

If the main interface to fixity checks is an email that something has failed -- then "all good" is indistinguishable from "checks aren't running at all" or "notifications of failures aren't working."

We wanted a **fixity health dashboard** that can be accessed on demand to show fixity summary across the collection.

All images are screenshots from our live and fully-functional system, with a public front end at <https://digital.sciencehistory.org>

Our system is a Rails application that does not use Hyrax or Fedora, but is informed by concepts and experiences from Samvera community

## Fixity Health Report

👍 Fixity health good.

- We intend to schedule checks such that every file is checked at least once every 7 days.
- Fixity checks run **nightly at 2:30 am** and will check 1/7 of our total assets (~4392 assets) each night, choosing the assets with missing or most stale checks to run checks on.
- Can only run fixity checks on assets that are *fully ingested* and have attached files.
- For assets created within last 1 day, we tolerate missing fixity checks, as they may not have run yet.

### Assets Currently Failing Fixity

👍 None

### Assets created more than 1 day ago

Expected to have a recent fixity check on record

Which assets	How many
Total	30,745
Not yet fully ingested	0
Fully ingested without <i>fresh</i> (within past 7 days) fixity check	0 <span>👍 Good</span>

### Asset and Fixity Decriptive Statistics

Which assets	How many
Total assets	30,745
Created in the last 1 day	0
Not yet fully ingested (no attached file) <i>These should be temporary, pending ingest, and can't be fixity checked yet.</i>	0
Fully ingested (attached file)	30,745
Fully ingested, no checks	0
Fully ingested, at least one check	30,745
Fully ingested, checked in the past 7 days	30,745
Fully ingested, stale checks (have not been checked for 7 days, or ever)	0
Fully-ingested Asset with oldest most-recent fixity check	<a href="#">2020-Sep-25 02:30:04 EDT</a>
Oldest fixity check on record	2019-Sep-25 02:30:06 EDT
Most recent fixity check on record	2020-Oct-01 15:40:44 EDT

The dashboard intends to let a Repository Manager, without needing to get a developer involved, know:

- ▶ A simple **green** or **red** summary of whether fixity health is all good or has problems.
- ▶ That **all** assets in repository have been fixity checked as recently as expected
- ▶ That they have all passed their checks
- ▶ Some overview description of the collection, to give further confidence that all is as expected, or clues to what might not be.

## Fixity Health Report

🗨️ There are fixity health problems!

- 1 asset currently failing fixity check
- 2 assets expected to have fresh fixity checks, but missing fresh fixity checks

- We intend to schedule checks such that every file is checked at least once every 7 days.
- Fixity checks run **nightly at 2:30 am** and will check 1/7 of our total assets (~4392 assets) each night, choosing the assets with missing or most stale checks to run checks on.
- Can only run fixity checks on assets that are *fully ingested* and have attached files.
- For assets created within last 1 day, we tolerate missing fixity checks, as they may not have run yet.

### Assets Currently Failing Fixity

🔴 1 failure

✖ <a href="#">b1063696_629.tif</a>
------------------------------------

### Assets created more than 1 day ago

Expected to have a recent fixity check on record

Which assets	How many
Total	30,746
Not yet fully ingested	1 <span>⚠️ Expect 0. Are these stuck in ingest process?</span>
Fully ingested without <i>fresh</i> (within past 7 days) fixity check	2 <span>🔴 Unexpected, should be 0</span>

### Asset and Fixity Decriptive Statistics

Which assets	How many
Total assets	30,747
Created in the last 1 day	1
Not yet fully ingested (no attached file) <i>These should be temporary, pending ingest, and can't be fixity checked yet.</i>	1
Fully ingested (attached file)	30,746
Fully ingested, no checks	2
Fully ingested, at least one check	30,744
Fully ingested, checked in the past 7 days	30,743
Fully ingested, stale checks (have not been checked for 7 days, or ever)	3
Fully-ingested Asset with oldest most-recent fixity check	<a href="#">2020-Sep-24 04:46:38 EDT</a>
Oldest fixity check on record	2019-Sep-25 02:30:06 EDT
Most recent fixity check on record	2020-Oct-01 04:52:36 EDT

**Fixity audits** Status ✔ OK

**Latest** 4 days ago  
**Earliest** about 1 year ago

[Fixity Check History](#)

- ✔ 2019-Sep-27 04:53:35 EDT SHA512:3a7c534816
- ✔ 2020-Aug-16 03:47:56 EDT SHA512:3a7c534816
- ✔ 2020-Aug-23 03:54:42 EDT SHA512:3a7c534816
- ✔ 2020-Aug-30 03:58:18 EDT SHA512:3a7c534816
- ✔ 2020-Sep-06 03:51:23 EDT SHA512:3a7c534816
- ✔ 2020-Sep-14 03:38:41 EDT SHA512:3a7c534816
- ✔ 2020-Sep-21 03:36:50 EDT SHA512:3a7c534816
- ✔ 2020-Sep-28 03:34:24 EDT SHA512:3a7c534816

[Schedule a check now](#)

For an individual asset, a repository manager can see some fixity check history, and trigger an additional immediate fixity check.

**More information for a currently fixity-failing individual asset.**

Including a link to the file in the S3 Console for more investigation. (Our files all live in S3).

**Fixity audits** Status ✖ FAILING

**Latest** 1 day ago  
**Earliest** about 1 year ago

[Fixity Check History](#)

- ✔ 2019-Sep-25 05:26:18 EDT SHA512:f8fd92e4f0
- ✔ 2020-Aug-11 03:03:26 EDT SHA512:f8fd92e4f0
- ✔ 2020-Aug-18 03:03:54 EDT SHA512:f8fd92e4f0
- ✔ 2020-Aug-25 03:06:12 EDT SHA512:f8fd92e4f0
- ✔ 2020-Sep-01 03:07:09 EDT SHA512:f8fd92e4f0
- ✔ 2020-Sep-09 02:51:46 EDT SHA512:f8fd92e4f0
- ✔ 2020-Sep-16 02:54:49 EDT SHA512:f8fd92e4f0
- ✔ 2020-Sep-23 02:54:27 EDT SHA512:f8fd92e4f0
- ✖ 2020-Sep-30 02:54:13 EDT SHA512:ASC55581CA

**Current Failing FixityCheck id**  
3778515

**Time**  
2020-Sep-30 02:54:13 EDT

**Expected**  
SHA512:

**Actual**  
SHA512:

**Checked file location**  
<https://scihist-digicoll-staging-originals.s3.amazonaws.com/asset/fca9e039-3e8b-4eb2-a9ef-0f930c21472c/clbea5b7dd57fe98d2dfbec99b85e23d.1>

[Schedule a check now](#)

**Some unresolved issues**

- ▶ For failing fixity checks we list/link to relevant assets. But for other problematic/of-concern categories, we only give a count, no way to see the actual assets, which would be nice.
- ▶ The logic for trying to fixity-check 1/7th of assets every night, and checking to make sure all assets are checked weekly -- doesn't add up under ingest of new assets, and at any given time we usually have a warning that a 10-20 assets are overdue for a check.