



Scaling Directories: Design & Deployment Considerations

GUUG 1st Int. Conference on LDAP
Cologne, 07-SEP-2007

Abdi Mohammadi, Sun Microsystems GmbH
Cengiz Tuztas, Sun Microsystems GmbH



Agenda

- Large Scale LDAP Requirements
- Large Scale LDAP Design Goals
- Scalable Architecture and Models
- Deployments
- Q&A

Large Scale LDAP Requirements

Requirements

- Directory is business critical
- High performance on read & write
 - > > 10000 reads/s & > 1000 writes/s
- High available service
- Fast setup & initializing
- Reliable fast backup & recovery
- Multiple backends
- Management & maintenance tools
- Single Point of Access
- Security

Parameters of Scalability

- Usage Pattern
 - > Type, number and size of entries
 - > Number & Type of indexes
 - > DIT and the Data Model
 - > Type & frequency of operations
 - > Duration of peak load
 - > Demand on availability
- Hard- & Software limits
 - > Data throughput (Disk & Network I/O)
 - > Operating System limits
 - > Application bottlenecks

Large Scale LDAP Design Goals

Design Goals (I)

- High performance on read & write operations
 - > Load balanced consumers for read operations
 - > Cache Optimization for read operations by load distribution
 - > Distribution of data for write operations
- High available service
 - > Highly redundant SDN architecture
- Fast setup & initializing
 - > Parallel bulk loading & binary copy
- Reliable fast backup & recovery
 - > Snapshot & binary copy

Design Goals (II)

- Multiple backends
 - > Split the DITs to Subtrees (Subsuffixes) with dedicated backends
 - > Cache Optimization by load distribution
 - > Data partitioning by distribution layer
- Management & maintenance tools
 - > Prepared scripts for setting up and maintaining of the whole environment
 - > SNMP agent
- Single Point of Access
 - > Directory as a single box of Service
 - > Single LDAP box by providing of a VIP for load balanced servers
 - > Dumb clients may not be able to fail-over or rebind
 - > Chaining instead of referrals
 - > Dumb clients may not be able to follow referrals

Design Goals (III)

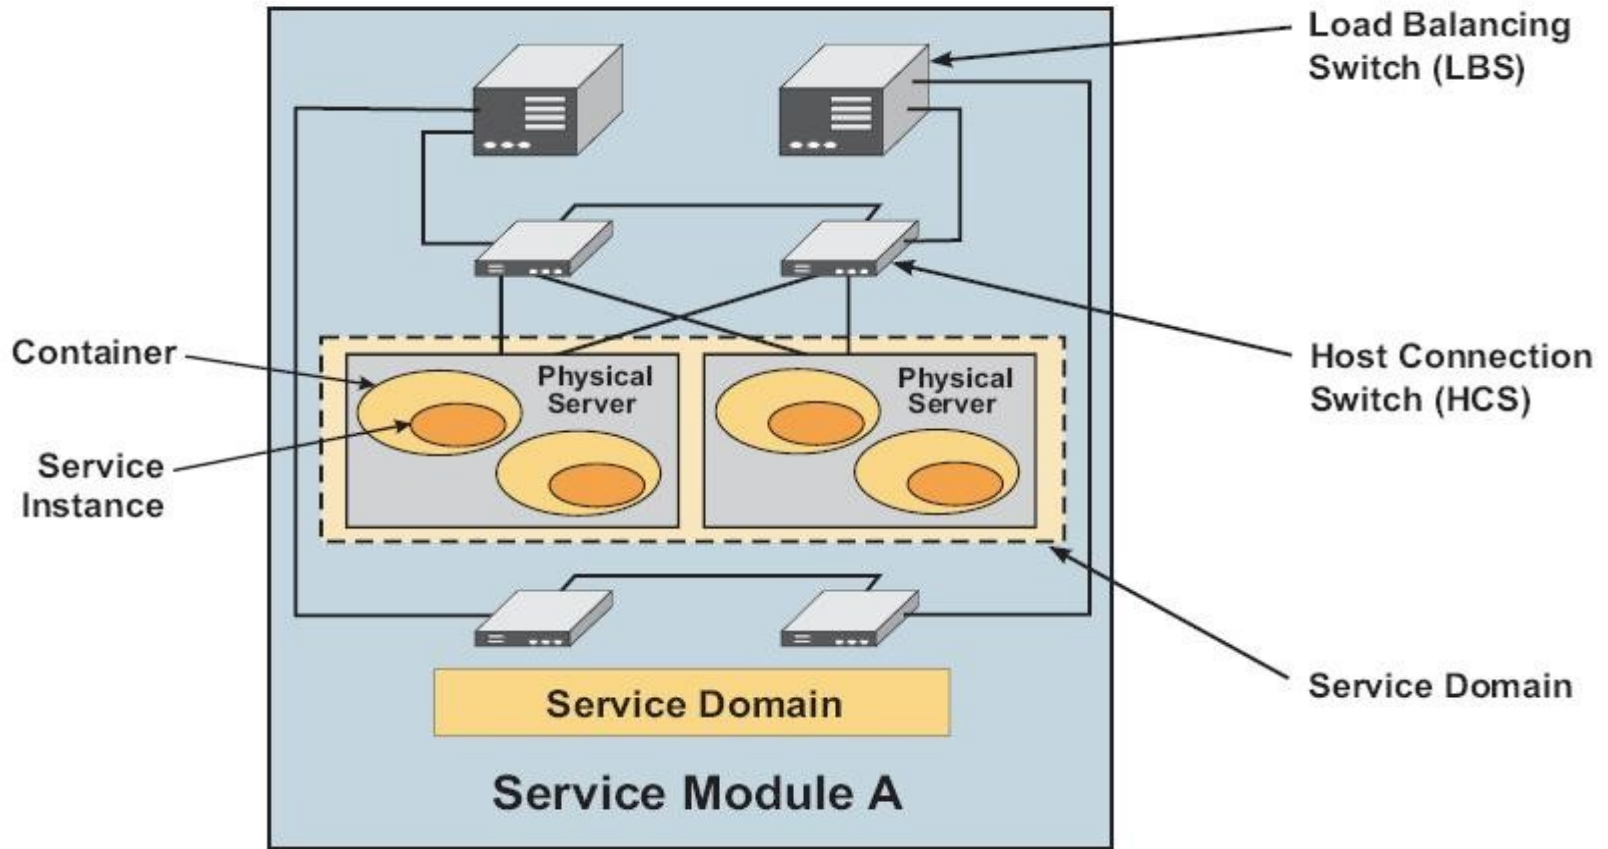
- Security
 - > Different service layers for each type of servers
 - > Port filtering on LB switches
 - > OS Hardening and OS minimization
 - > Dedicated replication network with no access from outside of the box

Scalable Architecture

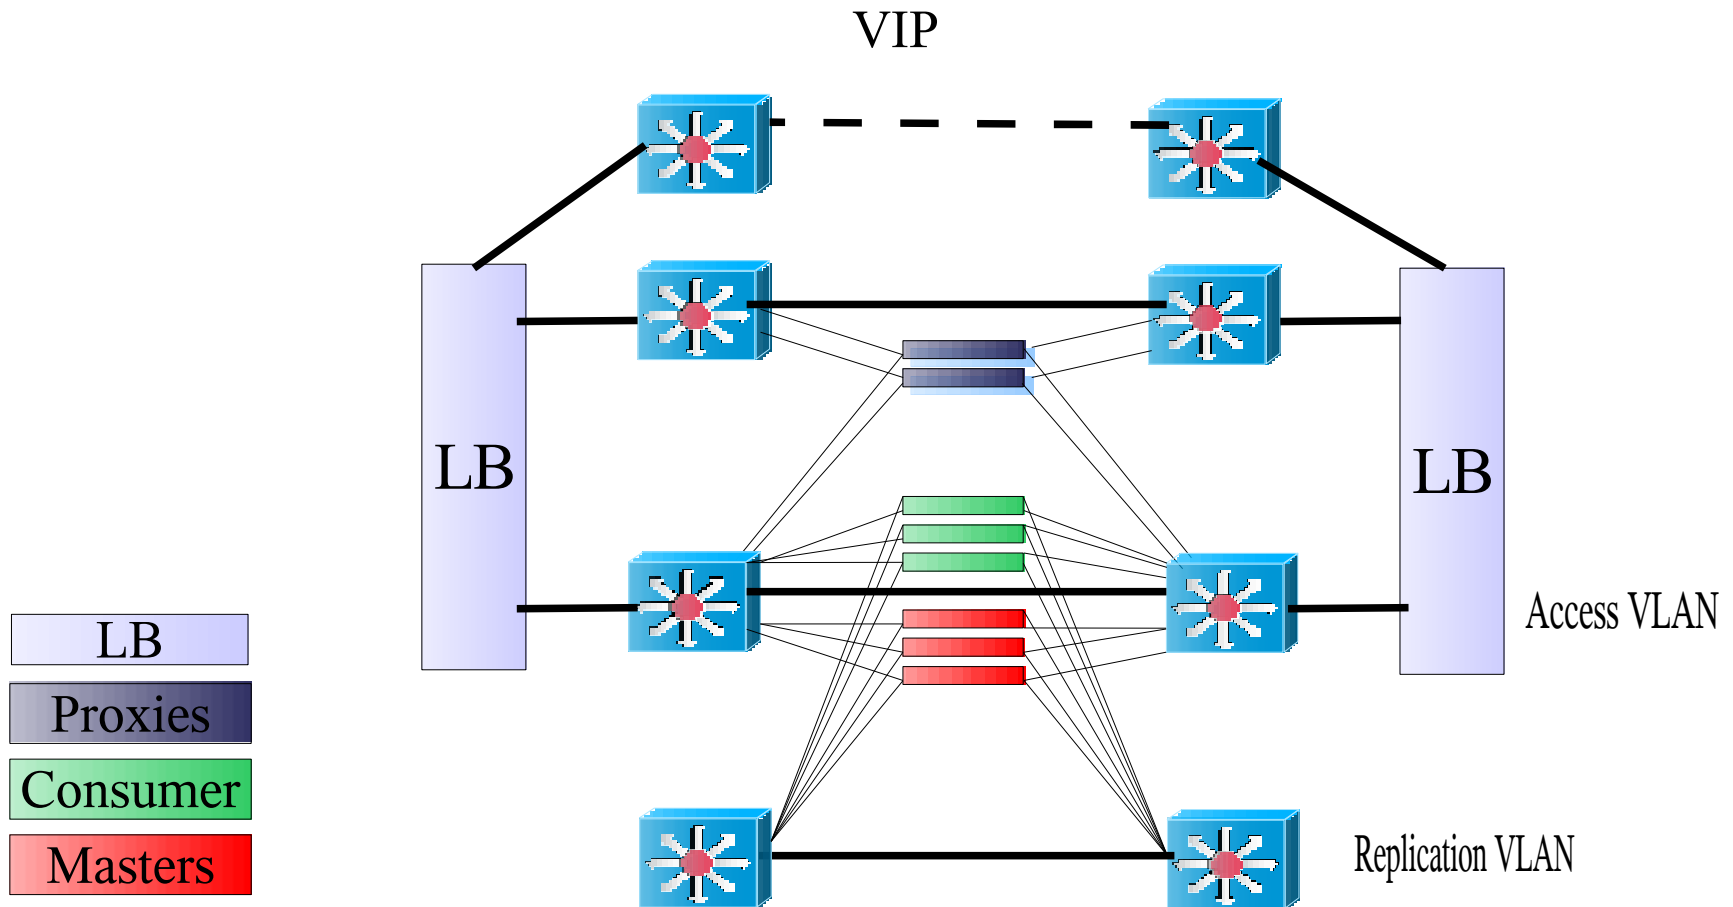
SDN Architecture

- The Service Delivery Network
 - > High availability by redundant LB switches and redundant Layer 2 switches
 - > IPMP
 - > Secured architecture by 'service domains'
 - > One VIP for each 'service domain'
 - > OS hardening and minimization
 - > <http://www.sun.com/blueprints/0905/819-4148.html>

SDN Service Module



SDN for Large Scale LDAP



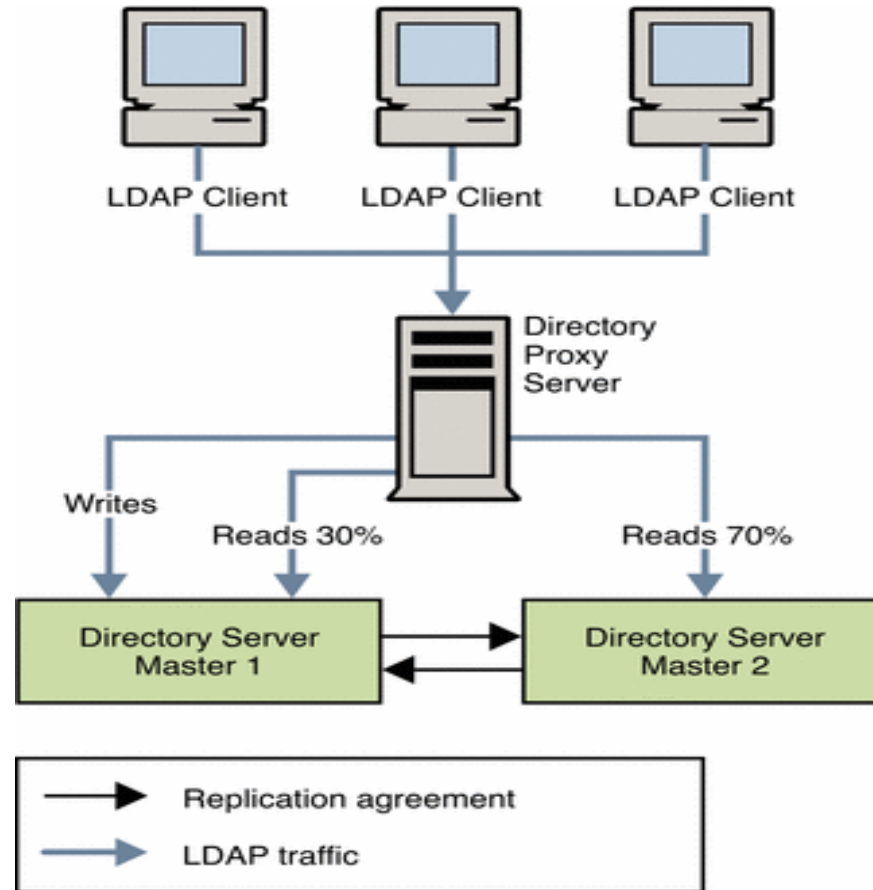
Architecture Considerations

- Multimaster Replication
 - > LAN & WAN
- Proxy follows referrals in behalf of the client
- Partitioned databases (sub-suffixes)
- Configuration server + scripting
- Load balancing
- Disk layout
- Network layout

Single Instance L-Model

- Horizontally scalability on read
 - > Load balanced replicas for read operations
 - > Cache Optimization by load distribution at proxy level
- High availability on write
 - > All writes go to the preferred master, avoids replication conflicts
 - > Second master as failover/backup server
- Load balancing Algorithm
 - > LDAP operational distribution for BIND, SEARCH....
 - > Load balancing and/or fail-over
- Same type of indexing


Load balancing



Distributed XL-Model

- High performance on read & write
 - > Load distribution on Proxy layer
 - > Data partitioning on backend layer
 - > Load balanced Proxy servers for read and write
- Distribution plugin
 - > Various algorithm (custom ?)
 - > Hashing, Integer, alphanumerical..... of the naming attribute
 - > Dealing with MODRDN ??
 - > Redistribution of entries by Delete then Add (Provisioning)
 - > Dealing with Uniqueness of attribute values
 - > Parallel Broadcasting of non matching search filters

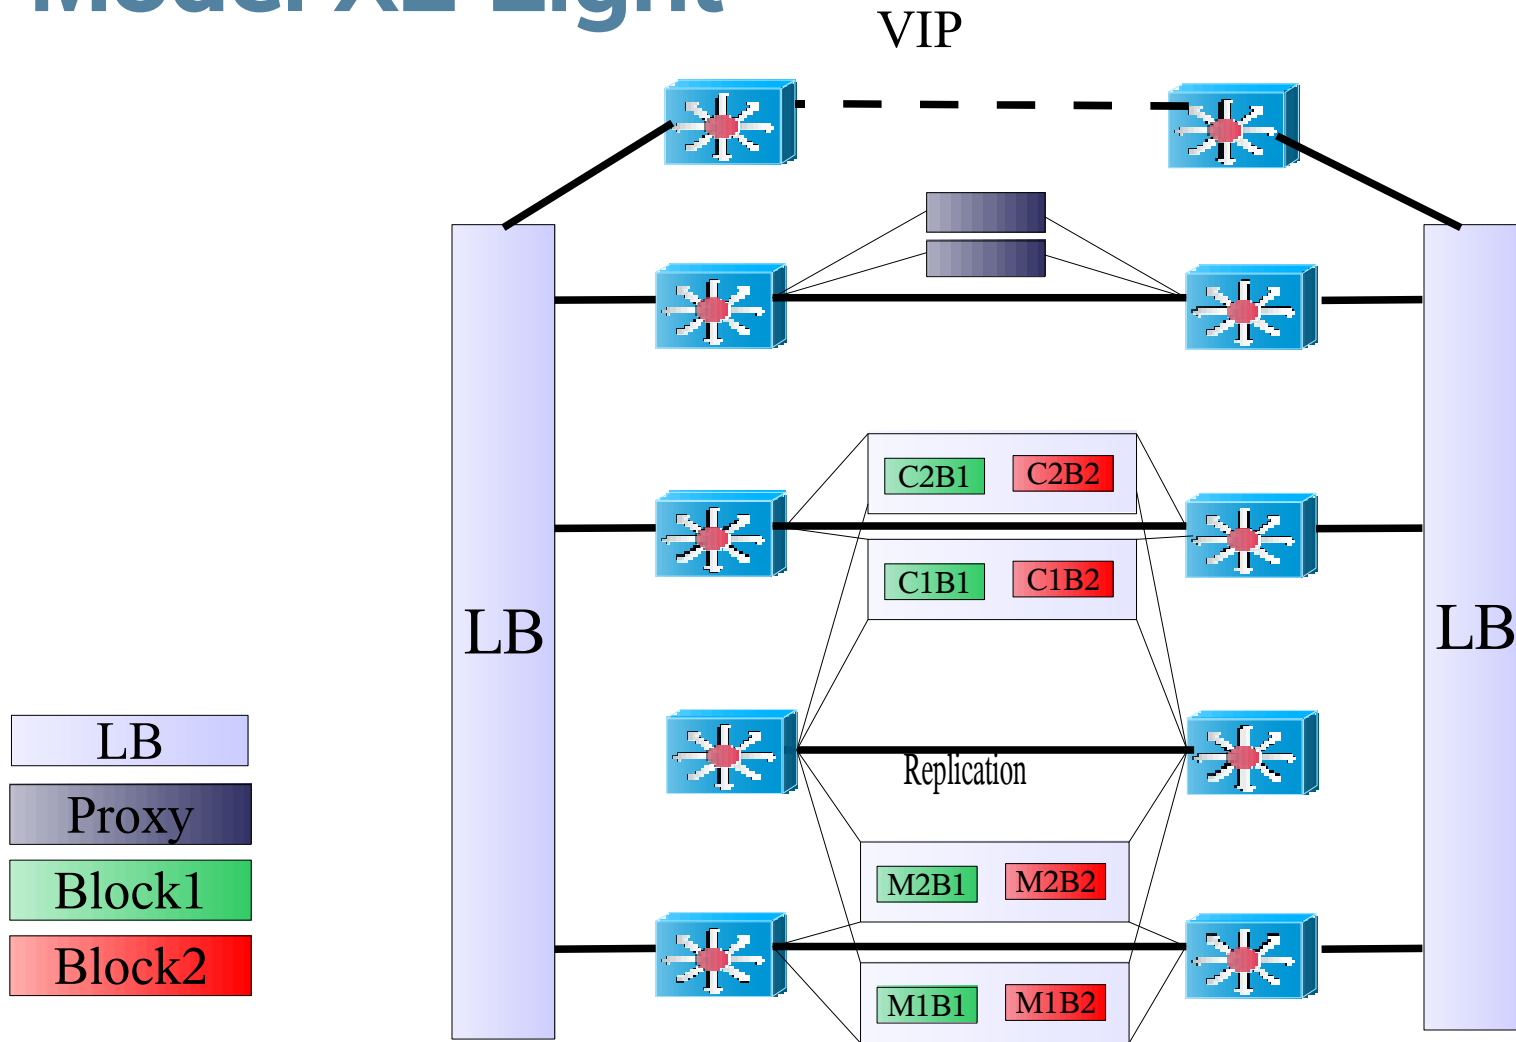
Model XL



Model XL-Light

- Multiple instances of DS on one server
- Higher write performance as Model L
- Proxy server points to multiple L-blocks hosted on a single physical L-block
- Each DS instance on a different port & dedicated mount point to its DB
- Easy upgrade to XL
 - > Stop the server & mount disks on new server
 - > Update LB switch: new real server
 - > Start new instance

Model XL-Light



Deployments

Deployment Examples (I)

- L – Model
 - > Up to 70 million entries (DS5.2SP4) in a single instance (Germany, France)
 - > Almost 100 % read access since years
 - > 99,99 % write access (scheduled down time)

Deployment Examples (II)

- > Many Hours to export/import LDIF file for 25 million database backend
- > < 2 minutes to stop, snapshot & start the server
- > 5 times re-initializing of all databases in staging environment in 12 months for testing purposes
- > Staging environment used for testing applications in different scenarios
- > Fast bulk load by binary copy from master to all replica servers
- > Instant Image for Point in Time Copy
- > Snapshot and in offline mode on DS5.x
 - > DS6.x provides frozen mode in online mode

Deployment Examples (III)

- XL-Light Model
 - > 4 partitions of data each around 10 million entries (Spain)
 - > Custom distribution algorithm
 - > Since 2003 in production
- XL Model
 - > Telco in Italy
 - > Custom distribution algorithm

Q & A

Your questions, please.


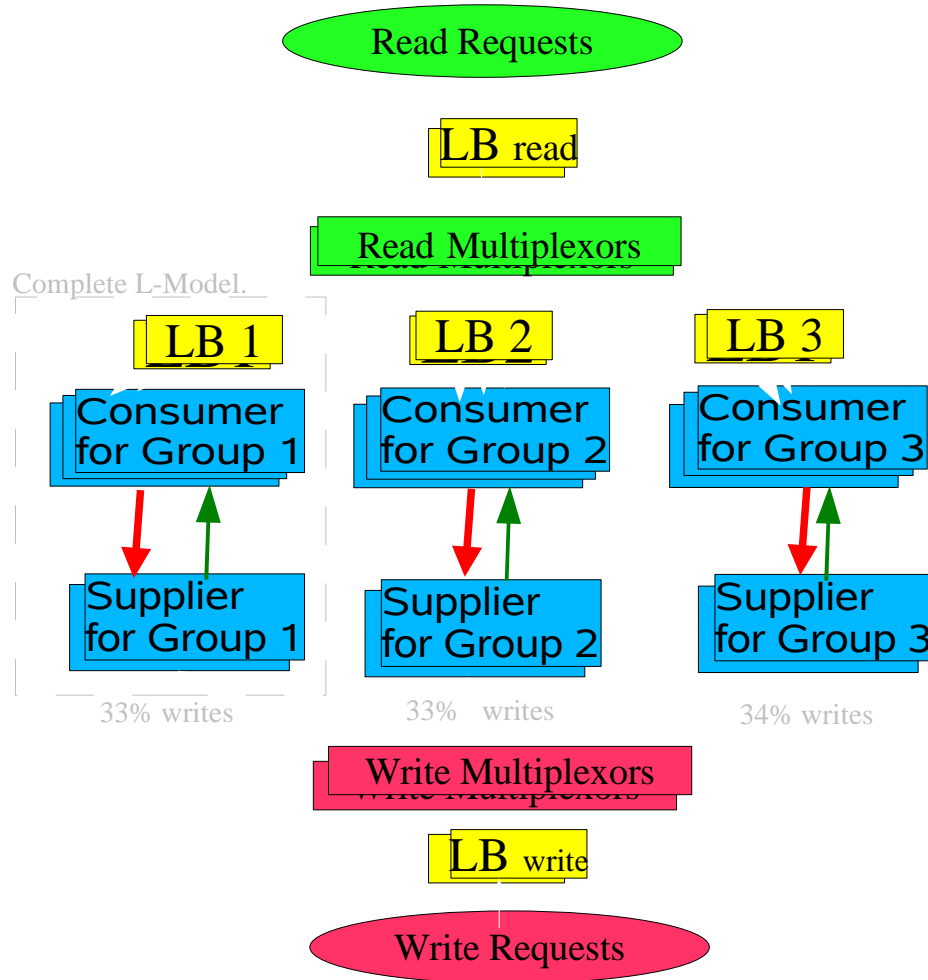


Thank you.

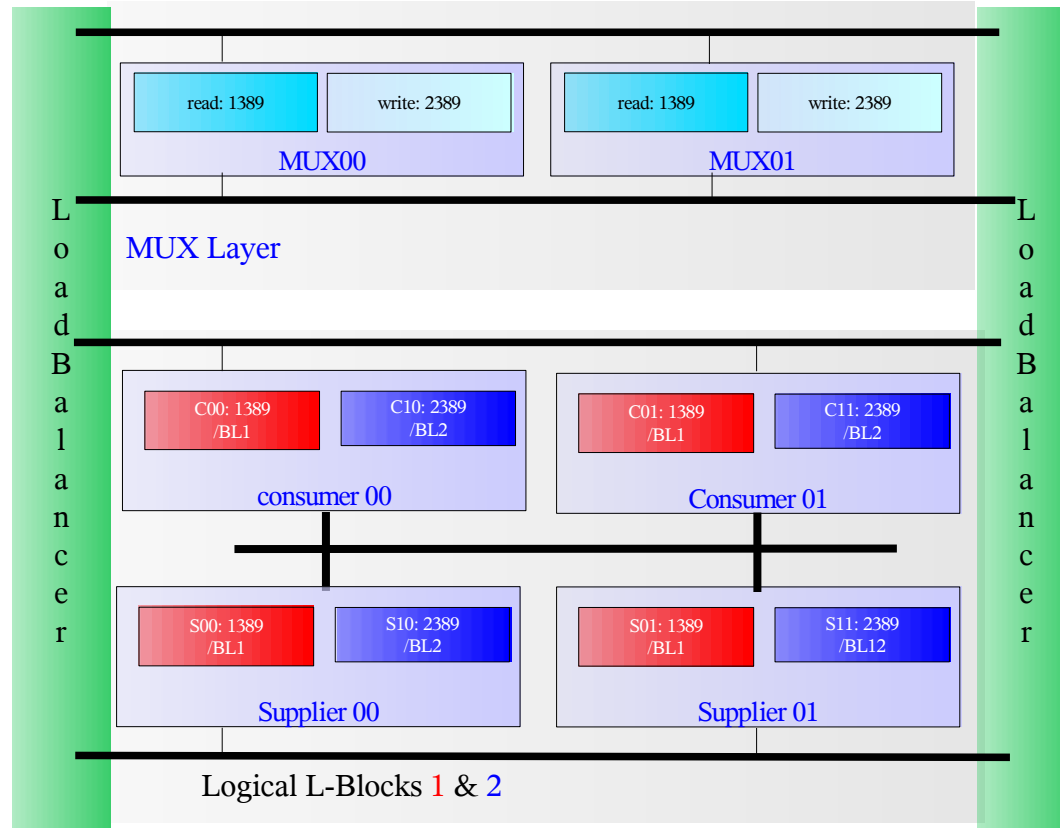
Abdi.Mohammadi@Sun.COM
Cengiz.Tuztas@Sun.COM



Model XL - Architecture



Model XL-Light



**End of
Overflow Slides**