
CMS Documentation

Release 1.2.0

The CMS development team

March 22, 2015

1	Introduction	1
1.1	General structure	1
1.2	Services	1
1.3	Security considerations	2
2	Installation	3
2.1	Dependencies	3
2.2	Installing CMS	6
2.3	Updating CMS	6
3	Running CMS	7
3.1	Configuring the DB	7
3.2	Configuring CMS	7
3.3	Running CMS	8
3.4	Recommended setup	9
3.5	Logs	9
4	Creating a contest	11
4.1	Creating a contest from scratch	11
4.2	Creating a contest from the filesystem	11
4.3	Creating a contest from an exported contest	12
5	Configuring a contest	13
5.1	Limitations	13
5.2	Feedback to contestants	13
5.3	Computation of the score	14
5.4	Primary statements	15
5.5	Timezone	15
5.6	User login	16
5.7	USACO-like contests	16
5.8	Extra time and delay time	16
5.9	Programming languages	17
6	Detailed timing configuration	19
6.1	Fixed-window contests	19
6.2	Customized-window contests	20
7	Task types	21
7.1	Introduction	21
7.2	Standard task types	21
8	Score types	25
8.1	Introduction	25

8.2	Standard score types	25
9	Task versioning	27
9.1	Introduction	27
9.2	Datasets	27
10	External contest formats	29
10.1	Italian import format	29
10.2	Polygon format	32
11	RankingWebServer	35
11.1	Description	35
11.2	Managing it	36
11.3	Securing the connection between PS and RWS	37
11.4	Using a proxy	38
11.5	Some final suggestions	39
12	Localization	41
12.1	For developers	41
12.2	For translators	41
13	Troubleshooting	43
13.1	Database	43
13.2	Servers	43
13.3	Sandbox	44
14	Internals	45
14.1	RPC protocol	45
14.2	Backdoor	46

Introduction

CMS (Contest Management System) is a software for organizing programming contests similar to well-known international contests like the IOI (International Olympiad in Informatics). It was written and it received contributions by people involved in the organization of similar contests on a local, national and international level, and it is regularly used for such contests in many different countries. It is meant to be secure, extendable, adaptable to different situations and easy to use.

CMS is a complete, tested and well proved solution for managing a contest. It does not, however, provide tools for the development of the task data belonging to the contest (task statements, solutions, testcases, etc.) or for configuring machines and network resources that host the contest itself. These are responsibility of the contest administrators, although there are tools that help to automate them.

1.1 General structure

The system is organized in a modular way, with different services running (potentially) on different machines, and providing extendability via service replications on several machines.

The state of the contest is wholly kept on a PostgreSQL database. At the moment, there is no way to use other SQL databases, because the Large Object (LO) feature of PostgreSQL is used. It is unlikely that in the future we will target different databases.

As long as the database is operating correctly, all other services can be started and stopped independently without problems. This means that if a machine goes down, then the administrator can quickly replace it with an identical one, which will take its roles (without having to move information from the broken machine). Of course, this also means that if the database goes down, the system is unable to work. In critical contexts, it is the necessary to configure the database redundantly and be prepared to rapidly do a fail-over in case something bad happens. The choice of PostgreSQL as the database to use should ease this part, since there are many different, mature and well-known solutions to provide such redundance and fail-over procedures.

1.2 Services

CMS is composed of several services, that can be run on a single or on many servers. The core services are:

- **LogService**: collects all log messages in a single place;
- **ResourceService**: collects data about the services running on the same server, and takes care of starting all of them with a single command;
- **Checker**: simple heartbeat monitor for all services;
- **EvaluationService**: organizes the queue of the submissions to compile or evaluate on the testcases, and dispatches these jobs to the workers;
- **Worker**: actually runs the jobs in a sandboxed environment;
- **ScoringService**: collects the outcomes of the submissions and computes the score;

- ProxyService: sends the computed scores to the rankings;
- PrintingService: processes files submitted for printing and sends them to a printer;
- ContestWebServer: the webserver that the contestants will be interacting with;
- AdminWebServer: the webserver to control and modify the parameters of the contests.

Finally, RankingWebServer, whose duty is of course to show the ranking. This webserver is - on purpose - separated from the inner core of CMS in order to ease the creation of mirrors and restrict the number of people that can access services that are directly connected to the database.

There are also other services for testing, importing and exporting contests.

Each of the core services is designed to be able to be killed and reactivated in a way that keeps the consistency of data, and does not block the functionalities provided by the other services.

Some of the services can be replicated on several machine: these are ResourceService (designed to be run on every machine), ContestWebServer and Worker.

1.3 Security considerations

With the exception of RWS, there are no cryptographic or authentication schemes between the various services or between the services and the database. Thus, it is mandatory to keep the services on a dedicated network, properly isolating it via firewalls from contestants or other people's computers. This sort of operations, like also preventing contestants from communicating and cheating, is responsibility of the administrator and is not managed by CMS itself.

Installation

2.1 Dependencies

These are our requirements (in particular we highlight those that are not usually installed by default) - previous versions may or may not work:

- build environment for the programming languages allowed in the competition;
- PostgreSQL ≥ 9.0 ;
- GNU compiler collection (in particular the C compiler `gcc`);
- `gettext` ≥ 0.18 ;
- Python $\geq 2.7, < 3.0$;
- `setuptools` ≥ 0.6 ;
- `Tornado` ≥ 2.0 ;
- `Psycogp` ≥ 2.4 ;
- `gevent` ≥ 1.0 ;
- `SQLAlchemy` ≥ 0.7 ;
- `libcg`;
- `psutil` ≥ 0.6 ;
- `netifaces` ≥ 0.5 ;
- `PyCrypto` ≥ 2.3 ;
- `pytz`;
- `six` ≥ 1.1 ;
- `requests` ≥ 1.1 ;
- `werkzeug` ≥ 0.8 ;
- `patool` ≥ 1.7 ;
- `iso-codes`;
- `shared-mime-info`;
- `PyYAML` ≥ 3.10 (only for some importers);
- `BeautifulSoup` ≥ 3.2 (only for running tests);
- `mechanize` ≥ 0.2 (only for running tests);
- `coverage` ≥ 3.4 (only for running tests);
- `mock` ≥ 1.0 (only for running tests);

- **Sphinx** (only for building documentation).
- **TeX Live** (only for printing)
- **pycups** (only for printing)
- **a2ps** (only for printing)
- **PyPDF2** (only for printing)

You will also require a Linux kernel with support for control groups and namespaces. Support has been in the Linux kernel since 2.6.32. Other distributions, or systems with custom kernels, may not have support enabled. At a minimum, you will need to enable the following Linux kernel options: `CONFIG_CGROUPS`, `CONFIG_CGROUP_CPUACCT`, `CONFIG_MEMCG` (previously called as `CONFIG_CGROUP_MEM_RES_CTLR`), `CONFIG_CPUSETS`, `CONFIG_PID_NS`, `CONFIG_IPC_NS`, `CONFIG_NET_NS`. It is anyway suggested to use Linux kernel version at least 3.8.

Then you require the compilation and execution environments for the languages you will use in your contest:

- **GNU compiler collection** (for C, C++ and Java, respectively with executables `gcc`, `g++` and `gcj`);
- **Free Pascal** (for Pascal, with executable `fpc`);
- **Python** ≥ 2.7 , < 3.0 (for Python, with executable `python2`; note though that this must be installed anyway because it is required by CMS itself);
- **PHP** ≥ 5 (for PHP, with executable `php5`).

All dependencies can be installed automatically on most Linux distributions.

2.1.1 Ubuntu

On Ubuntu 14.04, one will need to run the following script to satisfy all dependencies:

```
sudo apt-get install build-essential fpc postgresql postgresql-client \
  gettext python2.7 python-setuptools python-tornado python-psycpg2 \
  python-sqlalchemy python-psutil python-netifaces python-crypto \
  python-tz python-six iso-codes shared-mime-info stl-manual \
  python-beautifulsoup python-mechanize python-coverage python-mock \
  cgroup-lite python-requests python-werkzeug python-gevent patool

# Optional.
# sudo apt-get install nginx-full php5-cli php5-fpm phppgadmin \
#   python-yaml python-sphinx texlive-latex-base python-cups a2ps
# You can install PyPDF2 using Python Package Index.
```

2.1.2 Arch Linux

On Arch Linux, unofficial AUR packages can be found: `cms` or `cms-git`. However, if you don't want to use them, the following command will install almost all dependencies (some of them can be found in the AUR):

```
sudo pacman -S base-devel fpc postgresql postgresql-client python2 \
  setuptools python2-tornado python2-psycpg2 python2-sqlalchemy \
  python2-psutil python2-netifaces python2-crypto python2-pytz \
  python2-six iso-codes shared-mime-info python2-beautifulsoup3 \
  python2-mechanize python2-mock python2-requests python2-werkzeug \
  python2-gevent python2-coverage

# Install the following from AUR.
# https://aur.archlinux.org/packages/libcgroup/
# https://aur.archlinux.org/packages/patool/
# https://aur.archlinux.org/packages/sgi-stl-doc/

# Optional.
```

```
# sudo pacman -S nginx php php-fpm phppgadmin python2-yaml python-sphinx \
#     texlive-core python2-pycups a2ps
# Optionally install the following from AUR.
# https://aur.archlinux.org/packages/python2-pypdf2/
```

2.1.3 Debian

While Debian uses (almost) the same packages as Ubuntu, setting up cgroups is more involved. Debian requires the memory module of cgroups to be activated via a kernel command line parameter. Add `cgroup_enable=memory` to `GRUB_CMDLINE_LINUX_DEFAULT` in `/etc/default/grub` and then run `update-grub`.

Also, we need to mount the cgroup filesystems (under Ubuntu, the `cgroup-lite` package does this). To do this automatically, add the following file into `/etc/init.d`:

```
#!/bin/sh
# /etc/init.d/cgroup

# The following part carries out specific functions depending on arguments.
case "$1" in
start)
    mount -t tmpfs none /sys/fs/cgroup/
    mkdir /sys/fs/cgroup/memory
    mount -t cgroup none /sys/fs/cgroup/memory -o memory
    mkdir /sys/fs/cgroup/cpuacct
    mount -t cgroup none /sys/fs/cgroup/cpuacct -o cpuacct
    mkdir /sys/fs/cgroup/cpuset
    mount -t cgroup none /sys/fs/cgroup/cpuset -o cpuset
    ;;
stop)
    umount /sys/fs/cgroup/cpuset
    umount /sys/fs/cgroup/cpuacct
    umount /sys/fs/cgroup/memory
    umount /sys/fs/cgroup
    ;;
*)
    echo "Usage: /etc/init.d/foobar {start|stop}"
    exit 1
    ;;
esac

exit 0
```

Then execute `chmod 755 /etc/init.d/cgroup` as root and finally `update-rc.d cgroup defaults` to add the script to the default scripts. The following command should now mount the cgroup filesystem:

```
/etc/init.d/cgroup start
```

2.1.4 Python dependencies via pip

If you prefer using Python Package Index, you can retrieve all Python dependencies with this line:

```
sudo pip install -r REQUIREMENTS.txt
```

2.2 Installing CMS

You can download CMS 1.2.0 from [GitHub](#) and extract it on your filesystem. After that, you can install it (recommended, not necessary though):

```
./setup.py build
sudo ./setup.py install
```

If you install CMS, you also need to add your user to the `cmsuser` group and logout to make the change effective:

```
sudo usermod -a -G cmsuser <your user>
```

You can verify to be in the group by issuing the command:

```
groups
```

Warning: Users in the group `cmsuser` will be able to launch the `isolate` program with root permission. They may exploit this to gain root privileges. It is then imperative that no untrusted user is allowed in the group `cmsuser`.

2.3 Updating CMS

As CMS develops, the database schema it uses to represent its data may be updated and new versions may introduce changes that are incompatible with older versions.

To preserve the data stored on the database you need to dump it on the filesystem using `cmsContestExporter` **before you update CMS** (i.e. with the old version).

You can then update CMS and reset the database schema by running:

```
cmsDropDB
cmsInitDB
```

To load the previous data back into the database you can use `cmsContestImporter`: it will adapt the data model automatically on-the-fly (you can use `cmsDumpUpdater` to store the updated version back on disk and speed up future imports).

Running CMS

3.1 Configuring the DB

The first thing to do is to create the user and the database. For PostgreSQL, this is obtained with the following commands (note that the user doesn't need to be a superuser, nor be able to create databases nor roles):

```
sudo su - postgres
createuser cmsuser -P
createdb -O cmsuser database
psql database -c 'ALTER SCHEMA public OWNER TO cmsuser'
psql database -c 'GRANT SELECT ON pg_largeobject TO cmsuser'
```

The last two lines are required to give the PostgreSQL user some privileges which it doesn't have by default, despite being the database owner.

Then you may need to adjust the CMS configuration to contain the correct database parameters. See *Configuring CMS*.

Finally you have to create the database schema for CMS, by running:

```
cmsInitDB
```

Note: If you are going to use CMS services on different hosts from the one where PostgreSQL is running, you also need to instruct it to accept the connections from the services. To do so, you need to change the listening address of PostgreSQL in `postgresql.conf`, for example like this:

```
listen_addresses = '127.0.0.1,192.168.0.x'
```

Moreover, you need to change the HBA (a sort of access control list for PostgreSQL) to accept login requests from outside localhost. Open the file `pg_hba.conf` and add a line like this one:

```
host database cmsuser 192.168.0.0/24 md5
```

3.2 Configuring CMS

There are two configuration files, one for CMS itself and one for the rankings. Samples for both files are in the directory `config/`. You want to copy them to the same file names but without the `.sample` suffix (that is, to `config/cms.conf` and `config/cms.ranking.conf`) before modifying them.

- `cms.conf` is intended to be the same on all servers; all configurations are explained in the file; of particular importance is the definition of `core_services`, that specifies where the services are going to be run, and how many of them, and the connecting line for the database, in which you need to specify the name of the user created above and its password.

- `cms.ranking.conf` is not necessarily meant to be the same on each server that will host a ranking, since it just controls settings relevant for one single server. The addresses and log-in information of each ranking must be the same as the ones found in `cms.conf`.

These files are a pretty good starting point if you want to try CMS. There are some mandatory changes to do though:

- you must change the connection string given in `database`; this usually means to change username, password and database with the ones you chose before;
- if you are running low on disk space, you may want to change `keep_sandbox` to `false`;
- if you want to run CMS without installing it, you need to change `process_cmdline` to reflect that.

If you are organizing a real contest, you must change `secret_key` from the default, and also you will need to think about how to distribute your services and change accordingly `core_services`. Finally, you should change the ranking section of `cms.conf`, and `cms.ranking.conf`, to use a non-trivial username and password.

Warning: As the name implies, the value of `secret_key` must be kept confidential. If a contestant knows it (for example because you are using the default value), they may be easily able to log in as another contestant.

After having modified `cms.conf` and `cms.ranking.conf` in `config/`, you can reinstall CMS in order to make these changes effective, with

```
sudo ./setup.py install
```

3.3 Running CMS

Here we will assume you installed CMS. If not, you should replace all commands path with the appropriate local versions (for example, `cmsLogService` becomes `./scripts/cmsLogService`).

At this point, you should have CMS installed on all the machines you want run services on, with the same configuration file, and a running PostgreSQL instance. To run CMS, you need a contest in the database. To create a contest, follow [these instructions](#).

CMS is composed of a number of services, potentially replicated several times, and running on several machines. You can run all the services by hand, but this is a tedious task. Luckily, there is a service (`ResourceService`) that takes care of starting all the services on the machine it is running, limiting thus the number of binaries you have to run. Services started by `ResourceService` do not show their logs to the standard output; so it is expected that you run `LogService` to inspect the logs as they arrive (logs are also saved to disk). To start `LogService`, you need to issue, in the machine specified in `cms.conf` for `LogService`, this command:

```
cmsLogService 0
```

where `0` is the “shard” of `LogService` you want to run. Since there must be only one instance of `LogService`, it is safe to let CMS infer that the shard you want is the `0`-th, and so an equivalent command is

```
cmsLogService
```

After `LogService` is running, you can start `ResourceService` on each machine involved, instructing it to load all the other services:

```
cmsResourceService -a
```

The flag `-a` informs `ResourceService` that it has to start all other services, and we have omitted again the shard number since, even if `ResourceService` is replicated, there must be only one of it in each machine. If you have a funny network configuration that confuses CMS, just give explicitly the shard number. In any case, `ResourceService` will ask you the contest to load, and will start all the other services. You should start see logs flowing in the `LogService` terminal.

Note that it is your duty to keep CMS’s configuration synchronized among the machines.

3.4 Recommended setup

Of course, the number of servers one needs to run a contest depends on many factors (number of participants, length of the contest, economical issues, more technical matters...). We recommend that, for fairness, each Worker runs on a dedicated machine (i.e., without other CMS services beyond ResourceService).

As for the distribution of services, usually there is one ResourceService for each machine, one instance for each of LogService, ScoringService, Checker, EvaluationService, AdminWebServer, and one or more instances of ContestWebServer and Worker. Again, if there are more than one Worker, we recommend to run them on different machines.

We suggest and support out-of-the-box using CMS over Ubuntu 14.04. Yet, CMS can be successfully run on different Linux distributions. Non-Linux operating systems are not supported.

You can replicate the service handling the contestant-facing web server, `cmsContestWebServer`; in this case, you need to configure a load balancer in front of them. We suggest to use nginx for that, and provide a sample configuration for it at `config/nginx.conf.sample` (this file also configures nginx to act as a HTTPS endpoint and to force secure connections, by redirecting HTTP to HTTPS). This file probably needs to be adapted to your distribution if it's not Ubuntu: try to merge it with the file you find installed by default. For additional information see the official nginx [documentation](#) and [examples](#). Note that without the `ip_hash` option some features might not always work as expected.

3.5 Logs

When the services are running, log messages are streamed to the log service. This is the meaning of the log levels:

- **debug**: you can ignore them (in the default configuration, the log service does not show them);
- **info**: they inform you on what is going on in the system and that everything is fine;
- **warning**: something went wrong or was slightly unexpected, but CMS knew how to handle it, or someone fed inappropriate data to CMS (by error or on purpose); you may want to check these as they may evolve into errors or unexpected behaviors, or hint that a contestant is trying to cheat;
- **error**: an unexpected condition that should not have happened; you are really encouraged to take actions to fix them, but the service will continue to work (most of the time, ignoring the error and the data connected to it);
- **critical**: a condition so unexpected that the service is really startled and refuses to continue working; you are forced to take action because with high probability the service will continue having the same problem upon restarting.

Warning, error, and critical log messages are also displayed in the main page of AdminWebServer.

Creating a contest

4.1 Creating a contest from scratch

The most immediate (but often less practical) way to create a contest in CMS is using the admin interface. You can start the AdminWebServer using the command `cmsAdminWebServer` (or using the ResourceService).

After that, you can connect to the server using the address and port specified in `cms.conf`; typically, <http://localhost:8889/>.

Here, you can create a contest clicking the “+” next to the drop down on the left. After that, you must add the tasks and the users. Up to now, each of these operations is manual; plus, it is usually more practical to work, for example, on a file specifying the contestants’ details instead of using the web interface.

Luckily, there is another way to create a contest.

4.2 Creating a contest from the filesystem

Our idea is that CMS does not get in the way of how you create your contest and your tasks (unless you want to). We think that every national IOI selection team and every contest administrator has a preferred way of developing the tasks, and of storing their data in the filesystem, and we do not want to change the way you work.

Instead, we provided CMS with tools to import a contest from a custom filesystem description. The command `cmsImporter` reads a filesystem description and creates a new contest from it. The command `cmsReimporter` reads a filesystem description and updates an existing contest. Thus, with reimporting you can update, add or remove users or tasks to a contest without losing the existing submissions (unless, of course, they belong to a task or a user that is being deleted).

In order to make these tools compatible with your filesystem format, you have to write a simple Python module that converts your filesystem description to the internal CMS representation of the contest. This is not a hard task: you just have to write an extension of the class `Loader` in `cmscontrib/BaseLoader.py`, implementing missing methods as required by the docstrings. You can use the loader for the Italian format at `cmscontrib/YamlLoader.py` as a template.

You can also use one of the two formats for which CMS have already a loader.

- The Italian filesystem format supports all the features of CMS, but evolved in a rather messy way and is now full of legacy behaviors and shortcomings. No compatibility in time is guaranteed with this format. If you want to use it anyway, an example of a contest written in this format is in [this GitHub repository](#), while its explanation is [here](#).
- The **Polygon format**, which is the format used in several contests and by Codeforces. Polygon does not support all of CMS features, but having this importer is especially useful if you have a big repository of tasks in this format.

4.3 Creating a contest from an exported contest

This option is not really suited for creating new contests but to store and move contest already used in CMS. If you have the dump of a contest exported from CMS, you can import it with `cmsContestImporter <source>`, where `<source>` is the archive filename or directory of the contest.

Configuring a contest

In the following text “user” and “contestant” are used interchangeably.

Configuration parameters will be referred to using their internal name, but it should always be easy to infer what fields control them in the AWS interface by using their label.

5.1 Limitations

Contest administrators can limit the ability of users to submit submissions and user_tests, by setting the following parameters:

- `max_submission_number / max_user_test_number`

These set, respectively, the maximum number of submissions or user tests that will be accepted for a certain user. Any attempt to send in additional submissions or user tests after that limit has been reached will fail.

- `min_submission_interval / min_user_test_interval`

These set, respectively, the minimum amount of time the user is required to wait after a submission or user test has been submitted before they are allowed to send in new ones. Any attempt to submit a submission or user test before this timeout has expired will fail.

The limits can be set both for individual tasks and for the whole contest. A submission or user test is accepted if it verifies the conditions on both the task *and* the contest. This means that a submission or user test will be accepted if the number of submissions or user tests received so far for that task is strictly less than the task’s maximum number *and* the number of submissions or user tests received so far for the whole contest (i.e. in all tasks) is strictly less than the contest’s maximum number. The same holds for the minimum interval too: a submission or user test will be accepted if the time passed since the last submission or user test for that task is greater than the task’s minimum interval *and* the time passed since the last submission or user test received for the whole contest (i.e. in any of the tasks) is greater than the contest’s minimum interval.

Each of these fields can be left unset to prevent the corresponding limitation from being enforced.

5.2 Feedback to contestants

Each testcase can be marked as public or private. After sending a submission, a contestant can always see its results on the public testcases: a brief passed / partial / not passed status for each testcase, and the partial score that is computable from the public testcases only. Note that input and output data are always hidden.

Tokens were introduced to provide contestants with limited access to the detailed results of their submissions on the private testcases as well. If a contestant uses a token on a submission, then they will be able to see its result on all testcases, and the global score.

5.2.1 Tokens rules

Each contestant have a set of available tokens at their disposal; when they use a token it is taken from this set, and cannot be use again. These sets are managed by CMS according to rules defined by the contest administrators, as explained later in this section.

There are two types of tokens: contest-tokens and task-tokens. When a contestant uses a token to unlock a submission, they are really using one token of each type, and therefore needs to have both available. As the names suggest, contest-tokens are bound to the contest while task-tokens are bound to a specific task. That means that there is just one set of contest-tokens but there can be many sets of task-tokens (precisely one for every task). These sets are controlled independently by rules defined either on the contest or on the task.

A token set can be disabled (i.e. there will never be tokens available for use), infinite (i.e. there will always be tokens available for use) or finite. This setting is controlled by the `token_mode` parameter.

If the token set is finite it can be effectively represented by a non-negative integer counter: its cardinality. When the contest starts (or when the user starts its per-user time-frame, see *USACO-like contests*) the set will be filled with `token_gen_initial` tokens (i.e. the counter is set to `token_gen_initial`). If the set is not empty (i.e. the counter is not zero) the user can use a token. After that, the token is discarded (i.e. the counter is decremented by one). New tokens can be generated during the contest: `token_gen_number` new tokens will be given to the user after every `token_gen_interval` minutes from the start (note that `token_gen_number` can be zero, thus disabling token generation). If `token_gen_max` is set, the set cannot contain more than `token_gen_max` tokens (i.e. the counter is capped at that value). Generation will continue but will be ineffective until the contestant uses a token. Unset `token_gen_max` to disable this limit.

The use of tokens can be limited with `token_max_number` and `token_min_interval`: users cannot use more that `token_max_number` tokens in total (this parameter can be unset), and they have to wait at least `token_min_interval` seconds after they used a token before they can use another one (this parameter can be zero). These have no effect in case of infinite tokens.

Having a finite set of both contest- and task-tokens can be very confusing, for the contestants as well as for the contest administrators. Therefore it is common to limit just one type of tokens, setting the other type to be infinite, in order to make the general token availability depend only on the availability of that type (e.g. if you just want to enforce a contest-wide limit on tokens set the contest-token set to be finite and set all task-token sets to be infinite). CWS is aware of this “implementation details” and when one type is infinite it just shows information about the other type, calling it simply “token” (i.e. removing the “contest-” or “task-” prefix).

Note that “token sets” are “intangible”: they’re just a counter shown to the user, computed dynamically every time. Yet, once a token is used, a Token object will be created, stored in the database and associated with the submission it was used on.

Changing token rules during a contest may lead to inconsistencies. Do so at your own risk!

5.3 Computation of the score

5.3.1 Released submissions

The score of a contestant for the contest is always the sum of the score for each task. The score for a task is the best score among the set of “released” submissions.

Admins can use the configuration “Score mode” in AdminWebServer to change the way CMS defines the set of released submission. There are two ways, corresponding to the rules of IOI 2010-2012 and IOI 2013-.

In the first mode, used in IOI from 2010 to 2012, the released submissions are those on which the contestant used a token, plus the latest one submitted.

In the second mode, used since 2013, the released submissions are all submissions.

Usually, a task using the first mode will have a certain number of private testcases, and a limited sets of tokens. In this situation, you can think that contestants are required to “choose” the submission they want to use for grading, by submitting it last, or by using a token on it.

On the other hand, a task using the second mode usually has all testcases public, and therefore it would be silly to ask contestants to choose the submission (as they would always choose the one with the best score).

5.3.2 Score rounding

Based on the `ScoreTypes` in use and on how they are configured, some submissions may be given a floating-point score. Contest administrators will probably want to show only a small number of these decimal places in the scoreboard. This can be achieved with the `score_precision` fields on the contest and tasks.

The score of a user on a certain task is the maximum among the scores of the “tokened” submissions for that task, and the last one. This score is rounded to a number of decimal places equal to the `score_precision` field of the task. The score of a user on the whole contest is the sum of the *rounded* scores on each task. This score itself is then rounded to a number of decimal places equal to the `score_precision` field of the contest.

Note that some “internal” scores used by `ScoreTypes` (for example the subtask score) are not rounded using this procedure. At the moment the subtask scores are always rounded at two decimal places and there’s no way to configure that (note that the score of the submission is the sum of the *unrounded* scores of the subtasks). That will be changed soon. See [issue #33](#).

The unrounded score is stored in the database (and it’s rounded only at presentation level) so you can change the `score_precision` at any time without having to rescore any submissions. Yet, you have to make sure that these values are also updated on the `RankingWebServers`. To do that you can either restart `ScoringService` or update the data manually (see [RankingWebServer](#) for further information).

5.4 Primary statements

When there are many statements for a certain task (which are often different translations of the same statement) contest administrators may want to highlight some of them to the users. These may include, for example, the “official” version of the statement (the one that is considered the reference version in case of questions or appeals) or the translations for the languages understood by that particular user. To do that the `primary_statements` field of the tasks and the users has to be used.

The `primary_statements` field for the tasks is a JSON-encoded list of strings: it specifies the language codes of the statements that will be highlighted to all users. A valid example is `["en_US", "it"]`. The `primary_statements` field for the users is a JSON-encoded object of lists of strings. Each item in this object specifies a task by its name and provides a list of language codes of the statements to highlight. For example `{"task1": ["de"], "task2": ["de_CH"]}`.

Note that users will always be able to access all statements, regardless of the ones that are highlighted. Note also that language codes in the form `xx` or `xx_YY` (where `xx` is an [ISO 639-1 code](#) and `YY` is an [ISO 3166-1 code](#)) will be recognized and presented accordingly. For example `en_AU` will be shown as “English (Australia)”.

5.5 Timezone

CMS stores all times as UTC timestamps and converts them to an appropriate timezone when displaying them. This timezone can be specified on a per-user and per-contest basis with the `timezone` field. It needs to contain a string in the format `Europe/Rome` (actually, any string recognized by `pytz` will work).

When CWS needs to show a timestamp to the user it first tries to show it according to the user’s timezone. If the string defining the timezone is unrecognized (for example it is the empty string), CWS will fallback to the contest’s timezone. If it is again unable to interpret that string it will use the local time of the server.

5.6 User login

Users log into CWS using a username and a password. These have to be specified, respectively, in the `username` and `password` fields (in cleartext!). These credentials need to be inserted by the admins (i.e. there's no way to have an automatic login, a "guest" session, etc.). The user needs to login again if they do not navigate the site for `cookie_duration` seconds (specified in the `cms.conf` file).

In fact, there are other reasons that can cause the login to fail. If the `ip_lock` option (in `cms.conf`) is set to `true` then the login will fail if the IP address that attempted it doesn't match the address or subnet in the `ip` field of the specified user. If `ip` is not set then this check is skipped, even if `ip_lock` is `true`. Note that if a reverse-proxy (like `nginx`) is in use then it is necessary to set `is_proxy_used` (in `cms.conf`) to `true` and configure the proxy in order to properly pass the `X-Forwarded-For`-style headers (see *Recommended setup*).

The login can also fail if `block_hidden_users` (in `cms.conf`) is `true` and the user trying to login as has the `hidden` field set.

5.7 USACO-like contests

One trait of the `USACO` contests is that the contests themselves are many days long but each user is only able to compete for a few hours after their first login (after that they are not able to send any more submissions). This can be done in CMS too, using the `per_user_time` field of contests. If it is unset the contest will behave "normally", that is all users will be able to submit solutions from the contest's beginning until the contest's end. If, instead, `per_user_time` is set to a positive integer value, then a user will only have a limited amount of time. In particular, after they log in, they will be presented with an interface similar to the pre-contest one, with one additional "start" button. Clicking on this button starts the time frame in which the user can compete (i.e. read statements, download attachments, submit solutions, use tokens, send user tests, etc.). This time frame ends after `per_user_time` seconds or when the contest `stop` time is reached, whichever comes first. After that the interface will be identical to the post-contest one: the user won't be able to do anything. See [issue #61](#).

The time at which the user clicks the "start" button is recorded in the `starting_time` field of the user. You can change that to shift the user's time frame (but we suggest to use `extra_time` for that, explained in *Extra time and delay time*) or unset it to make the user able to start its time frame again. Do so at your own risk!

5.8 Extra time and delay time

Contest administrators may want to give some users a short additional amount of time in which they can compete to compensate for an incident (e.g. a hardware failure) that made them unable to compete for a while during the "intended" time frame. That's what the `extra_time` field of the users is for. The time frame in which the user is allowed to compete is expanded by its `extra_time`, even if this would lead the user to be able to submit after the end of the contest.

During extra time the user will continue to receive newly generated tokens. If you don't want them to have more tokens than other contestants, set the `token_max_number` parameter described above to the number of tokens you expect a user to have at their disposal during the whole contest (if it doesn't already have a value less than or equal to this). See also [issue #29](#).

Contest administrators can also alter the competition time of a contestant setting `delay_time`, which has the effect of translating the competition time window for that contestant of the specified number of seconds in the future. Thus, while setting `extra_time` *adds* some times at the end of the contest, setting `delay_time` *moves* the whole time window. As for `extra_time`, setting `delay_time` may extend the contestant time window beyond the end of the contest itself.

Both options have to be set to a non negative number. They can be used together, producing both their effects. Please read *Detailed timing configuration* for a more in-depth discussion of their exact effect.

Note also that submissions sent during the extra time will continue to be considered when computing the score, even if the `extra_time` field of the user is later reset to zero (for example in case the user loses the appeal): you need to completely delete them from the database.

5.9 Programming languages

It is possible to limit the set of programming languages available to contestants by setting the appropriate configuration in the contest page in AWS. By default, the historical set of IOI programming languages is allowed (C, C++, and Pascal). These languages have been used in several contests and with many different types of tasks, and are thus fully tested and safe.

Contestants may be also allowed to use Java, Python and PHP, but these languages have only been tested for Batch tasks, and have not been thoroughly analyzed for potential security and usability issues. Being run under the sandbox, they should be reasonably safe, but, for example, the libraries available to contestants might be hard to control.

5.9.1 Language details

- Pascal support is provided by `fpc`, and submissions are optimized with `-O2`.
- C/C++ support is provided by the GNU Compiler Collection. Submissions are optimized with `-O2`. The standards used by default by CMS are `gnu90` for C (that is, C90 with the GNU extension, the default for `gcc`) and C++11 for C++. Note that C++11 support in `g++` is still incomplete and experimental. Please refer to the [C++11 Support in GCC](#) page for more information.
- Java programs are first compiled using `gcj` (optimized with `-O3`), and then run as normal executables. Proper Java support using a JVM will most probably come in the next CMS version.
- Python submissions are interpreted using Python 2 (you need to have `/usr/bin/python2`).
- PHP submissions are interpreted by `/usr/bin/php5`.

The compilation lines can be inspected and amended in `cms/grading/__init__.py` (there is no way of configuring them apart from changing the source code). Possible amendments are changing the Python version from 2 to 3 (there are instructions in the file on how to do it) or changing the standard used by the GCC.

Detailed timing configuration

This section describes the exact meaning of CMS parameters for controlling the time window allocated to each contestant. Please see *Configuring a contest* for a more gentle introduction and the intended usage of the various parameters.

When setting up a contest, you will need to decide the time window in which contestants will be able to interact with the contest (by reading statements, submit solutions, ...). In CMS there are several parameters that allow to control this time window, and it is also possible to personalize it for each single user in case it is needed.

The first decision to choose among these two possibilities:

1. all contestants will start and end the contest at the same time (unless otherwise decided by the admins during the contest for fairness reasons);
2. each contestant will start the contest at the time they decide.

The first situation is that we will refer to as a fixed-window contest, whereas we will refer to the second situation as customized-window contest.

6.1 Fixed-window contests

These are quite simple to configure: you just need to set `start_time` and `end_time`, and by default all users will be able to interact with the contest between these two instants.

For fairness reasons, during the contest you may want to extend the time window for all or for particular users. In the first case, you just need to change the `end_time` parameter. In the latter case, you can use one of two slightly different per-contestant parameters: `extra_time` and `delay_time`.

You can use `extra_time` to award more time at the end of the contest for a specific contestant, whereas you can use `delay_time` to shift in the future the time window of the contest just for that user. There are two main practical differences between these two options.

1. If you set `extra_time` to S seconds, the contestant will be able to interact with the contest in the first S seconds of it, whereas if you use `delay_time`, they will not, as in the first case the time window is extended, in the second is shifted (if S seconds have already passed from the start of the contest, then there is no difference).
2. If tokens are generated every M minutes, and you set `extra_time` to S seconds, then tokens for that contestants are generated at $start_time + k * M$ (in particular, it might be possible that more tokens are generated for contestants with `extra_time`); if instead you set `delay_time` to S seconds, tokens for that contestants are generated at $start_time + S + k * M$ (i.e., they are shifted from the original, and the same amount of tokens as other contestants will be generated).

Of course it is possible to use both at the same time, but we do not see much value in doing so.

6.2 Customized-window contests

In these contests, contestants can use a time window of fixed length (`per_user_time`), starting from the first time they log in between `start_time` and `end_time`. Moreover, the time window is capped at `end_time` (so if `per_user_time` is 5 hours and a contestant logs in for the first time one minute before `end_time`, they will have just one minute).

Again, admins can change the time windows of specific contestants for fairness reasons. In addition to `extra_time` and `delay_time`, they can also use `starting_time`, which is automatically set by CMS when the contestant logs in for the first time.

The meaning of `extra_time` is to extend both the contestant time window (as defined by `starting_time` + `per_user_time`) and the contest time window (as defined by `end_time`) by the value of `extra_time`, but only for that contestant. Therefore, setting `extra_time` to S seconds effectively allows a contestant to use S seconds more than before (regardless of the time they started the contest).

Again, delay time is similar, but it shifts both contestant and contest time window by that value. The effect on available time similar to that achieved by setting `extra_time`, with the difference explained before in point 1. Also, there is a difference in token generation as explained in point 2 above.

Finally, changing `starting_time` is very similar to changing `delay_time`, but it shifts just the contestant time window, hence if that window was already going over `end_time`, at all effects advancing `starting_time` would not award more time to the contestant, because the end would still be capped at `end_time`. The effect on token generation is the same.

Again, there is probably no need to fiddle with more than one of these three parameters, and our suggestion is to just use `extra_time` or `delay_time` to award more time to a contestant.

Task types

7.1 Introduction

In the CMS terminology, the task type of a task describes how to compile and evaluate the submissions for that task. In particular, they may require additional files called managers, provided by the admins.

A submission goes through two steps involving the task type: the compilation, that usually creates an executable from the submitted files, and the evaluation, that runs this executable against the set of testcases and produces an outcome for each of them.

Note that the outcome doesn't need to be obviously tied to the score for the submission: typically, the outcome is computed by a grader (which is an executable or a program stub passed to CMS) or a comparator (a program that decides if the output of the contestant's program is correct) and not by the task type. Hence, the task type doesn't need to know the meaning of the outcome, which is instead known by the grader and by the *score type*.

7.2 Standard task types

CMS ships with four task types: Batch, OutputOnly, Communication, TwoSteps. The first two are well tested and reasonably strong against cheating attempts and stable with respect to the evaluation times. Communication should be usable but it is less tested than the first two. The last one, TwoSteps, is probably not ready for usage in a public competition. The first two task types cover all but three of the IOI tasks up to IOI 2012.

OutputOnly does not involve programming languages. Batch works with all supported languages (C, C++, Pascal, Java, Python, PHP), but only the first four if you are using a grader. The other task types have not been tested with Java, Python or PHP.

You can configure, for each task, the behavior of these task types on the task's page in AdminWebServer.

7.2.1 Batch

In a Batch task, the contestant submits a single source file, in one of the *allowed programming languages*.

The source file is either standalone or to be compiled with a grader provided by the contest admins. The resulting executable does I/O either on standard input and output or on two files with a specified name. The output produced by the contestant's program is then compared to the correct output either using a simple diff algorithm (that ignores whitespaces) or using a comparator, provided by the admins.

The three choices (standalone or with a grader, standard input and output or files, diff or comparator) are specified through parameters.

If the admins want to provide a grader that takes care of reading the input and writing the output (so that the contestants only need to write one or more functions), they must provide a manager for each allowed language, called `grader.ext`, where `ext` is the standard extension of a source file in that language. If header files for C/C++ or Pascal are needed, they can be provided with names `task_name.h` or `task_namelib.pas`. See the end of the section for specific issues of Java.

If the output is compared with a diff, the outcome will be a float, 0.0 if the output is not correct, 1.0 if it is. If the output is validated by a comparator, you need to provide a manager called `checker`. It must be an executable that:

- is compiled statically (e.g., with `-static` using `gcc` or `g++`);
- takes three filenames as arguments (input, correct output and contestant's output);
- writes on standard output the outcome (that is going to be used by the score type, and is usually a float between 0.0 and 1.0);
- writes on standard error a message to forward to the contestant.

The submission format must contain one filename ending with `.%1`. If there are additional files, the contestants are forced to submit them, the admins can inspect them, but they are not used towards the evaluation.

Batch tasks are supported also for Java, with some requirements. The solutions of the contestants must contain a class named like the short name of the task. A grader must have a class named `grader` that in turns contains the main method; whether in this case the contestants should write a static method or a class is up to the admins.

7.2.2 OutputOnly

In an OutputOnly task, the contestant submits a file for each testcase. Usually, the semantics is that the task specifies a task to be performed on an input file, and the admins provide a set of testcases composed of an input and an output file (as it is for a Batch task). The difference is that, instead of requiring a program that solves the task without knowing the input files, the contestant are required, given the input files, to provide the output files.

There is only one parameter for OutputOnly tasks, namely how correctness of the contestants' outputs is checked. Similarly to the Batch task type, these can be checked using a diff or using a comparator, that is an executable manager named `checker`, with the same properties of the one for Batch tasks.

OutputOnly tasks usually have many uncorrelated files to be submitted. Contestants may submit the first output in a submission, and the second in another submission, but it is easy to forget the first output in the other submission; it is also tedious to add every output every time. Hence, OutputOnly tasks have a feature that, if a submission lacks the output for a certain testcase, the current submission is completed with the most recently submitted output for that testcase (if it exists). This has the effect that contestants can work on a testcase at a time, submitting only what they did from the last submission.

The submission format must contain all the filenames of the form `output_num.txt` where `num` is a three digit decimal number (padded with zeroes, and goes from 0 (included) to the number of testcases (excluded). Again, you can add other files that are stored but ignored. For example, a valid submission format for an OutputOnly task with three testcases is `["output_000.txt", "output_001.txt", "output_002.txt"]`.

7.2.3 Communication

In a Communication task, a contestant must submit a source file implementing a function, similarly to what happens for a Batch task. The difference is that the admins must provide both a stub, that is a source file that is compiled together with the contestant's source, and a manager, that is an executable.

The two programs communicate through two fifo files. The manager receives the name of the two fifos as its arguments. It is supposed to read from standard input the input of the testcase, and to start communicating some data to the other program through the fifo. The two programs exchange data through the fifo, until the manager is able to assign an outcome to the evaluation. The manager then writes to standard output the outcome and to standard error the message to the user.

If the program linked to the user-provided file fails (for a timeout, or for a non-allowed syscall), the outcome is 0.0 and the message describes the problem to the user.

The submission format must contain one filename ending with `.%1`. If there are additional files, the contestants are forced to submit them, the admins can inspect them, but they are not used towards the evaluation.

7.2.4 TwoSteps

Warning: use this task type only if you know what are you doing.

In a TwoSteps task, contestants submit two source files implementing a function each (the idea is that the first function gets the input and compute some data from it with some restriction, and the second tries to retrieve the original data).

The admins must provide a manager compiled together with both files. The resulting executable is run twice (one acting as the computer, one acting as the retriever. The manager in the computer executable must take care of reading the input from standard input; the one in the retriever executable of writing the outcome and the explanation message to standard output and error respectively. Both must take responsibility of the communication between them through a pipe.

More precisely, the executable are called with two arguments: the first is an integer which is 0 if the executable is the computer, and 1 if it is the retriever; the second is the name of the pipe to be used for communication between the processes.

Score types

8.1 Introduction

For every submission, the score type of a task comes into play after the *task type* produced an outcome for each testcase. Indeed, the most important duty of the score type is to describe how to translate the list of outcomes into a single number: the score of the submission. The score type also produces a more informative output for the contestants, and the same information (score and detail) for contestants that did not use a token on the submission. In CMS, these latter set of information is called public, since the contestant can see them without using any tokens.

8.2 Standard score types

Like task types, CMS has the most common score types built in. They are Sum, GroupMin, GroupMul, GroupThreshold.

The first of the four well-tested score types, Sum, is the simplest you can imagine, just assigning a certain amount of points for each correct testcases. The other three are useful for grouping together testcases and assigning points for that group only if some conditions held. Groups are also known as subtasks in some contests. The group score types also allow test cases to be weighted, even for groups of size 1.

Also like task types, the behavior of score types is configurable from the task's page in AdminWebServer.

8.2.1 Sum

This score type interprets the outcome for each testcase as a floating-point number measuring how good the submission was in solving that testcase, where 0.0 means that the submission failed, and 1.0 that it solved the testcase correctly. The score of that submission will be the sum of all the outcomes for each testcase, multiplied by an integer parameter given in the Score type parameter field in AdminWebServer. The parameter field must contain only this integer. The public score is given by the same computation over the public testcases instead of over all testcases.

For example, if there are 20 testcases, 2 of which are public, and the parameter string is 5, a correct solution will score 100 points (20 times 5) out of 100, and its public score will be 10 points (2 times 5) out of 10.

8.2.2 GroupMin

With the GroupMin score type, outcomes are again treated as a measure of correctness, from 0.0 (incorrect) to 1.0 (correct); testcases are split into groups, and each group has an integral multiplier. The score is the sum, over all groups, of the minimum outcome for that group times the multiplier. The public score is computed over all groups in which all testcases within are public.

More precisely, the parameters string for GroupMin is of the form $[[m1, t1], [m2, t2], \dots]$, meaning that the first group comprises the first $t1$ testcases and has multiplier $m1$; the second group comprises the testcases from the $t1 + 1$ to the $t1 + t2$ and has multiplier $m2$; and so on.

8.2.3 GroupMul

GroupMul is almost the same as GroupMin; the only difference is that instead of taking the minimum outcome among the testcases in the group, it takes the product of all outcomes. It has the same behavior as GroupMin when all outcomes are either 0.0 or 1.0.

8.2.4 GroupThreshold

GroupThreshold thinks of the outcomes not as a measure of success, but as an amount of resources used by the submission to solve the testcase. The testcase is then successfully solved if the outcome is between 0.0 and a certain number, the threshold, specified separately for each group.

The parameter string is of the form $[[m1, t1, T1], [m2, t2, T2], \dots]$ where the additional parameter T for each group is the threshold.

Task versioning

9.1 Introduction

Task versioning allows admins to store several sets of parameters for each task at the same time, to decide which are graded and among these the one that is shown to the contestants. This is useful before the contest, to test different possibilities, but especially during the contest to investigate the impact of an error in the task preparation.

For example, it is quite common to realize that one input file is wrong. With task versioning, admins can clone the original dataset (the set of parameters describing the behavior of the task), change the wrong input file with another one, or delete it, launch the evaluation on the new dataset, see which contestants have been affected by the problem, and finally swap the two datasets to make the new one live and visible by the contestants.

The advantages over the situation without task versioning are several:

- there is no need to take down scores during the re-evaluation with the new input;
- it is possible to make sure that the new input works well without showing anything to the contestants;
- if the problem affects just a few contestants, it is possible to notify just them, and the others will be completely unaffected.

9.2 Datasets

A dataset is a version of the sets of parameters of a task that can be changed and tested in background. These parameters are:

- time and memory limits;
- input and output files;
- libraries and graders;
- task type and score type.

Datasets can be viewed and edited in the task page. They can be created from scratch or cloned from existing ones. Of course, during a contest cloning the live dataset is the most used way of creating a new one.

Submissions are evaluated as they arrive against the live dataset and all other datasets with background judging enabled, or on demand when the admins require it.

Each task has exactly one live dataset, whose evaluations and scores are shown to the contestants. To change the live dataset, just click on “Make live” on the desired dataset. Admins will then be prompted with a summary of what changed between the new dataset and the previously active, and can decide to cancel or go ahead, possibly notifying the contestants with a message.

Note: Remember that the summary looks at the scores currently stored for each submission. This means that if you cloned a dataset and changed an input, the scores will still be the old ones: you need to launch a recompilation, reevaluation, or rescoring, depending on what you changed, before seeing the new scores.

After switching live dataset, scores will be resent to RankingWebServer automatically.

External contest formats

There are two different sets of needs that external contest formats strive to satisfy.

- The first is that of contest admins, that for several reasons (storage of old contests, backup, distribution of data) want to export the contest original data (tasks, contestants, ...) together with all data generated during the contest (from the contestants, submissions, user tests, ... and from the system, evaluations, scores, ...). Once a contest has been exported in this format, CMS must be able to reimport it in such a way that the new instance is indistinguishable from the original.
- The second is that of contest creators, that want an environment that helps them design tasks, testcases, and insert the contest data (contestant names and so on). The format needs to be easy to write, understand and modify, and should provide tools to help developing and testing the tasks (automatic generation of testcases, testing of solutions, ...). CMS must be able to import it as a new contest, but also to import it over an already created contest (after updating some data).

CMS provides an exporter `cmsContestExporter` and an importer `cmsContestImporter` working with a format suitable for the first set of needs. This format comprises a dump of all serializable data regarding the contest in a JSON file, together with the files needed by the contest (testcases, statements, submissions, user tests, ...). The exporter and importer understand also compressed versions of this format (i.e., in a zip or tar file). For more information run

```
cmsContestExporter -h
cmsContestImporter -h
```

As for the second set of needs, the philosophy is that CMS should not force upon contest creators a particular environment to write contests and tasks. Therefore, CMS provides two general-purpose commands, `cmsImporter` (for importing a totally new contest) and `cmsReimporter` (for merging an already existing contest with the one being imported). These two programs have no knowledge of any specific on-disk format, so they must be complemented with a set of “loaders”, which actually interpret your files and directories. You can tell the importer or the reimported wich loader to use with the `-L` flag, or just rely and their autodetection capabilities. Running with `-h` flag will list the available loaders.

At the moment, the only loader distributed with CMS understand the format used within Italian Olympiad. It is not particularly suited for general use (see below for some details more), so we encourage you to write a loader for your favorite format and then get in touch with CMS authors to have it accepted in CMS. See files `cmscontrib/BaseLoader.py` and `cmscontrib/YamlLoader.py` for some hints.

10.1 Italian import format

You can follow this description looking at [this example](#). A contest is represented in one directory, containing:

- a YAML file named `contest.yaml`, that describes the general contest properties;
- for each task `task_name`, a directory `task_name` that contains the description of the task and all the files needed to build the statement of the problem, the input and output cases, the reference solution and (when used) the solution checker.

The exact structure of these files and directories is detailed below. Note that this loader is not particularly reliable and providing confusing input to it may lead to create inconsistent or strange data on the database. For confusing input we mean parameters and/or files from which it can infer no or multiple task types or score types.

As the name suggest, this format was born among the Italian trainers group, thus many of the keywords detailed below used to be in Italian. Now they have been translated to English, but Italian keys are still recognized for backward compatibility and are detailed below. Please note that, although so far this is the only format natively supported by CMS, it is far from ideal: in particular, it has grown in a rather untidy manner in the last few years (CMS authors are planning to develop a new, more general and more organic, format, but unfortunately it doesn't exist yet).

For the reasons above, instead of converting your tasks to the Italian format for importing into CMS, it is suggested to write a loader for the format you already have. Please get in touch with CMS authors to have support.

Warning: The authors offer no guarantee for future compatibility for this format. Again, if you use it, you do so at your own risk!

10.1.1 General contest description

The `contest.yaml` file is a plain YAML file, with at least the following keys.

- `name` (string; also accepted: `nome_breve`): the contest's short name, used for internal reference (and exposed in the URLs); it has to match the name of the directory that serves as contest root.
- `description` (string; also accepted: `nome`): the contest's name (description), shown to contestants in the web interface.
- `tasks` (list of strings; also accepted: `problemi`): a list of the tasks belonging to this contest; for each of these strings, say `task_name`, there must be a directory called `task_name` in the contest directory, with content as described *below*; the order in this list will be the order of the tasks in the web interface.
- `users` (list of associative arrays; also accepted: `utenti`): each of the elements of the list describes one user of the contest; the exact structure of the record is described *below*.
- `token_mode`: the token mode for the contest, as in *Tokens rules*; it can be disabled, `infinite` or `finite`; if this is not specified, the loader will try to infer it from the remaining token parameters (in order to retain compatibility with the past), but you are not advised to rely on this behavior.

The following are optional keys.

- `start` (integer; also accepted: `inizio`): the UNIX timestamp of the beginning of the contest (copied in the `start` field); defaults to zero, meaning that contest times haven't yet been decided.
- `stop` (integer; also accepted: `fine`): the UNIX timestamp of the end of the contest (copied in the `stop` field); defaults to zero, meaning that contest times haven't yet been decided.
- `per_user_time` (integer): if set, the contest will be USACO-like (as explained in *USACO-like contests*); if unset, the contest will be traditional (not USACO-like).
- `token_*`: additional token parameters for the contest, see *Tokens rules* (the names of the parameters are the same as the internal names described there).
- `max*_number` and `min*_interval` (integers): limitations for the whole contest, see *Limitations* (the names of the parameters are the same as the internal names described there); by default they're all unset.

10.1.2 User description

Each contest user (contestant) is described in one element of the `utenti` key in the `contest.yaml` file. Each record has to contains the following keys.

- `username` (string): obviously, the username.

- `password` (string): obviously as before, the user's password.

The following are optional keys.

- `first_name` (string; also accepted: `nome`): the user real first name; defaults to the empty string.
- `last_name` (string; also accepted: `cognome`): the user real last name; defaults to the value of `username`.
- `ip` (string): the IP address or subnet from which incoming connections for this user are accepted, see *User login*.
- `hidden` (boolean; also accepted: `fake`): when set to true set the `hidden` flag in the user, see *User login*; defaults to false (the case-sensitive *string* `True` is also accepted).

10.1.3 Task directory

The content of the task directory is used both to retrieve the task data and to infer the type of the task.

These are the required files.

- `task.yaml`: this file contains the name of the task and describes some of its properties; its content is detailed *below*; in order to retain backward compatibility, this file can also be provided in the file `task_name.yaml` in the root directory of the *contest*.
- `statement/statement.pdf` (also accepted: `testo/testo.pdf`): the main statement of the problem. It is not yet possible to import several statement associated to different languages: this (only) statement will be imported according to the language specified under the key `primary_language`.
- `input/input%d.txt` and `output/output%d.txt` for all integers `%d` between 0 (included) and `n_input` (excluded): these are of course the input and reference output files.

The following are optional files, that must be present for certain task types or score types.

- `gen/GEN`: in the Italian environment, this file describes the parameters for the input generator: each line not composed entirely by white spaces or comments (comments start with `#` and end with the end of the line) represents an input file. Here, it is used, in case it contains specially formatted comments, to signal that the score type is *GroupMin*. If a line contains only a comment of the form `# ST: score` then it marks the beginning of a new group assigning at most `score` points, containing all subsequent testcases until the next special comment. If the file does not exist, or does not contain any special comments, the task is given the *Sum* score type.
- `sol/grader.%l` (where `%l` here and after means a supported language extension): for tasks of type *Batch*, it is the piece of code that gets compiled together with the submitted solution, and usually takes care of reading the input and writing the output. If one grader is present, the graders for all supported languages must be provided.
- `sol/*.h` and `sol/*.lib.pas`: if a grader is present, all other files in the `sol` directory that end with `.h` or `lib.pas` are treated as auxiliary files needed by the compilation of the grader with the submitted solution.
- `check/checker` (also accepted: `cor/correttore`): for tasks of types *Batch* or *OutputOnly*, if this file is present, it must be the executable that examines the input and both the correct and the contestant's output files and assigns the outcome. It must be a statically linked executable (for example, if compiled from a C or C++ source, the `-static` option must be used) because otherwise the sandbox will prevent it from accessing its dependencies. It is going to be executed on the workers, so it must be compiled for their architecture. If instead the file is not present, a simple diff is used to compare the correct and the contestant's output files.
- `check/manager`: (also accepted: `cor/manager`) for tasks of type *Communication*, this executable is the program that reads the input and communicates with the user solution.
- `sol/stub.%l`: for tasks of type *Communication*, this is the piece of code that is compiled together with the user submitted code, and is usually used to manage the communication with `manager`. Again, all supported languages must be present.

- `att/*`: each file in this folder is added as an attachment to the task, named as the file's filename.

10.1.4 Task description

The task YAML files require the following keys.

- `name` (string; also accepted: `nome_breve`): the name used to reference internally to this task; it is exposed in the URLs.
- `title` (string; also accepted: `nome`): the long name (title) used in the web interface.
- `n_input` (integer): number of test cases to be evaluated for this task; the actual test cases are retrieved from the *task directory*.
- `score_mode`: the score mode for the task, as in *Computation of the score*; it can be `max_tokened_last` (for the legacy behavior), or `max` (for the modern behavior).
- `token_mode`: the token mode for the task, as in *Tokens rules*; it can be `disabled`, `infinite` or `finite`; if this is not specified, the loader will try to infer it from the remaining token parameters (in order to retain compatibility with the past), but you are not advised to rely on this behavior.

The following are optional keys.

- `time_limit` (float; also accepted: `timeout`): the timeout limit for this task in seconds; defaults to no limitations.
- `memory_limit` (integer; also accepted: `memlimit`): the memory limit for this task in megabytes; defaults to no limitations.
- `public_testcases` (string; also accepted: `risultati`): a comma-separated list of test cases (identified by their numbers, starting from 0) that are marked as public, hence their results are available to contestants even without using tokens.
- `token_*`: additional token parameters for the task, see *Tokens rules* (the names of the parameters are the same as the internal names described there).
- `max*_number` and `min*_interval` (integers): limitations for the task, see *Limitations* (the names of the parameters are the same as the internal names described there); by default they're all unset.
- `output_only` (boolean): if set to `True`, the task is created with the *OutputOnly* type; defaults to `False`.

The following are optional keys that must be present for some task type or score type.

- `total_value` (float): for tasks using the *Sum* score type, this is the maximum score for the task and defaults to 100.0; for other score types, the maximum score is computed from the *task directory*.
- `infile` and `outfile` (strings): for *Batch* tasks, these are the file names for the input and output files; default to `input.txt` and `output.txt`; if left empty, `stdin` and `stdout` are used.
- `primary_language` (string): the statement will be imported with this language code; defaults to `it` (Italian), in order to ensure backward compatibility.

10.2 Polygon format

Polygon is a popular platform for the creation of tasks, and a task format, used among others by Codeforces.

Since Polygon doesn't support CMS directly, some task parameters cannot be set using the standard Polygon configuration. The importer reads from an optional file `cms_conf.py` additional configuration specifics to CMS. Additionally, user can add file named `contestants.txt` to allow importing some set of users.

By default, all tasks are batch files, with custom checker and score type is *Sum*. Loaders assumes that checker is `check.cpp` and written with usage of `testlib.h`. It provides customized version of `testlib.h` which allows using Polygon checkers with CMS. Checkers will be compiled during importing the contest. This is important in case the architecture where the loading happens is different from the architecture of the workers.

Polygon (by now) doesn't allow custom contest-wide files, so general contest options should be hard-coded in the loader.

RankingWebServer

11.1 Description

The **RankingWebServer** (RWS for short) is the web server used to show a live scoreboard to the public.

RWS is designed to be completely separated from the rest of CMS: it has its own configuration file, it doesn't use the PostgreSQL database to store its data and it doesn't communicate with other services using the internal RPC protocol (its code is also in a different package: `cmsranking` instead of `cms`). This has been done to allow contest administrators to run RWS in a different location (on a different network) than the core of CMS, if they don't want to expose a public access to their core network on the internet (for security reasons) or if the on-site internet connection isn't good enough to serve a public website.

To start RWS you have to execute `cmsRankingWebServer`.

11.1.1 Configuring it

The configuration file is named `cmsranking.conf` and RWS will search for it in `/usr/local/etc` and in `/etc` (in this order!). In case it's not found in any of these, RWS will use a hard-coded default configuration that can be found in `cmsranking/Config.py`. If RWS is not installed then the `config` directory will also be checked for configuration files (note that for this to work your working directory needs to be root of the repository). In any case, as soon as you start it, RWS will tell you which configuration file it's using.

The configuration file is a JSON object. The most important parameters are:

- `bind_address`

It specifies the address this server will listen on. It can be either an IP address or a hostname (in the latter case the server will listen on all IP addresses associated with that name). Leave it blank or set it to `null` to listen on all available interfaces.

- `http_port`

It specifies which port to bind the HTTP server to. If set to `null` it will be disabled. We suggest to use a high port number (like 8080, or the default 8890) to avoid the need to start RWS as root, and then use a reverse proxy to map port 80 to it (see *Using a proxy* for additional information).

- `https_port`

It specifies which port to bind the HTTPS server to. If set to `null` it will be disabled, otherwise you need to set `https_certfile` and `https_keyfile` too. See *Securing the connection between PS and RWS* for additional information.

- `username` and `password`

They specify the credentials needed to alter the data of RWS. We suggest to set them to long random strings, for maximum security, since you won't need to remember them. `username` cannot contain a colon.

Warning: Remember to change the `username` and `password` every time you set up a RWS. Keeping the default ones will leave your scoreboard open to illegitimate access.

To connect the rest of CMS to your new RWS you need to add its connection parameters to the configuration file of CMS (i.e. `cms.conf`). Note that you can connect CMS to multiple RWSs, each on a different server and/or port. The parameter you need to change is `rankings`, a list of URLs in the form:

```
<scheme>://<username>:<password>@<hostname>:<port>/<prefix>
```

where `scheme` can be either `http` or `https`, `username`, `password` and `port` are the values specified in the configuration file of the RWS and `prefix` is explained in *Using a proxy* (it will generally be blank, otherwise it needs to end with a slash). If any of your RWSs uses the HTTPS protocol you also need to specify the `https_certfile` configuration parameter. More details on this in *Securing the connection between PS and RWS*.

You also need to make sure that RWS is able to keep enough simultaneously active connections by checking that the maximum number of open file descriptors is larger than the expected number of clients. You can see the current value with `ulimit -Sn` (or `-Sa` to see all limitations) and change it with `ulimit -Sn <value>`. This value will be reset when you open a new shell, so remember to run the command again. Note that there may be a hard limit that you cannot overcome (use `-H` instead of `-S` to see it). If that's still too low you can start multiple RWSs and use a proxy to distribute clients among them (see *Using a proxy*).

11.2 Managing it

RWS doesn't use the PostgreSQL database. Instead, it stores its data in `/var/local/lib/cms/ranking` (or whatever directory is given as `lib_dir` in the configuration file) as a collection of JSON files. Thus, if you want to backup the RWS data, just make a copy of that directory. RWS modifies this data in response to specific (authenticated) HTTP requests it receives.

The intended way to get data to RWS is to have the rest of CMS send it. The service responsible for that is ProxyService (PS for short). When PS is started for a certain contest, it will send the data for that contest to all RWSs it knows about (i.e. those in its configuration). This data includes the contest itself (its name, its begin and end times, etc.), its tasks, its users and the submissions received so far. Then it will continue to send new submissions as soon as they are scored and it will update them as needed (for example when a user uses a token). Note that hidden users (and their submissions) will not be sent to RWS.

There are also other ways to insert data into RWS: send custom HTTP requests or directly write JSON files. They are both discouraged but, at the moment, they are the only way to add team information to RWS (see *issue #65*).

11.2.1 Logo, flags and faces

RWS can also display a custom global logo, a flag for each team and a photo ("face") for each user. Again, the only way to add these is to put them directly in the data directory of RWS:

- the logo has to be saved right in the data directory, named "logo" with an appropriate extension (e.g. `logo.png`), with a recommended resolution of 200x160;
- the flag for a team has to be saved in the "flags" subdirectory, named as the team's name with an appropriate extension (e.g. `ITA.png`);
- the face for a user has to be saved in the "faces" subdirectory, named as the user's username with an appropriate extension (e.g. `ITAL.png`).

We support the following extensions: `.png`, `.jpg`, `.gif` and `.bmp`.

11.2.2 Removing data

PS is only able to create or update data on RWS, but not to delete it. This means that, for example, when a user or a task is removed from CMS it will continue to be shown on RWS. To fix this you will have to intervene manually. The `cmsRWSHelper` script is designed to make this operation straightforward. For example, calling `cmsRWSHelper delete user username` will cause the user `username` to be removed from all the RWSs that are specified in `cms.conf`. See `cmsRWSHelper --help` and `cmsRWSHelper action --help` for more usage details.

In case using `cmsRWSHelper` is impossible (for example because no `cms.conf` is available) there are alternative ways to achieve the same result, presented in decreasing order of difficulty and increasing order of downtime needed.

- You can send a hand-crafted HTTP request to RWS (a `DELETE` method on the `/entity_type/entity_id` resource, giving credentials by Basic Auth) and it will, all by itself, delete that object and all the ones that depend on it, recursively (that is, when deleting a task or a user it will delete its submissions and, for each of them, its subchanges).
- You can stop RWS, delete only the JSON files of the data you want to remove and start RWS again. In this case you have to *manually* determine the depending objects and delete them as well.
- You can stop RWS, remove *all* its data (either by deleting its data directory or by starting RWS with the `--drop` option), start RWS again and restart PS for the contest you're interested in, to have it send the data again.

Note: When you change the username of an user, the name of a task or the name of a contest in CMS and then restart PS, that user, task or contest will be duplicated in RWS and you will need to delete the old copy using this procedure.

11.2.3 Multiple contests

Since the data in RWS will persist even after the PS that sent it has been stopped it's possible to have many PS serve the same RWS, one after the other (or even simultaneously). This allows to have many contests inside the same RWS. The users of the contests will be merged by their username: that is, two users of two different contests will be shown as the same user if they have the same username. To show one contest at a time it's necessary to delete the previous one before adding the next one (the procedure to delete an object is the one described in [Removing data](#)).

Keeping the previous contests may seem annoying to contest administrators who want to run many different and independent contests one after the other, but it's indispensable for many-day contests like the IOI.

11.3 Securing the connection between PS and RWS

RWS accepts data only from clients that successfully authenticate themselves using the HTTP Basic Access Authentication. Thus an attacker that wants to alter the data on RWS needs the username and the password to authenticate its request. If they are random (and long) enough the attacker cannot guess them but may eavesdrop the plaintext HTTP request between PS and RWS. Therefore we suggest to use HTTPS, that encrypts the transmission with TLS/SSL, when the communication channel between PS and RWS is not secure.

HTTPS does not only protect against eavesdropping attacks but also against active attacks, like a man-in-the-middle. To do all of this it uses public-key cryptography based on so-called certificates. In our setting RWS has a public certificate (and its private key). PS has access to a copy to the same certificate and can use it to verify the identity of the receiver before sending any data (in particular before sending the username and the password!). The same certificate is then used to establish a secure communication channel.

The general public does not need to use HTTPS, since it is not sending nor receiving any sensitive information. We think the best solution is, for RWS, to listen on both HTTP and HTTPS ports, but to use HTTPS only for private

internal use. Not having final users use HTTPS also allows you to use home-made (i.e. self-signed) certificates without causing apocalyptic warnings in the users' browsers.

Note that users will still be able to connect to the HTTPS port if they discover its number, but that is of no harm. Note also that RWS will continue to accept incoming data even on the HTTP port; simply, PS will not send it.

To use HTTPS we suggest you to create a self-signed certificate, use that as both RWS's and PS's `https_certfile` and use its private key as RWS's `https_keyfile`. If your PS manages multiple RWSs we suggest you to use a different certificate for each of those and to create a new file, obtained by joining all certificates, as the `https_certfile` of PS. Alternatively you may want to use a Certificate Authority to sign the certificates of RWSs and just give its certificate to PS. Details on how to do this follow.

Note: Please note that, while the indications here are enough to make RWS work, computer security is a delicate subject; we urge you to be sure of what you are doing when setting up a contest in which "failure is not an option".

11.3.1 Creating certificates

A quick-and-dirty way to create a self-signed certificate, ready to be used with PS and RWS, is:

```
openssl req -newkey rsa:1024 -nodes -keyform PEM -keyout key.pem \  
            -new -x509 -days 365 -outform PEM -out cert.pem -utf8
```

You will be prompted to enter some information to be included in the certificate. After you do this you'll have two files, `key.pem` and `cert.pem`, to be used respectively as the `https_keyfile` and `https_certfile` for PS and RWS.

Once you have a self-signed certificate you can use it as a CA (Certificate Authority) to sign other certificates. If you have a `ca_key.pem/ca_cert.pem` pair that you want to use to create a `key.pem/cert.pem` pair signed by it, do:

```
openssl req -newkey rsa:1024 -nodes -keyform PEM -keyout key.pem \  
            -new -outform PEM -out cert_req.pem -utf8  
openssl x509 -req -in cert_req.pem -out cert.pem -days 365 \  
            -CA ca_cert.pem -CAkey ca_key.pem -set_serial <serial>  
rm cert_req.pem
```

Where `<serial>` is a number that has to be unique among all certificates signed by a certain CA.

For additional information on certificates see [the official Python documentation on SSL](#).

11.4 Using a proxy

As a security measure, we recommend not to run RWS as root but to run it as an unprivileged user instead. This means that RWS cannot listen on port 80 and 443 (the default HTTP and HTTPS ports) but it needs to listen on ports whose number is higher than or equal to 1024. This is not a big issue, since we can use a reverse proxy to map the default HTTP and HTTPS ports to the ones used by RWS. We suggest you to use nginx, since it has been already proved successfully for this purpose (some users have reported that other software, like Apache, has some issues, probably due to the use of long-polling HTTP requests by RWS).

A reverse proxy is most commonly used to map RWS from a high port number (say 8080) to the default HTTP port (i.e. 80), hence we will assume this scenario throughout this section.

With nginx it's also extremely easy to do some URL mapping. That is, you can make RWS "share" the URL space of port 80 with other servers by making it "live" inside a prefix. This means that you will access RWS using an URL like "`http://myserver/prefix`".

We'll provide here an example configuration file for nginx. This is just the "core" of the file, but other options need to be added in order for it to be complete and usable by nginx. These bits are different on each distribution, so the best is for you to take the default configuration file provided by your distribution and adapt it to contain the following code:

```

http {
    server {
        listen 80;
        location ^~ /prefix/ {
            proxy_pass http://127.0.0.1:8080/;
            proxy_buffering off;
        }
    }
}

```

The trailing slash is needed in the argument of both the `location` and the `proxy_pass` option. The `proxy_buffering` option is needed for the live-update feature to work correctly (this option can be moved into `server` or `http` to give it a larger scope). To better configure how the proxy connects to RWS you can add an `upstream` section inside the `http` module, named for example `rws`, and then use `proxy_pass http://rws/`. This also allows you to use `nginx` as a load balancer in case you have many RWSs.

If you decide to have HTTPS for private internal use only, as suggested above (that is, you want your users to use only HTTP), then it's perfectly fine to keep using a high port number for HTTPS and not map it to port 443, the standard HTTPS port. Note also that you could use `nginx` as an HTTPS endpoint, i.e. make `nginx` decrypt the HTTPS transmission and redirect it, as cleartext, into RWS's HTTP port. This allows to use two different certificates (one by `nginx`, one by RWS directly), although we don't see any real need for this.

The example configuration file provided in *Recommended setup* already contains sections for RWS.

11.4.1 Tuning nginx

If you're setting up a private RWS, for internal use only, and you expect just a handful of clients then you don't need to follow the advices given in this section. Otherwise please read on to see how to optimize `nginx` to handle many simultaneous connections, as required by RWS.

First, set the `worker_processes` option ¹ of the core module to the number of CPU or cores on your machine. Next you need to tweak the `events` module: set the `worker_connections` option ² to a large value, at least the double of the expected number of clients divided by `worker_processes`. You could also set the `use` option ³ to an efficient event-model for your platform (like `epoll` on linux), but having `nginx` automatically decide it for you is probably better. Then you also have to raise the maximum number of open file descriptors. Do this by setting the `worker_rlimit_nofile` option ⁴ of the core module to the same value of `worker_connections` (or greater). You could also consider setting the `keepalive_timeout` option ⁵ to a value like 30s. This option can be placed inside the `http` module or inside the `server` or `location` sections, based on the scope you want to give it.

For more information see the official `nginx` documentation:

11.5 Some final suggestions

The suggested setup (the one that we also used at the IOI 2012) is to make RWS listen on both HTTP and HTTPS ports (we used 8080 and 8443), to use `nginx` to map port 80 to port 8080, to make all three ports (80, 8080 and 8443) accessible from the internet, to make PS connect to RWS via HTTPS on port 8443 and to use a Certificate Authority to generate certificates (the last one is probably an overkill).

At the IOI 2012, we had only one server, running on a 2 GHz machine, and we were able to serve about 1500 clients simultaneously (and, probably, we were limited to this value by a misconfiguration of `nginx`). This is to say that you'll likely need only one public RWS server.

¹ http://wiki.nginx.org/CoreModule#worker_processes

² http://wiki.nginx.org/EventsModule#worker_connections

³ <http://wiki.nginx.org/EventsModule#use>

⁴ http://wiki.nginx.org/CoreModule#worker_rlimit_nofile

⁵ http://wiki.nginx.org/HttpCoreModule#keepalive_timeout

If you're starting RWS on your server remotely, for example via SSH, make sure the `screen` command is your friend :-).

Localization

12.1 For developers

When you change a string in a template or in a web server, you have to generate again the file `cms/server/po/messages.pot`. To do so, run this command from the root of the repository.

```
xgettext -o cms/server/po/messages.pot --language=Python --no-location \
  --keyword=_:1,2 --keyword=N_ --keyword=N_:1,2 --width=79 \
  cms/grading/*.py cms/grading/**/*.py cms/server/*.py \
  cms/server/templates/admin/*.html \
  cms/server/templates/contest/*.html
```

When you have a new translation, or an update of an old translation, you need to update the `.mo` files (the compiled versions of the `.po` files). You can run `./setup.py build` to update all translations (and also do a couple of other things, like compiling the sandbox). Alternatively, run the following inside `cms/server/`.

```
msgfmt po/<code>.po -o mo/<code>/LC_MESSAGES/cms.mo
```

If needed, create the tree. Note that to have the new strings, you need to restart the web server.

12.2 For translators

To begin translating to a new language, run this command, from `cms/server/po/`.

```
msginit --width=79 -l <two_letter_code_of_language>
```

Right after that, open `<code>.po` and fill the information in the header. To translate a string, simply fill the corresponding `msgstr` with the translations. You can also use specialized translation softwares such as `poEdit` and others.

When the developers update the `.pot` file, you do not need to start from scratch. Instead, you can create a new `.po` file that merges the old translated string with the new, to-be-translated ones. The command is the following, run inside `cms/server/po/`.

```
msgmerge --width=79 <code>.po messages.pot > <code>.new.po
```

You can now inspect `<code>.new.po` and, if satisfied, move it to `<code>.po` and finish the translation.

Troubleshooting

Subtle issues with CMS can arise from old versions of libraries or supporting software. Please ensure you are running the minimum versions of each dependency (described in *Dependencies*).

In the next sections we list some known symptoms and their possible causes.

13.1 Database

- *Symptom.* Error message “Cannot determine OID of function lo_create”

Possible cause. Your database must be at least PostgreSQL 8.x to support large objects used by CMS.

- *Symptom.* Exceptions regarding missing database fields or with the wrong type.

Possible cause. The version of CMS that created the schema in your database is different from the one you are using now. If the schema is older than the current version, you can update the schema as in *Updating CMS*.

- *Symptom.* Some components of CMS fail randomly and PostgreSQL complains about having too many connections.

Possible cause. The default configuration of PostgreSQL may allow insufficiently many incoming connections on the database engine. You can raise this limit by tweaking the `'max_connections'` parameter in `'postgresql.conf'` (see docs). This, in turn, requires more shared memory for the PostgreSQL process (see `'shared_buffers'` parameter in docs), which may overflow the maximum limit allowed by the operating system. In such case see the suggestions in <http://www.postgresql.org/docs/9.1/static/kernel-resources.html#SYSVIPC>. Users reported that another way to go is to use a connection pooler like PgBouncer.

Slightly different, but related, is another issue: CMS may be unable to create new connections to the database because its pool is exhausted. In this case you probably want to modify the `'pool_size'` argument in `cms/db/__init__.py` or try to spread your users over more instances of ContestWebServer.

13.2 Servers

- *Symptom.* Message from ContestWebServer such as: `WARNING:root:Invalid cookie signature KFZzdW5kdWRlCnAwCkxxMzI5MzQzNzIwCnRw...`

Possible cause. The contest secret key (defined in `cms.conf`) may have been changed and users' browsers are still attempting to use cookies signed with the old key. If this is the case, the problem should correct itself and won't be seen by users.

- *Symptom.* Ranking Web Server displays wrong data, or too much data.

Possible cause. RWS is designed to handle groups of contests, so it retains data about past contests. If you want to delete previous data, run RWS with the `-d` option. See [RankingWebServer](#) for more details

13.3 Sandbox

- *Symptom.* The Worker fails to evaluate a submission logging about an invalid (empty) output from the manager.

Possible cause. You might have been used a non-statically linked checker. The sandbox prevent dynamically linked executables to work. Try compiling the checker with `-static`. Also, make sure that the checker was compiled for the architecture of the workers (e.g., 32 or 64 bits).

- *Symptom.* The Worker fails to evaluate a submission with a generic failure.

Possible cause. Make sure that the isolate binary that CMS is using has the correct permissions (in particular, its owner is root and it has the suid bit set). Be careful of having multiple isolate binaries in your path. Another reason could be that you are using an old version of isolate.

- *Symptom.* Contestants' solutions fail when trying to write large outputs.

Possible cause. CMS limits the maximum output size from programs being evaluated for security reasons. Currently the limit is 1 GB and can be configured by changing the parameter `max_file_size` in `cms.conf`.

Internals

This section contains some details about some CMS internals. They are mostly meant for developers, not for users. However, if you are curious about what's under the hood, you will find something interesting here (though without any pretension of completeness). Moreover, these are not meant to be full specifications, but only useful notes for the future.

Oh, I was nearly forgetting: if you are curious about what happens inside CMS, you may actually be interested in helping us writing it. We can assure you it is a very rewarding task. After all, if you are hanging around here, you must have some interest in coding! In case, feel free [to get in touch with us](#).

14.1 RPC protocol

Different CMS processes communicate between them by mean of TCP sockets. Once a service has established a socket with another, it can write messages on the stream; each message is a JSON-encoded object, terminated by a `\r\n` string (this, of course, means that `\r\n` cannot be used in the JSON encoding: this is not a problem, since new lines inside string represented in the JSON have to be escaped anyway).

An RPC request must be of the form (it is pretty printed here, but it is sent in compact form inside CMS):

```
{
  "__method": <name of the requested method>,
  "__data": {
    <name of first arg>: <value of first arg>,
    ...
  },
  "__id": <random ID string>
}
```

The arguments in `__data` are (of course) not ordered: they have to be matched according to their names. In particular, this means that our protocol enables us to use a `kwargs`-like interface, but not a `args`-like one. That's not so terrible, anyway.

The `__id` is a random string that will be returned in the response, and it is useful (actually, it's the only way) to match requests with responses.

The response is of the form:

```
{
  "__data": <return value or null>,
  "__error": <null or error string>,
  "__id": <random ID string>
}
```

The value of `__id` must of course be the same as in the request. If `__error` is not null, then `__data` is expected to be null.

14.2 Backdoor

Setting the `backdoor` configuration key to `true` causes services to serve a Python console (accessible with `netcat`), running in the same interpreter instance as the service, allowing to inspect and modify its data, live. It will be bound to a local UNIX domain socket, usually at `/var/local/run/cms/service_shard`. Access is granted only to users belonging to the `cmsuser` group. Although there's no authentication mechanism to prevent unauthorized access, the restrictions on the file should make it safe to run the backdoor everywhere, even on workers that are used as contestants' machines. You can use `rlwrap` to add basic readline support. For example, the following is a complete working connection command:

```
rlwrap netcat -U /var/local/run/cms/EvaluationService_0
```

Substitute `netcat` with your implementation (`nc`, `ncat`, etc.) if needed.