Table of Contents

PandaOperationalSecurity

Introduction

This page describes some operational security procedures for PanDA.

This documentation is mainly concerned with investigating suspicious user jobs from a site report, passing additional information back to the site and preventing the user for submitting more jobs to panda.

Advice for Sites

Identifying the PanDA Job ID

If you need to identify the panda job ID for a job running in your batch system you can find this directly from the pilot's working directory or by querying the panda monitor:

1. Inside the pilot's running directory (e.g., `/tmp/condorg_hLI26201/pilot3/Panda_Pilot_26241_1274389488`) there will be a directory named, e.g., `PandaJob_1073623690_1274389590`. The middle section is the panda job ID. In this case 1073623690.
2. Inside the same directory in the pilotlog.txt file there is a line `PandaID=NNNNNNNNN`, which gives the panda ID.
3. If you have the batch system ID, you can search directly on the panda monitor for the corresponding jobs. Enter the batch system ID in the "Quick Search: Job" box. e.g., http://panda.cern.ch:25980/server/pandamon/query?job=9708243.svr016.gla.scotgrid.ac.uk

Once you have the panda job ID much more information about this job and its payload is available from the ATLAS PanDA Monitor: http://panda.cern.ch/server/pandamon/query. Enter the job ID in the "Quick Search: Panda job ID" box at the left of the page, or construct the following URL directly: `http://panda.cern.ch/server/pandamon/query?job=PANDAJOBID`.

Advice for AMOD

Note that the site is likely to pass information about the DN of the pilot factory owner, but this will not be the owner of the payload which was submitted. See the ADCoS#Pilot_Factories_and_methods for who submits pilots to which sites.

DDoS
: Reports might come from someone that are not ATLAS sites specifying IP addresses (or host names) where they observe heavy accesses. The hosts can be WNs, or can be **squid servers**

PanDA Job Information

If a site has a security question or other query about a user job run via panda, you need to identify the panda ID. If the site did not pass this to you, then see above to ask them for more information.

You can then pass them the following information to the site:

1. Payload submitter: the `prodUserID` field.
2. Submit host (i.e. the machine from which the job was submitted to panda): the `creationHost` field.
3. Submit time: the `creationTime` field.
4. Payload URL:
    1. For prun type jobs (which execute user scripts directly), in the `jobParameters` field contains the following:
        1. `--sourceURL https://voatlas59.cern.ch:25443` The panda server holding the payload code (panda servers are load balanced, but only one holds the code for each user job)
        2. `-a jobO.6941944f-2e26-4a18-91f0-cb594aac7883.tar.gz` The payload file.
        3. `-p "testme.sh"` The executed script.
        4. To get the payload URL concatenate the server URL with `/cache/` and the payload file. e.g., in the above example the payload URL is `https://voatlas59.cern.ch:25443/cache/jobO.6941944f-2e26-4a18-91f0-cb594aac78`
    2. For other job types:
        1. For other athena jobs the user code payload is contained in the associated build job. The build job is found from the run jobs via the "Associated build job" link just under the first box table on the panda monitor.
        2. The input payload tarball has argument `-i` instead of "-a", otherwise the procedure is the same.

Banning a User

If there is suspicion about a user's jobs then you should ban them from submitting more jobs to panda and cancel their running and queued jobs which they have.

Note that the user's 'name' for PanDA is usually the CN and is stored in the panda database as `prodUserName`. It is case sensitive. User names can also easily be seen here: http://panda.cern.ch:25980/server/pandamon/query?ui=users&days=7.

Setup Environment

1. Login to the panda-server (see https://twiki.cern.ch/twiki/bin/view/Atlas/CentralServiceExpertOnCall#How_to_connect_to_Panda_servers)
2. Sudo to the atlpan account: `sudo -s -u atlpan`.

3. Setup the environment as follows:

```
source /data/atlpan/srv/etc/sysconfig/panda_server-sysconfig
cd /data/atlpan/srv/lib/python2.6/site-packages/pandaserver/test
```

Ban Command

```
python banUser.py --user 'Naughty User'
```

Unban Command

```
python banUser.py --user 'Naughty User' --unban
```

Cancelling Jobs

```
python killUser.py --user 'Naughty User' --jobsetID=all
```

The output should be a list of jobIDs killed.

Note that, for running jobs, this registers the job to be killed the next time the pilot contacts the panda server. It does this every 30mins so there can be a delay of up to 30mins before the jobs is killed. The modificationTime is the last such pilot contact. One can see the cancel job is successful by seeing 'commandToPilot tobekilled'.

N.B. In the case of genuinely malicious code it's trivial for a user to circumvent the pilot kill. Sites should be warned of this.

General Guidelines

Ensure all discussion on security matters involves `atlas-adc-csirt@cern.ch`. This list includes the CERN security team. All security incidents need to be followed up by experts with upmost priority.

Example Email

Here is an example of an email sent back to a site after an enquiry:

```
Dear Worried Site

We have investigated. The DN you give is from a pilot factory (voatlas61.cern.ch) host here at CE
This particular job picked up a payload submitted to panda from this DN:

/O=somegrid/O=users/O=site/CN=Naughty User/CN=proxy

You may see some job details here:

http://panda.cern.ch:25980/server/pandamon/query?job=4343069.lcgbatch01.gridpp.rl.ac.uk

This job was submitted to ATLAS panda at 2010-04-28 14:30:30,652 UT from host 212-127-174-213.cab

You may examine the code run by the suspicious job here:

https://voatlas57.cern.ch:25443/cache/jobO.942c862c-44dd-4c7a-861d-815619ed64c4.tar.gz

The job executed this command:

gridssc.sh%20Site%3DRAL-LCG2%20ROC%3DUK

(Unescape the HTML!)
```

```
We will contact the user involved and investigate further.

Cheers

Graeme
```

Major updates:
-- GraemeStewart - 20-May-2010

Responsible: GraemeAStewart

Never reviewed

This topic: PanDA > PandaOperationalSecurity
Topic revision: r12 - 2016-04-24 - TorreWenaus