SCIENCE STARTS HERE

Speed, Flexibility, Security: The MySQL solution to a large-store XML dilemma

Mariella Di Giacomo

Beth Goldsmith

Miriam Blake

Mark Martinez

Los Alamos National Laboratory

Research Library

LA-UR-04-0190

Outline

- Introduction/Background
- Multiple installations of MySQL at our site
- SearchPlus, our XML project, known as “XML dilemma”
- Project, Application and Architecture Requirements
- Architecture
- MySQL Implementation and Optimization
- Short Term Future Directions

Los Alamos National Laboratory

The World's Greatest Science Protecting America

- Established 1943
- 43 Square mile site, Northern New Mexico
- Operated by University of California for the National Nuclear Security Administration of the Department of Energy

Who we are

- 12000 employees
- 6000 technical staff
- Technology application & development in:

Accelerators, Advanced Concepts, Astronomy, Biology,Chemistry, Computing, Energy, Environmental, Explosives, Genomics, Manufacturing, Mathematics, Modeling/Simulation, Nonproliferation, Nuclear Materials & Weapons, Physics, Space Sciences, Superconductivity

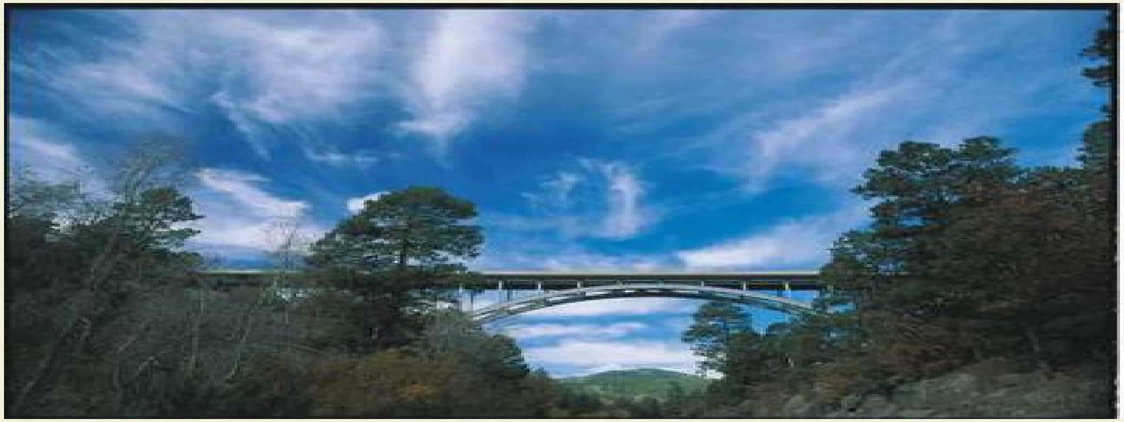

LANL Research Library (RL)

- Support the LANL research mission
- Increase the efficiency of Laboratory research staff through our information & knowledge products and services
- Enhance scientific communication through knowledge discovery, managing intellectual capital, and facilitating scholarly communication

Library Without Walls (LWW)

- Delivery of information to the researchers' desktops, wherever and whenever they need it, from digital library resources
- Goal: To increase the richness of the scientific literature available to scientists through the development of new capabilities that exploit our information repositories and further scientific collaboration

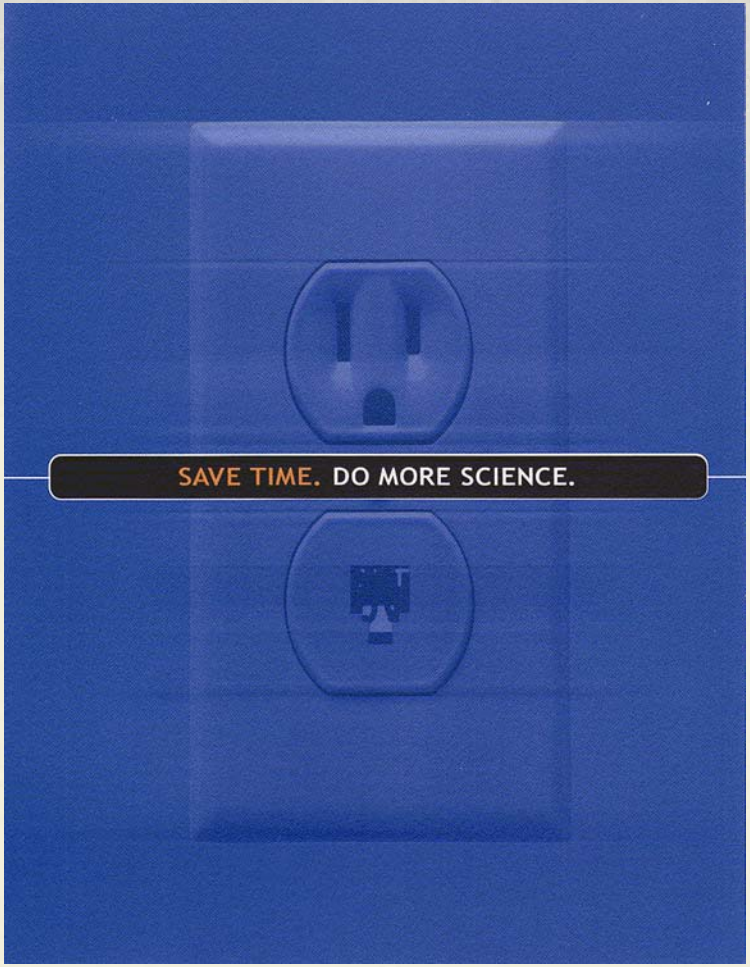

LWW Initiatives

- LANL Publications Digitization & Access
- Electronic journals (Science Server)
- Electronic Databases (SciSearch, Social SciSearch, BIOSIS, Arts & Humanities, DOE, Engineering Index, INSPEC, Nuclear Science Abstracts)
- FlashPoint – multi-database search tool
- LinkSeeker (SFX) – on-the-fly creation of service links
- MyLibrary – Web-based digital library portal

MySQL Installations at the RL

- Nearly 30 Server Installations
- Linux and Solaris Operating Systems
- Mostly Commodity Hardware
- Typical installation uses between 1-10GB
  The Biggest installation utilizes 0.5 TB
- Used for production, batch processing, replication and development

POWERED BY MySQL

Los Alamos NATIONAL LABORATORY

The Next Generation

"What we need now is a brilliant idea…"

SearchPlus

(aka "The XML Project")

- Searching across multiple databases by combining data from 7 databases
- Leverage power of XML for on-the-fly formatting
- De-Duplication
- Linkages between bibliographies and records we store
- Updated interface with new features & access points (links, browsing, counts)

The Team

- 8 developers

  Miriam Blake, Doug Chafe, Mariella Di Giacomo, Frances Knudson, Beth Goldsmith, Mark Martinez, Ming Yu, Jeff Scott

The Data

- 60+ million citations with multiple access points
- 7+ individual databases
  - ✓ SciSearch 1945-present: ~30 M + 4 k weekly
  - ✓ Social SciSearch 1973-present: ~15 M + 1 k weekly
  - ✓ Arts & Humanities 1975-present: ~5 M + .5 k weekly
  - ✓ ISI Conference Proceedings 1990-present: ~3 M + .5 k weekly
  - ✓ INSPEC: ~ 8 M
  - ✓ BIOSIS: ~15 M
  - ✓ Engineering Index: ~5.5 M
  - ✓ Other (DOE, LANL Tech Reports, etc)

RESEARCH LIBRARY

SCIENCE STARTS HERE

RL Data

Full Text

PERGAMON

Acta mater. 48 (2000) 2071–2079

Acta MATERIALIA

www.elsevier.com/locate/actamat

USE OF ABRUPT STRAIN PATH CHANGE FOR DETERMINING SUBSEQUENT YIELD SURFACE: EXPERIMENTAL STUDY WITH METAL SHEETS

T. KUWABARA[1], M. KURODA[2], V. TVERGAARD[3]† and K. NOMURA[1]

[1]Department of Mechanical Systems Engineering, Tokyo University of Agriculture and Technology, 2-24-16, Nakacho, Koganei, Tokyo 184-8588, Japan, [2]Department of Mechanical Systems Engineering, Yamagata University, 4-3-16 Jonan, Yonezawa, Yamagata 992-8510, Japan and [3]Department of Solid Mechanics, Technical University of Denmark, Bld. 404, DK-2800 Lyngby, Denmark

(Received 17 November 1999; accepted 31 January 2000)

Abstract—A basic idea for a method for determining the subsequent yield surface in the vicinity of a current loading point by using an abrupt strain path change has been proposed recently by Kuroda and Tvergaard (*Acta mater.*, 1999, **47**, 3879). The proposed method is applied to real experimental studies. In a biaxial tensile testing apparatus, a cruciform specimen is used, with the strains measured by a biaxial-strain gauge. Then, with the hydraulic pressure of two sets of opposing hydraulic cylinders servo-controlled independently, the testing apparatus can be used to prescribe an abrupt change of the strain path. Both a cold-rolled steel sheet and an aluminum alloy sheet are investigated. The differences between the yield surface shapes found by the strain path change procedure and the shapes found by probing the yield points from

165% 1 of 9 8.26 x 11 in

Los Alamos
NATIONAL LABORATORY

RL Data

Article Metadata

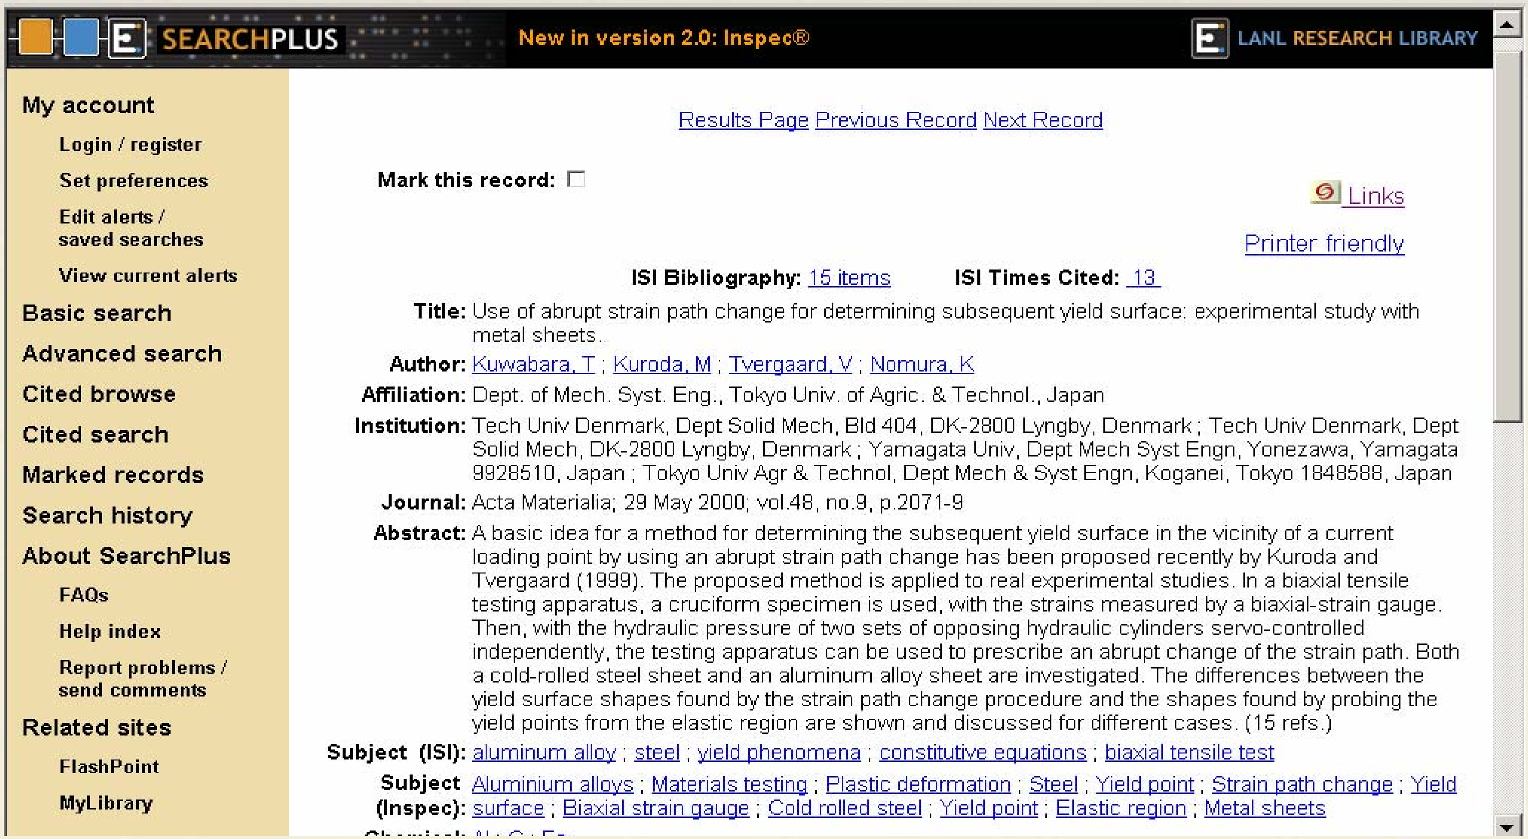

RL Data

Citation (Bibliography) Metadata

Edit alerts / saved searches
View current alerts
Basic search
Advanced search
Cited browse
Cited search
Marked records
Search history
About SearchPlus
FAQs
Help index
Help for this page
Report problems / send comments
Related sites
FlashPoint
MyLibrary

Printer friendly

ISI Bibliography

Title: Use of abrupt strain path change for determining subsequent yield surface: experimental study with metal sheets.
Author: Kuwabara, T.; Kuroda, M.; Tvergaard, V.; Nomura, K.
Source: Acta Materialia; 29 May 2000; vol.48, no.9, p.2071-9
ISI Bibliography: 15 item(s) Full record display

Mark all references:

1. HECKER, SS
CONSTITUTIVE EQUATIO; 1976, p.1
Links

2. **An investigation of plastic flow and differential work hardening in orthotropic brass tubes under fluid pressure and axial load.**
Hill, R.; Hecker, SS; Stout, MG
Source: International Journal of Solids and Structures; Nov. 1994; vol.31, no.21, p.2999-3021
Links Full record

3. **Shear band formation in plane strain.**
Hutchinson, JW; Tvergaard, V.
Source: International Journal of Solids and Structures; 1981; vol.17, no.5, p.451-70
Links Full record

4. HUTCHINSON, JW
P ROY SOC LOND A MAT; 1970, v.319, p.247
Links

5. IKEGAMI, K
MECH BEHAV ANISOTROP; 1982, v.115, p.201
Links

Done Internet

Los Alamos NATIONAL LABORATORY

SearchPlus Data Requirements

- Transform the data acquired into a common XML format and store it for indexing and retrieval.

- Process the data in a secure environment behind a firewall.

- Make the data available to users through a flexible, robust and fast web application outside the firewall.

- Build a scalable system, capable of handling weekly content updates and new data sources.

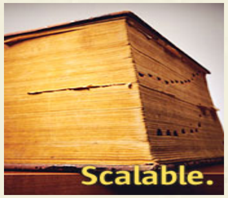

Application Needs

- Empower search, query and analysis across multiple data sources
- Provide links between article cited references and citation articles
- Enable article author browsing

Application Solution

- Search capability.
  Native XML search engines were not meeting our needs. After investigating several full text search engines, we settled on Verity K2 Enterprise

- Browse capability on authors, cited articles and citing articles. Linkages between bibliographies and article metadata.
  The XML data proved to be easily mapped into a Relational Database (DB). After some investigation we chose MySQL.

Browsing

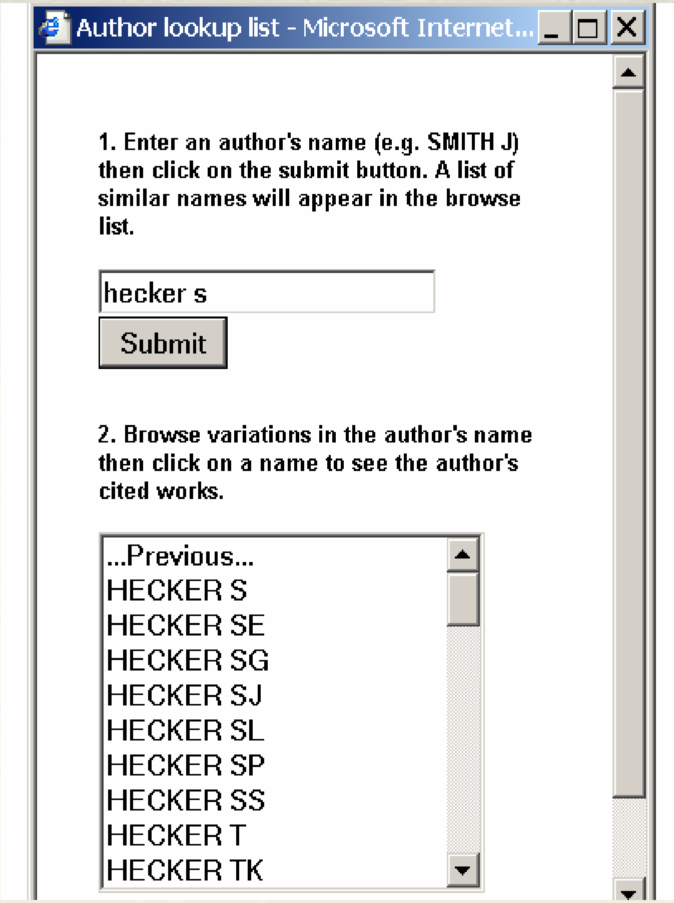
Author lookup list - Microsoft Internet...
1. Enter an author's name (e.g. SMITH J) then click on the submit button. A list of similar names will appear in the browse list.
hecker s
Submit
2. Browse variations in the author's name then click on a name to see the author's cited works.
...Previous...
HECKER S
HECKER SE
HECKER SG
HECKER SJ
HECKER SL
HECKER SP
HECKER SS
HECKER T
HECKER TK

Browse Details

SEARCHPLUS New in version 2.0: Inspec® LANL R

- My account
  - Login / register
  - Set preferences
  - Edit alerts / saved searches
  - View current alerts
- Basic search
- Advanced search
- Cited browse >>
- Cited search
- Marked records
- Search history
- About SearchPlus
  - FAQs
  - Help index
  - Report problems / send comments
- Related sites
  - FlashPoint
  - MyLibrary

Cited works: HECKER SS
(Limits: Year<2000)

Set/Clear Limits Printer friendly

Select All
Unclick to unselect

(90 displayed of 218 total citations)
Next Group
Total times these 90 articles were cited: 699

[2]	HECKER, SS	1999	v.	p.	INT C AG STUD LIF EX	
[3]	HECKER, SS	1997	v. 13	p.6	ISSUES SCI TECHNOL	Full record
[1]	HECKER, SS	1996	v. 13	p.619	METALLURGICAL T A	
[1]	HECKER, SS	1995	v. 20	p.4	MRS BULL	Full record
[18]	HILL, R	1994	v. 31	p.2999	INT J SOLIDS STRUCT	Full record
[2]	HECKER, SS	1994	v.	p.44	ISSUES SCI TECHN SPR	Full record
[1]	HECKER, SS	1994	v.	p.	ISSUES SCI TECHN SPR	
[2]	HECKER, SS	1994	v. 10	p.	ISSUES SCI TECHNOL	Full record
[1]	HECKER, SS	1994	v.	p.	LOS ALAMOS NATL LAB	
[1]	HECKER, SS	1993	v.	p.	LOS ALAMOS SCI	
[1]	HECKER, SS	1992	v. 13	p.	METALL T A	
[1]	HECKER, SS	1988	v.	p.	COMMUNICATION	
[4]	HECKER, SS	1988	v. 14	p.150	METALS HDB	
[26]	SMITH, JL	1985	v. 130	p.151	PHYSICA B #wwAmp; C	Full record
[1]	HECKER, SS	1984	v.	p.37	DEFORMATION PROCESSI	
[25]	HECKER, SS	1984	v.	p.1	DEFORMATION PROCESSI	
[1]	HECKER, SS	1984	v.	p.1	P SEM DEF PROC STRUC	

Times Cited Linkage

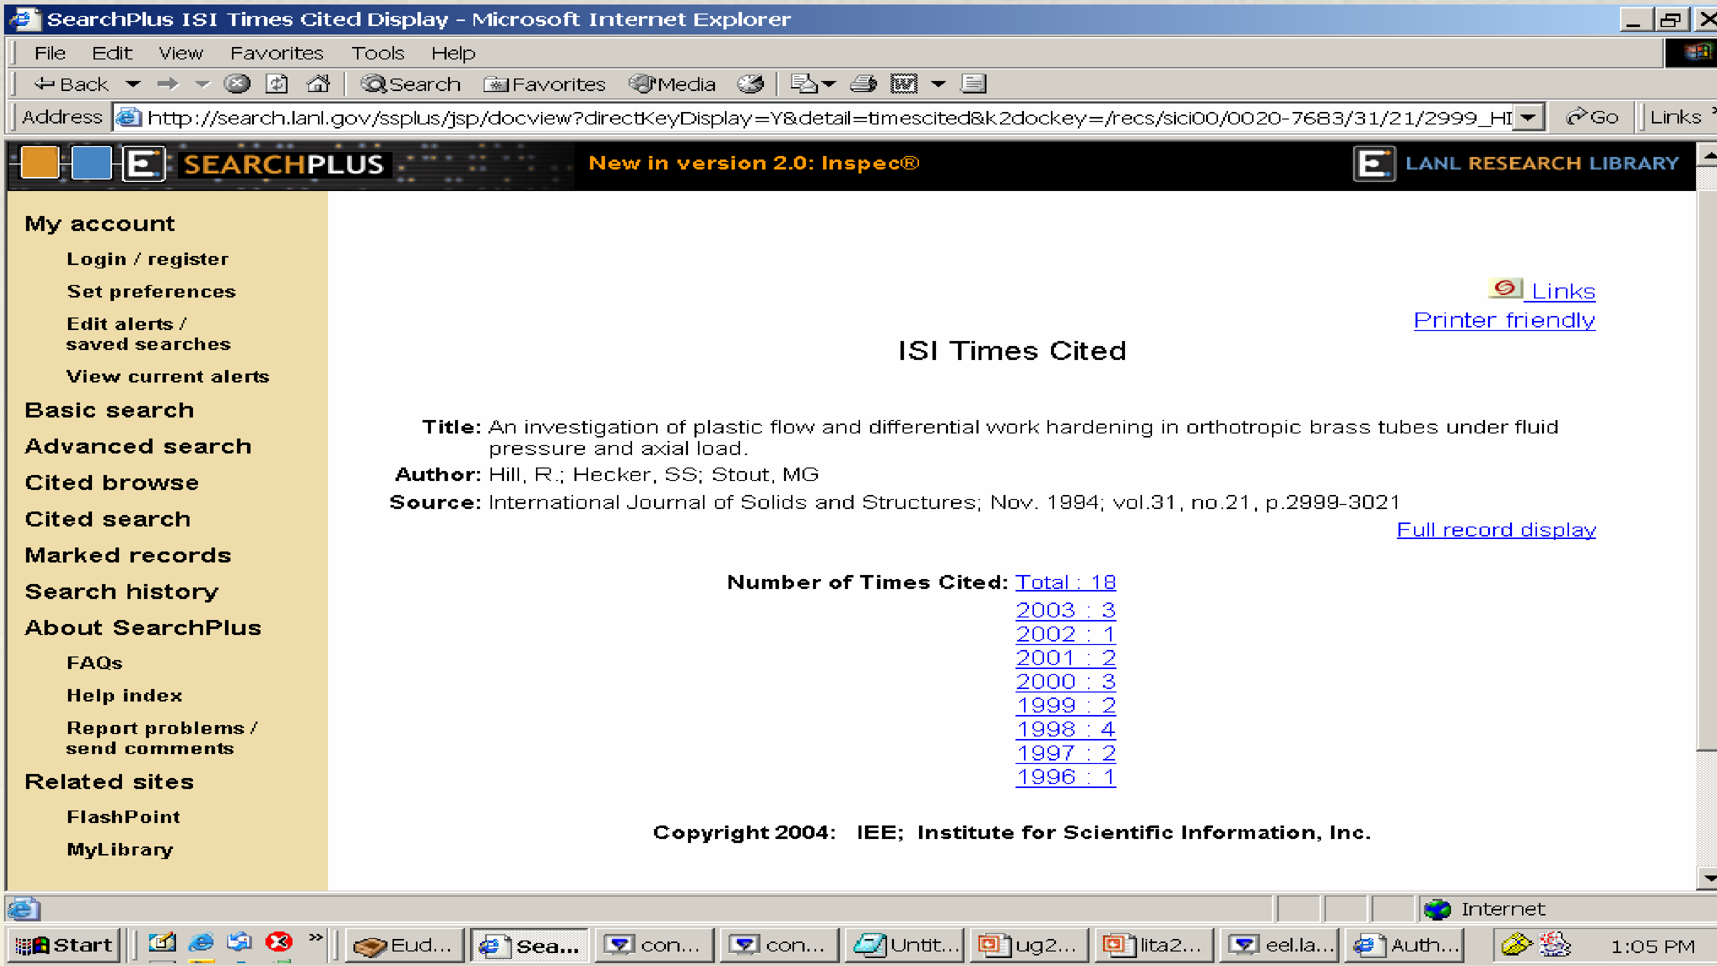

The Citation Data

- Cites – citation data (bibliographies) in each bibliographic record
  - ✓ Searchable separately from the articles which cite them
  - ✓ 500+ Million individual citations (~170M are unique)
  - ✓ Can be search by cited author, source, year, volume or a combination thereof
- One cites XML record can have multiple citations
  - ✓ One for each article cited
- Contain only the briefest of bibliographic (fuzzy matching) details

Citation (Bibliography) Data Example

```
<xml>
 <references>
      .....
 <refItem type="ref">
   <citAu src="cit">HILL, R</citAu>
   <citSo src="cit">INT J SOLIDS STRUCT</citSo>
   <citYear src="cit">1994</citYear>
   <citVol src="cit">31</citVol>
   <citPage src="cit">2999</citPage>
</refItem>
</references>
</xml>
```

Metadata Example

/recs/sici00/0020-7683/31/21/2999_HILL-IPFDWHOBTUFPAL


```
<xml>
<zauthor>
  <author db="Sci">HILL, R</author>
  <author db="Sci">HECKER, SS</author>
  <author db="Sci">STOUT, MG</author>
 </zauthor>
 <zsource>
  <zjournal>
   <journalAbbr db="Sci">Int. J. Solids Struct.</journalAbbr>
   <journalAbbrJ1 db="Sci">INT J SOL S</journalAbbrJ1>
   <journal db="Sci">INTERNATIONAL JOURNAL OF SOLIDS AND
     STRUCTURES; NOV 1994; v.31, no.21, p.2999-3021</journal>
…
</xml>
```

Making the Links

26 M records with bibliographies

Title: xxx

Citation 1

Citation 2

Citation 3

Citation x…

500 M individual citations <refItem>

2. **An investigation of plastic flow and differential work hardening in orthotropic brass tubes under fluid pressure and axial load.**
Hill, R.; Hecker, SS; Stout, MG
Source: International Journal of Solids and Structures; Nov. 1994; vol.31, no.21, p.2999-3021

3 authors
Hill, R (1st author)
Hecker, SS
Stout, MG

Source
(title, year, vol, issue, start page)
International Journal of Solids and Structures, 1994, v.31, #21, p.2999

FullKey for this item:
/recs/sici00/0020-7683/31/21/2999_HILL-IPFDWHOBTUFPAL

XML Cited Reference Example

```xml
<refItem type="ref">
  <fullKey>/recs/sici00/0020-7683/31/21/2999_HILL-PFDWHOBTUFPAL</fullKey>
  <starKey>/recs/sici00/0020-7683/31/*/2999_HILL*</starKey>
  <citAu src="cit">HILL, R</citAu>
  <citAu src="bib">HILL, R</citAu>
  <citAu src="bib">HECKER, SS</citAu>
  <citAu src="bib">STOUT, MG</citAu>
  <citSo src="cit">INT J SOLIDS STRUCT</citSo>
  <citSo src="bib">INTERNATIONAL JOURNAL OF SOLIDS AND STRUCTURES</citSo>
  <citSo src="bib">INT J SOL S</citSo>
  <citYear src="cit">1994</citYear>
  <citVol src="cit">31</citVol>
  <citIssue src="bib">21</citIssue>
  <citPage src="cit">2999</citPage>
  <citEndPg src="bib">3021</citEndPg>
  <citIssn src="bib">0020-7683</citIssn>
</refItem>
```

Cited Linkages

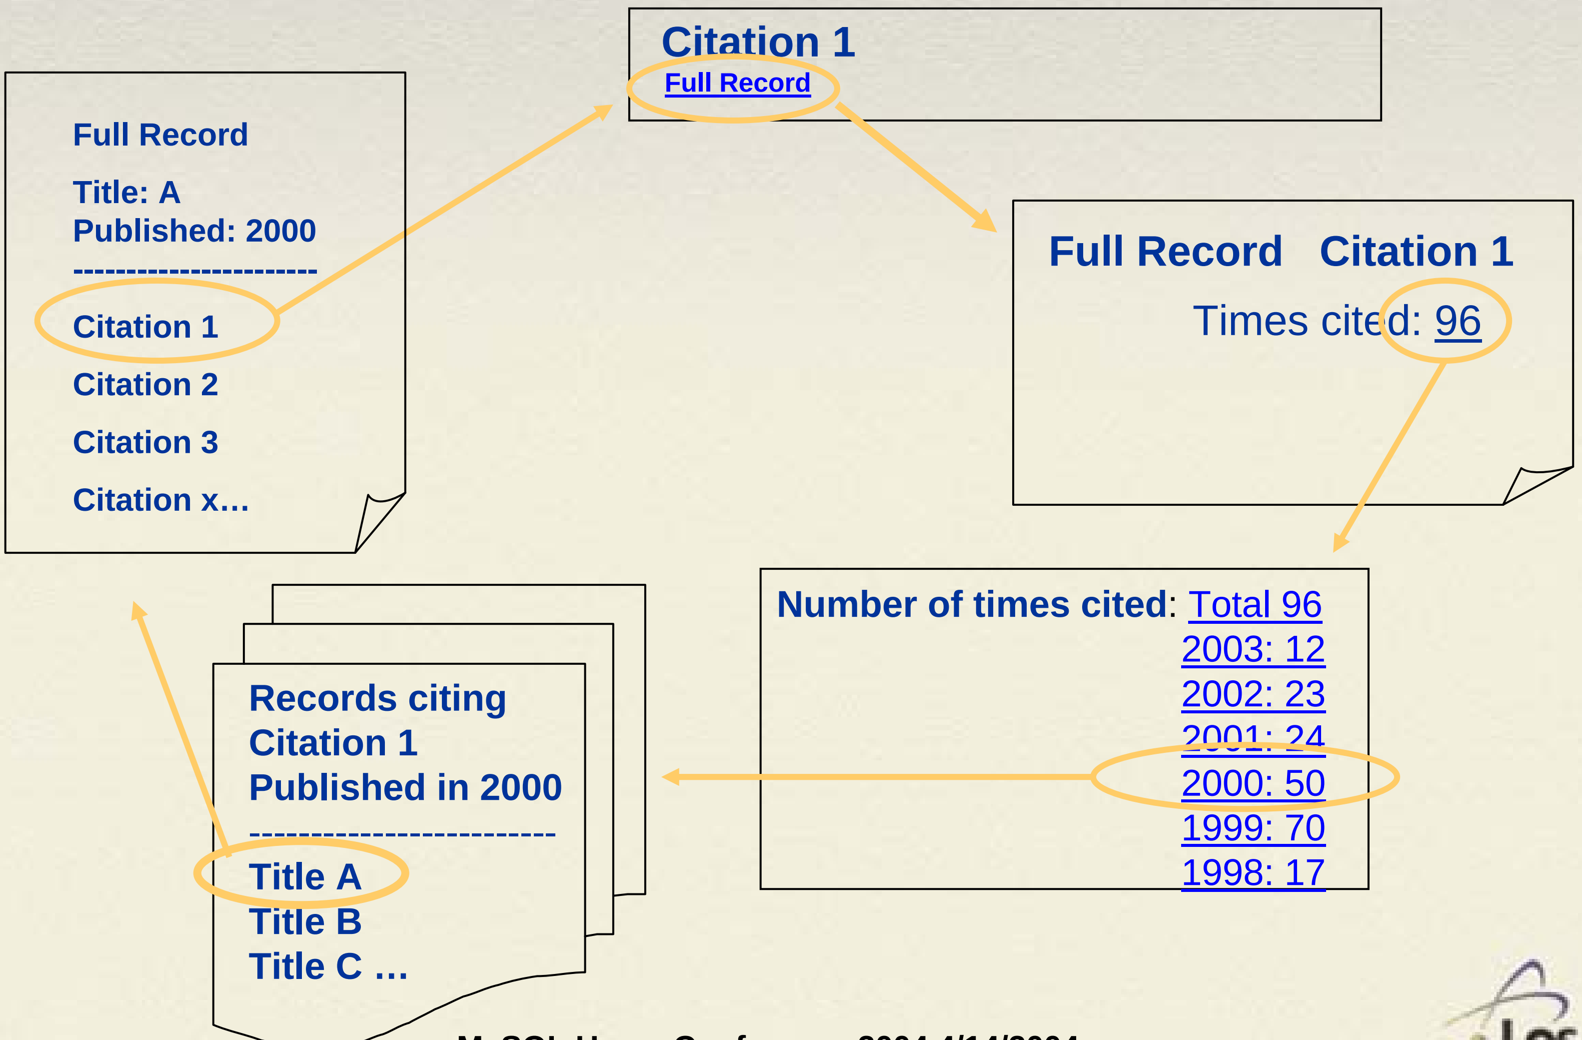

Why a Database ?

Knowledge Discovery & Visualization

- Browsing
- Browse details
- Links between cited references and citation records
- On-the-fly counts
- Backup for file system (bibliography rebuild)

Relational Database

- SEARCHPLUS uses a relational database for browsing.
- In addition to search and retrieval capabilities, we aimed a browse capability on authors, cited and citing articles, and dynamic cite counts.
- A relational database provides flexibility to build links and browsing, especially because our XML data maps very well to a normalized relational structure.

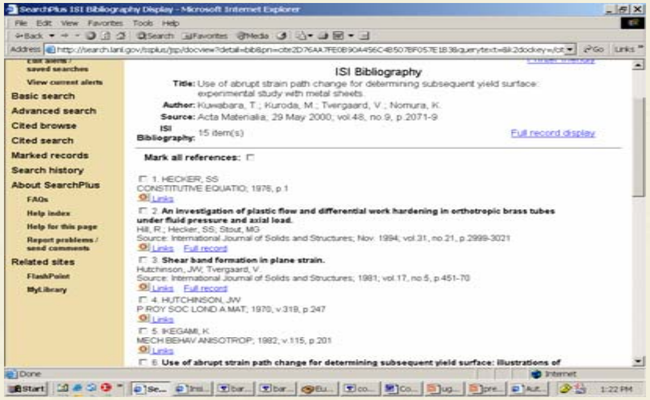

Relational Database

- ➢ SEARCHPLUS XML Data Repository consists of file systems with millions of small files. Backup and restore of such file systems can be problematic. Mirroring the data stored on disk in a relational database offers the additional benefit of faster backup time and useful data redundancy

Relational Database

- Using a relational database for storing user access information.

SEARCHPLUS uses a relational database to authenticate and authorize users accessing the application.

SearchPlus Architecture Requirements

- Storage for millions of XML data files
- A system that has as little service disruption as possible
- A robust, fast, flexible, scalable and secure system
- Process data behind a firewall, read data outside the firewall

SearchPlus Architecture Solution

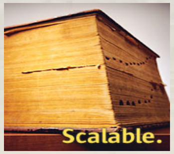



- Scalable, robust, fast and flexible.

Redundant Arrays of Inexpensive Disks (RAID) and Storage Area Network (SAN) technologies have been used to mitigate data failure and provide storage capacity. Redundant systems and MySQL have been used to provide system, application and data redundancy

- Secure environment.

The combination of a SAN and a shared-access file systems has given us the possibility of sharing data among servers located inside and outside the firewall

SearchPlus Hardware Architecture

The whole hardware architecture consists of:

- 12 Processing Nodes.
- 46 Processors.
- 234 GB of Main Memory.
- ~ 6 TB of Disk Storage on a Storage Area Network (SAN).

SearchPlus Software Architecture

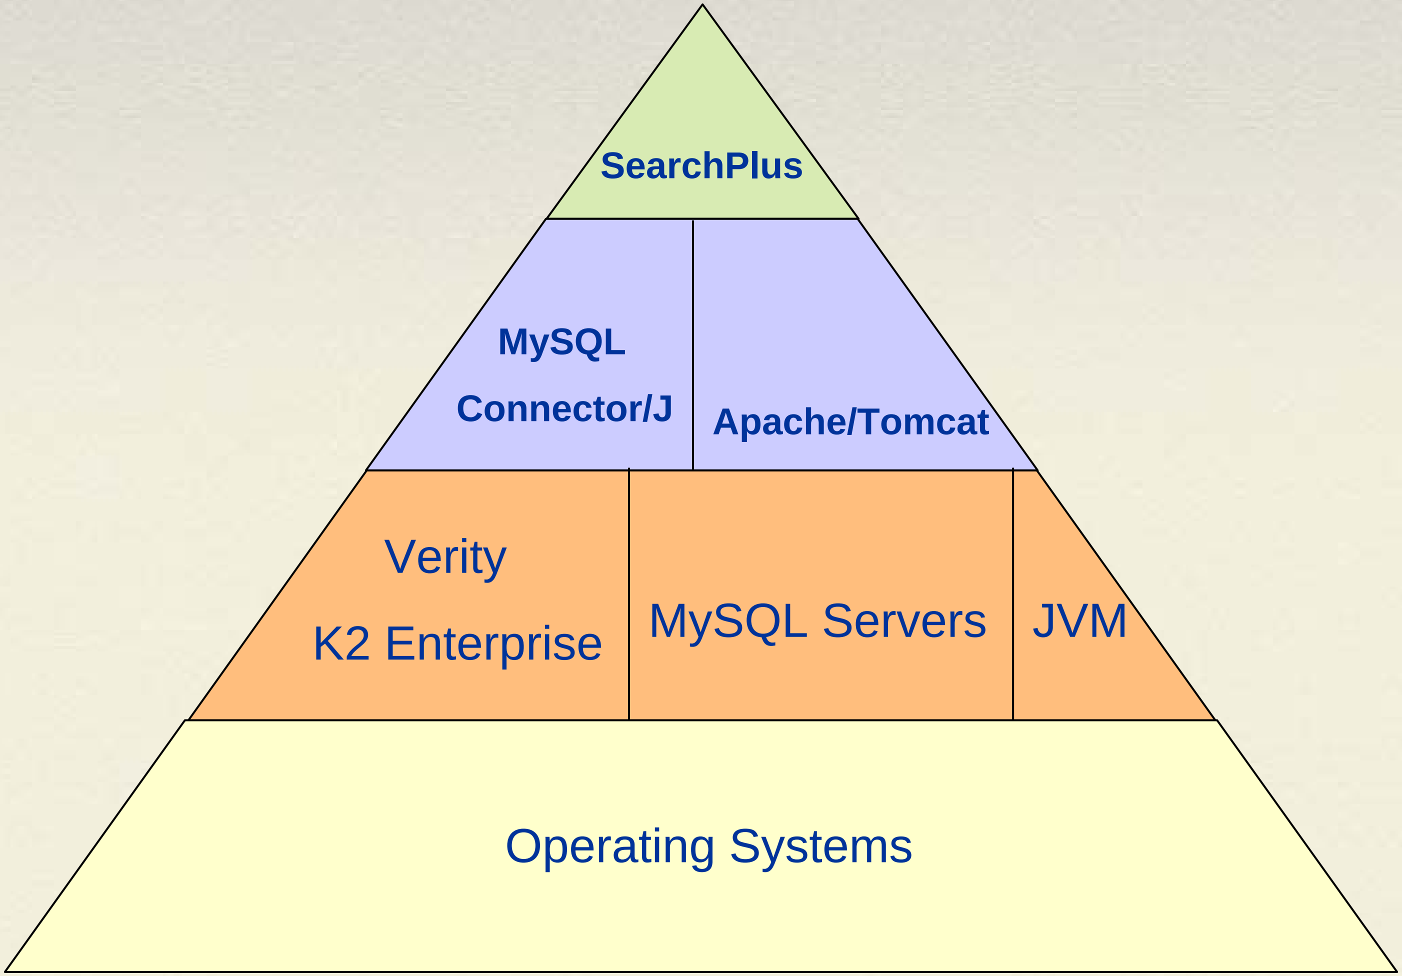

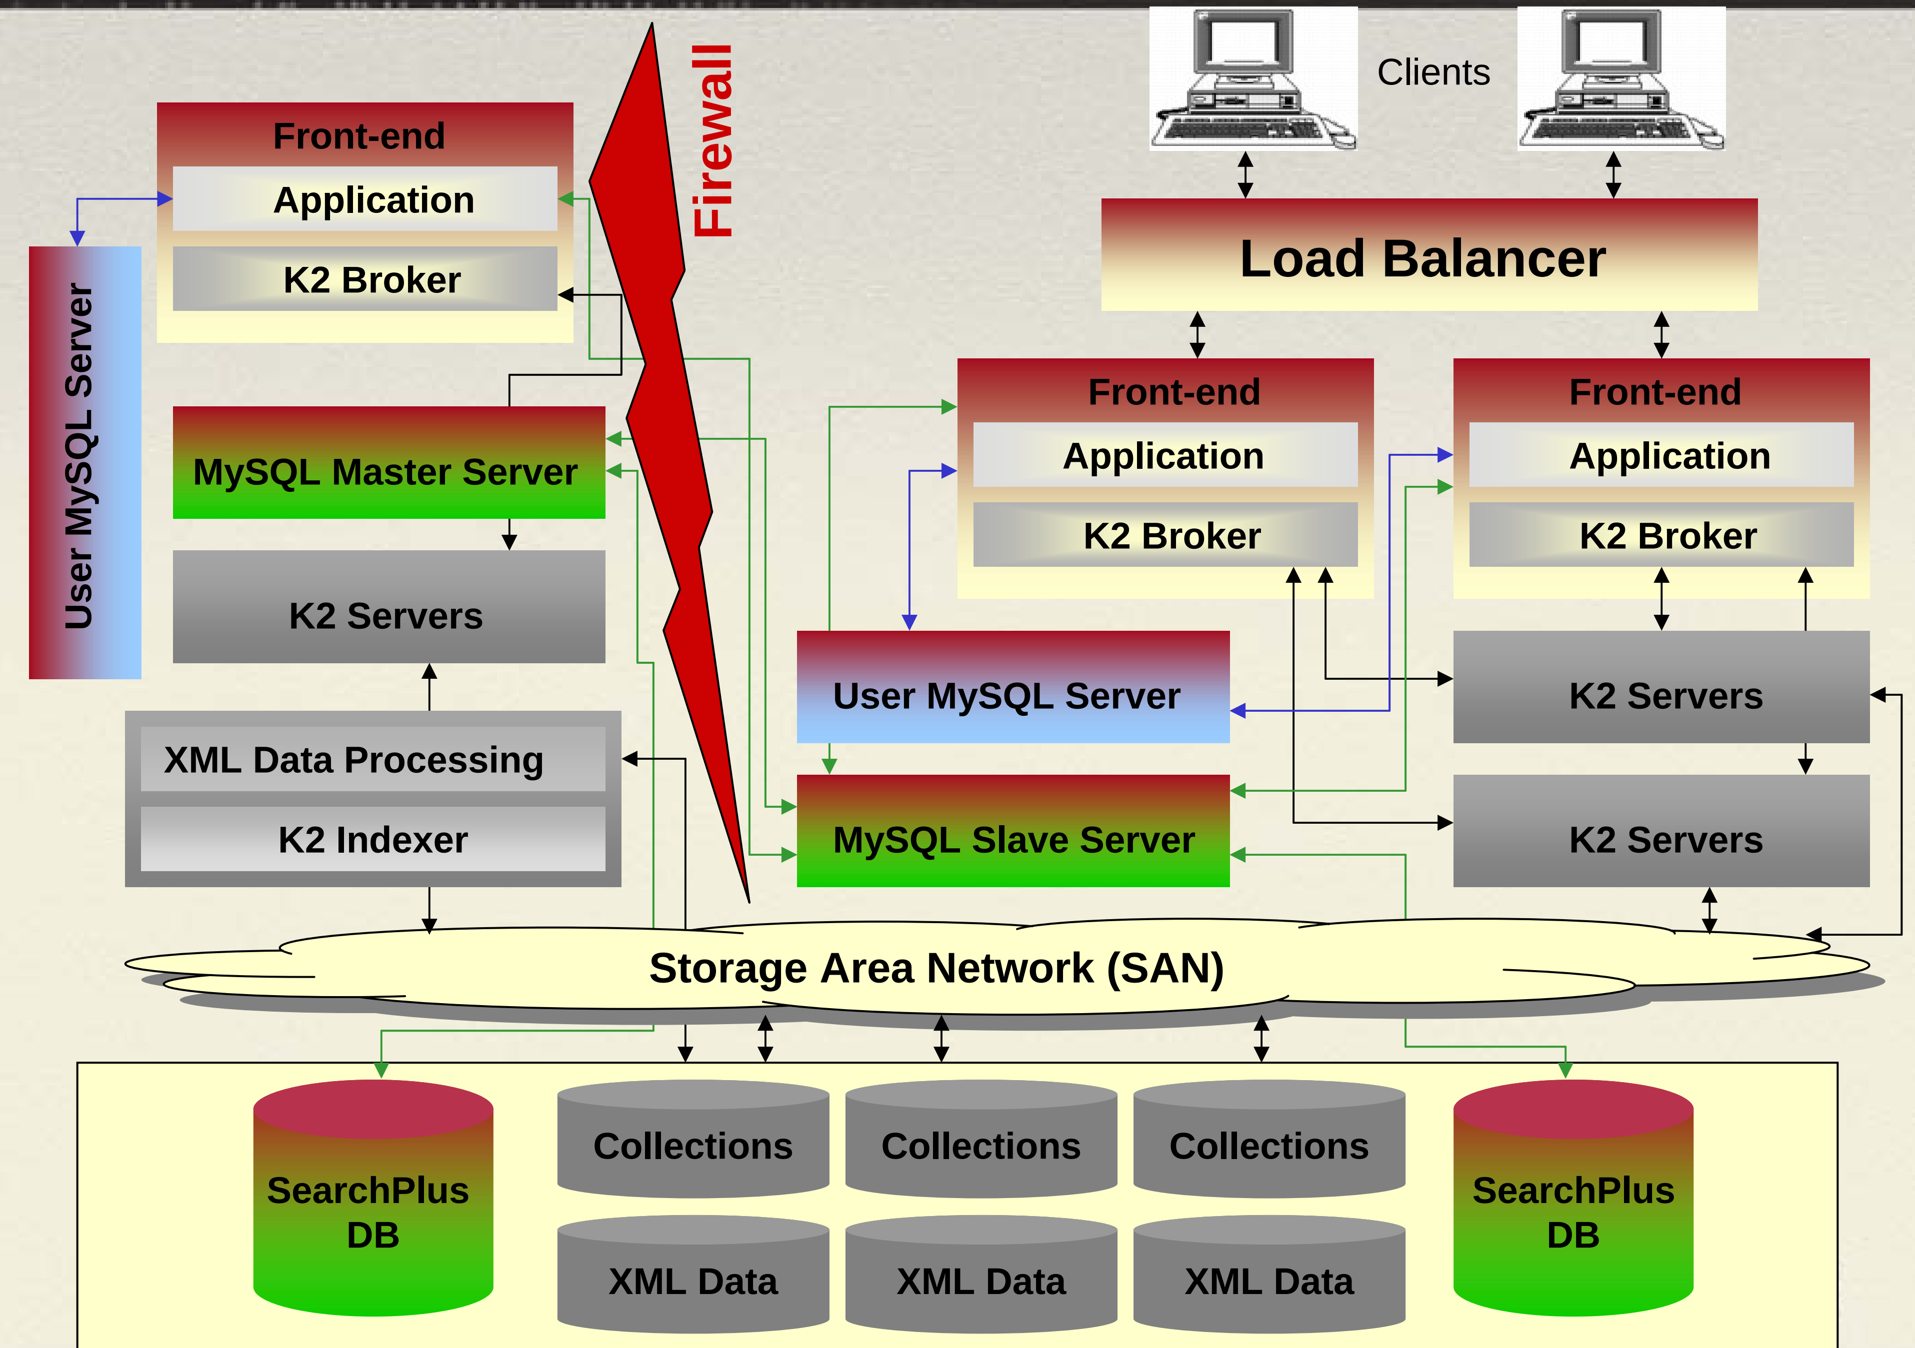
Clients
Load Balancer
Firewall
Front-end
Application
K2 Broker
User MySQL Server
MySQL Master Server
K2 Servers
XML Data Processing
K2 Indexer
Front-end
Application
K2 Broker
Front-end
Application
K2 Broker
User MySQL Server
K2 Servers
MySQL Slave Server
K2 Servers
Storage Area Network (SAN)
SearchPlus DB
Collections
Collections
Collections
XML Data
XML Data
XML Data
SearchPlus DB

MySQL For SearchPlus

MySQL has been chosen on SEARCHPLUS for the following main reasons:

- Open Source Relational Database
- Speed
- Data Storage Capabilities
- Database Design
- Fault Tolerance
- Security
- Replication

MySQL Open Source Software

- MySQL is an Open Source Relational Database.
- MySQL is easy to install, maintain and use.
- MySQL is very well documented.
- Good support through the user's group!

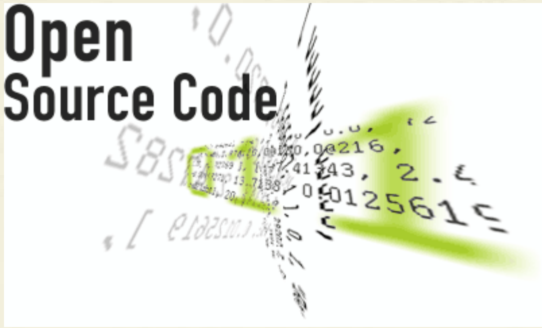

MySL & Speed

- Speed.
  MySQL has proven to be fast at handling links among 1,435,000,000 rows of data in several virtual tables.

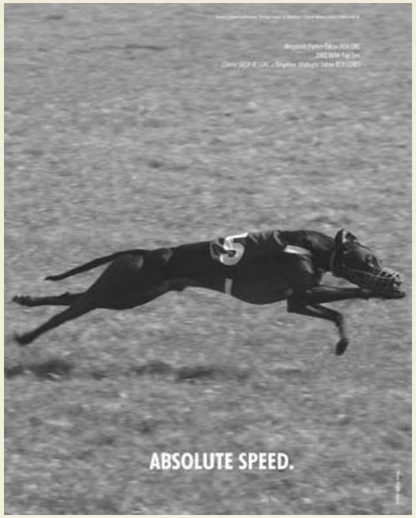

MySQL Data Storage Capabilities

- Data Storage Capabilities.
  We currently manage ~0.5 TB of data in a single database.
  SearchPlus MySQL Database, used as browse capability on authors, cited and citing articles, and dynamic cite counts, consists of 212 MyISAM physical tables and 10 virtual tables (MERGE)

MySQL Data Storage Capabilities

- The database contains a copy of all metadata articles.
  # Records is ~ 40 M
- The database contains a copy of all bibliographies for the articles.
  # Records is ~ 27 M
- The database contains all cited articles.
  # Cited Articles is ~ 500 M
  # Unique Cited Articles is ~ 170 M
- The database contains all cited authors.
  # Cited Authors is ~ 202 M
  # Unique Cited Authors is ~ 12 M

MySQL Database Design

MySQL supports several table types. Our table type choice has been MyISAM, for the following reasons:

- Each table data is stored in a MYD file
- Each table indices are stored in a MYI file
- Memory key buffer (key cache) for the indexed data
- Good for high volume of writes or reads, but not both concurrently. Locking happens at table level
- Variation used: Merge

MySQL & Fault Tolerance

- Fault Tolerance.
  We use MySQL in production and as disk-based backup for our data.

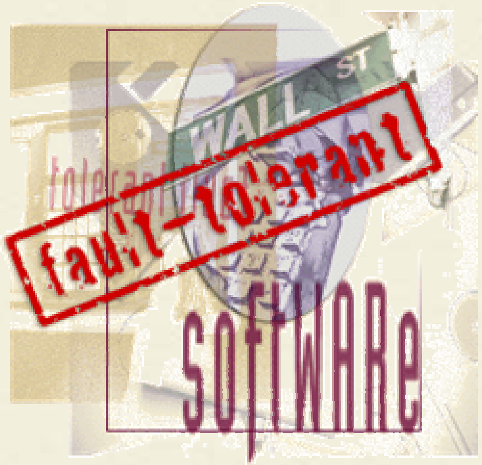

MySQL & High Availability

- Three very reliable servers are used for maintaining our SearchPlus MySQL database.
- We are using MySQL Database Replication for balancing the load and avoiding single point of failure outside the firewall.

MySQL & Security

- MySQL Replication allows us to update our data inside the firewall using a Master Server, while Slave Servers outside the firewall are used to provide the data with read only permission to customers

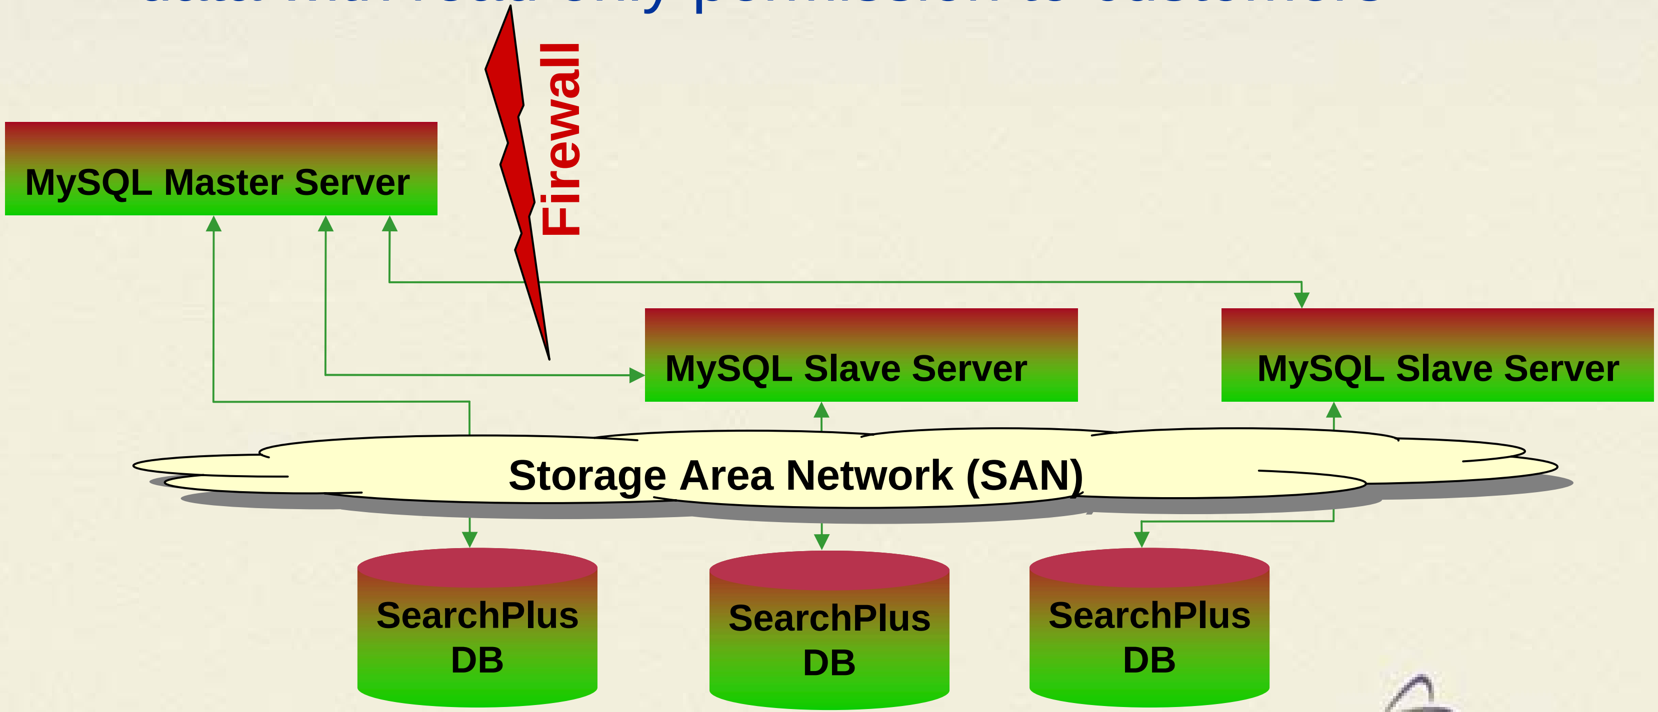

Hardware

MySQL Servers run on dedicated systems:

- Sun Solaris V480 servers
- Lots of memory 16-32 GB
- Good Disk I/O Performance
- Lots of Disk Space
- The servers are closely physically located

MySQL Replication

- What is MySQL Replication
- Why we use MySQL Replication on SearchPlus
- Replication Configuration and topology chosen
- Replication Setup
- Possible Problems

MySQL Replication is

- Streams of queries that allow the databases on one MySQL Server to be duplicated on another MySQL Server
- Very light-weight on the master server
- Relatively Fast
- Easy to configure
- Asynchronous

Why MySQL Replication on SearchPlus

- Isolate the main server from users (security)
- Provides Redundancy and Fault-Tolerance
- Make Backup Easier
- Scaling: Load-Balancing
- Test Environment

Replication Configuration on SearchPlus

- Topology:
  One Master Server and Two Slave Servers
- Master Server records all write/update queries in the binary log
- Slave Servers read the binary log from the Master and run the queries locally
- Master and Slaves selectively filter queries
- Slave operation is multi-threaded since version 4.0
  The SQL Thread uses the local relay log.
  The Relay Thread connects to master and copies queries to the relay log

Replication setup

- Configure replication account on the Master
- Enable binary log on the Master
- Snapshot of the Master databases and reset log
- Install snapshot on the Slave server
- Setup replication using MySQL configuration file
- Restart the Slave
- Check Slave status and error log
- Keep the servers “close” because of network latency

Possible Problems using MySQL Replication

- Communication between Master server and Slave servers may break.
  Check Slave Status

- Binary (or relay) logs use all the disk space

Optimization

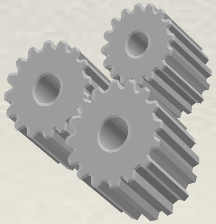

- **Why Optimize ?**
  Get more Performance with same Hardware.
  As your data grows, Performance may degrade

- **What Optimize ?**
  Operating System, Hardware Architecture, Application Components and MySQL

- **Where Optimize ?**
  Monitor your systems and applications and watch for possible bottlenecks

MySQL Optimization

- Server Compilation
- Server Configuration/Tuning
- Table Structure (Database Design)
- Table allocation
- Query Handling
- Data Loading
- Application Design
- Concurrency

MySQL Server Compilation

- Compilation using proper optimizations made improvements of 15% of speed
- Comparisons between MySQL server compiled for 32-bit architecture versus 64-bit architecture.
  64-bit MySQL server, even with some slight degradation compared to the 32-bit, allowed us to address 32 GB of memory
- By default MySQL retrieves indices in chunks of 1KB.
  Benchmarks were done using different chunk sizes, such as, 1KB, 8KB and 16KB. 8KB proved to be the best in our case

MySQL Server Configuration

- MySQL assumes little about your hardware system configuration
- The size of the key_buffer used to allocate indices in memory plays an important role
- By default MySQL retrieves indices in chunks of 1KB

MySQL Server Configuration

- Monitor Important Performance Numbers. Queries per second, bytes transferred per second, new connections per second
- Key Cache Efficiency
- Query Cache Efficiency
- Table Cache, Max Temporary Tables
- Max Connections, Max User Connections
- Long Query Time
- Thread Concurrency

MySQL Database Design

After the first Database design, some improvements were made to the table structure:

- Reduced the size of data inside a single table and the index space for some fields
- Reduced some indices on some long fields, to take fewer resources when indices applied on smaller columns
- Designed 10 Virtual Tables (MRG), 212 physical MyISAM tables
- Use of NOT NULL fields where possible

MySQL Database Design

- Use of fixed column sizes.
  MyISAM tables with fixed rows are faster
- Store compressed data when possible for long fields, that are not modified often
- Consider ratio of reads versus writes

MySQL Table Allocation

- Layout the tables that are accessed at the same time (in the same query) among several disks or different file systems
- On those tables for which has not been possible to define fixed rows, run MySQL table check after updates to remove fragmentation (OPTIMIZE or myisamchk) and re-sort the indices
- Store MySQL Logs separately from the database
- Using a journaling file system makes crash recovery faster

MySQL Query Handling

- Use EXPLAIN SELECT to improve queries
- Use Indexes
- Simplify some of the WHERE clauses
- Use UNIONs of SELECT instead of only one SELECT with several conditions in the WHERE clause
- Use query cache selectively if you have lots of writes or lots of data.
  SELECT SQL_CACHE
- Avoid Table Scan

MySQL Data Loading

- REPLACE queries for insertion and/or replacement
- Bulk-Loading (LOAD DATA or mysqlimport)
- Use prepared queries and placeholders when using Perl DBI

MySQL Application design

- Use Connection Pooling, whenever possible
- Use a fast MySQL driver under both Perl and Java
- Use prepared queries and placeholders when using Perl DBI

MySQL Concurrency

- While running updates (UPDATE, REPLACE) on the MySQL tables and retrieving the data (SELECT) from the same set, penalties in performance were high due to mutex locks

Impact of MySQL Optimizations

Compilation	Table Optimization	Block Size	Table Structure	Buffer Cache	Query Cache
15%	80%	5%	20%	25%	15%

Short Term Directions

- Re-Structuring SearchPlus Database Design
- MySQL Replication for SearchPlus Users Database

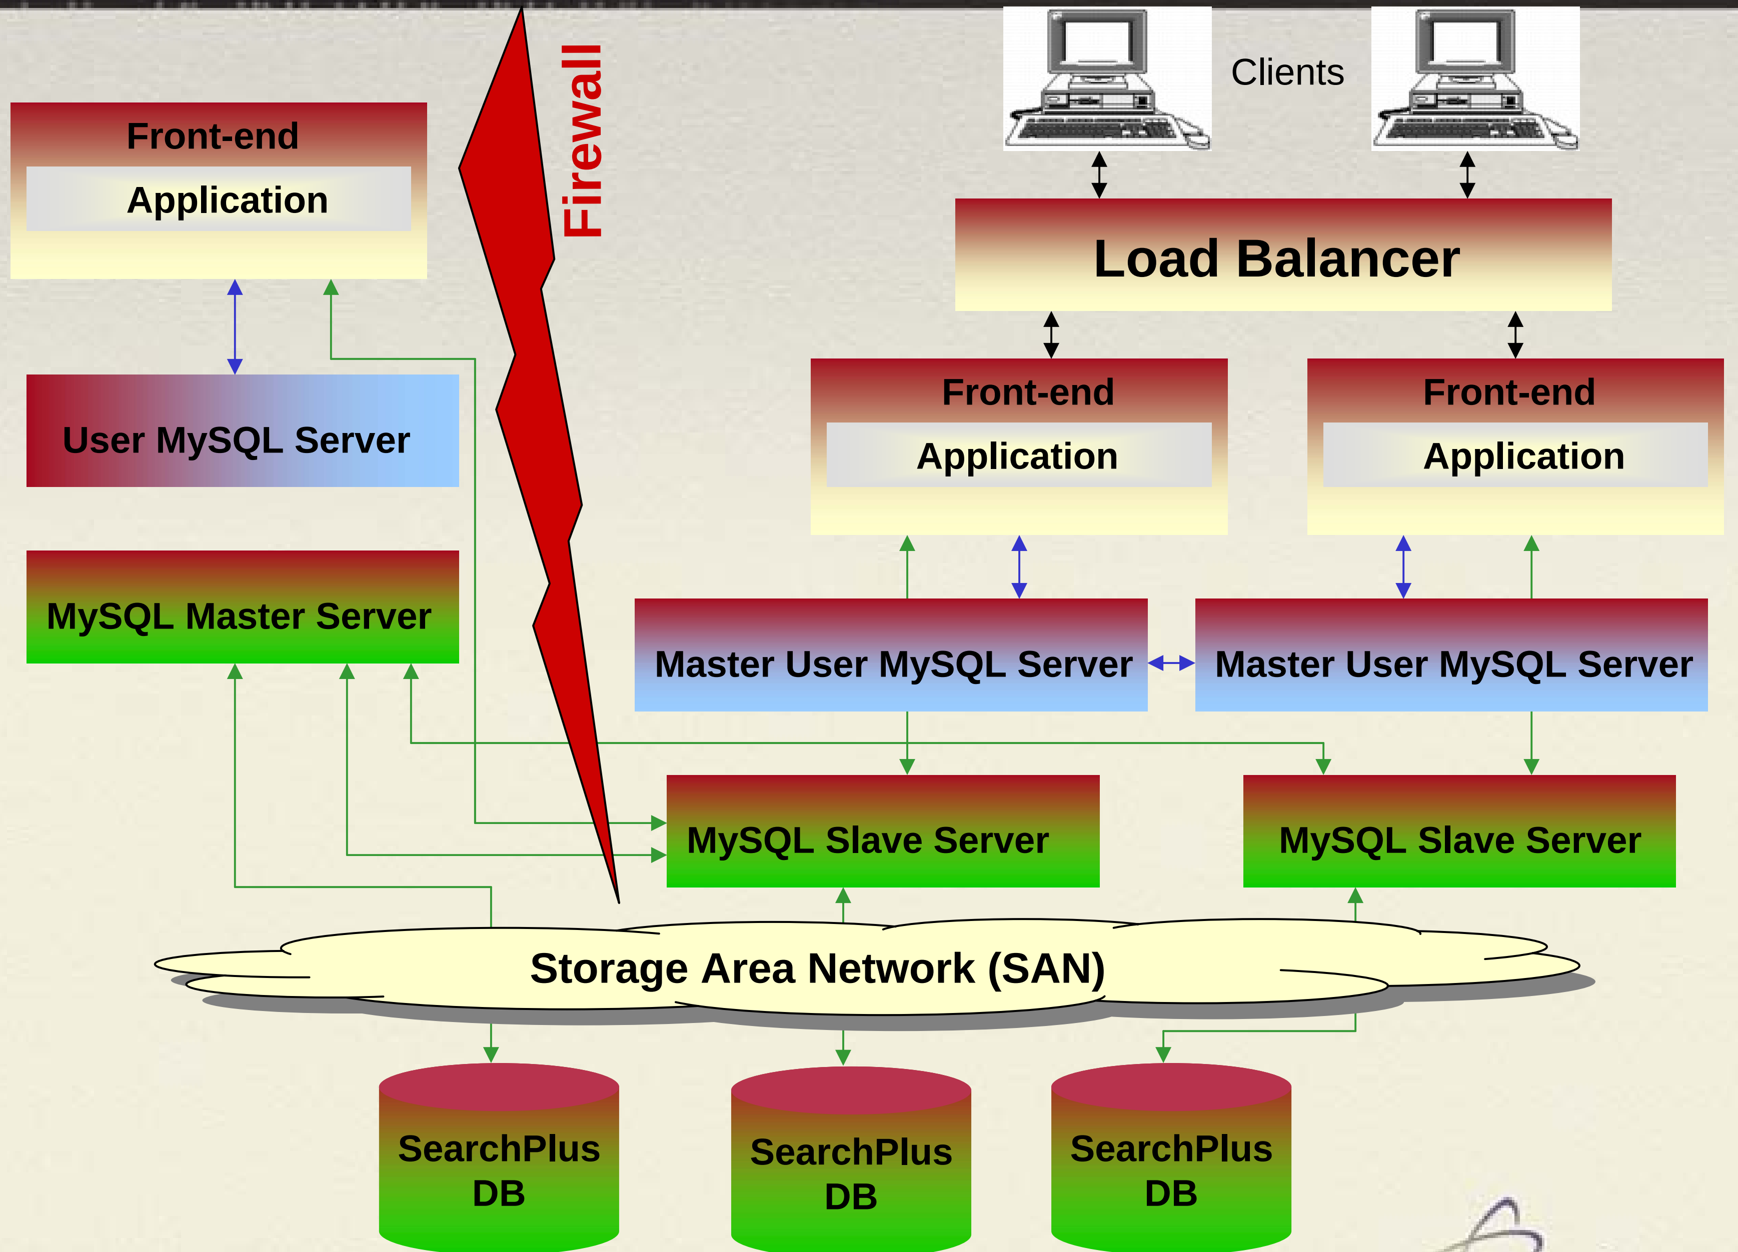
Clients
Load Balancer
Front-end
Application
Front-end
Application
Front-end
Application
Firewall
User MySQL Server
MySQL Master Server
Master User MySQL Server
Master User MySQL Server
MySQL Slave Server
MySQL Slave Server
Storage Area Network (SAN)
SearchPlus DB
SearchPlus DB
SearchPlus DB

Questions ?

Thanks