


Apache Software Foundation
Building Bridges
@ FOSDEM 2005

Apache Gump

Continuous Integration on Steroids


Apache Gump: Continuous integration on Steroids
By Leo Simons, leosimons@apache.org

26 February 2005




About LSD

- <http://www.leosimons.com/>
- leosimons@apache.org
- ASF Member
- Apache Excalibur, Gump, Wiki, Infrastructure
- Physics Student, Enschede




Presentations, Pragmatic Open Source Style™

- Please ask questions
- The bikeshed is definitely red
- SubEthaEdit, not Vi, not Emacs
- We don't talk about licensing
- We talk too much about everything else
- And do ask questions!



About the audience


- There are about this many people here:
- Many are F/OSS developers:
- And these people don't do java:




In this presentation

- Apache Cocoon (briefly)
 - XML publishing framework
- Apache Excalibur (briefer)
 - IoC framework
- Apache $\{foo\}$ (honerable mention)
 - There's a lot
- Apache Gump (kew!!)
 - Making all that work together
 - Past, present, future
 - Building bridges
 - And the steroids

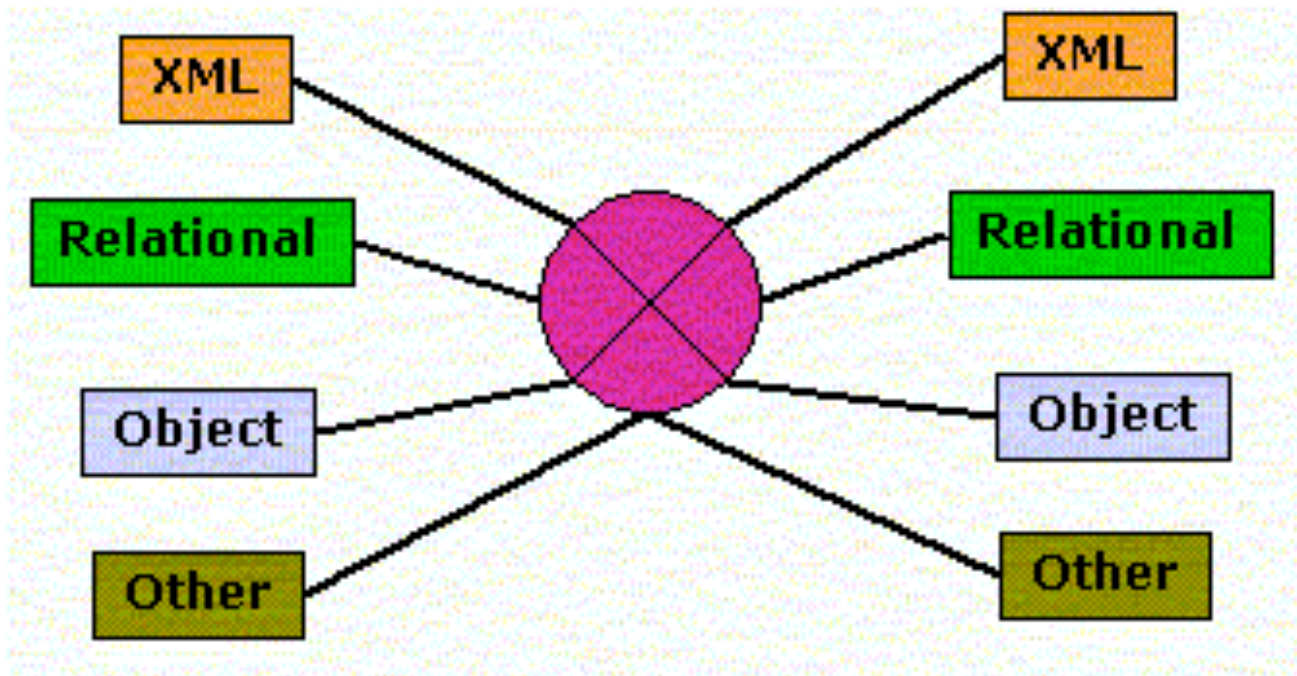
Apache Cocoon: Goal



This image ripped from cocoon documentation, © ASF and stuff

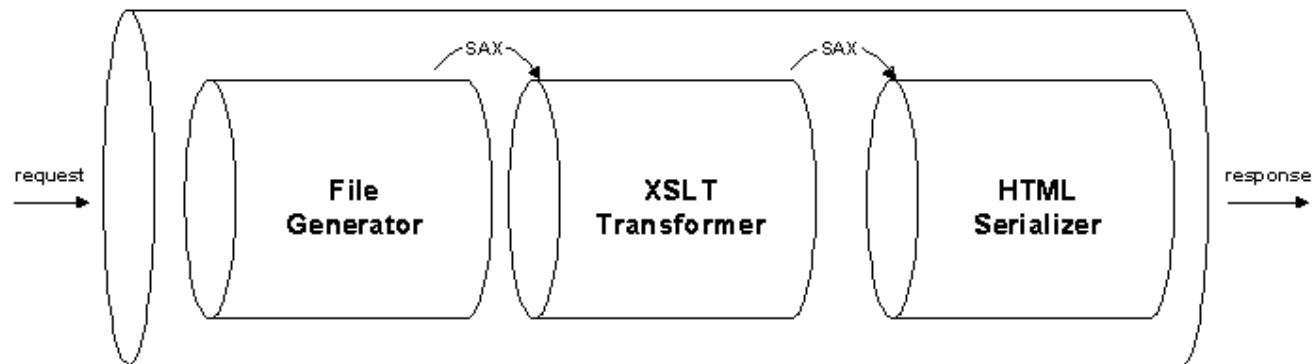


Apache Cocoon: Glue




This image ripped from cocoon documentation, © ASF and stuff


Apache Cocoon: Pipelines



This image ripped from cocoon documentation, © ASF and stuff


Apache Cocoon: Architecture






Apache Cocoon: Oh boy!

- 2400+ java sourcefiles
- 300.000+ lines of code (excluding comments)
- Integrates with dozens of other F/OSS projects using component architecture
- Cocoon “core” has 338 dependencies



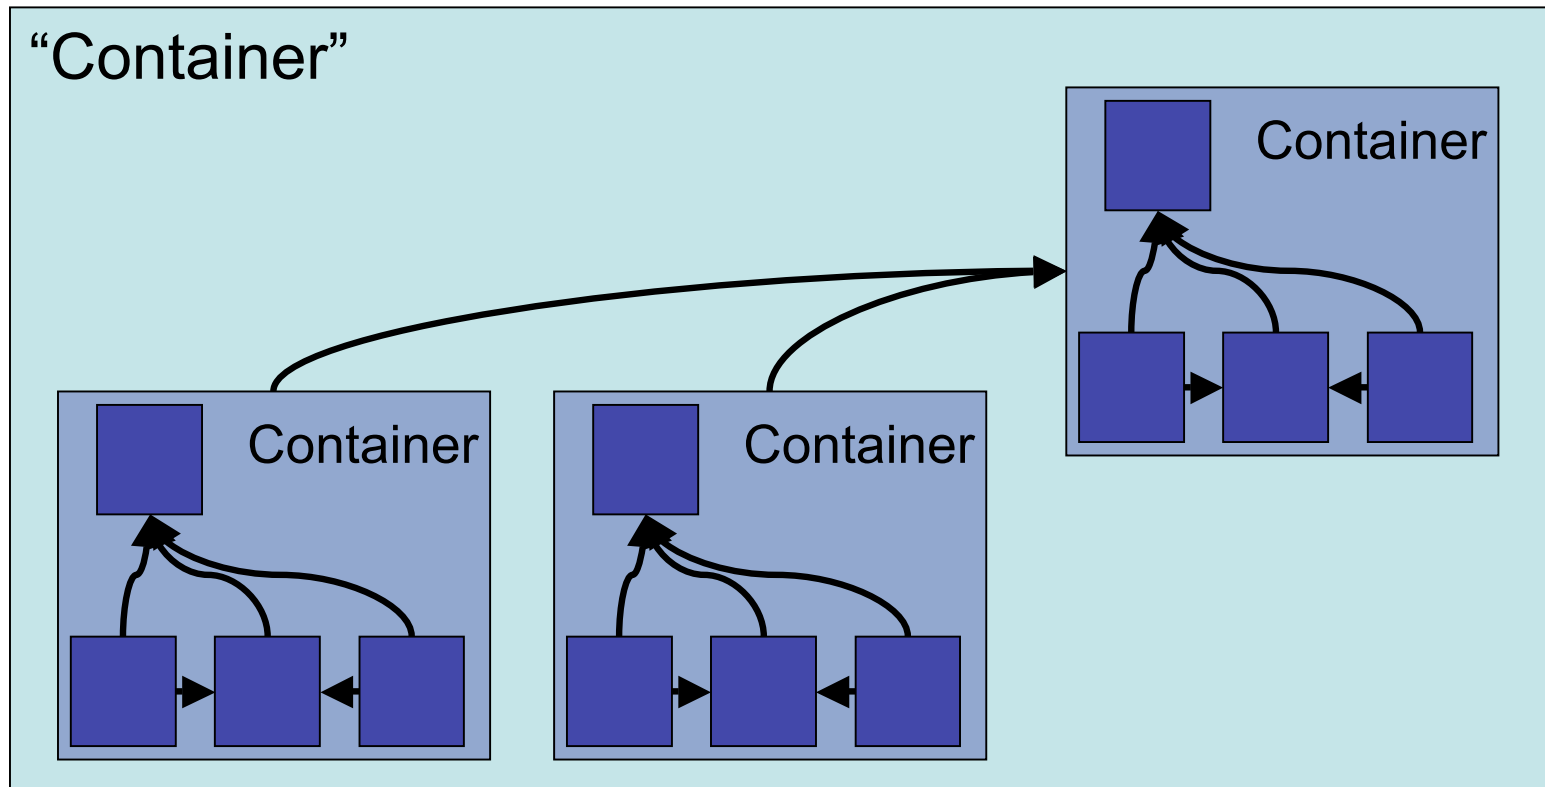
Apache Excalibur: Goal


Somewhere
In there...



This image ripped from cocoon documentation, © ASF and stuff


Apache Excalibur: Inversion of Control





Apache Excalibur: Umpf!

- 800+ java sourcefiles
- 90000+ lines of code (excluding comments)
- 50+ separate distributables
- Excalibur “core” has 100s of dependencies



Apache `{foo}`


- Majority of activity at ASF is java-based now
- 1000+ committers
- Apache is on the JCP
- A big chunk of server-side java happens at Apache
- Some of the de-facto standards such as Ant
- Millions of lines of java code

Artist impression




Apache Gump: Continuous integration on Steroids
By Leo Simons, leosimons@apache.org

26 February 2005




Apache Software Foundation
Building Bridges
@ FOSDEM 2005

Okay, so we do java...




Apache Software Foundation
Building Bridges
@ FOSDEM 2005

**...how on earth do we know
it all actually works?**




Introducing Continuous integration

- Get yourself a beefy server
- Automate all project builds
- Automate all project tests
- Generate reports and statistics
- Make it all part of your development process (i.e. send yourself e-mail)




Introducing Gump


- A batch-processing continuous integration package written in python
- Lots of useful features like e-mail notification, RSS syndication, etc.
- Understands ant, maven, make, etc
- XML descriptors in a central repositories for shared maintenance



When continuous integration fails...


- Project B depends on project A
- Project A makes a teeny weenie small incompatible change in CVS head
- Project B releases 1.1
- Project A releases 1.1
- 1.1 of B does not work with 1.1 of A






...it can fail miserably...


- Project B depends on project A, Project C on A and B
- Project A makes a small incompatible change in CVS head
- Project B adds a cool new feature
- Project C notices the change in CVS head of A in time, and adapts, using a “snapshot” version. Same for B.
- Project C releases 1.1
- Project B releases 1.1
- Project A releases 1.1






Welcome to jar hell!

- The problem is only noticed months later...
- So what's the best recovery solution?
Complex!
- Now imagine you've got 400 dependencies,
not 2...






Gump on Steroids


- Project B might declare it depends on Project A version 1.0
- We *override* that declaration, and instead compile it against A from CVS head *that we just built*

```
<!-- maven buildfile snippet -->  
<project>  
<name>B</name>  
<dependencies>  
  <depend>  
    <groupId>A</groupId>  
    <artifactId>A</artifactId>  
    <version>1.0</version>  
  </depend>  
</dependencies>  
</project>
```



The World on Steroids

- Gump builds 100s of projects (not just A and B)
- Gump compiles over 150.000 java sourcefiles several times a day
- We send out lots of e-mails to developer mailing lists when stuff breaks



Example: change to log4j

Subject: [human-assisted gump] log4j back incompatible change

Date: Fri, 29 Oct 2004 01:07:38 -0400


(...) change

http://cvs.apache.org/viewcvs.cgi/logging-log4j/src/java/org/apache/log4j/Category.java?r1=1.88&r2=1.89&diff_format=h

caused a compilation failure of the following projects:


- commons-logging
- velocity
- ant

and, as a result, caused a drop in the gump's success from 83% to 53%
(...)




It actually works! *(sometimes)*

- Without Gump, it would be near-impossible to upgrade one of Cocoon's 338 dependencies without a lot of stuff breaking
- In fairness, Gump is quite far from being perfect, and incompatibility problems still do occur frequently
- But Gump really does help to increase awareness and really does get people to talk early and often



Building Bridges: Kaffe on Steroids

- Each night, we do a gump run against a CVS snapshot of Kaffe
- <http://brutus.apache.org/gump/kaffe/index.html>
- Measure of compatibility between ASF projects and the free java implementations



Where we're at

Kaffe-CVS

Overall project success : 31.43%


Projects	Successes	Failures	Prereqs
821	158 (19.24%)	26 (3.17%)	537 (65.41%)

JDK 1.4

Overall project success : 82.95%

Projects	Successes	Failures	Prereqs
821	581 (70.77%)	31 (3.78%)	109 (13.28%)

(Numbers for Friday afternoon)




Example: Ant on kaffe

```
On Wed, 23 Feb 2005, Dalibor Topic <robilad@kaffe.org> wrote:  
> Stefan Bodewig <bodewig@apache.org> writes:  
>> ... since optional/condition/IsPingable.java shouldn't be compiled  
>> otherwise.  
>>  
>> Ant checks for the class java.lang.Readable to detect JDK 1.5, has  
>> this recently been added to Kaffe?  
>  
> Yeah, just yesterday. Merge of a few non-generics bits from the  
> generics branch within GNU classpath. I guess the  
> java.lang.management classes are a safe bet, though :)
```

*So far the only JDK 1.5 feature Ant uses is InetAddress.isReachable.
Do us a favor and implement it ;-)*


I'll modify Ant's build file to adapt, but I'd like to keep the
Readable class as marker. It may be good to revisit each JDK 1.5
feature we add (if we do) to see whether Kaffe can support it.

Stefan



What does gump do here?

- Help pinpoint exactly where to focus developer effort (both sides) to improve interop
- Provide a measure of progress
- Act as a really big integration test for Kaffe and classpath




Apache Software Foundation
Building Bridges
@ FOSDEM 2005

Gump3: where we're going


Fancy Web UI

(thanks to cocoon and the mozilla team :-)

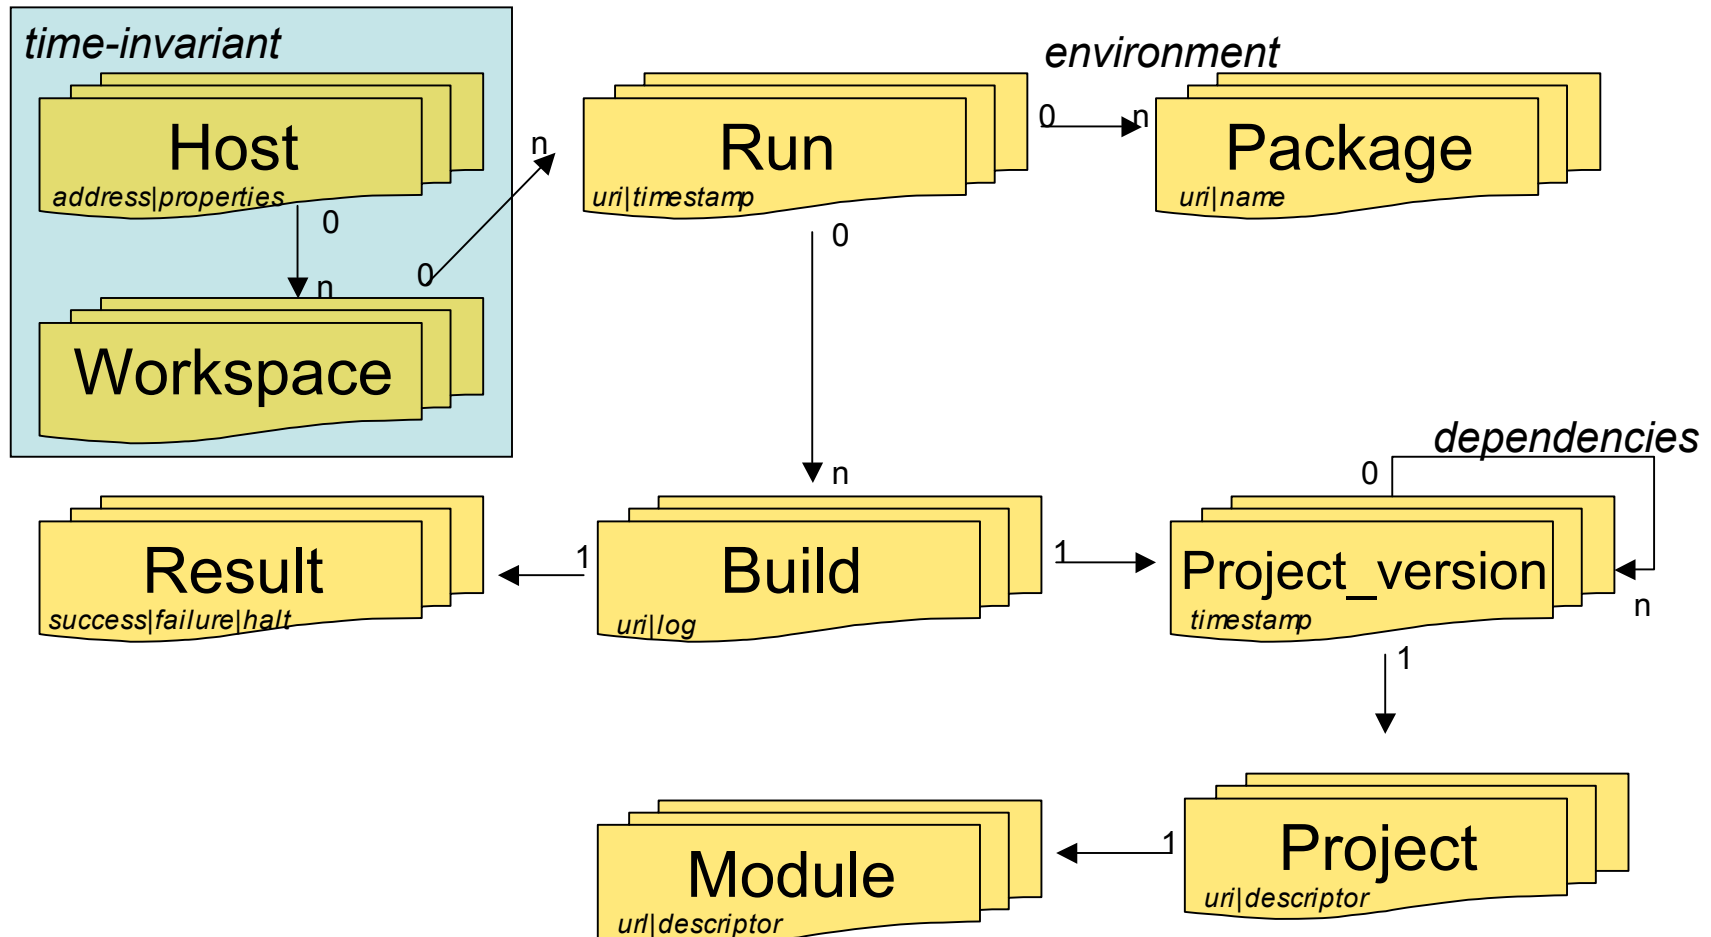



Gump3: a *lot* smarter

- We figured out the algorithm that shows *exactly* what change where caused what problems
- “Just” need to convert it to code



Gump3: Historical Database





Gump3: Fancy CLI

```
paddo:~/svn/gump/branches/Gump3 lsimons$ ./gump help
```

```
Utility commandline interface for Gump.
```


```
Usage:
```

```
./gump command [opts ...]
```

```
Available commands are:
```

```
run          -- run pygump
debug        -- run pygump in debug mode
test         -- run the pygump unit tests
dynagump     -- run the dynagump web application server
update-host -- update the configuration of a gump host
create-database -- create a new gump MySQL database
pycompile    -- compile all pyump source files
```

```
(...)
```




Gump3: Plugin architecture

```
class LoggingPlugin(AbstractPlugin):
    """Plugin that prints debug messages as it visits model objects."""
    def __init__(self, log):
        self.log = log

    def initialize(self):
        self.log.debug("Initializing...")


    def visit_workspace(self, workspace):
        self.log.debug("Visiting workspace.")

    def visit_repository(self, repository):
        self.log.debug("Visiting repository '%s'." % repository.name)
    (...)
    def finalize(self):
        self.log.debug("Finishing up...")
```




Apache Software Foundation
Building Bridges
@ FOSDEM 2005

Gump3: You tell us!



Apache Software Foundation
Building Bridges
@ FOSDEM 2005

**Questions, comments,
criticism, offers to buy
me beer?**



Apache Software Foundation
Building Bridges
@ FOSDEM 2005

ApacheCon Europe 2006

- Stuttgart, July 18-22
- <http://www.apachecon.com/>
- Everything you ever wanted to know about Apache and more!