



LOCKSS and 360 Link Integration Guide

Philip Gust

Stanford University LOCKSS Program

26 July 2012

Introduction

LOCKSS (*Lots of Copies Keep Stuff Safe*) and Serials Solutions have made content preserved in LOCKSS boxes available through on-line public access catalogs (OPACs) that use the 360 Link link resolver. This guide describes the steps required to customize your LOCKSS box and 360 Link instance.

These steps include:

1. Enabling the Content Server on your LOCKSS box
2. Configuring access conditions to the Content Server
3. Configuring 360 Link with information about the LOCKSS content server
4. Activating access to preserved titles in 360 Link

Customizing Your LOCKSS Box

A LOCKSS box has an optional Content Server that makes preserved content available from a web browser through OpenURL queries. The 360 Link provider uses OpenURL queries to request content from your local LOCKSS box. The Content Server is not enabled by default. To make content in your LOCKSS box available through 360 Link, you must use its configuration screens to enable and configure this feature.

Enabling the Content Server

To enable the Content Server, select “Content Access Options” from the main administration screen of your LOCKSS box, then select “Content Server Options” from the “Content Access Options” screen. This takes you to a screen for managing the LOCKSS box's content servers and proxies.

Select “Enable content server on port” and enter the port number to use for serving preserved content from a web browser. Port numbers below 1024 require the application to run as “root” and are therefore not available. It is recommended that you choose a port that is close to 8081, the default port used by the administration GUI of your LOCKSS box. Port 8082 is used in this example. Finally, select “Update Content Servers” to finalize this change. The operating system on your LOCKSS box host may require opening “non-standard” ports to enable outside access. Consult with the person who administers the host where your LOCKSS box is running to determine which “standard” ports are available.

To test your Content Server configuration, use your browser to contact the Content Server of your LOCKSS box through the Content Server port. By default, the Content Server displays a list of Archival Units (AUs) that are configured. In this example, if your

Content Access Options

Manage this box's content servers and proxies.

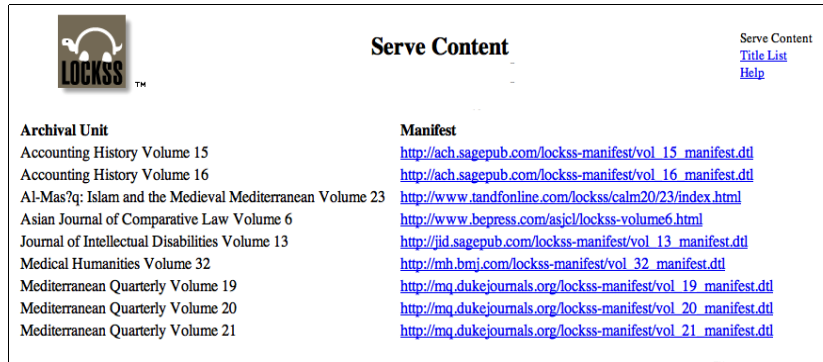
Enable content server² on port

Enable content proxy³ on port

Enable audit proxy⁴ on port

Enable ICP server⁵ on port

LOCKSS box administration GUI is accessed at <http://lockss.xyz.edu:8081/>, your Content Server is at <http://lockss.xyz.edu:8082/ServeContent>.



Archival Unit	Manifest
Accounting History Volume 15	http://ach.sagepub.com/lockss-manifest/vol_15_manifest.dtl
Accounting History Volume 16	http://ach.sagepub.com/lockss-manifest/vol_16_manifest.dtl
Al-Mas'q: Islam and the Medieval Mediterranean Volume 23	http://www.tandfonline.com/lockss/calm20/23/index.html
Asian Journal of Comparative Law Volume 6	http://www.bepress.com/asjcl/lockss-volume6.html
Journal of Intellectual Disabilities Volume 13	http://jid.sagepub.com/lockss-manifest/vol_13_manifest.dtl
Medical Humanities Volume 32	http://mh.bmj.com/lockss-manifest/vol_32_manifest.dtl
Mediterranean Quarterly Volume 19	http://mq.dukejournals.org/lockss-manifest/vol_19_manifest.dtl
Mediterranean Quarterly Volume 20	http://mq.dukejournals.org/lockss-manifest/vol_20_manifest.dtl
Mediterranean Quarterly Volume 21	http://mq.dukejournals.org/lockss-manifest/vol_21_manifest.dtl

You must ensure that the packet filters (firewall) configured on your LOCKSS box allow access to the port you chose for your Content Server. The default packet filter is 8082. Enabling the port for outside access requires administrative privileges, and depends on the operating system running your LOCKSS box. Ask the systems administrator of the host where your LOCKSS box is running to do this.

Configuring Content Server Access Control

Now that you have enabled the Content Server, you'll need to enable access to preserved content by your user community. To do this, return to the main administration page of your LOCKSS box and select "Content Access Control". You will see two lists side-by-side. The "Allow Access" list specifies which IP addresses or ranges can access preserved content through the Content Server. The "Deny Access" list allows you to deny access to specific IP addresses or ranges, typically ones within a wider range specified by the "Allow Access" list.

By default, the "Allow Access" list includes two address ranges: the initial access subnet that was specified when your LOCKSS box was installed, and an address range used by the LOCKSS organization to provide technical support.



Enter the list of IP addresses that should be allowed to use this LOCKSS box as a proxy server, and access the content preserved on it. To be allowed access, an IP address must match some entry on the allow list, and not match any entry on the deny list.

Allow Access ¹	Deny Access ¹
10.67.132.0/24 171.66.236.0/26	

Consult your contracts with the publishers whose content is preserved on your LOCKSS box to determine who is permitted to access it. You should configure the “Allow Access” list to include the same IP address ranges that you supplied to publishers. Note that the simple configuration that is supported by your LOCKSS box assumes that all publishers support the same IP address ranges.

For example, XYZZY University provides the following IP address ranges to its e-pub vendors:

IP range	CIDR Prefix	Netmask	Use
10.12.0.0	/16	255.255.0.0	Dorms
10.64.0.0	/16	255.255.0.0	Main campus
10.65.0.0	/16	255.255.0.0	Medical Center, Hospital

Your institution's contracts with publishers may exclude certain campus functions. Only those address ranges that are allowed should be included in the “Allow Access” list. Here is what XYZZY University would add to its list:

10.12.0.0/16
10.64.0.0/16
10.65.0.0/16

Once you are done configuring the access control for your LOCKSS box, select “Update” to finalize your configuration. Test your configuration by accessing the Content Server from various allowed and denied hosts within your institution.

Using Apache as a Front End to your LOCKSS Content Server

The configuration just described gives end users direct access to your LOCKSS box's Content Server through the port you specified. Some IT organizations limit the use of “non-standard” ports for web services because it requires opening additional ports in their institution's firewalls. These organizations may also prefer to administer access using the same web server access controls that they use throughout their institution. If this does not apply to your institution, you can skip this section.

One way to accomplish this is to direct users to a web server such as Apache that runs on the standard port 80, and configure the Apache server to act as a proxy. It relays requests to the LOCKSS box Content Server and returns responses from the LOCKSS box to the user. This allows the presence of the LOCKSS box to be hidden, and enables the IT staff to use existing access control configurations on the Apache server to limit access to preserved content, including user authentication using HTTPS and SSL certificates. Load balancing could also be done across several LOCKSS boxes in a similar way.

The Apache instance can be one that is shared by many applications within your institution, or it can be a custom instance that is installed on the same host as the LOCKSS box. The simplest configuration uses the standard 'mod_proxy' Apache module to forward content server traffic to and from the LOCKSS box. Here is a typical 'mod_proxy' configuration for an Apache server virtual host at port 80 on the same machine as the LOCKSS box:

```
<VirtualHost *:80>
ServerAdmin lockss-admin@xyzzzy.edu
DocumentRoot /var/www/html
ServerName lockss.xyzzzy.edu

<IfModule mod_proxy.c>
ProxyRequests Off
ProxyVia On
<Proxy lockss.xyzzzy.edu/*>
```

```

AddDefaultCharset off
Order deny,allow
Allow from all
</Proxy>
ProxyPassMatch ^/((ServeContent|images).*)$ http://localhost:8082/$1
</IfModule>

ErrorLog logs/error_log
CustomLog logs/access_log common
</VirtualHost>

```

If you use this technique, you should omit adding IP addresses to the Content Access Control of your LOCKSS boxes because it will only be accessed locally by the Apache server. If the Apache server is on a different subnet, you should add only the IP address of that host to the Content Access Control of your LOCKSS box. You should also add the following special configuration parameter to Expert Config screen of your LOCKSS box.

```
org.lockss.serveContent.absoluteLinks=false
```

This causes Content Server to rewrite links in pages it serves relative to the LOCKSS box address, which simplifies the 'mod_proxy' configuration.

Customizing your 360 Link Database

Once you have configured your LOCKSS box, the next step is to make its contents available through 360 Link. In your 360 Link Client Center, you will need to configure the LOCKSS provider for the titles in your database that correspond to titles that are available through your LOCKSS box.

You will work with your Serials Solutions account manager to configure your 360 Link account. First, ask your account manager to add the LOCKSS provider to your account. The LOCKSS provider enables your users to select LOCKSS to serve content to your end users for titles that are preserved on your LOCKSS box. Serials Solutions has a record of all the titles that your institution subscribes to. It also maintains an up-to-date database of the titles and ranges that are available for preservation by LOCKSS. Your Serials Solutions account manager will enable the LOCKSS provider for those titles you hold that are also available for preservation in your local LOCKSS box.

Next, ask your Serials Solutions account manager to customize the LOCKSS provider on your account for your local LOCKSS box. Your LOCKSS provider must be customized by modifying the provider URL to address your LOCKSS box Content Server. If you are giving users direct access to your LOCKSS box, your Serials Solutions account manager will create a custom URL by replacing “LOCKSS_RESOLVER” in your LOCKSS provider with the address and port number of your LOCKSS box Content Server.

For example, if your LOCKSS box is at `lockss.xyzy.edu` and its Content Server is configured for port 8082, your account manager will replace “LOCKSS_RESOLVER” with `lockss.xyzy.edu:8082`. If you are giving indirect access to your LOCKSS box through an Apache server as described earlier, you should provide your Serials Solutions account manager with the address of the Apache server. For example, if your Apache server is running on the same host at the default port 80, then your account manager will replace “LOCKSS_RESOLVER” with `lockss.xyzy.edu:80` or just `lockss.xyzy.edu`.

If you request an article, issue, or volume through the 360 Link LOCKSS provider that is preserved on your LOCKSS box, the LOCKSS box will present the the most up-to-date version from either the publisher or the copy that is cached on the LOCKSS box. For example, if you request a specific article, you will be presented with the article. If you request just the issue or volume, you will be taken to the issue or volume table of contents page. If you request just the title, or the volume you requested is not preserved on your LOCKSS box, you will be presented with a list of volumes for that title that are preserved on your LOCKSS box, along with a link to title page at the publisher. This enables you to navigate to content in a preserved volume, or go to the title table of contents page on the publisher's site.

Troubleshooting

If you cannot access an article from 360 Link, here is a sequence of steps that you can take to diagnose the problem.

1. Verify that the LOCKSS provider is properly configured in your 360 Link account.
2. Temporarily set the following configuration parameter in your LOCKSS box Expert Config screen. This causes the OpenURL resolver to output detailed information about how it handles queries.

```
org.lockss.log.OpenUrlResolver.level=debug3
```

3. Send an OpenURL query to your LOCKSS box Content Server for an article that is preserved. Here is a sample OpenURL query:

```
http://lockss.xyzy.org:8082/ServeContent?issn=1085-4908&volume=25&issue=1&spage=7
```

4. Check your LOCKSS log file by selecting the Logs link in the list of actions on the right of your LOCKSS box administrative GUI, and the the 'daemon' file from the log directory list. The most recent lines should display *OpenUrlResolver* log information about the query.

If you cannot determine the cause of the problem, contact LOCKSS support at support@lockss.org.