

2IS55 – Software Evolution

Revision Control Systems

Revision Control Systems

- Definition
- Quick recap of general use
- Terms used often
- History
- Distributed workflows
- Advanced concepts

Definition

- System used to reliably reproduce a specific revision of software over time

Quick recap of general use

- scm checkout software-1.0
- Edit and test in your local sandbox
- scm checkin
 - “I added a new feature that does X”
- scm branch mainline software-2.1
- scm merge feature-branch software-1.1

Terms used often

- Repository
- Revision
- Checkout
- Commit

Terms used often (continued)

- Mainline *or* Trunk
- Branch
- Merge
- Tag

History

- SCCS
- RCS/CVS
- Subversion
- First Wave DVCS: Arch, Monotone
- Current Wave DVCS: Bazaar, Mercurial, Git

SCCS

(~ 1975)

- First revision control system
- Single file based
- Control records
- Delta table
- Revision information: who, when, changes

RCS

(1982)

- Checkout needs locks
- Single file based, history stored in `filename,v`
- Tags for groups of files
- Reverse deltas and forward deltas
- Revision information: who, when, changes, branch, tags

CVS

(1986)

- No locks
- Client / server
- Anonymous readonly access
- First de-facto open source RCS



CVS, continued

- Revision information: who, when, changes, branch, tags
- `cvls log boot.c`

```
RCS file: /cvs/src/sys/arch/sparc/stand/boot/boot.c,v
Working file: boot.c
head: 1.6
branch:
locks: strict
access list:
symbolic names:
  OPENBSD_4_7: 1.6.0.26
  [...]
keyword substitution: kv
total revisions: 12;   selected revisions: 1
description:
-----
revision 1.3
date: 2002/03/14 01:26:44;  author: millert;  state: Exp;  lines: +5 -5
First round of __P removal in sys
=====
```

Subversion

(2000)

- Mission: “Build a better CVS”
- Atomic commits
- Version control for directories
- Copy files and directories
- Better branching, merging and tagging
- Previous (?) de-facto open source RCS



Subversion, continued

- Revision information: who, when, changes, branch, tags
- `svn log -v https://svn.win.tue.nl/repos/MCRL2/trunk`

```
-----  
r7524 | wieger | 2010-03-18 11:38:17 +0100 (Thu, 18 Mar 2010) | 1 line
```

```
Changed paths:
```

```
  M /trunk/tools/pbespareqelm/pbespareqelm.cpp
```

```
- Fixed typo in option string.
```

Arch

(2001)

- Mission: “Replace CVS”
- Bizarre command set and file naming conventions
 - `$ t1a changes -o ,,new-robot-comment`
- But: distributed version control!
- Not in use anymore



Monotone

(2003)

- Mission: “Replace CVS”
- Correctness, verifiable history
- Initially slow
- But: distributed version control!
- Still in use



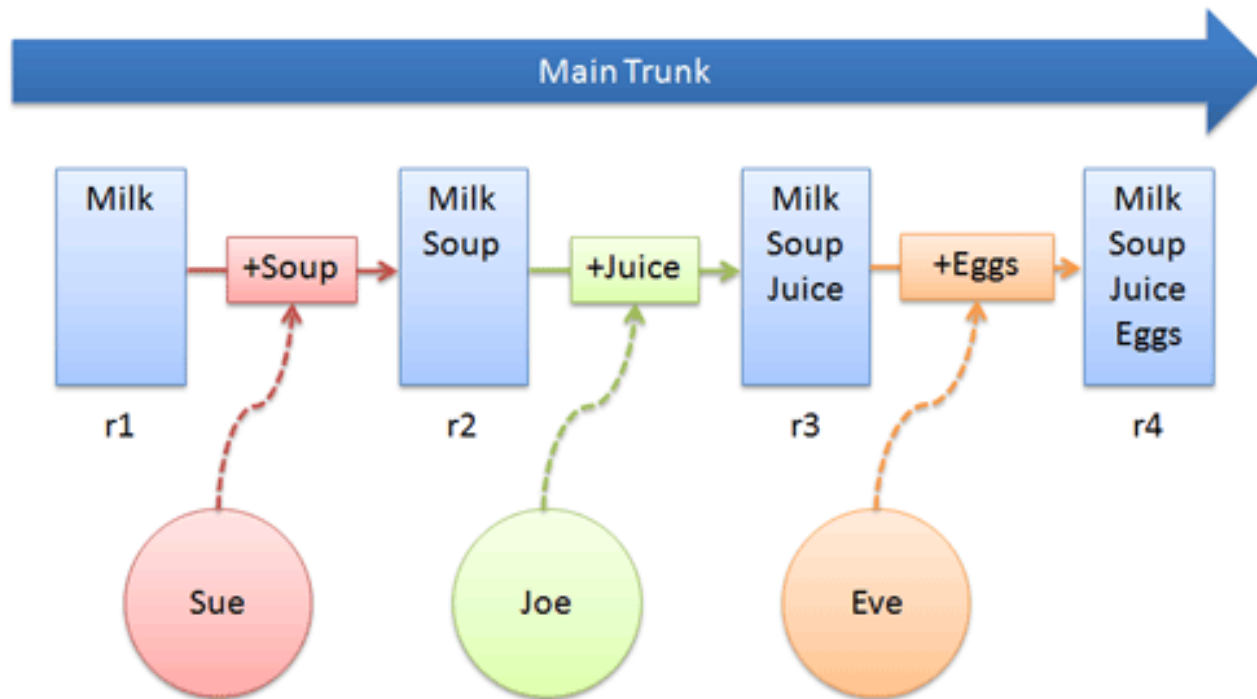
Distributed vs Centralized

- Differences
- Distributed is “chaos”
- Resistance to change
- “Chaos” is manageable

Centralized development

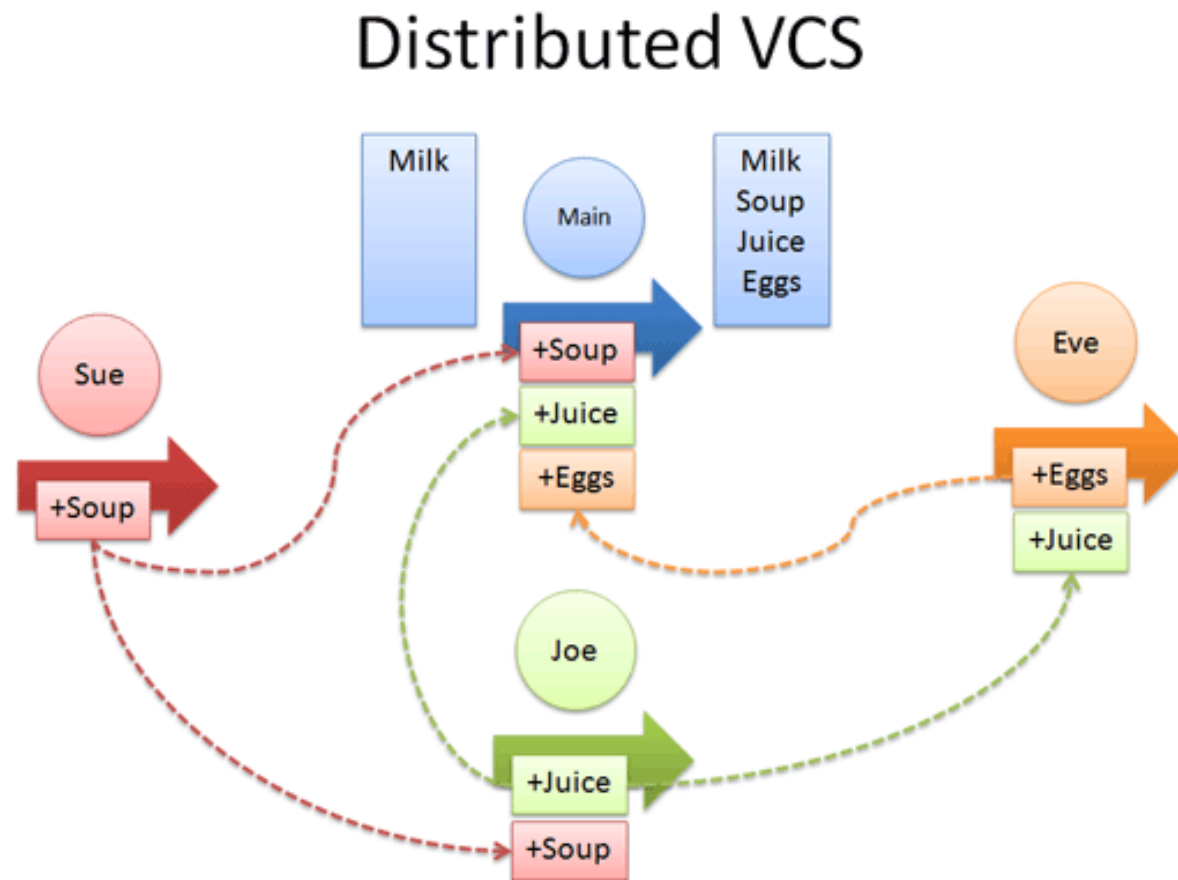
- “Serial” development

Centralized VCS



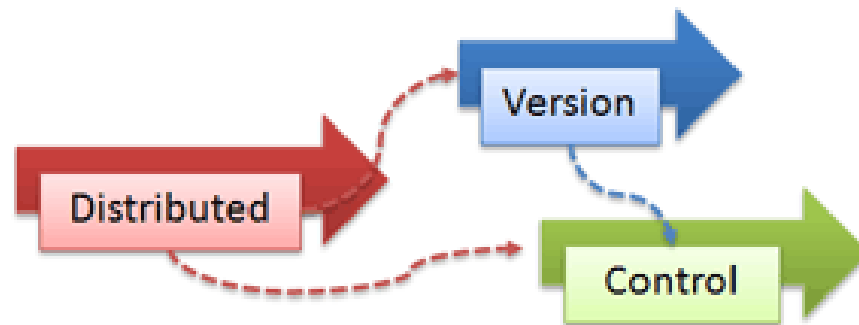
Distributed development

- “Parallel” development



Distributed vs Centralized, ct'd

- “Throw away” development
- Delegate subprojects
- “Central” release repository



The current big three

- Bazaar
- Mercurial
- Git
- A lot of overlap

Bazaar

(2005)

- Evolved from Arch
- Mission: “Accessible DRCS”
- Cross-platform
- Slow initially, better now
- Projects: MySQL, Debian APT



Unique features of Bazaar

- Branches are simply new clones of repositories
- Very good support for graphical UIs on all platforms
- Tracks both file and directory renames explicitly

Bazaar log

- `bzr log lp:bzr-svn`

```
-----  
revno: 3274  
committer: Jelmer Vernooij <jelmer@samba.org>  
branch nick: trunk  
timestamp: Tue 2010-03-09 13:56:24 +0100  
message:  
    Provide BranchConfig._get_change_editor().  
-----
```

```
revno: 3273  
committer: Jelmer Vernooij <jelmer@samba.org>  
branch nick: trunk  
timestamp: Tue 2010-03-09 00:47:01 +0100  
message:  
    Fix use of svn:author=author.
```

Mercurial

(2005)

- Mission: replace Bitkeeper
- Inspiration from Monotone
- Cross platform
- Projects: SUN Java, Mozilla Firefox, Python



Unique features of Mercurial

- Global branches
- Local revision numbers

Mercurial log

- hg log mozilla

```
changeset: 15262:263749849d0b
tag:       tip
user:      Marco Zehe <marco.zehe@googlemail.com>
date:      Thu Jun 05 11:49:51 2008 +0200
summary:   b=432467, firefox segfaults in plone kupu editor
           [@ nsDocAccessible::FlushPendingEvents], on Tablet PC
           [@arena_dalloc_small] (steps to reproduce in comment #26),
           p=Ginn Chen <ginn.chen@sun.com>, r=surkov
```

```
changeset: 15261:49cdeb4f8144
user:      Simon Montagu <smontagu@smontagu.org>
date:      Thu Jun 05 04:05:25 2008 -0700
summary:   Reftest for bug 229764
```

Git

(2005)

- Mission: replace Bitkeeper
- Inspiration from Monotone
- Linux dependent at first
- Projects: Linux Kernel, Android, Wine, Perl



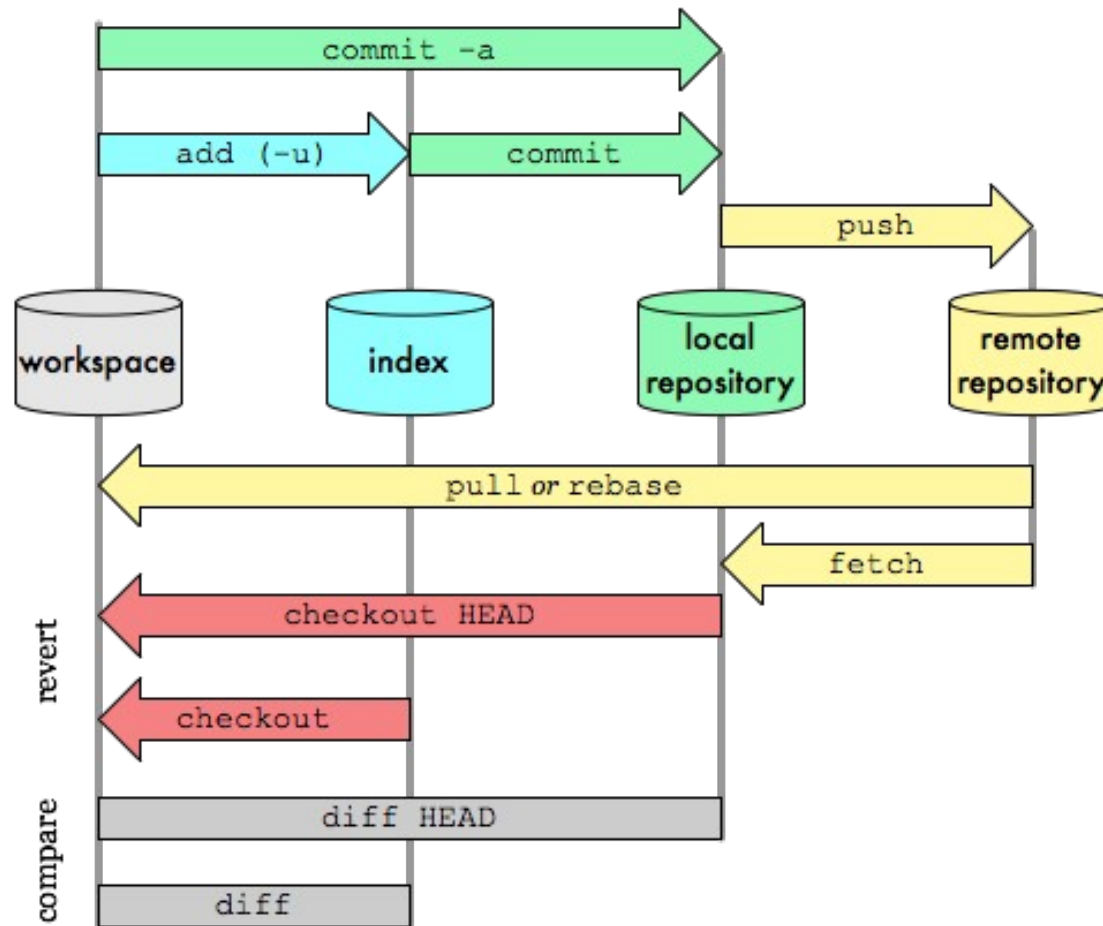
Unique features of Git

- Tracks file contents, not files
- Internal details visible
- The “index” or “staging area”

Git index

Git Data Transport Commands

<http://osteele.com>



Git log

- git log HEAD


```
commit ef36c9a3b2828f5a11feda9e4d2708bf3a4a7a52
Author: Eric Anholt <eric@anholt.net>
Date:   Wed Mar 17 12:46:21 2010 -0700
```

```
intel: Install the header file in the libdrm/ directory.
```

```
Suggested-by: Rémi Cardona <remi@gentoo.org>
Signed-off-by: Eric Anholt <eric@anholt.net>
```

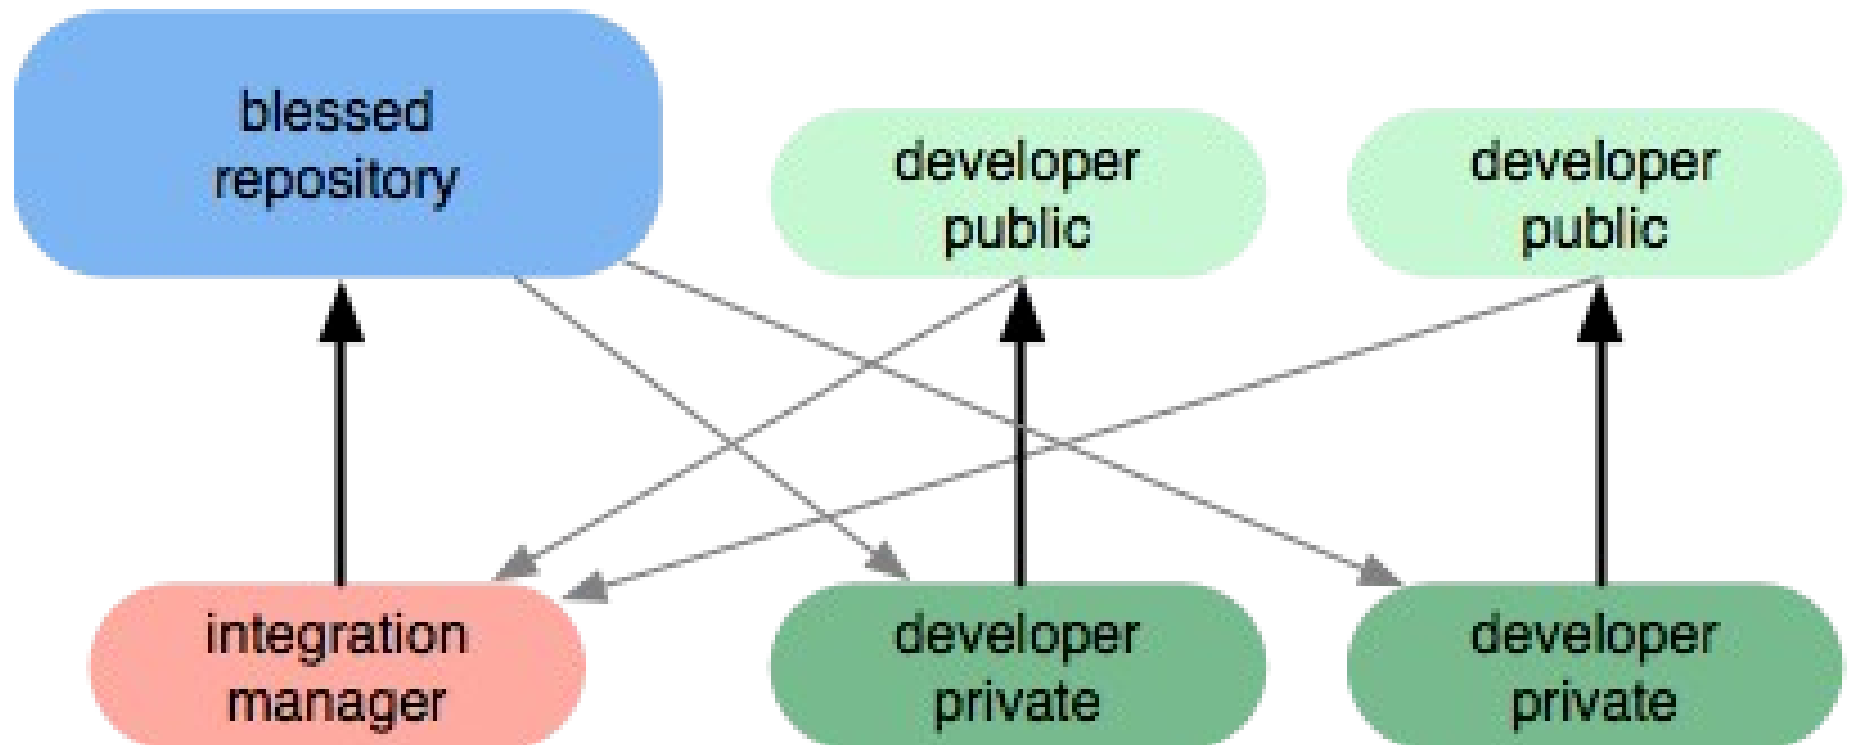
Distributed Workflows

- Centralized with benefits



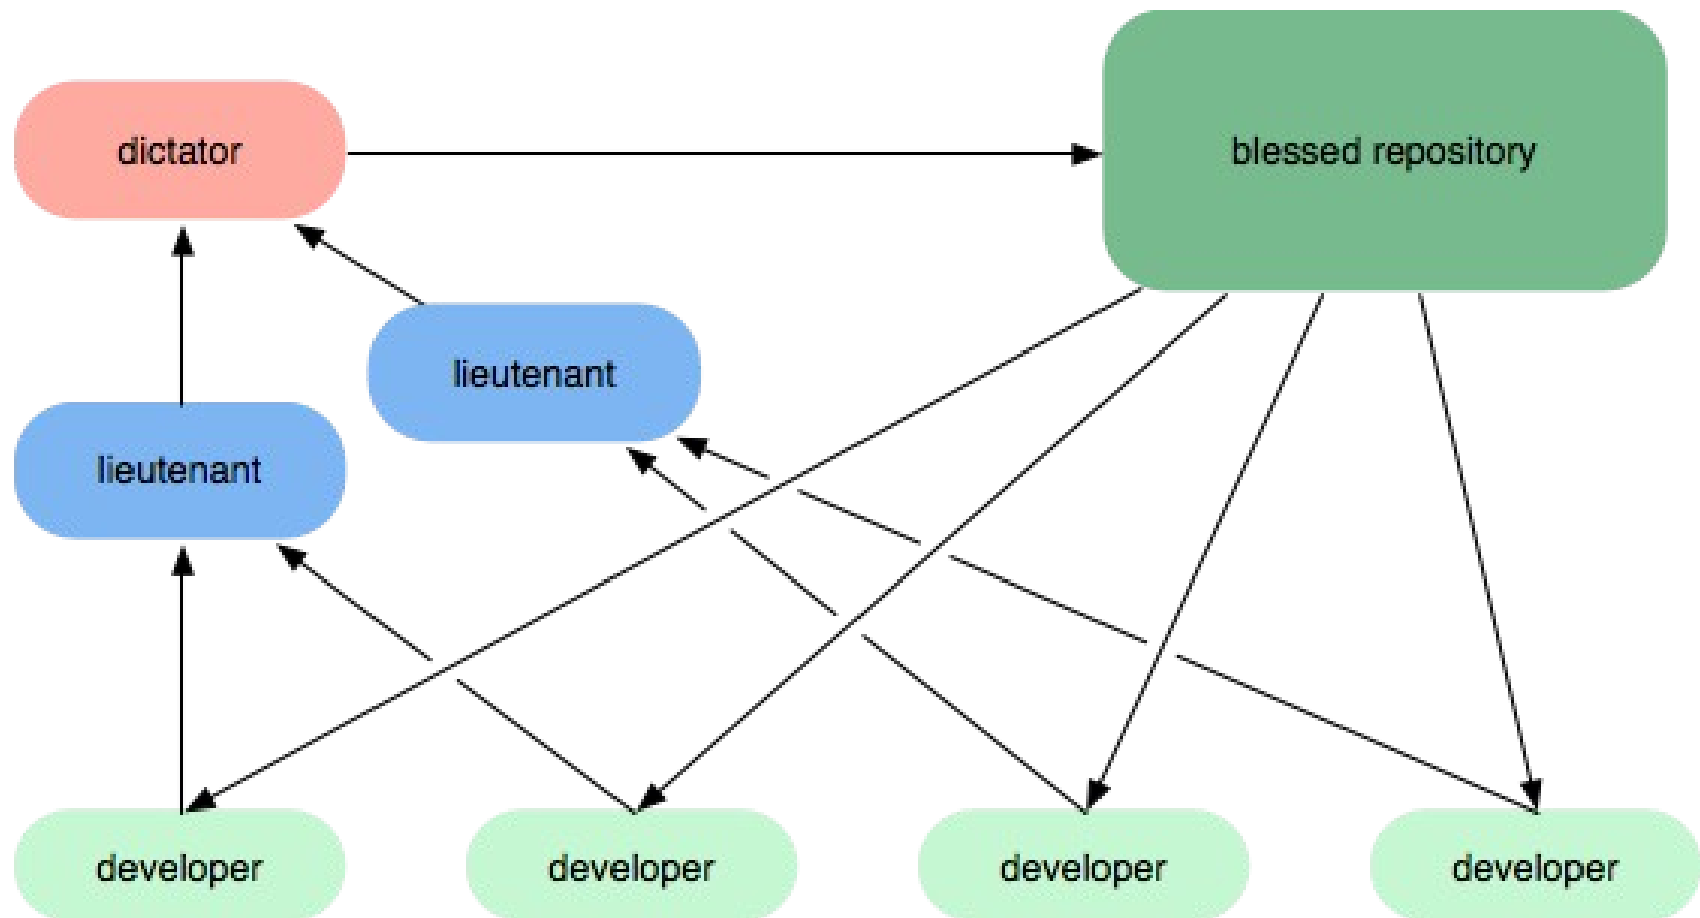
Distributed Workflows, ct'd

- Integration manager




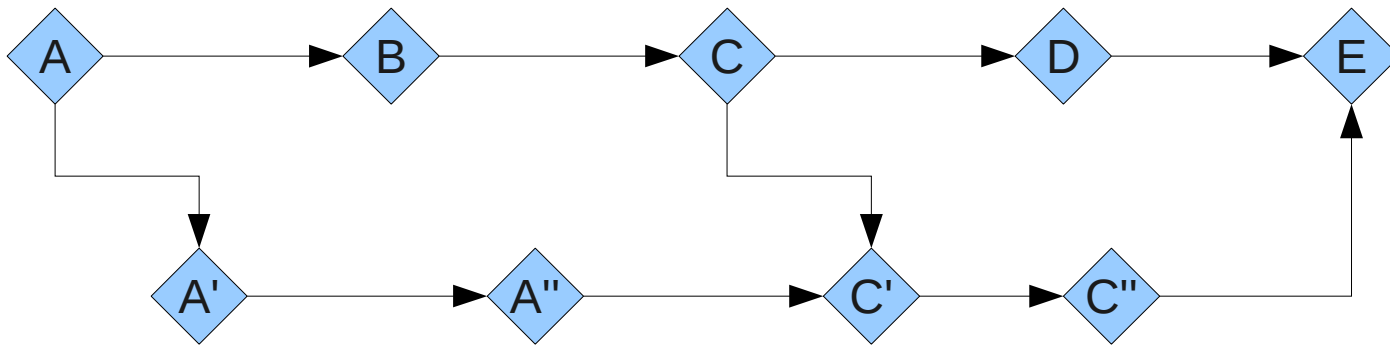
Distributed Workflows, ct'd

- Integration with “lieutenants”



Advanced concepts

- Rebase versus merge



Advanced concepts, ct'd

- Octopus merge

source: **linux-2.6.32**

[VIEW LATEST REVISION](#)

Added Modified Copied or renamed

Revision Log Mode:

- Stop on copy
 Follow copies
 Show only adds and deletes

View log starting at a

Show at most revisions

Show full log messages

View changes

Diff	Graph	Rev	Age	Author	Log Message
<input type="radio"/>		<input type="checkbox"/> @2490138c	14 months	torvalds	Merge branch 'for-linus' of ...
<input type="radio"/>		<input type="checkbox"/> @f6f79190	14 months	torvalds	Merge branch 'for-linus' of ...
<input type="radio"/>		<input type="checkbox"/> @73f526da	14 months	rolandd	Merge branch 'mad' into for-linus Conflicts: ...
<input type="radio"/>		<input type="checkbox"/> @45c448a1	14 months	rolandd	Merge branches 'cxgb3', 'ehca', 'ipath', 'ipoib', 'misc', 'mlx4', 'mthca' ..
<input type="radio"/>		<input type="checkbox"/> @a3c8b973	14 months	jmorris	Merge branch 'next' into for-linus
<input type="radio"/>		<input type="checkbox"/> @0d03d59	14 months	geert	md: Fix "strchr" [drivers/md/dm-log-userspace.ko] undefined! Comm
<input type="radio"/>		<input type="checkbox"/> @9f0ab4a3	14 months	roland	binfmt_elf: fix PT_INTERP bss handling In fs/binfmt_elf.c, ...
<input type="radio"/>		<input type="checkbox"/> @ec579358	14 months	jgunthorpe	TPM: Fixup boot probe timeout for tpm_tis driver When probing the de
<input type="radio"/>		<input type="checkbox"/> @526b6780	14 months	torvalds	Merge branch 'lookup-permissions-cleanup' * lookup-permissions-clea

Try it out!

- Bazaar:
<http://launchpad.net/>
- Mercurial:
<http://bitbucket.org/>
- Git:
<http://github.com/>
- Or just use it locally

Further reading

- **Original SCCS paper:**
<http://basepath.com/aup/talks/SCCS-Slideshow.pdf>
- **History of RCS:**
<http://www.burlingtontelecom.net/~ashawley/rcs/tichy1985rcs/rcs.html#id2504601>
- **Linux Kernel and BitKeeper:**
<http://lkml.indiana.edu/hypermail/linux/kernel/0504.0/1540.html>
<http://lkml.org/lkml/2005/4/7/150>
- **Distributed vs Centralized:**
<http://blog.ianbicking.org/distributed-vs-centralized-scm.html>
<http://web.mit.edu/ghudson/thoughts/bitkeeper.whynot>