

Aegis - An SCM system based on transactions

Walter Franzini

walter.franzini@gmail.com

2007-10-27

Outline

- 1 Intro
- 2 The model
- 3 Distributed development
- 4 Supported platforms
- 5 TODO

Terminology

Software Configuration Management is a methodology to control and manage a software development project.

A transaction is a transformation of state which has the properties of atomicity (all or nothing), durability (effects survive failures) and consistency (a correct transformation). The transaction concept is key to the structuring of data management application. The concept may have applicability to programming system in general.

Terminology

***Software Configuration Management** is a methodology to control and manage a software development project.*

A transaction is a transformation of state which has the properties of atomicity (all or nothing), durability (effects survive failures) and consistency (a correct transformation). The transaction concept is key to the structuring of data management application. The concept may have applicability to programming system in general.

History

- Author: Peter Miller
- License: GPL v2 (will be v3)
- First Public release: 1.4 (1991)
- 1997 - Branch support
- 1999 - Distributed development
- 2004 - UUID Introduction for files and changes
- Current stable release: 4.22.2 (released 2007-10-18)

Main Features

- Support for large projects and large development teams.
- Support change sets and atomic commits.
- *Designed* around incremental development.
- *Designed* for repository security (availability, integrity and confidentiality).
- Support multiple repositories


The model

Aegis uses a particular model of the development of software projects. This model has a master source (or **baseline**) of a project, and a team of developers creating changes to be made to this baseline.


When a change is complete, it is integrated with the baseline, to become the new baseline. Each change must be atomic and self-contained, no change is allowed to cause the baseline to cease to work.

Working is defined as passing it's own tests. The tests are considered part of the baseline.

The work-flow



Change States (simplified)



The roles

- developer
- reviewer
- integrator
- administrator

Test Driven Development

- Test-Driven Development (TDD) is a software development technique that involves repeatedly first writing a test case and then implementing only the code necessary to pass the test.
- Tests are collected within Aegis as files with `test` or `manual_test` usage.
- Test exemption:
 - `test_exempt = false` – the change must pass the test.
 - `baseline_test_exempt = false` – the baseline must *fail* the test;
 - `regression_test_exempt` – the change must pass all the tests.

Working in a LAN

- Aegis is *not* Client/Server
- Accessing the same repository from different machines.
- Install aegis on all the machines.
- Share the repository (NFS).
- Shared/unified user identities (pam_ldap with OpenLDAP).

Working on the Internet

- Support for multiple repositories. Aegis is not centralised.
- A copy of the repository is published on the Internet
 - install a copy of Aegis compiled with the SF hack.
 - requires a CGI capable web server.
 - a simple script to replicate the repository
- a remote developer will use aedist to consult/replicate changes from your repository.
- each change set is identified by an UUID
- change set inventories
- aedist(1) –pending
- aedist –missing
- aedist –replay

Working on the Internet, aedist --pending

```
$ aedist --pending -p aegis.cur -f aegis.sf.net -inver 4.23*
Project "aegis.4.23", Pending Change Set Inventory Page 1
http://aegis.sf.net/cgi-bin/aeget/aegis.4.23?inventory+all 2007/09/14
Change      UUID      Description
-----
4.23.D114   6549acc9-db96-4e17-aacb-ceb5acbde2b0 Improve t0246a-walt.sh
4.23.D272   80fef3a6-a400-4flc-b360-4dd3ffdc83f6 refactor time2change for speed
4.23.D312   f0678fca-3388-4f69-b070-40f97c3b81aa Deprecate non method acces
to change
4.23.D313   7db9d63c-8f27-4230-9490-e23d7464cfcf improve time2delta performance
4.23.D328   30601c66-6682-4b86-82e9-ab9ee8e674ab Fix licenses url
4.23.D329   95bb8417-3ce2-4ald-81f7-e213a09eb7dd aelock install, again
4.23.D330   0f3c459a-fece-4f91-a033-3391775273ce Improve the aepconf(5) man page
4.23.D331   7fb4f184-5334-4f98-895b-95alab001d0a file_usage_build vs.
derived files directory style
```

Remote change sets: 1234.

Pending 8 change sets.

Working on the Internet, aedist --pending

```
$ aedist --pending -p aegis.cur -f aegis.sf.net -inver 4.23*
Project "aegis.4.23", Pending Change Set Inventory Page 1
http://aegis.sf.net/cgi-bin/aeget/aegis.4.23?inventory+all 2007/09/14
```

Change	UUID	Description
4.23.D114	6549acc9-db96-4e17-aacb-ceb5acbde2b0	Improve t0246a-walt.sh
4.23.D272	80fef3a6-a400-4flc-b360-4dd3ffdc83f6	refactor time2change for speed
4.23.D312	f0678fca-3388-4f69-b070-40f97c3b81aa	Deprecate non method acces to change
4.23.D313	7db9d63c-8f27-4230-9490-e23d7464cfcf	improve time2delta performance
4.23.D328	30601c66-6682-4b86-82e9-ab9ee8e674ab	Fix licenses url
4.23.D329	95bb8417-3ce2-4ald-81f7-e213a09eb7dd	aelock install, again
4.23.D330	0f3c459a-fece-4f91-a033-3391775273ce	Improve the aepconf(5) man page
4.23.D331	7fb4f184-5334-4f98-895b-95alab001d0a	file_usage_build vs. derived files directory style

Remote change sets: 1234.

Pending 8 change sets.

Working on the Internet, aedist --pending

```
$ aedist --pending -p aegis.cur -f aegis.sf.net -inver 4.23*
Project "aegis.4.23", Pending Change Set Inventory Page 1
http://aegis.sf.net/cgi-bin/aeget/aegis.4.23?inventory+all 2007/09/14
```

Change	UUID	Description
4.23.D114	6549acc9-db96-4e17-aacb-ceb5acbde2b0	Improve t0246a-walt.sh
4.23.D272	80fef3a6-a400-4f1c-b360-4dd3ffdc83f6	refactor time2change for speed
4.23.D312	f0678fca-3388-4f69-b070-40f97c3b81aa	Deprecate non method acces to change
4.23.D313	7db9d63c-8f27-4230-9490-e23d7464cfcf	improve time2delta performance
4.23.D328	30601c66-6682-4b86-82e9-ab9ee8e674ab	Fix licenses url
4.23.D329	95bb8417-3ce2-4a1d-81f7-e213a09eb7dd	aelock install, again
4.23.D330	0f3c459a-fece-4f91-a033-3391775273ce	Improve the aepconf(5) man page
4.23.D331	7fb4f184-5334-4f98-895b-95a1ab001d0a	file_usage_build vs. derived files directory style

Remote change sets: 1234.

Pending 8 change sets.

Working on the Internet, aedist --pending

```
$ aedist --pending -p aegis.cur -f aegis.sf.net -inver 4.23*
Project "aegis.4.23", Pending Change Set Inventory Page 1
http://aegis.sf.net/cgi-bin/aeget/aegis.4.23?inventory+all 2007/09/14
```

Change	UUID	Description
4.23.D114	6549acc9-db96-4e17-aacb-ceb5acbde2b0	Improve t0246a-walt.sh
4.23.D272	80fef3a6-a400-4f1c-b360-4dd3ffdc83f6	refactor time2change for speed
4.23.D312	f0678fca-3388-4f69-b070-40f97c3b81aa	Deprecate non method acces to change
4.23.D313	7db9d63c-8f27-4230-9490-e23d7464cfcf	improve time2delta performance
4.23.D328	30601c66-6682-4b86-82e9-ab9ee8e674ab	Fix licenses url
4.23.D329	95bb8417-3ce2-4a1d-81f7-e213a09eb7dd	aelock install, again
4.23.D330	0f3c459a-fece-4f91-a033-3391775273ce	Improve the aepconf(5) man page
4.23.D331	7fb4f184-5334-4f98-895b-95a1ab001d0a	file_usage_build vs. derived files directory style

Remote change sets: 1234.

Pending 8 change sets.

Working on the Internet, aedist --missing

```
$ aedist --missing -p aegis.cur -f aegis.sf.net -inver 4.23*  
Project "aegis.4.23", Missing Change Set Inventory  
http://aegis.sf.net/cgi-bin/aeget/aegis.4.23?inventory
```

Page 1
2007/09/11

Change	UUID	Description
4.23.D165	56e43e08-7de7-42e3-b81e-c17d08b16e10	ael cd segfault
4.23.D268	7136d79b-f138-464f-bcfb-bbdebab093f4	rescind 556
		Remote change sets: 1110.
		Missing 2 change sets.

```
$
```

Working on the Internet, aedist --missing

```
$ aedist --missing -p aegis.cur -f aegis.sf.net -inver 4.23*  
Project "aegis.4.23", Missing Change Set Inventory  
http://aegis.sf.net/cgi-bin/aeget/aegis.4.23?inventory
```

Page 1
2007/09/11

Change	UUID	Description
4.23.D165	56e43e08-7de7-42e3-b81e-c17d08b16e10	ael cd segfault
4.23.D268	7136d79b-f138-464f-bcfb-bbdebab093f4	rescind 556
		Remote change sets: 1110.
		Missing 2 change sets.

```
$
```

Working on the Internet, aedist --missing

```
$ aedist --missing -p aegis.cur -f aegis.sf.net -inver 4.23*  
Project "aegis.4.23", Missing Change Set Inventory  
http://aegis.sf.net/cgi-bin/aeget/aegis.4.23?inventory
```

Page 1
2007/09/11

Change	UUID	Description
-----	-----	-----
4.23.D165	56e43e08-7de7-42e3-b81e-c17d08b16e10	ael cd segfault
4.23.D268	7136d79b-f138-464f-bcfb-bbdebab093f4	rescind 556
		Remote change sets: 1110.
		Missing 2 change sets.

\$

Working on the Internet, `aedist --replay`

`aedist --replay` is *not* a synchronization mechanism. The usual model is applied also on changes from remote repositories.

A developer use `aedist (1) --replay` to fetch missing change-sets from a remote repository.

```
$ aedist --replay -p aegis.cur -f aegis.sf.net -inver 4.23*
aedist: aedist --receive -project=aegis.4.23 -change=283
      -file 'http://aegis.sf.net/cgi-bin/aeget/aegis.4.23.C283/?aedist'
aedist: aegis --new-change 283 --project=aegis.4.23 --file=/tmp/aedist-...
aegis: project "aegis.4.23": change 283: created
aedist: aegis --develop-begin 283--project aegis.4.23
project "aegis.4.23": change 283: user "walter" has begun development
aedist: aegis --copy-file --project=aegis.4.23 \
      --change=283 --branch=4.23 --delta=121 \
      libaegis/wide_output/column.cc libaegis/wide_output/wrap.cc
aedist: aegis --copy-file --project=aegis.4.23 --change=283 \
      --branch=4.23 --delta=120 libaegis/output/to_wide.cc
aedist: aegis --new-test --project=aegis.4.23 --change=283 test/02/t0232a-pmil.sh
```

Working on the Internet, `aedist --replay`

`aedist --replay` is *not* a synchronization mechanism. The usual model is applied also on changes from remote repositories.

A developer use `aedist (1) --replay` to fetch missing change-sets from a remote repository.

```
$ aedist --replay -p aegis.cur -f aegis.sf.net -inver 4.23*
aedist: aedist --receive -project=aegis.4.23 -change=283
      -file 'http://aegis.sf.net/cgi-bin/aeget/aegis.4.23.C283/?aedist'
aedist: aegis --new-change 283 --project=aegis.4.23 --file=/tmp/aedist-...
aedist: project "aegis.4.23": change 283: created
aedist: aegis --develop-begin 283--project aegis.4.23
project "aegis.4.23": change 283: user "walter" has begun development
aedist: aegis --copy-file --project=aegis.4.23 \
      --change=283 --branch=4.23 --delta=121 \
      libaegis/wide_output/column.cc libaegis/wide_output/wrap.cc
aedist: aegis --copy-file --project=aegis.4.23 --change=283 \
      --branch=4.23 --delta=120 libaegis/output/to_wide.cc
aedist: aegis --new-test --project=aegis.4.23 --change=283 test/02/t0232a-pmil.sh
```

Working on the Internet, `aedist --replay`

`aedist --replay` is *not* a synchronization mechanism. The usual model is applied also on changes from remote repositories.

A developer use `aedist (1) --replay` to fetch missing change-sets from a remote repository.

```
$ aedist --replay -p aegis.cur -f aegis.sf.net -inver 4.23*
aedist: aedist --receive -project=aegis.4.23 -change=283
      -file 'http://aegis.sf.net/cgi-bin/aeget/aegis.4.23.C283/?aedist'
aedist: aegis --new-change 283 --project=aegis.4.23 --file=/tmp/aedist-...
aedist: project "aegis.4.23": change 283: created
aedist: aegis --develop-begin 283--project aegis.4.23
project "aegis.4.23": change 283: user "walter" has begun development
aedist: aegis --copy-file --project=aegis.4.23 \
      --change=283 --branch=4.23 --delta=121 \
      libaegis/wide_output/column.cc libaegis/wide_output/wrap.cc
aedist: aegis --copy-file --project=aegis.4.23 --change=283 \
      --branch=4.23 --delta=120 libaegis/output/to_wide.cc
aedist: aegis --new-test --project=aegis.4.23 --change=283 test/02/t0232a-pmil.sh
```

Working on the Internet, `aedist --replay`

`aedist --replay` is *not* a synchronization mechanism. The usual model is applied also on changes from remote repositories.

A developer use `aedist (1) --replay` to fetch missing change-sets from a remote repository.

```
$ aedist --replay -p aegis.cur -f aegis.sf.net -inver 4.23*
aedist: aedist --receive -project=aegis.4.23 -change=283
      -file 'http://aegis.sf.net/cgi-bin/aeget/aegis.4.23.C283/?aedist'
aedist: aegis --new-change 283 --project=aegis.4.23 --file=/tmp/aedist-...
aedist: project "aegis.4.23": change 283: created
aedist: aegis --develop-begin 283--project aegis.4.23
project "aegis.4.23": change 283: user "walter" has begun development
aedist: aegis --copy-file --project=aegis.4.23 \
      --change=283 --branch=4.23 --delta=121 \
      libaegis/wide_output/column.cc libaegis/wide_output/wrap.cc
aedist: aegis --copy-file --project=aegis.4.23 --change=283 \
      --branch=4.23 --delta=120 libaegis/output/to_wide.cc
aedist: aegis --new-test --project=aegis.4.23 --change=283 test/02/t0232a-pmil.sh
```

Working on the Internet, `aedist --replay`

`aedist --replay` is *not* a synchronization mechanism. The usual model is applied also on changes from remote repositories.

A developer use `aedist (1) --replay` to fetch missing change-sets from a remote repository.

```
$ aedist --replay -p aegis.cur -f aegis.sf.net -inver 4.23*
aedist: aedist --receive -project=aegis.4.23 -change=283
      -file 'http://aegis.sf.net/cgi-bin/aeget/aegis.4.23.C283/?aedist'
aedist: aegis --new-change 283 --project=aegis.4.23 --file=/tmp/aedist-...
aedist: project "aegis.4.23": change 283: created
aedist: aegis --develop-begin 283 --project aegis.4.23
project "aegis.4.23": change 283: user "walter" has begun development
aedist: aegis --copy-file --project=aegis.4.23 \
      --change=283 --branch=4.23 --delta=121 \
      libaegis/wide_output/column.cc libaegis/wide_output/wrap.cc
aedist: aegis --copy-file --project=aegis.4.23 --change=283 \
      --branch=4.23 --delta=120 libaegis/output/to_wide.cc
aedist: aegis --new-test --project=aegis.4.23 --change=283 test/02/t0232a-pmil.sh
```

Working on the Internet, `aedist --replay`

`aedist --replay` is *not* a synchronization mechanism. The usual model is applied also on changes from remote repositories.

A developer use `aedist (1) --replay` to fetch missing change-sets from a remote repository.

```
$ aedist --replay -p aegis.cur -f aegis.sf.net -inver 4.23*
aedist: aedist --receive -project=aegis.4.23 -change=283
      -file 'http://aegis.sf.net/cgi-bin/aeget/aegis.4.23.C283/?aedist'
aedist: aegis --new-change 283 --project=aegis.4.23 --file=/tmp/aedist-...
aedist: project "aegis.4.23": change 283: created
aedist: aegis --develop-begin 283 --project aegis.4.23
project "aegis.4.23": change 283: user "walter" has begun development
aedist: aegis --copy-file --project=aegis.4.23 \
      --change=283 --branch=4.23 --delta=121 \
      libaegis/wide_output/column.cc libaegis/wide_output/wrap.cc
aedist: aegis --copy-file --project=aegis.4.23 --change=283 \
      --branch=4.23 --delta=120 libaegis/output/to_wide.cc
aedist: aegis --new-test --project=aegis.4.23 --change=283 test/02/t0232a-pmil.sh
```

Working on the Internet, `aedist --replay`

`aedist --replay` is *not* a synchronization mechanism. The usual model is applied also on changes from remote repositories.

A developer use `aedist (1) --replay` to fetch missing change-sets from a remote repository.

```
$ aedist --replay -p aegis.cur -f aegis.sf.net -inver 4.23*
aedist: aedist --receive -project=aegis.4.23 -change=283
      -file 'http://aegis.sf.net/cgi-bin/aeget/aegis.4.23.C283/?aedist'
aedist: aegis --new-change 283 --project=aegis.4.23 --file=/tmp/aedist-...
aedist: project "aegis.4.23": change 283: created
aedist: aegis --develop-begin 283 --project aegis.4.23
project "aegis.4.23": change 283: user "walter" has begun development
aedist: aegis --copy-file --project=aegis.4.23 \
      --change=283 --branch=4.23 --delta=121 \
      libaegis/wide_output/column.cc libaegis/wide_output/wrap.cc
aedist: aegis --copy-file --project=aegis.4.23 --change=283 \
      --branch=4.23 --delta=120 libaegis/output/to_wide.cc
aedist: aegis --new-test --project=aegis.4.23 --change=283 test/02/t0232a-pmil.sh
```

Working on the Internet, `aedist --replay`

`aedist --replay` is *not* a synchronization mechanism. The usual model is applied also on changes from remote repositories.

A developer use `aedist (1) --replay` to fetch missing change-sets from a remote repository.

```
$ aedist --replay -p aegis.cur -f aegis.sf.net -inver 4.23*
aedist: aedist --receive -project=aegis.4.23 -change=283
      -file 'http://aegis.sf.net/cgi-bin/aeget/aegis.4.23.C283/?aedist'
aedist: aegis --new-change 283 --project=aegis.4.23 --file=/tmp/aedist-...
aedist: project "aegis.4.23": change 283: created
aedist: aegis --develop-begin 283 --project aegis.4.23
project "aegis.4.23": change 283: user "walter" has begun development
aedist: aegis --copy-file --project=aegis.4.23 \
      --change=283 --branch=4.23 --delta=121 \
      libaegis/wide_output/column.cc libaegis/wide_output/wrap.cc
aedist: aegis --copy-file --project=aegis.4.23 --change=283 \
      --branch=4.23 --delta=120 libaegis/output/to_wide.cc
aedist: aegis --new-test --project=aegis.4.23 --change=283 test/02/t0232a-pmil.sh
```

Working on the Internet, aedist --replay

```
aedist: /usr/bin/patch -p0 -f -s libaegis/output/to_wide.cc \  
    /tmp/aedist-6f35ab91-fbf6-4b7a-9556-7ccbbdc43cc9  
aedist: /usr/bin/patch -p0 -f -s libaegis/wide_output/column.cc \  
    /tmp/aedist-6bbe8e50-c5e2-4be2-80f6-ff98f9595070  
aedist: /usr/bin/patch -p0 -f -s libaegis/wide_output/wrap.cc \  
    /tmp/aedist-047b0c98-3cec-4c81-a771-6c5da74484cd  
aedist: aegis --diff --only-merge --change=283 --project=aegis.4.23  
aegis: cd /prep/baseline/aegis/history  
aegis: cd /prep/work-in-progress/aegis/aegis.4.23.C283  
aegis: fmerge -ignore_identical_conflicts /tmp/aegis-146... \  
.../libaegis/output/to_wide.cc \  
.../aegis.4.23.C283/libaegis/output/to_wide.cc,B \  
    -o .../aegis.4.23.C283/libaegis/output/to_wide.cc \  
-c .../aegis.4.23.C283/libaegis/output/to_wide.cc,C  
...
```

Working on the Internet, aedist --replay

```
aedist: /usr/bin/patch -p0 -f -s libaegis/output/to_wide.cc \  
    /tmp/aedist-6f35ab91-fbf6-4b7a-9556-7ccbbdc43cc9  
aedist: /usr/bin/patch -p0 -f -s libaegis/wide_output/column.cc \  
    /tmp/aedist-6bbe8e50-c5e2-4be2-80f6-ff98f9595070  
aedist: /usr/bin/patch -p0 -f -s libaegis/wide_output/wrap.cc \  
    /tmp/aedist-047b0c98-3cec-4c81-a771-6c5da74484cd  
aedist: aegis --diff --only-merge --change=283 --project=aegis.4.23  
aegis: cd /prep/baseline/aegis/history  
aegis: cd /prep/work-in-progress/aegis/aegis.4.23.C283  
aegis: fmerge -ignore_identical_conflicts /tmp/aegis-146... \  
.../libaegis/output/to_wide.cc \  
.../aegis.4.23.C283/libaegis/output/to_wide.cc,B \  
    -o .../aegis.4.23.C283/libaegis/output/to_wide.cc \  
-c .../aegis.4.23.C283/libaegis/output/to_wide.cc,C  
...
```

Working on the Internet, aedist --replay

```
aedist: /usr/bin/patch -p0 -f -s libaegis/output/to_wide.cc \  
    /tmp/aedist-6f35ab91-fbf6-4b7a-9556-7ccbbdc43cc9  
aedist: /usr/bin/patch -p0 -f -s libaegis/wide_output/column.cc \  
    /tmp/aedist-6bbe8e50-c5e2-4be2-80f6-ff98f9595070  
aedist: /usr/bin/patch -p0 -f -s libaegis/wide_output/wrap.cc \  
    /tmp/aedist-047b0c98-3cec-4c81-a771-6c5da74484cd  
aedist: aegis --diff --only-merge --change=283 --project=aegis.4.23  
aegis: cd /prep/baseline/aegis/history  
aegis: cd /prep/work-in-progress/aegis/aegis.4.23.C283  
aegis: fmerge -ignore_identical_conflicts /tmp/aegis-146... \  
.../libaegis/output/to_wide.cc \  
.../aegis.4.23.C283/libaegis/output/to_wide.cc,B \  
    -o .../aegis.4.23.C283/libaegis/output/to_wide.cc \  
-c .../aegis.4.23.C283/libaegis/output/to_wide.cc,C  
...
```

Working on the Internet, `aedist --replay`

```
aedist: aegis --diff --no-merge --change=283 --project=aegis.4.23
...
...
aegis: project "aegis.4.23": change 283: difference complete
aedist: Warning: This change contains files which could host a Trojan
horse attack. You should review it before building it or testing it
or completing development. This change remains in the being_developed state.
```

Working on the Internet, `aedist --replay`

```
aedist: aegis --diff --no-merge --change=283 --project=aegis.4.23
...
...
aegis: project "aegis.4.23": change 283: difference complete
aedist: Warning: This change contains files which could host a Trojan
horse attack. You should review it before building it or testing it
or completing development. This change remains in the being_developed state.
```

Working on the Internet, `aedist --replay`

```
aedist: aegis --diff --no-merge --change=283 --project=aegis.4.23
...
...
aegis: project "aegis.4.23": change 283: difference complete
aedist: Warning: This change contains files which could host a Trojan
horse attack. You should review it before building it or testing it
or completing development. This change remains in the being_developed state.
```

POSIX compliant systems

No problem :-)

Microsoft Windows

Aegis depends on the underlying security provided by the operating system (rather than re-invent yet another security mechanism). However, in order to do this, Aegis uses the POSIX seteuid system call, which has no direct equivalent on Windows NT. This makes porting difficult. Single-user ports are possible (e.g. using Cygwin), but are not usually what folks want.

Some ideas

- A plugin for Eclipse
- Integration with GForge/Trac
- Integration with tailor.
Tailor is a tool to migrate changesets between different SCM systems repositories.

The last slide

- For more info see <http://aegis.sourceforge.net/>
- Questions?????