

An Oracle White Paper  
February 2012

# Benefits of an Exclusive Multimaster Deployment of Oracle Directory Server Enterprise Edition

## Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Introduction .....	1
Legacy Master Read Servicing Model .....	2
Limited Number of Servers .....	2
Applications Conducting Both Read and Write Operations .....	2
Enhanced Directory Server Features for Write Operations .....	2
Write Referral Risk and Mitigation .....	3
Considerations for Directory Server Multimaster Deployments .....	3
Consumer Directory Deployment Models (Read Only) .....	4
Highly Secured Consumer Directory Server Model .....	4
Lightweight Consumer Directory Server Model .....	4
Comparison of Replication Topologies .....	5
Example of an Exclusively Multimaster Replication Topology .....	7
For More Information .....	8

## Introduction

One of the significant features introduced since Sun Directory Server Enterprise Edition 6 (hereafter referred to as "Directory Server 6") is that it allows unlimited master directory servers. This document describes the benefits of deploying an exclusively multimaster topology instead of a master, consumer, and replication hub directory topology as Directory Server 5.2 customers used to do. While this document uses Directory Server 11g as example for exclusive multimaster deployments, it is applicable to any versions starting with Directory Server 6.0.

Using an exclusively multimaster deployment provides the following benefits:

- Mitigates risk in isolating existing application write behavior
- Provides the greatest aggregate performance in a highly distributed and balanced model
- Lowers administrative burden in failover and recovery requirements

Customers using Directory Server 5.x masters for read traffic typically use one or more of the models described in this document. To define read and write operations, write operational traffic is considered to be any of the following: creation of new directory objects (in a typical user model, creation of new users), object modification (password, telephone, surname changes), and object deletion (user purging).

Read operational traffic is search operations (in a typical user model, authentications, authorizations, and entitlements)

## Legacy Master Read Servicing Model

### Limited Number of Servers

Customers can combine read and write traffic on the same physical servers. Typically, within the Directory Server 5.2 framework, this model is used to save in hardware costs. Services are potentially impacted during peak usage and maintenance. Degraded performance might occur for provisioning (write) applications during heavy read operation load and, alternatively, when heavier write demand impacts read or search availability.

High availability is generally not a key business requirement in these Directory Server 5.2 models.

### Applications Conducting Both Read and Write Operations

This model is typically used in companies when there is limited knowledge of the behavior or capabilities of their business applications. These applications might direct all operations to masters to avoid write referral requirements.

Application owners might independently direct new applications to the masters without the consent of the directory owners or administrators. These applications are sometimes identified only when observing directory server read operation statistics.

Bulk provisioning applications directed to the masters conduct minimal required read operations to manipulate and update target directory objects.

Within this model, multiple Directory Server 5.x Directory Proxy Server instances are sometimes deployed. One or more are designated for bulk provisioning applications using weighted or failover routing. The others are designated for everything else, including random writes distributed across a balanced routing table.

## Enhanced Directory Server Features for Write Operations

The following enhancements were introduced in Directory Server 6 (and thus in Directory Server 7 and Oracle Directory Server Enterprise Edition 11g). They are specific to write operation performance and agility, and they should be considered during the architectural definition stage.

- Faster write acceptance
- Support for unlimited master directories
- Increased write speed by distributing write operations
- Write affinity ability
- Enhanced replication ability and speed
- Support for full and partially meshed replication topologies
- Increased management and administration of replication

- Operational-based routing ability
- Enhanced proxy routing based on availability
- Enhanced memory and cache management
- Global user lockout on failed password attempts

### Write Referral Risk and Mitigation

Some Directory Server 5.x users experience risk related to applications and clients that are unable to effectively follow write referrals when write operations are requested from a consumer (read-only) directory server. This risk is mostly mitigated by using Oracle Directory Proxy Servers, because they can be defined to follow write referral requests on behalf of clients and applications. This transparent proxy operation is successful except for applications that do not generate the appropriate return codes. In these cases, a few options exist to mitigate risk; most involve small adjustments in the application code or the use of application environmental variables. Another mitigation that is used is to direct these specific applications and clients to master directory servers on which write referrals are not presented or required. Risk can be further mitigated by using Directory Server 11g.

Directory Server 11g supports an exclusive multimaster based architecture with unlimited masters, which can avoid write referral requirements because each master has the ability to accept localized write operations. The Directory Proxy Server also has the ability to direct write operations to the master directory servers. This ability can be defined according to design preferences, so a primary master can be thought of as the target or it can be routed using a balanced or weighted definition across the master directory servers. Directory Server 5.x also supports an exclusive multimaster architecture, but it supports a maximum of four multimaster directory servers.

### Considerations for Directory Server Multimaster Deployments

Full, exclusive use of master directory servers can be a deployment practice with Directory Server 11g. Read-only consumer directories can be deployed only to address specific business requirements. An exclusive master topology provides such high availability, increased aggregate performance advantages, and growth flexibility that read-only consumer directories might be less desirable and might be used only to fulfill specific business requirements.

Replication hubs, which are frequently used in Directory Server 5.x architectures, might be less significant in Directory Server 11g deployments. Support of full and meshed replication topologies and the enhanced management of replication agreements might make replication hub components less relevant and without cost advantage. Companies can replace the replication hub architecture in their legacy deployments with an additional multimaster directory to gain

additional aggregate performance in a load-balanced strategy and also to provide additional unattended failover redundancy within their framework.

Synchronizing failed password attempts is achieved by writeable master directory servers. Consumer directory servers can be leveraged in a failed-password, retry-count model when operational routing is used to direct authentications to masters. Comparisons between an exclusive multimaster model and a master-to-consumer model are provided later in this document.

## Consumer Directory Deployment Models (Read Only)

### Fractional Replication Consumer Directory Server Model

Fractionally replicated consumer directory servers provide directories that are deployed with a minimal set of user attributes in a read-only state, and these servers can be positioned in unsecured openly accessible areas (DMZs). In addition, these servers can be configured with fewer hardware requirements than the masters due to their limited data sets, potential for lower indexing and cache, and lack of a change log database.

### Highly Secured Consumer Directory Server Model

Consumer directory servers are not capable of accepting write operations. These servers can be further protected by prohibiting indirect client or application writes by blocked write referrals when Directory Proxy Servers are used. Consistent protections can be provided in an exclusively multimaster deployment as well through options available within the Directory Proxy Servers and native directory server.

A typical use case for a restricted consumer read-only topology in which writes cannot be introduced, even indirectly by referrals, might be when the authoritative source of data exists in a back-end data repository pushed to the LDAP directory. Manipulation of data is never initiated within the LDAP architecture by LDAP clients or applications. The LDAP architecture in this model is basically a consumer of changes that are made from other data sources and pushed to the LDAP directory.

### Lightweight Consumer Directory Server Model

Consumers can be deployed in a combined master-to-consumer model consistent with the Directory Server 5.x model. There is little gained in this two-directory-server-profile model (three profiles when replication hubs are used), and it includes a loss of some of the benefits gained in an exclusively master deployment. The emphasis in a consumer model, relative to performance, is typically to reduce directory server overhead by not maintaining a localized change log database. Read-only consumer directory servers do not maintain replication synchronization change logs.

A master-to-consumer or multimaster-to-consumer model can include Directory Server 11g operation-based routing or write referral handling through the ODSEE 11g Directory Proxy Services.

## Comparison of Replication Topologies

	EXCLUSIVELY MULTIMASTER TOPOLOGY	MULTIMASTER-TO-CONSUMER TOPOLOGY
<b>HIGH AVAILABILITY</b>	High availability is realized because all masters can be configured consistently. Any master can be used to restore or initialize another master or build an additional master.	High availability is provided by redundancy within each of the server profiles, masters, consumers, and replication hubs. Greater hardware investment is typically required in definition of the different server profiles because redundancy and sustained service levels are required for each profile.
<b>AGGREGATE PERFORMANCE</b>	When all exclusive master servers are configured consistently throughout and balanced through use of the ODSEE 11g Directory Proxy Servers, aggregate performance can be realized and grown inserting additional servers without any adverse service impact.	Having two or three profiles of servers, masters, and consumers, and potentially having replication hubs, and incorporating required redundancy within each of the profiles typically results in only similar servers being balanced and used for aggregate usage. There is greater likelihood that one profile or another will be less utilized, and aggregate performance potential from the whole architecture is not available.
<b>WRITE REFERRALS AND LOCALIZED WRITE OPERATIONS</b>	All masters have the ability to accept localized writes and replicate these writes to other directory servers. No write referrals are required when only masters are used. This increases the speed of the write commitment because additional operational overhead is not required. Write affinity can also be leveraged so that immediate searches after writes are fulfilled successfully.	Only master directory servers are capable of accepting local writes. Consumer directories require write referral handling or operation-based routing to a master for write acceptance. Write referral handling and operation-based routing can be facilitated through the ODSEE 11g Directory Proxy Services.
<b>CHANGE LOG DATABASES</b>	<p>Change log databases exist only on master directory servers and are used to maintain replication synchronization to peer or subordinate servers.</p> <p>The advantage is that all masters can be identically configured and maintain synchronization. A master server can be restored or initialized quickly from another master and maintain quick synchronization based on the use of change log databases.</p> <p>A small disadvantage of having local change log databases on each master is the additional server disk write I/O. This disadvantage is mitigated by</p>	<p>Change log databases exist only on master directory servers and are used to maintain replication synchronization to peer or subordinate servers.</p> <p>Consumers can be promoted to masters upon a recovery requirement but this requires administrative intervention, and a consumer might require significant time to establish itself as a master or peer with other servers.</p> <p>Best practices do not depend on consumer promotion to master strategies. The best practice master recovery model uses an alternate multimaster in an unattended failover. Best practice deployments include a minimum of three master directory servers</p>



	EXCLUSIVELY MULTIMASTER TOPOLOGY	MULTIMASTER-TO-CONSUMER TOPOLOGY
	faster directory write capabilities in Directory Server 11g and by potentially having fewer servers than in a typical redundant master, redundant consumers model (that possibly also has replication hubs) and by having fewer but higher-performance servers.	for high availability of writes. Consumers without change log database write overhead consume less disk I/O.
<b>PASSWORD RETRY LOCKOUT</b>	Only writable servers (masters) are capable of maintaining and synchronizing the failed password attempts. Lockout state is replicated automatically regardless of whether the directory is a master or a consumer.	Only writable servers (masters) are capable of maintaining and synchronizing the failed password attempts, so when consumers are used, administrators can direct all authentication requests to the master servers and redirect to consumers their normal search and entitlement operations. This operation-based routing is available through the Directory Server 6 Directory Proxy Server. Lockout state is replicated automatically regardless of whether the directory is a master or a consumer.
<b>BINARY BACKUP AND RESTORE</b>	Binary recovery is supported only by "like" servers. A master with a consistent configuration can be used to initialize or recover another master server. Any master can be used in a consistent model.	Binary recovery is supported only by "like" servers. A master with a consistent configuration can be used to initialize or recover another master server. A consumer directory can be used only to initialize or recover another consumer directory server.
<b>HARDWARE REQUIREMENTS</b>	Fewer servers can be defined in this architecture because each master can have a consistent definition and fulfill both read and write operations. New master servers can be included in aggregate use through proxy and load balancing.	Typically, more servers are used because specific servers are assigned to specific assignments. Adequate server redundancy is required for consumers and masters.
<b>REPLICATION HUB REQUIREMENTS</b>	Typically, in exclusive multimaster topologies, replication hubs are not required because flexible replication routing definitions across masters can be achieved.	Replication hubs might be required in some master-to-consumer topologies to offload master replication overhead to the subordinate consumer servers.
<b>LOAD-BALANCED WRITES</b>	Performance benchmark results indicate that ODSEE 11g enhanced collision avoidance and the speed of directory writes contribute to significant gains in distributing directory writes. Customer preference will dictate this write routing strategy and decisions can be made virtually on the fly between strategies within the Directory Proxy Server.	In a master-to-consumer model, only the masters accept the balanced or target writes. Not utilizing all servers as masters dictates that fewer masters are available than in an exclusive model, and this subsequently reduces the aggregate benefit realized in benchmarks distributing the write load.

## Example of an Exclusively Multimaster Replication Topology

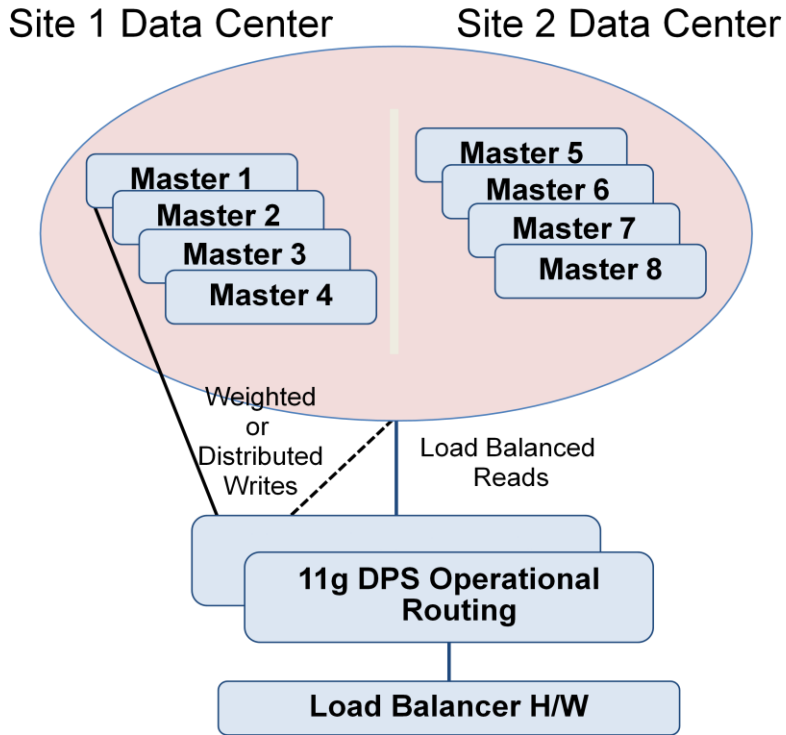


Figure 1: Oracle Directory Server Enterprise Edition 11g Topology with two data centers

## For More Information

Here are additional resources:

- Oracle Directory Server Enterprise Edition web site:  
<http://www.oracle.com/technetwork/middleware/id-mgmt/overview/index-085178.html>
- Oracle Directory Server Enterprise Edition downloads:  
<http://www.oracle.com/technetwork/middleware/id-mgmt/downloads/index.html>
- Oracle Directory Server Enterprise Edition documentation:  
[http://docs.oracle.com/cd/E20295\\_01/index.htm](http://docs.oracle.com/cd/E20295_01/index.htm)
- Oracle Directory Server Enterprise Edition training courses:  
[Oracle Directory Server Enterprise Edition 11g: Maintenance and Operations](#)
- Oracle Directory Server Enterprise Edition blog:  
<http://blogs.oracle.com/directoryservices>



Benefits of an Exclusive Multimaster  
Deployment of Oracle Directory Server  
Enterprise Edition  
February 2012  
Author: Michael Melore  
Contributing Authors: Etienne Remillon, Forest  
Yin, Olaf Stullich

Oracle Corporation  
World Headquarters  
500 Oracle Parkway  
Redwood Shores, CA 94065  
U.S.A.

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200  
oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2012, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

**Hardware and Software, Engineered to Work Together**