Introducing Host Factory

Why manage a host?

Have your machines been broken into?
How do you know?
How would you repair?

How do you manage a host?

- You have no means to audit or inspect hosts.
- You depend on malware being indicated by anomalous behavior.
- You don't have anomalous behavior.
- Yet, there exists quality malware which does not show anomalous behavior.

How do you manage a host?

- You have existing damage today.
- From technical difficulty enforcing consistency, mistakes, missed updates, disk failures, entropy
- Maybe even malware
- Your vendors never tested your collection of versions against each other.

How could you manage a host?

- A featureful Linux system has 500,000 files.
- Managing the contents and ownership/modes of these files is a good way to manage a host.
- But some of those files are supposed to be different.
- No previous tool does this.

How could you manage a host?

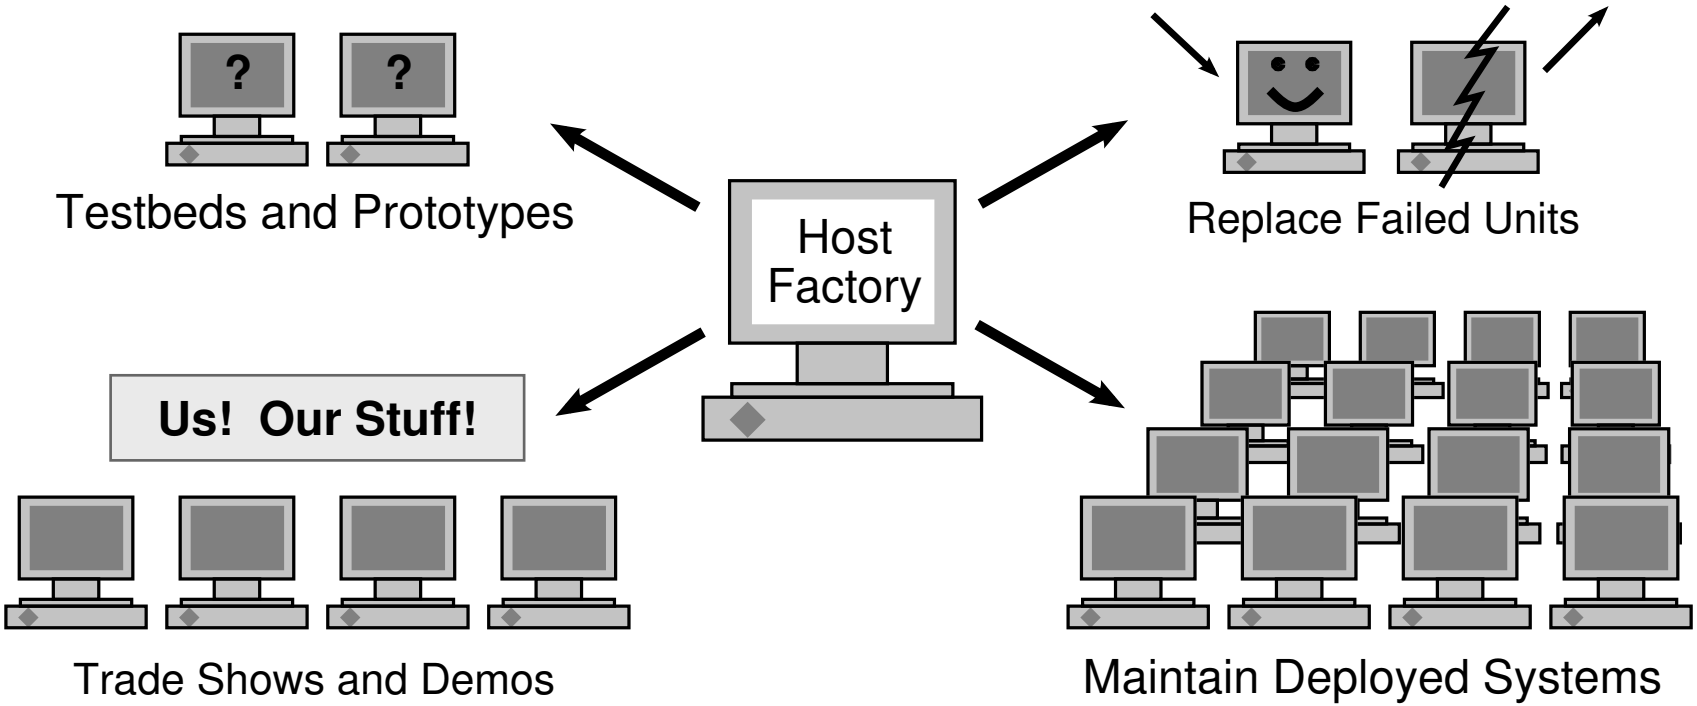

Then you can also do this easily

Market category analysis

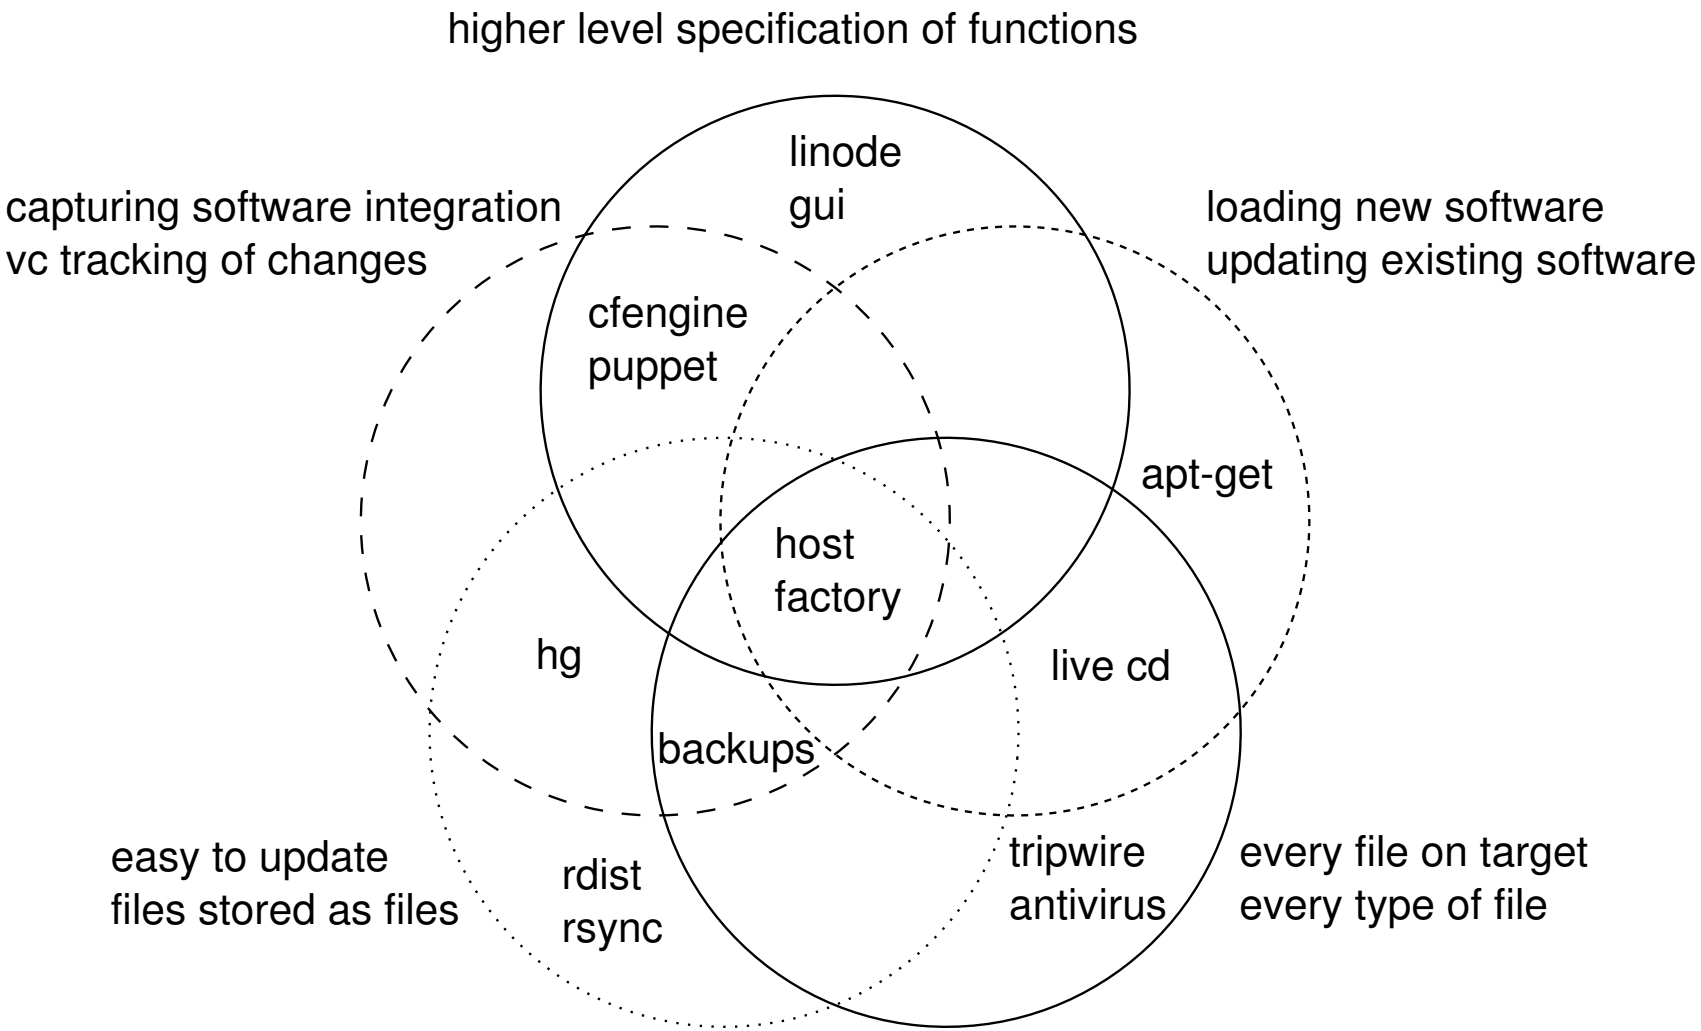

Market category analysis

Higher level specification of functions

- linode gui
- cfengine
- puppet

except ...

Market category analysis

Loading new software, updating existing software

- apt-get
- live cd

except ...

Market category analysis

Capturing software integration, vc tracking of changes

- cfengine
- puppet
- hg
- backups

except ...

Market category analysis

Every file on target, every type of file

- tripwire
- antivirus
- live cd
- backups

except ...

Market category analysis

Easy to update, files stored as files

- rdist
- rsync
- hg
- backups

except ...

Closed-loop with negative feedback

Open-loop differential updates aren't reliable and will lose position. Murphy's law requires closed-loop, otherwise host correctness is vulnerable to single error in single version of single package, or a powered-off host missing an update.

Closed-loop provides rollback capability. Lacking the ability to roll back, sysadmins develop unrealistic expectations of every other piece in the sysadmin software development waterfall such that the rollback capability will never be needed. *Testing will be perfect, so bugs will never be pushed onto production where they would need to be completely located and rolled back. Upstream vendor packages will be perfectly complete and accurate and handle every dependency, so pushing an old package will roll back every new change,* and so on.

Feature requirements

1. Solving 80% of the problem isn't good enough.
2. Audit running hosts without disrupting them.
3. Predict changes in advance of making them.
4. Detect changes made in unanticipated places.
5. Audit, correction, and update of master pattern are variations of same mechanism.
6. Intermediate stages of generation are inspectable.
7. Use real SQL database for database, not flat file.
8. Use real version control system for version control.
9. Store 500,000 files as file tree, not named in flat file.
10. Keep upstream package system working, don't duplicate package rollout/upgrade logic.

Host Factory block diagram

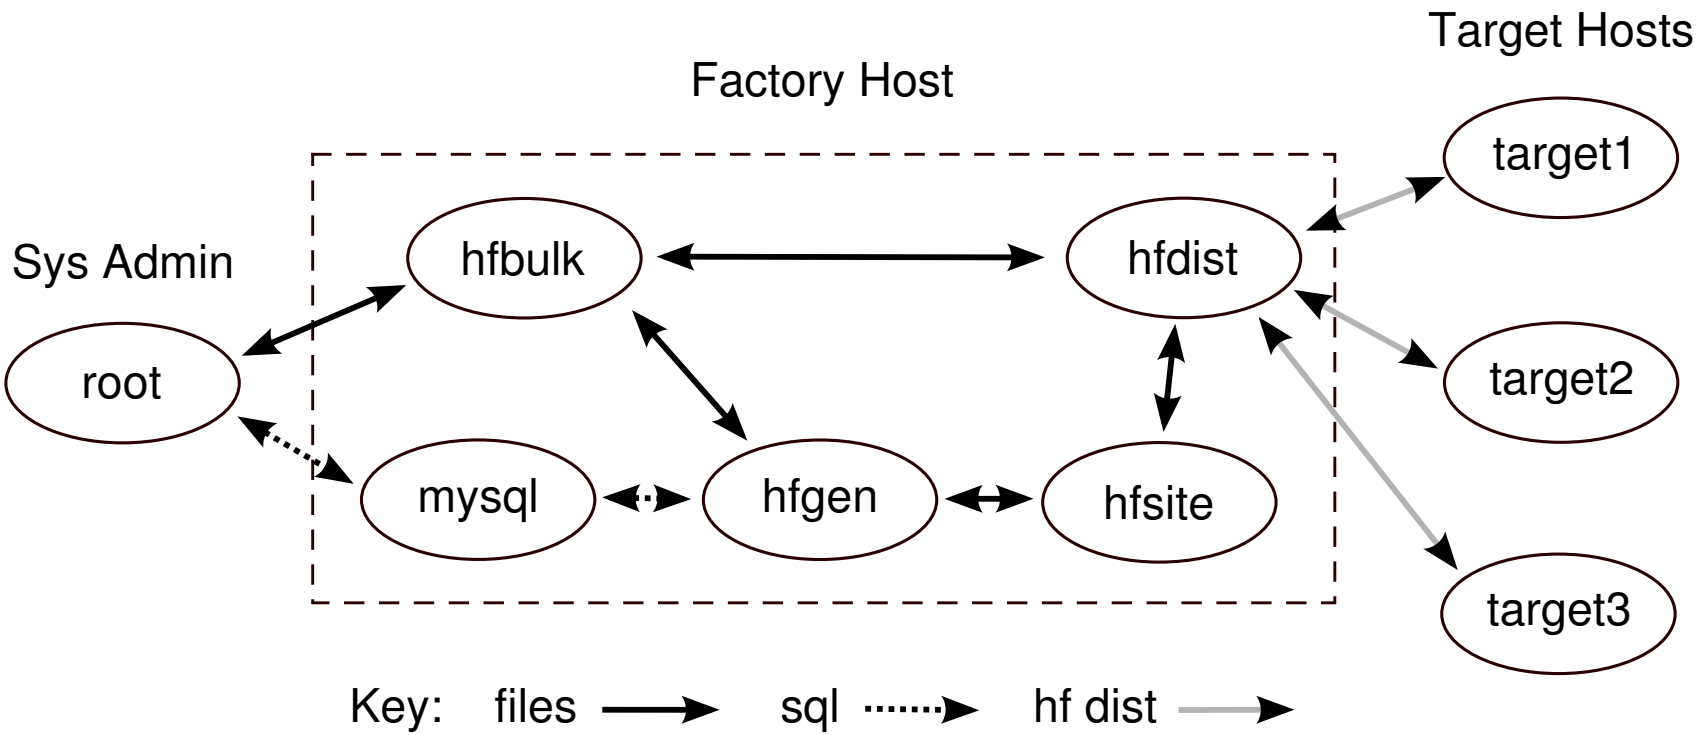

Commands implementing blocks

- Rdist-like remote filesystem update program, with expanded exception system to avoid marking some differences as errors, inspired by layers and transparency of gimp, named `hfdist`.
- Mercurial version control for all file types, named `oh`.
- /etc files and control of rdist-like program generated as reports from a SQL database, named `hfgen`.
- Framework to check hosts, named `hfcheck`.
- Framework to run `hfcheck` jobs in parallel to groups of hosts, named `oe`, `oel`, `oewatch`.
- Versioning of SQL database, named `lsdb`, `forkdb`, `diffdb`, `dropdb`.

Classes of files

- `/bin/ls`
  Identical on every host
- `/etc/hostname`
  Different on every host, predictable contents
- `/dev/tty`, `/var/log/syslog`
  Some stats identical, some different; different stats and contents are unpredictable
- `/var/spool/mqueue/foo123`
  Some stats predictable, names and contents unpredictable
- `/home/bb/Downloads/foo.pdf`
  Not under control of Host Factory. Managed by applications, user, and backups

File classes get hfdist treatments

- `/bin/ls`

  000: Copy identically.
- `/opt/hf`

  101: For HF development, don't check owner/group.
- `/dev/tty`

  102: Ignore owner/group/mode/mtime.
- `/etc/hostname`

  103: Take file content from one layer, file content from other layer, merge the two.
- `/run/resolvconf/resolv.conf`

  113: Like 103, but don't check mtime because target re-writes it.

File classes get hfdist treatments

- `/dev`, `/var/log`
  104: Don't check size/mtime.
- `/var/log/dmesg.1.gz`
  105: Like 104, no error if nonexist.
- `/etc/ssh/ssh_host_ecdsa_key`
  106: Only update files.
- `/dev/input/by-path/*`
  108: No check symlink targets.
- For tidying up mtime differences from apt-get
  109: Match up filenames and update mtimes.
- `/var/spool/mqueue/123`
  207: Get name from target walk, generate owner/group/mode, check it.

hfdist layers terms

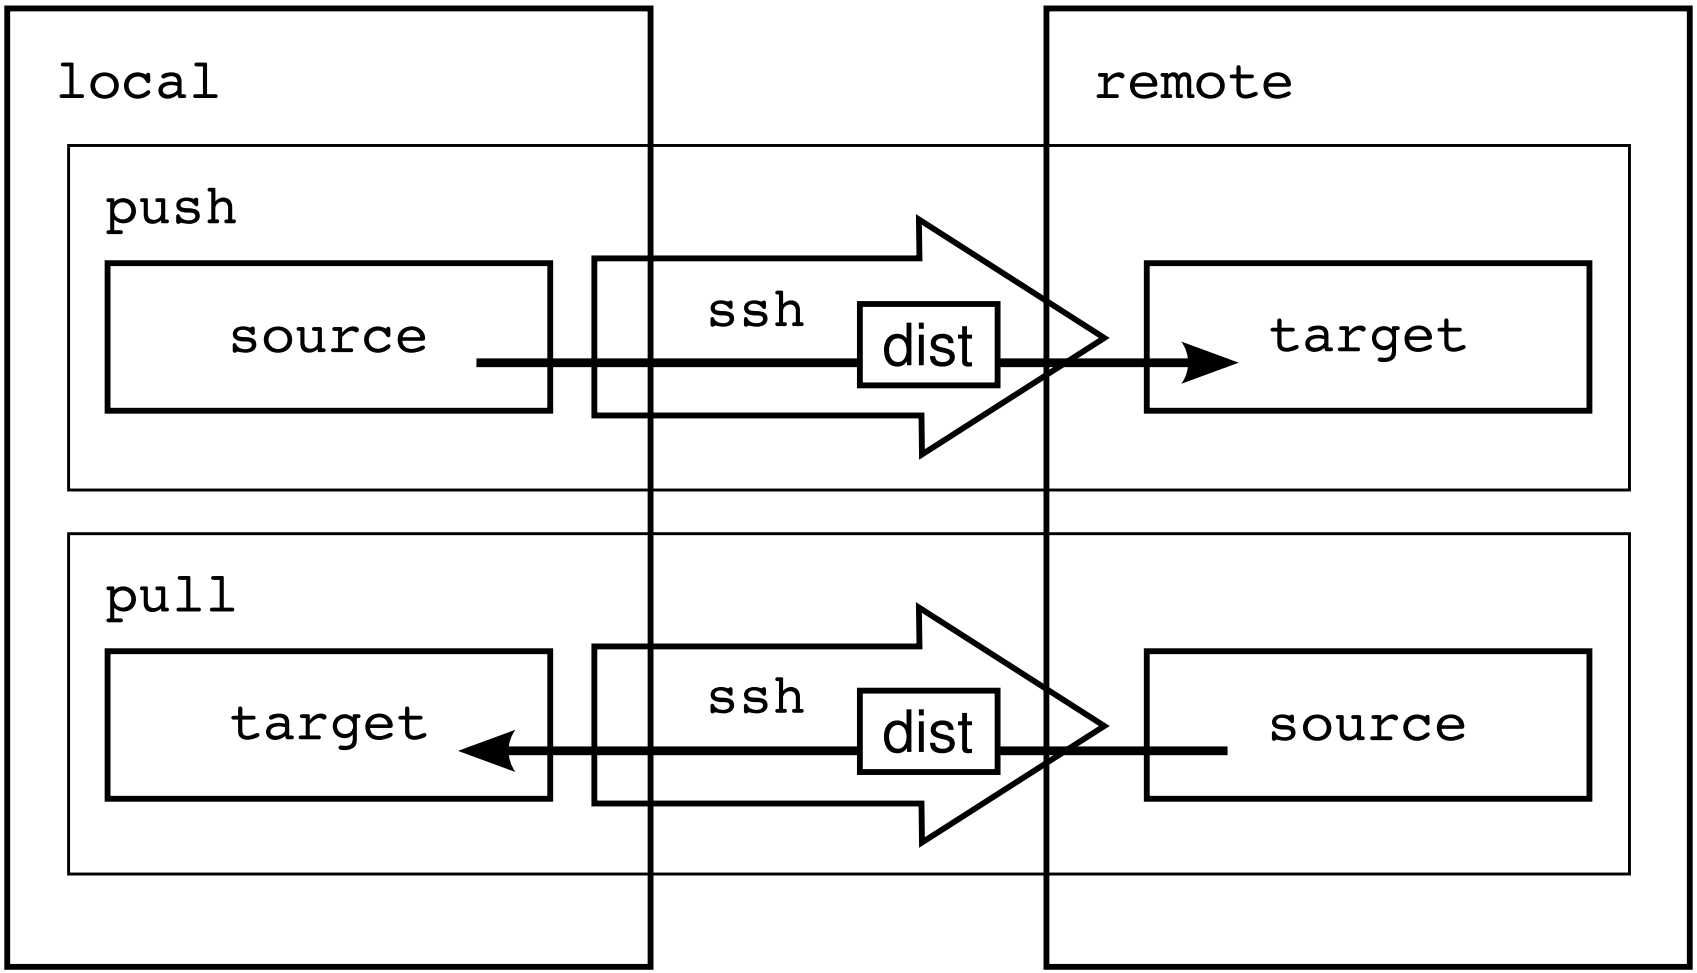

hfdist layers stack

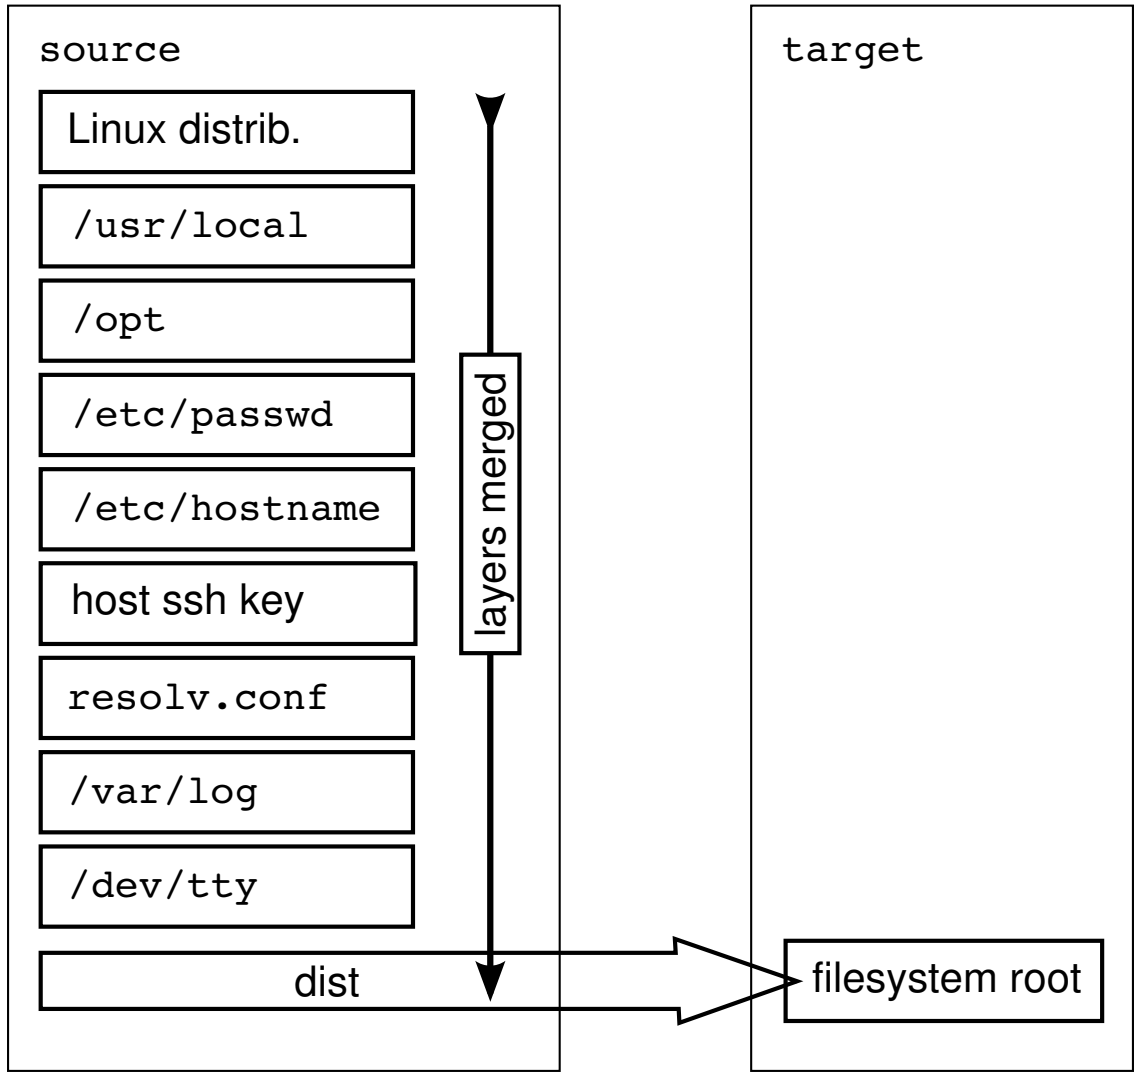

hfdist layers example

```
{
'source_layers' => [

        # source layer 0
        {
                'mount' => [ '', '/opt/hfbulk/main3/m/g', 'dirwalk' ],

                'filetree_skip' => [

                        # layer glue
                        '^depot$',
                        '^dev$',
                        '^opt$',
                        '^var/www$',

                        # unpredictable insides
                        '^var/cache/bind(/|$)',
                        '^var/cache/cups(/|$)',
                        '^var/lib/alsa(/|$)',
                        '^var/lib/amavis(/|$)',
                        '^var/lib/clamav(/|$)',
                        '^var/lib/dhcp(/|$)',
                        '^var/lib/dovecot(/|$)',
                        '^var/lib/lightdm(/|$)',
                        '^var/lib/logrotate(/|$)',
                        '^var/lib/mailman(/|$)',
                        '^var/lib/mlocate(/|$)',
                        '^var/lib/mysql(/|$)',
                        '^var/lib/NetworkManager(/|$)',
```

Inventory database schema

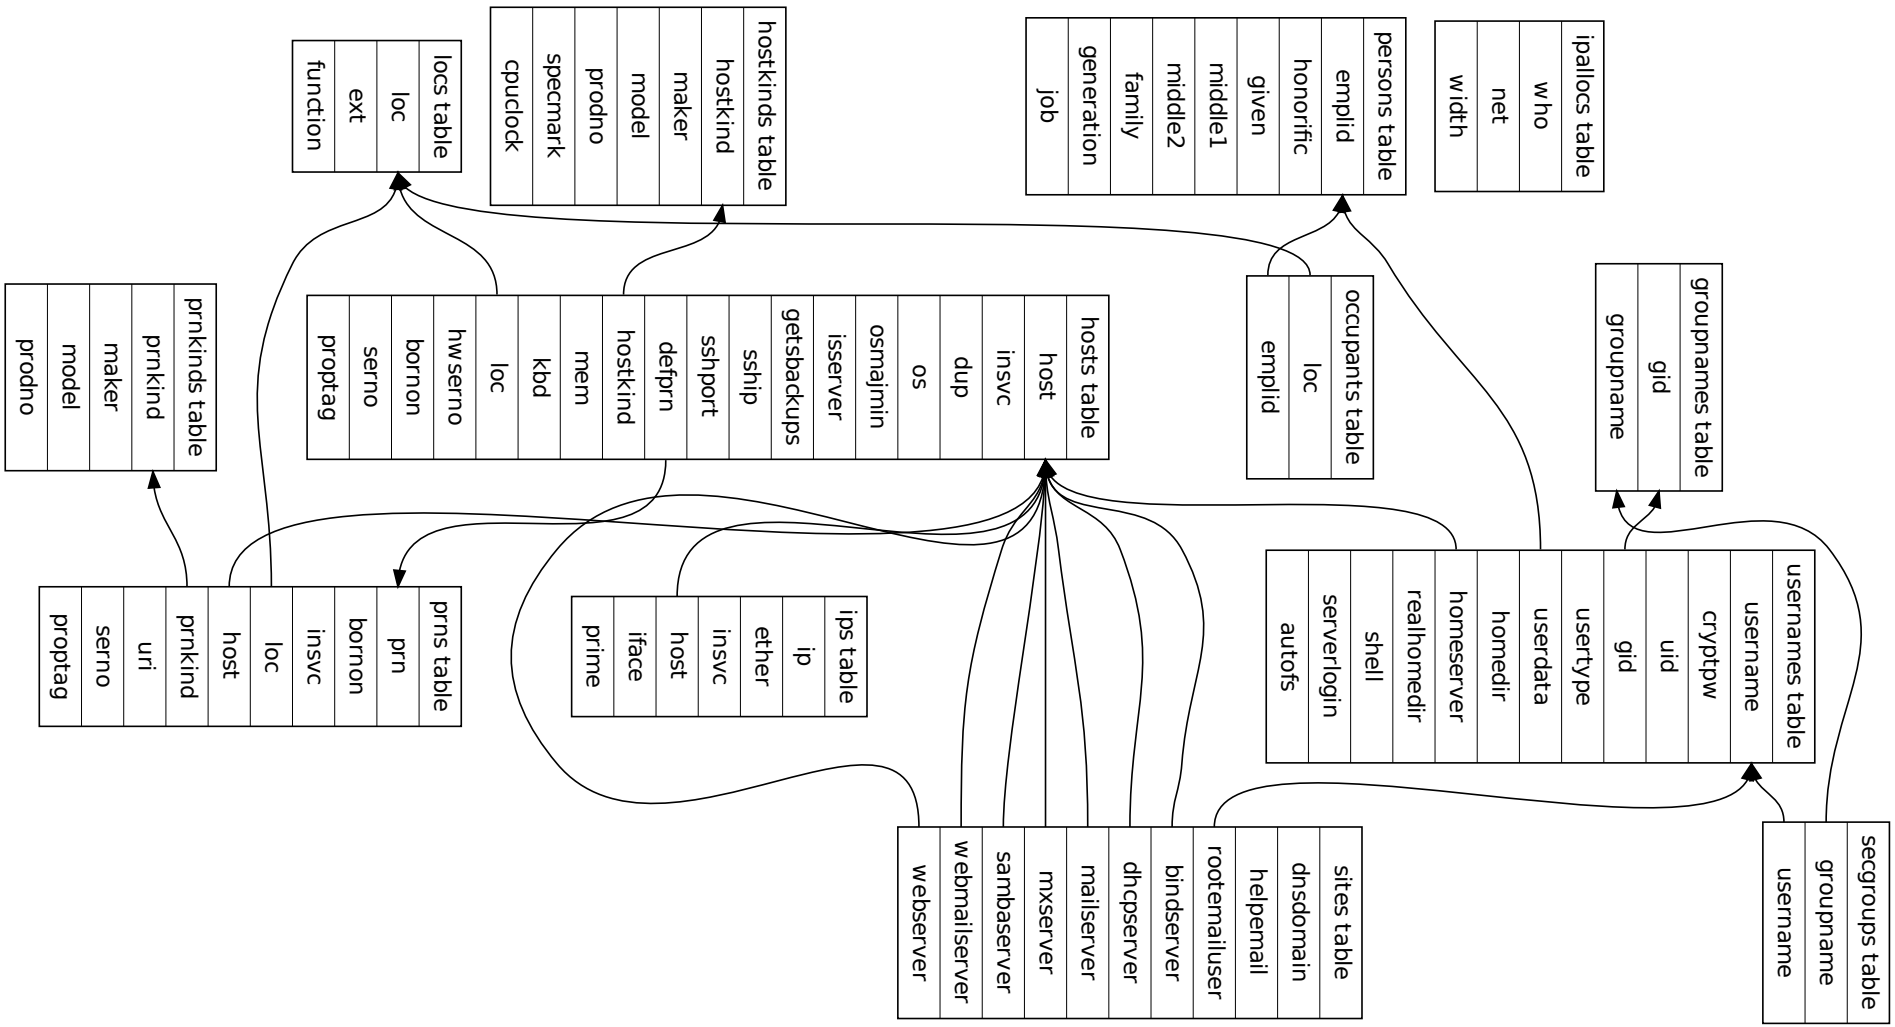

Getting /dev into Mercurial

Mercurial tracks file contents. It does not track empty directories, char devices, block devices, pipes, sockets, hard links, or suid modes. It also rejects Mercurial repositories.

Therefore, every non-tracked detail is recorded, deleted, and restored by a script which runs before and after Mercurial. This process is managed by a wrapper script called `oh`.

You can put anything you want into `oh`, including Mercurial repostories and `oh` repositories, except files larger than 2 Gig.

Running jobs in parallel

The program `oe` (On Each) runs hfdist and other jobs in parallel. The lists of hostnames for `oe` come out of the database, and there are other ways in which it integrates.

The program `oel` (On Each List) generates and manipulates lists of hostnames, including subtractions like `servers -thisserver`. `oe` accepts the same syntax.

The program `oewatch` (oe Watch) tails the logfiles of an `oe`.

Data dependencies and limits

The highest level of data independence is in hfdist. hfdist knows that a filename can contain any byte except a NUL or a forward slash. Funny characters are quoted in the error messages. There is a regression test suite which covers historical bugs, and repeatedly permuting one random filetree to another with partially overlapping filenames.

Data independence reduces outward from hfdist. If you put a space in a hostname, you will get what you deserve.

Log file output for programs to read is formatted as a data declaration structure in perl, and extractions are done with `hflog` or perl.

Multiple site models at once

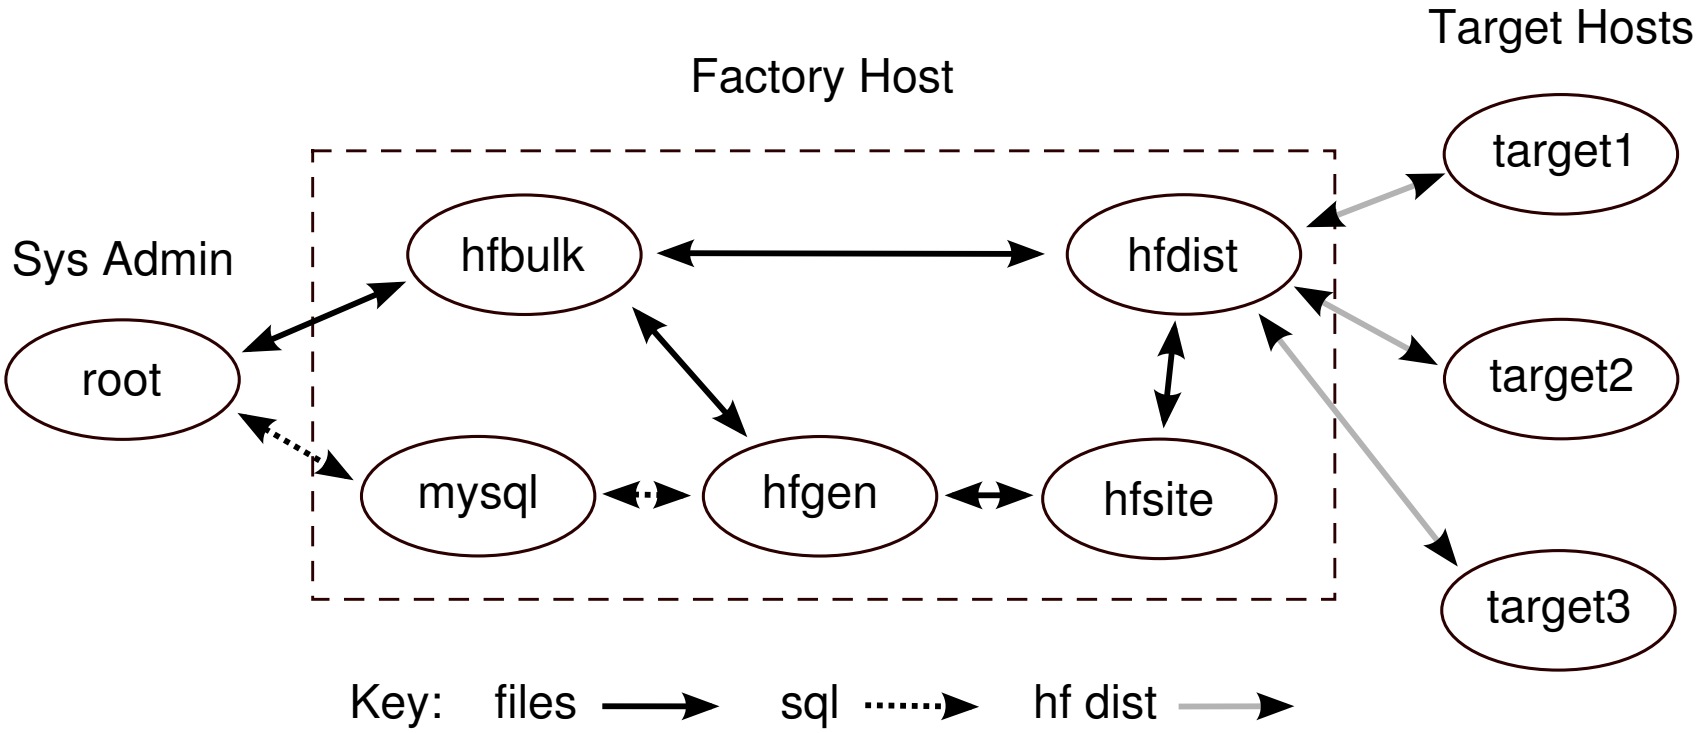

To support sysadmin software development without endangering production, each of the factory bubbles can be forked. The variant bubbles can be reconnected any way you want to create multiple independent site models.

Multiple site models at once

- `hfdist` This client-server program has an MD5 checksum in the install path, and a protocol version number.
- `hfsite` Output location selected by `HFGENDB`.
- `hfgen` Different versions in different directories. Each version can read from a different database and write to a different `hfsite` by setting `HFGENDB`.
- `hfbulk` Evolutions of filetrees like /usr are recorded as Mercurial versions within `oh`. Unrelated distributions are in different directories.
- `mysql` The SQL databases are tiny, and can be forked and diffed with `forkdb example1 example2` and `diffdb example1 example2`.

Updating the master model

When the shellshock bash bug was announced I knew my web server had not been compromised, because I regularly check every file on it. I did apt-get upgrades on my web server and two other machines. This downloaded about 200 Meg of updates onto each machine. Then I ran my scripts to collect and boil down those three instances to one uniform set of files I want everywhere. As usual, I discovered new places where apt-get allows meaningless differences. I added these to my boil down scripts, and the new boil down run completed. I pushed the merged version out to my three hosts, which mostly adjusted mtimes on files.

Performance

4.7 minutes to `hfcheck` 9 Gigabytes of 440K total files and dirs on my web server, checked from Gainesville Hackerspace to Linode in Atlanta.

17 seconds to `hfcheck -l etc` 10 Megabytes of 4K total files and dirs, from Hackerspace to Atlanta.

2.8 minutes to `oh snapshot` after trivial update to 10 Gig of 476K total files and dirs.

4.2 minutes to `oh snapshot` after 250 Meg of writes to same.

Run from my Dell Inspiron N7110 laptop containing 8 Gig of memory and Intel© Core™ i7-2670QM CPU @ 2.20GHz.

Web report by username

Accounts by Username

username	honorific	family	given	middle1	middle2	generation	job	loc	homeserver	emplid
aaa		Abrikosov	Alexei	A.			physics	ofc-105	brig	empl492
aam		Michelson	Albert	A.			physics	ofc-363	brig	empl333
acarrel		Carrel	Alexis				medicine	ofc-371	brig	empl157
aclaude		Claude	Albert				medicine	ofc-127	fluke	empl254
adh		Hershey	Alfred	D.			medicine	ofc-252	fluke	empl243
adolfb		Butenandt	Adolf				chemistry	ofc-385	fluke	empl039
adolfv		von Baeyer	Adolf				chemistry	ofc-109	brig	empl005
afc		Cournand	Andre	F.			medicine	ofc-362	fluke	empl211
afh		Huxley	Andrew	F.			medicine	ofc-409	fluke	empl228
afleming	Sir	Fleming	Alexander				medicine	ofc-371	batten	empl191
agg		Gilman	Alfred	G.			medicine	ofc-114	brig	empl301
agm		MacDiarmid	Alan	G.			chemistry	ofc-243	fluke	empl134
aharden		Harden	Arthur				chemistry	ofc-117	brig	empl026
ahc		Compton	Arthur	H.			physics	ofc-235	brig	empl355

Web report by room

Accounts by Room

loc	function	ext	honorific	given	middle1	middle2	family	generation	job	host	username	homeserver
ofc-101	Office	121		J.	Michael		Bishop		medicine	abaft	jmb	brig
ofc-101	Office	121		Paul	C.		Lauterbur		medicine	abaft	pcl	fluke
ofc-101	Office	121		Rita			Levi-Montalcini		medicine	abaft	rital	brig
ofc-101	Office	121		Kary	B.		Mullis		chemistry	abaft	kbm	fluke
ofc-101	Office	121		Louis			Neel		physics	abaft	lneel	brig
ofc-101	Office	121		Harold	C.		Urey		chemistry	abaft	hcu	batten
ofc-102	Office	107		Sidney			Altman		chemistry	abeam	saltman	brig
ofc-102	Office	107		Edmond	H.		Fischer		medicine	abeam	ehf	fluke
ofc-102	Office	107		Francois			Jacob		medicine	abeam	fjacob	fluke
ofc-102	Office	107		Theodore	W.		Richards		chemistry	abeam	twr	fluke
ofc-102	Office	107		Thomas	H.		Weller		medicine	abeam	thw	batten
ofc-103	Office	349		Edward	A.		Doisy		medicine	beam	ead	brig
ofc-103	Office	349		Niels	K.		Jerne		medicine	beam	nkj	fluke
ofc-103	Office	349		Peter			Mitchell		chemistry	beam	mitchell	fluke

Web report by host

Building Table

loc	function	honorific	given	middle1	middle2	family	generation	host	os	osmajmin	hostkind	mem	bornon	serno
ofc-101	Office		J.	Michael		Bishop		abaft	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	19663
ofc-101	Office		Paul	C.		Lauterbur		abaft	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	19663
ofc-101	Office		Rita			Levi-Montalcini		abaft	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	19663
ofc-101	Office		Kary	B.		Mullis		abaft	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	19663
ofc-101	Office		Louis			Neel		abaft	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	19663
ofc-101	Office		Harold	C.		Urey		abaft	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	19663
ofc-102	Office		Sidney			Altman		abeam	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	12267
ofc-102	Office		Edmond	H.		Fischer		abeam	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	12267
ofc-102	Office		Francois			Jacob		abeam	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	12267
ofc-102	Office		Theodore	W.		Richards		abeam	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	12267
ofc-102	Office		Thomas	H.		Weller		abeam	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	12267
ofc-103	Office		Edward	A.		Doisy		beam	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	45486
ofc-103	Office		Niels	K.		Jerne		beam	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	45486
ofc-103	Office		Peter			Mitchell		beam	Ubuntu	12.04.5 LTS	PC	2048	2013-03-10	45486

Web report on floorplan

Headquarters 1st floor

107 Office x365 amc jhn klorenz nikolaas skb bow 107prn	109 Office x202 adolfv efermi fjoliot ghooft thunt bowline	115 Men's RR	108 Janitor	110 Office x360 brm eac hgd hkh mwilkins vfh bowsprit
105 Office x279 aaa fcr freds gporter hkroemer jew boat				112 Office x166 cda cgb emoniz gwittig wms bridge
103 Office x349 ead mitchell nkj res vprelog beam		106 Conference x139 boom		113 Office x220 barujb louisd mtheiler pberg peterm bulkhead
102 Office x107 ehf fjacob saltman thw twr abeam				114 Office x297 agg cde duve gds hos juliusw rjr
101 Office x121 hcu jmb kbm lneel pcl rital abaft				126 Office x468 gcori hev hsg irvingl scct clew
104 Reception x329 prous belay				120 Elevator

Screenshot of hfcheck

```
# hfcheck -wd -l var main-linode
NOTICE hfgen.out is not in cd,
NOTICE doing chdir to /opt/hf/hf-8.0-12345678901234567890123456789012/gen
# STATUS hostnames came from: literal main-linode
STATUS starting host main-linode numbered  0 originally named main-linode
CHECKPOINT waiting on 1 children with results in /tmp/oe.11321
 main-linode 11323
STATUS pid 11323 host main-linode exited with value 0
STATUS exits 0
STATUS inputs more /tmp/oe.11321/{commandtail,hostlist,stdin}
STATUS outputs more /tmp/oe.11321/done /tmp/oe.11321/hosts/*/oe_{exit,pid,stdout,stderr}
STATUS outputs retained, to delete do rm -r /tmp/oe.11321
---------- logfilename /tmp/oe.11321/hosts/main-linode/log.pl
----- host 1 direction 'push' channel 'ssh' ip '10.20.30.40'
[... hflog output ...]
stdout and exits are:
main-linode      HFDIST_DIFF_TOTAL=414      HFDIST_FAIL_TOTAL=0        exit 0
bash commands for detailed logs:
hflog /tmp/oe.11321/hosts/main-linode/log.pl | less
hflog /tmp/oe.11321/hosts/*/log.pl | less
less /tmp/oe.11321/hosts/*/oe_stderr
To cleanup rm -r /tmp/oe.11321
```

Screenshot of oe

```
$ oe 'server -batten newhost' cat /etc/hostname
# STATUS hostnames came from: file /opt/hf/hf-8.0-12345678901234567890123456789012/gen/hfgen.out/example1/
# STATUS hostnames came from: neg literal batten
# STATUS hostnames came from: literal newhost
STATUS starting host brig    numbered  0 originally named brig
STATUS starting host cat     numbered  1 originally named cat
STATUS starting host fluke   numbered  2 originally named fluke
STATUS starting host gear    numbered  3 originally named gear
STATUS starting host newhost numbered  4 originally named newhost
CHECKPOINT waiting on 5 children:
        brig 32635           cat 32636         fluke 32638          gear 32640
     newhost 32642
```

Screenshot of oewatch

```
# oewatch
[...]
main-linode  sev 0 loc | PROGRESS total sent 186.3 MB; 36360000 of 42504005 bytes, 85% of var/cache/apt/sr
main-linode  sev 0 loc | PROGRESS total sent 190.5 MB; 40550000 of 42504005 bytes, 95% of var/cache/apt/sr
main-linode  sev 0 loc | PROGRESS total sent 194.4 MB; 1970000 of 3114987 bytes, 63% of var/lib/apt-xapian
main-linode  sev 0 loc | PROGRESS total sent 198.3 MB; 1560000 of 2862813 bytes, 54% of var/lib/apt/lists/
main-linode  sev 0 loc | PROGRESS total sent 202.2 MB; 2620000 of 3149441 bytes, 83% of var/lib/apt/lists/
main-linode  sev 0 loc | PROGRESS total sent 206.4 MB; 210000 of 582513 bytes, 36% of var/lib/apt/lists/se
main-linode  sev 0 loc | PROGRESS total sent 210.3 MB; 2300000 of 5783153 bytes, 39% of var/lib/apt/lists/
main-linode  sev 0 loc | PROGRESS total sent 214.4 MB; 590000 of 6085813 bytes, 9% of var/lib/apt/lists/us
main-linode  sev 0 loc | PROGRESS total sent 218.4 MB; 4640000 of 6085813 bytes, 76% of var/lib/apt/lists/
main-linode  sev 0 loc | PROGRESS total sent 222.2 MB; 2350000 of 3461778 bytes, 67% of var/lib/apt/lists/
main-linode  sev 0 loc | PROGRESS total sent 226.2 MB; 120000 of 203077 bytes, 59% of var/lib/apt/lists/us
main-linode  sev 0 loc | PROGRESS total sent 230.4 MB; 400000 of 455082 bytes, 87% of var/lib/apt/lists/us
main-linode  sev 0 rem | PROGRESS fixing hardlinks
main-linode  sev 0 rem | PROGRESS fixing dir mtimes
main-linode  sev 0 rem | PROGRESS sending reports and logs
main-linode  sev 0 rem | PROGRESS remote shutting down
main-linode  sev 0 loc | PROGRESS local shutting down
main-linode  HFDIST_DIFF_TOTAL=414
main-linode  HFDIST_FAIL_TOTAL=0
```

Screenshot of hflog

```
# hflog /tmp/oe.11321/hosts/main-linode/log.pl
---------- logfilename /tmp/oe.11321/hosts/main-linode/log.pl
----- host 1 direction 'push' channel 'ssh' ip '10.20.30.40'
what     why       want        have        t  name
-------  --------  ----------  ----------  -  --------------------------------
fix      mtime     2147397246  2147398480  d  var/cache/apt
fix      mtime     2147397246  2147398480  -  var/cache/apt/pkgcache.bin
replace  size        42566659    42710616  -  var/cache/apt/pkgcache.bin
fix      mtime     2147397246  2147398480  -  var/cache/apt/srcpkgcache.bin
replace  size        42504005    42620443  -  var/cache/apt/srcpkgcache.bin
fix      mtime     2147397246  2147398480  d  var/cache/apt-xapian-index
create   dir                               d  var/cache/apt-xapian-index/index.1
create   file                              -  var/cache/apt-xapian-index/index.1/flintlock
create   file                              -  var/cache/apt-xapian-index/index.1/iamchert
create   file                              -  var/cache/apt-xapian-index/index.1/postlist.DB
create   file                              -  var/cache/apt-xapian-index/index.1/postlist.baseA
create   file                              -  var/cache/apt-xapian-index/index.1/postlist.baseB
create   file                              -  var/cache/apt-xapian-index/index.1/record.DB
create   file                              -  var/cache/apt-xapian-index/index.1/record.baseA
create   file                              -  var/cache/apt-xapian-index/index.1/record.baseB
create   file                              -  var/cache/apt-xapian-index/index.1/spelling.DB
create   file                              -  var/cache/apt-xapian-index/index.1/spelling.baseA
```

Screenshot of hflog -l

```
# hflog -l /tmp/oe.11321/hosts/main-linode/log.pl
---------- logfilename /tmp/oe.15633/hosts/main-linode/log.pl
----- host 1 direction 'push' channel 'ssh' ip '10.20.30.40'
what     why       want        have        t  name
-------  --------  ----------  ----------  -  --------------------------------
         t  mode lnk user group size/rdev       mtime      mtime                      target/checksum
         - ----- -- ----- ----- -------------- ---------- ------------------------ ---------------
fix      mtime     2147397246  2147398480  d  var/cache/apt
    want d 00755  3     0     0                2147397246 Sun Jan 17 22:14:06 2038
    have d 00755  3     0     0                2147398480 Sun Jan 17 22:34:40 2038
fix      mtime     2147397246  2147398480  -  var/cache/apt/pkgcache.bin
    want - 00644  1     0     0      42,566,659 2147397246 Sun Jan 17 22:14:06 2038
    have - 00644  1     0     0      42,711,624 2147398480 Sun Jan 17 22:34:40 2038
replace  size        42566659    42711624  -  var/cache/apt/pkgcache.bin
    want - 00644  1     0     0      42,566,659 2147397246 Sun Jan 17 22:14:06 2038
    have - 00644  1     0     0      42,711,624 2147398480 Sun Jan 17 22:34:40 2038
fix      mtime     2147397246  2147398480  -  var/cache/apt/srcpkgcache.bin
    want - 00644  1     0     0      42,504,005 2147397246 Sun Jan 17 22:14:06 2038
    have - 00644  1     0     0      42,621,901 2147398480 Sun Jan 17 22:34:40 2038
replace  size        42504005    42621901  -  var/cache/apt/srcpkgcache.bin
    want - 00644  1     0     0      42,504,005 2147397246 Sun Jan 17 22:14:06 2038
    have - 00644  1     0     0      42,621,901 2147398480 Sun Jan 17 22:34:40 2038
```

Screenshot of hflog -s

```
# hflog -s /tmp/oe.13166/hosts/main-linode/log.pl
---------- logfilename /tmp/oe.13166/hosts/main-linode/log.pl
----- host 1 direction 'push' channel 'ssh' ip '10.20.30.40'
dowrite 0 dodelete 0 dochecksums 0 donotfixhardlinks 0 opt_x 0 euid 0
fail_no_remote_answer 0 fail_wrong_remote_version 0 checkreadiness 1 fail_wire_read 0 fail_wire_write 0

time_begin 2147397246 Sun Jan 17 22:14:06 2038
time_end   2147397528 Sun Jan 17 22:18:48 2038  seconds 282

             file       dir   symlink      pipe   chardev    blkdev    socket       all    file byte seen
source    320,556    35,041    84,170         0       646        43         1   440,457     9,383,407,634
target    320,556    35,041    84,170         0       646        43         1   440,457     9,383,408,017

   want byte sent want file    file byte sent file sent   fixlink dir mtime  diff tot  fail tot
                0         0                 0         0         0         0         0         0
```

Screenshot of hflog -c

```
# hflog -c /tmp/oe.11321/hosts/main-linode/log.pl
---------- logfilename /tmp/oe.11321/hosts/main-linode/log.pl
----- host 1 direction 'push' channel 'ssh' ip '10.20.30.40'
size            name
--------------- ----
     55,328,768 var/cache/apt-xapian-index/index.3/postlist.DB
     55,304,192 var/cache/apt-xapian-index/index.1/postlist.DB
     45,146,112 var/cache/apt-xapian-index/index.3/termlist.DB
     44,523,520 var/cache/apt-xapian-index/index.1/termlist.DB
     42,566,659 var/cache/apt/pkgcache.bin
     42,504,005 var/cache/apt/srcpkgcache.bin
      6,085,813 var/lib/apt/lists/us.archive.ubuntu.com_ubuntu_dists_precise-updates_main_binary-i386_Pack
      5,783,153 var/lib/apt/lists/us.archive.ubuntu.com_ubuntu_dists_precise-updates_main_binary-amd64_Pac
      5,636,096 var/cache/apt-xapian-index/index.3/spelling.DB
      5,611,520 var/cache/apt-xapian-index/index.1/spelling.DB
      3,461,778 var/lib/apt/lists/us.archive.ubuntu.com_ubuntu_dists_precise-updates_main_i18n_Translatio
      3,149,441 var/lib/apt/lists/security.ubuntu.com_ubuntu_dists_precise-security_main_binary-i386_Packa
      3,114,987 var/lib/apt-xapian-index/cataloged_times.p
      2,862,813 var/lib/apt/lists/security.ubuntu.com_ubuntu_dists_precise-security_main_binary-amd64_Pack
      2,509,804 var/lib/apt/lists/us.archive.ubuntu.com_ubuntu_dists_precise-updates_main_source_Sources
      2,198,666 var/lib/apt/lists/security.ubuntu.com_ubuntu_dists_precise-security_main_i18n_Translation-
      2,015,232 var/cache/apt-xapian-index/index.3/record.DB
```

Screenshot of oh

```
# oh init
# (cd /dev ; tar -cf - .) | (cd oh ; tar -xpf -)
tar: ./log: socket ignored
# statlist -k oh > /tmp/before-oh.sl
# oh wrap
# oh commit save this /dev copy
# rm -r *
# oh fix
# oh unwrap
# statlist -k oh > /tmp/after-oh.sl
# diff /tmp/before-oh.sl /tmp/after-oh.sl
```

Screenshot of database versioning

```
$ forkdb example1 example2
$ mysql example2
mysql> update hosts set osmajmin = '12.04.5 LTS' where host = 'latitude';
$ diffdb example1 example2
3c3
< -- Host: localhost    Database: example1
---
> -- Host: localhost    Database: example2
170c170
< INSERT INTO `hosts` VALUES ('latitude','Y','Ubuntu','12.04 LTS','N','N','',0,'372prn','PC',2048,'PC','of
---
> INSERT INTO `hosts` VALUES ('latitude','Y','Ubuntu','12.04.5 LTS','N','N','',0,'372prn','PC',2048,'PC',
$ dropdb example2
```

Screenshot of statlist

```
$ statlist
dir     00775       2    1000    1000   2147483646                                .
file    00664       1    1000    1000   2147483646             279                GNUmakefile
file    00755       1    1000    1000   2147483646            6959                demo.sh
file    00755       1    1000    1000   2147483646            1196                dopush
file    00644       1    1000    1000   2147483646             781                initial-pull.layer
file    00664       1    1000    1000   2147483646              33                md5
file    00664       1    1000    1000   2147483646             779                module.mk
file    00644       1    1000    1000   2147483646            2199                push.layers
file    00755       1    1000    1000   2147483646             958                vagrant-ssh
$ statlist | sed $'s/\t/@/g'
dir@00775@    2@ 1000@ 1000@2147483646@           @@     @.
file@00664@    1@ 1000@ 1000@2147483646@         279@@     @GNUmakefile
file@00755@    1@ 1000@ 1000@2147483646@        6959@@     @demo.sh
file@00755@    1@ 1000@ 1000@2147483646@        1196@@     @dopush
file@00644@    1@ 1000@ 1000@2147483646@         781@@     @initial-pull.layers
file@00664@    1@ 1000@ 1000@2147483646@          33@@     @md5
file@00664@    1@ 1000@ 1000@2147483646@         779@@     @module.mk
file@00644@    1@ 1000@ 1000@2147483646@        2199@@     @push.layers
file@00755@    1@ 1000@ 1000@2147483646@         958@@     @vagrant-ssh
```

How to contact me

Brian Bartholomew ⟨bb@workver.com⟩

http://www.workver.com/Community

About Brian Bartholomew

Brian Bartholomew has been a Unix system administrator for 25 years. His first system ran SCO Xenix on a Compaq desktop PC running a 286 (not a 386). Half his work experience has been in the commercial world in Boston, and half at the University of Florida. In Boston he supported financial trading floors for mutual fund firms and banks; at UF he supported two departments of mathematicians, both theoretical and applied, and the PeopleSoft enterprise accounting system. His career speciality in version control of running Unix systems started in 1990 with rdist. The first version of the dist program appeared in 1994, written in perl 4. A Linux Journal article appeared in 1997, featuring a versioning filesystem implemented as a user space NFS server talking to PostGres.