2.xlarge instances	20
Running On-Demand g2.2xlarge instances	5
Running On-Demand g2.8xlarge instances	2
Running On-Demand hi1.4xlarge instances	2
Running On-Demand hs1.8xlarge instances	2
Running On-Demand i2.2xlarge instances	8
Running On-Demand i2.4xlarge instances	4
Running On-Demand i2.8xlarge instances	2
Running On-Demand i2.xlarge instances	8
Running On-Demand m1.large instances	20
Running On-Demand m1.medium instances	20
Running On-Demand m1.small instances	20
Running On-Demand m1.xlarge instances	20
Running On-Demand m2.2xlarge instances	20
Running On-Demand m2.4xlarge instances	20
Running On-Demand m2.xlarge instances	20
Running On-Demand m3.2xlarge instances	20
Running On-Demand m3.large instances	20
Running On-Demand m3.medium instances	20
Running On-Demand m3.xlarge instances	20
Running On-Demand m4.2xlarge instances	20
Running On-Demand m4.4xlarge instances	20
Running On-Demand m4.8xlarge instances	20
Running On-Demand m4.large instances	20
Running On-Demand m4.xlarge instances	20
Running On-Demand r3.2xlarge instances	20
Running On-Demand r3.4xlarge instances	10
Running On-Demand r3.8xlarge instances	5
Running On-Demand r3.large instances	20
Running On-Demand r3.xlarge instances	20
Running On-Demand t1.micro instances	20
Running On-Demand t2.medium instances	20
Running On-Demand t2.micro instances	20
Running On-Demand t2.small instances	20
Security groups per VPC	100
VPC security groups per elastic network interface	5

ELB

Limit	Default
Active load balancers ^(TA)	20
Listeners per load balancer	100

ElastiCache

Limit	Default
Clusters	50
Nodes	50
Nodes per Cluster	20
Parameter Groups	20
Security Groups	50
Subnet Groups	50

RDS

Limit	Default
DB instances ^(TA)	40
DB security groups ^(TA)	25
DB snapshots per user ^(TA)	50
Event Subscriptions	20
Max auths per security group ^(TA)	20
Option Groups	20
Parameter Groups	50
Read Replicas per Master	5
Reserved Instances	40
Storage quota (GB) ^(TA)	100000
Subnet Groups	20
Subnets per Subnet Group	20
VPC Security Groups	5

VPC

Limit	Default
Entries per route table	50
Internet gateways ^(TA)	5
Network ACLs per VPC	200
Route tables per VPC	200
Rules per network ACL	20
Subnets per VPC	200
VPCs ^(TA)	5

7.6 Development

7.6.1 Pull Requests

Please cut all pull requests against the “develop” branch. I’ll do my best to merge them as quickly as possible. If they pass all unit tests and have 100% coverage, it’ll certainly be easier. I work on this project only in my personal time, so I can’t always get things merged as quickly as I’d like. That being said, I’m committed to doing my best, and please call me out on it if you feel like I’m not.

7.6.2 Installing for Development

To setup awslimitchecker for development:

1. Fork the `awslimitchecker` repository on GitHub
2. Create a `virtualenv` to run the code in:

```
$ virtualenv awslimitchecker
$ cd awslimitchecker
$ source bin/activate
```

3. Install your fork in the virtualenv as an editable git clone

```
$ pip install -e git+git@github.com:YOUR_NAME/awslimitchecker.git#egg=awslimitchecker
$ cd src/awslimitchecker
```

4. Check out a new git branch. If you’re working on a GitHub issue you opened, your branch should be called “issues/N” where N is the issue number.

7.6.3 Guidelines

- pep8 compliant with some exceptions (see `pytest.ini`)
- 100% test coverage with `pytest` (with valid tests)
- each `_AwsService` subclass should only connect to `boto` once, and should save the connection as `self.conn`. They *must not* connect in the class constructor.
- All modules should have (and use) module-level loggers.
- See the section on the AGPL license below.

7.6.4 Adding New Limits and Checks to Existing Services

1. Add a new `AwsLimit` instance to the return value of the Service class’s `get_limits()` method.
2. In the Service class’s `find_usage()` method (or a method called by that, in the case of large or complex services), get the usage information via `boto` and pass it to the appropriate `AwsLimit` object via its `_add_current_usage()` method. For anything more than trivial services (those with only 2-3 limits), `find_usage()` should be broken into multiple methods, generally one per AWS API call.
3. Ensure complete test coverage for the above.

7.6.5 Adding New Services

All Services are subclasses of `_AwsService` using the `abc` module.

1. The new service name should be in CamelCase, preferably one word (if not one word, it should be underscore-separated). In `awslimitchecker/services`, use the `addservice` script; this will create a templated service class in the current directory, and create a templated (but far from complete) unit test file in `awslimitchecker/tests/services`:

```
./addservice ServiceName
```

2. Find all “TODO” comments in the newly-created files; these have instructions on things to change for new services. Add yourself to the Authors section in the header if desired.
3. Add an import line for the new service in `awslimitchecker/services/__init__.py`.
4. Write at least high-level tests; TDD is greatly preferred.
5. Implement all abstract methods from `_AwsService` and any other methods you need; small, easily-testable methods are preferred. Ensure all methods have full documentation. For simple services, you need only to search for “TODO” in the new service class you created (#1). See [Adding New Limits](#) for further information.
6. Test your code; 100% test coverage is expected, and mocks should be using `autospec` or `spec_set`.
7. Ensure the `required_iam_permissions()` method of your new class returns a list of all IAM permissions required for it to work.
8. Write integration tests. (currently not implemented; see [issue #21](#))
9. Run all tox jobs, or at least one python version, docs and coverage.
10. Commit the updated documentation to the repository.
11. As there is no programmatic way to validate IAM policies, once you are done writing your service, grab the output of `awslimitchecker --iam-policy`, login to your AWS account, and navigate to the IAM page. Click through to create a new policy, paste the output of the `--iam-policy` command, and click the “Validate Policy” button. Correct any errors that occur; for more information, see the AWS IAM docs on [Using Policy Validator](#). It would also be a good idea to run any policy changes through the [Policy Simulator](#).
12. Submit your pull request.

7.6.6 Trusted Advisor Checks

So long as the `Service` and `Limit` name strings returned by the Trusted Advisor (Support) API exactly match how they are set on the corresponding `_AwsService` and `AwsLimit` objects, no code changes are needed to support new limit checks from TA.

For further information, see [Internals / Trusted Advisor](#).

7.6.7 Unit Testing

Testing is done via `pytest`, driven by `tox`.

- testing is as simple as:
 - `pip install tox`
 - `tox`
- If you want to see code coverage: `tox -e cov`

- this produces two coverage reports - a summary on STDOUT and a full report in the `htmlcov/` directory
- If you want to pass additional arguments to pytest, add them to the tox command line after “-”. i.e., for verbose pytest output on py27 tests: `tox -e py27 -- -v`

Note that while boto currently doesn’t have python3 support, we still run tests against py3 to ensure that this package is ready for it when boto is.

7.6.8 Integration Testing

currently not implemented; see [issue #21](#)

7.6.9 Building Docs

Much like the test suite, documentation is build using tox:

```
$ tox -e docs
```

Output will be in the `docs/build/html` directory under the project root.

7.6.10 AGPL License

awslimitchecker is licensed under the [GNU Affero General Public License, version 3 or later](#).

Pursuant to Sections 5(b) and 13 of the license, all users of awslimitchecker - including those interacting with it remotely over a network - have a right to obtain the exact, unmodified running source code. We have done as much as possible to make this transparent to developers, with no additional work needed. See the guidelines below for information.

- If you’re simply *running* awslimitchecker via the command line, there’s nothing to worry about; just use it like any other software.
- If you’re using awslimitchecker in your own software in a way that allows users to interact with it over the network (i.e. in your deployment or monitoring systems), but not modifying it, you also don’t need to do anything special; awslimitchecker will log a WARNING-level message indicating where the source code of the currently-running version can be obtained. So long as you’ve installed awslimitchecker via Python’s packaging system (i.e. with *pip*), its current version and source will be automatically detected. This suffices for the AGPL source code offer provision, so long as it’s displayed to users and the currently-running source is unmodified.
- If you wish to modify the source code of awslimitchecker, you need to do is ensure that `_get_version_info()` always returns correct and accurate information (a publicly-accessible URL to the exact version of the running source code, and a version number). If you install your modified version directly from an editable (i.e. `pip install -e`), publicly-accessible Git repository, and ensure that changes are available in the repository before they are present in the code running for your users, this should be automatically detected by awslimitchecker and the correct URL provided. It is strongly recommended that any such repository is a fork of the project’s original GitHub repository. It is solely your responsibility to ensure that the URL and version information presented to users is accurate and reflects source code identical to what is running.
- If you’re distributing awslimitchecker with modifications or as part of your own software (as opposed to simply an editable requirement that gets installed with pip), please read the license and ensure that you comply with its terms.
- If you are running awslimitchecker as part of a hosted service that users somehow interact with, please ensure that the source code URL and version is correct and visible in the output given to users.

7.6.11 Release Checklist

1. Open an issue for the release; cut a branch off `develop` for that issue.
2. Build docs (`tox -e docs`) and ensure they're current; commit any changes.
3. Ensure that Travis tests are passing in all environments. If there were any changes to `awslimitchecker.versioncheck`, manually run the `-versioncheck` tox environments (these are problematic in Travis and with PRs).
4. Ensure that test coverage is no less than the last release (ideally, 100%).
5. Create or update an actual IAM user with the policy from `awslimitchecker --iam-policy`; run the command line wrapper and ensure that the policy works and contains all needed permissions.
6. Build docs for the branch (locally) and ensure they look correct.
7. Increment the version number in `awslimitchecker/version.py` and add version and release date to `CHANGES.rst`. Ensure that there are `CHANGES.rst` entries for all major changes since the last release. Mention the issue in the commit for this, and push to GitHub.
8. Confirm that `README.rst` renders correctly on GitHub.
9. Upload package to testpypi, confirm that `README.rst` renders correctly.
 - Make sure your `~/.pypirc` file is correct
 - `python setup.py register -r https://testpypi.python.org/pypi`
 - `python setup.py sdist upload -r https://testpypi.python.org/pypi`
 - Check that the `README` renders at <https://testpypi.python.org/pypi/awslimitchecker>
10. Create a pull request for the release to be merge into master. Upon successful Travis build, merge it.
11. Tag the release in Git, push tag to GitHub:
 - tag the release. for now the message is quite simple: `git tag -a vX.Y.Z -m 'X.Y.Z released YYYY-MM-DD'`
 - push the tag to GitHub: `git push origin vX.Y.Z`
12. Upload package to live pypi:
 - `python setup.py sdist upload`
13. make sure any GH issues fixed in the release were closed.
14. merge master back into develop
15. Blog, tweet, etc. about the new version.

7.7 Internals

7.7.1 Overall Program Flow

`AwsLimitChecker` provides the full and only public interface to this project; it's used by the `awslimitchecker` command line script (entry point to `runner`) and should be the only portion directly used by external code.

Each AWS Service is represented by a subclass of the `_AwsService` abstract base class; these Service Classes are responsible for knowing which limits exist for the service they represent, what the default values for these limits are, and how to check the current usage via the AWS API (via `boto`). When the Service Classes are instantiated, they build a dict of all of their limits, correlating a string key (the "limit name") with an `AwsLimit` object. The Service

Class constructors *must not* make any network connections; connections are created lazily as needed and stored as a class attribute. This allows us to inspect the services, limits and default limit values without ever connecting to AWS (this is also used to generate the *Supported Limits* documentation automatically).

When `AwsLimitChecker` is instantiated, it imports `services` which in turn creates instances of all `awslimitchecker.services.*` classes and adds them to a dict mapping the string Service Name to the Service Class instance. These instances are used for all interaction with the services.

So, once an instance of `AwsLimitChecker` is created, we should have instant access to the services and limits without any connection to AWS. This is utilized by the `--list-services` and `--list-defaults` options for the *command line client*.

7.7.2 Trusted Advisor

When `AwsLimitChecker` is initialized, it also initializes an instance of `TrustedAdvisor`. In `get_limits()`, `find_usage()` and `check_thresholds()`, when called with `use_ta == True` (the default), `update_limits()` is called on the `TrustedAdvisor` instance.

`update_limits()` polls Trusted Advisor data from the Support API via `_poll()`; this will retrieve the limits for all “flaggedResources” items in the Service Limits Trusted Advisor check result for the current AWS account. It then calls `_update_services()`, passing in the Trusted Advisor check results and the dict of `_AwsService` objects it was called with (from `AwsLimitChecker`).

`_update_services()` iterates over the Services in the Trusted Advisor check result and attempts to find a matching `_AwsService` (by string service name) in the dict passed in from `AwsLimitChecker`. If a match is found, it iterates over all limits for that service in the TA result and attempts to call the Service’s `_set_ta_limit()` method. If a matching Service is not found, or if `_set_ta_limit` raises a `ValueError` (matching Limit not found for that Service), an error is logged.

Using this methodology, no additional code is needed to support new/additional Trusted Advisor limit checks; *so long as* the Service and Limit name strings match between the Trusted Advisor API response and their corresponding `_AwsService` and `AwsLimit` instances, the TA limits will be automatically added to the corresponding `AwsLimit`.

7.8 awslimitchecker

7.8.1 awslimitchecker package

Subpackages

`awslimitchecker.services` package

Submodules

`awslimitchecker.services.autoscaling` module

`class awslimitchecker.services.autoscaling._AutoscalingService` (*warning_threshold*, *critical_threshold*)

Bases: `awslimitchecker.services.base._AwsService`

Describes an AWS service and its limits, and provides methods to query current utilization.

Constructors of `_AwsService` subclasses *must not* make any external connections; these must be made lazily as needed in other methods. `_AwsService` subclasses should be usable without any external network connections.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

`__abstractmethods__ = frozenset([])`

`__module__ = 'awslimitchecker.services.autoscaling'`

`__abc_cache = <_weakrefset.WeakSet object>`

`__abc_negative_cache = <_weakrefset.WeakSet object>`

`__abc_negative_cache_version = 25`

`__abc_registry = <_weakrefset.WeakSet object>`

`connect ()`

connect to API if not already connected; set self.conn

`find_usage ()`

Determine the current usage for each limit of this service, and update corresponding Limit via `__add_current_usage ()`.

`get_limits ()`

Return all known limits for this service, as a dict of their names to `AwsLimit` objects.

Returns dict of limit names to `AwsLimit` objects

Return type dict

`required_iam_permissions ()`

Return a list of IAM Actions required for this Service to function properly. All Actions will be shown with an Effect of “Allow” and a Resource of “*”.

Returns list of IAM Action strings

Return type list

`service_name = 'AutoScaling'`

awslimitchecker.services.base module

class `awslimitchecker.services.base._AwsService (warning_threshold, critical_threshold)`

Bases: `object`

Describes an AWS service and its limits, and provides methods to query current utilization.

Constructors of `_AwsService` subclasses *must not* make any external connections; these must be made lazily as needed in other methods. `_AwsService` subclasses should be usable without any external network connections.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

`__abstractmethods__ = frozenset(['get_limits', 'find_usage', 'required_iam_permissions', 'connect'])`

`__dict__ = dict_proxy({'__abc_cache': <_weakrefset.WeakSet object at 0x7f50b52c6410>, 'set_limit_override': <function`

`__init__` (*warning_threshold*, *critical_threshold*)

Describes an AWS service and its limits, and provides methods to query current utilization.

Constructors of `_AwsService` subclasses *must not* make any external connections; these must be made lazily as needed in other methods. `_AwsService` subclasses should be usable without any external network connections.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

`__metaclass__`

alias of `ABCMeta`

`__module__` = 'awslimitchecker.services.base'

`__weakref__`

list of weak references to the object (if defined)

`_abc_cache` = `<_weakrefset.WeakSet object>`

`_abc_negative_cache` = `<_weakrefset.WeakSet object>`

`_abc_negative_cache_version` = 25

`_abc_registry` = `<_weakrefset.WeakSet object>`

`_set_ta_limit` (*limit_name*, *value*)

Set the value for the limit as reported by Trusted Advisor, for the specified limit.

This method should only be called by `TrustedAdvisor`.

Parameters

- **limit_name** (*string*) – the name of the limit to override the value for
- **value** (*int*) – the Trusted Advisor limit value

Raises `ValueError` if `limit_name` is not known to this service

`check_thresholds` ()

Checks current usage against configured thresholds for all limits for this service.

Returns a dict of limit name to `AwsLimit` instance for all limits that crossed one or more of their thresholds.

Return type `dict` of `AwsLimit`

`connect` ()

If not already done, establish a connection to the relevant AWS service and save as `self.conn`.

`find_usage` ()

Determine the current usage for each limit of this service, and update the `current_usage` property of each corresponding `AwsLimit` instance.

This method must set `self._have_usage = True`

`get_limits` ()

Return all known limits for this service, as a dict of their names to `AwsLimit` objects.

All limits must have `self.warning_threshold` and `self.critical_threshold` passed into them.

Returns dict of limit names to *AwsLimit* objects

Return type dict

required_iam_permissions ()

Return a list of IAM Actions required for this Service to function properly. All Actions will be shown with an Effect of “Allow” and a Resource of “*”.

Returns list of IAM Action strings

Return type list

service_name = ‘baseclass’

set_limit_override (*limit_name*, *value*, *override_ta=True*)

Set a new limit *value* for the specified limit, overriding the default. If *override_ta* is True, also use this value instead of any found by Trusted Advisor. This method simply passes the data through to the *set_limit_override()* method of the underlying *AwsLimit* instance.

Parameters

- **limit_name** (*string*) – the name of the limit to override the value for
- **value** (*int*) – the new value to set for the limit
- **override_ta** (*bool*) – whether or not to also override Trusted Advisor information

Raises ValueError if *limit_name* is not known to this service

set_threshold_override (*limit_name*, *warn_percent=None*, *warn_count=None*, *crit_percent=None*, *crit_count=None*)

Override the default warning and critical thresholds used to evaluate the specified limit’s usage. Thresholds can be specified as a percentage of the limit, or as a usage count, or both.

Parameters

- **warn_percent** (*int*) – new warning threshold, percentage used
- **warn_count** (*int*) – new warning threshold, actual count/number
- **crit_percent** (*int*) – new critical threshold, percentage used
- **crit_count** (*int*) – new critical threshold, actual count/number

awslimitchecker.services.ebs module

class *awslimitchecker.services.ebs._EbsService* (*warning_threshold*, *critical_threshold*)

Bases: *awslimitchecker.services.base._AwsService*

Describes an AWS service and its limits, and provides methods to query current utilization.

Constructors of *_AwsService* subclasses *must not* make any external connections; these must be made lazily as needed in other methods. *_AwsService* subclasses should be usable without any external network connections.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

__abstractmethods__ = frozenset([])

__module__ = ‘awslimitchecker.services.ebs’

__abc_cache = <_weakrefset.WeakSet object>


```

_abc_negative_cache = <_weakrefset.WeakSet object>
_abc_negative_cache_version = 25
_abc_registry = <_weakrefset.WeakSet object>
_find_usage_ebs ()
    calculate usage for all EBS limits and update Limits
_find_usage_snapshots ()
    find snapshot usage
_get_limits_ebs ()
    Return a dict of EBS-related limits only. This method should only be used internally by
    :py:meth:`~.get_limits`.

    Return type dict

connect ()
    connect to API if not already connected; set self.conn

find_usage ()
    Determine the current usage for each limit of this service, and update corresponding Limit via
    _add_current_usage().

get_limits ()
    Return all known limits for this service, as a dict of their names to AwsLimit objects.

    Returns dict of limit names to AwsLimit objects

    Return type dict

required_iam_permissions ()
    Return a list of IAM Actions required for this Service to function properly. All Actions will be shown with
    an Effect of “Allow” and a Resource of “*”.

    Returns list of IAM Action strings

    Return type list

service_name = ‘EBS’

```

awslimitchecker.services.ec2 module

```

class awslimitchecker.services.ec2._Ec2Service (warning_threshold, critical_threshold)
    Bases: awslimitchecker.services.base._AwsService

```

Describes an AWS service and its limits, and provides methods to query current utilization.

Constructors of `_AwsService` subclasses *must not* make any external connections; these must be made lazily as needed in other methods. `_AwsService` subclasses should be usable without any external network connections.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

```

__abstractmethods__ = frozenset([])
__module__ = ‘awslimitchecker.services.ec2’
_abc_cache = <_weakrefset.WeakSet object>

```

`_abc_negative_cache = <_weakrefset.WeakSet object>`

`_abc_negative_cache_version = 25`

`_abc_registry = <_weakrefset.WeakSet object>`

`_find_usage_instances ()`
calculate On-Demand instance usage for all types and update Limits

`_find_usage_networking_eips ()`

`_find_usage_networking_eni_sg ()`

`_find_usage_networking_sgs ()`
calculate usage for VPC-related things

`_get_limits_instances ()`
Return a dict of limits for EC2 instances only. This method should only be used internally by :py:meth:`~.get_limits`.

Return type dict

`_get_limits_networking ()`
Return a dict of VPC-related limits only. This method should only be used internally by :py:meth:`~.get_limits`.

Return type dict

`_get_reserved_instance_count ()`
For each availability zone, get the count of current instance reservations of each instance type. Return as a nested dict of AZ name to dict of instance type to reservation count.

Return type dict

`_instance_types ()`
Return a list of all known EC2 instance types

Returns list of all valid known EC2 instance types

Return type list

`_instance_usage ()`
Find counts of currently-running EC2 Instances (On-Demand or Reserved) by placement (Availability Zone) and instance type (size). Return as a nested dict of AZ name to dict of instance type to count.

Return type dict

`connect ()`
connect to API if not already connected; set self.conn

`find_usage ()`
Determine the current usage for each limit of this service, and update corresponding Limit via `_add_current_usage ()`.

`get_limits ()`
Return all known limits for this service, as a dict of their names to `AwsLimit` objects.

Returns dict of limit names to `AwsLimit` objects

Return type dict

`required_iam_permissions ()`
Return a list of IAM Actions required for this Service to function properly. All Actions will be shown with an Effect of "Allow" and a Resource of "*".

Returns list of IAM Action strings

Return type `list`

`service_name = 'EC2'`

awslimitchecker.services.elasticache module

class `awslimitchecker.services.elasticache._ElasticCacheService` (*warning_threshold*, *critical_threshold*)

Bases: `awslimitchecker.services.base._AwsService`

Describes an AWS service and its limits, and provides methods to query current utilization.

Constructors of `_AwsService` subclasses *must not* make any external connections; these must be made lazily as needed in other methods. `_AwsService` subclasses should be usable without any external network connections.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

`__abstractmethods__ = frozenset([])`

`__module__ = 'awslimitchecker.services.elasticache'`

`_abc_cache = <_weakrefset.WeakSet object>`

`_abc_negative_cache = <_weakrefset.WeakSet object>`

`_abc_negative_cache_version = 25`

`_abc_registry = <_weakrefset.WeakSet object>`

`_find_usage_nodes ()`

find usage for cache nodes

`_find_usage_parameter_groups ()`

find usage for elasticache parameter groups

`_find_usage_security_groups ()`

find usage for elasticache security groups

`_find_usage_subnet_groups ()`

find usage for elasticache subnet groups

`connect ()`

`find_usage ()`

Determine the current usage for each limit of this service, and update corresponding Limit via `_add_current_usage ()`.

`get_limits ()`

Return all known limits for this service, as a dict of their names to `AwsLimit` objects.

Returns dict of limit names to `AwsLimit` objects

Return type `dict`

`required_iam_permissions ()`

Return a list of IAM Actions required for this Service to function properly. All Actions will be shown with an Effect of “Allow” and a Resource of “*”.

Returns list of IAM Action strings

Return type `list`

`service_name = 'ElastiCache'`

awslimitchecker.services.elb module

class `awslimitchecker.services.elb._ElbService` (*warning_threshold, critical_threshold*)

Bases: `awslimitchecker.services.base._AwsService`

Describes an AWS service and its limits, and provides methods to query current utilization.

Constructors of `_AwsService` subclasses *must not* make any external connections; these must be made lazily as needed in other methods. `_AwsService` subclasses should be usable without any external network connections.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

`__abstractmethods__ = frozenset([])`

`__module__ = 'awslimitchecker.services.elb'`

`_abc_cache = <_weakrefset.WeakSet object>`

`_abc_negative_cache = <_weakrefset.WeakSet object>`

`_abc_negative_cache_version = 25`

`_abc_registry = <_weakrefset.WeakSet object>`

`connect ()`

`find_usage ()`

Determine the current usage for each limit of this service, and update corresponding Limit via `_add_current_usage ()`.

`get_limits ()`

Return all known limits for this service, as a dict of their names to `AwsLimit` objects.

Returns dict of limit names to `AwsLimit` objects

Return type `dict`

`required_iam_permissions ()`

Return a list of IAM Actions required for this Service to function properly. All Actions will be shown with an Effect of “Allow” and a Resource of “*”.

Returns list of IAM Action strings

Return type `list`

`service_name = 'ELB'`

awslimitchecker.services.rds module

class `awslimitchecker.services.rds._RDSService` (*warning_threshold, critical_threshold*)

Bases: `awslimitchecker.services.base._AwsService`

Describes an AWS service and its limits, and provides methods to query current utilization.

Constructors of `_AwsService` subclasses *must not* make any external connections; these must be made lazily as needed in other methods. `_AwsService` subclasses should be usable without any external network connections.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

__abstractmethods__ = frozenset([])

__module__ = 'awslimitchecker.services.rds'

__abc_cache = <_weakrefset.WeakSet object>

__abc_negative_cache = <_weakrefset.WeakSet object>

__abc_negative_cache_version = 25

__abc_registry = <_weakrefset.WeakSet object>

__find_usage_event_subscriptions ()
find usage for event subscriptions

__find_usage_instances ()
find usage for DB Instances and related limits

__find_usage_option_groups ()
find usage for option groups

__find_usage_param_groups ()
find usage for parameter groups

__find_usage_reserved_instances ()
find usage for reserved instances

__find_usage_security_groups ()
find usage for security groups

__find_usage_snapshots ()
find usage for (manual) DB snapshots

__find_usage_subnet_groups ()
find usage for subnet groups

connect ()

find_usage ()
Determine the current usage for each limit of this service, and update corresponding Limit via [__add_current_usage](#) ().

get_limits ()
Return all known limits for this service, as a dict of their names to *AwsLimit* objects.

Returns dict of limit names to *AwsLimit* objects

Return type dict

required_iam_permissions ()

Return a list of IAM Actions required for this Service to function properly. All Actions will be shown with an Effect of “Allow” and a Resource of “*”.

Returns list of IAM Action strings

Return type list

service_name = 'RDS'

awslimitchecker.services.vpc module

class `awslimitchecker.services.vpc._VpcService` (*warning_threshold*, *critical_threshold*)

Bases: `awslimitchecker.services.base._AwsService`

Describes an AWS service and its limits, and provides methods to query current utilization.

Constructors of `_AwsService` subclasses *must not* make any external connections; these must be made lazily as needed in other methods. `_AwsService` subclasses should be usable without any external network connections.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

`__abstractmethods__` = `frozenset([])`

`__module__` = `'awslimitchecker.services.vpc'`

`_abc_cache` = `<_weakrefset.WeakSet object>`

`_abc_negative_cache` = `<_weakrefset.WeakSet object>`

`_abc_negative_cache_version` = `25`

`_abc_registry` = `<_weakrefset.WeakSet object>`

`_find_usage_ACLs` ()

find usage for ACLs

`_find_usage_gateways` ()

find usage for Internet Gateways

`_find_usage_route_tables` ()

find usage for route tables

`_find_usage_subnets` ()

find usage for Subnets

`_find_usage_vpcs` ()

find usage for VPCs

`connect` ()

connect to API if not already connected; set `self.conn`

`find_usage` ()

Determine the current usage for each limit of this service, and update corresponding Limit via `_add_current_usage` ().

`get_limits` ()

Return all known limits for this service, as a dict of their names to `AwsLimit` objects.

Returns dict of limit names to `AwsLimit` objects

Return type `dict`

`required_iam_permissions` ()

Return a list of IAM Actions required for this Service to function properly. All Actions will be shown with an Effect of “Allow” and a Resource of “*”.

Returns list of IAM Action strings

Return type `list`

```
service_name = 'VPC'
```

Submodules

awslimitchecker.checker module

```
class awslimitchecker.checker.AwsLimitChecker (warning_threshold=80, critical_threshold=99)
```

Bases: `object`

Main `AwsLimitChecker` class - this should be the only externally-used portion of `awslimitchecker`.

Constructor builds `self.services` as a dict of `service_name` (str) to `__AwsService` instance, and sets limit thresholds.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

```
__dict__ = dict_proxy({'set_limit_override': <function set_limit_override at 0x7f50b51d0500>, '__module__': 'awslimitchecker.checker'})
```

```
__init__ (warning_threshold=80, critical_threshold=99)
```

Main `AwsLimitChecker` class - this should be the only externally-used portion of `awslimitchecker`.

Constructor builds `self.services` as a dict of `service_name` (str) to `__AwsService` instance, and sets limit thresholds.

Parameters

- **warning_threshold** (*int*) – the default warning threshold, as an integer percentage, for any limits without a specifically-set threshold.
- **critical_threshold** (*int*) – the default critical threshold, as an integer percentage, for any limits without a specifically-set threshold.

```
__module__ = 'awslimitchecker.checker'
```

```
__weakref__
```

list of weak references to the object (if defined)

```
check_thresholds (service=None, use_ta=True)
```

Check all limits and current usage against their specified thresholds; return all `AwsLimit` instances that have crossed one or more of their thresholds.

If `service` is specified, the returned dict has one element, the service name, whose value is a nested dict as described below; otherwise it includes all known services.

The returned `AwsLimit` objects can be interrogated for their limits (`get_limit()`) as well as the details of usage that crossed the thresholds (`get_warnings()` and `get_criticals()`).

Parameters

- **service** (*string*) – the name of one service to return results for
- **use_ta** (*bool*) – check Trusted Advisor for information on limits

Returns dict of service name (string) to nested dict of limit name (string) to limit (`AwsLimit`)

Return type dict

find_usage (*service=None, use_ta=True*)

For each limit in the specified service (or all services if *service* is *None*), query the AWS API via *boto* and find the current usage amounts for that limit.

This method updates the *current_usage* attribute of the *AwsLimit* objects for each service, which can then be queried using *get_limits()*.

Parameters

- **service** (*None*, or *string* service name to get) – *__AwsService* name, or *None* to check all services.
- **use_ta** (*bool*) – check Trusted Advisor for information on limits

get_limits (*service=None, use_ta=True*)

Return all *AwsLimit* objects for the given service name, or for all services if *service* is *None*.

If *service* is specified, the returned dict has one element, the service name, whose value is a nested dict as described below.

Parameters

- **service** (*string*) – the name of one service to return limits for
- **use_ta** (*bool*) – check Trusted Advisor for information on limits

Returns dict of service name (*string*) to nested dict of limit name (*string*) to limit (*AwsLimit*)

Return type *dict*

get_project_url ()

Return the URL for the awslimitchecker project.

Returns URL of where to find awslimitchecker

Return type *string*

get_required_iam_policy ()

Return an IAM policy granting all of the permissions needed for awslimitchecker to fully function. This returns a dict suitable for json serialization to a valid IAM policy.

Internally, this calls *required_iam_permissions()* on each *__AwsService* instance.

Returns dict representation of IAM Policy

Return type *dict*

get_service_names ()

Return a list of all known service names

Returns list of service names

Return type *list*

get_version ()

Return the version of awslimitchecker currently running.

Returns current awslimitchecker version

Return type *string*

set_limit_override (*service_name, limit_name, value, override_ta=True*)

Set a manual override on an AWS service limits, i.e. if you had limits increased by AWS support.

This method calls *set_limit_override()* on the corresponding *__AwsService* instance.

Explicitly set limit overrides using this method will take precedence over default limits. They will also take precedence over limit information obtained via Trusted Advisor, unless `override_ta` is set to `False`.

Parameters

- **service_name** (*string*) – the name of the service to override limit for
- **limit_name** (*string*) – the name of the limit to override:
- **value** (*int*) – the new (overridden) limit value)
- **override_ta** (*bool*) – whether or not to use this value even if Trusted Advisor supplies limit information

Raises `ValueError` if `limit_name` is not known to the service instance

set_limit_overrides (*override_dict*, *override_ta=True*)

Set manual overrides on AWS service limits, i.e. if you had limits increased by AWS support. This takes a dict in the same form as that returned by `get_limits()`, i.e. `service_name` (str) keys to nested dict of `limit_name` (str) to limit value (int) like:

```
{
  'EC2': {
    'Running On-Demand t2.micro Instances': 1000,
    'Running On-Demand r3.4xlarge Instances': 1000,
  }
}
```

Internally, for each limit override for each service in `override_dict`, this method calls `set_limit_override()` on the corresponding `_AwsService` instance.

Explicitly set limit overrides using this method will take precedence over default limits. They will also take precedence over limit information obtained via Trusted Advisor, unless `override_ta` is set to `False`.

Parameters

- **override_dict** (*dict*) – dict of overrides to default limits
- **override_ta** (*bool*) – whether or not to use this value even if Trusted Advisor supplies limit information

Raises `exceptions.ValueError` if `limit_name` is not known to the service instance

set_threshold_override (*service_name*, *limit_name*, *warn_percent=None*, *warn_count=None*, *crit_percent=None*, *crit_count=None*)

Set a manual override on the threshold (used for determining warning/critical status) for a specific limit. See `AwsLimitChecker` for information on Warning and Critical thresholds.

See `set_threshold_override()`.

Parameters

- **service_name** (*string*) – the name of the service to override limit for
- **limit_name** (*string*) – the name of the limit to override:
- **warn_percent** (*int*) – new warning threshold, percentage used
- **warn_count** (*int*) – new warning threshold, actual count/number
- **crit_percent** (*int*) – new critical threshold, percentage used
- **crit_count** (*int*) – new critical threshold, actual count/number

set_threshold_overrides (*override_dict*)

Set manual overrides on the threshold (used for determining warning/critical status) a dict of limits. See *AwsLimitChecker* for information on Warning and Critical thresholds.

Dict is composed of service name keys (string) to dict of limit names (string), to dict of threshold specifications. Each threshold specification dict can contain keys ‘warning’ or ‘critical’, each having a value of a dict containing keys ‘percent’ or ‘count’, to an integer value.

Example:

```
{
  'EC2': {
    'SomeLimit': {
      'warning': {
        'percent': 80,
        'count': 8,
      },
      'critical': {
        'percent': 90,
        'count': 9,
      }
    }
  }
}
```

See *set_threshold_override()*.

Parameters *override_dict* (*dict*) – nested dict of threshold overrides

awslimitchecker.limit module

class *awslimitchecker.limit.AwsLimit* (*name*, *service*, *default_limit*, *def_warning_threshold*, *def_critical_threshold*, *limit_type=None*, *limit_subtype=None*)

Bases: *object*

Describes one specific AWS service limit, as well as its current utilization, default limit, thresholds, and any Trusted Advisor information about this limit.

Parameters

- **name** (*string*) – the name of this limit (may contain spaces); if possible, this should be the name used by AWS, i.e. TrustedAdvisor
- **service** – the *_AwsService* class that this limit is for
- **default_limit** (*int*) – the default value of this limit for new accounts
- **def_warning_threshold** (*int*) – the default warning threshold, as an integer percentage.
- **def_critical_threshold** (*int*) – the default critical threshold, as an integer percentage.
- **limit_type** – the type of resource this limit describes, specified as one of the type names used in [CloudFormation](#) # noqa such as “AWS::EC2::Instance” or “AWS::RDS::DBSubnetGroup”.
- **limit_subtype** – resource sub-type for this limit, if applicable, such as “t2.micro” or “SecurityGroup”

Raises ValueError

`__dict__ = dict_proxy({'set_limit_override': <function set_limit_override at 0x7f50b522b5f0>, '__module__': 'awslimitchecker.limit'})`

`__init__(name, service, default_limit, def_warning_threshold, def_critical_threshold, limit_type=None, limit_subtype=None)`

Describes one specific AWS service limit, as well as its current utilization, default limit, thresholds, and any Trusted Advisor information about this limit.

Parameters

- **name** (*string*) – the name of this limit (may contain spaces); if possible, this should be the name used by AWS, i.e. TrustedAdvisor
- **service** – the `_AwsService` class that this limit is for
- **default_limit** (*int*) – the default value of this limit for new accounts
- **def_warning_threshold** (*int*) – the default warning threshold, as an integer percentage.
- **def_critical_threshold** (*int*) – the default critical threshold, as an integer percentage.
- **limit_type** – the type of resource this limit describes, specified as one of the type names used in `CloudFormation # noqa` such as “AWS::EC2::Instance” or “AWS::RDS::DBSubnetGroup”.
- **limit_subtype** – resource sub-type for this limit, if applicable, such as “t2.micro” or “SecurityGroup”

Raises ValueError

`__module__ = 'awslimitchecker.limit'`

`__weakref__`

list of weak references to the object (if defined)

`__add_current_usage(value, resource_id=None, aws_type=None)`

Add a new current usage value for this limit.

Creates a new `AwsLimitUsage` instance and appends it to the internal list. If more than one usage value is given to this service, they should have `id` and `aws_type` set.

This method should only be called from the `_AwsService` instance that created and manages this Limit.

Parameters

- **value** (*int* or *float*) – the numeric usage value
- **resource_id** (*string*) – If there can be multiple usage values for one limit, an AWS ID for the resource this instance describes
- **aws_type** (*string*) – if `id` is not `None`, the AWS resource type that ID represents. As a convention, we use the AWS Resource Type names used by `CloudFormation # noqa`

`__get_thresholds()`

Get the warning and critical thresholds for this Limit.

Return type is a 4-tuple of:

- 1.warning integer (usage) threshold, or None
- 2.warning percent threshold
- 3.critical integer (usage) threshold, or None

4.critical percent threshold

Return type tuple

`__reset_usage()`

Discard all current usage data.

`__set_ta_limit(limit_value)`

Set the value for the limit as reported by Trusted Advisor.

This method should only be called by *TrustedAdvisor*.

Parameters `limit_value` (*int*) – the Trusted Advisor limit value

`check_thresholds()`

Check this limit's current usage against the specified default thresholds, and any custom thresholds that have been set on the instance. Return True if usage is within thresholds, or false if warning or critical thresholds have been surpassed.

This method sets internal variables in this instance which can be queried via *get_warnings()* and *get_criticals()* to obtain further details about the thresholds that were crossed.

Parameters

- **`default_warning`** (*int* or *float* percentage) – default warning threshold in percentage; usage higher than this percent of the limit will be considered a warning
- **`default_critical`** (*int* or *float* percentage) – default critical threshold in percentage; usage higher than this percent of the limit will be considered a critical

Returns False if any thresholds crossed, True otherwise

Return type bool

`get_criticals()`

Return a list of *AwsLimitUsage* instances that crossed the critical threshold. These objects are comparable and can be sorted.

Return type list

`get_current_usage()`

Get the current usage for this limit, as a list of *AwsLimitUsage* instances.

Returns list of current usage values

Return type list of *AwsLimitUsage*

`get_current_usage_str()`

Get the a string describing the current usage for this limit.

If no usage has been added for this limit, the result will be “<unknown>”.

If the limit has only one current usage instance, this will be that instance's *__str__()* value.

If the limit has more than one current usage instance, this will be the a string of the form `max: X (Y)` where X is the *__str__()* value of the instance with the maximum value, and Y is a comma-separated list of the *__str__()* values of all usage instances in ascending order.

Returns representation of current usage

Return type string

`get_limit()`

Returns the effective limit value for this Limit, taking into account limit overrides and Trusted Advisor data.

Returns effective limit value

Return type `int`

get_limit_source ()

Return `SOURCE_DEFAULT` if `get_limit()` returns the default limit, `SOURCE_OVERRIDE` if it returns a manually-overridden limit, or `SOURCE_TA` if it returns a limit from Trusted Advisor.

Returns one of `SOURCE_DEFAULT`, `SOURCE_OVERRIDE`, or `SOURCE_TA`

Return type `int`

get_warnings ()

Return a list of `AwsLimitUsage` instances that crossed the warning threshold. These objects are comparable and can be sorted.

Return type `list`

set_limit_override (*limit_value*, *override_ta=True*)

Set a new value for this limit, to override the default (such as when AWS Support has increased a limit of yours). If `override_ta` is `True`, this value will also supersede any found through Trusted Advisor.

Parameters

- **limit_value** (*int*) – the new limit value
- **override_ta** (*bool*) – whether or not to also override Trusted Advisor information

set_threshold_override (*warn_percent=None*, *warn_count=None*, *crit_percent=None*, *crit_count=None*)

Override the default warning and critical thresholds used to evaluate this limit’s usage. Thresholds can be specified as a percentage of the limit, or as a usage count, or both.

Parameters

- **warn_percent** (*int*) – new warning threshold, percentage used
- **warn_count** (*int*) – new warning threshold, actual count/number
- **crit_percent** (*int*) – new critical threshold, percentage used
- **crit_count** (*int*) – new critical threshold, actual count/number

class `awslimitchecker.limit.AwsLimitUsage` (*limit*, *value*, *resource_id=None*, *aws_type=None*)

Bases: `object`

This object describes the usage of an AWS resource, with the capability of containing information about the resource beyond an integer usage.

The simplest case is an account- / region-wide count, such as the number of running EC2 Instances, in which case a simple integer value is sufficient. In this case, the `AwsLimit` would have one instance of this class for the single value.

In more complex cases, such as the “Subnets per VPC”, the limit is applied by AWS on multiple resources (once per VPC). In this case, the `AwsLimit` should have one instance of this class per VPC, so we can determine *which* VPCs have crossed thresholds.

`AwsLimitUsage` objects are comparable based on their numeric `value`.

Parameters

- **limit** (`AwsLimit`) – the `AwsLimit` that this instance describes
- **value** (`int` or `float`) – the numeric usage value
- **resource_id** (`string`) – If there can be multiple usage values for one limit, an AWS ID for the resource this instance describes

- **aws_type** (*string*) – if `id` is not `None`, the AWS resource type that ID represents. As a convention, we use the AWS Resource Type names used by [CloudFormation](#) # noqa

```
__dict__ = dict_proxy({'__ne__': <function __ne__ at 0x7f50b522bde8>, '__module__': 'awslimitchecker.limit', '__wea
```

```
__eq__ (other)
```

```
__ge__ (other)
```

```
__gt__ (other)
```

```
__init__ (limit, value, resource_id=None, aws_type=None)
```

This object describes the usage of an AWS resource, with the capability of containing information about the resource beyond an integer usage.

The simplest case is an account- / region-wide count, such as the number of running EC2 Instances, in which case a simple integer value is sufficient. In this case, the *AwsLimit* would have one instance of this class for the single value.

In more complex cases, such as the “Subnets per VPC”, the limit is applied by AWS on multiple resources (once per VPC). In this case, the *AwsLimit* should have one instance of this class per VPC, so we can determine *which* VPCs have crossed thresholds.

AwsLimitUsage objects are comparable based on their numeric `value`.

Parameters

- **limit** (*AwsLimit*) – the *AwsLimit* that this instance describes
- **value** (*int* or *float*) – the numeric usage value
- **resource_id** (*string*) – If there can be multiple usage values for one limit, an AWS ID for the resource this instance describes
- **aws_type** (*string*) – if `id` is not `None`, the AWS resource type that ID represents. As a convention, we use the AWS Resource Type names used by [CloudFormation](#) # noqa

```
__lt__ (other)
```

```
__module__ = 'awslimitchecker.limit'
```

```
__ne__ (other)
```

```
__str__ ()
```

Return a string representation of this object.

If `id` is not set, return `value` formatted as a string; otherwise, return a string of the format `id=value`.

Return type *string*

```
__weakref__
```

list of weak references to the object (if defined)

```
get_value ()
```

Get the current usage value

Returns current usage value

Return type *int* or *float*

```
awslimitchecker.limit.SOURCE_DEFAULT = 0
```

indicates a limit value that came from hard-coded defaults in *awslimitchecker*

```
awslimitchecker.limit.SOURCE_OVERRIDE = 1
```

indicates a limit value that came from user-defined limit overrides

`awslimitchecker.limit.SOURCE_TA = 2`
 indicates a limit value that came from Trusted Advisor

awslimitchecker.runner module

class `awslimitchecker.runner.Runner`

Bases: `object`

`__dict__` = `dict_proxy({'__module__': 'awslimitchecker.runner', 'color_output': <function color_output at 0x7f50b4cf>})`

`__init__` ()

`__module__` = 'awslimitchecker.runner'

`__weakref__`

list of weak references to the object (if defined)

`check_thresholds` ()

`color_output` (*s*, *color*)

`console_entry_point` ()

`iam_policy` ()

`list_defaults` ()

`list_limits` ()

`list_services` ()

`parse_args` (*argv*)

parse arguments/options

Parameters `argv` (*list*) – argument list to parse, usually `sys.argv[1:]`

Returns parsed arguments

Return type `argparse.Namespace`

`print_issue` (*service_name*, *limit*, *crits*, *warns*)

Parameters

- **service_name** (*str*) – the name of the service
- **limit** (*AWSLimit*) – the Limit this relates to
- **crits** – the specific usage values that crossed the critical threshold
- **warns** – the specific usage values that crossed the warning threshold

`set_limit_overrides` (*overrides*)

`show_usage` ()

`awslimitchecker.runner.console_entry_point` ()

awslimitchecker.trustedadvisor module

class `awslimitchecker.trustedadvisor.TrustedAdvisor`

Bases: `object`

Class to contain all TrustedAdvisor-related logic.

`__dict__ = dict_proxy({'__module__': 'awslimitchecker.trustedadvisor', 'get_limit_check_id': <function get_limit_c`

`__init__()`

Class to contain all TrustedAdvisor-related logic.

`__module__ = 'awslimitchecker.trustedadvisor'`

`__weakref__`

list of weak references to the object (if defined)

`get_limit_check_id()`

Query currently-available TA checks, return the check ID and metadata of the 'performance/Service Limits' check.

Returns 2-tuple of Service Limits TA check ID (string), metadata (list), or (None, None).

Return type tuple

`poll()`

Poll Trusted Advisor (Support) API for limit checks.

Return a dict of service name (string) keys to nested dict vals, where each key is a limit name and each value the current numeric limit.

e.g.:

```

{
    'EC2': {
        'SomeLimit': 10,
    }
}
```

`update_services(ta_results, services)`

Given a dict of TrustedAdvisor check results from `poll()` and a dict of Service objects passed in to `update_limits()`, updated the TrustedAdvisor limits for all services.

Parameters

- **ta_results** (*dict*) – results returned by `poll()`
- **services** (*dict*) – dict of service names to `_AwsService` objects

`connect()`

`update_limits(services)`

Poll 'Service Limits' check results from Trusted Advisor, if possible. Iterate over all `AwsLimit` objects for the given services and update their limits from TA if present in TA checks.

Parameters **services** (*dict*) – dict of service name (string) to `_AwsService` objects

awslimitchecker.utils module

`class awslimitchecker.utils.StoreKeyValuePair(option_strings, dest, nargs=None, const=None, default=None, type=None, choices=None, required=False, help=None, metavar=None)`

Bases: `argparse.Action`

Store key=value options in a dict as {'key': 'value'}.

Supports specifying the option multiple times, but NOT with `nargs`.

See `Action`.

`__call__` (*parser, namespace, values, option_string=None*)

`__init__` (*option_strings, dest, nargs=None, const=None, default=None, type=None, choices=None, required=False, help=None, metavar=None*)

`__module__` = 'awslimitchecker.utils'

`awslimitchecker.utils.dict2cols` (*d, spaces=2, separator=' '*)

Take a dict of string keys and string values, and return a string with them formatted as two columns separated by at least spaces number of separator characters.

Parameters

- **d** (*dict*) – dict of string keys, string values
- **spaces** (*int*) – number of spaces to separate columns by
- **separator** (*string*) – character to fill in between columns

awslimitchecker.version module

`class awslimitchecker.version.AWSLimitCheckerVersion` (*release, url, commit=None, tag=None*)

Bases: `object`

`__dict__` = `dict_proxy({'__module__': 'awslimitchecker.version', '__str__': <function __str__ at 0x7f50b6cf5aa0>, ...`

`__init__` (*release, url, commit=None, tag=None*)

`__module__` = 'awslimitchecker.version'

`__repr__` ()

Return a representation of this object that is valid Python and will create an identical `AWSLimitCheckerVersion` object.

Return type `string`

`__str__` ()

Human-readable string representation of this version object, in the format: "version_str <url>"

Return type `string`

`__weakref__`

list of weak references to the object (if defined)

`version_str`

The version string for the currently-running awslimitchecker; includes git branch and tag information.

Return type `string`

`awslimitchecker.version._get_version_info` ()

Returns the currently-installed awslimitchecker version, and a best-effort attempt at finding the origin URL and commit/tag if installed from an editable git clone.

Returns awslimitchecker version

Return type `string`

awslimitchecker.versioncheck module

`class awslimitchecker.versioncheck.AGPLVersionChecker`

Bases: `object`

`__dict__ = dict_proxy({'__dict__': <attribute '__dict__' of 'AGPLVersionChecker' objects>, '__module__': 'awslimitc`

`__module__ = 'awslimitchecker.versioncheck'`

`__weakref__`

list of weak references to the object (if defined)

`__dist_version_url (dist)`

Get version and homepage for a `pkg_resources.Distribution`

Parameters `dist` – the `pkg_resources.Distribution` to get information for

Returns 2-tuple of (version, homepage URL)

Return type `tuple`

`__find_git_info ()`

Find information about the git repository, if this file is in a clone.

Returns information about the git clone

Return type `dict`

`__find_pip_info ()`

Try to find information about the installed awslimitchecker from pip. This should be wrapped in a try/except.

Returns information from pip about 'awslimitchecker'

Return type `dict`

`__find_pkg_info ()`

Find information about the installed awslimitchecker from `pkg_resources`.

Returns information from `pkg_resources` about 'awslimitchecker'

Return type `dict`

`__is_git_dirty ()`

Determine if the git clone has uncommitted changes or is behind origin

Returns True if clone is dirty, False otherwise

Return type `bool`

`find_package_version ()`

Find the installed version of the specified package, and as much information about it as possible (source URL, git ref or tag, etc.)

This attempts, to the best of our ability, to find out if the package was installed from git, and if so, provide information on the origin of that git repository and status of the clone. Otherwise, it uses pip and `pkg_resources` to find the version and homepage of the installed distribution.

This class is not a sure-fire method of identifying the source of the distribution or ensuring AGPL compliance; it simply helps with this process `_iff_` a modified version is installed from an editable git URL `_and_` all changes are pushed up to the publicly-visible origin.

Returns a dict with keys 'version', 'tag', 'commit', and 'url'. Values are strings or None.

Parameters `package_name (str)` – name of the package to find information for

Returns information about the installed version of the package

Return type `dict`

`awslimitchecker.versioncheck.__check_output (args, stderr=None)`

Python version compatibility wrapper for `subprocess.check_output`

Parameters `stderr` – what to do with STDERR - None or an appropriate argument to `subprocess.check_output / subprocess.Popen`

Raises `subprocess.CalledProcessError`

Returns command output

Return type `string`

`awslimitchecker.versioncheck._get_git_commit()`

Get the current (short) git commit hash of the current directory.

Returns short git hash

Return type `string`

`awslimitchecker.versioncheck._get_git_tag(commit)`

Get the git tag for the specified commit, or None

Parameters `commit` (*string*) – git commit hash to get the tag for

Returns tag name pointing to commit

Return type `string`

`awslimitchecker.versioncheck._get_git_url()`

Get the origin URL for the git repository.

Returns repository origin URL

Return type `string`

7.9 Getting Help

7.9.1 Reporting Bugs and Asking Questions

Questions, bug reports and feature requests are happily accepted via the [GitHub Issue Tracker](#). Pull requests are welcome. Issues that don't have an accompanying pull request will be worked on as my time and priority allows, but I'll do my best to complete feature requests as quickly as possible. Please take into account that I work on this project solely in my personal time, I don't get paid to work on it and I can't work on it for my day job, so there may be some delay in getting things implemented.

7.9.2 IRC

I can often be found on [freenode](#) as “jantman”.

Indices and tables

- [genindex](#)
- [modindex](#)
- [search](#)

8.1 License

awslimitchecker is licensed under the [GNU Affero General Public License, version 3 or later](#). This shouldn't be much of a concern to most people; see *Development / AGPL* for more information.

a

- awslimitchecker, 33
- awslimitchecker.checker, 43
- awslimitchecker.limit, 46
- awslimitchecker.runner, 51
- awslimitchecker.services, 33
- awslimitchecker.services.autoscaling, 33
- awslimitchecker.services.base, 34
- awslimitchecker.services.ebs, 36
- awslimitchecker.services.ec2, 37
- awslimitchecker.services.elasticache, 39
- awslimitchecker.services.elb, 40
- awslimitchecker.services.rds, 40
- awslimitchecker.services.vpc, 42
- awslimitchecker.trustedadvisor, 51
- awslimitchecker.utils, 52
- awslimitchecker.version, 53
- awslimitchecker.versioncheck, 53

Symbols

- `_AutoscalingService` (class in `awslimitchecker.services.autoscaling`), 33
- `_AwsService` (class in `awslimitchecker.services.base`), 34
- `_EbsService` (class in `awslimitchecker.services.ebs`), 36
- `_Ec2Service` (class in `awslimitchecker.services.ec2`), 37
- `_ElastiCacheService` (class in `awslimitchecker.services.elasticache`), 39
- `_ElbService` (class in `awslimitchecker.services.elb`), 40
- `_RDSService` (class in `awslimitchecker.services.rds`), 40
- `_VpcService` (class in `awslimitchecker.services.vpc`), 42
- `__abstractmethods__` (`awslimitchecker.services.autoscaling._AutoscalingService` attribute), 34
- `__abstractmethods__` (`awslimitchecker.services.base._AwsService` attribute), 34
- `__abstractmethods__` (`awslimitchecker.services.ebs._EbsService` attribute), 36
- `__abstractmethods__` (`awslimitchecker.services.ec2._Ec2Service` attribute), 37
- `__abstractmethods__` (`awslimitchecker.services.elasticache._ElastiCacheService` attribute), 39
- `__abstractmethods__` (`awslimitchecker.services.elb._ElbService` attribute), 40
- `__abstractmethods__` (`awslimitchecker.services.rds._RDSService` attribute), 41
- `__abstractmethods__` (`awslimitchecker.services.vpc._VpcService` attribute), 42
- `__call__()` (`awslimitchecker.utils.StoreKeyValuePair` method), 52
- `__dict__` (`awslimitchecker.checker.AwsLimitChecker` attribute), 43
- `__dict__` (`awslimitchecker.limit.AwsLimit` attribute), 47
- `__dict__` (`awslimitchecker.limit.AwsLimitUsage` attribute), 50
- `__dict__` (`awslimitchecker.runner.Runner` attribute), 51
- `__dict__` (`awslimitchecker.services.base._AwsService` attribute), 34
- `__dict__` (`awslimitchecker.trustedadvisor.TrustedAdvisor` attribute), 51
- `__dict__` (`awslimitchecker.version.AWSLimitCheckerVersion` attribute), 53
- `__dict__` (`awslimitchecker.versioncheck.AGPLVersionChecker` attribute), 53
- `__eq__()` (`awslimitchecker.limit.AwsLimitUsage` method), 50
- `__ge__()` (`awslimitchecker.limit.AwsLimitUsage` method), 50
- `__gt__()` (`awslimitchecker.limit.AwsLimitUsage` method), 50
- `__init__()` (`awslimitchecker.checker.AwsLimitChecker` method), 43
- `__init__()` (`awslimitchecker.limit.AwsLimit` method), 47
- `__init__()` (`awslimitchecker.limit.AwsLimitUsage` method), 50
- `__init__()` (`awslimitchecker.runner.Runner` method), 51
- `__init__()` (`awslimitchecker.services.base._AwsService` method), 34
- `__init__()` (`awslimitchecker.trustedadvisor.TrustedAdvisor` method), 52
- `__init__()` (`awslimitchecker.utils.StoreKeyValuePair` method), 53
- `__init__()` (`awslimitchecker.version.AWSLimitCheckerVersion` method), 53
- `__lt__()` (`awslimitchecker.limit.AwsLimitUsage` method), 50
- `__metaclass__` (`awslimitchecker.services.base._AwsService` attribute), 35
- `__module__` (`awslimitchecker.checker.AwsLimitChecker` attribute), 43
- `__module__` (`awslimitchecker.limit.AwsLimit` attribute), 47
- `__module__` (`awslimitchecker.limit.AwsLimitUsage` attribute), 50

tribute), 50

__module__ (awslimitchecker.runner.Runner attribute), 51

__module__ (awslimitchecker.services.autoscaling._AutoscalingService attribute), 34

__module__ (awslimitchecker.services.base._AwsService attribute), 35

__module__ (awslimitchecker.services.ebs._EbsService attribute), 36

__module__ (awslimitchecker.services.ec2._Ec2Service attribute), 37

__module__ (awslimitchecker.services.elasticache._ElasticCacheService attribute), 39

__module__ (awslimitchecker.services.elb._ElbService attribute), 40

__module__ (awslimitchecker.services.rds._RDSService attribute), 41

__module__ (awslimitchecker.services.vpc._VpcService attribute), 42

__module__ (awslimitchecker.trustedadvisor.TrustedAdvisor attribute), 52

__module__ (awslimitchecker.utils.StoreKeyValuePair attribute), 53

__module__ (awslimitchecker.version.AWSLimitCheckerVersion attribute), 53

__module__ (awslimitchecker.versioncheck.AGPLVersionChecker attribute), 54

__ne__() (awslimitchecker.limit.AwsLimitUsage method), 50

__repr__() (awslimitchecker.version.AWSLimitCheckerVersion method), 53

__str__() (awslimitchecker.limit.AwsLimitUsage method), 50

__str__() (awslimitchecker.version.AWSLimitCheckerVersion method), 53

__weakref__ (awslimitchecker.checker.AwsLimitChecker attribute), 43

__weakref__ (awslimitchecker.limit.AwsLimit attribute), 47

__weakref__ (awslimitchecker.limit.AwsLimitUsage attribute), 50

__weakref__ (awslimitchecker.runner.Runner attribute), 51

__weakref__ (awslimitchecker.services.base._AwsService attribute), 35

__weakref__ (awslimitchecker.trustedadvisor.TrustedAdvisor attribute), 52

__weakref__ (awslimitchecker.version.AWSLimitCheckerVersion attribute), 53

__weakref__ (awslimitchecker.versioncheck.AGPLVersionChecker attribute), 54

_abc_cache (awslimitchecker.services.autoscaling._AutoscalingService attribute), 34

_abc_cache (awslimitchecker.services.base._AwsService attribute), 35

_abc_cache (awslimitchecker.services.ebs._EbsService attribute), 36

_abc_cache (awslimitchecker.services.ec2._Ec2Service attribute), 37

_abc_cache (awslimitchecker.services.elasticache._ElasticCacheService attribute), 39

_abc_cache (awslimitchecker.services.elb._ElbService attribute), 40

_abc_cache (awslimitchecker.services.rds._RDSService attribute), 41

_abc_cache (awslimitchecker.services.vpc._VpcService attribute), 42

_abc_negative_cache (awslimitchecker.services.autoscaling._AutoscalingService attribute), 34

_abc_negative_cache (awslimitchecker.services.base._AwsService attribute), 35

_abc_negative_cache (awslimitchecker.services.ebs._EbsService attribute), 36

_abc_negative_cache (awslimitchecker.services.ec2._Ec2Service attribute), 37

_abc_negative_cache (awslimitchecker.services.elasticache._ElasticCacheService attribute), 39

_abc_negative_cache (awslimitchecker.services.elb._ElbService attribute), 40

_abc_negative_cache (awslimitchecker.services.rds._RDSService attribute), 41

_abc_negative_cache (awslimitchecker.services.vpc._VpcService attribute), 42

_abc_negative_cache_version (awslimitchecker.services.autoscaling._AutoscalingService attribute), 34

_abc_negative_cache_version (awslimitchecker.services.base._AwsService attribute), 35

_abc_negative_cache_version (awslimitchecker.services.ebs._EbsService attribute), 37

_abc_negative_cache_version (awslimitchecker.services.ec2._Ec2Service attribute), 38

_abc_negative_cache_version (awslimitchecker.services.elasticache._ElasticCacheService attribute), 39

_abc_negative_cache_version (awslimitchecker.services.elb._ElbService attribute), 40

40					
<code>_abc_negative_cache_version</code>	(awslimitchecker.services.rds._RDSService attribute),	<code>_find_usage_instances()</code>	(awslimitchecker.services.rds._RDSService method),		
41		<code>_find_usage_networking_eips()</code>	(awslimitchecker.services.ec2._Ec2Service method),		
<code>_abc_negative_cache_version</code>	(awslimitchecker.services.vpc._VpcService attribute),	38			
42		<code>_find_usage_networking_eni_sg()</code>	(awslimitchecker.services.ec2._Ec2Service method),		
<code>_abc_registry</code> (awslimitchecker.services.autoscaling._AutoscalingService attribute),	34	38			
<code>_abc_registry</code> (awslimitchecker.services.base._AwsService attribute),	35	<code>_find_usage_networking_sgs()</code>	(awslimitchecker.services.ec2._Ec2Service method),		
<code>_abc_registry</code> (awslimitchecker.services.ebs._EbsService attribute),	37	38			
<code>_abc_registry</code> (awslimitchecker.services.ec2._Ec2Service attribute),	38	<code>_find_usage_nodes()</code>	(awslimitchecker.services.elasticache._ElastiCacheService method),		
<code>_abc_registry</code> (awslimitchecker.services.elasticache._ElastiCacheService attribute),	39	<code>_find_usage_option_groups()</code>	(awslimitchecker.services.rds._RDSService method),		
<code>_abc_registry</code> (awslimitchecker.services.elb._ElbService attribute),	40	41			
<code>_abc_registry</code> (awslimitchecker.services.rds._RDSService attribute),	41	<code>_find_usage_param_groups()</code>	(awslimitchecker.services.rds._RDSService method),		
<code>_abc_registry</code> (awslimitchecker.services.vpc._VpcService attribute),	42	41			
<code>_add_current_usage()</code> (awslimitchecker.limit.AwsLimit method),	47	<code>_find_usage_parameter_groups()</code>	(awslimitchecker.services.elasticache._ElastiCacheService method),		
<code>_check_output()</code> (in module awslimitchecker.versioncheck),	54	39			
<code>_dist_version_url()</code> (awslimitchecker.versioncheck.AGPLVersionChecker method),	54	<code>_find_usage_reserved_instances()</code>	(awslimitchecker.services.rds._RDSService method),		
<code>_find_git_info()</code> (awslimitchecker.versioncheck.AGPLVersionChecker method),	54	41			
<code>_find_pip_info()</code> (awslimitchecker.versioncheck.AGPLVersionChecker method),	54	<code>_find_usage_route_tables()</code>	(awslimitchecker.services.vpc._VpcService method),		
<code>_find_pkg_info()</code> (awslimitchecker.versioncheck.AGPLVersionChecker method),	54	42			
<code>_find_usage_ACLs()</code> (awslimitchecker.services.vpc._VpcService method),	42	<code>_find_usage_security_groups()</code>	(awslimitchecker.services.elasticache._ElastiCacheService method),		
<code>_find_usage_ebs()</code> (awslimitchecker.services.ebs._EbsService method),	37	39			
<code>_find_usage_event_subscriptions()</code> (awslimitchecker.services.rds._RDSService method),	41	<code>_find_usage_security_groups()</code>	(awslimitchecker.services.rds._RDSService method),		
<code>_find_usage_gateways()</code> (awslimitchecker.services.vpc._VpcService method),	42	41			
<code>_find_usage_instances()</code> (awslimitchecker.services.ec2._Ec2Service method),	38	<code>_find_usage_snapshots()</code>	(awslimitchecker.services.ebs._EbsService method),		
		37			
		<code>_find_usage_snapshots()</code>	(awslimitchecker.services.rds._RDSService method),		
		41			
		<code>_find_usage_subnet_groups()</code>	(awslimitchecker.services.elasticache._ElastiCacheService method),		
		39			
		<code>_find_usage_subnet_groups()</code>	(awslimitchecker.services.rds._RDSService method),		
		41			
		<code>_find_usage_subnets()</code>	(awslimitchecker.services.vpc._VpcService method),		
		42			
		<code>_find_usage_vpces()</code>	(awslimitchecker.services.vpc._VpcService method),		
		42			

- `_get_git_commit()` (in module `awslimitchecker.versioncheck`), 55
 - `_get_git_tag()` (in module `awslimitchecker.versioncheck`), 55
 - `_get_git_url()` (in module `awslimitchecker.versioncheck`), 55
 - `_get_limit_check_id()` (`awslimitchecker.trustedadvisor.TrustedAdvisor` method), 52
 - `_get_limits_ebs()` (`awslimitchecker.services.ebs._EbsService` method), 37
 - `_get_limits_instances()` (`awslimitchecker.services.ec2._Ec2Service` method), 38
 - `_get_limits_networking()` (`awslimitchecker.services.ec2._Ec2Service` method), 38
 - `_get_reserved_instance_count()` (`awslimitchecker.services.ec2._Ec2Service` method), 38
 - `_get_thresholds()` (`awslimitchecker.limit.AwsLimit` method), 47
 - `_get_version_info()` (in module `awslimitchecker.version`), 53
 - `_instance_types()` (`awslimitchecker.services.ec2._Ec2Service` method), 38
 - `_instance_usage()` (`awslimitchecker.services.ec2._Ec2Service` method), 38
 - `_is_git_dirty()` (`awslimitchecker.versioncheck.AGPLVersionChecker` method), 54
 - `_poll()` (`awslimitchecker.trustedadvisor.TrustedAdvisor` method), 52
 - `_reset_usage()` (`awslimitchecker.limit.AwsLimit` method), 48
 - `_set_ta_limit()` (`awslimitchecker.limit.AwsLimit` method), 48
 - `_set_ta_limit()` (`awslimitchecker.services.base._AwsService` method), 35
 - `_update_services()` (`awslimitchecker.trustedadvisor.TrustedAdvisor` method), 52
- A**
- `AGPLVersionChecker` (class in `awslimitchecker.versioncheck`), 53
 - `AwsLimit` (class in `awslimitchecker.limit`), 46
 - `AwsLimitChecker` (class in `awslimitchecker.checker`), 43
 - `awslimitchecker` (module), 33
 - `awslimitchecker.checker` (module), 43
 - `awslimitchecker.limit` (module), 46
 - `awslimitchecker.runner` (module), 51
 - `awslimitchecker.services` (module), 33
 - `awslimitchecker.services.autoscaling` (module), 33
 - `awslimitchecker.services.base` (module), 34
 - `awslimitchecker.services.ebs` (module), 36
 - `awslimitchecker.services.ec2` (module), 37
 - `awslimitchecker.services.elasticache` (module), 39
 - `awslimitchecker.services.elb` (module), 40
 - `awslimitchecker.services.rds` (module), 40
 - `awslimitchecker.services.vpc` (module), 42
 - `awslimitchecker.trustedadvisor` (module), 51
 - `awslimitchecker.utils` (module), 52
 - `awslimitchecker.version` (module), 53
 - `awslimitchecker.versioncheck` (module), 53
 - `AWSLimitCheckerVersion` (class in `awslimitchecker.version`), 53
 - `AwsLimitUsage` (class in `awslimitchecker.limit`), 49
- C**
- `check_thresholds()` (`awslimitchecker.checker.AwsLimitChecker` method), 43
 - `check_thresholds()` (`awslimitchecker.limit.AwsLimit` method), 48
 - `check_thresholds()` (`awslimitchecker.runner.Runner` method), 51
 - `check_thresholds()` (`awslimitchecker.services.base._AwsService` method), 35
 - `color_output()` (`awslimitchecker.runner.Runner` method), 51
 - `connect()` (`awslimitchecker.services.autoscaling._AutoscalingService` method), 34
 - `connect()` (`awslimitchecker.services.base._AwsService` method), 35
 - `connect()` (`awslimitchecker.services.ebs._EbsService` method), 37
 - `connect()` (`awslimitchecker.services.ec2._Ec2Service` method), 38
 - `connect()` (`awslimitchecker.services.elasticache._ElastiCacheService` method), 39
 - `connect()` (`awslimitchecker.services.elb._ElbService` method), 40
 - `connect()` (`awslimitchecker.services.rds._RDSService` method), 41
 - `connect()` (`awslimitchecker.services.vpc._VpcService` method), 42
 - `connect()` (`awslimitchecker.trustedadvisor.TrustedAdvisor` method), 52
 - `console_entry_point()` (`awslimitchecker.runner.Runner` method), 51
 - `console_entry_point()` (in module `awslimitchecker.runner`), 51

D

dict2cols() (in module awslimitchecker.utils), 53

F

find_package_version() (awslimitchecker.versioncheck.AGPLVersionChecker method), 54

find_usage() (awslimitchecker.checker.AwsLimitChecker method), 43

find_usage() (awslimitchecker.services.autoscaling._AutoscalingService method), 34

find_usage() (awslimitchecker.services.base._AwsService method), 35

find_usage() (awslimitchecker.services.ebs._EbsService method), 37

find_usage() (awslimitchecker.services.ec2._Ec2Service method), 38

find_usage() (awslimitchecker.services.elasticache._ElasticacheService method), 39

find_usage() (awslimitchecker.services.elb._ElbService method), 40

find_usage() (awslimitchecker.services.rds._RDSService method), 41

find_usage() (awslimitchecker.services.vpc._VpcService method), 42

G

get_criticals() (awslimitchecker.limit.AwsLimit method), 48

get_current_usage() (awslimitchecker.limit.AwsLimit method), 48

get_current_usage_str() (awslimitchecker.limit.AwsLimit method), 48

get_limit() (awslimitchecker.limit.AwsLimit method), 48

get_limit_source() (awslimitchecker.limit.AwsLimit method), 49

get_limits() (awslimitchecker.checker.AwsLimitChecker method), 44

get_limits() (awslimitchecker.services.autoscaling._AutoscalingService method), 34

get_limits() (awslimitchecker.services.base._AwsService method), 35

get_limits() (awslimitchecker.services.ebs._EbsService method), 37

get_limits() (awslimitchecker.services.ec2._Ec2Service method), 38

get_limits() (awslimitchecker.services.elasticache._ElasticacheService method), 39

get_limits() (awslimitchecker.services.elb._ElbService method), 40

get_limits() (awslimitchecker.services.rds._RDSService method), 41

get_limits() (awslimitchecker.services.vpc._VpcService method), 42

get_project_url() (awslimitchecker.checker.AwsLimitChecker method), 44

get_required_iam_policy() (awslimitchecker.checker.AwsLimitChecker method), 44

get_service_names() (awslimitchecker.checker.AwsLimitChecker method), 44

get_value() (awslimitchecker.limit.AwsLimitUsage method), 50

get_version() (awslimitchecker.checker.AwsLimitChecker method), 44

get_warnings() (awslimitchecker.limit.AwsLimit method), 49

I

iam_policy() (awslimitchecker.runner.Runner method), 51

L

list_defaults() (awslimitchecker.runner.Runner method), 51

list_limits() (awslimitchecker.runner.Runner method), 51

list_services() (awslimitchecker.runner.Runner method), 51

P

parse_args() (awslimitchecker.runner.Runner method), 51

print_issue() (awslimitchecker.runner.Runner method), 51

R

required_iam_permissions() (awslimitchecker.services.autoscaling._AutoscalingService method), 34

required_iam_permissions() (awslimitchecker.services.base._AwsService method), 36

required_iam_permissions() (awslimitchecker.services.ebs._EbsService method), 37

required_iam_permissions() (awslimitchecker.services.ec2._Ec2Service method), 38

required_iam_permissions() (awslimitchecker.services.elasticache._ElasticacheService method), 39

required_iam_permissions() (awslimitchecker.services.elb._ElbService method), 40

required_iam_permissions() (awslimitchecker.services.rds._RDSService method), 41

required_iam_permissions() (awslimitchecker.services.vpc._VpcService attribute), 42

Runner (class in awslimitchecker.runner), 51

S

service_name (awslimitchecker.services.autoscaling._AutoscalingService attribute), 34

service_name (awslimitchecker.services.base._AwsService attribute), 36

service_name (awslimitchecker.services.ebs._EbsService attribute), 37

service_name (awslimitchecker.services.ec2._Ec2Service attribute), 39

service_name (awslimitchecker.services.elasticache._ElastiCacheService attribute), 40

service_name (awslimitchecker.services.elb._ElbService attribute), 40

service_name (awslimitchecker.services.rds._RDSService attribute), 41

service_name (awslimitchecker.services.vpc._VpcService attribute), 42

set_limit_override() (awslimitchecker.checker.AwsLimitChecker method), 44

set_limit_override() (awslimitchecker.limit.AwsLimit method), 49

set_limit_override() (awslimitchecker.services.base._AwsService method), 36

set_limit_overrides() (awslimitchecker.checker.AwsLimitChecker method), 45

set_limit_overrides() (awslimitchecker.runner.Runner method), 51

set_threshold_override() (awslimitchecker.checker.AwsLimitChecker method), 45

set_threshold_override() (awslimitchecker.limit.AwsLimit method), 49

set_threshold_override() (awslimitchecker.services.base._AwsService method), 36

set_threshold_overrides() (awslimitchecker.checker.AwsLimitChecker method), 45

show_usage() (awslimitchecker.runner.Runner method), 51

SOURCE_DEFAULT (in module awslimitchecker.limit), 50

SOURCE_OVERRIDE (in module awslimitchecker.limit), 50

SOURCE_TA (in module awslimitchecker.limit), 50

StoreKeyValuePair (class in awslimitchecker.utils), 52

T

TrustedAdvisor (class in awslimitchecker.trustedadvisor), 51

U

update_limits() (awslimitchecker.trustedadvisor.TrustedAdvisor method), 52

V

version_str (awslimitchecker.version.AWSLimitCheckerVersion attribute), 53