

Summary of the z39.50 stress test on the Catalog of U.S. Government Publications, held on December 9, 2009

Background:

Library Services and Content Management (LSCM) opened z39.50 access to GPO's Aleph 500 Integrated Library System (ILS) in March, 2008. The level of access began at the Aleph default per session limit of 1,000 records. The z39.50 per session limit was raised to 10,000 on October 1, 2009, in response to requests from the depository library community.

Can z39.50 access be opened to all interested users? Currently access to the Aleph ILS via z39.50 is limited by password to participants in the Federal Depository Library Program (FDLP). There has been a question in the mind of automation librarian staff in Library Technical Information Services (LTIS) within LSCM and contract ILS support staff on the ability of existing, four-year old ILS hardware to support the current per session level of z39.50 access and possible future open access. This had never been put to the test.

Summary of the test:

On December 9, 2009 at approximately 2:30 pm until approximately 3:00 pm EST, LSCM staff in coordination with volunteers in five FDLP libraries conducted a stress test on our Aleph production server primarily to test what impact if any the 10,000 record limit in z39.50 has on internal staff's workflow and on the public using the CGP, on success and speed of transactions and continued connectivity. The internal testers were coordinated to do regular searches shown to retrieve large result sets in both the CGP and the Aleph GUI, connect to PURLs from records in CGP results lists, create brief bib records in Aleph, and export full cataloging records from OCLC to Aleph. The FDLP volunteers were asked to do a number of z39.50 searches during the test period that were known to return large result sets. This test activity was comparable to what could actually occur during normal usage and was not exaggerated; however, there is not a way in which this test could potentially mimic maximum usage, so the results may fall on the conservative side.

Methodology of the test:

External users from five depository libraries conducted searches that would return large results sets, near or over the 10,000 limit, for the duration of the test using a z39.50 client of their choice. Automation librarians designed searches for the testers that met this criterion. Testers also performed searches of their own construction.

Internal users (sixteen participants) in LTIS performed tasks in Aleph that they would engage in as part of their normal workday: searching the CGP; searching the Aleph GUI;

creating brief bibliographic records; exporting records from OCLC; editing/updating Aleph bibliographic records. These were directed to work in Aleph specifically during the test period. Other staff were unaware of the test period and may or may not have been working in Aleph during that time. Internal testers assigned to search were coordinated to perform regular searches that were known to retrieve large result sets through the Aleph GUI and the Web OPAC.

Test Results:

- The group searching the CGP was the hardest hit in terms of reduced functionality, a concern since the CGP has replaced the print Monthly Catalog as the access point for the public to federal publications cataloged for the FDLP and the Cataloging and Indexing Program. CGP searchers mainly experienced slow response times that continued to slow down as the end of the test period approached for actions including connecting to PURLs, searching, retrieving results sets and navigating through results, and some failed searches. The slowest times were experienced near the end of the test, after 2:55. At 3:01, navigation through a results list speeded up noticeably. See table below for comparison PURL connect times during and one hour after the test and again one month post-test.
- Aleph staff client searches experienced slightly slower response times, few failed searches;
- OCLC export: staff exporting records from OCLC into the ILS reported normal response times for the most part when conducting the export function, but slow response times verifying records in the CGP;
- Other: one automation librarian monitoring at the server level lost connectivity with the server; one cataloger lost connectivity while using the Aleph GUI.

Post-test Findings and Recommendations:

LTIS automation librarians, ILS contractor support, and GPO Information Technology staff met on December 16. The test was an agenda item. All parties agreed, based on the results, that the test showed that the ILS servers, now four years old, lack robustness to support users at increasing levels. There was no server failure, although the server load rose significantly during the test period and the CGP users experienced slowdown as described above. In view of the test results, the group concluded that z39.50 access should be held at the current level, or even fall back to 5,000 records per session. Implementation of any further increases should be put on hold until the replacement of the ILS hardware, scheduled for the 2010 federal fiscal year. The replacement hardware has been ordered, received, and is currently being configured as of the date of this report. The Aleph software will also be upgraded this fiscal year from version 18.01 to version 20.

LSCM Recommendations for the short term:

- Maintain access only by password to z39.50 so that only depository libraries have access;
- Maintain the status quo z39.50 limit of 10,000 records per session while monitoring server loads for demonstration of any need to reduce the number of records per session;
- Conduct another z39.50 stress test once the ILS hardware upgrade is complete and the upgrade of the Aleph software to version 20 is complete, projected for July 2010;
- Take second test results into consideration to plan the way forward for z39.50 access to the Aleph ILS
- Continue to keep the FDLP community informed of z39.50 plans and developments via the FDLP Desktop

Tables:

PURL connect times in CGP during the z39.50 test, one hour post-test and one month post-test

| PURL # | Test connect time (2:50–3:02 pm) | One hour post-test connect time (4:02– 4:07 pm) | One month post- test connect time (1–8–2010, 4:05 pm) |
|-----------|-------------------------------------|---|--|
| LPS115497 | 14 secs | 4 secs | 4 secs |
| FDLP423 | 13 secs | 7 secs | 6 secs |
| LPS117157 | 59 secs | 37 secs | 23 secs |
| LPS112230 | 21 secs | 2 secs | 2 secs |
| LPS116001 | 12 secs | 4 secs | 3 secs |
| LPS103490 | 25 secs | 2 secs | 5 secs |

z39.50 user sessions per month, beginning with the second full month of availability

| Year | Month | Z39.50 Sessions |
|------|-----------|-----------------|
| 2008 | May | 2331 |
| 2008 | June | 2236 |
| 2008 | July | 2507 |
| 2008 | August | 2338 |
| 2008 | September | 2772 |
| 2009 | March | 3891 |
| 2009 | April | 4366 |
| 2009 | May | 2964 |
| 2009 | June | 3553 |
| 2009 | July | 3372 |
| 2009 | August | 3498 |
| 2009 | September | 4333 |
| 2009 | October | 4116 |
| 2009 | November | 4803 |
| 2009 | December | 5132 |
| 2010 | January | 4014 |
| 2010 | February | 4773 |

Thanks to our external testers:

James Jacobs (and colleagues), Stanford U., Jonsson Library

Robbie Sittel, Tulsa City-County Library

Chris Brown, University of Denver

Kathleen Amen, St. Mary's University

Jennie Burroughs, University of Montana