

Hands-off testing for networked filesystems

Daria Phoebe Brashear

AuriStor, Inc

Introduction

- Distributed Filesystems
 - What AuriStor is and how it works
- Continuous integration systems
 - Buildbot, TAP, creating and running tests
- Testing from the wrong side of the kernel abstraction
- Lessons and future work

Distributed network filesystems

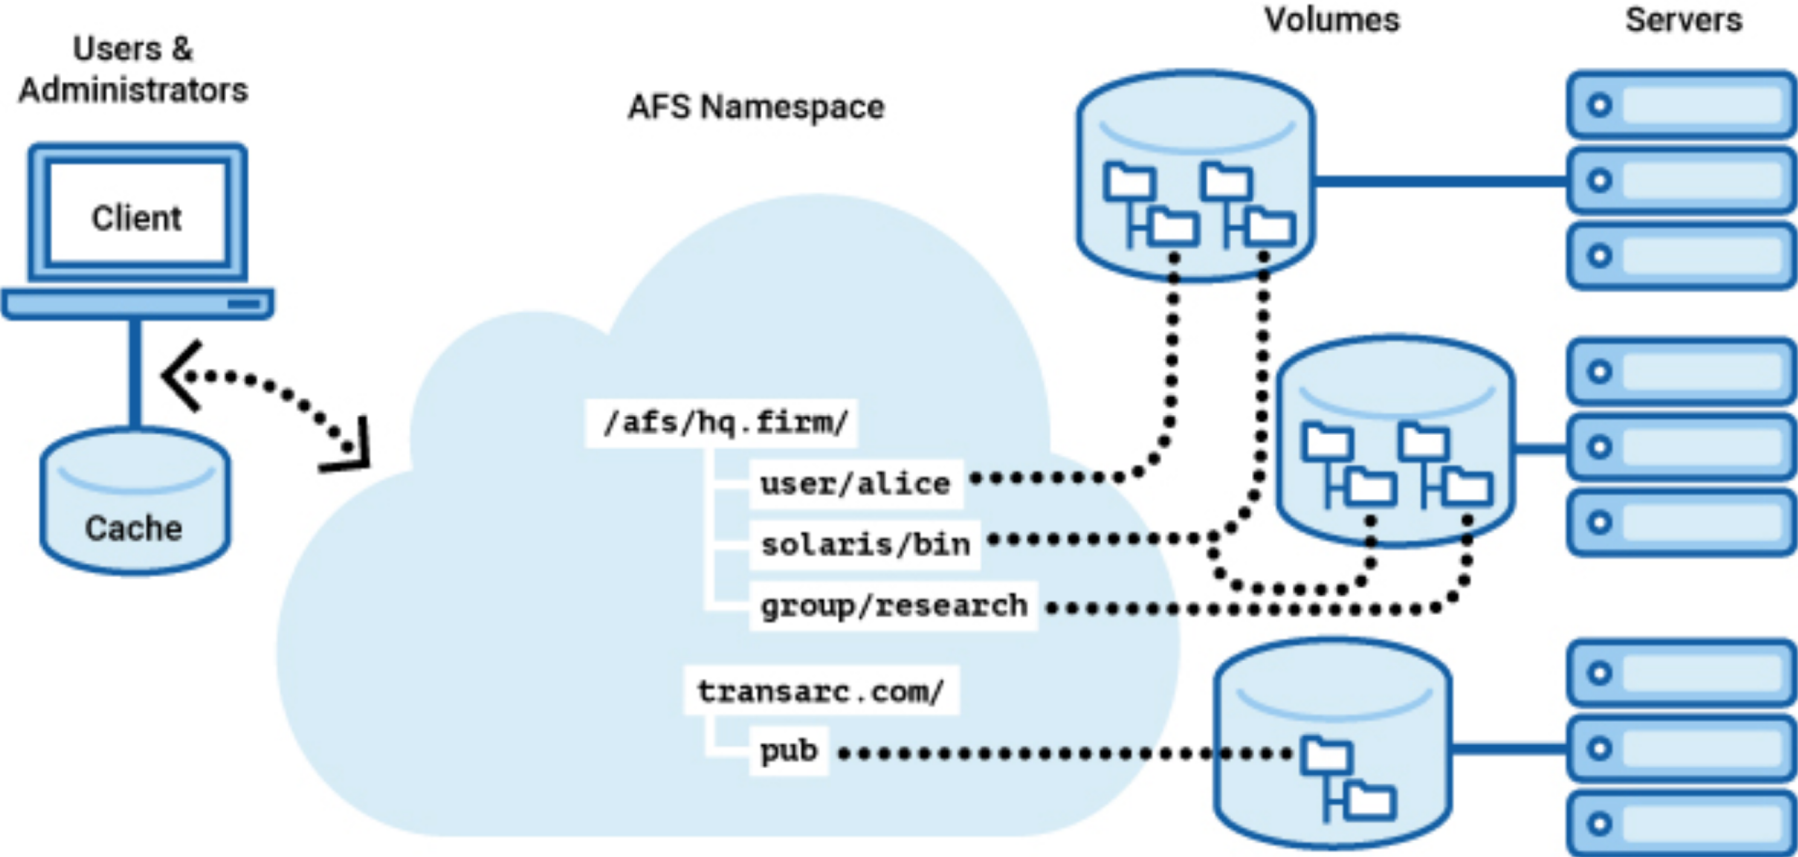
No matter where you go, there you(r files) are

- Carnegie Mellon's Andrew File System (1983)
 - Required patching the kernel
 - VFS concept introduced later, with NFS
- Sun's NFS (1984)
 - Introduced SunRPC, later adopted/co-opted by Andrew File System
- AppleShare
- SMB
- CIFS
- ZFS

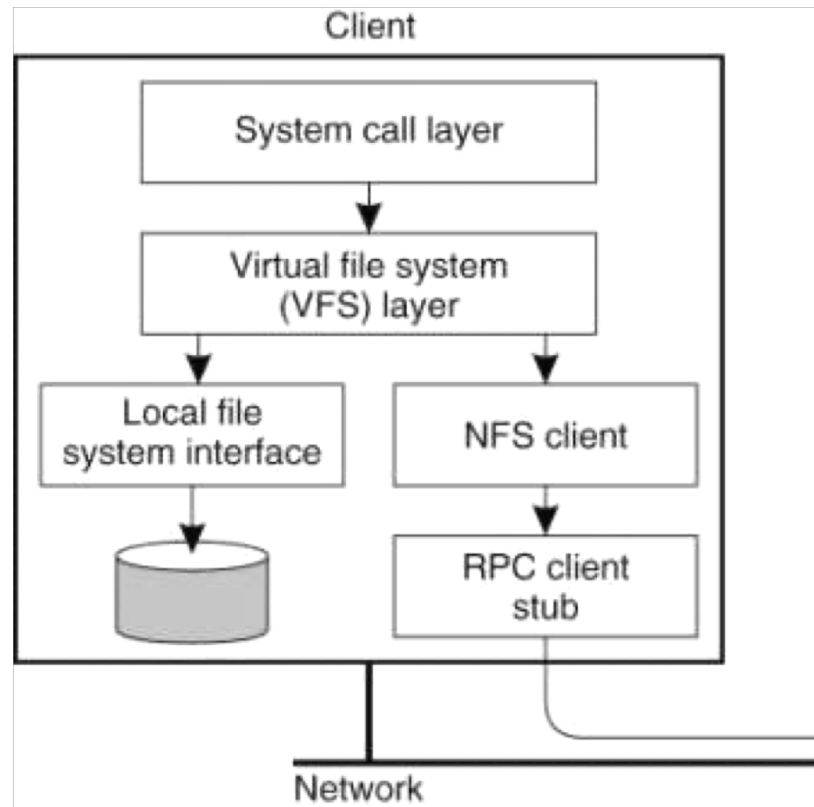
About AuriStor

- Founded in 2007
- From the lineage of OpenAFS (and IBM/Transarc AFS)
- Commercially developed and supported
 - Cross-platform (Windows, macOS, Linux, Solaris, web)
 - Highly scalable
 - Highly available
 - Secure transport
 - Granular self-administered access controls
 - Unified namespace (same path everywhere)
- In short, your files, but everywhere

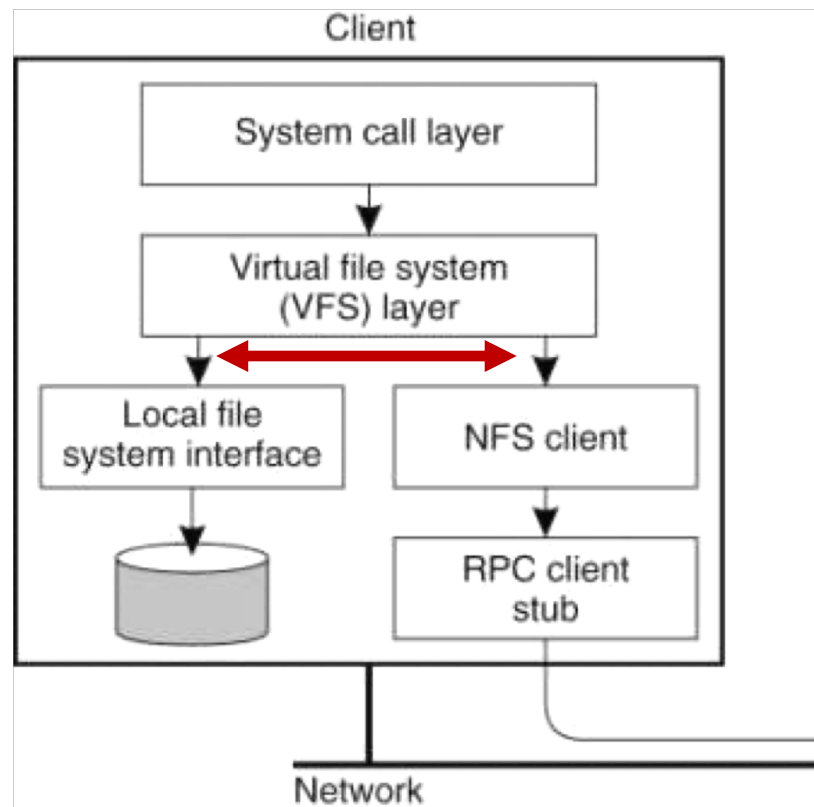
AuriStor (and AFS) architecture



Plugging into the VFS



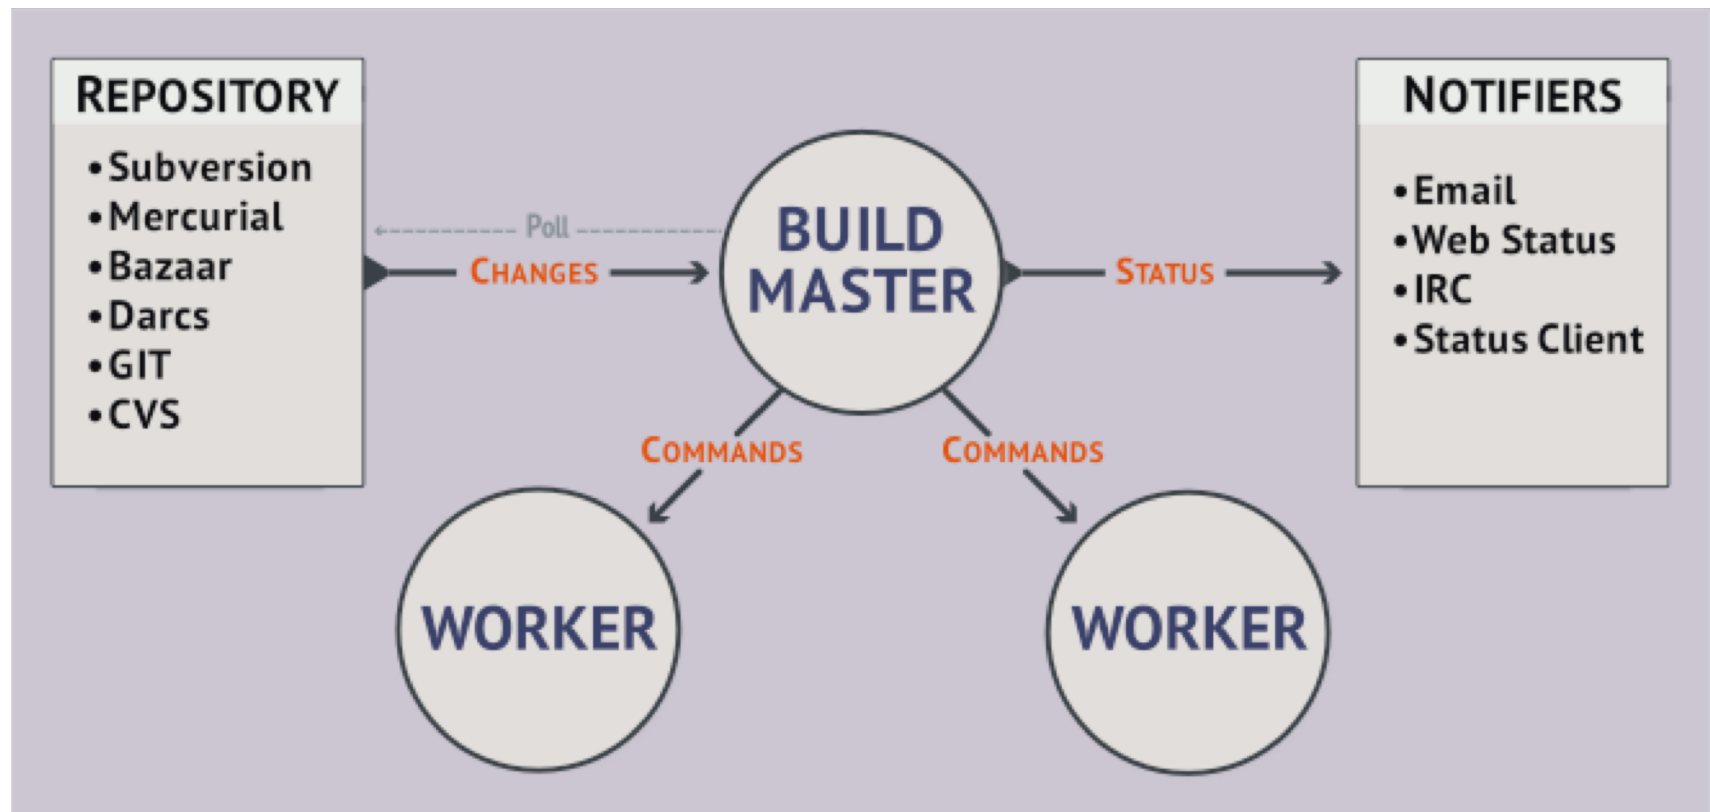
AuriStor and the VFS



Continuous integration systems

- Attached to your version control system
 - Pulls in changes
 - Pushes back results
- Basic cases
 - Compiling each change for all supported platforms
- Complex cases
 - Regression tests
 - Unit tests
- Different options depending on your needs
 - Self-hosted, cloud or hybrid?

Buildbot architecture



Buildbot web status



Buildbot with Gerrit

| Subject | Status | Owner | Assignee | Repo | Branch | Updated | Size | CR | T | V |
|--|----------------|-----------------|----------|---------|-----------------------------|----------|------|----|---|---|
| ☆ redhat: build debuginfo for kmods | -- | Daria Brashear | -- | openafs | master (redhat-kmod-deb.. | 12:48 PM | S | | | |
| ☆ rx: hide rx_endpoint prototypes when not present for KERNEL | -- | Jeffrey Altman | -- | openafs | master (rx_endpoint_asSt.. | 9:45 AM | XS | +1 | | ✓ |
| ☆ rzgk: SRXGK_GSSNegotiate handle rx_wireconn_setSpecific failures | -- | Jeffrey Altman | -- | openafs | master (rx_SetSpecific-e.. | 7:55 AM | S | +1 | | ✓ |
| ☆ viced: handle failures of getClientForConnection | -- | Jeffrey Altman | -- | openafs | master (rx_SetSpecific-e.. | 7:52 AM | XS | +1 | | ✓ |
| ☆ rx: rx_wireconn_setSpecific signature change | -- | Jeffrey Altman | -- | openafs | master (rx_SetSpecific-e.. | 7:51 AM | XS | +1 | | ✓ |
| ☆ rx: rx_SetSpecific return NULL if memory allocation failure | -- | Jeffrey Altman | -- | openafs | master (rx_SetSpecific-e.. | 7:28 AM | S | +1 | | ✓ |
| ☆ DARWIN: Update MKShim to Apple OpenSource MITKerberosShim-79 | -- | Jeffrey Altman | -- | openafs | master (Apple-MITKerber.. | 12:34 AM | L | | | ✓ |
| ☆ Unix CM: print_internet_addresses no GLOCK across rx_conn_networkError | -- | Jeffrey Altman | -- | openafs | master (GLOCK-vs-RX) | May 03 | S | | | ✓ |
| ☆ auth: introduce cell default time-to-live cellservdb info value | -- | Jeffrey Altman | -- | openafs | master (cellservdb-defaul.. | May 03 | M | | | ✓ |
| ☆ volser: Add support for volume features | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | L | -1 | | ✓ |
| ☆ volser: Add library plumbing for features | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | M | +1 | | ✓ |
| ☆ volser: Add volume feature RPCs | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | M | -1 | | ✓ |
| ☆ viced: Enforce file and hard link features | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | M | -1 | | ✓ |
| ☆ vos: Add a -features option to vos create | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | S | +1 | | ✓ |
| ☆ volser: Add an RPC to create a volume and set its features | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | M | +1 | | ✓ |
| ☆ vos: Add selffeatures command to set volume features | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | L | -1 | | ✓ |
| ☆ vos: Add listfeatures command to list a volume's current feature set | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | M | -1 | | ✓ |
| ☆ volser: Add the -features option | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | M | +1 | | ✓ |
| ☆ volser: Add support for dumping volume features | -- | Simon Wilkinson | -- | openafs | master (vol-features) | May 03 | M | -1 | | ✓ |
| ☆ DARWIN: AFSPreference correct asprintf format string | -- | Jeffrey Altman | -- | openafs | master (DARWIN) | May 02 | XS | | | ✓ |
| ☆ AFSBackgrounder: define OTHER_CFLAGS | -- | Jeffrey Altman | -- | openafs | master (DARWIN) | May 01 | XS | | | ✗ |
| ☆ platform/DARWIN: remove unused variables | -- | Jeffrey Altman | -- | openafs | master (DARWIN) | May 01 | XS | | | ✓ |
| ☆ memcmp, strcmp, and friends do not return a boolean value | -- | Jeffrey Altman | -- | openafs | master (comparison-func.. | Apr 27 | XL | | | ✓ |
| ☆ rx: extend acknowledgements from 255 to 2040 using bit fields | -- | Jeffrey Altman | -- | openafs | master (rx-extended-sack) | Apr 26 | M | -1 | | ✓ |
| ☆ rx: compute and advertise per-connection maxSkew | Merge Conflict | Jeffrey Altman | -- | openafs | master (rx-extended-sack) | Apr 26 | M | -1 | | ✓ |

Buildbot Verified +1 9:04 AM ^

Patch Set 5: Verified+1

Build success
: Win64-Free
: <https://buildbot.auristor.com/builders/Win64-Free/builds/19055>

Build success
: Win64-Checked
: <https://buildbot.auristor.com/builders/Win64-Checked/builds/19063>

Build success
: Win32-Free
: <https://buildbot.auristor.com/builders/Win32-Free/builds/18756>

Build success
: Win32-Checked
: <https://buildbot.auristor.com/builders/Win32-Checked/builds/18762>

Build success
: Fedora32_64
: https://buildbot.auristor.com/builders/Fedora32_64/builds/6107

Build success
: Fedora32_64libtree
: https://buildbot.auristor.com/builders/Fedora32_64libtree/builds/4555

Build success
: Fedora33_64
: https://buildbot.auristor.com/builders/Fedora33_64/builds/3253

Build success
: Fedora34_64
: https://buildbot.auristor.com/builders/Fedora34_64/builds/584

Build success
: RHEL6_64
: https://buildbot.auristor.com/builders/RHEL6_64/builds/15644


Build success
: plus k

Screenshot

Running tests

- A test harness will make life easier
- **Test Anything Protocol**
 - Simple, text based protocol (from Perl, circa 1987)
 - Offers a very simple dependency model
 - Originally in Perl but other integrations exist (C-TAP-Harness)

Testing with TAP



1..3

ok 1 - we can foo

ok 2 - we can bar

not ok 3 - we can ni

Testing with TAP

```
setprograme(argv[0]);

plan(9);

tmpdir = getenv("YFS_TMPDIR");
if (!tmpdir)
    tmpdir = gettmpdir();

ok(afstest_GetAddr(NULL) != -1, "Got IP address for building config");

is_int(0, namei_testXattr(tmpdir),
      "tmpdir supports xattrs");

asprintf(&cmd, "perl -e \"use experimental\"");
is_int(0, system(cmd), "perl experimental installed");
free(cmd);

asprintf(&cmd, "perl -e \"use File::Touch\"");
is_int(0, system(cmd), "perl File::Touch installed");
free(cmd);

asprintf(&cmd, "perl -e \"use File::Which\"");
is_int(0, system(cmd), "perl File::Which installed");
```

Building out a test suite

- Start small
 - API tests
 - Simple tests first
 - Group related functionality
- Grow gradually
 - RPC client/server tests

Building up a test suite

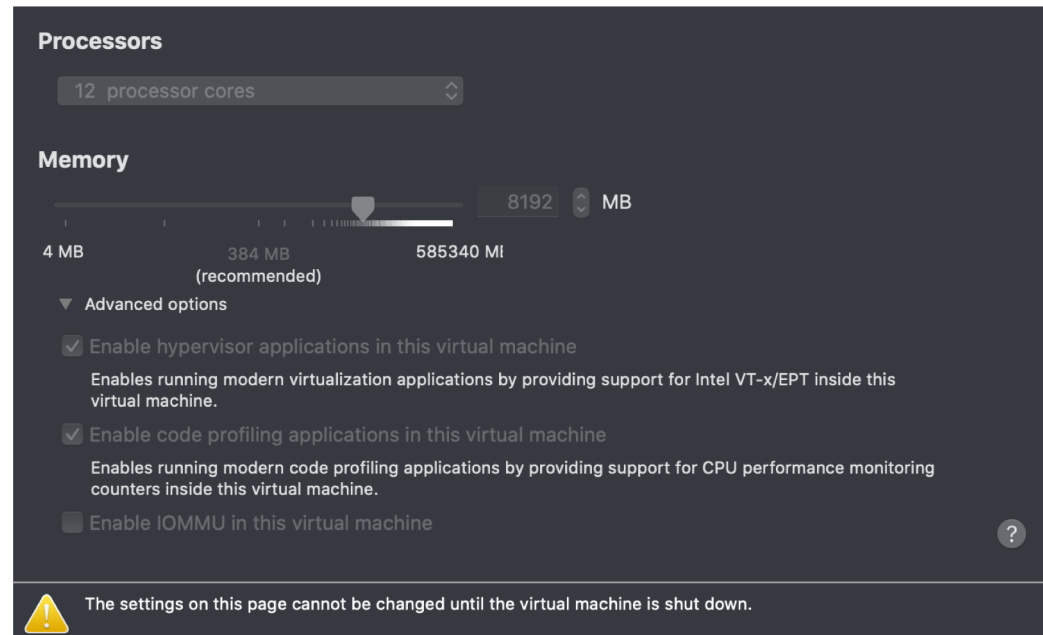
- Test everything
 - A version which spins up, locally, the entire system and manipulates it
- Test EVERYTHING
 - A client which runs in userspace can be used to test much of the client code
- TEST EVERYTHING
 - Virtualization can be used to test kernelspace client code, and multiple server nodes
- TEST EVERYTHING!!
 - Not every platform is easily virtualizable in a way which can be harnessed!

Virtualizing tests

- Docker
 - Basic multi-server functionality tests
 - Uses prebuilt OS container
 - Cross-system compat tests with OpenAFS with prebuilt OpenAFS container
- KVM
 - A client which runs in userspace can be used to test much of the client code, but you want to test the kernelspace code too!

Virtualizing tests

- Nested virtualization
 - Virtualization can be used to test kernelspace client code, and multiple server nodes



Emulating the VFS

- FUSE (Filesystem in Userspace) library and associated kernel module
 - Provides a VFS abstraction stub, calls out to userspace to do work
 - Calls the same vnode operation backend code the kernel does, BUT
 - Uses different memory allocation and pools
 - Doesn't use VFS linkage code
 - Doesn't use same locking or resource management
 - Doesn't use an in-kernel path for file caching

Testing the VFS

- Virtualization with KVM
 - Allows same-platform builds to be tested on the host
 - But it only works if you can virtualize the target with KVM
- Direct testing on virtual iron
 - Works with any platform capable of being virtualized
 - Typically X86/X64 (Linux, BSDs, Solaris, macOS)
 - VM snapshots as reset points

A typical run

- Triggered nightly
 - Tests take long enough to run that every patchset would take too long
- Full build
 - git checkout from current tip of main branch
 - Generate autoconf artifacts
 - ./configure --with...
 - Full build
 - make check

Seeing old build logs

Builder Fedora34_64Check

[\(view in waterfall\)](#)

No current builds

No Pending Build Requests

Recent Builds:

| Time | Revision | Result | Build # | Reason | Info |
|--------------|-----------------|-----------|---------|---|-----------------------|
| May 24 03:00 | d2258006abdd... | failure | #80 | The Nightly scheduler named 'nightly' triggered this build | Failed shell_3 |
| May 23 17:17 | d2258006abdd... | success | #79 | The web-page 'rebuild' button was pressed by '<unknown>': | Build successful |
| May 23 03:00 | d2258006abdd... | failure | #78 | The Nightly scheduler named 'nightly' triggered this build | Failed shell_3 |
| May 22 03:00 | d2258006abdd... | success | #77 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 21 03:00 | c8798d1f46e7... | success | #76 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 20 03:00 | f501bf424063... | success | #75 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 19 03:00 | 82a8202a443e... | failure | #74 | The Nightly scheduler named 'nightly' triggered this build | Failed shell_3 |
| May 18 03:00 | 992a4836af91... | exception | #73 | The Nightly scheduler named 'nightly' triggered this build | Exception interrupted |
| May 17 03:00 | ba3bcbc76b54... | success | #72 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 16 03:00 | ba3bcbc76b54... | failure | #71 | The Nightly scheduler named 'nightly' triggered this build | Failed shell_3 |
| May 15 03:00 | ba3bcbc76b54... | success | #70 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 14 03:00 | a1f658abfa3d... | success | #69 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 13 03:00 | 12e132c640f6... | success | #68 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 12 03:00 | fb8b17a51ca1... | success | #67 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 11 19:56 | 5173fa26c08d... | success | #66 | A build was forced by 'jaltman@auristor.com': force buildjaltman@auristor.com | Build successful |
| May 11 03:00 | e20983b02de6... | success | #65 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 10 03:00 | e20983b02de6... | success | #64 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 09 03:00 | e20983b02de6... | success | #63 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 08 03:00 | e20983b02de6... | success | #62 | The Nightly scheduler named 'nightly' triggered this build | Build successful |
| May 07 03:00 | 83cbcb713500... | failure | #61 | The Nightly scheduler named 'nightly' triggered this build | Failed shell_3 |

[Show more](#)

Drilling down

Builder Fedora34_64Check Build #77

Results:

Build successful

SourceStamp:

| | |
|--------------|--|
| Branch | master |
| Got Revision | d2258006abdd379c147e5dd0a1a21d8ebd5eac31 |

Builder:

Fedora34_64

Reason:

The Nightly scheduler named 'nightly' triggered this build

Steps and Logfiles:

1. `shell_1 'git clean ...'` (0 secs)
 1. `stdio`
2. `git update` (3 secs)
 1. `stdio`
3. `shell_2 'sh regen.sh'` (6 secs)
 1. `stdio`
4. `configure configure` (4 mins, 10 secs)
 1. `stdio`
5. `compile compile` (1 mins, 12 secs)
 1. `stdio`
6. `shell_3 extended-tests` (1 hrs, 51 mins, 5 secs)
 1. `stdio`

Build Properties:

| Name | Value | Source |
|----------------|--|-----------------------------|
| branch | master | Build |
| builddir | /home/buildbot/worker/Fedora34_64Check | Build |
| buildname | Fedora34_64Check | Builder |
| buildnumber | 77 | Build |
| codebase | | Build |
| got_revision | d2258006abdd379c147e5dd0a1a21d8ebd5eac31 | Git |
| project | | Build |
| repository | | Build |
| revision | None | Build |
| scheduler | nightly | Scheduler |
| slavename | Fedora34_64 | Builder |
| warnings-count | 0 | WarningCountingShellCommand |
| workdir | /home/buildbot/worker/Fedora34_64Check | Builder |

Forced Build Properties:

| Name | Label | Value |
|------|-------|-------|
|------|-------|-------|

Responsible Users:

no responsible users

Timing:

| | |
|---------|--------------------------|
| Start | Sat May 22 03:00:01 2021 |
| End | Sat May 22 04:56:40 2021 |
| Elapsed | 1 hrs, 56 mins, 39 secs |

Resubmit Build:

To force a build, fill out the following fields and push the 'Force Build' button

Reason for re-running build:

Rebuild using: 

Seeing the results

```
yfs: Stopping background daemons
yfs: Stopping background threads
yfs: Waiting for background threads to complete
yfs: Shutting down background threads
yfs: Stopping callback service
yfs: Stopping main daemon
yfs: Stopping check server daemon
yfs: Stopping cache truncate daemon
yfs: Stopping cellinfo upcall queue
yfs: Draining cellinfo upcall queue
yfs: Stopping thread broker
yfs: Waiting for running threads to complete
yfs: Shutdown complete

ok

solaris/solaris-kld

1..7
ok 1 - Copied kld
ok 2 - Installed kld
ok 3 - Mounted filesystem
amd64 linux26
autorun.inf
backups
Desktop.ini
local
project
public
service
support
test
u
user
usr
www
ok 4 - Listed sample directory
ok 5 - Unmounted filesystem
afsd: Shutting down all afs processes and afs state
ok 6 - Stopped afsd
ok 7 - Unloaded module
ok

All tests successful, 459 tests skipped.
Files=101, Tests=7000, 678.96 seconds (44.73 usr + 36.03 sys = 80.76 CPU)
gmake[1]: Leaving directory `/export/home/buildbot/worker/SunOS11_X64Check/build/test
```

Screenshot

What we've learned

- Adding regression tests helps
- Pay attention to nightly builds
- Retain logs where possible
- Snapshot known-working configurations

Future work

- Triggered builds of Docker or KVM images
 - Currently Docker images always prebuilt, KVM always ad-hoc
- Docker to run services, KVM or local to test kernel extensions
 - Prebuilt binary images for docker to avoid building for 2 systems
- Watchdog systems (via APIs or cloud service)

Hands-off testing for networked filesystems

Thanks for listening!

Questions?

dariaphoebe@auristor.com