

[LOCKSS](#) > [Technical Support](#) > [Use A LOCKSS Box](#) > [Daemon Release Notes](#) > Daemon Release Notes 1.3 – 1.36

Daemon Release Notes 1.3 – 1.36

Daemon 1.36.3

- Features
 - LCAP messages for unreachable peers are queued and retried at intervals.
 - Linux/Solaris: if `/etc/lockss/runssl` script exists, run it before starting daemon (to set up SSL keystores and environment).
 - Send crawl end report with list of URLs and hashes, if `org.lockss.crawler.crawlEndReportEmail` is set to email address. Hash algorithm specified by `org.lockss.crawler.crawlEndReportHashAlg`, defaults to SHA-1. This is for PLNs staging content to ingest server, to detect when content may be removed.
- Bug fixes
 - If daemon can't start because it can't load config, it once again displays a message to that effect in the UI. [The actual change is that it sets up the usernames and passwords so it's possible to see the message.
 - Linux/Solaris scripts set up tmp dir on cache partition, so there's enough space.
 - Linux/Solaris hostconfig does more checking of required user inputs.
 - Linux/Solaris: Changed default proxy port from 9090 to 8080.
 - Link extractor parses files using the charset specified in the HTTP headers.

Daemon 1.35.4

- Bug fixes
 - Network errors could cause temporary file to accumulate on disk.

Daemon 1.35.3

This is a Linux-only release, with changes to the Linux configuration and startup scripts. The daemon itself is identical to 1.35.2.

- Features
 - NAT is now supported. hostconfig prompts for external IP address of NAT router. See [NAT](#) for port forwarding instructions.

- Bug fixes
 - Avoid asking JVM for too much heap on large memory machines.
 - hostconfig no longer erroneously reports that cache dirs aren't writable.

Daemon 1.35.2

- Features
 - More improvements in poll scheduling to create willing repairers more quickly.
 - Improve odds of achieving quorum in polls on sparsely held AUs by keeping track of which peers hold which AUs.
 - Differential peer addressing allows LOCKSS networks to comprise an arbitrary mix of peers in public and private (NATted) networks.
 - Global config files (e.g., lockss.xml) may include additional files (by URL).
- Bug fixes
 - Fixed several bugs in URL rewriting. ServeContent is usable for more sites.
 - Files larger than 2GB can now be repaired from a peer. This necessitated a backwards-incompatible change to the LCAP message header; daemon 1.35 will not poll with previous daemon versions.
 - Fixed "direct first" variant of generated PAC files.
 - Work around a problem in some environments where mkdirs fails intermittently.

Daemon 1.34.3

- Features
 - Sites that require Basic or Digest authentication can now be crawled. Username and password are configured as AU params, either in title DB or daemon's AU config.
 - AUs with content spread across an unpredictable set of hosts can now be crawled. Plugins may supply a path which will be used as the absolute path of a permission page on any host for which no permission page is otherwise specified.
 - Several enhancements designed to concentrate polling resources where they're needed most, *eg*, AUs with few willing repairers or low agreement.
 - Inquorate polls are tallied anyway if doing so might establish agreement (repair-ability) where none existed.
- Bug fixes
 - Sorting AU Status column by percent agreement works correctly.
 - Don't count "down" AUs as needing to be crawled.
- Parameter changes
 - `org.lockss.baseau.toplevel.poll.interval.min` and

`org.lockss.baseau.toplevel.poll.interval.max` have been replaced by `org.lockss.poll.v3.toplevelPollInterval`. AUs with low agreement can be polled more frequently by using `org.lockss.poll.v3.pollIntervalAgreementCurve` and `org.lockss.poll.v3.pollIntervalAgreementLastResult`. See <http://www.lockss.org/lockssdoc/gamma/daemon/paramdoc.html>

Daemon 1.33.5

- Features
 - Call polls more frequently on AUs with low agreement.
- Bug fixes
 - Backup mail was not being sent.

Daemon 1.33.4

- Features
 - Runaway crawls caused by URLs with ever-lengthening substring repetition are now detected and prevented.
 - Peer invitation and nomination probability better account for unresponsive peers.
- Bug fixes
 - CSS parsing errors issue warnings, don't abort the crawl.
 - Long-running crawls no longer disappear from the Crawl Status table.
- UI changes:
 - The Repair Candidates table now displays blanks (instead of 0.0) for unknown agreement percentages.

Daemon 1.33.3

- Features
 - Read-only UI logins are supported for PLNs. They will be available to all sites in an upcoming release.
 - Plugins can specify multiple start URLs.
 - Set- or range-valued AU config params used in start URL or permission page specs generate multiple URLs by enumerating the sets or (integer) ranges.
 - Set-valued AU config params used in a regexp context (crawl rules, login-URL pattern) match more precisely.
 - The v3 poll state directory can be moved off the first disk by setting `org.lockss.v3.statePath`.
- Bug fixes

- Links are followed from more variants of `<meta http-equiv="refresh" ...>` tags.
- UI changes:
 - The table for peers holding an AU, in addition to the columns showing the highest and last agreement in a poll this peer called, now has highest and last agreement hint columns showing the agreement in the last poll in which this peer voted. Similarly, the display for a completed vote now includes an agreement field. These fields will only become credible gradually as daemons are updated and new votes are completed. Until then they will display 0.0.

Daemon 1.32.4

- Features
 - Direct content serving (see note below).
- Bug fixes
 - Crawler is more tenacious about retrying after socket timeouts, and displays a better message.
 - Disk space tables properly display filesystems larger than 4TB.
 - Optimized VoteBlock lookup when checking repaired files. (Was causing excessively long single-URL hashes).
- UI changes:
 - Rename “Proxy Options” page to “Content Access Options”, rename “Proxy Server Options” page to “Content Server Options”, add facility to enable the content server.

Direct Content Serving: As promised, this release includes the first part of LOCKSS support for access to content via URL resolvers such as SFX, namely the ability for your LOCKSS box to *serve* the content it is preserving as well as *proxying* it. To do this, the daemon now supports URL rewriting. Some of your box’s AUs cannot yet be served, because we have not yet updated all plugins with support for URL rewriting. Among the AUs awaiting URL rewriting support are those for HighWire journals; these will appear in future daemon releases. These releases will also be able to auto-generate the metadata needed to describe your LOCKSS box’s holdings to SFX.

To enable serving content, use the administrative UI, choose “Content Access Options” then “Content Server Options” and enable the content server on one of the available ports, for example 8082. If you need to disable another service to free up a port on which to run the content server, disable it on this page, then visit the administrative UI home page, wait a couple minutes, then follow the instructions above.

Then visit the page: <http://your.lockss.box:8082/> – you will see a list of those of your box’s AUs that

support content serving, each with a link to the AU's start page. Follow that link to see the AU's content.

At this early stage there may be some glitches. A link that points elsewhere in the journal site but not to a collected URL will get a 404 error. A link that should have been rewritten but wasn't will get you to the publisher's site – this may well happen with links generated by Javascript, which are difficult to rewrite. Please help us improve content serving by reporting any pages that don't look right to

lockss-support@lockss.org

Publishers for which content serving is currently enabled include Annual Reviews, AnthroSource, BioMed Central, Project Muse, and the open access publications 19th Century Art Worldwide, Absinthe Literary Review, Applied Semiotics / Sémiotique Appliquée, Arkivoc, Blackbird, CLCWeb, Electronic Journal of Contemporary Japanese Studies, Evergreen Review, Exquisite Corpse, Genders, Invisible Culture and Journal of Buddhist Ethics.

Daemon 1.31.2

- Features
 - Crawler supports SSL (https:).
 - The ARC exploder has been re-implemented to use the new infrastructure introduced for the ZIP and TAR exploders. It is now capable of ingesting ARC files that contain crawls covering multiple hosts, and organizing them into an ArchivalUnit per host.
- Bug fixes
 - Measure actual hash duration when necessary to recover from excessive hash estimates (fixes persistent “No Time” vote errors).

Daemon 1.30.3

This is a Linux and Solaris build which is feature-identical to 1.30.2.

Daemon 1.30.2

- Features
 - Crawler now retries (after delay) after socket timeout, socket reset, etc. Plugin can customize retry count and delay.
 - Serve repairs of open access AUs without prior agreement (unless `org.lockss.poll.v3.openAccessRepairNeedsAgreement = true`).
 - Display AU creation date.
 - Sort numeric and date columns in status displays descending by default (when clicked).
 - `org.lockss.plugin.acceptExpiredCertificates` can be set false to reject plugin certificates that are either expired or not yet valid.

- Bug fixes
 - Don't vote on uncrawled AU, even if marked pub-down.
 - Fixed scheduler bug that allowed tasks to be scheduled when there was no time available – caused many hash timeouts.
 - Fixed race condition in comm code that caused some peers to become unreachable until restart.
 - Various poller bug fixes and optimizations.
 - Expired plugin certificates were accepted regardless of validity.

Daemon 1.29.4

- Bug fixes
 - Fixed bad HTML that caused display problems in IE.
 - Additional invitation mechanism was stopping too early, not inviting enough participants into polls.

Daemon 1.29.2

- Features
 - Polls are called according to AU's priority, not randomly as before. Initial priority implementation is based on time since AU last successfully polled.
 - Crawl rules that reference unassigned optional AU config params are now explicitly ignored.
 - Empty query strings in URLs are normalized away. Existing files with empty query strings are renamed.
 - Recognize Open Journal System's LOCKSS permission statement.
 - Scheduler now compensates for background load.
 - Added Reload Config button to DebugPanel.
 - Status tables sort numeric columns in descending order by default.
 - Poll tables show more info about participant state.
 - AU tables display repository dir, more poll state.
 - "Disabled" status of crawler, poller, voter are displayed in status tables and overview.
 - RPM install scripts support multiple groups, other improvements.
- Bug fixes
 - Resolved poll threading issues that caused bottlenecks, deadlocks and frequent watchdog exits.
 - Fixed poll timing bugs that caused polls to have unreasonable durations.
 - Fixed scheduler bug that prevented adding new tasks in some circumstances.
 - Peers in different groups aren't invited into polls or displayed in Peer IDs table.
 - Polls don't delete files from exploded AUs.
 - Prevent AUs from polling too frequently.

- Reenabled polling on plugin registry AUs.
- Fixed race creating state, data dirs.
- Fixed bug evaluating nested conditionals in XML config files.

Daemon 1.28.2

- Features
 - Crawl rules are now case-independent by default, can be overridden by plugin.
 - Message is now clearer when Admin UI pages or proxy access are forbidden because of access control.
 - RPM now includes info for chkconfig to create links in rc.N dirs so daemon is started when system boots. (chkconfig must still be run manually after install.)
 - Added more detailed communications status information to UI.
 - Minor improvements to crawl scheduling and status reporting.
- Bug Fixes
 - Fixed bug that blocked communication with some peers until restart. More polls should now achieve quorum.
 - Proxy “not found” error page didn’t include index of likely related AU in all cases.
 - Fixed occasional infinite loop updating title db, leading to daemon restart.

Daemon 1.27.4

- Features
 - New framework for collecting archive files and exploding content into AUs. Implementations for zip, arc and tar.
 - Crawler improvements: make more efficient use of available crawl time, spread load more evenly across sites, more sensibly prioritize AUs to crawl.
 - Poller improvements: better prioritize AUs to poll, simplify and rationalize poll duration calculation, more lenient (and more correct) repair policy.
 - Open access journals can identify themselves as such with new permission statement. (“LOCKSS system has permission to collect, preserve, and serve this open access Archival Unit”).
 - Plugins can specify operator-action-needed message when AU configured.
 - UI responses are gzipped if the client request includes Accept-Encoding: gzip.
 - Status tables can be output as csv, allows import into spreadsheets. Specify output=csv as query arg.
 - Various status info improvements.
 - Eliminated obsolete treewalk mechanism.
- Bug fixes

- Voters now correctly report and display “No Time Available” condition.
- Voters that run out of time before their hashing is done display the correct “Expired w/o Voting” status.
- Avoid multiple fetches of permission pages where possible.
- Hash errors no longer leave streams open or prevent other versions of file from being voted on.
- Reduced hasher memory requirements.

Daemon 1.26.8

- Features
 - New daemon status overview page.
 - Plugins can specify fetch rates for specific MIME-types.
 - Display publishing platform for (possibly) multi-hosted titles.
 - Supply version and identity info to props server on first request.
 - Warn on low disk space, always display disk space on AU config pages.
 - Improvements to polling.
- Bug fixes
 - Ensure http connections closed promptly.
 - Startup race could result in no polls, no defined titlesets.
 - Redirect outside crawl spec now properly recorded as excluded URL.
 - Proxy now transparent to cookies.
 - Login page checker invoked on all fetches.
 - Ensure correct ClassLoader for all plugin auxilliary classes.
 - Polling should achieve Quorum more often.

Daemon 1.25.3

- Features
 - Content can now be harvested from CONTENTdm and other Dublin-Core metadata sources.
 - Plugins may now specify URL regular expression to detect redirects to login pages.
 - Proxy server to be used for crawling is now user settable.
 - Unadorned list of all URLs in AU now available from AU status page.
 - Single cache can now be in multiple PLN groups.
 - PLN group name(s) now displayed in UI.
 - New Platform Config status shows IP addr, SMTP server, group name(s), disks, etc.
 - Initial Solaris support.
 - Daemon may be configured with multiple Title Database URLs.
 - Peer Identities that are not part of the Daemon’s PLN will be pruned from the Identity Database over time.
- Bug fixes

- List of available ports in ICP config now correct.
- Lists of URLs excluded from crawl now exclude most off-site URLs, to avoid lists getting too large.
- Unexpected errors during crawl (e.g., illegal CSS syntax) now have message in Errors list.
- Content from other than port 80 now stored with explicit port number.

Daemon 1.24.2

- Features
 - Add Titles dialogue by default places new AUs on disk with most free spece.
 - Plugins may specify reminder to be displayed to user when AUs configured.
 - Overly permissive entries in access lists require confirmation; ridiculously permissive entries are disallowed.
 - PLN operators may specify their own plugin registries, title databases and keystores.
 - CLOCKSS no longer maintains "subscribed" status of AUs.
- Bug fixes
 - Polling now correctly prunes peers that have declined to vote.
 - Plugin validation code is now completely Java 1.5 compliant.

Daemon 1.23.1

- Features
 - Support alternate charset encodings in HTML and other content.
 - Performance improvements in proxy and ICP.
 - URL(s) of manifest page(s) highlighted in AU URL list.
 - Proxy blocks CONNECT method. Not needed with LOCKSS and can be misused if proxy access control is too lax.
- Bug fixes
 - Fixed inconsistency in temporary redirect handling that could cause proxy to not serve file.
- CLOCKSS changes
 - Proxy and ICP search for subscribed AU if more than one has URL.

Daemon 1.22.6

- Features
 - Seamless proxy. If user requests URL not in cache (e.g., journal home or search, etc.) and site is gone or not providing content, proxy will offer index of cached content from that site.
 - Crawler now follows links from CSS (in CSS files and <style> fragments in HTML).
 - PAC files instruct browsers to failover to direct connection if LOCKSS proxy doesn't respond.
 - V3 polls invite additional peers as needed, resulting in fewer no-quorum polls.
 - Upgraded to PDFBox 0.7.3.
- Bug fixes

- Proxy generates error page if unexpected processing error.
- Now possible to re-add AUs on 2nd disk.
- Netmask 0/0 now interpreted correctly.
- In rare circumstances inconsistent URL-encoding on a site caused files to appear to be missing.
- AU and repository status tables come up much more promptly. Any sizes that aren't known aren't displayed, and are scheduled for background (re)calculation.
- Fixed the way poll duration was calculated.

Daemon 1.21.3

- Features
 - Upgraded to Commons Lang 2.2.
- Bug fixes
 - No longer parse for URLs in the body of <script> tags.
 - Parsing <option> tags for URLs is off by default.
 - Crawls blocked by crawl windows have better status message in UI.
- Memory/Performance improvements
 - The daemon's memory requirements have been greatly reduced, and maximum memory usage is now scaled to physical memory size. This will reduce performance problems for machines with limited memory (<512 MB) that are preserving a large number of AUs.
 - The handling of voteblocks temporarily stored on disk has been improved, which should prevent the /dist file system filling up.
- CLOCKSS changes
 - The remaining part of the SSL support, which allows the use of a permanent keystore protected by a random password. These changes should allow CLOCKSS and LOCKSS boxes to use exactly the same daemon package in future releases.

Daemon 1.20.2

Released 10/18/06

- Features
 - New framework based on PDFBox to filter dynamic content in PDF files.
 - Upgraded to Getopt 1.0.13.
 - V3 improvements.
- Bug fixes
 - Several fixes and improvements for V3 Poll scheduling.
 - Poll status is now preserved between daemon restarts.
 - Various fixes to crawl status display.
 - Reduced maximum heap size to avoid thrashing on small memory machines.

Daemon 1.19.6**Released 10/5/06**

- Bug fixes
 - Incorrect base calculation in unusual redirect situation could lead to spurious 404s in crawl.
 - Prevent V3 poller's temporary reservation task from wasting CPU time.
 - Simultaneous crawls sharing a rate limiter now get scheduled fairly.

Daemon 1.19.2**Released 9/6/06**

- Features
 - Framework based on HTMLParser to filter dynamic content in HTML files. Allows more sophisticated filters than existing text pattern matchers.
 - Initial support for protecting V3 poll communication using SSL. This should be completed in the next platform release with support for the necessary keys and certificates. In the meantime it is being tested in the CLOCKSS environment.
 - CLOCKSS: Crawler requires CLOCKSS permission if running as CLOCKSS.
 - CLOCKSS: Crawls from one of two IP addresses used to collect non-subscribed material, detect and record lack of subscription.
- Bug fixes
 - Avoid queueing crawls that can't start due to crawl window.
 - Poll history files were getting too large, taking much too much time to load and store. Now trimmed to configurable max size.

Daemon 1.18.3**Released 8/7/06**

- Features
 - Resource-specific rate limiters allow more accurate control over load placed on publishers' sites by crawler.
 - Enhancements to V3 to make it ready for rolling out to the LOCKSS community.
 - Various status table improvements.
 - More accurately report V3 polls in the AU status table.
 - Upgraded to Commons Logging 1.1, Commons Collections 3.2 and Log4j 1.2.13.
 - Increased runtime heap size from 250MB to 500MB.
- Bug fixes
 - Fixed bugs related to fetching repairs with V3.
 - Fixed bugs related to restoring V3 polls after a daemon restart.

- Fixed bugs and inconsistencies in the handling of string ranges and numeric ranges in plugins.
- The resource manager now properly manages ports that were specified by the platform as range strings.
- Manual AU config doesn't offer "down" titles.

Daemon 1.17.4

Released 6/29/06

- Features
 - A list of (and links to) the manifest pages for all AUs is available at <http://<cache-name>:8080/>
 - Partial support for accessing previous versions of content.
 - Configured ISO image (for config-on-CD) is available for download via UI.
 - Plugin registry crawls have higher priority than other crawls.
 - Improved the serialization facilities (fine-grained error reporting, new individual error recovery modes, better test coverage).
 - Normalized the format of time zones in serialized form.
 - Pooled most config param objects between plugins.
- Bug fixes
 - Crawl status Bytes Fetched column now works for repair crawls.

Daemon 1.16.10

Released 6/7/06

- Bug fixes
 - Fix for null pointer exception in the ICP manager.
 - Improved performance for new content crawls.
 - Improvements in V3 polling.

Daemon 1.16.7/1.16.8

Released 5/3/06

- Features
 - Improved memory profile of V3 polling by keeping some state on disk.
 - Better view of V3 polling status.
 - Enabled repair of V3 content from publishers.
 - PAC files and EZProxy/Squid configuration fragments now featuring detailed comments for each domain name.
 - Increased heap size allows more AUs to be preserved.

- Adding a large number of AUs through the UI is much faster and uses less memory.
- Daemon starts faster.
- New content crawls are managed by a thread pool and queue:
 - Crawls start much more promptly when needed,
 - Number of simultaneous crawls can be controlled.
- Crawl status displays number of content bytes fetched.
- Configuration properties are fetched from property server in compressed form.
- If daemon can't start because the cache's IP address isn't in our access list, a more helpful message is displayed.
- Config backup email suppressed if no AUs have been configured.
- Bug fixes
 - Prevent pollers from inviting themselves into polls.
 - Improvements to the serialization code. (More improvements underway)
 - Interpretation of malformed URLs that contain possible path-traversal attacks changed to agree with popular browsers. (Extra ".." components are now removed, rather than causing entire URL to be ignored.)
 - Hung connection to property server no longer causes daemon to hang.
 - Fixed bug that prevented repairs from being fetched from any but the first of a list of candidate caches.
 - Daemon RPM installation no longer complains about unknown user.

Daemon 1.15.7/1.15.8

Released 3/23/06

- Bug fixes
 - Increased OpenBSD ulimits for stability on 3.8 and -current.
 - Fix for parsing HTML with null or empty base tags.

Daemon 1.15.5/1.15.6

Released 3/22/06

- Features
 - Daemon now reports the number of bytes of content fetched during a crawl.
 - Java heap size increased to give better daemon stability.
 - Upgraded to XStream 1.1.3.
 - Improved documentation for proxy integration.
 - V3 Polling now ready for wider testing.
- Bug fixes
 - URL path normalization fixed to comply with browser behavior.

Daemon 1.14.6/1.14.7**Released 2/27/06**

- Bug fixes
 - Fixed a bug that prevented available titles from being selected.

Daemon 1.14.4/1.14.5**Released 2/16/06**

- Bug fixes
 - Increased (and parameterized) max uploaded file size.

Daemon 1.14.2/1.14.3**Released 2/9/06**

- Features
 - Platform access subnet can be list of subnets.
 - Ability to receive and apply repairs using the V3 Protocol.
 - New content releases.
- Bug fixes
 - Syntax of V3 identity changed to TCP: [ip]:port to remove ambiguity.
 - URL processing fixed to remove newlines and leading whitespace, and not normalize query part.
 - Numerous V3 stability fixes.

Daemon 1.13.5**Released 12/14/05**

- Bug fixes
 - Fixed a bug which prevented caches from sending unicast packets

Daemon 1.13.3**Released 12/7/05**

- Features
 - Added detailed poll status tables for V3.
- Code cleanup
 - Many improvements to V3 polling stability.

- Permission checker is now separate from Crawler

Daemon 1.12.3

Released 11/4/05

- Features
 - Ability to call and participate in V3 protocol polls, suitable for internal testing.
 - Arbitrary title config props allow greater flexibility in defining title sets.
 - Repair crawls can now be set to require permission before crawling the publisher site
 - When collecting content, we can now generate a single "Cookie" header (combining all the cookies)
 - Config files now support nested <if>
 - The LOCKSS cache now implements ICP, a protocol used by caches and proxies to communicate and to locate content. See [pubwiki:IntegratingWithSquid](#) for details.
- Bug fixes
 - OAI harvests now properly update the AU state, allowing polls to run.
 - All files now closed when hashes abort.
- Code cleanup
 - Updated HTTPClient jar to its current version
 - Updated xml processing and jakarta jars to its current version
 - Updated OCLC Harvester2 jar to its current version

Daemon 1.11.5

Released 9/27/05

- Features
 - Configuration backup file (used for repair if disk fails) can be emailed to admin monthly.
 - ICP support (beta, disabled)
 - New object persistence framework using XStream instead of Castor. Migration of existing files will begin with the next release.
 - New hashing framework for V3 polling.
 - Added persistence for V3 polls.
- Bug fixes
 - Disk and repository errors during crawl correctly reported in UI
 - Login page checkers now reset the Reader for non-login pages
- Code cleanup
 - updated various 3rd party jars to current version

Daemon 1.10.6

Released 8/22/05

- Features
 - AUs can be marked “publisher down” when content no longer available from original source. Suppresses crawls, alters repair behavior.
 - Backup file now contains all info necessary to recover from total loss (e.g., disk failure) when publisher is down.
 - Crawler now aware of damaged nodes list: force unconditional fetch (suppress if-modified-since), clear damage after fetch from publisher.
 - Crawler status/history more accurate, includes more details.
 - Hostname map makes proxy check for cached URL more efficient.
 - Daemon version numbers are interpreted more intuitively (base 10).
 - Simplified V3 state machine
 - Began implementation of V3 poll controller
- Bug fixes
 - Internal AUs (e.g., plugin registries) can't be deleted or deactivated.
 - Crawler couldn't get permission for publisher-down repair, now doesn't try.
 - V1 tree-based repair couldn't fetch slashless dir-node if out of crawl spec.
 - OAI handler was holding onto all metadata records, causing out of memory errors.
 - Failure to fetch prop file on startup results in explanatory UI page, not auth failure.
 - Fixed file descriptor leak in proxy, sometimes caused file open failures until next GC.
 - Transient open errors caused repository to deactivate node.
 - Multicast quench on non-multicast-connected machines caused machine to not record its own votes, increasing chance of no-quorum.
- Testing
 - Added functional tests for total loss recovery, repair from cache.
 - V3 testing framework begun

Daemon 1.9.4**Released 7/26/05**

- Fixed situations where file descriptors could be leaked

Daemon 1.9.2**Released 6/20/05**

- Crawler
 - Supports Creative Commons licenses
 - Added checking for login pages during crawling

- Added “probe” permission checking
- String searches are now done w/ BoyerMoore
- UI
 - Plug-in Detail table added

Daemon 1.8.2

Released 5/10/05

- Crawler
 - Collects cookies and passes them along
 - Better handling of javascript URLs
- Plug-ins
 - pulled out into their own hierarchy
- UI
 - Available ports listed in UI adult proxy config page
- General bug fixes
- Internal restructuring of message classes, V3 additions

Daemon 1.7.4

Released 3/23/05

- UI
 - access control pages now usable from lynx on console (workaround for lynx textarea bug)
 - unselectable title sets are now greyed out rather than omitted
 - bug fixes for odd disk usage output
 - HashCUS servlet protected against divide by zero
- Crawler
 - handles temporary redirects (302, 303, 307) by only writing content at original URL
 - html parser extracts links in <embed> and <applet> tags
- Communication layer
 - improved packet rate limiting
 - added measures to reduce likelihood of “islands” forming
 - better handling of bad packets
- Polling code
 - use only mime-type portion of content-type to determine parser
 - properly close all filter readers
 - bug fix in [White Space Filter](#)
 - new Protocol State Machine (infrastructure work for V3 polling)
- Configuration

- better handling of malformed config files
- General
 - another findbugs audit and clean up
 - more protection against null pointer exceptions in the code

Daemon 1.6.5

Released 2/24/05

- Fixed a stream closing bug

Daemon 1.6.4

Released 2/7/05

- Polling code
 - General bug fixes
- Repository
 - Fixed not to hold so many open files
- Plug-in framework
 - Title information can now be bundled with plug-ins
 - Plug-ins now have size estimates for their AUs
- UI
 - Crawl status now lists number of errors the crawl had
 - Crawl status tracks and lists URL that were fetched, parsed, not fetched or had errors
 - AU status table now lists disk usage
 - New UI for adding titles
- Logging
 - Each log message now has thread id
- Crawler
 - Keeps much more info about each crawl
 - Temporary redirects are now not written to each redirected URL, but only to the original

Daemon 1.5.7

Released 12/6/04

- General
 - FindBugs cleanups
 - Memory optimizations to allow more AUs
 - Initial format migration framework
- Daemon

- Generalized PeerIdentity replaces LcapIdentity, for V3 polling support.
- Restructured daemon startup options
- Configuration
 - Moved configuration code into org.lockss.config
 - Optimized to avoid unnecessary rereading of config files
 - Bug fix to write local config files atomically
- Crawler / Parser
 - Beginnings of OAI Harvesting
 - Added per-AU crawl-start rate limiters
 - Crawls can be disabled by a parameter
 - Extended crawl rules to allow for issue sets and ranges
 - Support for Creative Commons licenses
 - Crawler can use a proxy
 - Unambiguously identify repair requests.
 - Bug fix for parsing pages multiple times
 - Properly decode HTML entities in links
- Plug-in framework
 - CachedUrlSet, CachedUrl and UrlCacher now point directly at owning ArchivalUnit instead of (happenstance) CachedUrlSet.
 - Loadable plug-ins revetted only when they change
- Polling code
 - Numerous bug fixes
 - Infrastructure work for V3 polling code
 - Poll lock restructuring
- Communication layer
 - can disable datagram comm layer
- Proxy
 - Enforce new repair-fetch identification (configurable)
- Library
 - Upgraded to commons-collections-3.1 jar
 - Added htmlparser.jar
- Testing
 - Beginning of automated multidaemon testing (STF)

Daemon 1.4.8

Released 10/26/04

- Crawling can now be disabled by a property value
- The number of crawls on an AU is now rate limited
- Fixed a bug where repairs could result in a loop (if the node had many children or children with long

names).

Daemon 1.4.6

Released 10/7/04

- further support for downloadable plug-ins
- URL canonicalization (for IOP and HighWire)
- define AU by issue ranges
- refactoring of plug-in code
- LockssApp pulled out of LockssDaemon
- Crawler will not follow cross-host redirects without checking for permission statement on new host
- Crawl permission check moved out of plug-in framework, into crawler
- Improvements to handling of polls in which all caches have damage
- removed need for org.lockss.plugin.xml Plugins configuration setting
- list of agreeing caches exported to UI

Daemon 1.3.6

Released 8/11/2004

- addition of the audit proxy port
- plug-ins are now downloadable separately
- new xml prop files are supported
- UI comes up with warning until all AUs are started
- tiny UI comes up when daemon is unable to reach prop server
- initial cut of alerts framework