

**ZLOT Project Phase 2:
System Implementation of the LOT
Resource Discovery Service**

**Z39.50 Server Implementation
Issues and Recommendations**

This report is submitted in partial fulfillment of contract # 306-03-8322.

William E. Moen, Ph.D.

Kathleen R. Murray, Ph.D.

Texas Center for Digital Knowledge

University of North Texas
Denton, Texas

June 23, 2003

The following research assistants contributed to this report:

Barbara Schultz-Jones & Scott J. Simon



Contents

Introduction.....	3
LOT RDS Overview.....	3
Implementation Issues	4
Configuration Issues	5
Performance Issues	6
Conclusion.....	7

Introduction

According to the ZLOT Technology Inventory and Assessment survey in 2002, eighty percent (number = 278) of responding libraries (number = 346) indicated they would like citizens located anywhere in Texas to be able to discover their library's resources. Local library implementation of a Z39.50 server to provide access to its online catalog and offering its catalog as a search target for the Library of Texas (LOT) Resource Discovery Service (RDS) will enable libraries to achieve this vision. This report addresses common issues associated with the implementation of a library automation system's Z39.50 server product or module.

The Texas State Library and Archives Commission (TSLAC) requested that the ZLOT Project produce this report for use by the TSLAC and Texas libraries. The intended audience is Texas libraries that implement and manage Z39.50 servers.

The focus of this report is on Z39.50 server issues; local library implementation of Z39.50 client software is not addressed. The ZLOT Project staff gathered information for this report through a review of the literature and interviews with library professionals who have considerable implementation experience with Z39.50 products. The report begins with a brief introduction to the LOT RDS. The remainder of the report describes issues related to Z39.50 server implementation, configuration, and performance. For a brief and useful overview of the Z39.50 standard and its use in Texas, see the online article by Christine Peterson (2003) titled "Metasearching in the Lone Star State".

LOT RDS Overview

The LOT RDS enables Texas library patrons, through a single search interface, to easily send the same search to multiple Texas library catalogs, TexShare databases, and TRAIL (Texas Record and Information Locator). When the LOT RDS begins service in September 2003, it will include many library catalogs as search targets, thus making visible libraries' collections as represented in their online catalogs. To participate in the LOT as a search target, a library catalog must be accessible via a Z39.50 server (Z-server). Additionally, the Z-server should be compliant with the Bath Protocol specification.¹

The goal of the LOT RDS is to enable searching across multiple library catalogs (and other search targets such as TexShare databases) from a single interface without regard for the geographical location of either the searcher or the library catalog. Figure 1 depicts the Z39.50 Distributed Search Model for multiple library catalogs. The TSLAC will host the web-accessible RDS. The RDS application includes a Z39.50 client that will communicate directly with the Z-servers located in Texas libraries. From any web browser, a Texas library patron can access the LOT RDS and type in a search request [1] and the RDS will issue the search to multiple catalogs (or other LOT resources such as TexShare databases) [2]. Results from searches of the multiple catalogs are presented to the user in an easy-to-understand manner [3].

A significant benefit of this architecture is that only the centralized Z39.50 client located within the RDS application must be configured to successfully search the participating libraries' catalogs via their Z39.50 servers. Because client configuration can be quite complex when a variety of non-Bath Profile compliant Z-servers are included, the LOT RDS architecture effectively removes this complexity from local libraries in configuring their local Z-clients while leveraging the RDS investment for all Texas libraries. It is important to note that the LOT RDS intends to be Bath Profile compliant but it is recognized that not all Z39.50 products in Texas libraries meet this goal at this time.

¹ "The Bath Profile is an international Z39.50 specification supporting library applications and resource discovery. It describes and specifies a subset of ANSI/NISO Z39.50-1995, Information Retrieval (Z39.50): Application Service Definition and Protocol Specification (ISO 23950)." -- <http://www.nlc-bnc.ca/bath/ap-bath-e.htm>

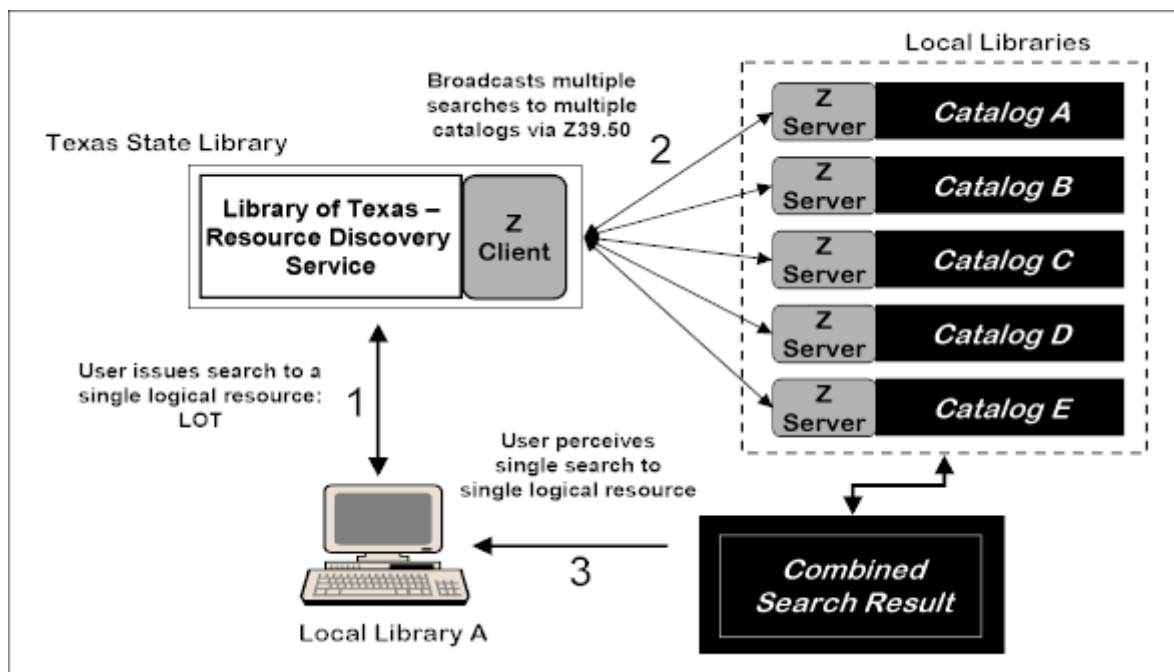


Figure 1. Z39.50 Distributed Search Model

The remainder of this report addresses common Z39.50 server issues and suggests recommendations for dealing with them. The issues fall into three categories: implementation, configuration, and performance.

Implementation Issues

Typically, implementation of Z39.50 server software is straightforward, requires little customization, and can be part of the implementation of an integrated library system. As with any technology implementation, ongoing technical support as well as training for staff and library professionals is important for success. It is helpful to contact other libraries that have active Z39.50 servers and make use of their knowledge and experience. The TSLAC maintains a partial list of active Z39.50 servers in Texas that includes contact information.²

1. Firewalls

Depending on the city or academic institution in which a library is located, certain library automation systems and Z39.50 servers will be located inside the organization's firewall. If the library is inside an effective firewall, the Z39.50 server may not be accessible to Z39.50 clients outside the firewall. If this is the case, no access to the library collection will be available via the LOT RDS. Typically, the solution to this situation is to ask that the firewall 'opens' the port used by the Z39.50 server.

Recommendation

Each library should ensure that their server is accessible from outside of their organizational firewall. Accomplishing this may require some information transfer between the library and the information technology (IT) department responsible for firewall security (e.g., providing the IT staff with the Z-server port number). IT groups may not recognize the unique nature of library services, particularly in regard to information access and resource sharing. Communication of the library's

² <http://www.tsl.state.tx.us/ld/projects/z3950/serverinfo.html>

needs is essential. The IT staff should be made aware that the use of Z39.50 through the firewall does not jeopardize security of the automation system or other library computer systems.

On a related point, there may be occasions when technical changes within the organizational infrastructure impact the operation of the library's systems. It is important that IT support groups have an appreciation for the impact of their decisions and actions on the operations of the automated library. In a networked environment, servers are accessed increasingly from outside the library.

For example, there is an expectation that LOT RDS Z39.50 servers be available as specified by the participating library. Closing hours for the physical library do not indicate that the library is "closed" and therefore that it is a 'good time for' IT to perform maintenance and upgrades on the libraries' systems and servers. Communication with IT support groups is very important in the networked library environment since users are not constrained from accessing library resources based on the open hours of the physical library.

2. Documentation

In most cases a vendor is responsible for configuring and activating a Z39.50 server. There will be occasions, however, when configuration questions will arise and reference material will be required.

Recommendation

Documentation from the vendor should provide adequate information for installation and maintenance of the server. If it is not provided, it should be requested.

3. Bandwidth

Insufficient bandwidth between a library's Z39.50 server and the Internet will result in poor performance, if not in failed server queries.

Recommendation

Access to a library's Z39.50 server requires a broadband (i.e., ISDN or better) Internet connection. It is important to ensure that adequate bandwidth is available prior to allowing access to the library's Z39.50 server. Libraries with lower speed, dial-up Internet connections should not consider implementing Z39.50 servers.

4. Technical Support

Generally, once a Z39.50 server is operational it requires little maintenance. However, problems and questions do arise from time to time.

Recommendation:

It is important to plan for adequate technical support, either from the vendor or an in-house technical group.

Configuration Issues

1. Client Interaction

As previously mentioned, most implementation challenges are encountered when configuring a Z39.50 client to communicate successfully with a variety of Z39.50 servers. From the Z39.50 server perspective, there are times when a Z39.50 client will require a library's Z39.50 server configuration information in order to successfully query the server.

Recommendation

Whenever possible, libraries should request documentation of server configuration defaults and available configuration options. This will minimize the potential complexity involved in configuring a Z39.50 server and make available the configuration information for Z39.50 client implementors (including the RDS application) to correctly configure the client software. Many configuration problems result from poor or insufficient documentation. Additionally, communicating with technical staff at other libraries that have active Z39.50 servers is recommended.

2. Bath Profile

The LOT RDS requires support for the Bath Profile in library Z39.50 servers. Libraries need to know if their Z39.50 server is pre-configured to support the Bath Profile and, further, to what extent the Profile is supported. For example, are all Bath Profile Level 0 queries supported but only some of the Level 1 queries?

Recommendation

It is important that a library identify the level of their Z39.50 server's conformance with the Bath Profile. While some vendors may only offer limited support for the Bath Profile at this time, they may be planning to become fully compliant in the future. A statement of direction from the vendor in regard to support for the Bath Profile might assist a library in making technology plans.

Performance Issues

1. Impact of Increased Catalog Searches

Once a library's catalog and Z39.50 server are available for queries in a distributed Z39.50 searching environment, an increase in the number of queries to the local catalog can be expected. If the hardware platforms for the catalog and server are not optimized for the increased load, performance degradation might result.

Recommendation

Both the library automation system and its Z39.50 server need to have adequate capacity to handle a potential increase in the number of searches to the online catalog. Prior to implementing a Z39.50 server, it is wise to consult with the vendor and verify the platform requirements.

2. Response Time

In the past, Z39.50 servers, operating as search targets in the distributed searching environment, had a reputation of slow response times when faced with heavy traffic loads. However, in recent tests Hammer and Andresen (2002) found that the cause of slow response times was not network congestion or server overload. In each test that resulted in a distributed search operation appearing sluggish, the slowest server in the group being searched caused the delay. However in no case was that Z-server's response time significantly slower than when accessed individually. These results suggest that it is not an inherent weakness in the Z39.50 protocol that causes slow response times but rather sub-optimal Z39.50 server implementations.

Recommendation

Performance implications for Z39.50 queries are not likely to be more resource intensive than web queries of a library's online catalog. Z39.50 server operations should not overwhelm the existing hardware unless it is inadequate. Similar to the previous recommendation, both library automation systems and Z39.50 servers need to have adequate computing capacity to handle anticipated increases in traffic volume.

3. Server Availability

In general an online catalog (via its Z39.50 server) that is listed as a search target of the LOT RDS should be operational seven days a week, 24 hours a day. Unavailable or inconsistently available servers will result in the RDS providing a lower level of quality service than is desirable and possible.

Recommendation

Responsibility for monitoring server operations should be assigned within the library. Management guidelines should be established to outline the actions to take if either the server is not available or if unavailability is a repeated problem. Ensure that there is a regular review of Z39.50 server reports and compare actual performance with established performance standards. Additionally, vendor maintenance and support agreements should reflect the importance of Z39.50 server availability. The RDS application will incorporate functionality to routinely check that Z39.50 servers are running. In the event that a server is not responding, a designated staff person at the library will be notified via email. It is important for libraries to designate a staff member who will receive and take appropriate actions to restore the server when notified.

Conclusion

An effective three-fold strategy to promote a successful Z39.50 implementation consists of communicating library requirements to IT support groups, working closely with library automation vendors, and contacting other Texas libraries that have active Z39.50 servers. Adoption of a common Z39.50 specification, such as the Bath Profile adopted by the LOT RDS, solves many of the configuration and complexity problems encountered in early Z39.50 server implementations. Today's Z39.50 servers can be implemented and maintained with relative ease.

It is important for libraries to identify capacity requirements for their Z39.50 servers as well as for underlying library automation systems. Increased traffic can be managed with adequate capacity planning. Management practices need to be established for ongoing operational and performance issues. In particular, ensuring server availability is an important responsibility of the local library.

The primary goal of the LOT RDS is to expand the resource discovery options of Texas library patrons by enabling searching across multiple library catalogs and databases from a single interface. Texas libraries have the opportunity to participate in the LOT RDS by making their online catalogs accessible through the implementation of a local Z39.50 server. The benefits of implementing a Z39.50 server and identifying it to the LOT RDS include improved exposure of Texas library resources and, hopefully, increased circulation of collections. In short, a successful resource discovery service for Texans is consistent with the service mission of Texas libraries.

Resource List

Abbas, J., Antonelli, M., Gilman, M., Hight, P., Hoski, V., Kearns, J., et al. (1999). *An overview of Z39.50, supplemented by a case study of implementing the zebra server under the linux operating system.* Retrieved May 5, 2003, from <http://www.unt.edu/wmoen/Z3950/GIZMO/section1.htm>.

ANSI/NISO Z39.50-1995, *Information Retrieval (Z39.50): Application Service Definition and Protocol Specification (ISO 23950)*. Available at <http://lcweb.loc.gov/z3950/agency/>.

Gregory, Steve. (2002, December 16-17). Z39.50 interoperability: A view from the virtual trenches. Retrieved May 5, 2003, from http://www.unt.edu/zlot/EOP/MeetingFiles/ZLOT_overview_special_meeting_16dec2002.ppt.

Hammer, S. & Andresen, L. (2002). *Issues in Z39.50 parallel searching.* Retrieved May 4th, 2003, from http://www.indexdata.dk/paraz/parallel_search.html

Kunze, J.A. & Warling, B.N. (1996, March). Recent developments in GALEN II: Evolution of a digital library for the health sciences. *D-Lib Magazine*. Retrieved May 5, 2003, from <http://www.dlib.org/dlib/march96/03galen2.html>.

Lynch, C.A. (1997, April). The Z39.50 information retrieval standard part I: A strategic view of its past, present and future. *D-Lib Magazine*. Retrieved May 5, 2003, from <http://www.dlib.org/dlib/april97/04lynch.html>.

Murray, K. R., Moen, W.E., & Simon, S.J. (2002). ZLOT Project deliverable D: Architecture design. Retrieved May 6th, 2003 from <http://www.unt.edu/zlot/ProjectDocs.htm>

Peterson, Christine. (2003). Metasearching in the lone star state. Retrieved June 22, 2003 from <http://libraryjournal.reviewsnews.com/index.asp?layout=article&articleid=CA266434&publication=libraryjournal>